Policy Management

CMP Wireline User Guide

910-6895-001 Revision A

October 2013

Tekelec

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 wireline 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) device, which manages multiple network-based client sessions; the network in which the MPE device operates; policies; and the Configuration Management Platform (CMP) system, which controls MPE devices and associated applications.
- *Configuring the Policy Management Topology* describes how to set the topology configuration.
- *Managing MPE Devices* describes how to use the CMP system to configure and manage the MPE devices in a network.
- *Managing Network Elements* describes how to manage network elements.
- *Managing Application Profiles* describes how to manage application profiles.
- *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 Subscribers* describes how to manage subscriber tiers and accounts within the CMP system.
- *System-Wide Reports* describes the reports available on the function of Policy Managementsystems 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 Policy Management 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	**Danger**: (This icon and text indicate the possibility of *personal injury*.)
WARNING	**Warning**: (This icon and text indicate the possibility of *equipment damage*.)
CAUTION	**Caution**: (This icon and text indicate the possibility of *service interruption*.)
TOPPLE	**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

Note: Some of the documents that were released in support of Release 6.4 have since been replaced in other releases. These changes are reflected in the documents listed below.

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*
- *CMP Cable User Guide*
- *Troubleshooting Reference Guide*
- *SNMP User Guide*
- *OSSI XML Interface Definitions Reference Guide*
- *Platform Configuration User Guide*
- *Bandwidth on Demand Application Manager User Guide*
- *PCMM specification PKT-SP-MM-I06* (third-party document, used as reference material for PCMM)

Wireless

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

Wireline

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

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 provides a policy and charging rules function (PCRF) as defined in the 3rd Generation Partnership Project (3GPP) technical specification “Policy and charging control architecture” (TS 23.203). The 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.

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

The MPE device provides Call Admission Control (CAC) to support video on demand (VoD). *Figure 1: The Policy Management Solution and MPE Devices* shows how the Tekelec Policy Management solution fits into a wireline network containing VoD servers and bandwidth remote access server (B-RAS) routers.

The major elements of a Policy Management network are:

- MPE devices — Provide policy control decisions and flow-based charging control. When a request for a policy decision is received for a subscriber session, the MPE device obtains subscriber information, evaluates the applicable policies, and directs the enforcement device to handle the session based on policy rules.
- Configuration Management Platform (CMP) — Provides the policy console. The CMP system contains a centralized repository for configuration information that is used by MPE devices to make CAC decisions, including policy rules, policy objects, network elements, bandwidth allocation per interface, element links, and subscriber data. Carriers can exchange database information in eXtensible Markup Language (XML) format with an office support system (OSS). The Policy Management network can communicate with a network management station (NMS) using Simple Network Management Protocol (SNMP).

The Policy and Charging Enforcement Function (PCEF) receives requests to start new sessions for subscribers. An example of a PCEF is a B-RAS device. MPE devices communicate with PCEFs to receive requests for policy decisions and send those policy decisions to PCEFs for implementation.

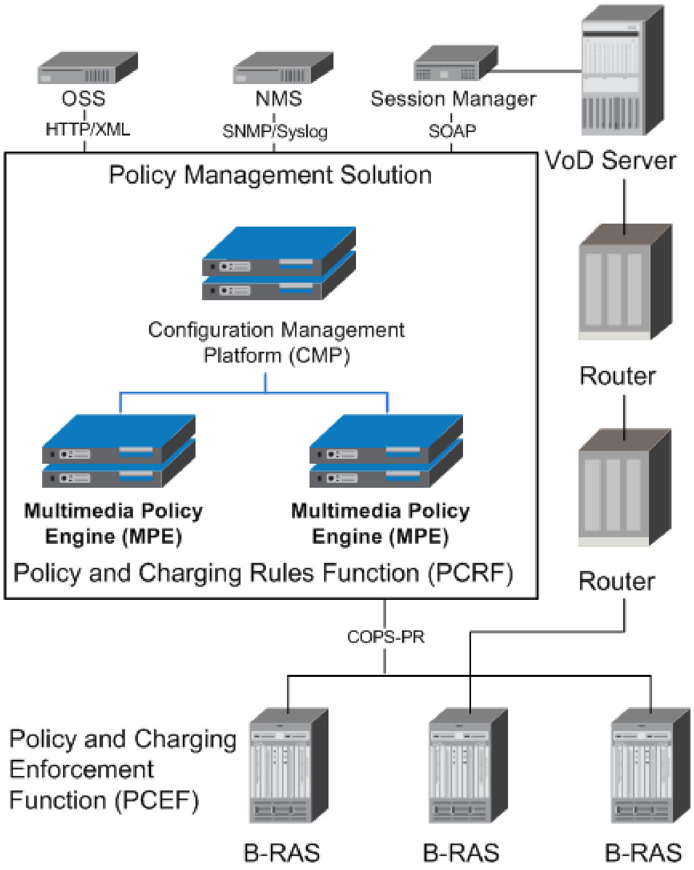
OSS
HTTP/XML
NMS
SNMP/Syslog
Session Manager
SOAP
VoD Server
Policy Management Solution
Configuration Management Platform (CMP)
Multimedia Policy Engine (MPE)
Multimedia Policy Engine (MPE)
Policy and Charging Rules Function (PCRF)
Router
Router
COPS-PR
Policy and Charging Enforcement Function (PCEF)
B-RAS
B-RAS
B-RAS

Figure 1: The Policy Management Solution and MPE Devices

Configuration information is pushed to the CMP database by an OSS system or manually configured using the CMP Graphical User Interface (GUI). The CMP device is responsible for pushing this data to MPE devices in the network.

A B-RAS router connects to the MPE device using the COPS-PR protocol, and once a connection is established, the MPE device synchronizes with the B-RAS router configuration information, providing the MPE device with the current list of subscribers connected to the B-RAS router.

As users connect to or disconnect from the network, the B-RAS router notifies the MPE device about the users' IP addresses. This information is later used by the MPE device in making CAC decisions.

When a user generates a VoD request, a Session Manager device requests a CAC decision from an MPE devices. The MPE device, using the IP address of the subscriber and the VoD server in the request, determines a path between the source (the VoD server) and the destination (the B-RAS router) that would be used for the video data stream. This path could consist of multiple network segment hops; for example, the path could be from the VoD server to a video distribution router, to a video aggregation router, to a gateway router (GWR). Once it has the path, the MPE device runs the path through configured policies that limit the usage on the segments in the path. If any one of the segments fails a policy test, the MPE device returns a CAC failure message to the Session Manager. If all the segments in the path pass the policy test, the MPE device returns a CAC success message to the Session Manager.

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.
- 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.

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:

- 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*.

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 2: Structure of the CMP GUI* shows the structure of the CMP GUI.

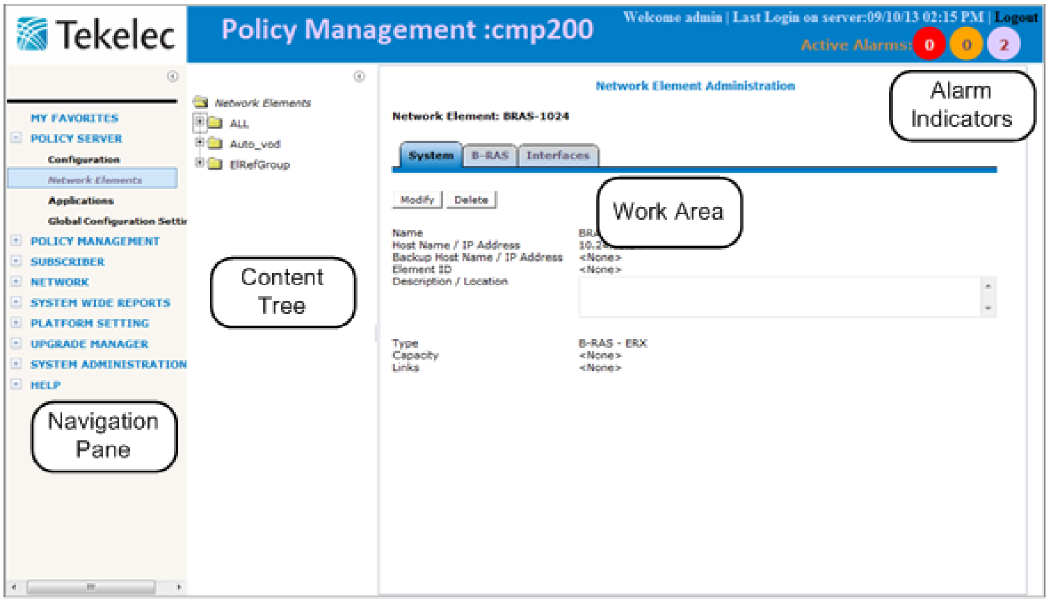

Network Element: BRAS-1024

Name BRA
Host Name / IP Address 10.2
Backup Host Name / IP Address <None>
Element ID <None>
Description / Location

Type B-RAS - ERX
Capacity <None>
Links <None>

Figure 2: 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**.

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

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

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

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

- **Work Area** — Contains information that relates to choices in both the navigation pane and the content tree. This is the area in which you perform all work.
- **Alarm Indicators** — Provides visual indicators that show the number of active alarms.

Specifications for Using the GUI

Tekelec recommends the following:

- **Web Browsers** —
  - Mozilla Firefox release 23.0.1 or higher
  - Microsoft Internet Explorer 9.0 or higher
- **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 3: 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.

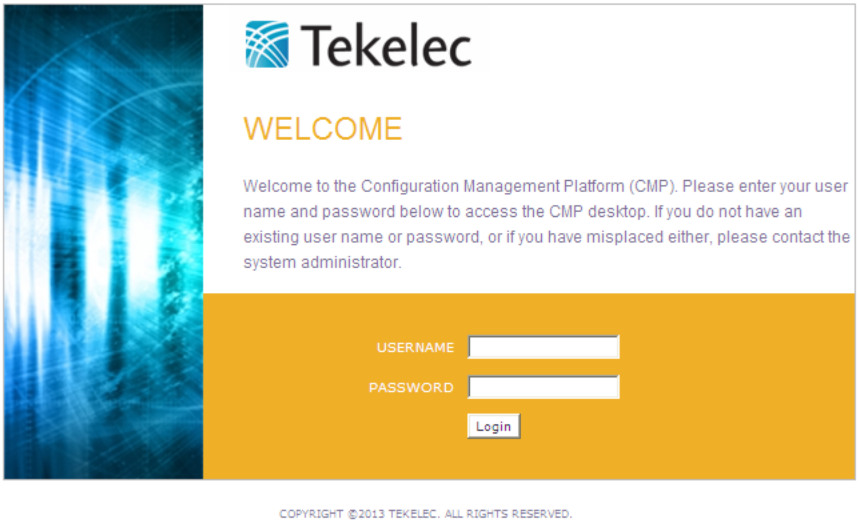

WELCOME

Welcome to the Configuration Management Platform (CMP). Please enter your user name and password below to access the CMP desktop. If you do not have an existing user name or password, or if you have misplaced either, please contact the system administrator.

COPYRIGHT ©2013 TEKELEC. ALL RIGHTS RESERVED.

Figure 3: CMP Login Page

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.

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:

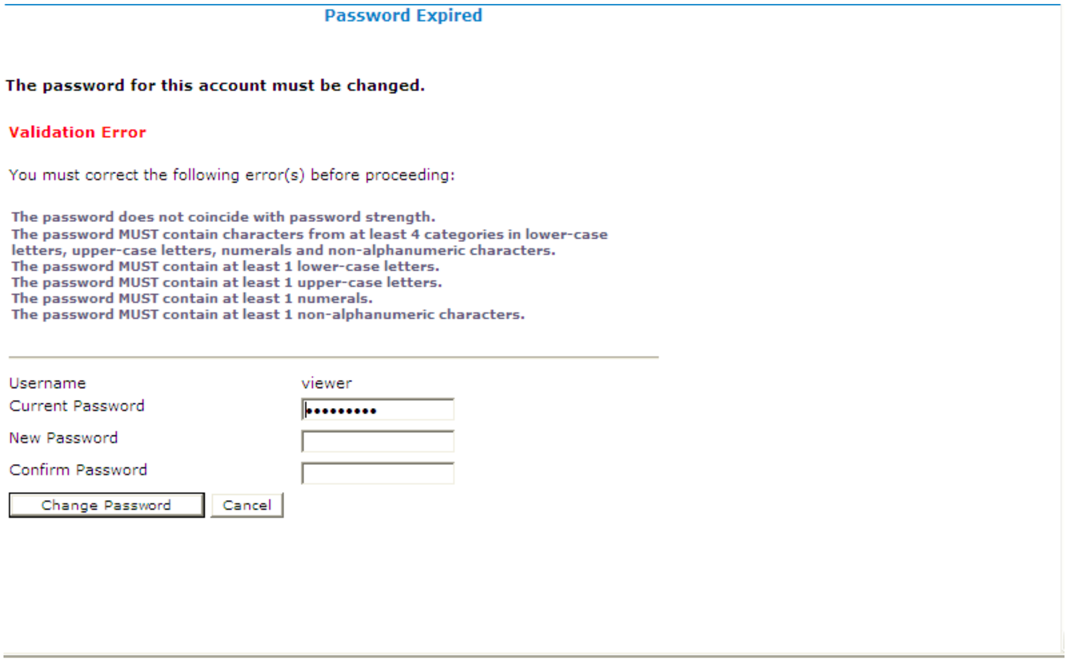

Password Expired

The password for this account must be changed.

Validation Error

You must correct the following error(s) before proceeding:

The password does not coincide with password strength.
The password MUST contain characters from at least 4 categories in lower-case letters, upper-case letters, numerals and non-alphanumeric characters.
The password MUST contain at least 1 lower-case letters.
The password MUST contain at least 1 upper-case letters.
The password MUST contain at least 1 numerals.
The password MUST contain at least 1 non-alphanumeric characters.

Username viewer
Current Password •••••••••
New Password
Confirm Password

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

Your password is changed.

Overview of Main Tasks

The major tasks involved in using MPE devices are configuration, defining manageable elements and profiles, creating and deploying policy rules, 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, which defines the addresses of Policy Management clusters in your network. These steps are described in *Configuring the Policy Management Topology*.

The element and profile definition tasks you need to perform depend on what exists on your network. They can be defined in any order at any time as needed. Once elements and profiles are defined, you can refer to them in policy rules. 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*.

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

1. Create policy rules in the CMP database. This step is described in *Understanding and Creating Policy Rules*.
2. Deploy the policy rules from the CMP database 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:

- Manage subscriber tiers and accounts. This task is described in *Managing Subscribers*.
- View reports the function of the Policy Management systems in your network. This task is described in *System-Wide Reports*.
- Manage CMP users, accounts, access, authorization, and operation. This task is described in *System Administration*.
- Upgrade software using the Upgrade Manager GUI page. This page is described in *Upgrade Manager*.

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 and MPE devices). The topology determines the following:

- How clusters are set up
- How configuration data is replicated
- How incidents (events and alarms) get reported to the CMP system that controls the Policy Management network.

Figure 4: Policy Management Topology illustrates a Policy Management topology consisting of a CMP cluster and two MPE clusters.

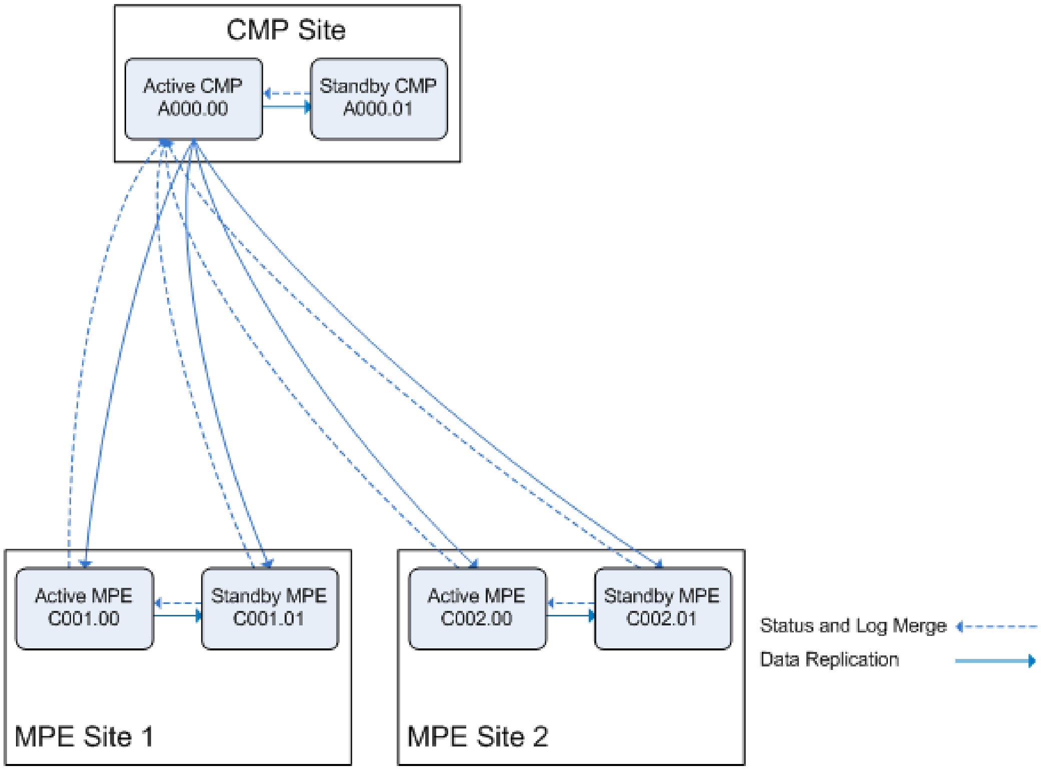

Figure 4: Policy Management Topology

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 alarms.

Separating OAM and signaling traffic allows the ability to shut down one network without affecting the other, and also the opportunity to include separate and redundant signaling (SIG-A and SIG-B) networks.

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 in-memory 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.

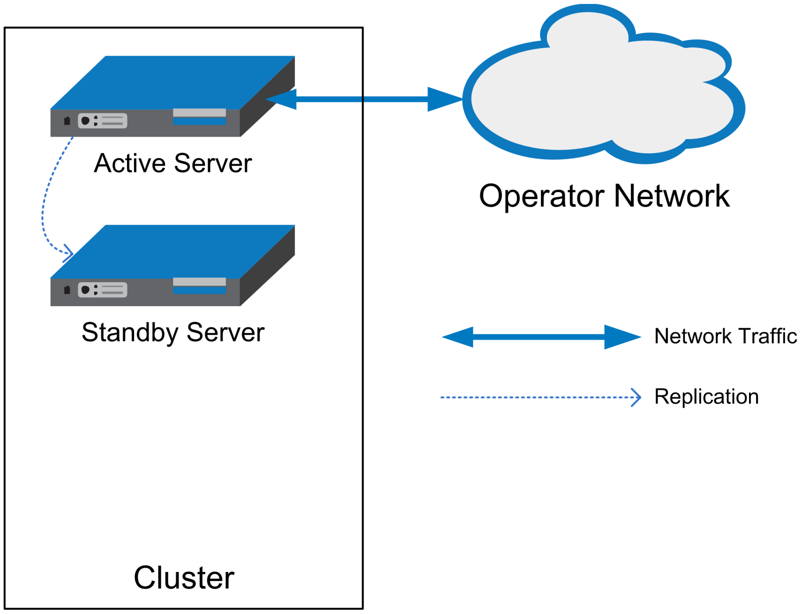

Figure 5: High Availability

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.

- **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.
- **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

Setting Up the Topology

Topology configuration consists of defining Policy Management sites and clusters, including their addresses and hierarchy. You can add MPE 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. Configure the CMP cluster settings. The settings are replicated (pushed) to the standby CMP server. Together, the two servers form the primary CMP site for the whole topology network. The primary site cannot be deleted from the topology.
2. Configure MPE clusters — Enter MPE 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 the active and standby servers. These servers form an MPE 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.

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

Setting Up the CMP Cluster

You must define a CMP cluster before continuing with the topology. The site you define will be the primary (Site 1) cluster.

Before defining the 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 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.
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 6: Cluster Settings Page for CMP Cluster* shows an example):
   a) **HW Type** — Select **RMS** (for a rack-mounted server).
   b) **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.

   c) **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.

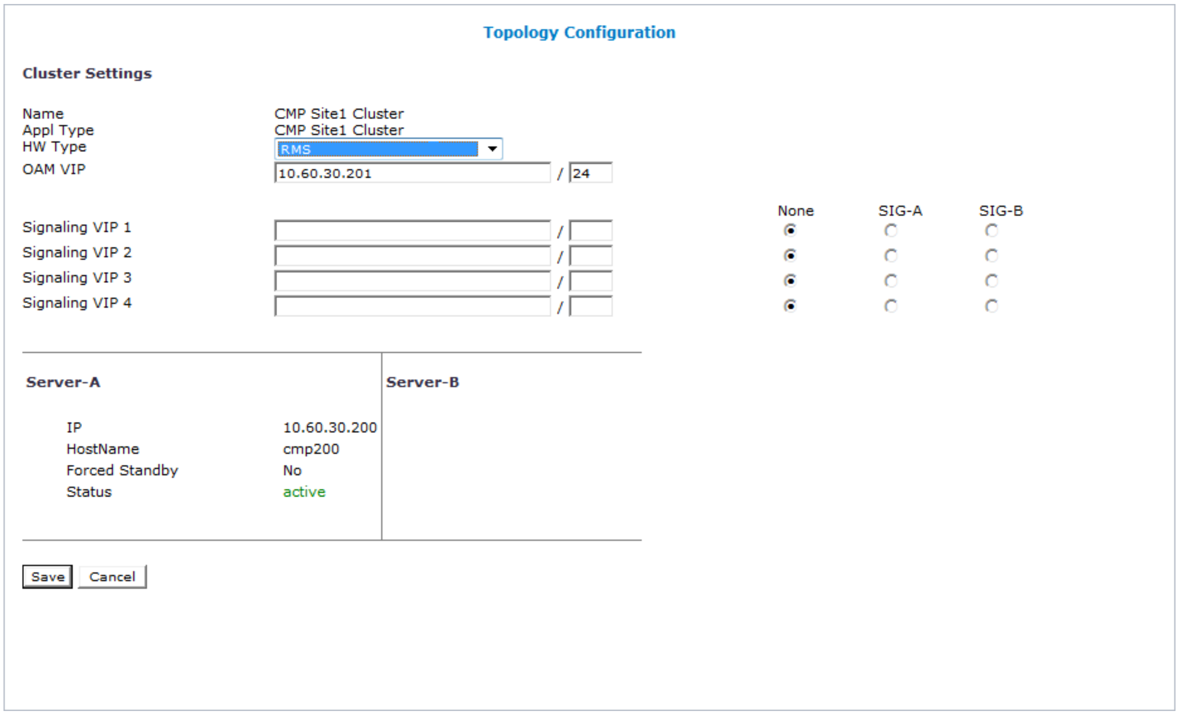

Topology Configuration

Cluster Settings

Name	CMP Site1 Cluster				
Appl Type	CMP Site1 Cluster				
HW Type	RMS				
OAM VIP	10.60.30.201 / 24				
			None	SIG-A	SIG-B
Signaling VIP 1	/		●	○	○
Signaling VIP 2	/		●	○	○
Signaling VIP 3	/		●	○	○
Signaling VIP 4	/		●	○	○

Server-A

IP	10.60.30.200
HostName	cmp200
Forced Standby	No
Status	active

Server-B

Figure 6: Cluster Settings Page for CMP Cluster

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

Setting Up an MPE Cluster

Before defining an MPE cluster, ensure the following:

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

To define an MPE cluster:

1. From the **Platform Setting** section of the navigation pane, select **Topology Setting**.
   The Cluster Configuration page opens.
2. Click **Add MPE/MRA Cluster**.
   The Topology Configuration Page opens.
3. Enter the following information (*Figure 7: Cluster Settings Page for MPE Cluster* shows an example):
   a) **Name** (required) — Name of the cluster. Enter up to 255 characters, excluding quotation marks (") and commas (,).
   b) **Appl Type** — Select **MPE** (the default).
   c) **HW Type** — Select **RMS** (for a rack-mounted server).

   d) **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 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.

   e) **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 cluster. (To support redundant communication channels, an MPE 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 cluster, at least one signaling VIP, either SIG-A or SIG-B, is required.

