



PolicyCenter

CLI Command Reference

Legal Notice

Copyright © 2018 Symantec Corp. All rights reserved. Symantec, the Symantec Logo, the Checkmark Logo, Blue Coat, and the Blue Coat logo are trademarks or registered trademarks of Symantec Corp. or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners. This document is provided for informational purposes only and is not intended as advertising. All warranties relating to the information in this document, either express or implied, are disclaimed to the maximum extent allowed by law. The information in this document is subject to change without notice.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE. SYMANTEC CORPORATION PRODUCTS, TECHNICAL SERVICES, AND ANY OTHER TECHNICAL DATA REFERENCED IN THIS DOCUMENT ARE SUBJECT TO U.S. EXPORT CONTROL AND SANCTIONS LAWS, REGULATIONS AND REQUIREMENTS, AND MAY BE SUBJECT TO EXPORT OR IMPORT REGULATIONS IN OTHER COUNTRIES. YOU AGREE TO COMPLY STRICTLY WITH THESE LAWS, REGULATIONS AND REQUIREMENTS, AND ACKNOWLEDGE THAT YOU HAVE THE RESPONSIBILITY TO OBTAIN ANY LICENSES, PERMITS OR OTHER APPROVALS THAT MAY BE REQUIRED IN ORDER TO EXPORT, RE-EXPORT, TRANSFER IN COUNTRY OR IMPORT AFTER DELIVERY TO YOU.

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

www.symantec.com

10/4/2018

Command Line Overview

The command-line interface (CLI) provides a UNIX-like interface for accessing PolicyCenter.

Access the CLI

You can access the PolicyCenter CLI using one of the following techniques:

- Use the CLI Commands utility available in the browser interface. See "Access the CLI with the CLI Commands Utility" on page 5.
- Directly connect a computer to the PolicyCenter S-Series serial port, and then use a terminal emulation program to log in and access the CLI. See "Access the CLI Using a Direct Console Connection" on page 5.
- Connect the PolicyCenter S-Series serial port to a serial console/terminal server, and then telnet to the server. See "Access the CLI Using a Serial Console/Terminal Server" on page 6.



You cannot SSH or telnet to the PC-S200; only the above methods are supported ways for accessing the PolicyCenter CLI.

Look vs. Touch Access

Users who log in with the touch user name have access to all PolicyCenter commands and will see a # sign in the command prompt:

```
/default#
```

PolicyCenter lists the current configuration name before the # prompt. In the above example, /default is the current configuration. You can use the [config show](#) to see a list of configuration names and use the [config view](#) command to work on a specific configuration. For example:

```
/default# config view /legacy
```

```
/Legacy#
```

Users who log in with look access will see the following prompt:

```
Look Mode>
```

Look users cannot enter commands that modify the PolicyCenter configuration.

Note that a touch user is allowed touch access even if there is currently a touch UI session. This exception is in place to allow the serial console user the ability to end a UI session in which a user failed to log out and has locked out other users from logging in with touch access; use the **authentication session end** CLI command to close a hung UI session.



A touch user logging in to the serial console should always use the **authentication session show** command before making any configuration changes. If there are any current touch sessions in the UI, they should be closed so that two users won't be able to create conflicting configuration changes.

Command Usage Conventions

A few basic conventions apply to commands:

PolicyCenter CLI Commands

- Commands are not case sensitive – that is, you can use either uppercase or lowercase characters.
- A command can be abbreviated by entering the minimum number of characters required to uniquely distinguish it from other commands. For example, you can type *set pass* instead of *setup password*.
- Command syntax can be verified by typing one of the following:
help <command>
or
<command> ?
where <command> is the name of the command for which you want help.
- To issue multiple commands from a single command line, separate the commands with a semicolon (;) – for example, *setup show;version verbose*. The semicolon is the equivalent of pressing the Enter key.
- You can scroll through the command history by pressing the up and down arrows. You can also edit previously entered commands.


Typographical Conventions

Convention	Description	Example
Boldface	Commands	class show
[Square brackets]	Optional argument in a command line	version [verbose]
<angle brackets>	Required argument for which you will supply the value	set images remove <release_ID>
Pipe character ()	The "or" symbol in a command line— choose one of the options separated by the symbol	setup discover on off

Editing Previously Entered Commands

If you make a typing mistake in your command, you don't need to retype it – you can redisplay the command and edit it. This capability is available via SSH clients and with a direct console connection.

Function	Technique
Display a previously entered command	Press up arrow until the command you want is displayed
Scroll down through the command history	Press down arrow
Move cursor to the left	Press left arrow
Move cursor to the right	Press right arrow
Insert characters	Position cursor with arrows, start typing
Delete character to left of cursor	Press Backspace

 If the arrow keys aren't working, make sure your SSH client is emulating VT100 arrows. You may need to enable this option in your client.

Access the CLI with the CLI Commands Utility

The CLI Commands utility provides a command line into which you can enter CLI commands for loading PolicyCenter software and licenses, as well as enter other CLI commands.

To run CLI commands:

1. Click the **CLI Commands** link at the bottom of the PolicyCenter screen. The utility opens in a new tab or window.
2. Select the **Enable any CLI Command** check box.
3. In the **Command** field, type the CLI command.
4. Click **Run**.
5. When finished with the utility, close the tab or window.

Access the CLI Using a Direct Console Connection



You cannot SSH or telnet to the PC-S200; you must connect a null-modem cable to the PC-S200 to access the PolicyCenter CLI.

To access the PolicyCenter S-Series directly with a null-modem cable:

1. Using the provided null-modem serial cable, attach a workstation or laptop to the PolicyCenter S-Series serial port.
2. Start your terminal emulation program (such as HyperTerminal).
3. Verify that you have configured the program with the following values to communicate with the unit's console serial port:
9600 bps, 8 data bits, 1 stop bit, no parity, hardware flow control
4. If you are using a modem connected to the serial port, the modem must be set to: 9600 bps, 8 data bits, 1 stop bit, no parity, auto-answer (usually ATH1 in the standard Hayes command set), and DTR always on (usually a DIP switch setting). Check the modem manual for details.
5. Power on the unit, if you have not already done so. If the unit was already turned on, you will need to press Enter several times to make the connection.
6. When prompted, enter the user name (**touch** or **look**) and password (**touch** or **look** by default) .

When you successfully log in, you will see the unit's command-line prompt, for example `/default#`. The prompt indicates the name of the current configuration (`/default`, in this example). Use the [config view](#) command to switch to another configuration and [config edit](#) to modify a draft of the configuration.

Note that a touch user is allowed touch access even if there is currently an active touch UI session. This exception is in place to allow the serial console user the ability to end a UI session in which a user failed to log out and has locked out other users from logging in with touch access; use the **authentication session end** CLI command to close a hung UI session.



A touch user logging in to the serial console should always use the [authentication session show](#) command before making any configuration changes. If there are any current touch sessions in the UI, they should be closed so that two users won't be able to create conflicting configuration changes.

Access the CLI Using a Serial Console/Terminal Server

To access the PolicyCenter command-line interface (CLI) using a serial console/terminal server:

1. Using the provided null-modem serial cable, attach a serial console/terminal server to the PolicyCenter S-Series serial port on the back of the appliance. Make note of the identifying number assigned to the PolicyCenter S-Series appliance.
2. Telnet to the serial terminal, including the identifying number assigned to the PolicyCenter S-Series, for example, *telnet 192.0.2.41 10032*

When you connect successfully, you will be prompted to log in.

3. Enter the user name (**touch** or **look**) and password (**touch** or **look** by default).

When you successfully log in, the look user sees the following command-line prompt: `Look Mode>` and the touch user sees a `#` prompt that includes the current configuration name, for example `/default#`. Use the [config view](#) command to switch to another configuration and [config edit](#) to modify a draft of the configuration.

Note that a touch user is allowed touch access even if there is currently an active touch UI session. This exception is in place to allow the serial console user the ability to end a UI session in which a user failed to log out and has locked out other users from logging in with touch access; use the **authentication session end** CLI command to close a hung UI session.



A touch user logging in to the serial console should always use the [authentication session show](#) command before making any configuration changes. If there are any current touch sessions in the UI, they should be closed so that two users won't be able to create conflicting configuration changes.

actionfile library

Show the current portfolios of adaptive response action files available for distribution from PolicyCenter to individual PacketShapers.

```
actionfile library [verbose]
```

The **actionfile library** command shows the name of the available portfolios only. Use **actionfile library verbose** to view the names of all the action files within each portfolio.

actionfile prescribe

Prescribe a group of adaptive response action files by portfolio name. Use the **actionfile library** command to determine available action file portfolios.

```
actionfile prescribe <portfolio> default|none|show
```

<i><portfolio></i>	Name of portfolio. A portfolio is any sub-folder of <i>publish/action</i> that contains a group of action files.
default none show	On a child configuration, the default option allows that child configuration to inherit its portfolio of action files from its parent configuration. (On a parent configuration, the default option sets the prescription to unconfigured.) Specify none if the configuration should not inherit its portfolio. The show option shows the configuration's current prescribed portfolio of action files.

actionfile subscribe

Configure when and how often PacketShapers assigned to a PolicyCenter configuration update their portfolio of adaptive response action files.

```
actionfile subscribe asap|scheduled|default
```


The **actionfile subscribe** command has the following options:

<i>asap</i>	PacketShapers assigned to the configuration will automatically update their action file portfolio as soon as an updated portfolio is prescribed.
<i>scheduled</i>	PacketShapers assigned to the configuration will wait for the actionfile sync command before downloading the prescribed portfolio of files.
<i>default</i>	If a child configuration is set to default , the child configuration inherits its action file subscription behavior from its parent. If a parent configuration is set to default , units assigned to the parent configuration will automatically update their action file portfolio as soon as an updated portfolio is prescribed.

actionfile sync

For units in shared mode only

Issue this command from an individual PacketShaper to immediately download adaptive response action files prescribed for the unit's PolicyCenter configuration. This command is only required when the PolicyCenter configuration prescription mode has been set to **scheduled** with the [actionfile subscribe](#) command.