4. Select **Server-A** and enter the following information for the first server of the cluster:
   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`).
5. Once you define Server A, you can optionally click **Add Server-B** and enter the appropriate information for the second server of the cluster.
6. When you finish, click **Save** (or **Cancel** to discard your changes).
7. If you are setting up multiple clusters, repeat the above steps as often as necessary.

The MPE cluster is defined.

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

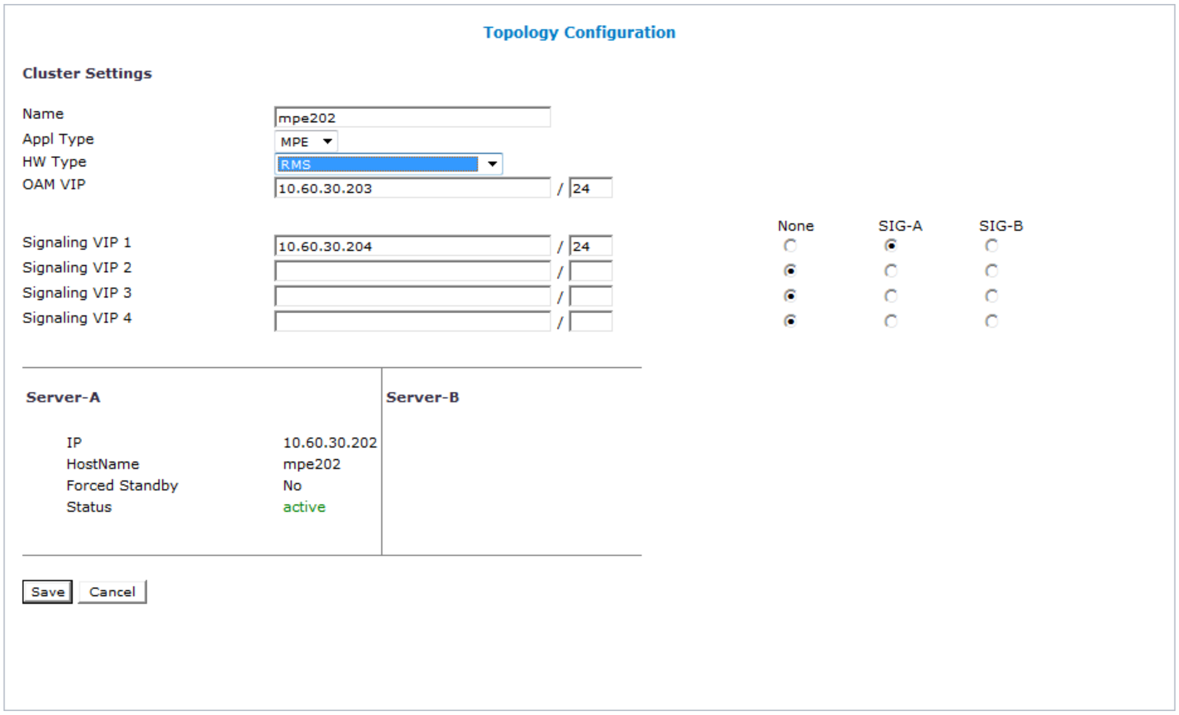

Topology Configuration

Cluster Settings

Name mpe202

Appl Type MPE

HW Type RMS

OAM VIP 10.60.30.203 / 24

			None	SIG-A	SIG-B
Signaling VIP 1	10.60.30.204	/ 24	○	●	○
Signaling VIP 2		/	●	○	○
Signaling VIP 3		/	●	○	○
Signaling VIP 4		/	●	○	○

Server-A

IP 10.60.30.202

HostName mpe202

Forced Standby No

Status active

Server-B

Figure 7: Cluster Settings Page for MPE 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 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 an MPE Cluster*
- *Modifying a CMP Cluster*
- *Removing a Cluster from the Topology*
- *Forcing a Server into Standby Status*

Modifying an MPE Cluster

To modify an MPE 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 server, click **Modify Server-A**.
   - To modify the secondary server, click **Modify Server-B**.

   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 either server from a cluster, but not both. You can select **Forced Standby** on one or more servers of an MPE cluster.

   CAUTION

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

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

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 *Setting Up the CMP Cluster*.
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 cluster from the topology. (You cannot remove the Site 1 (primary) CMP cluster from the topology.)

CAUTION **Caution:** Contact Tekelec Technical Support before restoring a cluster deleted from the topology.

Before removing an MPE 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.

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.”

CAUTION **Caution:** If you force all servers in a cluster into Standby status, you can trigger a site outage.

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. Click either **Modify Server-A** or **Modify Server-B**, as appropriate.
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.

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. The field is required to allow the Manager to receive traps. By default, these fields are empty.

Field Name	Description
	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”.
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”.

Field Name	Description
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 Stats Settings

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

To change stats settings, do the following:

1. From the **Policy Server** 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

   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. 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, Policy Log, Syslog, and session synchronization log configurations.
   - **Policy Server** — Lets you associate applications and network elements with the MPE device and configure protocol information.
   - **Policies** — Lets you manage policies that are deployed on the policy server.
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 an MPE device profile from the ALL group also deletes it from any associated group.

To delete an MPE device profile:

1. From the **Policy Server** section of the navigation pane, select **Configuration**.
   The content tree displays a list of policy server groups; the initial group is **ALL**.
2. From the content tree, select the **ALL** group.

   The Policy Server Administration page opens in the work area, displaying all defined MPE devices; for example:

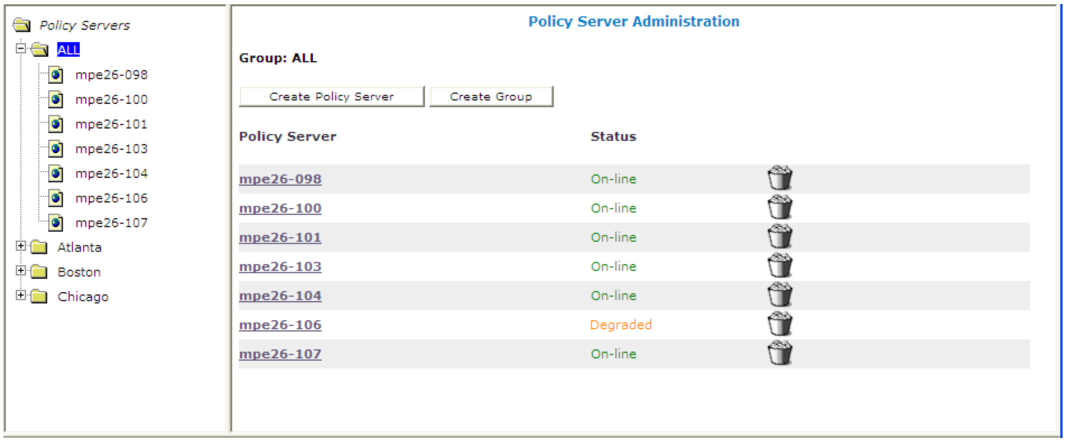

Policy Server Administration

Group: ALL

Policy Server	Status	
mpe26-098	On-line	🗑
mpe26-100	On-line	🗑
mpe26-101	On-line	🗑
mpe26-103	On-line	🗑
mpe26-104	On-line	🗑
mpe26-106	Degraded	🗑
mpe26-107	On-line	🗑

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.
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`).
Validate wireline user	If enabled, sessions for unknown users are rejected.
VoD Server Synchronization	
Tandberg Interval (minutes)	The interval of time specified for synchronization between a TANDBERG VoD server and the MPE device. When a gateway establishes a new connection to the MPE device, the MPE device initiates either a full or incremental synchronization. This field accepts a numeric value of 1–99999. **Note:** If your attributes are defined and you want to synchronize the VoD server, click **Sync Tandberg VoD Server Now**. Synchronization is initiated with all known TANDBERG VoD servers within ten seconds. When the synchronization is completed, the gateway sends Address Management messages to the MPE device whenever a new user connects to or disconnects from the network.
Tandberg Application Name	Set to match the name that the TANDBERG VoD server sends. The default is **OpenStream**.
Tandberg Keep Alive (seconds; 0 is disabled)	The interval of time, in seconds, before the MPE device issues a KEEPALIVE status request to the TANDBERG server in the absence of any other traffic. The default value is 0 (no keepalive messages are sent). **Note:** The MPE device logs, but does not take any other action, if the TANDBERG server does not respond to a keepalive request.
Bras Server Synchronization	

Attribute	Description
Concurrent Bras synchronizations	When a Juniper B-RAS server connects to an MPE device, the MPE device issues a synchronization request to it. This causes the B-RAS server to send a COPS-PR request for each attached subscriber that it knows about. This can potentially result in thousands of messages being returned to the MPE device. This option limits the number of outstanding synchronization requests that the MPE device will have active at any given time. The default is 8.
Fast Sync Enabled	When enabled, when an ERX device connects to the MPE device, if the IP address it reports of the last policy server to which it connected matches the MPE device address, the MPE device sends an unsolicited Decision (DEC) message to the ERX device, which replies with just Request (REQ) and Delete (DRQ) messages instead of full state information. The default is disabled.
Bras Buffers Configuration	
TCP Send Buffer (bytes)	The default is 0 bytes.
Max Size of TCP Send Buffer (bytes)	The default is 4,194,304 bytes (4 MB).
Shrink Wait Time (milliseconds)	The default is 3,000,000 ms (50 minutes).
Load Shedding Configuration	
Enabled	Select to enable 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

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.

- **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:

    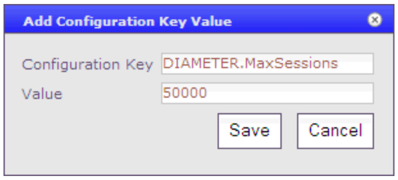

    Configuration Key DIAMETER.MaxSessions

    Value 50000

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

    CAUTION

    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.

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 (,).

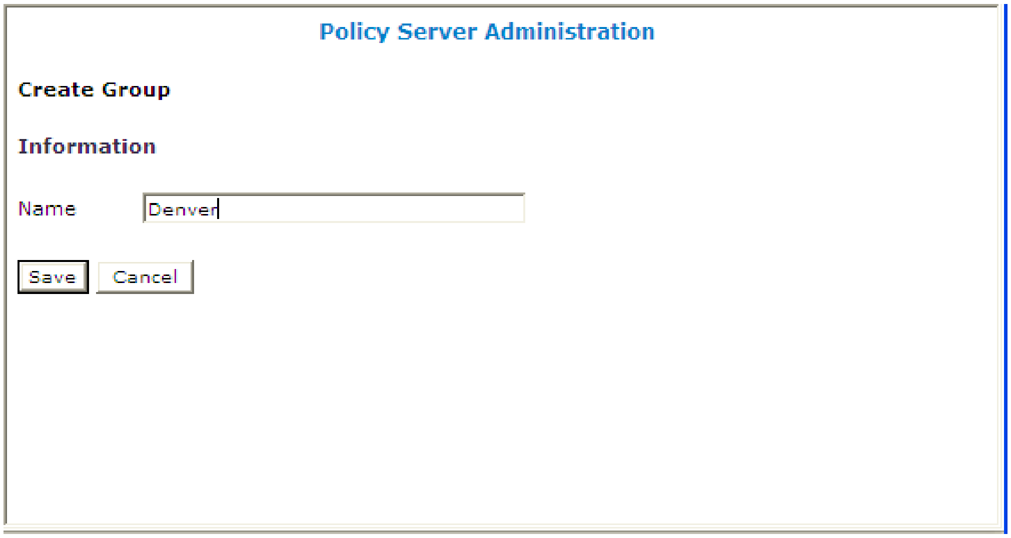

Policy Server Administration

Create Group

Information

Name: Denver

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 and links), 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.
5. Click **Reapply Subscriber Configuration**.
   Subscriber 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** — The servers in the cluster have completed startup, and their database services are synchronized.
  - **Degraded** — At least one server is not functioning properly (its database services are not synchronized or it has not completed startup) or has failed, but the cluster continues to function with the active 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.
  - **Out of Service** — Communication to the cluster has been lost.
  - **No Data:** Communication to the cluster has been lost. This status value provides backward compatibility with previous Policy Management releases. It can be observed during the upgrade process.
  - **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 8: Group View shows an example of a Group View in which one of the servers is degraded.

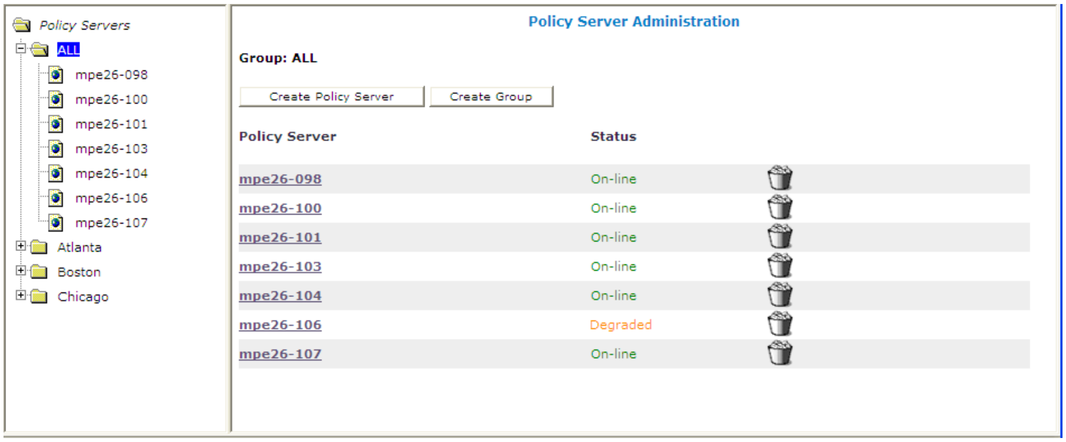

Group: ALL

Policy Server	Status
mpe26-098	On-line
mpe26-100	On-line
mpe26-101	On-line
mpe26-103	On-line
mpe26-104	On-line
mpe26-106	Degraded
mpe26-107	On-line

Figure 8: 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.

Report pages provide the following information:

- **Mode** — Shows whether data collection is currently **Active** or **Paused**, **Absolute** (displaying statistics since the last reset) or **Delta** (displaying changes in the statistics during the last 10-second refresh period).
- **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:
  - **Show Absolute/Show Deltas** — Switches between absolute mode (statistics since last reset) and delta mode (statistics since last display).
  - **Reset Counters/Reset All Counters** — Resets counters on the current page, or 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.
  - **Cancel** — Returns to previous page.

The CMP system also displays various statistics and counters related to the following:

- **Cluster** — Information about the cluster.
- **Policy Statistics** — Information about the execution of policy rules.
- **Protocol Statistics** — Information about the active network protocols.

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 **Cluster Status**:

- **On-line**: If one server, it is active; if two servers, one is active and one is standby. No server is in forced-standby mode or out of service.
- **Degraded** (two servers only): One server is active, but the other server is not available due to an ongoing database synchronization, being in forced-standby mode, being out of service, or loss of both bond interfaces.
- **Offline**: No server is active.
- **Inconsistent** (two servers only): Both servers are in the active role. This is a “split brain” error condition, and can only happen when the backplane link fails.

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 (that is, which blade is the active server). The report also lists the following server-specific information:

- **Overall** — Displays the current topology state (Active, Standby, or Out-Of-Service), 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), average value for the CPU utilization, and memory.
- **Actions** — Buttons in this section let you restart the Policy Management software on the server (**Restart**) or restart the server itself (**Reboot**).

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.

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.

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 .

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.

VoD Server Statistics

The VoD Statistics section summarizes VoD server activity within the MPE device. This information is presented as a table of summary statistics plus statistics for any selected network elements.

The summary represents the aggregated statistics for every VoD server associated with this MPE device, and includes the following:

- **Total messages in / out —** The total number of incoming and outgoing messages sent and received by this MPE device.
- **Total VOD sessions —** The total number of VoD session requests received by this MPE device, whether successful or not, since the last time the counters were reset. A session teardown does not decrement this value.
- **Success VOD sessions —** The total number of successful reserve requests (defined as a single reserve request followed by an ACK from this MPE device) since the last time the counters were reset.
- **Failure VOD sessions —** The total number of failed session requests (defined as a single reserve request from a VoD server followed by a NAK from this MPE device) since the last time the counters were reset.
- **Active VOD sessions —** The number of currently active session requests.

The Reports tab also includes a table that summarizes activity through network elements with which this MPE device is communicating. This table lets you select specific network elements to further examine detailed statistics that are specific to that network element.

You can search for specific network elements of one or more types. Select **B-RAS**, **Subscriber Group**, **Router**, **Server**, or **Wireline Gateway**, enter the name of the network element (up to 250 characters;

use “*” or “?” as wildcard characters, or leave the field blank to search for all elements of that type), and click **Search**. Information for the selected network element(s) is displayed.

Tip: To display information for all network elements, select the element class and click **Search**.

The resulting table displays the following information:

- **Network Element** — Unique identifier for this device.
- **Session count** — Number of active sessions handled by this device.
- **Session success count** — Number of successful reserve requests (defined as a single reserve request followed by an ACK from the MPE device) by this device since the MPE device was last started or the counters reset.
- **Session count HS/SD** — Number of sessions in high definition (HD) and standard definition (SD).
- **Upstream bandwidth** — Current reserved upstream bandwidth allocated for this network element.
- **Downstream bandwidth** — Current reserved downstream bandwidth allocated for this network element.

For statistics on an individual VoD server, click on its network element name.

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.

Policy Server Logs

The log files trace the activity of a Policy Management device. The system handles log file writing, compression, forwarding, and rotation automatically. 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 9: 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.
   - **Session Synchronization log** — Contains information on VoD session synchronization.

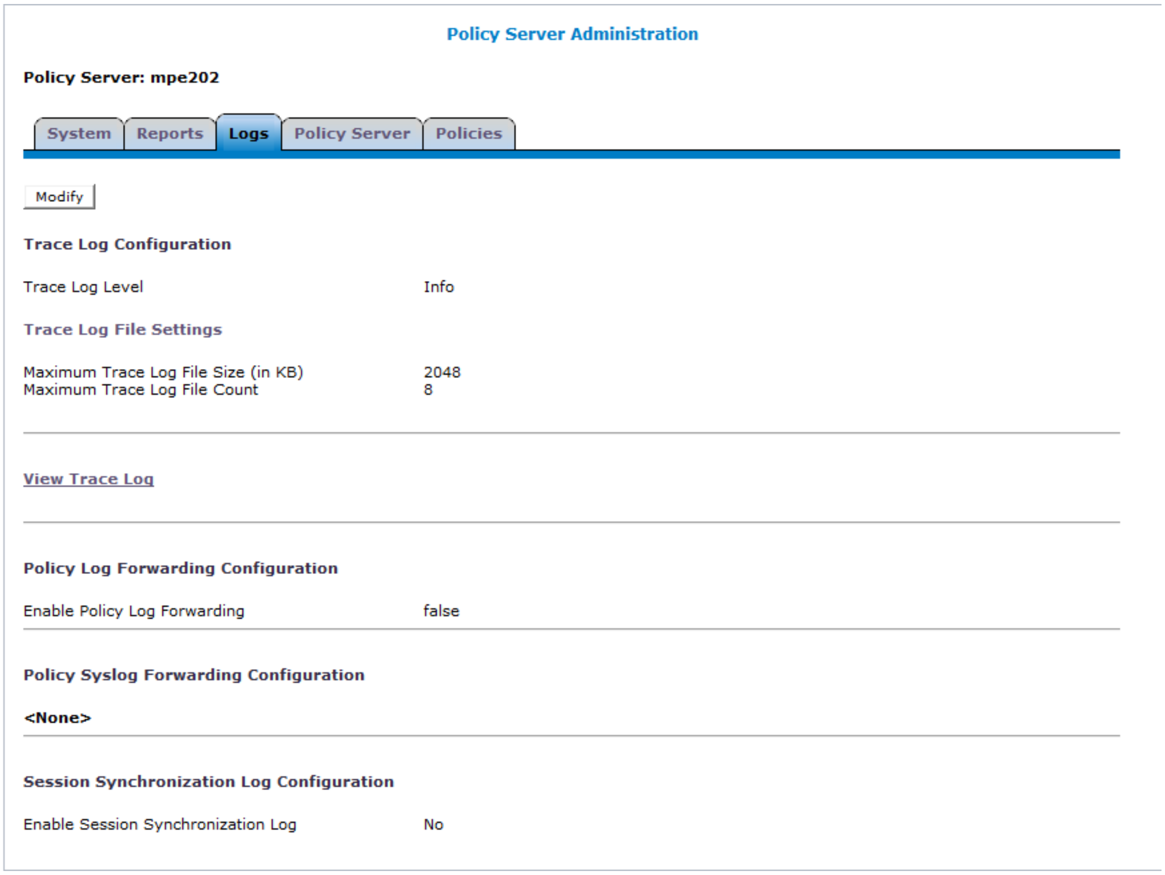

Policy Server Administration

Policy Server: mpe202

Trace Log Configuration

Trace Log Level Info

Trace Log File Settings

Maximum Trace Log File Size (in KB) 2048
Maximum Trace Log File Count 8

View Trace Log

Policy Log Forwarding Configuration

Enable Policy Log Forwarding false

Policy Syslog Forwarding Configuration

<None>

Session Synchronization Log Configuration

Enable Session Synchronization Log No

Figure 9: Policy Server Logs Tab

Viewing the Trace Log

The trace log records Policy Management application notifications, such as protocol 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.

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.

The Session Synchronization Log

The session synchronization log records VoD synchronization operations performed by the MPE device. Log files are stored in compressed format.

When the logging function is enabled, the following information is logged during every VoD synchronization operation by the MPE device:

- VoD session count list
- VoD session server ID list
- MPE device session count
- MPE device session ID list
- Sessions missing on MPE device
- Sessions missing on VoD server (with associated subscriber information as available)
- Any delete or recreate actions taken for session mismatches
- Session count and IDs included in synch response sent by MPE device
- Session count and IDs included in synch response received by MPE device

The format of a log file record is as follows:

```
timestamp|servertype|IP_addr|operation|(Count: nnnnn) session_list|
```

where:

- *timestamp* is a date/time stamp
- *servertype* is Tandberg
- *IP_addr* is the IP address of the server
- *operation* is one of the following:
  - LS — local sessions
  - RS — remote sessions (VoD server sessions)
  - LD — local deleted sessions
  - RD — remote deleted sessions
  - LR — local recreated sessions
  - SR — status response to VoD server

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** (the default) — Designates informational messages highlighting overall progress of the application.
   - **Debug** — Designates information events of lower importance.

   CAUTION

   Caution: Before changing the default logging level, consider the implications. Lowering the trace log level setting from its default value 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. Configure the maximum trace log file size (in KB).

   The system will maintain up to this number of trace log files, removing old files when it reaches this limit. The choices are 512, 1,024, 2,048, 4,096, 8,192, 16,384, or 32,678 KB. The default is 2,048 KB.

4. Configure the maximum trace log file count. The system manages rotation of log files automatically.

   The range is 2–8 files. The default is 8 files.

5. Configure the trace log 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 IPv4 address.

      CAUTION

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

      - **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** (the default) — Designates informational messages highlighting overall progress of the application.
- **Debug** — Designates information events of lower importance.

6. In the **Modify Log Forwarding Configuration** section of the page, select **Enable Policy Log Forwarding** to forward the policy log to remote locations.
7. 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 IPv4 address.

   CAUTION

   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** (the default) — Designates informational messages highlighting overall progress of the application.
   - **Debug** — Designates information events of lower importance.

8. In the **Modify Session Synchronization Log Settings** section of the page, select **Enable Session Synchronization Log** to enable the session synchronization log.
   The Number of Session Synchronization Log Files field appears. Enter the number of session synchronization log files. The system manages rotation of log files automatically. The range is 2–10 files. The default is 10 files.
9. When you finish, click **Save** (or **Cancel** to discard your changes).

The log configurations are changed.

VoD Session Flow Scenarios

The following is a general list of session flow scenarios and the action(s) of the MPE device for each.

Scenario 1

The MPE device has knowledge of a user's account, but the user’s IP address was not received by the MPE device from the B-RAS server. **MPE Action:** An internal switch, set by Tekelec Technical Support personnel only, is available to allow the MPE device to accept this request. Otherwise, if the Session

Manager sends a request for a video session to the MPE device, the request is **rejected** as the MPE device has no knowledge of the user's IP address.

Scenario 2

The MPE device has no knowledge of a user's account, but does know the user's IP address from messaging with the B-RAS server. **MPE Action:** The request is **processed normally**.

Scenario 3

A VoD request destination IP address does not match configured subnets on the B-RAS, and an entry does not exist in the COPS-PR database. **MPE Action:** The request is **rejected**. A 6203 message is written into the trace log.

Scenario 4

A VoD request destination IP address does not match configured subnets on the B-RAS, but an entry exists in the COPS-PR database. **MPE Action:** The request is **rejected**. A 6203 message is written into the trace log.

Scenario 5

A VoD request destination IP address matches configured subnets on the B-RAS server, but the entry does not exist in the COPS-PR database. **MPE Action:** An internal switch, set by Tekelec Technical Support only, is available to allow the MPE device to accept this request. Otherwise, the request is **rejected**. A 6203 message is written into the trace log.

Scenario 6

A VoD request destination IP address matches configured subnets on the B-RAS server, but the entry exists in the COPS-PR database under a different gateway router (GWR). **MPE Action:** The request is **rejected**. A 6203 message is written into the trace log.

If a request is not rejected by any of the scenarios above, it is presented to the customer-defined policy rules.

Note: During installation and transition of MPE devices, rejections due to scenarios 1 and 2 must be disabled. This is to allow for the staged startup of B-RAS devices. During the transition only a subset of B-RAS servers will be sending the MPE device COPS-PR information, but session requests will be received from all B-RAS devices within a given video hub office (VHO). Rejection mode is disabled and enabled using an internal switch set by Tekelec Technical Support personnel only.

Synchronizing a TANDBERG VoD Server

The MPE device communicates with multiple VoD servers, but will not synchronize with a TANDBERG server until it first receives an allocate resource or status request from it. The MPE device uses SSL when synchronizing with a TANDBERG server if the last allocate request was received over an HTTPS connection. The synchronization interval is expressed in minutes.

To force a synchronization with all connected TANDBERG VoD servers, select the Policy Server tab and click **Sync Tandberg VoD Server Now**. The page displays the message “The VoD Server synchronization initiated by user.”

TANDBERG synchronization proceeds as follows:

1. The MPE device builds a local list of sessions that it knows about.

2. The MPE device sends a SESSIONLIST status request to the TANDBERG server to obtain its list of sessions.
3. The MPE device compares the two lists and creates three new lists:
   a) VoD server only
   b) MPE device only
   c) Common to both
4. Ideally, all sessions are in the "common to both" list and the other two lists are empty, in which case no further action is required.
5. For each session in the VoD server only list:
   a) The MPE device checks to see whether the session was created while waiting for the session list response from the VoD server.

      If the session now exists locally, no further action is required for this session.
   b) If the session still does not exist locally, the MPE device sends a SESSIONID status request to the VoD server to request session details for the missing session.
   c) If the VoD server responds with "session not found," then no further action is required for this session.
   d) If the VoD server responds with details for the missing session, then the MPE device attempts to recreate the session by allocating resources for it.
   e) If recreating the session succeeds, no further action is required.
   f) If recreating the session fails, then the MPE device sends a "release resources" request to the TANDBERG server to tear the session down.
6. For each session in the MPE only list:
   a) The MPE device sends a SESSIONID status request to the VoD server to request session details for the missing session.
   b) If the VoD server responds with details for the session, then no further action is required.
   c) If the VoD server responds with "session not found," then the session is removed locally from the MPE device.

Synchronization Operations and Failover

After a blade failover, the MPE device loads the list of VoD servers that the previously active blade knew about and starts a synchronizer for each of them. The new active blade bases its next synchronization time on the last successful synchronization that the previously active blade completed. The synchronization failure counts are intentionally not transferred between blades; thus the newly active blade tries for another 150 minutes to communicate with a dead VoD server before removing it from the synchronization list.

Synchronizing a B-RAS Server

An ERX device needs to be synchronized with MPE devices. MPE devices support both full synchronization and the fast synch feature over the COPS-PR interface.

Fast Synch: When an ERX device connects to the MPE device, it includes information on the last policy server to which it connected as part of an OPN message. When fast synch is enabled, and the IP address sent matches the MPE device address, the MPE device sends an unsolicited Decision (DEC) message to the ERX device, which replies with just Request (REQ) and Delete (DRQ) messages instead of full state information. This makes it easier to recover from a temporary connection loss or a warm restart. (For information on configuring fast synch, see *Configuring Protocol Options on the Policy Server*.)

Full Synch: If an MPE device detects that the last policy server IP address reported by the ERX device is different from its own, it sends a Synchronization (SSQ) message, which triggers a full state synchronization. If the ERX device determines that a full state synchronization is required (for example, after a cold restart or if the threshold hold time for fast synch recovery has expired), it reports the last policy server IP address as 0.0.0.0. In this case, the MPE device sends an SSQ message.

Force Synch: To force a full synchronization with a B-RAS server:

1. From the **Policy Server** section of the navigation pane, select **Network Elements**.
   The content tree displays a list of network element groups; the initial group is **ALL**.
2. From the content tree, select the desired B-RAS server.
   The Network Element Administration page opens in the work area.
3. On the B-RAS tab, click **Force Full Synch Now**.

B-RAS synchronization proceeds as follows:

1. The MPE device closes the connection to the ERX device and waits for it to reconnect.
2. The MPE device determines whether a forced synchronization was requested, as since it was, sends an SSQ message to the server and ignores the policy server IP address reported back.
3. Once the full synchronization is complete, the forced synchronization override is reset, and fast synchronizations can resume (if the feature is enabled).

Chapter 5

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.

About Network Elements

A network element is a high-level device, server, or other entity within your network for which you would like to use an MPE device to manage Quality of Service (QoS). Examples include the following:

- Broadband remote access server (B-RAS)
- Router
- Server

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 an interface on a router, or devices that are connected directly to network elements. Typically, there is no need to define these lower-level entities, because once a network element is associated with an MPE device the lower-level devices related to that network element are discovered and associated automatically.

Create a network element profile for each device you are associating with an MPE device. After defining a network element in the CMP database, configure its protocol options. The options available depend on the network element type.

For ease of management, once you define network elements, you can combine them into network element groups.

Defining a Network Element

You must define a network element for each device associated with any of the MPE devices within the network. To define a network element:

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

   (See *Creating a Network Element Group* for information on creating network element groups.)

   The Network Element Administration page opens in the work area.
3. On the Network Element Administration page, click **Create Network Element**.
   The New Network Element page opens.
4. Enter information as appropriate for the network element:
   a) **Name** (required) — The name you assign to the network element.

      Enter up to 255 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) **Element ID** — Alternate unique ID for a network element.

Enter up to 250 characters.

e) **Description/Location** — Free-form text.

Enter up to 250 characters.

f) **Type** (required) — Select the type of network element.

The supported types are:

- **B-RAS** (the default) — Broadband Remote Access Server (with the subtypes **ERX** or **E320**)
- **Subscriber Group** — a subscriber group (for more information, see *Creating an Account*)
- **Router**
- **Server**
- **Wireline Gateway** — a gateway router (with the subtype **MX Series**)

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

5. Select one or more policy servers (MPE devices) to associate with this network element.
6. To add a network element to a network element group, select the desired group (see *Adding a Network Element to a Network Element Group*).
7. 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**. If there are more than 50 network elements, only the number of network elements is displayed, not the elements themselves.
2. On the Network Element Administration page, select the desired network element.

   If there are more than 50 network elements, the display is paginated; select the page number or search for the network element by name (see *Finding a Network Element*).
3. On the System tab, 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 **Policy Server** section of the navigation pane, select **Network Elements**.

The content tree displays a list of network element groups; the initial group is **ALL**. If there are more than 50 network elements, only the number of network elements is displayed, not the elements themselves.

2. On the Network Element Administration page, select the desired network element.

   If there are more than 50 network elements, the display is paginated; select the page number or search for the network element by name (see *Finding a Network Element*).

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 **Policy Server** section of the navigation pane, select **Network Elements**.
   The content tree displays a list of network element groups; the initial group is **ALL**.
2. From the content tree, select **ALL**.
   The Network Element Administration page opens in the work area.
3. 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, `star*`, `*pGw`, or `*-*`), click **Filter**, and select from the filtered results.

5. Click **Bulk Delete** (or **Cancel** to cancel the request).
   You are prompted: "Are you sure you want to delete all the selected Network Elements?"
6. Click **OK** to delete the network elements (or **Cancel** to cancel the request).
   The system displays the message "m Folder(s) and n Network Element(s) were deleted successfully."

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

Finding a Network Element

The Search function lets you find a specific network element within a large configuration. To search the CMP database for a specific network element:

1. From the **Policy Server** section of the navigation pane, select **Network Elements**.
   The content tree displays a list of network element groups; the initial group is **ALL**. If there are more than 50 network elements in a group, only the number of network elements is displayed, not the elements themselves.
2. From the content tree, select **ALL**.
   The Network Element Administration page opens in the work area. If there are more than 50 network elements, the display is paginated.

3. On the Network Element Administration page, click **Search**.
   The Network Element Search Criteria window opens.
4. Enter the desired search criteria. Searches are not case sensitive. You can use the wildcard characters '*' and '?'.

   - **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.
   - **ID** — The network element ID (an alternate unique ID for a network element). Enter up to 250 characters.
   - **Description** — The information pertaining to the network element that helps identify it within the network. Enter up to 250 characters.

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 elements types available depend on the operating mode in which your CMP system is configured, and may differ from the list given here.

B-RAS, Router, and Server

To configure interface information for a B-RAS, Router, or Server network element:

1. From the **Policy Server** section of the navigation pane, select **Network Elements**.
   The content tree displays a list of network element groups; the initial group is **ALL**. If there are more than 50 network elements, only the number of network elements is displayed, not the elements themselves.
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 Interfaces tab and click **Create Interface**.
   The Create Network Element Interface page opens.
4. Configure the following information:
   a) **Name** — The name assigned to the network element.

      The name can be up to 32 characters in length. The name cannot contain the character string "::" (doubled colons). Synchronization requests are processed based on the network element name.
   b) **Capacity** — The bandwidth capacity of this interface.
   c) **Description / Location** — The information pertaining to the network element that helps identify it within the network.
   d) **Links** — Specifies the links to other network elements.
5. When you finish, click **Save** (or **Cancel** to discard your changes).

The interface information is configured.

Creating Subnets

A B-RAS server can contain subnets, which can be provisioned from an operations support system (OSS) or configured manually.

To create subnets associated with a device:

1. From the **Policy Server** section of the navigation pane, select **Network Elements**.
   The content tree displays a list of network element groups; the initial group is **ALL**. If there are more than 50 network elements only the number of network elements is displayed, not the elements themselves.
2. On the Network Element Administration page, select the desired network element. If there are more than 50 network elements, the display is paginated; select the page number or search for the network element by name (see *Finding a Network Element*).
3. Select the B-RAS tab.

   Subnets are displayed in two categories:

   a) **Subnets Configured Manually** — You can add to or delete from this list.
   b) **Subnets Obtained from the OSS** — This read-only field displays subnets that were imported via the OSS interface to the CMP database.
4. Click **Modify**.
   The Modify Network Element page opens.
5. Modify the subnet list as required:

   - To add a subnet, type the address block in CIDR (Classless Inter-Domain Routing) format and click **Add**. The subnet is added to the list.
   - To delete a subnet, select it from the list and click **Delete**. The subnet is removed from the list.
6. When you finish, click **Save** (or **Cancel** to discard your changes).

The subnets are configured.

To synchronize subnet changes throughout the Policy Management network, click **Force Full Synch Now**.

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:

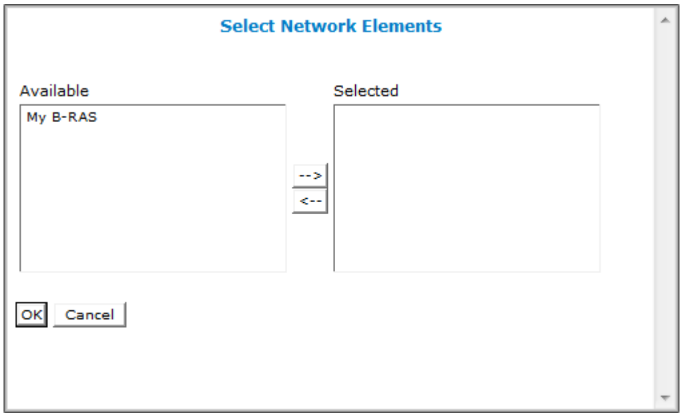

6. Select the desired network elements from the **Available** list and click **-->**.

   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 **Policy Server** section of the navigation pane, select **Network Elements**.
   The content tree displays a list of network element groups; the initial group is **ALL**.
2. From the content tree, select the **ALL** group.
   The Network Element Administration page opens in the work area.
3. 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 255 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 **Policy Server** section of the navigation pane, select **Network Elements**.
   The content tree displays a list of network element groups; the initial group is **ALL**.
2. From the content tree, select the 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 10: 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.”

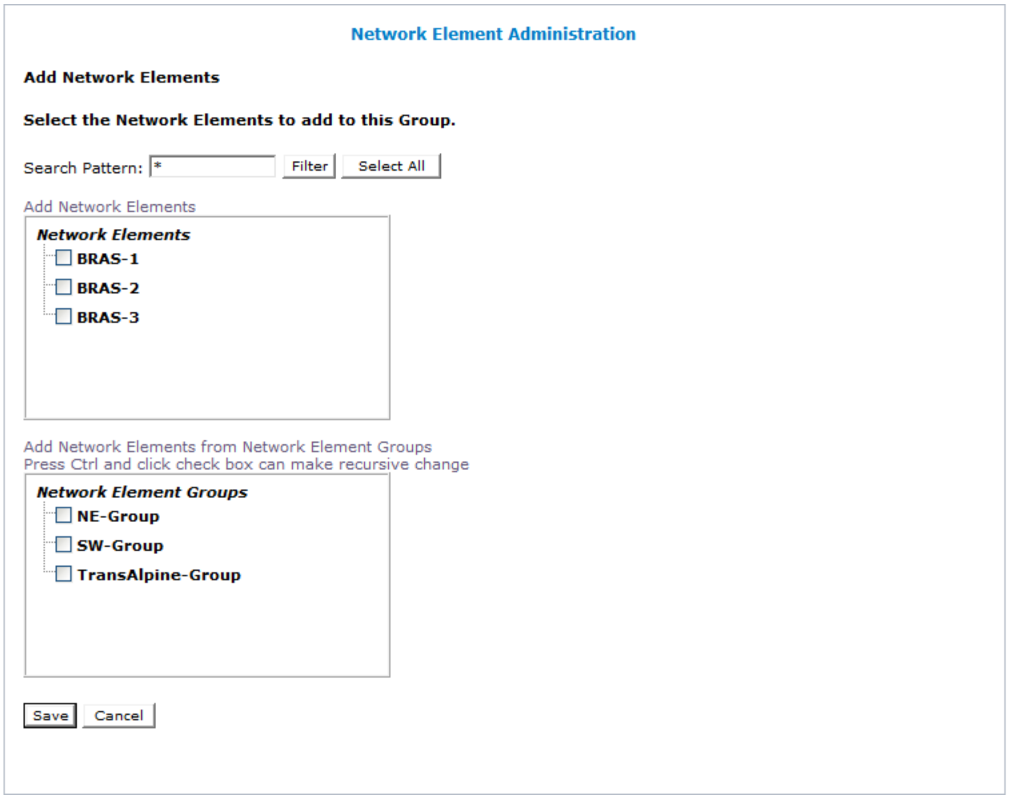

Network Element Administration

Add Network Elements

Select the Network Elements to add to this Group.

Search Pattern: *

Add Network Elements

Network Elements

- □ BRAS-1
- □ BRAS-2
- □ BRAS-3

Add Network Elements from Network Element Groups
Press Ctrl and click check box can make recursive change

Network Element Groups

- □ NE-Group
- □ SW-Group
- □ TransAlpine-Group

Figure 10: 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 **Policy Server** section of the navigation pane, select **Network Elements**.
   The content tree displays a list of network element groups; the initial group is **ALL**.
2. From the content tree, select the 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 **Policy Server** section of the navigation pane, select **Network Elements**.
   The content tree displays a list of network element groups; the initial group is **ALL**.
2. From the content tree, select the network element group or sub-group.
   The Network Element Administration page opens in the work area.
3. 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 **Policy Server** section of the navigation pane, select **Network Elements**.
   The content tree displays a list of network element groups.

2. From the content tree, select the network element group or sub-group.
   The Network Element Administration page opens in the work area, displaying the contents of the selected network element group or sub-group.
3. 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.

Importing VoD Configuration Information

The following procedure describes how to import VoD configuration information. It assumes that you have all components installed and operational, and that all subscriber information is available.

For additional information, see the *OSSI XML Interface Definitions Reference Guide*. This document describes the CMP OSSI XML interface, which lets you programmatically provision the system and retrieve operational statistics from the policy managers.

Provisioning Topology and Subscriber Data

The following subsections describe how to provision the topology and subscriber data. See *Managing Subscribers* for additional information about subscribers.

Manually within the CMP

To provision the topology and subscriber data manually:

1. Define the VoD server, including manually configured subnets and all legitimate VoD *pump* addresses.
2. Define a set of policy rules and deploy them to the policy server.
3. Define a network element group, add the VoD server into the group, and associate the group with the policy server (all of which is sent automatically to the policy server).

Using the OSSI XML Interface

To provision the topology and subscriber data using the OSSI XML interface:

1. Add network elements, network element interfaces, and network element links to represent the routers, B-RASs, and links in the network; then associate them with the same network element groups defined earlier so they are sent to the policy server automatically.

   Over time, these network resources are updated via the OSSI XML interface to add new subnets and/or delete existing subnets.

2. Add paths that refer to the server, routers, B-RASs, links, and network element Interfaces defined earlier so they are sent to the MPE device automatically.
3. Add tiers.

   Each tier should include a <TierRef> tag that refers to an associated tier so it is sent to the MPE device automatically.

4. Add accounts.

   Each account should include a <NetworkElementName> tag that refers to its associated network element, as well as a <SubscriberData> tag that defines interface information, so they are sent to the MPE device automatically.

Path Definitions

Path definitions define the sequence the data transmission elements used by video sessions originating on VoD servers and terminating on gateway routers.

At a minimum, these path definitions consist of a series of interface definitions connecting the VoD server with a specific gateway router.

The following example shows a sample base path definition:

```
<?xml version="1.0" encoding="UTF-8"?>
<XmlInterfaceRequest>
<AddPath>
  <Path>
 <Name>VOD1-GWR</Name>
 <Description/>
 <Hops>
  <Hop>
   <NeName>VoD1</NeName>
  </Hop>
  <Hop>
   <NeName>VDR1</NeName>
   <IfName>if1</IfName>
  </Hop>
  <Hop>
   <NeName>VAR1</NeName>
   <IfName>if2</IfName>
  </Hop>
  <Hop>
   <NeName>GWR</NeName>
  </Hop>
 </Hops>
</Path>
</AddPath>
</XmlInterfaceRequest>
```

If resource tracking and policy rule execution is desired against the router device (rather the interface on a router), you can use the router definition itself in the path definition. In the following example, the entries are added to track resources at the VDR1 and VAR1 routers:

```
<?xml version="1.0" encoding="UTF-8"?>
<XmlInterfaceRequest>
<AddPath>
 <Path>
 <Name>VOD1-GWR</Name>
 <Description/>
 <Hops>
  <Hop>
   <NeName>VoD1</NeName>
  </Hop>
  <Hop>
  <NeName>VDR1</NeName>
  </Hop>
  <Hop>
   <NeName>VDR1</NeName>
   <IfName>if1</IfName>
  </Hop>
  <Hop>
   <NeName>VAR1</NeName>
  </Hop>
  <Hop>
   <NeName>VAR1</NeName>
   <IfName>if2</IfName>
  </Hop>
  <Hop>
   <NeName>GWR</NeName>
  </Hop>
 </Hops>
</Path>
</AddPath>
</XmlInterfaceRequest>
```

Operational statistics are available for the routers defined in the path.

Importing a Large Number of Subscribers

The following procedure is recommended when using the OSSI XML interface to import a large number of subscribers:

1. Break up the entire collection of subscribers into subsets of 10,000.
2. Within the XML, include all 10,000 accounts in a single <AddAccount> tag (as opposed to using 10,000 separate <AddAccount> tags).
3. Use a separate HTTP POST command to push each subset to the CMP database.

   Note: Do not specify a *DistributeImmediately* attribute of `no` in these commands. Either specify `yes` or do not include the attribute at all (the default value is `yes`). The subscriber data is distributed immediately from the CMP database to the MPE devices. In addition, the CMP system is configured to allow only post file sizes of up to 20MB.

Chapter

6

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 for a TANDBERG Server

An application profile is associated with a request received from a TANDBERG server based on the application name in the request. (If the application name is absent, the server IP address in the request is used.) This application profile can be used in policy conditions.