 It is not necessary to issue this command if the prescription mode is currently in its default state, or has been set to **asap** with the **actionfile subscribe** command.

```
actionfile sync <seconds>
```

If you include the optional *<seconds>* value, the **actionfile sync** operation runs for the specified number of seconds.

agent action

Delete an adaptive response action file, temporarily disable or re-enable an existing action file, or modify the value of an existing parameter in the current configuration. Note that this command will not create a new action file, or add a new parameter to an existing action file.

```
agent action <name> green|red [on <filename>] | [off] | [delete] | [parm <parm-name> <parm-value>] | [resetparms]
```

<name>	Name of the agent. If the agent name has a space, the words must be entered within quotation marks, for example "My Agent". If the agent name is a single word, the quotation marks are not necessary.
green red	Action file will trigger when the green or red threshold is crossed
<[on <filename>] [off] [delete]	Specify one of the following: on: Enables the action file. Specify the name of the action file you want to associate with the agent with the <i><filename></i> variable. off: Disables the action file delete: Deletes the action file specification for the agent. The action file is no longer associated with the agent, but the action file is not removed from the unit or PolicyCenter.
[parm <parm-name><parm-value>]	Specify the following: <i><parm-name></i> : The name of the action file parameter being modified <i><parm-value></i> : The new value of the parameter
[resetparms]	Specify this operation only when action file parameters have been edited and need updating. agents will not recognizes new action file parameters unless the action file is reset with this variable.

Before you can issue any other agent actionfile commands, you must first issue the command **agent actionfile <name> green|red on <filename>**

to associate an action file with the agent. You may then issue any of the following commands (see the table above for an explanation of variables):

```
agent action <name> green|red off
```

```
agent action <name> green|red delete
```

```
agent action <name> green|red resetparms
```

```
agent action <name> green|red parm <parm-name> <parm-value>
```

For Example:

```
agent action "Packet Drops" green on actnfile.cmd
```

```
agent action "Packet Drops" green parm ClassName /outbound/Citrix
```

agent createdefaults

Recreate the default set of agents in the current configuration. The adaptive response feature must be [enabled](#) before you can create default agents with this command.

```
agent createdefaults
```

Note that this command will not overwrite any existing default agents that you may have customized, nor does it remove any new agents you may have created.

agent delete

Delete an existing adaptive response agent in the current configuration. Scoring and status information for the agent will no longer appear in the agent pop-up window on the unit's **info** page or the PolicyCenter **Configurations** page.

```
agent delete <name>
```

where *<name>* is the name of the agent. If the agent name has a space, the words must be entered within quotation marks, for example, *High Bandwidth Host* . If the agent name is a single word, the quotation marks are not necessary.

agent interval


Set an evaluation interval (in minutes) for an adaptive response agent in the current configuration. An evaluation interval determines how often the agent checks the status of its target.

```
agent interval <name> <interval> | default
```

where *<name>* is the name of the agent. If the agent name has a space, the words must be entered within quotation marks, for example, *High Bandwidth Host* . If the agent name is a single word, the quotation marks are not necessary. Specify the interval in minutes, or enter **default** for the default evaluation interval. The maximum evaluation interval allowed is 99999 minutes; the minimum is 1 minute.

agent new

Create a new adaptive response agent in the current configuration. Note that this command creates a new agent based on one of the agent templates, but does not allow you to specify parameter values. Once you have created a new agent, issue the command [agent parm](#) to change the parameter values from their default settings. Each PacketShaper or PolicyCenter configuration can have a maximum of 32 agents.

 Some agent templates do not allow multiple instances. If you want to create a new agent from the following templates, first delete the existing agent from that template from your unit or PolicyCenter configuration.

- High Bandwidth Host
- New Application
- High Bandwidth New App
- Memory Allocation
- Unit Limits
- System Load

agent new <name> <template>

<name>	Name you want to assign to the agent. An agent name can have up to 32 alphanumeric characters, including -, _, and . (period). If the agent name has a space, the words must be entered within quotation marks, for example, My Agent. If the agent name is a single word, the quotation marks are not necessary.
<template>	<div>Specify one of the following agent templates:</div> <div>Class ME Variables</div> <div>Default Traffic</div> <div>Failed Flow Ratio</div> <div>High Bandwidth Host</div> <div>High Bandwidth New App</div> <div>Host Info Variables</div> <div>Link ME Variables</div> <div>Memory Allocation</div> <div>New Application</div> <div>NFPM Failed Flows</div> <div>NFPM Side Unknown</div> <div>Partition ME Variables</div> <div>Partition Utilization</div> <div>System Load</div> <div>Traffic Performance</div> <div>Unit Limits</div>

Example:

agent new testagent "Class ME Variables"
agent new "agent two" "Class ME Variables"

agent off

Disable an existing adaptive response agent in the current configuration, without deleting it. The agent will no longer return values or create new reports, yet it can be re-enabled at any time with the [agent on](#) command.

```
agent off <name>
```

where <name> is the name of the agent. If the agent name has a space, the words must be entered within quotation marks, for example, *My Agent*.

agent on

Enable an existing adaptive response agent that has been disabled in the current configuration. The agent will once again return values and create new reports.

```
agent on <name>
```

where <name> is the name of the agent to be turned on. If the agent name has a space, the words must be entered within quotation marks, for example *"My Agent."*

agent override

For PolicyCenter/ PacketShapers in Shared Configuration Mode

Override an adaptive response agent that a child configuration inherits from a parent configuration, so the agent may be modified on the child configuration. Inherited agents cannot be modified until they are overridden.

```
agent override <name>
```

where <name> is the name of the agent. If the agent name has a space, the words must be entered within quotation marks, for example, *My Agent*.

agent parm

Specify the parameter values for an adaptive response agent. The agent must have been already defined with the [agent new](#) command.

```
agent parm <name> [<parm-name> <parm-value> | default]
```

<name>	Name of the agent. If the agent name has a space, the words must be entered within quotation marks, for example, <i>My Agent</i> . If the agent name is a single word, the quotation marks are not necessary.
<parm-name>	The name of the parameter or threshold to be set. Each agent is based on a template which has its own parameters.
<parm-value>	The parameter value for <parm-value>, or enter default for the parameter's default value. For information on the acceptable and default parameter values, see the links above.
default	Return the agent to its default values

PolicyCenter CLI Commands

Examples:

The first example shown below changes the *ClassName* parameter for the agent *testagent* so that agent will now monitor the class */Inbound/Citrix*.

agent parm testagent ClassName /Inbound/Citrix

If you don't specify any parameters, the **agent parm** command shows current and default parameter settings for the specified agent.

agent parm "System Load"

Score Parms
RedThreshold 95(Default: 95)
GreenThreshold 90(Default: 90)

agent show

Show data for one or many adaptive response agents, including information on the agent type and category, the corresponding plug-in file, any incident report files, and the agent version number.

The **agent show** command will show values with a timestamp based upon the end of the evaluation interval. This is different from the **measure dump** command, which shows values with a timestamp that reflects the beginning of the time interval.

```
agent show [name <name> | templates | [result <score-result>] | [feedback [<unitSN> <name>]]]
```

<name>	Name of the agent. If the agent name has a space, the words must be entered within quotation marks, for example <i>My Agent</i> . If the agent name is a single word, the quotation marks are not necessary.
templates	Show a list of available adaptive response agent templates
result <score- result>	If the agent is unable to measure its target, the output of the agent show name "<name>" command will display additional <i>Result</i> category of data showing an explanation of the error and the error code. You can determine the meaning of an error code by issuing the command agent show result <score-result> . See the example below.
<unitSN> <name>	Used with the agent show feedback command, the <unitSN> parameter is serial number of the unit for which want to view agent data and feedback. This parameter is optional—if you do not specify a unit, the agent show feedback command will show data for all agents (and when issued from PolicyCenter, all agents for all configurations). The <name> parameter is the name of the agent. If no data exists for a new agent, or there is no agent with the specified name, this command will return the output "No feedback available."

Examples:

The **agent show** command displays agent information for an individual PacketShaper, or when issued from the PolicyCenter CLI, agent information for the configuration you are editing. This information includes data on whether or not the agent has been enabled, the name of the agent, and the last score information.