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 (for example, **OpenStream**). Most TANDBERG allocation requests have ApplicationName = “OpenStream” and include MediaType = “VIDEO,” and the MPE device concatenates the fields with a period (.), so the application name is usually “OpenStream.VIDEO.” You must define an application with this name, or else all allocation requests from TANDBERG servers will generate “416 Invalid Application Name” errors.
      - **Description/Location** (optional) — Any information that helps identify the application.
      - **Connection IP Address(es)** (optional) — Enter the Connection Manager IP address(es) 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**. It is not necessary to include an IP address if the allocation request includes the application name.
      - **SD/HD Threshold** — Enter the bitrate threshold between standard definition (SD) and high definition (HD). Bitrate requests below the threshold are assigned the service type “SD,”

and bitrate requests above the threshold are asssigned the service type "HD." Type a numeric value in the range 0–2147483647 (2^{31} –1).

The traffic classes "SD" and HD" are available as conditions in the policy wizard (see *Conditions Available for Writing Policy Rules*).

b) **Policy Servers associated with this Application**: Select a policy server (MPE device) to associate it with this network element.

5. When you finish, click **Save** (or **Cancel** to discard your changes).
   The TANDBERG 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 TANDBERG Server* 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 7

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 downstream sessions

set policy context property SessionClass to large
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 (*downstream*)

The first action, **set policy context property ...**, specifies two parameters that represent the name and value of a policy context property 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 for downstream bandwidth
     where the requested guaranteed downstream bandwidth is greater than 2M bps
.
.
.
```

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 current time is between 18:00 and 23:59 using USER LOCAL TIME
     where today is a weekend day using USER LOCAL TIME
.
.
.
```

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
          where the request is for downstream bandwidth
          where the requested guaranteed downstream bandwidth is greater than 2M
 bps
     where the session is an application session

continue processing message
```

The policy wizard validates condition trees.

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.

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.

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.

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:

```
where the device type is B-RAS
   and where the device will be handling greater than 95 percent of downstream
capacity

   reject 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 Conditions page opens.
6. 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**, 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.
   - 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:

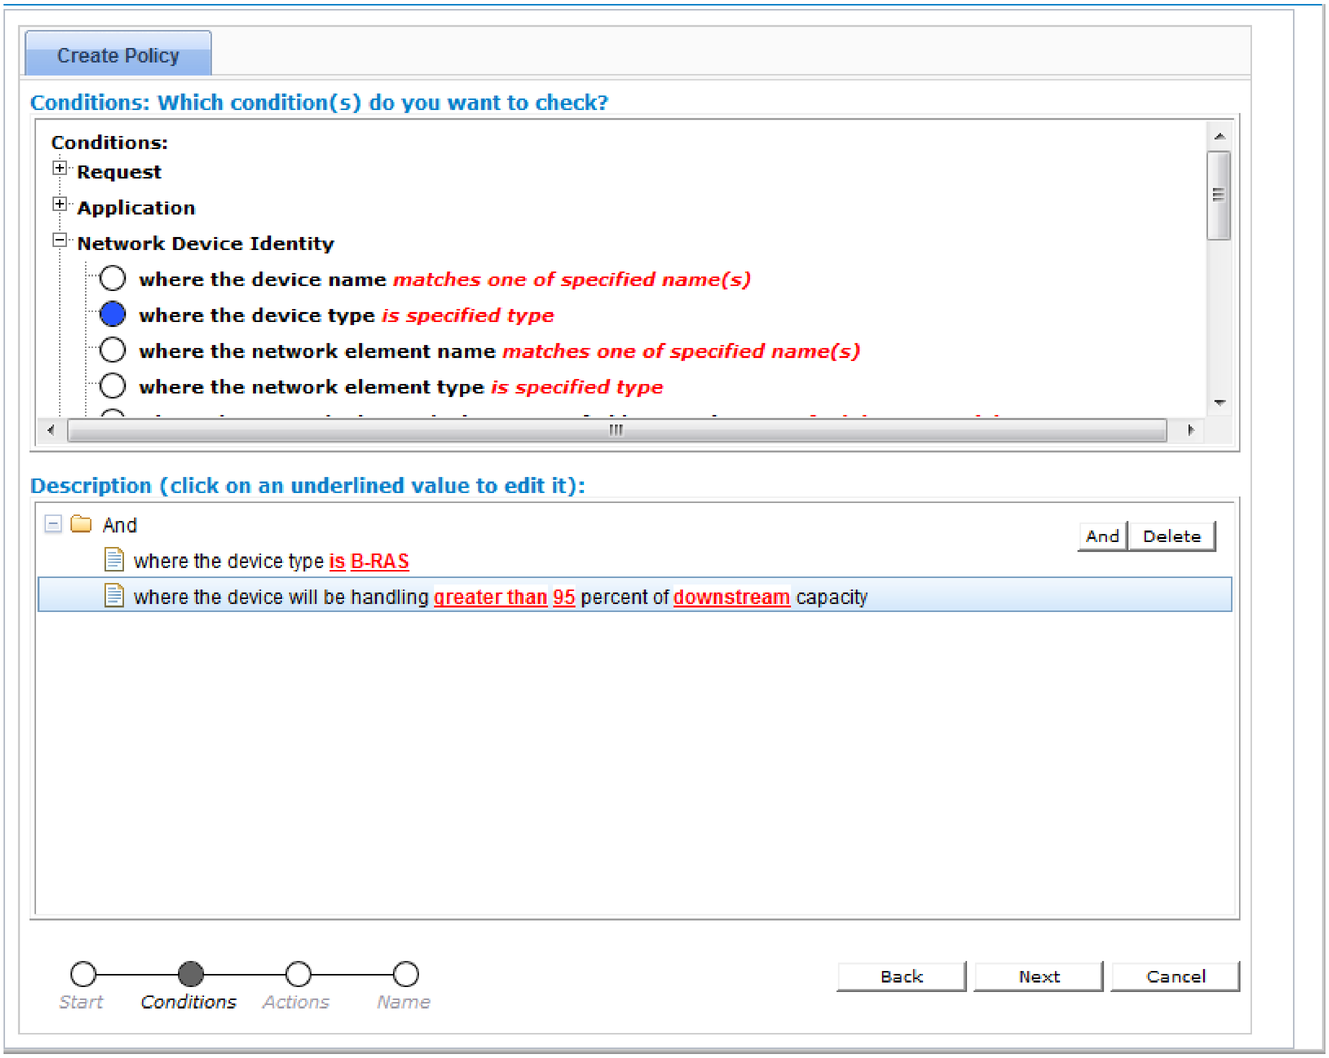

Conditions: Which condition(s) do you want to check?

Conditions:
- Request
- Application
- Network Device Identity
  - ○ where the device name *matches one of specified name(s)*
  - ● where the device type *is specified type*
  - ○ where the network element name *matches one of specified name(s)*
  - ○ where the network element type *is specified type*

Description (click on an underlined value to edit it):

- And
  - where the device type is B-RAS
  - where the device will be handling greater than 95 percent of downstream capacity

7. 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; for example:

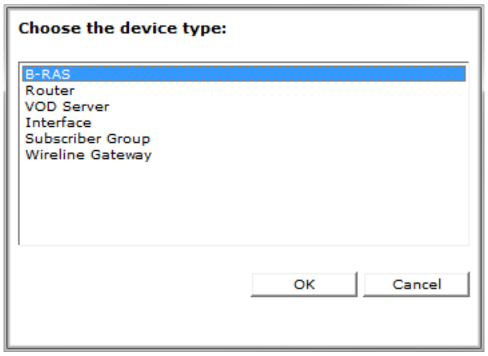

Choose the device type:

B-RAS
Router
VOD Server
Interface
Subscriber Group
Wireline Gateway

   - Enter a value (such as a traffic bit rate or percentage); for example:

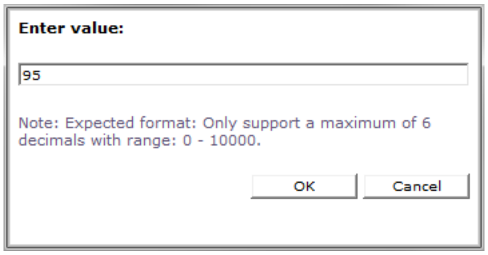

Enter value:

95

Note: Expected format: Only support a maximum of 6 decimals with range: 0 - 10000.

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

8. When you finish defining policy conditions, click **Next** (or **Cancel** to close the wizard without saving the policy).
   The Actions page opens.
9. 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:

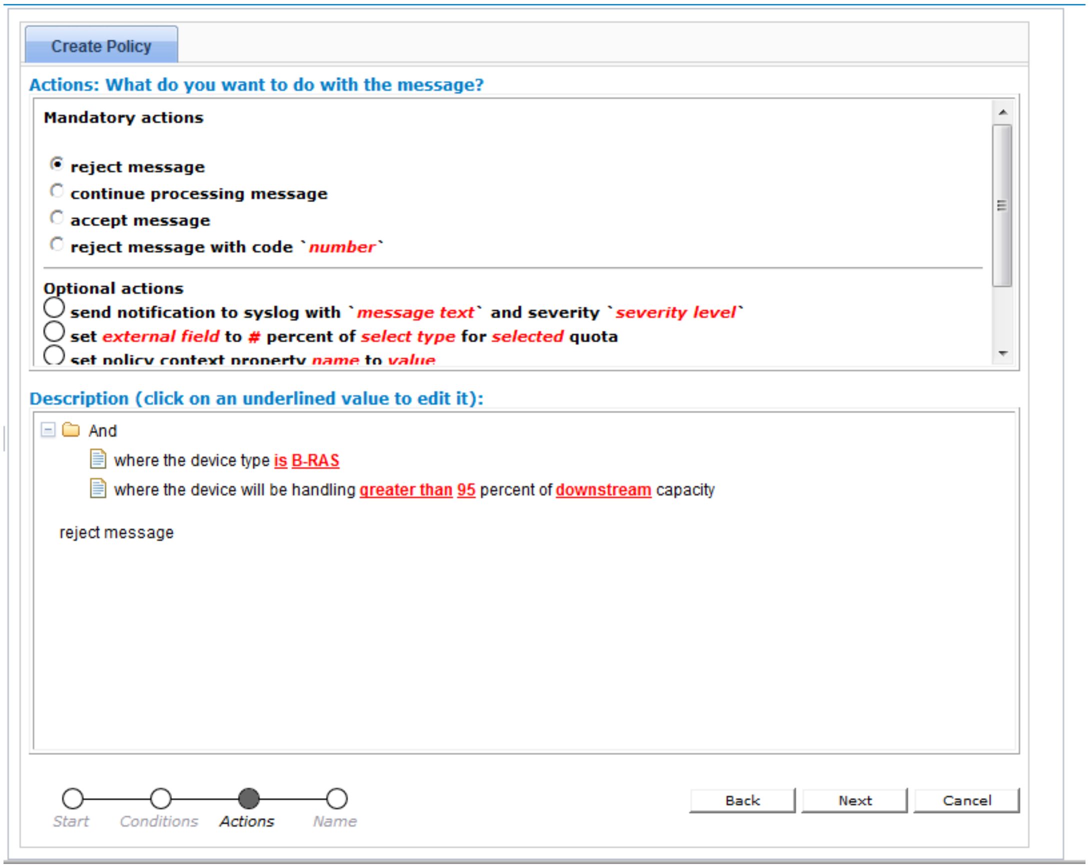

Actions: What do you want to do with the message?

Mandatory actions

● reject message
○ continue processing message
○ accept message
○ reject message with code `number`

Optional actions

○ send notification to syslog with `message text` and severity `severity level`
○ set external field to # percent of select type for selected quota
○ set policy context property name to value

Description (click on an underlined value to edit it):

And
where the device type is B-RAS
where the device will be handling greater than 95 percent of downstream capacity

reject message

- 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 (✕)

10. When you finish, click **Next** (or **Cancel** to close the wizard without saving the policy).
    The Name page opens.
11. Assign a unique name (where uniqueness is not case sensitive) to the new policy rule (the name cannot contain the characters < > & ' " = % \ ;). For example:

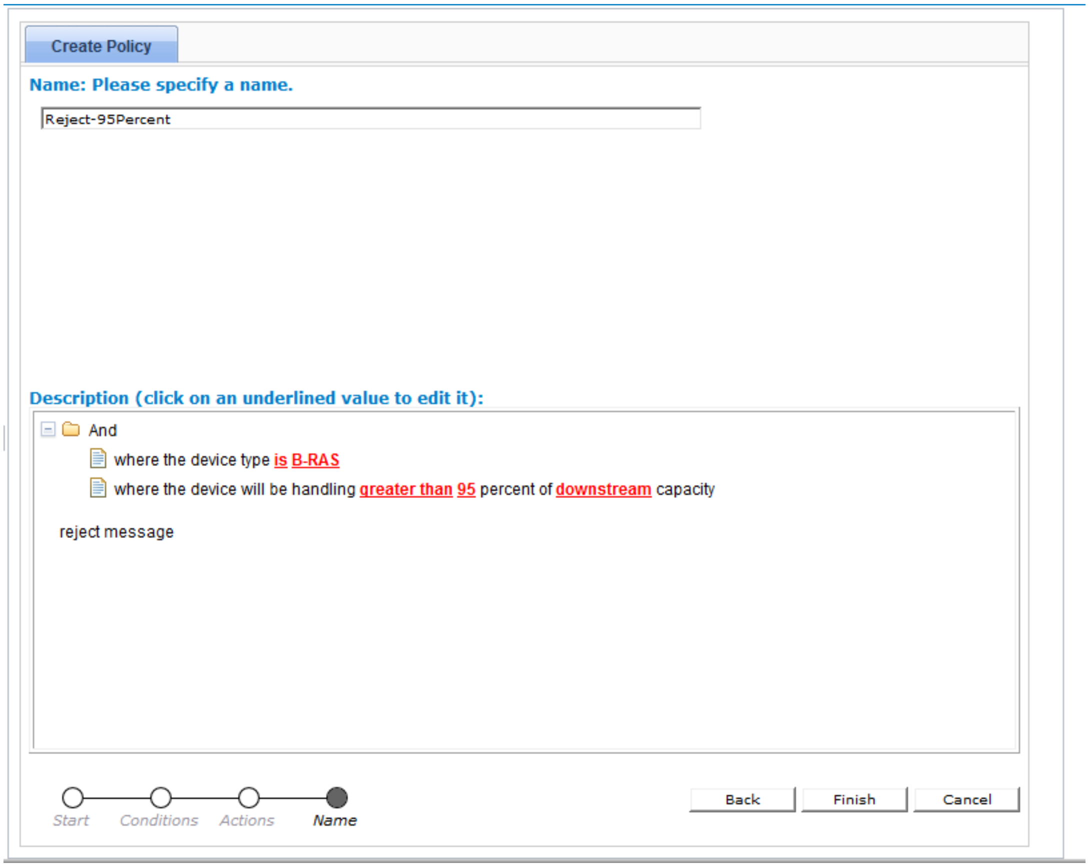

Name: Please specify a name.

Reject-95Percent

Description (click on an underlined value to edit it):

And
where the device type is B-RAS
where the device will be handling greater than 95 percent of downstream capacity

reject message

Note: The name cannot contain the following characters: `< > \ ; & ' " =`

12. 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. See *Managing Policy Rules*.

Modes Within the Policy Wizard

The behavior of the policy wizard varies depending on the mode in which your CMP system 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:

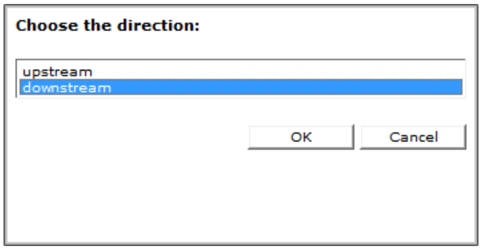

Choose the direction:

upstream
downstream

Table 4: 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 4: 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.

Parameter Type	Default Text	Description of Values
class-of-service	*specified class of*	One (or more) of the following: - **Standard Definition** - **High Definition**
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 and 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** - **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**

Parameter Type	Default Text	Description of Values
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 5: Policy Condition Categories*.

Table 5: 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.
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.

Category	Description
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 request is for *downstream* 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 *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:

- **Standard Definition**
- **High Definition**

Description

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

where the requested *downstream* bandwidth is *greater than #* and *less than #* bps

Syntax
where the requested *qos-direction* bandwidth is *operator-greater bandwidth* and *operator-less bandwidth* bps

Parameters
qos-direction

See common parameters.

operator-greater

See common parameters.

bandwidth

See common parameters.

operator-less

See common parameters.

Description

Selects protocol messages based on the direction and amount of bandwidth being requested, relative to a numeric value range.

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

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

Parameters
qos-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.

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 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 #* and *less than #* bps *specified class of* bandwidth

Syntax

where the application will be using *operator-greater bandwidth* and *operator-less bandwidth* bps *class-of-service* bandwidth

Parameters

operator-greater

See common parameters.

bandwidth

See common parameters.

operator-less

See common parameters.

class-of-service

One of the following:

- **Standard Definition**
- **High Definition**

Description

Triggers a policy based on the total amount of bandwidth used by the associated application as it relates to a defined range. This can be further qualified by the allocation class of service 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 #* and *less than # downstream* sessions

Syntax

where the application will be using *operator-greater number* and *operator-less number qos-direction* sessions

Parameters

operator-greater

See common parameters.

number

See common parameters.

operator-less

See common parameters.

qos-direction

See common parameters.

Description

Triggers a policy based on the total number of sessions used by the associated application as it relates to a defined range and direction. The total represents the number of sessions that are allocated if the current request is approved.

where the application will be using *greater than #* and *less than # specified class of* sessions

Syntax

where the application will be using *operator-greater number* and *operator-less number class-of-service* sessions

Parameters

operator-greater

See common parameters.

number

See common parameters.

operator-less

See common parameters.

class-of-service

One of the following:

- **Standard Definition**
- **High Definition**

Description

Triggers a policy based on the total number of sessions used by the associated application as it relates to a defined range. The total represents the number of sessions that are allocated if the current request is approved.

where the application will be using *greater than # specified class of* sessions

Syntax
where the application will be using *operator-greater number class-of-service* sessions

Parameters
operator-greater

See common parameters.

number

See common parameters.

class-of-service

One of the following:

- **Standard Definition**
- **High Definition**

Description

Triggers a policy based on the total number of sessions used by the associated application as it relates to a defined threshold. The total represents the number of sessions that are allocated if the current request is approved.

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 #* bps of *specified class* of bandwidth

Syntax

where the application will be using *operator-greater bandwidth* bps of *class-of-service* bandwidth

Parameters

operator-greater

See common parameters.

bandwidth

See common parameters.

class-of-service

One of the following:

- **Standard Definition**
- **High Definition**

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 the allocation class of service 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 #* sessions

Syntax

where the application will be using *operator-greater number* sessions

Parameters

operator-greater

See common parameters.

number

See common parameters.

Description

Triggers a policy based on the total number of sessions used by the associated application as it relates to a defined threshold. The total represents the number of sessions that are 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 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:

- **B-RAS**
- **Router**
- **VOD Server**
- **Interface**

Description

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

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 of the following:

- **B-RAS**
- **Router**
- **VOD Server**
- **Subscriber Group**
- **Wireline Gateway**

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 network element's description field is equal to *specified description(s)*

Syntax
where the network element's description field is equal to *string*

Parameters

string

See common parameters.

Description

Triggers a policy that is only evaluated if the Description field of the network element matches the specified string.

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 #* and *less than #* bps of *specified class of* sessions

Syntax

where the device will be handling *operator-greater number* and *operator-less number* bps of *class-of-service* sessions

Parameters

operator-greater

See common parameters.

number

See common parameters.

operator-less

See common parameters.

class-of-service

One or more of the following:

- **Standard Definition**
- **High Definition**

Description

Triggers a policy based on the total number of sessions used by the device as it relates to a defined range. This can be further qualified by the class of service of the sessions. The total represents the number of sessions that are allocated if the current request is approved.

where the device will be handling *greater than #* and *less than #* bps of *specified class of* bandwidth

Syntax

where the device will be handling *operator-greater bandwidth* and *operator-less bandwidth* bps of *class-of-service* bandwidth

Parameters

operator-greater

See common parameters.

bandwidth

See common parameters.

operator-less

See common parameters.

class-of-service

One or more of the following:

- **Standard Definition**
- **High Definition**

Description

Triggers a policy based on the total amount of bandwidth used by the current device as it relates to a defined range. This can be further qualified by the class of service 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 #* and *less than #* percent of *downstream* capacity

Syntax

where the device will be handling *operator-greater bandwidth* and *operator-less bandwidth* percent of *qos-direction* bandwidth

Parameters

operator-greater

See common parameters.

bandwidth

See common parameters.

operator-less

See common parameters.

qos-direction

See common parameters.

Description

Triggers a policy based on the percentage of capacity used by the current device as it relates to a defined range. This can be further qualified by the direction 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 #* and *less than # specified class of* sessions

Syntax

where the device will be handling *operator-greater number* and *operator-less number class-of-service* sessions

Parameters

operator-greater

See common parameters.

number

See common parameters.

operator-less

See common parameters.

class-of-service

One or more of the following:

- **Standard Definition**
- **High Definition**

Description

Triggers a policy based on the total number of sessions used by the device as it relates to a defined range. This can be further qualified by the class of service of the sessions. The total represents the number of sessions that are allocated if the current request is approved.

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

Syntax

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

Parameters

operator-greater

See common parameters.

bandwidth

See common parameters.

qos-direction

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 and direction. The total represents the bandwidth that is allocated if the current request is approved.

where the device will be handling *greater than #* bps of *specified class of* bandwidth

Syntax

where the device will be handling *operator-greater bandwidth* bps of *class-of-service* bandwidth

Parameters

operator-greater

See common parameters.

bandwidth

See common parameters.

class-of-service

One or more of the following:

- **Standard Definition**
- **High Definition**

Description

Triggers a policy based on the total amount of bandwidth used by the current device as it relates to a defined threshold. This is further qualified by the class of service 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 # downstream* sessions

Syntax

where the device will be handling *operator-greater number qos_direction* sessions

Parameters

operator-greater

See common parameters.

number

See common parameters.

qos-direction

See common parameters.

Description

Triggers a policy based on the total number of sessions used by the device as it relates to a defined direction and threshold. The total represents the number of sessions that are allocated if the current request is approved.

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

Syntax

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

Parameters

operator

See common parameters.

percent

See common parameters.

qos-direction

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 the direction 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 # specified class of* sessions

Syntax

where the device will be handling *operator-greater number class-of-service* sessions

Parameters

operator-greater

See common parameters.

number

See common parameters.

class-of-service

One or more of the following:

- **Standard Definition**
- **High Definition**

Description

Triggers a policy based on the total number of sessions used by the device as it relates to a defined threshold. The total represents the number of sessions 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). See *Managing Subscribers* for information on accounts.

where the account will be handling *greater than #* and *less than #* percent of *downstream* limit

Syntax

where the account will be handling *operator-greater percent* and *operator-less percent* percent of *qos-direction* limit

Parameters

operator-greater

See common parameters.

operator-less

See common parameters.

percent

See common parameters.

qos-direction

See common parameters.

Description

Triggers a policy based on the percent of the bandwidth limit used by the account related to a defined range. This can be further qualified by the direction of the bandwidth. The total is the bandwidth allocated if the request is approved. See *Managing Subscribers* for information on accounts.

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

Syntax

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

Parameters

operator

See common parameters.

percent

See common parameters.

qos-direction

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 the direction of the bandwidth. The total is the bandwidth allocated if the request is approved. See *Managing Subscribers* for information on accounts.

where the account will be using *greater than #* and *less than #* bps *downstream* bandwidth

Syntax

where the account will be using *operator-greater bandwidth* and *operator-less bandwidth* bps *qos-direction* bandwidth

Parameters

operator-greater

See common parameters.

bandwidth

See common parameters.

operator-less

See common parameters.

qos-direction

See common parameters.

Description

Triggers a policy based on the total amount of bandwidth used by the account as it relates to a defined range. This can be further qualified by the direction of the bandwidth. The total is the bandwidth allocated if the request is approved. See *Managing Subscribers* for information on accounts.

where the account will be using *greater than #* and *less than # downstream* sessions

Syntax
where the account will be handling *operator-greater number* and *operator-less number qos-direction* sessions

Parameters
operator-greater

See common parameters.

number

See common parameters.

operator-less

See common parameters.

qos-direction

See common parameters.

Description

Triggers a policy based on the number of sessions for a specific direction of service used by the account as it relates to a defined range. The total is the number of sessions allocated if the request is approved. See *Managing Subscribers* for information on accounts.

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

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

Parameters
operator

See common parameters.

bandwidth

See common parameters.

qos-direction

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 the direction of the bandwidth. The total is the bandwidth allocated if the request is approved. See *Managing Subscribers* for information on accounts.

where the account will be using *greater than #* bps of *specified class of* bandwidth

Syntax
where the account will be using *operator-greater number* bps of *class-of-service* bandwidth

Parameters

operator-greater

See common parameters.

number

See common parameters.

class-of-service

One of the following:

- **Standard Definition**
- **High Definition**

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 the class of service of the bandwidth. The total is the amount of bandwidth allocated if the request is approved.

where the account will be using *greater than # downstream* sessions

Syntax

where the account will be using *operator number qos-direction* sessions

Parameters

operator

See common parameters.

number

See common parameters.

qos-direction

See common parameters.

Description

Triggers a policy based on the total number of sessions used by the associated account as it relates to a defined threshold. This can be further qualified by the direction of the sessions. The total represents the number of sessions that are allocated if the current request is approved. See *Managing Subscribers* for information on accounts.

where the account will be using *greater than # specified class of* sessions

Syntax

where the account will be using *operator-greater number class-of-service* sessions

Parameters

operator-greater

See common parameters.

number

See common parameters.

class-of-service

One of the following:

- **Standard Definition**
- **High Definition**

Description

Triggers a policy based on the total number of sessions for specific classes of service used by the account as it relates to a defined threshold. This can be further qualified by the class of the sessions. The total is the number of sessions allocated if the request is approved. See *Managing Subscribers* for information on accounts.

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. See *Managing Subscribers* for information on tiers.

where the tier will be handling *greater than #* and *less than # specified class of* sessions

Syntax

where the tier will be handling *operator-greater number* and *operator-less number class-of-service* session

Parameters

operator-greater

See common parameters.

number

See common parameters.

operator-less

See common parameters.

class-of-service

One of the following:

- **Standard Definition**
- **High Definition**

Description

Triggers a policy based on the total number of sessions for a specific class of service used by the tier as it relates to a defined range. The total is the number of sessions allocated if the request is approved. See *Managing Subscribers* for information on tiers.

where the tier will be handling *greater than # specified class of* sessions

Syntax
where the tier will be handling *operator-greater number class-of-service* sessions

Parameters
operator-greater

See common parameters.

number

See common parameters.

class-of-service

One of the following:

- **Standard Definition**
- **High Definition**

Description

Triggers a policy based on the total number of sessions for a specific class of service used by the tier as it relates to a defined threshold. The total is the number of sessions allocated if the request is approved. See *Managing Subscribers* for information on tiers.

where the tier will be using *greater than #* bps of *specified class of* bandwidth

Syntax
where the tier will be using *operator-greater number* and *operator-less number* bps of *class-of-service* bandwidth

Parameters
operator-greater

See common parameters.

number

See common parameters.

operator-less

See common parameters.

class-of-service

One of the following:

- **Standard Definition**
- **High Definition**

Description

Triggers a policy based on the total amount of bandwidth used by the tier as it relates to a defined threshold. This is further qualified by the class of service of the bandwidth. The total is the amount of bandwidth allocated if the request is approved. See *Managing Subscribers* for information on tiers.

where the tier will be using *greater than* # and *less than* # bps of *specified class of* bandwidth

Syntax

where the tier will be using *operator-greater number* and *operator-less number* bps of *class-of-service* bandwidth

Parameters

operator-greater

See common parameters.

number

See common parameters.

operator-less

See common parameters.

class-of-service

One of the following:

- **Standard Definition**
- **High Definition**

Description

Triggers a policy based on the total amount of bandwidth used by the tier as it relates to a defined range. This can be further qualified by the class of service of the bandwidth. The total is the amount of bandwidth allocated if the request is approved. See *Managing Subscribers* for information on 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. See *Managing Subscribers* for information on tiers.

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. See *Managing Subscribers* for information on tiers.

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 user-defined name/value string pairs that can be created from policy actions and evaluated from policy conditions. By using policy context properties, one policy can influence the execution of other policies. Policy context properties exist across multiple policy executions on the same request, but are not persistent across requests.

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 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.

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.

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.

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.

reject message with code `*number*`

Description

After executing the current policy rule, the generated code is propagated back to the VoD server. This code is an integer from 1–2000000000.

Optional Policy-Processing Actions

The following optional policy-processing actions are available.

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 policy context property.

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.

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.

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 or an SPR. 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.

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.

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 6: 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, wireline) 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 6: 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, in the format *MMM*[*M*]/*dd* [/*yyyy*], where *MMM* is “Jan,” “Feb,” “Mar,” ..., or “Dec”, and *MM* is “01,” “02,” “03,” ..., or “12.”	--
{Time}	The time when the policy rule is executed, in the format *hh*:*mm*:*ss*.*SSS*.	--

Variable Name	Description	Modes, Protocols, Device Type
{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
{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	Any

Variable Name	Description	Modes, Protocols, Device Type
	so forth), then this will have a different value.	
{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
{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

Policy Rule Examples

The following are examples of policy rules.

Orange_GWR

```
where the device type is Interface
  and where the network element type is B-RAS
  and where the device will be handling greater than 200M bps
downstream bandwidth
  and where the device will be handling less than 2500M bps
```

```
downstream bandwidth
  and where the device will be handling greater than 25 percent of
 downstream capacity

send notification with `LC002, {Element.Name},{Interface.Name},
{Interface.DownstreamCapacity},{Device.FlowCount},{Account.AccountId},
 {AccountTier.Name},{Bandwidth},{Account.EndpointId}` and severity
 `Warning`
continue processing message
```

Red_GWR

```
where the device type is Interface
  and where the network element type is B-RAS
  and where the device will be handling greater than 2500M bps
downstream bandwidth
  and where the device will be handling greater than 50 percent of
 downstream capacity

send notification with `LC002, {Element.Name},{Interface.Name},
{Interface.DownstreamCapacity},{Device.FlowCount},{Account.AccountId},
 {AccountTier.Name},{Bandwidth},{Account.EndpointId}` and severity
 `Alert`
reject message
```

Reject_Subscriber_Session

```
where the account will be handling greater than 100 percent of
downstream limit

send notification with `SR002, {Account.AccountId},
{AccountTier.Name}, {Bandwidth}, {Account.EndpointId}` and severity
 `Alert`
reject message
```

Policy rule variable

The following example illustrates the use of a policy rule variable.

```
where the device type is Interface
  and where the policy context property donotreject does not exist
  and where the device will be handling greater than 70 percent of
 downstream capacity

reject message
```

Chapter 8

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**.
2. From the content tree, select the desired policy. The policy is displayed. *Figure 11: Sample Policy Description* shows an example.

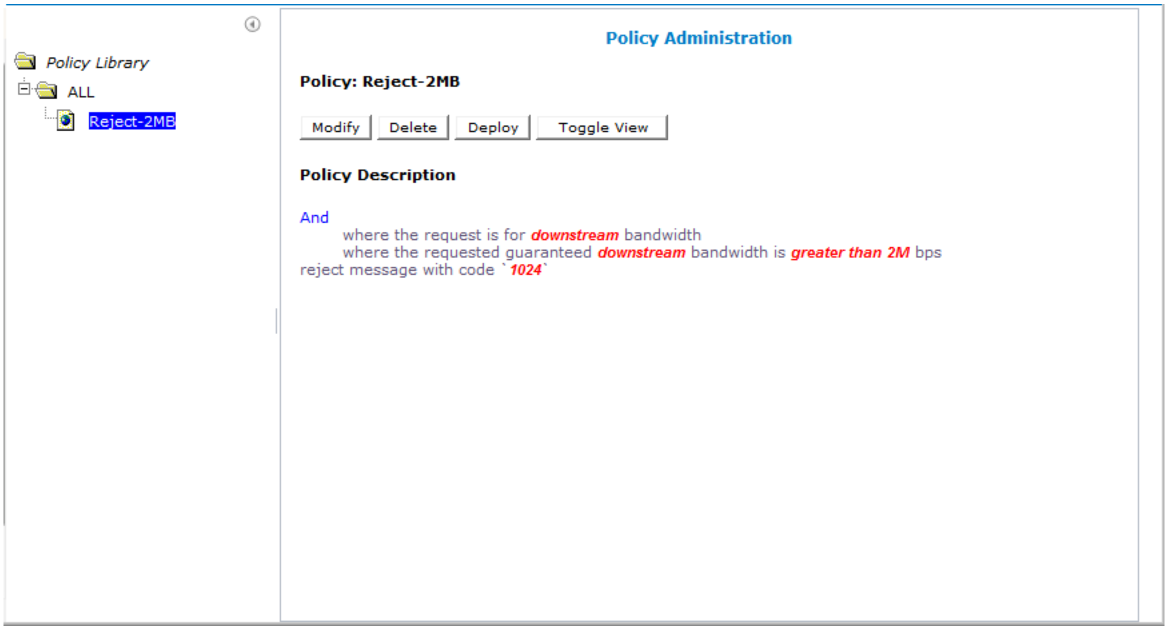

Policy: Reject-2MB

Policy Description

```
And
        where the request is for downstream bandwidth
        where the requested guaranteed downstream bandwidth is greater than 2M bps
reject message with code `1024`
```

Figure 11: 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**.

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 12: 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.

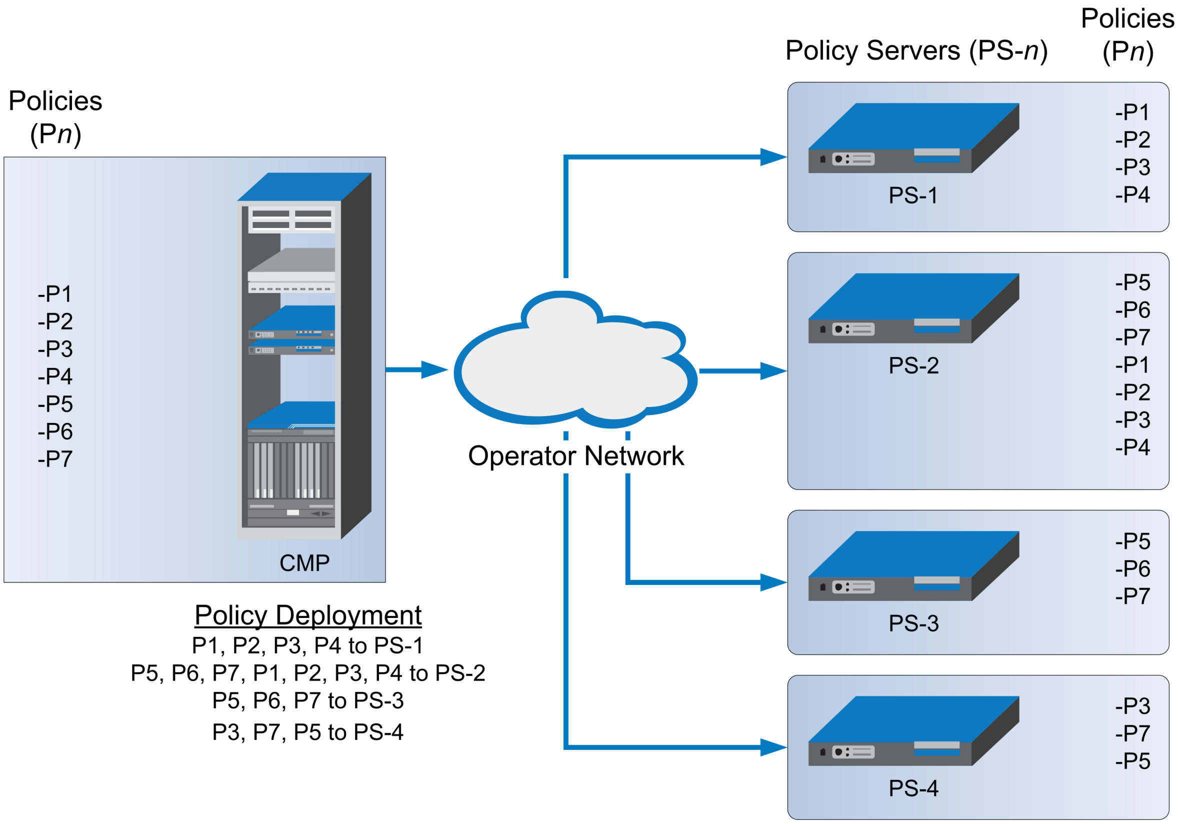

Figure 12: Policy Deployment

Figure 13: 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.

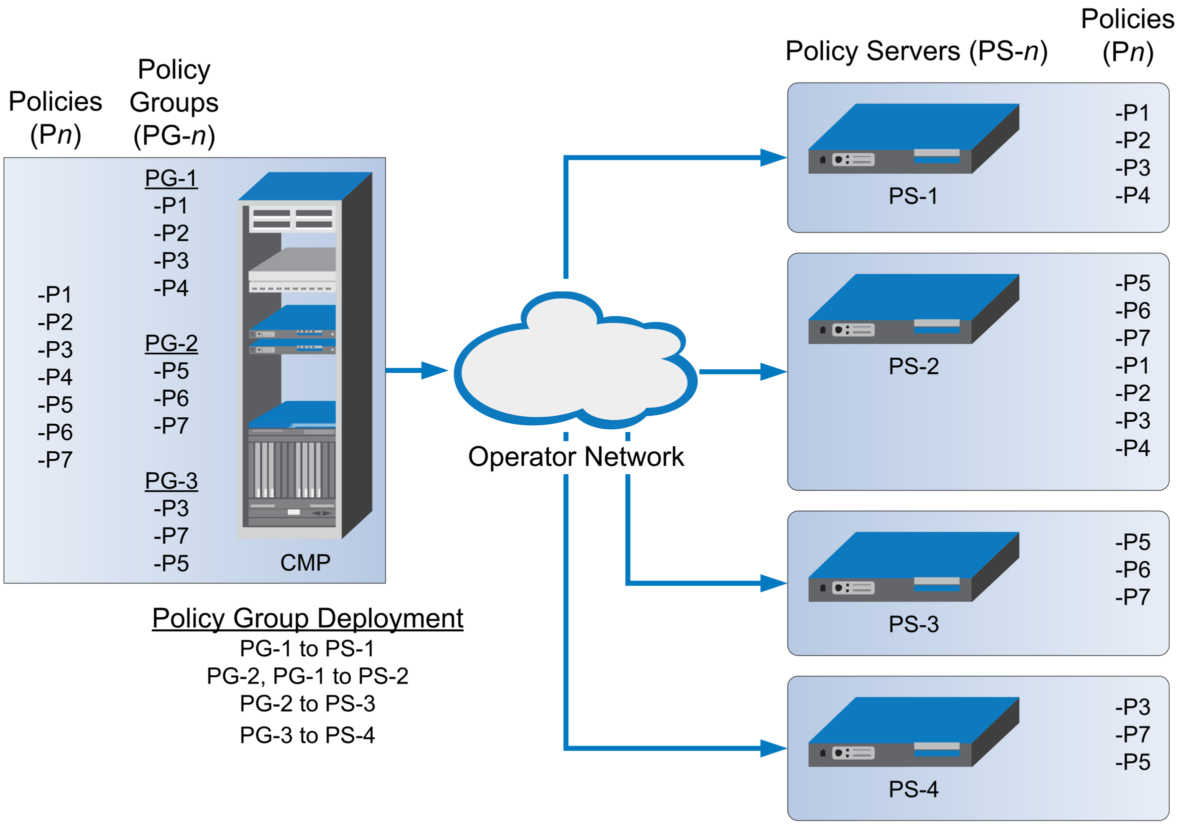

Figure 13: Policy Group Deployment

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

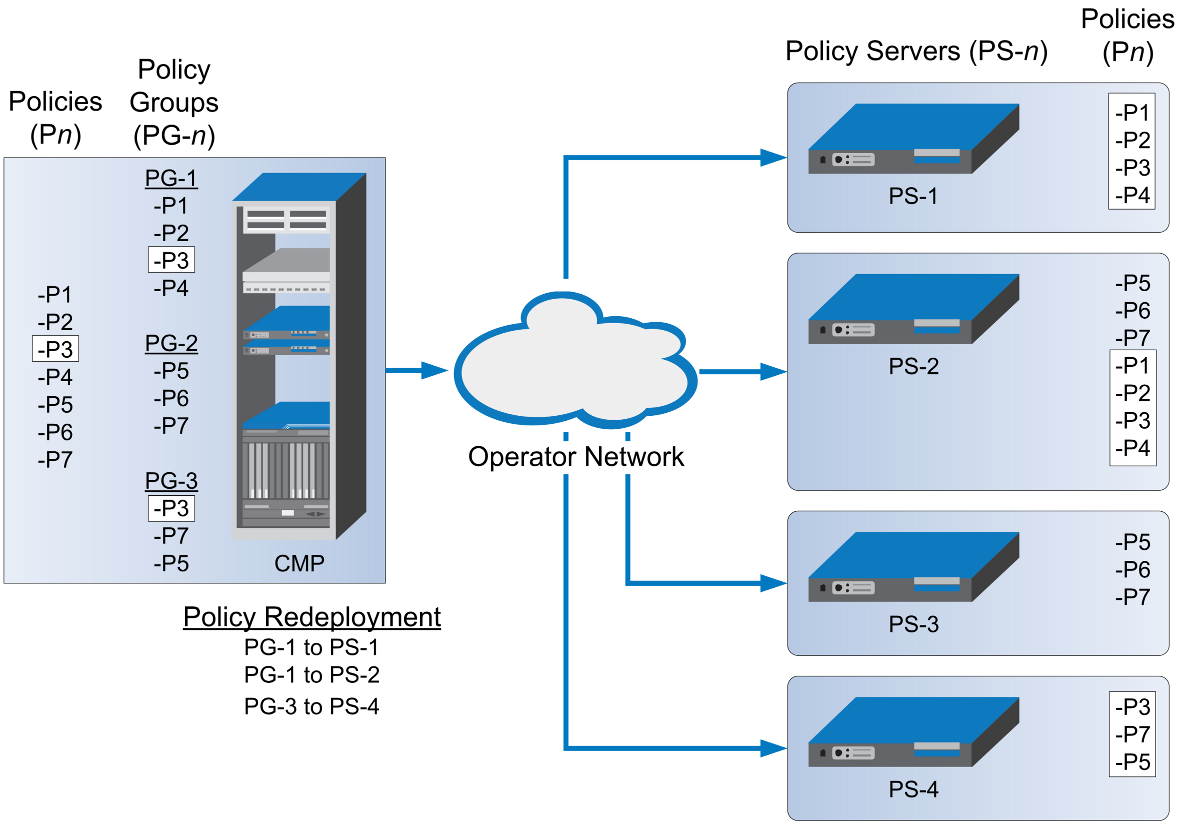

Figure 14: Policy Redeployment

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 (*Figure 15: Create New Template Window*).
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.

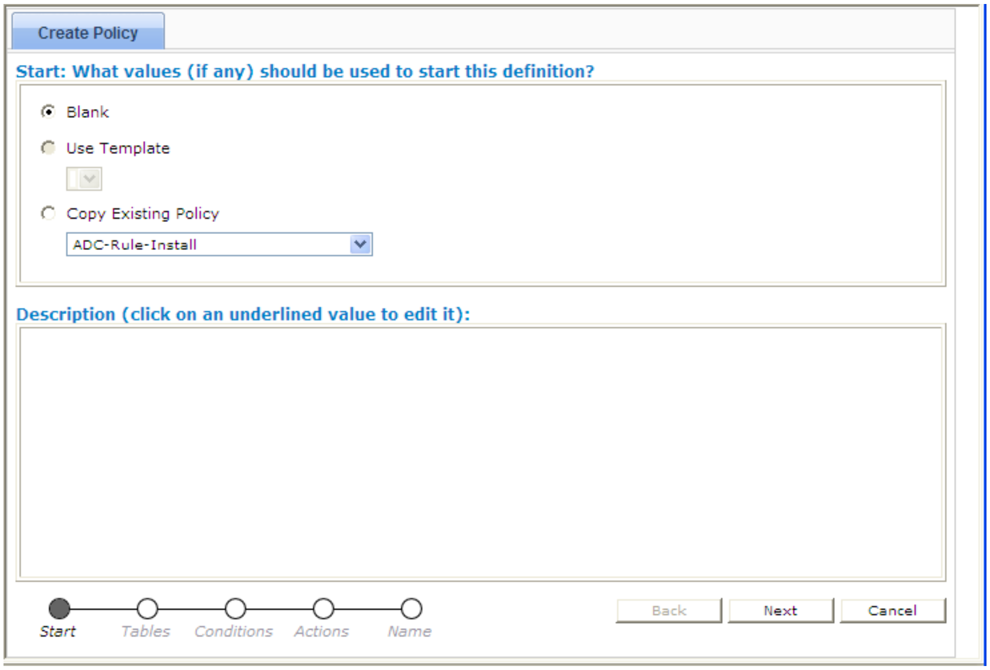

Create Policy

Start: What values (if any) should be used to start this definition?

● Blank

○ Use Template

○ Copy Existing Policy

ADC-Rule-Install

Description (click on an underlined value to edit it):

Figure 15: Create New Template Window

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. *Figure 16: Modify Policy Template Window* shows an example.
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.