For PolicyCenter configurations, an **I** to the left of the agent name indicates that the configuration has inherited that agent from a parent configuration. An **O** to the left of the agent name indicates that the configuration has a local override of an agent that supersedes the agent it inherits from its parent. An exclamation point (!) beside the agent name indicates a configuration error. Last Score Information includes the latest value measured by the agent, its status color, and the time and date of the measurement.

agent show

Status On

Agent Name	Status	Last Score Information

--		
Class ME Variables agent	On	0 Yellow Wed Jan 12 02:03:00
2005 PST		
High Bandwidth New App	On	New score value in 51m 13s.
Inbound Default Traffic	On	1 Green Wed Jan 12 02:03:00
2005 PST		
Outbound Default Traffic	On	0 Green Wed Jan 12 02:03:00
2005 PST		
Partition Utilization agent	On	0 Green Wed Jan 12 02:03:00
2005 PST		
Spoofing - Client	On	0 Green Wed Jan 12 02:03:00
2005 PST		
Spoofing - Server	On	0 Green Wed Jan 12 02:03:00
2005 PST		
Syn Attack - Failed Flows	On	0 Green Wed Jan 12 02:03:00
2005 PST		
Traffic Performance agent	On	1 Red Wed Jan 12 02:03:00 2005
PST		
*NT = No template found for agent.		
*NF = Either an action or incident file not found.		

agent show templates

PacketShaper# agent show templates

Template Name	PlugIn File	Incident Report File	Ver Num	Category

--				
Quota Bandwidth Host	-	hostquot.cmd	1.0	Hosts
Host Info Variables	-	hostvar.cmd	1.0	Hosts
Failed Flow Ratio	-	ffratio.cmd	1.0	Hosts
NFPM Failed Flow	-	syn.cmd	1.0	Hosts

PolicyCenter CLI Commands

NFPM Side Unknown	-	spoof.cmd	1.0	Hosts
Link ME Variables	-	melink.cmd	1.0	User Event
Emulation				
Partition ME Variables	-	meptn.cmd	1.0	User Event
Emulation				
Class ME Variables	-	mecclass.cmd	1.0	User Event
Emulation				
High Bandwidth New App	-	susapp.cmd	2.0	Application Health
New Application	-	newapp.cmd	1.0	Application Health
Default Traffic	-	dflttraf.cmd	2.0	Application Health
High Bandwidth Host	-	sushost.cmd	3.0	Hosts
Traffic Performance	-	trafperf.cmd	2.0	Network Health
Partition Utilization	-	ptnutl.cmd	3.0	Network Health
Memory Allocation	-	sysmem.cmd	1.0	Unit Health
System Load	-	sysload.cmd	1.0	Unit Health
Unit Limits	-	syslimit.cmd	1.0	Unit Health



The incident report files described in the above output above are the files used by each agent to create incident reports. Incident report files are different from action files, as they are used only to generate drill-down incident reports. *Do not edit or modify incident report files in any way.* Any modifications to an agent's incident report file could stop new reports from being generated for that agent.

agent show name "inbound default traffic"

Agent Name	Inbound Default Traffic
Status	On
Template Info	
Template Name	Default Traffic
VerNum	2.0
Category	Application Health
Description	This agent monitors the rate (avg-bps) of the default traffic
traffic	class. This agent can alert you when the amount of
traffic not	classified (falling into 'default') is too great. This
agent	must be used with a 'default' (i.e., /Inbound/Default)
traffic	class.
	Threshold Units: % of bandwidth on the partition
	Action File Variables: \$class-id, \$avg-bps
Plugin File	-
Incident File	9.258/agent/cmd/dflttraf.cmd
MultiInstance	Allowed
Interval	1 minute(s)
Score Params	

```

RedThreshold      15(Default: 15)
GreenThreshold    7(Default: 7)
ClassName         /Inbound/default(Default: /Inbound/default)
Color Mappings
  Green           Score < 7
  Red             Score > 15
Last Score Status
  Value           1
  Color           Green
  Start time      Wed Jan 12 02:05:00 2005 PST
  Finish time     Wed Jan 12 02:06:00 2005 PST
New score value in 39s.

```

If the agent in the example above had a status color of blue, the *Last Score Status* category would display additional *Result* information with an explanation of the error and an error code. The example below shows the *Last Score Status* displaying this additional *Result* output.

```

Last Score Status
Value           0
Color           Blue
Result          Agent score parm not found. (score-
                result: 4569) <-----
Start time      Mon Jun 19 08:19:00 2005 PST
Finish time     Mon Jun 19 08:20:00 2005 PST
New score value in 45s.

```

You can determine the meaning of a *Result* error code with the **agent show result <score-result>** command. The following example displays information for error code 4569.

```

agent show result 4569
Agent score parm not found.

```

This next example shows the resultant output when the command **agent show feedback** is issued for a PolicyCenter configuration. (If this command was issued for a unit configuration, it will show only the agents on that individual unit.) The Feedback Information includes the latest value measured by the agent, its status color, and the time and date of the measurement.

agent show feedback

Unit	Agent Name	Feedback Information
065-10000193	ClassMeVar	1 Green Mon Jul 19 22:33:01 2005 LST
065-10000193	Hosts	1 Green Mon Jul 19 22:38:01 2005 LST
065-10000179	ClassMeVar	2 Yellow Mon Jul 19 22:33:01 2005 LST

PolicyCenter CLI Commands

```
065-      7 Yellow Mon Jul 19
10000179 Hosts      22:38:01 2005 LST

065-      0 Green Mon Jul 19
10000238 PacketDrops 23:00:06 2005 LST

065-      1 Green Mon Jul 19
10000238 Hosts      23:00:06 2005 LST
```

Issue the **agent show feedback** command with the `<unit#>` and `<name>` parameters to display data for one agent on a single unit.

agent show feedback 025-10000210 "FTP Partition Over Limit"

```
Score Feedback:
  Score      18073
  Color      Red
  Category   User Event Emulation
  Start Time  Fri Oct 15 08:56:02 2005 PDT
  Finish Time Fri Oct 15 09:56:02 2005 PDT

Incident Report Feedback:
  File Output 9.258/agent/cmd/complete/155646.htm
  Result      Success.
  Finish Time  Fri Oct 15 09:56:02 2005 PDT
```

authentication organization add

PolicyCenter allows network administrators to define up to 256 different *organizations*, groups of configurations, and a list of *users* that can access those configurations.

Only PolicyCenter administrators can view and manage all units and configurations in the PolicyCenter configuration tree. If you want every PolicyCenter user to have complete access to all PolicyCenter configurations and units, you can make every user a PolicyCenter administrator. However, you may find that not all users need such a complete level of access.

You can restrict a user's access to a specific set of PolicyCenter configurations and units by creating a new organization, specifying the configurations and units the users in that organization are allowed to view or manage, then adding users to the organization.

PolicyCenter administrators can issue this command to add new organizations.

```
authentication organization add <organization>
```

where `<organization>` is the name of the new organization. An organization name can be comprised of up to 32 alphanumeric characters, periods, underscores, and dashes. The first character of the name must be a letter.

Once you have created a new organization, you can add new users to the organization with the [authentication user add](#) command.

authentication organization delete

PolicyCenter administrators can issue this command to permanently delete existing organizations. This command will also delete the user records of all users assigned to this organization, so they will no longer be able to access PolicyCenter. To temporarily disable all users in this organization while retaining their user information, issue the command [authentication organization disable](#).

```
authentication organization delete <organization>
```

where *<organization>* is the name of the organization to be deleted.

Example:

```
authentication organization delete org_2
```

Deleting "org_2" would also delete roles and users within this organization.
Continue with the deletion of this organization? (YES): yes

authentication organization disable

PolicyCenter administrators can issue this command to temporarily disable one or all other existing organizations. None of the users in a disabled organization will be allowed to access PolicyCenter configurations or units, but their user information will be retained. You can re-enable the organization and restore its users' access to PolicyCenter at any time with the [authentication organization enable](#) command.

```
authentication organization disable <organization>|all
```

where *<organization>* is the name of the organization to be temporarily disabled. Select the `all` option to disable all PolicyCenter organizations except for the default PC organization, which cannot be disabled or deleted.

PolicyCenter also allows you to disable individual users in an organization, while keeping the rest of the organization active. To disable individual users in an organization, use the command [authentication user disable](#).

See also:

[authentication organization delete](#)

authentication organization enable

PolicyCenter administrators can issue this command to re-enable one or more organizations that were temporarily disabled with the [authentication organization disable](#) command.

```
authentication organization enable <organization>|all
```

where *<organization>* is the name of the new organization. Select the `all` option to enable all PolicyCenter organizations.



PolicyCenter allows you to disable an entire organization of users and also disable specific individual users within an organization. This command will re-enable an organization, but will not re-enable a user that was individually disabled with the [authentication user disable](#) command.

authentication organization rename

Rename an existing organization. Any users assigned to the configuration will remain assigned to the organization after it is renamed. You must have touch role access to the default *PC* configuration to issue this command.

```
authentication organization rename <name> <newname>
```

where:

<name>	Current organization name
<newname>	New name for the organization. An organization name can be comprised of up to 32 alphanumeric characters, periods, underscores, and dashes. The first character of the name must be a letter.

See also: [authentication organization show](#)

authentication organization show

View details for a PolicyCenter organization, or all organizations.

```
authentication organization show [<organization>]
```

PolicyCenter administrators can issue this command to view details for all PolicyCenter organizations. Organization managers with touch access to any other organization can view details for that one organization only.

```
authentication organization show
```

```
Organization: PC
State: enabled
```

```
Organization: Marketing_2
State: enabled
```

```
Organization: Sales_1
State: enabled
```

```
Found 3 organizations
```

authentication session end

Terminate the current active session of another PolicyCenter user. Only PolicyCenter administrators may end a user session.

```
authentication session end <id> | <username>
```

where

<id>	User's unique session id. To view the current session IDs for each user currently logged in to PolicyCenter, issue the command authentication session show .
<username>	The user name of the user whose session you want to terminate.

authentication session show

Display information for current user sessions and attempted logins. Only organization managers with touch role access to an organization may view session information for to that organization. PolicyCenter administrators can view information for all users.

```
authentication session show
```

For example:

```
auth session show
```

ID	Stat	Age	Idle	Limit	Type	Access	User Name
44c9349b	logged in	30 secs	2 secs	60 mins	WUI	look	pbosten
44c93480	logged in	112 secs	45 secs	60 mins	WUI	look	lrose
44c93353	logged off	411 secs	0 secs	60 mins	CLI	touch	(admin)

Column	Description
ID	Identification given to the user session
Stat	The status of the session: logged in– the user has logged in logged out–the user has logged out
Age	Length of time the session has been active that is, the amount of time since the user logged in
Idle	Amount of time since the user gave a command; whenever a user gives a command, the idle value is reset to zero
Limit	Amount of time a session is idle before the user will be timed out and logged off; for example, if the limit is 60 minutes, a user will get logged off when no commands are given for a 60-minute period.
Type	Type of interface used: CLI (command-line interface), or WUI (web user interface)
Access	User's role for accessing PolicyCenter; Look or Touch
User Name	Name of the user who logged into the session

authentication show

Display information about the current user session.

```
authentication show
```

For example:

```
auth show
User name:      admin (Admin Admin)
Login time:     2015-10-06 11:06:07 PDT
Prior logout:   2015-10-05 14:01:17 PDT
Organization:   PC
Role:           Touch
Access Level:   look (for this session)
```

authentication user add

Add a new user to an organization. Only PolicyCenter administrators and organization managers with touch role access to their organization may add a new user to that organization. PolicyCenter supports up to 512 different user accounts.

```
authentication user add <username> <organization> <role> <firstname> <lastname>
[<password>]
```

where:

<username>	Login user name for the user. A user name can be comprised of up to 32 alphanumeric characters, periods, underscores, and dashes. The first character of the user name must be a letter. Each PolicyCenter user name must be unique; users in different organizations cannot have the same user name.
<organization>	Name of the organization to which the user will be added
<role>	Specify either look or touch to select a role for the new user. Users with touch access can view and modify settings for the configurations and units assigned to their organization through the PolicyCenter web-browser or command-line interfaces, or via the web-browser or command-line interfaces of their individual assigned units. Users with look access can only view these settings in PolicyCenter, but cannot modify them or access the individual units.
<firstname> <lastname>	New first and last names for the user. Names cannot have spaces; compound names will require a dash or underscore character (e.g., <i>Ann-Marie</i> or <i>Van_Patten</i>).
<password>	Specify a login password for the user. A password can be up to 19 characters long and include all printable characters, including spaces, periods, underscores, and dashes.

authentication user delete

Permanently delete a user from an organization. You must have touch role access to an organization in order to delete any of its users.

When you delete a user currently logged in to PolicyCenter, that user's session is terminated immediately. Note, however, that immediately terminating another user's PolicyCenter session can cause configuration errors if the user was in the process of making a configuration change.

```
authentication user delete <username>
```

where *<username>* is the login name user name for the user you want to delete. This command completely removes the user's personal record from PolicyCenter. To temporarily disable an individual user while retaining his or her user information, use [authentication user disable](#).

authentication user disable

Temporarily disable a user's login access to PolicyCenter by disabling their user name and password. This command does not delete user records from an organization, so you can re-enable these users at any time without having to recreate their user records. (To permanently remove a user record from an organization, issue the command [authentication user delete](#)) You must have touch role access to the user's organization to issue this command.

```
authentication user disable <username>|[all <organization> <role>]
```

where:

<i><username></i>	User's login user name for accessing PolicyCenter
<i><organization></i>	Name of the user's organization
<i><role></i>	Specify either look or touch to disable all look or touch users within an organization

See also: [authentication user enable](#)

authentication user enable

Enable either an individual user or all users with a specific role within the organization. This command reactivates users who were temporarily disabled with the [authentication user disable](#) command. You must have touch role access to the user's organization to issue this command.

```
authentication user enable <username>|[all <organization> <role>]
```

where:

<i><username></i>	User's login user name for accessing PolicyCenter
<i><organization></i>	Name of the user's organization

PolicyCenter CLI Commands

<role>	Specify either look or touch to disable all look or touch users within an organization
--------	--

For example:

```
authentication user enable jsmith
authentication user enable all org_2 look
```

Though this command will enable individual users or all users with a specific role, if the organization itself is disabled, these users will still be unable to access PolicyCenter.

See also: [authentication organization enable](#)

authentication user name

Change an existing user's first and last names in their user record. You must have touch role access to the user's organization to issue this command.

```
authentication user name <username> <firstname> <lastname>
```

where:

<username>	User's login user name for accessing PolicyCenter
<firstname> <lastname>	New first and last names for the user. Names cannot have spaces; compound names will require a dash or underscore character (e.g., <i>Ann-Marie</i> or <i>Van_Patten</i>).


To add a new user to an organization, use the command [authentication user new](#)

authentication user password

Modify a user's login password. A password can be up to 19 characters long and include all printable characters, including spaces, periods, underscores, and dashes.

```
authentication user password <username> [<password>]
```

where <username> is the login name for the user, and <password> is the new login password for the user.



If you modify the admin user password and forget the password, you can restore the default password (admin) by logging in from the serial console as the touch user and then issuing the following command:

```
authentication user password admin admin
```

authentication user rename

Change the user name for an existing PolicyCenter user. You must be a PolicyCenter administrator or have touch role access to the user's organization to issue this command.

```
authentication user rename <user> <newname>
```

where:

<user>	User's current login user name for accessing PolicyCenter
<newname>	New user name for the user. A user name can be comprised of up to 32 alphanumeric characters, periods, underscores, and dashes. The first character of the user name must be a letter.

To change a user's first and last names in their user record, use the command [authentication user name](#).

authentication user set

Assign a new role for the user. Users with touch access can view and modify settings for the configurations and units assigned to their organization through the PolicyCenter web-browser or command-line interfaces, or via the web-browser or command-line interfaces of their individual assigned units. Users with look access can only view their configuration settings, but can neither modify them nor access the individual units via PolicyCenter.

```
authentication user set <username> <role>
```

where:

<username>	User's login user name for accessing PolicyCenter
<role>	Specify either look or touch to select a role for the user

authentication user show

Show detailed user records for an entire organization, or a single user. You must have touch role access to the configuration to issue this command.

```
authentication user show [<username>] [{all <organization> [<role>]}
```

where:

<username>	User's login user name for accessing PolicyCenter
<organization>	Name of the user's organization
<role>	Specify either look or touch to view just those users within an organization with a look or touch role.

```
authentication user show exampleuser
```

```
login name: exampleuser (Joe Smith)
Login time: 2006-03-13 12:30:56 Pacific Standard Time
Logout time: 2006-07-18 18:06:17 Pacific Daylight Time
Organization: Retailer2
Role: Touch
```

banner show

Display the messages, such as file distribution errors, that are initially shown after logging into PolicyCenter. You can use the **banner show** command to display all of the appliance's configuration errors, warning messages, and notices. (This same information is displayed in the Info tab of the management console.)

```
banner show [verbose]
```

The **verbose** option displays additional information, such as the date and time and the type of message (*notice*, *warn*, etc.).

For example:

```
banner show verbose
```

```
/filedist/reboot, 29 Oct 15 07:02, notice: Rebooting because of image or plugin change.
```

cat

Display the contents of a file.

```
cat [-n] <filename>
```

where **-n** numbers the output lines.

cd

Change your current directory.

```
cd <dir>
```

For example, type **cd 9.258/** to change to the 9.258 folder.

class capture-ids

Creates a text file named classids.txt that contains a list of all well-known class identification values. The classids.txt file is located in the 9.256/log folder. This command is useful when using SNMP—the class ID is the index into tables of real-time class and partition data. For example, the well-known ID for /Inbound is 1 and the ID for /Outbound is 2.

```
class capture-ids
```

If you use the **more** command to view the contents of this file, a list appears with the class ID next to each class name. This list includes all classes that can be auto-discovered not just the ones currently in the traffic tree. Part of the ID list appears below.

```
1 /Inbound
2 /Outbound
3 /Inbound/Inside
4 /Inbound/Outside
5 /Inbound/Default
6 /Inbound/Global
7 /Inbound/Global/IP
```



```

8 /Inbound/Global/TCP
9 /Inbound/Global/UDP
10 /Inbound/Global/Miscellaneous
11 /Inbound/Global/DECnet
12 /Inbound/Localhost
13 /Inbound/SameSide
15 /Inbound/OutsideVPNTunnel
50 /Outbound/Inside
51 /Outbound/Outside
52 /Outbound/Default
53 /Outbound/Global
54 /Outbound/Global/IP
55 /Outbound/Global/TCP

```

class category

Assign a traffic class in the current configuration to a host accounting category. (See [host accounting categories](#) for details on creating the categories.) Once you have assigned a class to a category, the bytes sent and received for the class will get tallied into the assigned category for both the source and destination hosts.

```
class category <tclass> none|<category-name>
```

You can assign multiple classes to each category, if you like. The <tclass> must be a leaf class; that is, you cannot assign a category to a class that has any child classes.