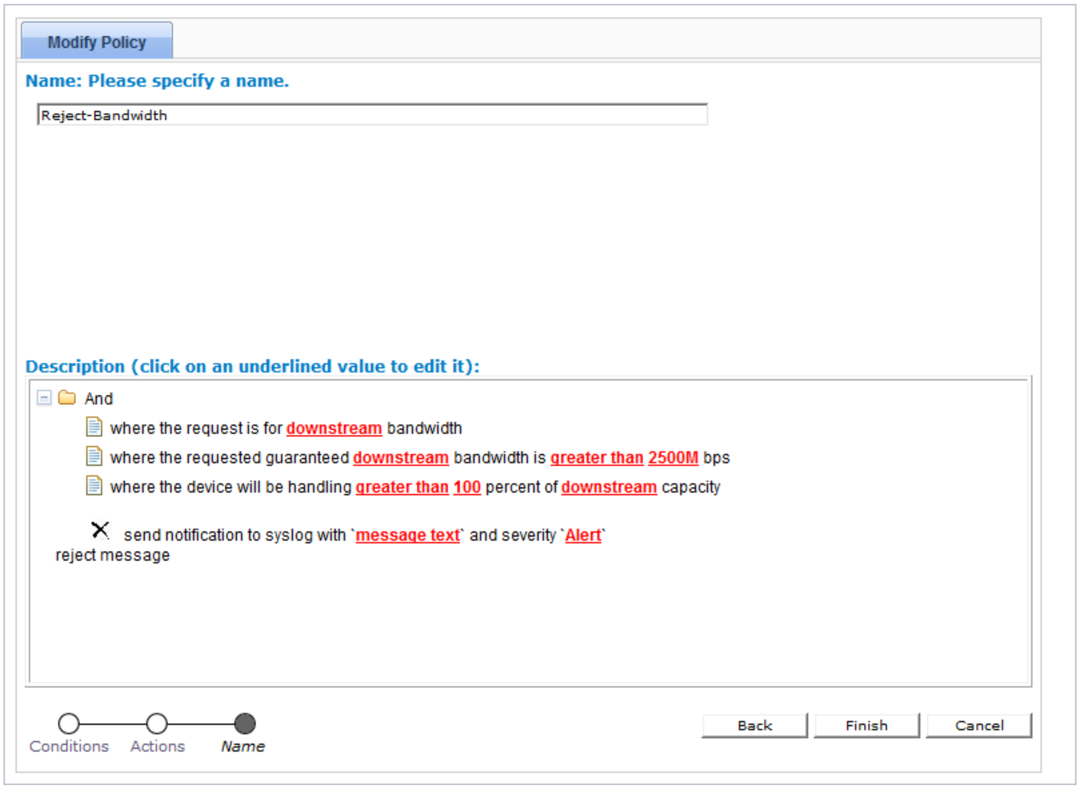

Name: Please specify a name.

Reject-Bandwidth

Description (click on an underlined value to edit it):

And
- where the request is for downstream bandwidth
- where the requested guaranteed downstream bandwidth is greater than 2500M bps
- where the device will be handling greater than 100 percent of downstream capacity

send notification to syslog with `message text` and severity `Alert`
reject message

Figure 16: Modify Policy Template Window

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 system lets you create policy groups. Policy groups are an organizational aid that provide for flexible policy management and deployment. You save policies to a group in the order in which you want an MPE device to apply 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. You can also group, or nest, policy groups.

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:

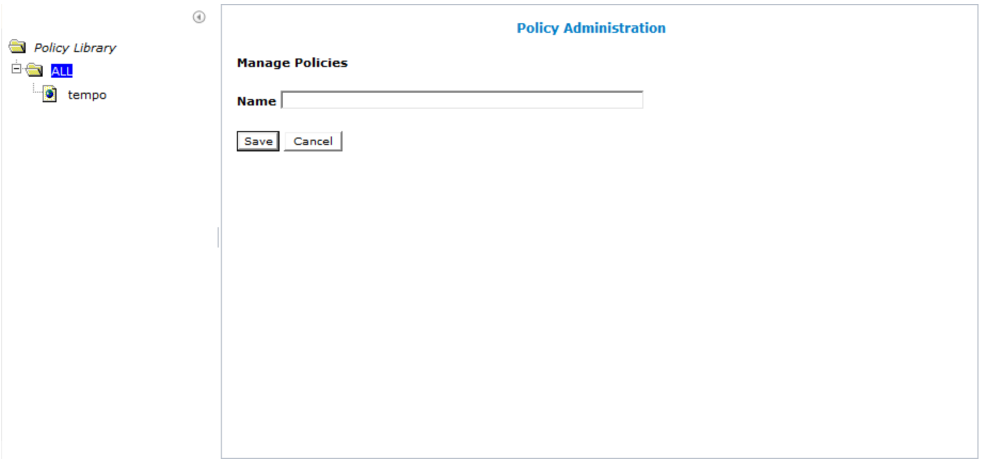

4. Enter the name to assign to the new group, then click **Save** (or **Cancel** to discard your changes).

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; for example:

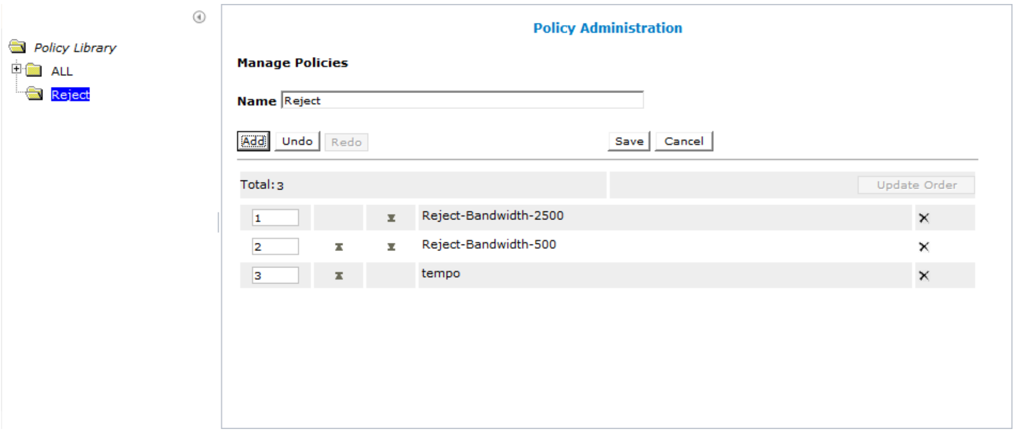

Name Reject

Total: 3

1	Reject-Bandwidth-2500
2	Reject-Bandwidth-500
3	tempo

4. Click **Add**. A window opens, displaying the policies and policy groups available; for example:

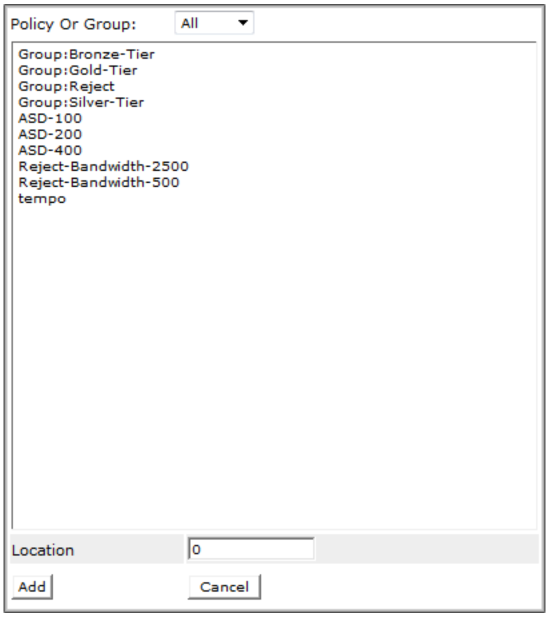

Policy Or Group: All

Group:Bronze-Tier
Group:Gold-Tier
Group:Reject
Group:Silver-Tier
ASD-100
ASD-200
ASD-400
Reject-Bandwidth-2500
Reject-Bandwidth-500
tempo

Location: 0

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. By default policies and policy groups are added after the first item in the group; to change the insert position, change the value in the **Location** field.
   The policies or policy groups are added to the policy group in the specified location 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 Policies or Policy Groups Within a Policy Group*).

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*).

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, located to the left of policies or policy groups. The arrow icon for the top item moves it to the bottom of the list; the arrow icon for the bottom item moves it to the top of the list.
   - Drag and drop policies or policy groups to a different position in the sequence.
   - Change the sequence numbers, located to the left of policies or policy groups. Click **Update Order** to refresh the display.
   - Optionally, you can click **Undo** or **Redo** to step back and forth through your changes.

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.

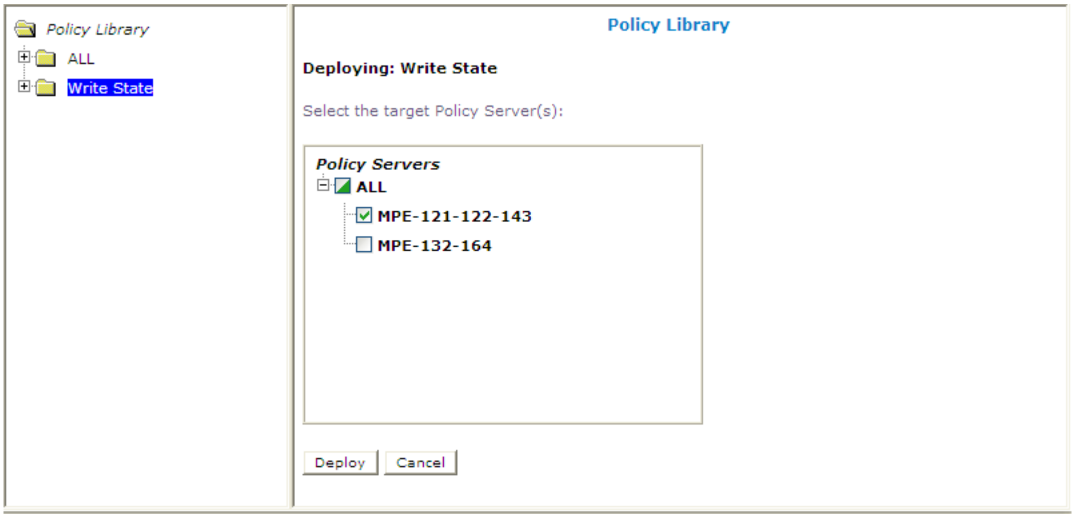

Deploying: Write State

Select the target Policy Server(s):

☑ MPE-121-122-143

□ MPE-132-164

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 from a Policy Group on an MPE Device

To remove a policy from within a policy group that was deployed to an MPE device, modify the policy group on the CMP system; the policy group is automatically redeployed. (To remove an entire policy group from an MPE device, 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 the group contains.
3. Remove the desired policy using one of the following methods:
   - From the Policy Library tree, select the policy. The Policy Administration page displays the profile information. Click **Remove**.
   - On the Policy Administration page, click **Modify** and then select the Remove icon located next to 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 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.

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, located to the left of policies or policy groups. The arrow icon for the top item moves it to the bottom of the list; the arrow icon for the bottom item moves it to the top of the list.
   - Drag and drop policies or policy groups to a different position in the sequence.
   - Change the sequence numbers, located to the left of policies or policy groups. Click **Update Order** to refresh the display.
   - Optionally, you can click **Undo** or **Redo** to step back and forth through your changes.
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.

Managing Policy Checkpoints

A policy checkpoint is a method of saving the records in the CMP database at a specific point in time. Records saved are policies, policy groups, and policy templates. You can save up to ten checkpoints.

Once a checkpoint is created, you can return to this set of records at any time by restoring the checkpoint.

CAUTION

Caution: When you restore a checkpoint, all existing data is permanently removed.

The checkpoint function is different from the export/import function in these ways:

- Checkpoints are saved on the CMP system rather than to a file.
- A checkpoint saves all records mentioned above; the import/export feature allows you to select which records to import or export.
- A checkpoint can only be used on a specific CMP system, and cannot be migrated to another CMP system.

To see this feature on the GUI menu and be able to use it, specify a value greater than 0 for the **Allow policy backup and rollback** field on the System Settings page (see *Configuring System Settings*). This field also controls the maximum number of checkpoints that can be saved.

Viewing and Comparing Policy Checkpoints

Use this procedure to view all checkpoints and/or compare a selected checkpoint's records to the current CMP records. You can also view the records saved for a specific checkpoint.

To view/compare policy checkpoints in the CMP database:

1. From the **Policy Management** section of the navigation pane, select **Policy Checkpoint/Restore**. The Checkpoint/Restore page opens.
2. Click **Diff** to view a report that compares the selected checkpoint's records to the current CMP records.
3. Click **More Info** to view a list of all required profile names for this checkpoint. These profiles must exist in the system before a checkpoint is restored, otherwise the restore will fail.

Creating a Policy Checkpoint

Use this procedure to create a new checkpoint.

Note: The maximum number of checkpoints that can be created is defined on the System Settings page. If you create more than the number defined, the oldest checkpoint is deleted.

To create a new policy checkpoint:

1. From the **Policy Management** section of the navigation pane, select **Policy Checkpoint/Restore**. The Checkpoint/Restore page opens.
2. Click **Create a new checkpoint**.
   If the maximum number of checkpoints already exists, you are prompted, "*n* checkpoints already exist, by creating this checkpoint the oldest one will be deleted. Continue?" (where *n* is the maximum number of checkpoints).
3. In the **Input your memo** field, type a description of the checkpoint.
4. Click **Save** (or **Cancel** to abandon the request).
   The message "Checkpoint successfully added" appears in green on the page.

The checkpoint is created.

Restoring a Policy Checkpoint

CAUTION

Caution: All current records are lost when a restore is performed. It is recommended that you save a checkpoint before restoring a previous checkpoint.

Use this procedure to return to a saved checkpoint.

To restore to a checkpoint in the CMP database without autodeployment to the MPE devices:

1. From the **Policy Management** section of the navigation pane, select **Policy Checkpoint/Restore**.
   The Checkpoint/Restore page opens.
2. Click **Restore**.
3. Select the checkpoint you are restoring.
4. Click **Restore**.
   You are prompted, "Caution: You had better save a checkpoint before any restoration. Are you sure that you want to restore to this Checkpoint?"
5. Click **OK** (or **Cancel** to exit if you need to create a checkpoint).
   If you click **OK**, you are prompted, "All existing deployed policies will be removed from on-line MPE. Select OK to continue."
6. Click **OK** (or **Cancel** to abandon your request).
   The selected checkpoint is restored.

A checkpoint report appears, listing which policies and policy groups were restored and which were removed.

Restoring a Policy Checkpoint to MPE Devices

CAUTION

Caution: All current records are lost when a restore is performed. It is recommended that you save a checkpoint before restoring a previous checkpoint.

To restore to a checkpoint in the CMP database and autodeploy to all MPE devices in the system:

1. From the **Policy Management** section of the navigation pane, select **Policy Checkpoint/Restore**.
   The Checkpoint/Restore page opens.
2. Click **Restore**.
3. Select the checkpoint you are restoring.
4. Click **Restore and Deploy**.
   You are prompted, "Caution: You had better save a checkpoint before any restoration. Are you sure that you want to restore to this Checkpoint and deploy it to MPEs?"
5. Click **OK** (or **Cancel** to exit if you need to create a checkpoint).
   If you click **OK**, you are prompted, "All existing deployed policies will be removed from on-line MPE. Select OK to continue."
6. Click **OK** (or **Cancel** to abandon your request).
   The selected checkpoint is restored and deployed to the MPE devices.

A checkpoint report appears, listing which policies and policy groups were restored, which were removed, and to which MPE devices the deployment succeeded.

Deleting a Policy Checkpoint

To delete a saved checkpoint from the CMP system:

1. From the **Policy Management** section of the navigation pane, select **Policy Checkpoint/Restore**.
   The Checkpoint/Restore page opens.
2. Select the checkpoint you are deleting.
3. Click **Delete the selected checkpoint** to remove the checkpoint from the system.
   You are prompted, "Are you sure you want to delete this Checkpoint?"
4. Click **OK**.
   The message "Checkpoint deleted successfully" appears in green on the page.

The selected checkpoint is deleted from the CMP database.

Chapter 9

Managing Subscribers

Topics:

Managing Subscribers describes how to create and manage subscriber tiers and accounts within the CMP system.

Creating a Tier

Tiers are categories that you can define and then apply to groups of subscribers. For example, you can create a series of tiers with different bandwidth limits. Once you define tiers, you can use them in policy rules.

To create a subscriber tier:

1. From the **Subscriber** section of the navigation pane, select **Tiers**.
   The content tree displays the **Tiers** folder.
2. Select the **Tiers** folder.
   The Tier Administration page opens.
3. Click **Create Tier**.
   The New Tier page opens.
4. Enter information as follows:
   a) **Name** (required) — Name of the tier.

      The name can be up to 255 characters long and must not contain quotation marks (") or commas (,).

   b) **Description/Location** — Free-form text.

      Enter up to 250 characters.

   c) **Downstream bandwidth limit (bps)** — The maximum amount of bandwidth capacity available in the downstream direction in bits per second.

      You can enter a value followed by M or G; for example, 4G for 4 gigabits per second.

   d) **Upstream bandwidth limit (bps)** — The maximum amount of bandwidth capacity available in the upstream direction in bits per second.

      You can enter a value followed by M or G; for example, 10M for 10 megabits per second.
5. When you finish, click **Save** (or **Cancel** to cancel the request).
   The tier is added to the CMP database, and the message "Tier created successfully" is displayed.

You can now use the tier in policy rules.

Displaying Subscriber Activity History

To display subscriber account activity:

1. From the **Subscriber** section of the navigation pane, select **Accounts**.
   The content tree displays the Subscriber Account page.
2. On the Subscriber Account page, type in search terms in one or both of the **Search** fields:
   - **Account ID** — You can use as wildcard characters an asterisk (*) to represent any string or a question mark (?) to represent any single character. Searches are case insensitive. The search string must represent the entire account ID; for example, if the account ID is "Account50619," you can find it using the search strings "`Account*`" or "`*50619`" or "`Account506??`," but not using the search strings "`Account`" or "`506`" or "`619`."

Note: Using wildcard characters can result in longer search times.

- **Static IP Address** — An IP address in the form *n.n.n.n*. You can use as wildcard characters an asterisk (*) to represent any string or a question mark (?) to represent any single character. Searches are case insensitive.

3. Click **Search**.
   The Subscriber Account Search page opens. The **Status** column describes the state of synchronization between each account and the MPE device(s) with which it is associated. The status values are as follows:

   - **up-to-date** — The account information is current on the MPE device(s) with which it is associated.
   - **Update Pending** — The account is either new, not yet associated with any MPE device, or has been changed, but the changes have not been sent to any MPE device.
   - **Delete Pending** — The account is marked for deletion, but has not yet been removed from the MPE device(s).

4. Click on an account ID.
   The Subscriber Account page opens. From this page you can do the following:

   - To modify account information, click **Modify**. (See *Modifying an Account*.)
   - To delete the account, click **Delete**. (See *Deleting an Account*.)
   - To display the static IP address of the account, click the Static IP Address tab; the Static IP Address page opens. From this page you can modify or delete the static IP address of this account. (See *Static IP Addresses*.)
   - To display real-time statistics for the account, click the Real-time Statistic tab; the Real-Time Statistics page opens. To refresh the data on this page, click **Refresh**. (See *Displaying Real-time Subscriber Statistics*.)

5. When you finish, click **Back to Search Results**.
   The previous page is displayed.

Displaying Real-time Subscriber Statistics

Dynamic statistical information for a subscriber account appears on the Real-time Statistic tab of the Subscriber Account page. The following information is displayed:

- **Policy Server** — the MPE device associated with the subscriber
- **Current IP Address** — the IP address associated with the subscriber (if applicable)
- **Total Bandwidth Allocated** — the total bandwidth, in Kbps, of all VoD sessions associated with the subscriber
- **Total VOD Session Count** — the number of VoD sessions associated with the subscriber

To refresh the display, click **Refresh**. The MPE device is queried.

To export data in XML format, click **Export**; the browser file download window opens. The data is exported in the same format used by the OSSI XML interface. (For format information, see the *OSSI XML Interface Definitions Reference*.)

Deleting a Tier

To delete a tier:

1. From the **Subscriber** section of the navigation pane, select **Tiers**.
   The **Tiers** folder appears in the content tree.
2. Delete the tier using one of the following methods:
   - From the work area, click the Delete icon, located to the right of the tier you wish to delete.
   - From the content tree, select the tier and click **Delete**.

   You are prompted, "Are you sure you want to delete this Tier?"
3. Click **OK** (or **Cancel** to cancel the request).
   The message "Tier deleted successfully" is displayed in green on the page.

You have deleted the tier.

Creating an Account

Subscriber accounts are usually created external to the CMP system. However, you can create or delete an individual subscriber account.

To create a subscriber account:

1. From the **Subscriber** section of the navigation pane, select **Accounts**.
   The Subscriber Account page opens.
2. Click **Create Account**.
   The New Account page opens.
3. Enter information as follows:
   a) **Account ID** (required) — Name of the account.
   b) **Subscriber Data** — Additional data associated with the account, such as the specific router interface.

      Enter up to 250 characters.
   c) **Network Element** (required) — The network element associated with this subscriber.
   d) **Subscriber Group** — If during a VoD reserve request the subscriber's account record contains a reference to an existing subscriber group network element, the network element is dynamically added to the path that the VoD request used. This charges all of the VoD session's resources against all network elements in the defined path as well as the associated subscriber group.
   e) **Downstream bandwidth limit (bps)** — The maximum amount of bandwidth capacity available in the downstream direction in bits per second.

      You can enter a value followed by M or G; for example, 4G for 4 gigabits per second. Leave blank to use the tier value.
   f) **Upstream bandwidth limit (bps)** — The maximum amount of bandwidth capacity available in the upstream direction in bits per second.

      You can enter a value followed by M or G; for example, 10M for 10 megabits per second. Leave blank to use the tier value.

g) **Tier** — The tier associated with this account.

4. When you finish, click **Save** (or **Cancel** to cancel the request).
   The message "Account created successfully" is displayed.

The account is added to the CMP database.

Modifying an Account

To modify an account:

1. From the **Subscriber** section of the navigation pane, select **Accounts**.
   The Subscriber Account page opens.
2. Use the search function to locate and display an account (see *Displaying Subscriber Activity History*).
   The account is displayed.
3. Click on the account.
   The Subscriber Account page displays information about that account.
4. Click **Modify**.
   The Modify Account page opens.
5. Make changes as required.
   (See *Creating an Account* for information on the fields on this page.)
6. When you finish, click **Save** (or **Cancel** to discard your changes).

The account information is modified, and the change is deployed to all associated MPE devices.

Note: If an associated MPE device is offline or unreachable, the subscriber account information is not deployed. See *Updating Accounts* for information on updating accounts.

Updating Accounts

Subscriber account information can be imported to, created on, or modified on a CMP system, but the information may not immediately be associated with an MPE device. For example, MPE devices can be offline when subscriber accounts are being deployed. This leaves accounts in a state where their new information is pending update, and the changes must be deployed (or "pushed") to MPE devices. You can use the CMP system to deploy subscriber account information.

The update status of accounts is displayed in on the Subscriber Account Search page. A status of "Update Pending" indicates that an account update is pending.

To update all pending accounts:

1. From the **Subscriber** section of the navigation pane, select **Accounts**.
   The Subscriber Account page opens.
2. Click **Push Pending Accounts to MPE**.
   All pending subscriber data updates are deployed to all MPE devices controlled by this CMP system. If the operation takes more than five seconds, a progress bar appears.

If a hardware failure or other such event results in a new MPE device coming online to replace an older one, then the topology and subscriber data need to reapplied to the new system. The **Reapply**

Configuration button, on the System tab of the Policy Server Administration page, redistributes topology information to the new MPE device. The **Reapply Subscriber Configuration** button, on the same page, redistributes subscriber account information. This function completely deploys all associated accounts (not just pending changes) to an MPE device, and provides a way to resynchronize the CMP and MPE device subscriber databases.