You cannot create a child class after the parent has been assigned to a host accounting category.

To remove a traffic class from a host accounting category, use:

```
class category <tclass> none
```

class copy

Copy a traffic class in the current configuration, and its children, to another parent in the traffic tree.

```
class copy <tclass><new_parent> [children]
```

Specify the explicit path and class name for the traffic class to be copied and the receiving parent traffic class. For example:

```
class copy /inbound/HTTP/Gifs /inbound/HTTP/Graphics
```



Any defined top talkers and top listeners are not copied with a traffic class.

class criteria

Certain services, such as Citrix and Oracle, can be further classified by application-specific criteria. For example, you can create a traffic class for a specific Citrix application or an Oracle database. You can use the **class criteria** commands to

PolicyCenter CLI Commands

display the attributes that can be specified in a matching rule for these applications and to discover the values that can be specified for the attributes.

class criteria attributes|recent|track

attributes	Display the available application-specific criteria
recent	Show recently tracked criteria values for a class
track	Enable or disable criteria tracking for a class

The application-specific criteria format in a matching rule is:

`<application>:<attribute>:<value>`

where `<application>` and `<attribute>` are as described in the table below, and `<value>` is specific to your configuration and classification requirements.

This table shows available applications, attributes, and sample values.

Application	Service Type	Attribute	Example of Value
Citrix	Citrix-ICA	application client priority	PeopleSoft pat-pc 0
DCOM	DCOM	UUID	1cbcad78-df0b-4934-b558-87839ea501c9
DICOM	DICOM	server client	DICOM_STORAGE DICOM_ECHO
FTP	FTP-Data-Clear	FileName	*.mp3
Web	HTTP	host url content-type user-agent	207.78.98.18 /Images/*.jpeg image/gif Mozilla/4.0
HTTP-Tunnel	HTTP-Tunnel	host port	207.78.98.18 80
ICMP	ICMP	type	echo
NNTP	NNTP-Clear	GroupName	alt.binaries.* microsoft.public.games

Application	Service Type	Attribute	Example of Value
Oracle	Oracle-netv2	dbname	corp
PostgreSQL	PostgreSQL	dbname	corp
RTCP	RTCP-I	encoding media clock	GSM v 8000
RTP	RTP-I	encoding media clock to from user-agent source destination	dynamic, GSM v 8000 +12125551212@server.phone2net.com agb@bell-telephone.com *X*Lite* 207.78.98.18 207.78.98.18 Additional Information <ul style="list-style-type: none"> For SIP attributes, you can enter a substring of the attribute. For example, to match all Motorola models, you can enter Motorola for the user-agent criteria. The asterisk (*) wildcard is supported for user-agent.
SMTP	SMTP-Clear	SenderEmail	*@bluecoat.com
SOAP	SOAP-HTTP	host url content-type user-agent	207.78.98.18 /Images/*.jpeg image/gif Mozilla/4.0
SSL	SSL	common name	optionslink.etrade.com
WAP	WAP	WAPURI	*.bluecoat.com



In order to add an application-specific matching rule to a class, the class' service type must be the one indicated in the Service Type column above. For example, to classify by Oracle database name, the class must be based on the Oracle-netv2 service.

You can use the **class criteria** commands to identify the specific values to use in application-specific matching rules. First, you use [class criteria attributes](#) to get a list of applications and attributes that can be used in matching rules. Next, you use [class criteria track](#) to enable tracking on a specific class. Then, you use [class criteria recent](#) to see a list of recent

PolicyCenter CLI Commands

values for the class; the output will provide you with the information you need to create an application-specific matching rule.

The following example shows how you can use the **class criteria track** and **class criteria recent** commands to identify the SIP user-agent for RTP-I traffic:

```
PacketShaper# class criteria track /inbound/RTP-I/Default RTP user-agent
```

After a period of time in which VoIP calls are made, issue the following command:

```
PacketShaper# class criteria recent inbound/RTP-I/Default
Traffic Class: /Inbound/RTP-I/Default
Application: RTP
Attribute: user-agent (SIP User-Agent)
Recent Attribute Values (most recent first)
-----
1. *X*Lite*
```

The above output indicates that **X*Lite** is the string that should be specified as the SIP user-agent criterion.

class criteria attributes

Display a list of applications that can be further classified using application-specific criteria. This information can be used to create a matching rule for a traffic class.

class criteria attributes

Application	Attribute	Example of Value

-		
Citrix	application	PeopleSoft
	client	pat-pc
	priority	3
DCOM	UUID	1cbcad78-df0b-4934-b558-87839ea501c9
DICOM	server	DICOM_STORAGE
	client	DICOM_ECHO
FTP	FileName	*.mp3
Web	host	207.78.98.18
	url	/Images/*.jpeg
	content-type	image/gif
	user-agent	Mozilla/4.0
HTTP-Tunnel	host	207.78.98.18
	port	80
ICMP	type	echo
NNTP	GroupName	alt.binaries.*
Oracle	dbname	corp
PostgreSQL	dbname	corp
ProxySG	app/port	HTTP/6000
Identity		
RTCP	encoding	GSM

	media	v
	clock	8000
RTP	encoding	GSM
	media	v
	clock	8000
	to	+12125551212@server.phone2net.com
	from	agb@bell-telephone.com
	user-agent	Motorola VT1000
	source	207.78.98.18
	destination	207.78.98.18
SMTP	SenderEmail	*@packeteer.com
SOAP	host	207.78.98.18
	url	/Images/*.jpeg
	content-type	image/gif
	user-agent	Mozilla/4.0
SSL	commonName	www.packeteer.com
WAP	WAPURI	*.packeteer.com

The application-specific criteria format in a matching rule is:

<application>:<attribute>:<value>

where *<application>* and *<attribute>* are as described in the table above, and *<value>* is specific to your configuration and classification requirements. (The values can be determined with the [class criteria track](#) command.) For example, the criterion for creating a matching rule that matches traffic to the PeopleSoft application running over Citrix is:

citrix:application:PeopleSoft

Examples of other valid matching rules are:

ftp:FileName:*.mp3

icmp:type:echo

web:host:207.78.98.18

oracle:dbname:corp



These matching rules can be specified in the [class new](#) or [class rule](#) commands.

class criteria recent

Display a list of values that have been tracked for a specific traffic class in the current configuration. This command is used after criteria tracking has been enabled.

class criteria recent *<tc class>*

where *<tc class>* is the traffic class that has criteria tracking enabled with the [class criteria track](#) command.

Example:

class criteria recent /inbound/http

PolicyCenter CLI Commands

Traffic Class: /Inbound/HTTP
Application: Web
Attribute: content-type (Content Type)

Recent Attribute Values (most recent first)

-
- 1. text/html
 - 2. image/gif
 - 3. application/x-shockwave-flash
 - 4. image/jpeg
 - 5. application/x-javascript
 - 6. application/x-www-form-urlencoded
 - 7. text/css
 - 8. application/octet-stream
 - 9. text/plain
 - 10. text/xml

class criteria track

Enable application-specific criteria tracking on a traffic class in the current configuration. (Use [class criteria attributes](#) to see a list of valid applications and attributes.)

class criteria track <iclass> <application>|off <attribute>

where <iclass> is the traffic class, <application> is the application name, and <attribute> is the type of values you want to track. For example, to set up criteria tracking of content-types for the HTTP class:

class criteria track /inbound/HTTP web content-type

After you have enabled criteria tracking for a class, use the [class criteria recent](#) command to see a list of recent values for the class.

When you are done, turn off tracking using the following command:

class criteria track <iclass> off

For example, to turn off criteria tracking on the /inbound/http class:

class criteria track /inbound/http off

class delete

Remove a class from the traffic tree in the current configuration.

class delete <iclass> [children]

<iclass>	The name of the traffic class to delete. The class' explicit hierarchical path must be supplied only if the class name itself is not unique.
[children]	Specify to delete all of the class' child classes; this parameter is required in order to delete a class that has children

If you delete a class that was created by traffic discovery and you have traffic discovery turned on, the class is likely to appear again in your traffic class tree.

class discover

Enable or disable traffic discovery within a specific class in the current configuration. For class discovery to take effect, traffic discovery must be enabled at a global level using the **setup discover** command.

```
class discover <tc> [inside|outside|both|off]
```

<tc>	The name of the traffic class within which you are enabling or disabling traffic discovery. The class' explicit hierarchical path must be supplied only if the class name itself is not unique.
[inside outside both off]	Specify the location of the server for which you want traffic to be discovered, or off to turn off discovery for this class. If you don't specify one of these options, the action defaults to turning on traffic discovery effectively using the both setting.

class group delete

Delete a custom service group in the current configuration. When you delete a group, the services in that group are moved into the Unassigned group. Note that you cannot delete the built-in groups.

```
class group delete <group_name>
```

where **<group_name>** is the name of the custom group.

Example:

```
class group delete CorpApps
```

- The child configuration inherits the parent's group of the same name.
- This group will contain all the services defined in the inherited group except for any services that were moved out of the group when the group had been overridden. These services will stay in the group that they had been moved into.
- If the overridden group contained other services that were moved into it, these services will go into Unassigned after the group is deleted.

class group move

Move all services from one service group in the current configuration into another group, or move one service into a different group. You can move services into a built-in or custom group.

```
class group move {<group_name> | <group_name>:<service_name>} <group_name>
```

Examples:

To move all the services in the Multimedia group into Mygroup:

```
class group move multimedia mygroup
```

To move service Citrix from the RemoteAccess group into the ClientServer group:

PolicyCenter CLI Commands

```
class group move RemoteAccess:citrix ClientServer
```

Additional Information

- To move a service back into its default group, use the [class group reset](#) command.
- Services in the NonIPv4 group cannot be moved to other groups.

class group new

Create a custom service group in the current configuration. PacketShaper includes 25 built-in service groups, but if these don't suit your needs, you can create your own groups. For example, if you have created user-defined services for your custom applications, you may want to create a custom group for them. You can create up to 25 custom service groups.

```
class group new <group_name> <description>
```

where *<group_name>* can be up to 31 characters (including hyphens, underscores, and periods) and *<description>* can contain up to 80 characters. If the description contains spaces, you must enclose the text string in quotes.

Example:

```
class group new CorpApps "Corporate applications"
```


class group override

For PolicyCenter / PacketShapers in Shared Configuration Mode

Override a service group that a child configuration has inherited from a parent configuration. Use this command if you don't want the service group to inherit any more changes from the parent.

```
class group override <group_name>
```

where *<group_name>* is the name of the built-in or custom service group that has been inherited.

 After overriding a service group, if you then want to re-inherit it from the parent configuration, you can [delete](#) the overridden group. Or, you can re-inherit *all* service groups from the parent with the [class group reinherit all](#) command.

class group reinherit

For PolicyCenter / PacketShapers in Shared Configuration Mode

Delete all service groups from the current configuration and re-inherit the service groups from the parent configuration. Use this command when a child configuration contains a number of local overridden service groups and you decide that you want the configuration to go back to inheriting the parent's groups.

```
class group reinherit all
```

Additional Information

- You may decide to perform this operation if, after modifying service groups in a child configuration, you end up with configuration errors (service conflicts) that you can't resolve.

- Use the [class group show](#) command to confirm that each service group has the I (Inherited) marker, indicating the group is inherited from the parent configuration. Any local custom groups you had in the child configuration will no longer appear on the group list.

class group reset

Return all services in the current configuration back to their default groups or return a single service back to its default group. Use this command if you have moved services around to different groups and then discover you made a mistake or have changed your mind.

```
class group reset <service_name> | all
```

where *<service_name>* is the name of the service that you want to return to its default built-in group.

Examples:

To move the Citrix service back into its original, default group:

```
class group reset citrix
```

Service citrix successfully reset to its default group (RemoteAccess).

To return all moved services back to their default built-in groups:

```
class group reset all
```

All services reset to their default group.



If you had created any custom groups, these groups will remain after the **reset all**, although they will no longer contain any services.

class group show

Display a list of all group names in the current configuration and their descriptions, details about a particular group, or a list of all services and the group to which each belongs. You can also use this command to find out to which group a particular service belongs.

```
class group show [<group_name> | service | {service <service_name>}]
```

Examples:

To display a list of all groups and their descriptions:

```
class group show
```

To list all groups and all the services that belong to each group:

```
class group show service
```

To display details for a particular group:

PolicyCenter CLI Commands

class group show healthcare

```
Name           : Healthcare
Description    : Healthcare related applications
Num services   : 2
Services in Healthcare group
    DICOM              Digital Imaging and Communications in Me
    HL7                Health Level Seven (HL7)
```

To find out to which group a service belongs:

```
class group show service ftp
  service:"ftp" belongs to group:"Internet"
```

class guaranteed

Display guaranteed rate statistics of a traffic class in the current configuration.

```
class guaranteed <iclass>
```

class hosts

Displays a list of all host references in matching rules and host lists in the current configuration. A host may be listed as an IP address, a DNS name, or an LDAP DN (Lightweight Directory Access Protocol domain name) for a host list entry. If more than one matching rule contains the same host reference, the host is shown only once.

```
class hosts
```

Host reference	DNS name	IP address
127.0.0.3	-	127.0.0.3
	www.lycos.com	206.79.171.51...
	www.excite.com	198.3.98.99

If there is a problem resolving a DNS name, the third column shows the DNS error message.

class id

Change or view a traffic class identification number in the current configuration. The numeric ID of a class is used for Simple Network Management Protocol (SNMP) and the measurement engine. It must be unique and does not change when the class is renamed.

```
class id <iclass> [<number>]
```

<code><tcld></code>	The name of the traffic class whose ID you are changing. The class' explicit hierarchical path must be supplied only if the class name itself is not unique.
[<code><number></code>]	The new unique number for the traffic class



Class IDs should be changed in special circumstances only, for example when you want class IDs to be the same across multiple PacketShapers. Changing class IDs can lead to erroneous reporting of data if you choose an ID value that was previously used by another class.

To see the current ID for a traffic class, type **class id** `<tcld>`. To see the ID for all services, use the [class services id](#) command.

class licenses

Limit the number of TCP flows allowed simultaneously in the given class in the current configuration, where the number of flows admitted to a class is based on a fixed number instead of the available bandwidth.

```
class licenses <tcld> off|<number>
```

where `<number>` is the maximum number of TCP flows to admit.

After `<number>` flows are active on the specified traffic class, new flows are given the admission control treatment defined by [policy admit](#).

Specify the **off** option to remove the limit on the number of flows.

After you have limited the flows with the **class licenses** command, you can use the **traffic licenses** command to see the number of flows currently in use.

class load

Load a new traffic configuration file. This command will load the traffic tree and everything related to the classes in the tree, such as policies and partitions. This feature can be used to share configurations with other units. You can upload a saved configuration to the 9.256/ folder of another PacketShaper unit and then activate it with the **class load** command.



Issuing the **class load** command will revert a unit in shared mode back to local mode.

```
class load <path>
```

The following example loads a config.ldi file from the 9.256/ folder:

```
class load 9.256/config.ldi
```


The **class load** command prompts for confirmation, then overwrites the existing 9.256/cfg/config.ldi file with the file you specify.

class move

Relocate a traffic class in the current configuration by assigning it to a new parent class. Unlike using the [class copy](#) command, the class will no longer reside under its original parent, but will be moved to a new location in the tree structure.

```
class move <tclass><new parent> [children]
```

Use the literal **children** to move all of the class' children as well; otherwise, only the parent class will be moved and the children will be promoted a level.



When moving a traffic class, you cannot change the direction. For example, you cannot move a traffic class from /Inbound to /Outbound.


class new

Create a new traffic class in the current configuration.

```
class new <parent_name><name> [nodefault] <rule>
```

<parent_name>	The parent class for the new traffic class. You must use the explicit hierarchical pathname if the class name is not unique - for example, /inbound/http.
<name>	A unique name for the new traffic class, up to 31 characters long. Use only alphanumeric characters and the following special characters: underscore (_), hyphen (-), and period (.). Specify only the class name, without the leading tree hierarchy pathname.
[nodefault]	A Default match-all class will not be created (applicable when creating a child class). For example, if you don't specify the nodefault parameter when creating the Inbound/HTTP/WebSurfing class, PacketShaper will also create an Inbound/HTTP/Default class. If you do specify the nodefault parameter, the Inbound/HTTP/Default class will <i>not</i> be created.
<rule>	A matching rule defines a traffic class' attributes. A class can contain multiple matching rules, which are treated as separate, distinct rules. To define one or more rules for a traffic class, see class rule . For matching rule details, see Matching Rule Details in PacketGuide.

You cannot create a child class if the parent has been assigned a host accounting category.



If your unit is within one traffic class of its capacity, PacketShaper will not let you create any more classes. This is due to the possibility that two classes will be created in some circumstances. For example, when you create the first child class for a parent, a Default class automatically gets created.

Creating a Class for a Specific File Type

Specify GIF file downloads:

```
class new inbound/http graphics outside service:http web:url:"*.gif"
```

Specify MP3 files downloaded via FTP:

```
class new inbound/ftp ftp_mp3_downloads outside service:ftp-data-clear
ftp:filename:*.mp3
```

Creating a Class for a Specific Host or Port

Target any traffic from an external host:

```
class new inbound competitor outside host:145.34.0.2 service:http
```

Specify web traffic to a port other than port 80, the normal web port:

```
class new inbound web_in inside service:http port:8080
```

Creating a Class for a Specific URL, IP Address, or Host List

Specify a URL (<http://altman.com/support/support.htm>):

```
class new inbound altman outside service:http host:altman.com
web:url: "/support/support.htm"
```

For security purposes, you can classify TCP traffic based on the origin of the connection. To do this, create a traffic class that specifies an outside TCP client. Create this type of class only after you are satisfied that traffic discovery has sufficiently identified traffic on your network. Otherwise, it will prevent the discovery of more specific services.

```
class new inbound mystuff outside tcp client
```

Specify an IP address if you do not have a DNS server configured:

```
class new inbound server_guru inside 203.160.106.3
```

Specify a host list (a set of IP addresses and/or DNS names):

```
class new inbound/servers inside host:any outside list:servers
```

Creating a Class for an IPv6 Subnet

```
class new inbound ipv6-2 inside net:2001:db8:1234:5678::/64
```

Creating a Class for ICMP or IGMP Traffic

When creating a symmetrical traffic class for the ICMP or IGMP protocols, we recommend that you explicitly specify the protocol for both the inside and outside interface. For example:

```
class new /Inbound/ABQ ICMP inside ICMP outside ICMP
```

To create an asymmetrical traffic class for ICMP or IGMP, where traffic is classified on either the inside or the outside interface:

```
class new /Inbound/ABQ ICMP inside ICMP
```

or

```
class new /Inbound/ABQ ICMP outside ICMP
```

class note

Annotate a traffic class in the current configuration.

```
class note <aclass>"<note>"
```

PolicyCenter CLI Commands

This note appears in the [class show](#) display. Non-printing characters are not allowed.

class override

Override an inherited traffic class in the current configuration by creating a local copy of the traffic class.

```
class override <tclass>
```

You must make a local copy of an inherited traffic class before you can change the class on the individual unit.

class owner

Specify an owner name for a traffic class in the current configuration.

```
class owner <tclass> <ownername>
```

The owner name can be up to 32 characters and the following special characters are not allowed: quote (), ampersand (&), backslash (\), and non-printing characters.