To update all subscriber account information:

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 to update.

   The Policy Server Administration page opens in the work area.

3. On the System tab, click **Reapply Subscriber Configuration**.

   All subscriber data is deployed to the MPE device, and an entry is written to the audit log. If the operation takes more than five seconds, a progress bar appears.

Deleting an Account

To delete an account:

1. From the **Subscriber** section of the navigation pane, select **Accounts**.
   The Subscriber Account page opens.
2. Use the search function to locate and display an account (see *Displaying Subscriber Activity History*).
   The account is displayed.
3. Click **Delete**.
   You are prompted, “Are you sure you want to delete this Account?”
4. Click **OK** (or **Cancel** to cancel the request).
   The message “Account deleted successfully” is displayed.

The account is deleted.

Static IP Addresses

The Policy Management product supports static IP addresses in the subscriber account definitions that are defined on (or provisioned to) the CMP system, and in the session handling function within MPE devices.

The OSSI XML code includes an optional IP address range definition within the subscriber account definition. This definition consists of a single starting IP address and the number of addresses in the range. A single account can have up to 125 static IP addresses.

Once provisioned, these subscriber records are pushed to the responsible MPE device by the CMP system.

The static IP address is manageable through the CMP system. Access to static IP address management is based on user permissions. The following functions are available:

- You can update and modify accounts with an IP address range. Changes are validated to ensure that duplicate static IP address are not created. (There is no provision to validate against a gateway router subnet definition.)
- You can search for accounts by IP address string. This search uses the same wildcards as the account ID search string and can be used either alone or in combination with the current account ID search string to retrieve accounts.
- Static IP objects are included in exported XML. The Static IP XML block is additive and existing Account XML imports/exports are backward compatible with the system.
- Modifications or exportations generate an audit log message for each operation indicating the operation that was performed, the username of the operator performing the action, and the generic status of the operation.

The MPE device maps static IP address fields to account records. Static IP addresses are stored in a separate database table. This separation ensures that static IP addresses are not disturbed during COPS-PR re-synchronization.

If a subscriber account has both static IP and dynamic IP addresses, all addresses are valid for processing VoD session requests. Also, if a subscriber account is configured with a static IP address, and a B-RAS server sends the same address as a dynamic IP address, then the B-RAS dynamic IP address is used.

When an IP address is learned (from the dynamic COPS-PR information flow) that is already defined as a static address (from the subscriber database update), a trace log entry is generated (Error level). This check is also performed at every video session request, and the same event is generated if necessary.

Configuring a Static IP Address

Account information is normally imported from an XML file. However, the CMP system lets you create an account manually. Static IP address information can only be added to an existing account.

To modify the static IP address of a subscriber account:

1. From the **Subscriber** section of the navigation pane, select **Accounts**.
   The Subscriber Account page opens.
2. Use the search function to locate and display an account (see *Displaying Subscriber Activity History*).
   The account is displayed.
3. Click on an account ID.
   The Subscriber Account page opens.
4. On the Static IP Address tab, click **Modify**.
   The Modify Account page opens.
5. Type the following information:
   - **Static IP Address** — An IP address, in the format *n.n.n.n*.
   - **Static IP Count** — The number of addresses in the range. A single account can have up to 125 static IP addresses.
6. When you finish, click **Save** (or **Cancel** to discard your changes).

The account information is modified.

Deleting a Static IP Address from a Subscriber Account

To delete static IP addresses from a subscriber account:

1. From the **Subscriber** section of the navigation pane, select **Accounts**.
   The Subscriber Account page opens.
2. Use the search function to locate and display an account (see *Displaying Subscriber Activity History*).
   The account is displayed.
3. Click on an account ID.
   The Subscriber Account page opens.
4. On the Static IP Address tab, click **Delete**.
   You are prompted, “Are you sure you want to delete these IP Addresses?”
5. Click **OK** (or **Cancel** to cancel the request).
   When the operation finishes, the page displays “<None>.”

The static IP addresses are removed from this subscriber account.

Provisioning Static IP Addresses Using XML

You can add, update, or delete static IP addresses through the OSSI XML Interface. The <AddAccount> and <UpdateAccount> tags include optional Static IP addresses within the subscriber account definition. This definition consists of a single starting IP address and the number of addresses in the range. A maximum of 125 static IP addresses can be associated with a single account.

For example, the following definition specifies an IP address range of 10.0.1.1, 10.0.1.2, 10.0.1.3, ..., 10.0.1.125:

```
<?xml version="1.0" encoding="UTF-8"?>
<XmlInterfaceRequest>
  <AddAccount>
    <Account>
      <AccountId>47/VAXA/261188/ /VZVA</AccountId>
      <SubscriberData>11/3.30251</SubscriberData>
      <NetworkElementName>FRBGVAFB1FW
      </NetworkElementName>
      <StaticIp>
        <IpAddress>10.0.1.1</IpAddress>
        <IpCount>125</IpCount>
      </StaticIp>
      <TierRef>
        <Name>INFOSPEED_FTTPV_AA</Name>
      </TierRef>
    </Account>
  </AddAccount>
</XmlInterfaceRequest>
```

Only one IP address range is allowed per subscriber. Validation is performed to ensure that duplicate static IP addresses are not provisioned.

Note: There is no validation against the gateway router subnet definition.

The <QueryAccount> tag includes XML for static IP objects as part of the response.

To change the address range, use the <UpdateAccountType> tag to redefine the range, which replaces the previous definition. Modifying the <IpCount> value separately is not supported.

For more information, see the *OSSI XML Interface Definitions Reference Guide*.

Chapter 10

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 summary provides an aggregate view of 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 17: Sample Active Alarms Report shows a sample active alarm report.

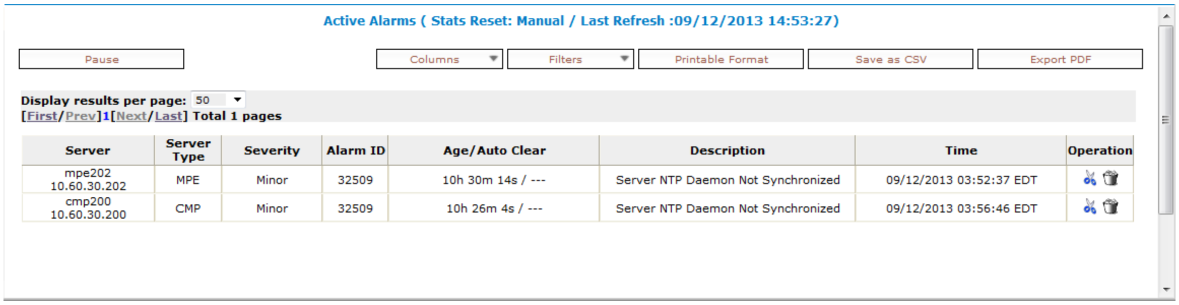

Active Alarms (Stats Reset: Manual / Last Refresh :09/12/2013 14:53:27)

Display results per page: 50

[First/Prev]1[Next/Last] Total 1 pages

Server	Server Type	Severity	Alarm ID	Age/Auto Clear	Description	Time	Operation
mpe202 10.60.30.202	MPE	Minor	32509	10h 30m 14s / ---	Server NTP Daemon Not Synchronized	09/12/2013 03:52:37 EDT	✂ 🗑
cmp200 10.60.30.200	CMP	Minor	32509	10h 26m 4s / ---	Server NTP Daemon Not Synchronized	09/12/2013 03:56:46 EDT	✂ 🗑

Figure 17: 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 *Viewing the Trace Log*.

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.

The Age/Auto Clear column shows how long an alarm has been active (that is, how long since it was raised) and how long the alarm will display before being automatically cleared. The Auto Clear time is shown as “---” if the alarm is not automatically cleared.

The following options are available:

- To sort the report on any column, click the column title.
- To hide an alarm, click the hide icon (scissors), located to the right of each row. All instances of alarms with that ID reported from that server are hidden from display (but shown in the Hidden Filter, which you can use to restore the display of those alarms).

  Note: Hiding an alarm only affects the current user. Other users will see the alarm if they display the Active Alarms page.

- To manually clear an alarm, click the clear icon (trash can), located to the right of each row. You are prompted, “This alarm will be cleared. Are you sure?” Click **OK** (or **Cancel** to abandon your request).

- 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 menu:
  - The **Search Filter** tab has three controls. 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.
  - The **Hidden Filter** tab shows alarms, by server and alarm ID, that are currently hidden from display. Click the delete icon, to the right of an entry, to remove it from the list of hidden items and display it in the page again.
- **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.

Note: If you wish to view the most recent alarms, and there are more than 50,000 alarms in the database, specify a start date/time that includes the present.

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 in IPv4 or IPv6 format.
- **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:

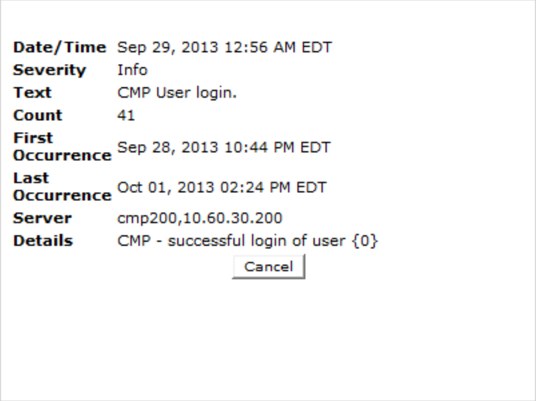

Date/Time	Sep 29, 2013 12:56 AM EDT
Severity	Info
Text	CMP User login.
Count	41
First Occurrence	Sep 28, 2013 10:44 PM EDT
Last Occurrence	Oct 01, 2013 02:24 PM EDT
Server	cmp200,10.60.30.200
Details	CMP - successful login of user {0}

Click **Cancel** to close the window.

Chapter 11

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 7: 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**, **MPE**, 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

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

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 to the configuration (such as creating or editing network elements or policies) may be lost.

WARNING

Warning: Contact the Tekelec Customer Care Center and inform them of your upgrade plans prior to beginning this or any upgrade procedure.

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.

About Preparing for an Upgrade

Upgrading a server requires a large amount of preparation. For detailed information about preparing for an upgrade, please access the Tekelec Customer Support site.

CAUTION

Caution: Use only the upgrade procedure provided by the Tekelec Customer Care Center. Also, before upgrading any system, please access the Tekelec Customer Support site and review any Technical Service Bulletins (TSBs) that relate to this upgrade.

WARNING

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 8: 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 application running on each server. The **Filters** pulldown menu lets you select **CMP Site1 Cluster**, **MPE**, or **All** servers.
IP	Displays the IP address of each server. The **Filters** pulldown menu lets you 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 lets you display **Active**, **Standby**, **Out-Of-Service**, **Force Standby**, or **All** servers.
ISO	Displays the ISOs or CD-ROMs 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 Management software release of each server. The **Filters** pulldown menu lets you display a specific release only or **All** releases.
Replication Status	Displays 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**, **IP**, **State**, **Prev Release**, and **Running Release**. These filters are set to **All** by default, so initially all servers appear. 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 selected (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.

	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 (you can click **OK** or **Cancel**). If you click **OK**, a progress bar displays the status of the command completion in a pop-up window. Once an operation is confirmed, it cannot be cancelled.

Operation	Effect on Selected Server(s)
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).
Force Standby	Forces server to standby status. A server must be in Forced Standby status before you can complete an upgrade. ⚠ CAUTION **Caution:** Setting Force Standby for all servers in a cluster effectively removes the cluster from service. **Note:** You cannot force both servers of a CMP cluster into standby status.
Turn Off Replication	Turns off replication.
Prepare Upgrade	Turns off COMCOL replication of database tables.
Switch ForceStandby	Switches the upgraded server to Active and the previously active server to Forced Standby to upgrade it.
Accept Upgrade	Completes the upgrade process. This operation is available for servers in the Pending state. A server must be in Forced Standby status before an upgrade can be completed. Once this operation is performed, the upgrade cannot be backed out.
Reject Upgrade	Backs out of the upgrade process. This operation is available for servers in the Pending state. A server must be in Forced Standby status before an upgrade can be rejected.

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.

Chapter 12

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.

      If more than the configured number of concurrent users try to log in (for example, a second user if this value is set to 1), they are blocked at the login page with the message "Your account already has the maximum number of concurrent sessions."

   d) **Password Expiration Period (days; 0=never)** — The number of days a password can be used before it expires. Enter a value from 7 to 365, or 0 to indicate that the password never expires.
   e) **Password Expiration Warning Period (days; default=3)** — The number of days before a password expires to begin displaying a window to users after login warning that their password is expiring.
   f) **Admin User Password Expiration** — By default, the password for the admin user never expires.

      If you select this option, the **admin** user is subject to the same password expiration policies as other users.

   g) **Block users when password expires** — By default, once a password expires, the user must immediately change it at the next login.

      If you select this option, if their password expires, users cannot log in at all. (If you select **Admin User Password Expiration** and the **admin** user's password expires, the user can still log in but must immediately 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 255 characters.
   k) **Top Banner Text** — The text that displays in the banner at the top of the GUI page. By default, the text displayed is "Policy Management: *hostname*" (where *hostname* is the name of the system).

If you enter text in this field, it will overwrite this default.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** — 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 18: Sample Password Strength Policy shows an example of settings that establish a password strength policy requiring user passwords to contain at least one uppercase letter, four numerals, and one non-alphabetic character. (A password that would satisfy this policy is P@ssword1357.) Users whose passwords do not meet these requirements will be forced to change their passwords the next time they log in.

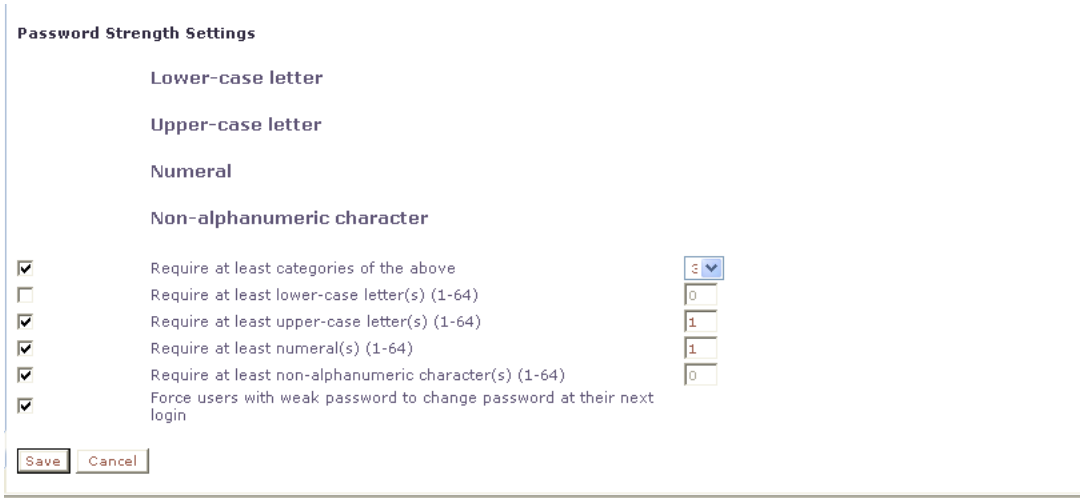

Password Strength Settings

Lower-case letter

Upper-case letter

Numeral

Non-alphanumeric character

☑	Require at least categories of the above	3
□	Require at least lower-case letter(s) (1-64)	0
☑	Require at least upper-case letter(s) (1-64)	1
☑	Require at least numeral(s) (1-64)	1
☑	Require at least non-alphanumeric character(s) (1-64)	0
☑	Force users with weak password to change password at their next login	

Figure 18: 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

For detailed information, see the document *OSSI XML Interface Definitions Reference Guide*.

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: * (specifies import all types), **Network Elements**, **Paths**, **Accounts**, **Tiers**, **Applications**, **Roles**, **Scopes**, or **Users**. * is the default value.
   If you select **Network Elements**, additional filtering fields appear to help you manage the volume of data being imported; you can filter by network element name, element ID, or Diameter identifier. If you select **Accounts**, an additional filtering field appears to help you manage the volume of data being imported; you can filter by account ID. 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), **Paths**, **Accounts**, **Tiers**, **Applications**, **Roles**, **Scopes**, or **Users**.
   If you select **Network Elements**, additional filtering fields appear to help you manage the volume of data being exported; you can filter by network element name, element ID, or Diameter identifier. If you select **Accounts**, additional filtering fields appear to help you manage the volume of data being exported; you can filter by account ID, network element name, or MPE device. 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 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 following information is displayed in the Manager Report:

- **Cluster Name and Designation** — The name of the cluster, whether it is the primary **(P)** or secondary **(S)** site, and its mode:

  - **Active**: The cluster is currently managing the Policy Management network.
  - **Standby**: The cluster is not currently managing the Policy Management network.

- **Cluster Type and Status** — A CMP system displays **Manager**. The possible values of the cluster status are the following:
  - **On-line**: If one server, it is active; if two servers, one is active and one is standby.
  - **Degraded** (two servers only): One server is active, but at least one other server is not available.
  - **Offline**: No server is active.
  - **Inconsistent** (two servers only): Both servers are active. This is a "split brain" error condition, and can only happen when the backplane link fails.
- **Blades** — The status of the servers (blades) contained within the cluster. A symbol (🔌) indicates which server currently has the external connection (that is, which blade is the active server). The report also lists the following server-specific information:
  - **Overall**: Displays the current topology state (Active, Standby, or Out-Of-Service), 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), average value for the CPU utilization, and memory.

The **Actions** buttons let you restart the CMP software on a server or restart (reboot) the server itself.

To pause refreshing the display, click **Pause**. To resume refreshing, click **Resume**. To reset the display counters, click **Reset All Counters**.

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.

Modifying the Trace Log Configuration

To configure the trace log display:

1. From the **System Administration** section of the navigation pane, select **Trace Log**.
   The Trace Log page opens in the work area, displaying the current trace log configuration.
2. On the Trace Log page, click **Modify**.

The Modify Trace Log Settings page opens.

3. Define the settings.
   For a description of the settings, see *Configuring Log Settings*.
4. When you finish, click **Save** (or **Cancel** to discard your changes).
   The Modify Trace Log Setting page closes.

The trace log configuration is modified.

Viewing the Audit Log

You can track and view configuration changes within the CMP 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 19: Audit Log* shows an example.)

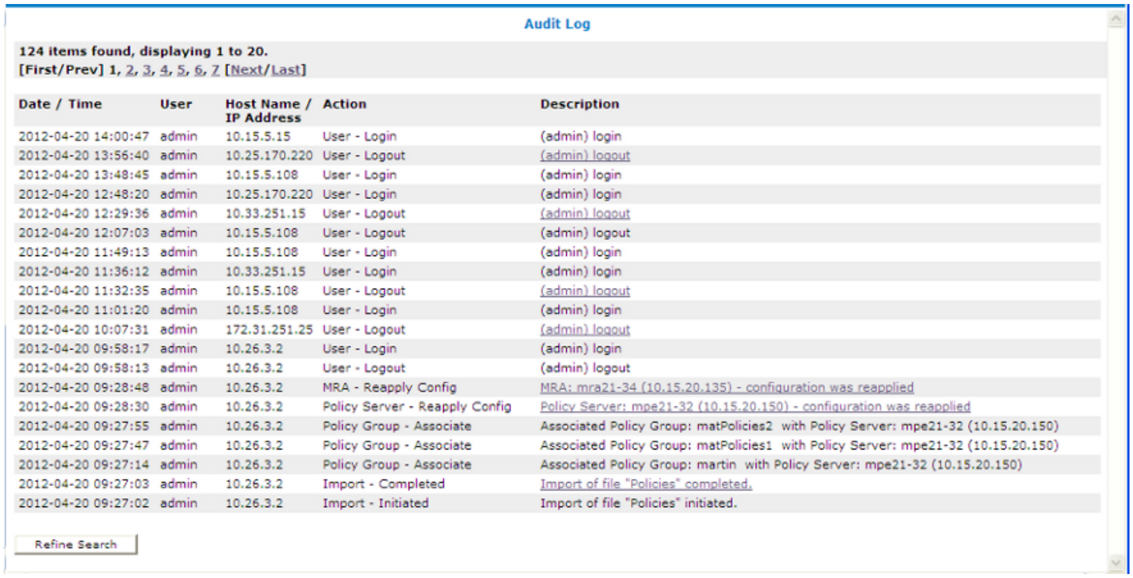

124 items found, displaying 1 to 20.
[First/Prev] 1, 2, 3, 4, 5, 6, 7 [Next/Last]

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.

Figure 19: Audit Log

For a detailed description of an item, click the underlined description. The details of the event display. (*Figure 20: 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*.)

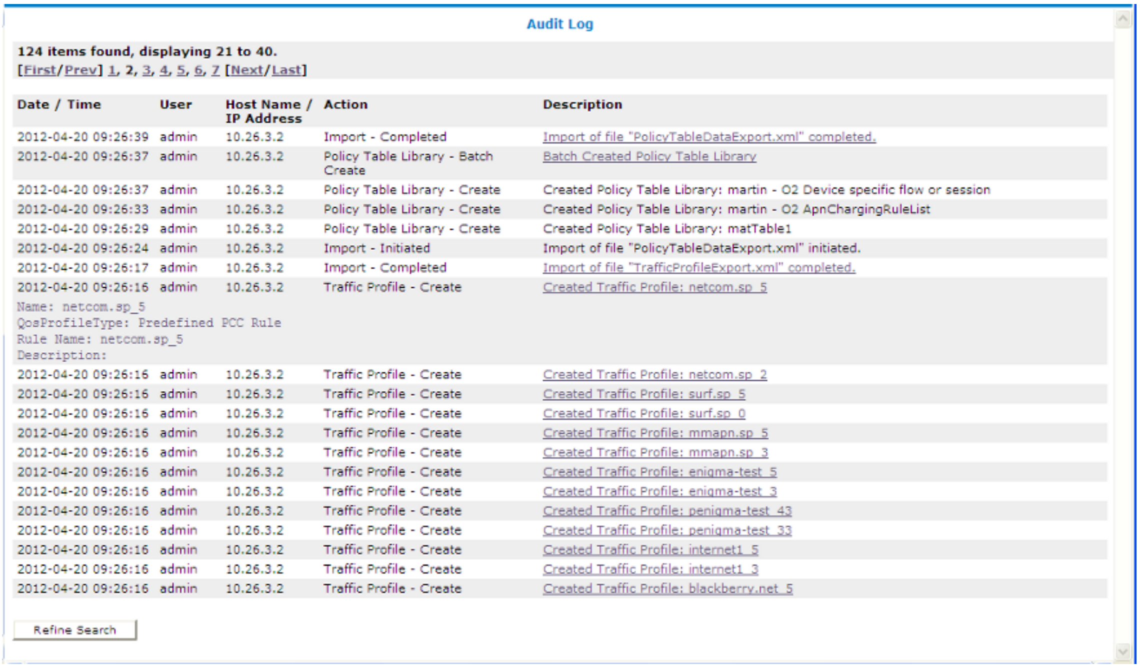

Audit Log

124 items found, displaying 21 to 40.
[First/Prev] 1, 2, 3, 4, 5, 6, 7 [Next/Last]

Date / Time	User	Host Name / IP Address	Action	Description
2012-04-20 09:26:39	admin	10.26.3.2	Import - Completed	Import of file "PolicyTableDataExport.xml" completed.
2012-04-20 09:26:37	admin	10.26.3.2	Policy Table Library - Batch Create	Batch Created Policy Table Library
2012-04-20 09:26:37	admin	10.26.3.2	Policy Table Library - Create	Created Policy Table Library: martin - O2 Device specific flow or session
2012-04-20 09:26:33	admin	10.26.3.2	Policy Table Library - Create	Created Policy Table Library: martin - O2 ApnChargingRuleList
2012-04-20 09:26:29	admin	10.26.3.2	Policy Table Library - Create	Created Policy Table Library: matTable1
2012-04-20 09:26:24	admin	10.26.3.2	Import - Initiated	Import of file "PolicyTableDataExport.xml" initiated.
2012-04-20 09:26:17	admin	10.26.3.2	Import - Completed	Import of file "TrafficProfileExport.xml" completed.
2012-04-20 09:26:16	admin	10.26.3.2	Traffic Profile - Create	Created Traffic Profile: netcom.sp_5
Name: netcom.sp_5 QosProfileType: Predefined PCC Rule Rule Name: netcom.sp_5 Description:				
2012-04-20 09:26:16	admin	10.26.3.2	Traffic Profile - Create	Created Traffic Profile: netcom.sp_2
2012-04-20 09:26:16	admin	10.26.3.2	Traffic Profile - Create	Created Traffic Profile: surf.sp_5
2012-04-20 09:26:16	admin	10.26.3.2	Traffic Profile - Create	Created Traffic Profile: surf.sp_0
2012-04-20 09:26:16	admin	10.26.3.2	Traffic Profile - Create	Created Traffic Profile: mmapn.sp_5
2012-04-20 09:26:16	admin	10.26.3.2	Traffic Profile - Create	Created Traffic Profile: mmapn.sp_3
2012-04-20 09:26:16	admin	10.26.3.2	Traffic Profile - Create	Created Traffic Profile: enigma-test_5
2012-04-20 09:26:16	admin	10.26.3.2	Traffic Profile - Create	Created Traffic Profile: enigma-test_3
2012-04-20 09:26:16	admin	10.26.3.2	Traffic Profile - Create	Created Traffic Profile: penigma-test_43
2012-04-20 09:26:16	admin	10.26.3.2	Traffic Profile - Create	Created Traffic Profile: penigma-test_33
2012-04-20 09:26:16	admin	10.26.3.2	Traffic Profile - Create	Created Traffic Profile: internet1_5
2012-04-20 09:26:16	admin	10.26.3.2	Traffic Profile - Create	Created Traffic Profile: internet1_3
2012-04-20 09:26:16	admin	10.26.3.2	Traffic Profile - Create	Created Traffic Profile: blackberry.net_5

Figure 20: 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) / MRA(s)** — Enter the name of the Policy Management device to audit.
   - **Audit Log Items to Show** — Specifies a category of items to audit for display: **Policy Server**, **Scheduled Task**, **Network Element**, **Network Element Group**, **Network Element Link**, **Application**, **Policy**, **Policy Group**, **Account**, **Tier**, **Path**, **User**, **Audit**, **Alarm**, **OM Statistics**, **MPE Manager**, **Upgrade Manager**, **Topology Setting**, or **Global Configuration Settings**. When you select some categories, a **Name** field appears, which lets you enter a search string; leave the field blank to include all items. When you select any category, an **Action(s)** link appears, which lets you select individual audit log items within the category. By default all items in the category are selected, but you can select individual items instead. By default you can specify three item categories; click **More Lines** to add an additional item category.
   - **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 defining the search parameters, click **Search**.
   The Audit Log displays search results.

Exporting Audit Log Data

You can export audit log data to a text file. The filename is `AuditLogExport.txt`.

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 —** Click the calendar icon, located to the right of the field; a calendar window opens. Select a date.
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 Audit Log Data

To purge data from the audit log:

1. From the **System Administration** section of the navigation pane, select **Audit Log**.

   The Audit Log page opens in the work area.
2. On the Audit Log page, click **Export/Purge**.

   The Export and Purge Audit Log Items page opens.
3. In the **Items to Purge** section, click the calendar icon, located to the right of the field; a calendar window opens. Select a date.
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 abandon the request).

The data is purged from the audit log.

Managing Scheduled Tasks

The CMP system runs batch jobs to complete certain operations. These tasks are scheduled to run at regular intervals, with some tasks scheduled to run in a certain order. You can change the scheduling of these tasks to better manage network load or to propagate a network element change to the Policy Management devices on demand. You can also abort a running task.

CAUTION

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 system. 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 before you can obtain any statistical data from the OSSI XML interface. The *OSSI XML Interface Definitions Guide* describes the OM Interface and the OM Statistics in detail.

- **Legacy OM Statistics** — Periodically retrieves OM statistics from MPE devices executing the previous release of Policy Management software. This task should be run only during migration between software releases.

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:

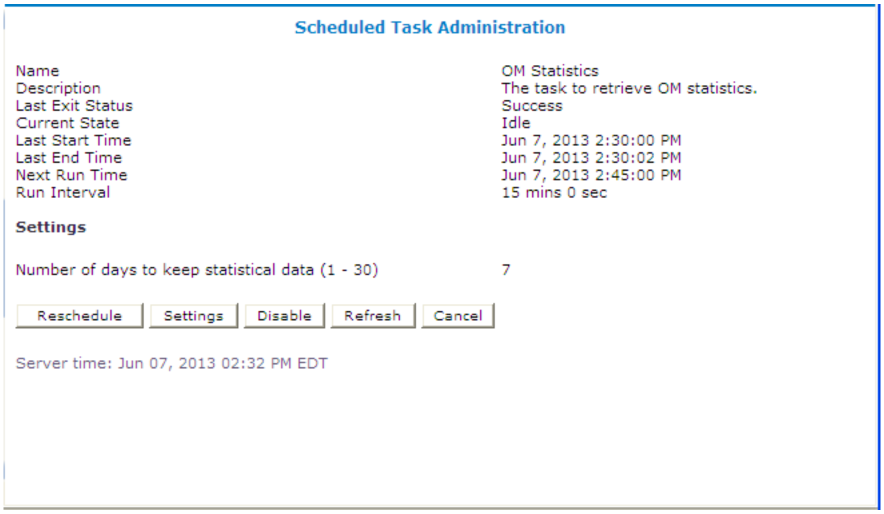
Scheduled Task Administration

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

Settings

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

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

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:

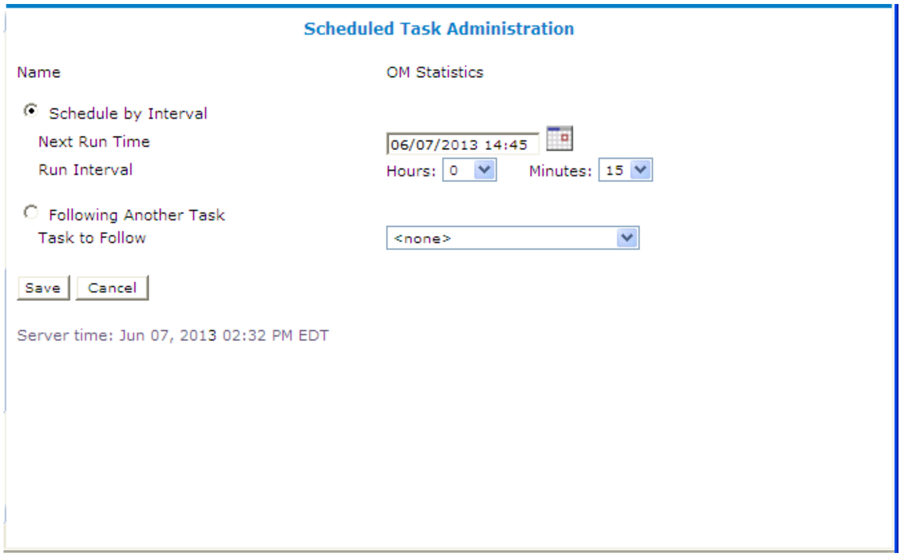
Scheduled Task Administration

Name OM Statistics

● Schedule by Interval

Next Run Time 06/07/2013 14:45

Run Interval Hours: 0 Minutes: 15

○ Following Another Task

Task to Follow <none>

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

   - **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.
   - **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** — What network element groups and Policy Management device groups a user can control, which provides 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** (that is, information can be displayed but not changed).
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

AVP Definition (not supported)

Global Configuration Settings

d) **Subscriber Privileges —** Defines access to the subscriber functions (assigning the privilege **Hide**, **Read-Only**, or **Read-Write**):

Account

Subscriber Tier

e) **Network Privileges —** Defines access to the network management functions (assigning the privilege **Hide**, **Read-Only**, or **Read-Write**):

Topology

Path

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**) (not supported)

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 (not supported)

h) **Platform Setting Privileges —** Defines access to the platform setting functions:

Platform Setting (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**) (not supported)

If set to **Read-Only**, the following fields are displayed for the Stats File Generator setting:

- **Name**

- **Description**
- **Last Exit Status**
- **Current State**
- **Last Start Time**
- **Last End Time**
- **Follows Task**

Task Settings

- **Local Repository—** Root directory of the local repository.
- **Maximum age to keep files (hours)—** Stats file retention period. Defaults to 72 hours.
- **File Format—** Any format can be selected. Defaults to XML.
- **Stats Type—** Any stats type can be selected to generate stats. Defaults to No one. If you do not select a stats type, the task will not run normally.

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 abandon the request).

The role's information is deleted from the CMP database.

Creating a New Scope

You can 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 database. 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. 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 21: 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.

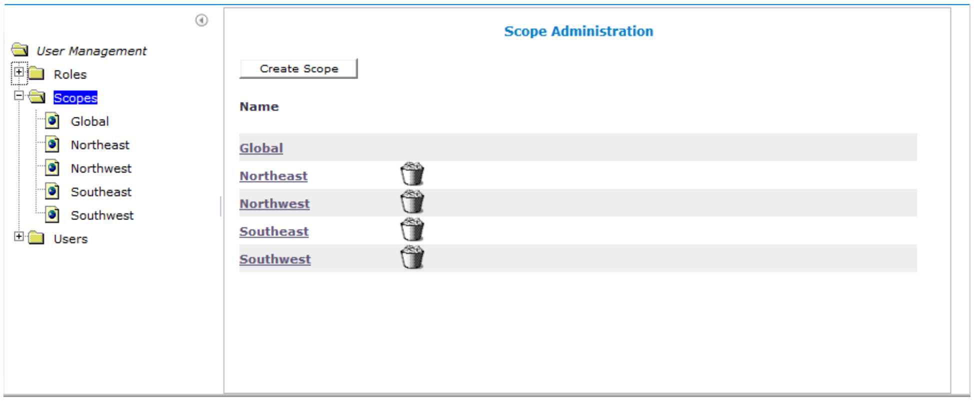

Figure 21: Deleting a Scope

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 (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).
   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 22: 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.

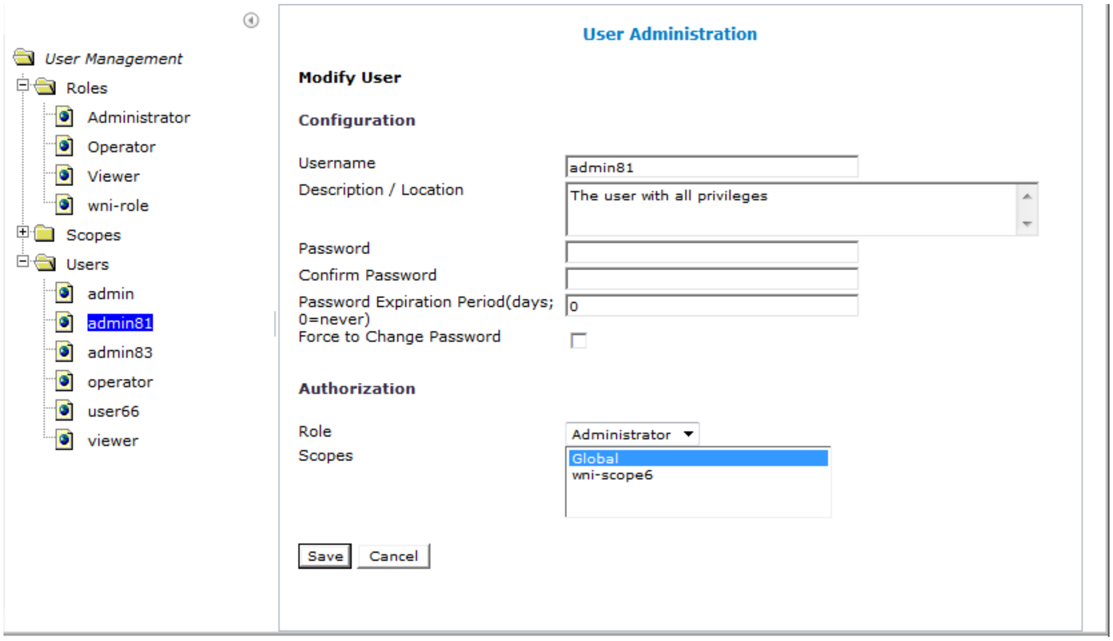

User Administration

Modify User

Configuration

Username: admin81

Description / Location: The user with all privileges

Password:

Confirm Password:

Password Expiration Period(days; 0=never): 0

Force to Change Password: □

Authorization

Role: Administrator

Scopes: Global (selected), wni-scope6

Figure 22: 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:

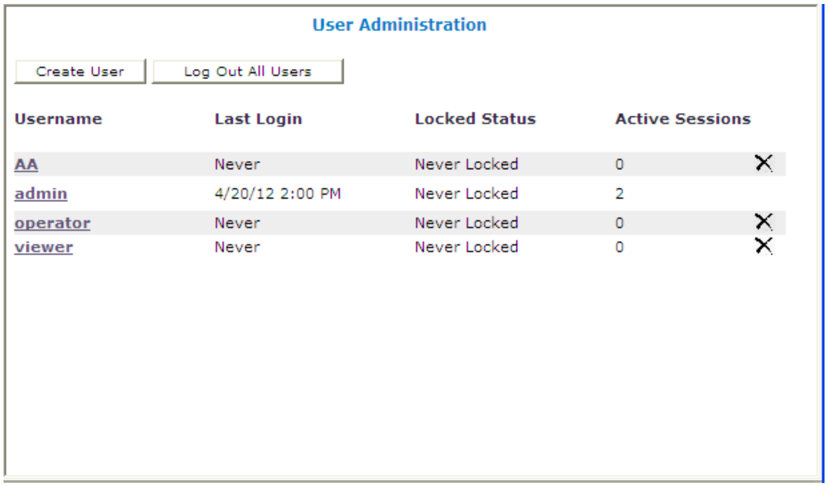

Username	Last Login	Locked Status	Active Sessions	
AA	Never	Never Locked	0	✕
admin	4/20/12 2:00 PM	Never Locked	2	
operator	Never	Never Locked	0	✕
viewer	Never	Never Locked	0	✕

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 the message “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 the message “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.

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:

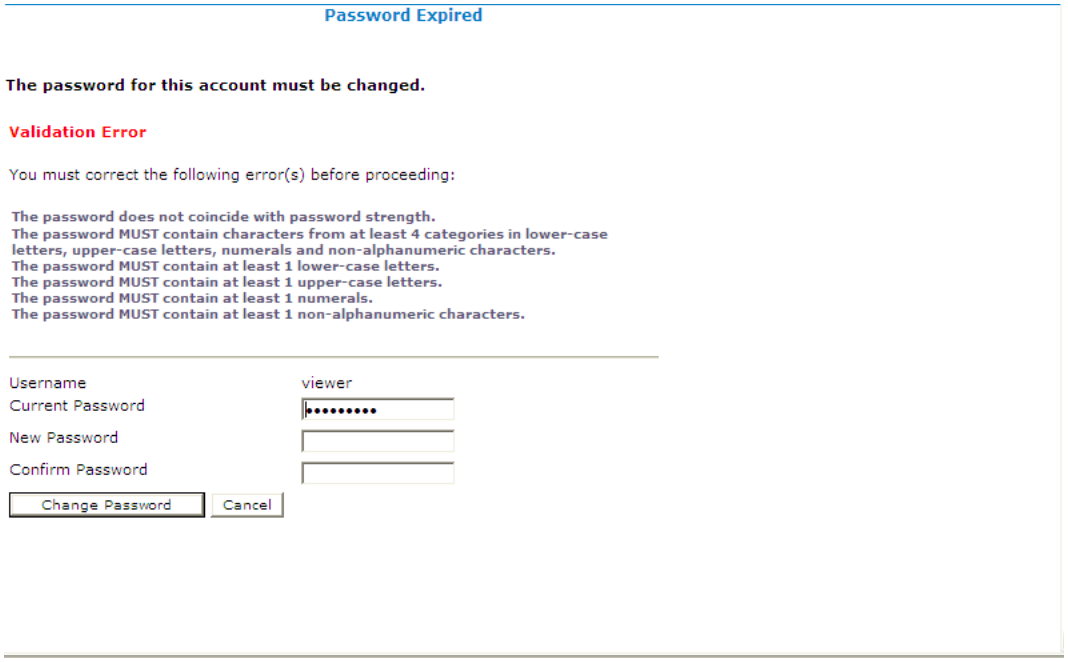

Password Expired

The password for this account must be changed.

Validation Error

You must correct the following error(s) before proceeding:

The password does not coincide with password strength.
The password MUST contain characters from at least 4 categories in lower-case letters, upper-case letters, numerals and non-alphanumeric characters.
The password MUST contain at least 1 lower-case letters.
The password MUST contain at least 1 upper-case letters.
The password MUST contain at least 1 numerals.
The password MUST contain at least 1 non-alphanumeric characters.

Username	viewer
Current Password	•••••••••
New Password	
Confirm Password	

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

Your password is changed.

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

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 23: 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 9: 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 SPR Subscriber Data** — Manage Subscriber Profile Repository servers
- **Manage Geo-Redundant MPE/MRA** — Manage georedundant MPE 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

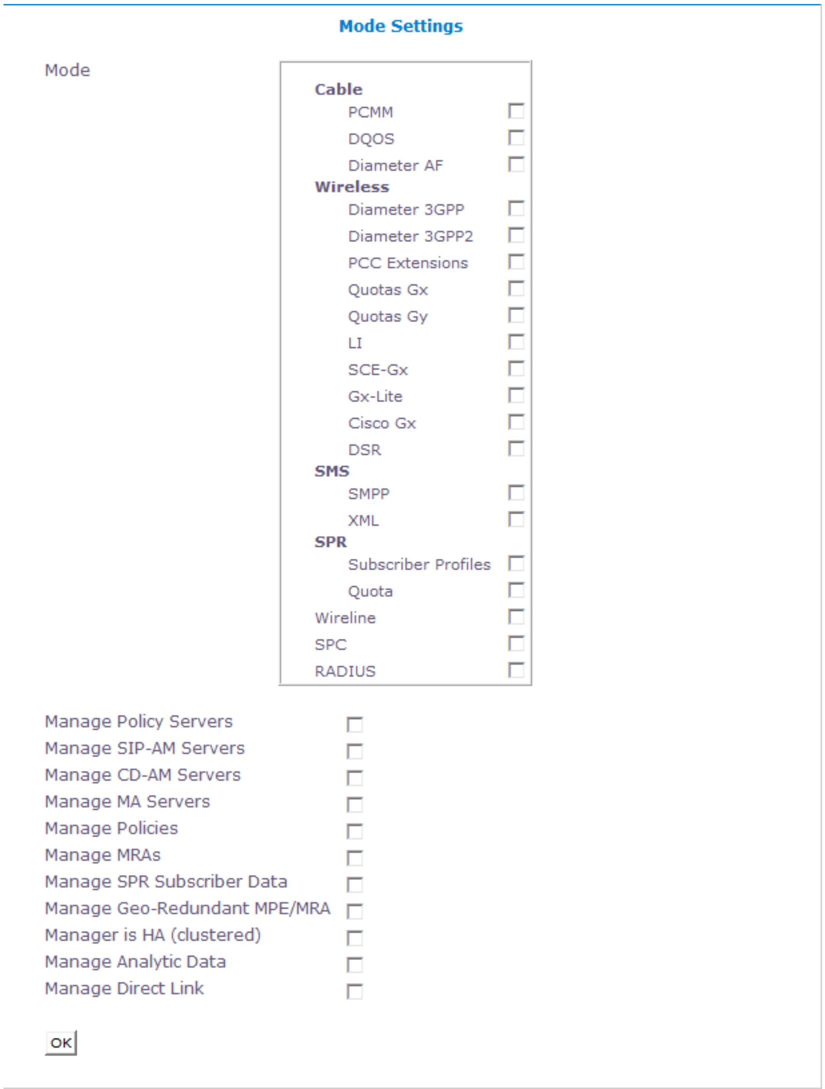

Mode Settings

Mode

Cable
- PCMM □
- DQOS □
- Diameter AF □

Wireless
- Diameter 3GPP □
- Diameter 3GPP2 □
- PCC Extensions □
- Quotas Gx □
- Quotas Gy □
- LI □
- SCE-Gx □
- Gx-Lite □
- Cisco Gx □
- DSR □

SMS
- SMPP □
- XML □

SPR
- Subscriber Profiles □
- Quota □

Wireline □

SPC □

RADIUS □

Manage Policy Servers □

Manage SIP-AM Servers □

Manage CD-AM Servers □

Manage MA Servers □

Manage Policies □

Manage MRAs □

Manage SPR Subscriber Data □

Manage Geo-Redundant MPE/MRA □

Manager is HA (clustered) □

Manage Analytic Data □

Manage Direct Link □

Figure 23: Mode Settings Page

Table 9: 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.
	Diameter AF	Supports Diameter AF functions.

Mode	Sub-Mode	Description
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.
	XML	Supports SMS using XML.

Mode	Sub-Mode	Description
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.

3GPP2

3rd Generation Partnership Project 2

A

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.

B

B-RAS

broadband remote access server

C

CMP

Configuration Management Platform

A centralized management interface to create policies, maintain policy libraries, configure, provision, and manage multiple distributed MPE policy server devices, and deploy policy rules to MPE devices. The CMP has a web-based interface.

D

Diameter

Protocol that provides an Authentication, Authorization, and Accounting (AAA) framework for applications such as network access or IP mobility. Diameter works in

D

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.

E

event

In Policy Management, an expected incident that is logged. Events can be used for debugging purposes.

F

failover

The capability to automatically switch to a redundant or backup server, system, or network when the previously active server, system, or network fails or terminates abnormally. In certain instances, however, automatic failover may not be desirable, and human intervention may be required to initiate the failover manually.

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

G

user with a graphic means for manipulating screen data rather than being limited to character based commands.

M

Multimedia Policy Engine

See MPE.

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

P

PCRF

Policy and Charging Rules Function

The ability to dynamically control access, services, network capacity, and charges in a network.

Maintains rules regarding a subscriber’s use of network resources. Responds to CCR and AAR messages. Periodically sends RAR messages. All policy sessions for a given subscriber, originating anywhere in the network, must be processed by the same PCRF.

policy and charging rules function

See PCRF.

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.

S

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.

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 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.