The owner name appears in the **class show** display.

class publish

This command publishes a traffic class on a child configuration to the traffic tree of its parent. The traffic class is then cleared from the child configuration, so it will inherit that class from its parent configuration. Include the **children** parameter to also publish all child classes of the selected traffic class. If the published class uses a host list, that host list is also published to its parent.

```
class publish <tclass> [children]
```



Classes based on service groups can be published only if the parent configuration has the service group in its configuration.

class rename

Rename a traffic class in the current configuration.

```
class rename <tclass><new tclass>
```

The class to be renamed must be specified with its full pathname; do not specify the path for the new class name. (The path from the original name is used.) For example:

```
class rename inbound/test sap
```

When renaming a class you are not allowed to change just the case; for example, you cannot rename HTTP to http.



If you rename a class and that class has an event associated with it, the class name is not automatically updated in the event registration. Therefore, after renaming a class, you will need to re-register the event with the new class name.

class reset

Clears all classes, policies, and partitions in the current configuration and reverts to either the default or model tree.

```
class reset [model]
```

The **model** option resets the configuration to a pre-configured tree that can be used as is, or modified to suit your needs. This tree organizes network traffic into folders of common categories, such as VoIP, risky websites, business-critical applications and data, and recreational websites and applications. It includes classes based on service groups or URL categories.

If you reset the tree without the **model** option, the tree is reset to the default—a bare-bones traffic tree that includes /Inbound and /Outbound classes with a Default class for each, and a Localhost class for the inbound and outbound directions. You can build out this tree by [turning on traffic discovery](#) or by [manually creating classes](#) for the type of traffic you want to track.

Issue this command from PolicyCenter to clear the class tree of any regular unit or sharable configuration; no draft configuration is required.

Notes:

- Use the [config save](#) CLI command to back up your configuration before resetting the tree; this gives you the capability of restoring the traffic tree and configuration if you change your mind.
- If a number of classes have top talkers/listeners enabled, a class reset could take several minutes when many hosts are being tracked. Before issuing the class reset command, you may want to disable top talkers/listeners to avoid this delay. See [hostdb topusers](#).

class rule

Add or delete matching rules for a traffic class in the current configuration.

```
class rule add <tclass> <rule>
```

```
class rule delete <tclass> <rule_id>
```

The maximum number of matching rules per traffic class depends on the PacketShaper model. If a traffic class has more than one matching rule, PacketShaper compares the flow to the first specification. If it doesn't find a match, it moves to the class' next matching rule.

Matching rules are identified by a rule ID in brackets []. You can determine the rule ID by using the command: **class show <tclass>**

See Matching Rule Details in PacketGuide for additional information.

Examples:

Create a new Oracle class with three matching rules. The first matches on an inside host IP address of 190.160.0.207, the second matches on 190.160.0.208, and the third on 190.169.0.254.

```
class new /outbound oracle inside service:oracle host:190.160.0.207
class rule add /outbound/oracle inside service:oracle host:190.160.0.208
class rule add /outbound/oracle inside service:oracle host:190.169.0.254
```

Create a new FTP class with two matching rules, one for the outside and the other for the inside.

PolicyCenter CLI Commands

```
class new inbound/ftp ftp_mp3_downloads outside service:ftp-data-clear
ftp:filename:*.mp3
  class rule add inbound/ftp/ftp_mp3_downloads inside service:ftp-data-clear
ftp:filename:*.mp3
```

Recall that if a traffic class has more than one matching rule definition, PacketShaper compares the flow to the first specification. If it doesn't find a match, it moves to the class' next rule. Traffic that matches any of a class' matching rules will fall into the class.

If the info page has flagged one or more of your classes with the configuration error message *attrib iqosMatchingRule = ??? , Failed to add matching rule to traffic class*, you have exceeded the maximum number of matching rules available on your PacketShaper model. (In the CLI, you can display configuration error messages with the **class show <tclass>** command.) To free up resources, you need to remove one or more classes or matching rules. Configuration errors will disappear once the total number of matching rules is less than the unit's limit. If you find that you are consistently exceeding your unit's maximum configuration limits, you should consider upgrading your PacketShaper.

class services

List the services available in the current configuration. These services are also listed in PacketGuide for PacketShaper (9.2, 11.x).

```
class services [<service name>][[plug-ins] [id]
```

<service name>	The name of a service; you can type the complete name, or just the first few letters
[plug-ins]	List only services that were individually added to (plugged into) the software that is, services <i>not</i> built into PacketShaper
[id]	List the internal ID numbers associated with each service name. Service ID numbers are recorded in flow detail records (FDRs). Since FDRs record the service ID, not the service name, the class services id command would be useful for someone interpreting FDR data with a protocol analyzer or other tool that displays FDR data.

The <service name> option is useful for narrowing down the service list to a particular name you are looking for. The following example lists all the services that start with AOL:

```
class services aol

AOL-IM AOL - Instant Messenger & ICQ Client-Server
AOL-IM-File AOL-IM - Point to Point File Transfer
AOL-IM-ICQ AOL - Instant Messenger & ICQ2000
AOL-IM-IMAGE AOL-IM-Image - Point to Point Chat
AOL-IM-Talk AOL-IM - Point to Point Talk
```

class set

Make a traffic class in the current configuration an exception class, or configure a class to allow its policy to be inheritable.

```
class set <tclass> inherit|standard|exception
```


inherit	Inheritable traffic classes have policies that can be applied to other classes when the other class doesn't have its own policy. Specific rules apply to how PacketShaper decides which policy a class should inherit; see Inheritance Rules in PacketGuide for details. The output of the class show command indicates (with an I flag) which classes have an inheritable policy.
standard	Standard traffic types have no exception or inheritable attributes.
exception	Exception traffic classes are always positioned above non-exception classes in the tree. When you make a class an exception class, you redefine the search order that PacketShaper uses to find a match for traffic flow. The exception attribute can be applied to all classes except /Inbound, /Outbound, and any default match-all classes. Marking a traffic class as an exception ensures that it is ordered first in the subtree, overriding the tree's built-in hierarchical order.

class show

Display traffic class information for a specific class or the entire traffic tree in the current configuration.

```
class show [<tclass> | verbose <tclass> | since <seconds>] | [id]
```

Use the **verbose** option to list all host lists referenced by a traffic class. The **since** option shows only classes auto-discovered within the last number of **<seconds>**.

When you specify a class, configuration details such as matching rule and policy information are displayed. Each matching rule is prefaced by a rule ID number. The class ID is also displayed as the last line of the output. For example:

```
class show dhcp
```

Traffic Class: /Outbound/DHCP

Partition: /Outbound

Class Flags: autocreated

Rule Types: optimized

Current guaranteed rate 0 excess rate 0

Matching Rules:

[52]	inside	any host	service:Client	any port	UDP
	outside	any host	service:DHCP-S	any port	
[54]	inside	any host	service:Client	any port	UDP
	outside	any host	service:DHCP-C	any port	
[53]	inside	any host	service:DHCP-S	any port	UDP
	outside	any host	service:Client	any port	
[55]	inside	any host	service:DHCP-C	any port	UDP
	outside	any host	service:Client	any port	

no policy

Class id (for SNMP and Measurement Engine): 1069

The **Class Flags** indicate class attributes:

PolicyCenter CLI Commands

autocreated—The class was created with the traffic discovery feature.

built-in—One of the classes built into PacketShaper (such as Inbound and Outbound). Built-in classes cannot be deleted.

cacheable—The class is cacheable (that is, a class based on an IP address that is on the same side as the cache).

discovering—Traffic discovery is turned on for this class.

exception—The class is treated as an exception, overriding PacketShaper's default ordering.

inherited—The policy for the class is inheritable.

policy—The class has a policy. (The specific policy type is shown next to Policy Flags near the bottom of the output.)

The **Rule Types** indicate the type of matching rule:

optimized—The class is optimized. An optimized class is one that was auto-discovered or one that was manually created with a simple matching rule (service type, IP address, or port number).

address-is-cacheable—The class has a pure IP address-based matching rule that is on the same side as the cache (on the inside, by default). It can be an individual IP address, a range of IP addresses, an address with a mask, or host lists. These classes can be cached unless an error in the tree configuration is causing cacheability problems.

match-all—This class is a match-all class (protocol = any, service = any; for example, a Default bucket).

If you don't specify a class, all classes in the traffic tree are displayed, but with less detail. When displaying the entire traffic tree with the **class show** command (as shown in the following example), several flags indicate class attributes, type of matching rule..

```
class show
```

Derivation: (I)nherited (O)verride (U)nderride (L)ocal

Class Flags: (A)utocreated (D)iscovering (E)xception (I)nherit (P)olicy
(C)acheable

Rule Types: (o)ptimized (m)atch-all (a)ddress is cacheable

Class Name	Flags	Partition Name
------------	-------	----------------

```

Inbound                               m /Inbound
Localhost                           E P  /Inbound
10.7.38.0                           a /Inbound
CUSTOMER                           P ma /Inbound
mysite.org                          C a /Inbound
Default                             IP m /Inbound
Outbound                            m /Outbound
Localhost                           E P  /Outbound
10.7.38.0                           a /Outbound
CUSTOMER                           ma /Outbound
mysite.org                          C a /Outbound
Default                             IP m /Outbound

```

class test

Test a traffic flow against the present classification tree in the current configuration in order to determine the flow's class, partition, and policy.

```
class test <direction> <protocol> [<inhost:inport> <outhost:outport>] [<device>]
```

<direction>	inbound or outbound
<protocol>	tcp, udp, icmp, netbeui, ipx, appletalk, decnet, fna, sna, lat, or misc
<inhost:inport> <outhost:outport>	<p>The inside and outside IPv4 or IPv6 addresses and port numbers to test (required for IP protocols only: TCP, UDP, ICMP)</p> <p>You must supply both an inside and an outside address. Use 0.0.0.0:0 as a placeholder if you don't have an address to test on one of the sides.</p> <p>If the hosts are IPv6 addresses, surround the IPv6 address with square brackets. For example, [2000:1:2::1]:3456 where 3456 is the port number.</p>
<device>	A valid device name, such as slot1 or slot3_pair1

This information simulates a flow, returning the following information:

Traffic Class	The traffic class in the current traffic tree into which the flow would be classified
Partition	The partition associated with the matching traffic class. If the traffic doesn't have its own partition, the parent partition is used.
Policy	The matching policy. If the matching traffic class has no applied policy, the policy is inherited. See Inheritance Rules in PacketGuide .



The **class test** command will only match traffic classes that have "any" for the server location.

Examples:

PolicyCenter CLI Commands

class test inbound appletalk

```
Traffic class --> /Inbound/AppleTalk
Partition    --> /Inbound
Policy       --> /Inbound/Default
```

class test inbound tcp 216.110.182.168:80 0.0.0.0:0

```
Traffic class --> /Inbound/HTTP
Partition    --> /Inbound
Policy       --> /Inbound/Default
```

Additional Information

- The **class test** command can be used to test basic classification for IP protocols, but is not intended to test every type of classification PacketShaper offers. Its purpose is to check a particular IP address or port number to determine how the traffic is classified into existing port-based and IP address-based classes in the traffic tree. The command does not include fields for specifying more complex types of classification such as MAC address.
- The **class test** command requires touch access.

class undelete

Issue this command to restore a class marked for deletion from a draft configuration. If the class has any child classes, they will also be restored.

```
class undelete <tclass>
```

class user-group

List the names of user groups in a specific Active Directory domain or in all domains.

```
class user-group <domain_name>|all
```

This command is part of the user awareness feature and requires that a BCAAA server be installed and [configured](#).

If you have a long list of user groups, some may scroll off the screen; if you want to be able to scroll through the list, use one of the following techniques:

- Output the list to a text file, for example: **class user-group all > grouplist**. When the command prompt redisplay, the file has finished saving. (This might take a while.) To display the list a page at a time: **more grouplist**
- Turn on session logging in your remote login utility (such as Putty or SecureCRT) before issuing the **class user-group all** command. You can then open the log file in a text editor.

class users

List the names of users in a specific Active Directory domain or all domains.

```
class users <domain_name>|all
```

This command is part of the user awareness feature and requires that a BCAA server be installed and configured. For details, see PacketGuide for PacketShaper ([9.2](#), [11.x](#)).

A long list of users will scroll off the screen; if you want to be able to scroll through the list, use one of the following techniques:

- Output the list to a text file, for example: **class users all > userlist**. When the command prompt redisplay, the file has finished saving. (This might take a while.) To display the list a page at a time: **more userlist**
- Turn on session logging in your remote login utility (such as Putty or SecureCRT) before issuing the **class users all** command. You can then open the log file in a text editor.

class user-services delete

Remove a user-defined service from the current configuration.

```
class user-services delete <serviceName>|all
```

where *<serviceName>* is the name of the service you want to delete. Use the **all** parameter to delete all user-defined services.

Example:

```
class user-services delete TDEmployees
```

Additional Information

- Service names are case sensitive. You must enter the service name with the same upper/lower case with which it was created.
- Use the [class user-services show](#) command to see a list of services that have been user-defined.

class user-services new

Create a custom service in the current configuration in order to identify and categorize traffic that is not currently classified by PacketShaper, or that is classified into a different service. This command allows you to create services for in-house applications on your network. The service can be defined by a signature (hex or string) and/or by port numbers.

```
class user-services new <serviceName> [signature:<hex>|<string> offset:<offset_value>]
[port:<nnn[-nnn]>] [packets:<packet_value>] [ipproto:TCP|UDP] [description:<string>]
```

<i><serviceName></i>	The name of the service, up to 30 characters long. Use only alphanumeric characters and the following special characters: underscore (_), hyphen (-), and period (.). The service name is case sensitive.
signature	The signature can be specified in hexadecimal format or as a quoted string. The string can be up to 30 characters long, is case sensitive, and must be enclosed in quotation marks. The hex representation can be up to 30 characters long. It must begin with 0x.
offset	Starting position of the signature in the payload (after the header). Valid values for the offset are 0-1499.

PolicyCenter CLI Commands

packets	Number of inbound or outbound data packets in each new flow that will be inspected for the signature. Up to 10 packets in each direction can be inspected. Note: Packets in each direction are counted separately. For example, a value of 8 tells the PacketShaper to look for the signature in the first eight inbound packets and first eight outbound packets of each new flow.
port	The port number or a range of port numbers. If the port option is not specified, the PacketShaper will inspect traffic on all ports.
ipproto	Type of IP protocol (UDP or TCP)
description	A description of the user-defined service, enclosed in quotation marks; up to 80 characters long.

The following types of traffic are candidates to be classified as a user-defined service:

- 1) traffic that PacketShaper has identified as an unknown service,
- 2) applications that have user-configurable ports (such as peer-to-peer and instant messaging)

Services that are associated with well-known ports (such as HTTP on port 80, FTP on port 21, and NNTP on port 119) cannot be classified into a user-defined service.

Examples:

```
class user-services new TDemployees signature:"TD Employee" offset:6 packets:1
description:"TD Employee Database"
```

```
class user-services new BCpayroll signature:0x424320706179726F6C6C offset:0
description:"BC Payroll application"
```

Additional Information

- You can use a third-party network protocol analyzer, such as EtherPeek or Wireshark, to analyze a trace to get the signature.
- You can create up to 10 user-defined services (UDS).
- The user-defined services are auto-discoverable.
- The user-defined services are stored in the config.idi configuration file.

class user-services show

Display a list of user-defined services in the current configuration.

```
class user-services show [<serviceName>]
```

Example:

```
class user-services show
```

User Defined Services

1. Name:BCpayroll
 serviceid:645 signature:0x424320706179726f6c6c offset:0 packets:2
 ipproto:TCP/UDP description:"BC Payroll application"
2. Name:TDemployees

```
serviceid:647 signature:"TD Employee" offset:6 packets:1
ipproto:TCP/UDP description:"TD Employee Database"
```

class web-app disable

Disables the service of a web-based application (such as Facebook or YouTube) in the current configuration. After a service is disabled, the traffic will get classified as HTTP or SSL. You might want to disable a service when you prefer to control the traffic by its URL category. For example, you can disable the Facebook service and then control all social networking traffic with a single Social Networking URL category class. This technique helps conserve classes and provides an easy way to report on and control how much HTTP is on the network.

```
class web-app disable <service>
```

where <service> is the name of the web application to disable. To see a list of web services that can be disabled, use the **help class web-app** command. For example, the help output includes the following:

The following values are currently supported for this argument:

```
Facebook
Youtube
MySpace
Orkut
Flickr
Meebo
GoogleVideo
Ogg
Smugmug
Ofoto
Motion
WebShots
```

PolicyCenter Support

Web services cannot be disabled in PolicyCenter; this feature is supported in local mode only. If PolicyCenter pushes a class to a PacketShaper that is in shared mode, and that class uses a service that has been disabled locally on the PacketShaper, the class will still be created. However, traffic will not get classified into the class as long as the service is disabled.

Additional Information

- If the PolicyCenter management console is open when you enable/disable the service, you will need to refresh the browser window to load the configuration change.
- When a web application is disabled, you cannot create a class for that service in the CLI, Advanced UI, or Sky UI. The application will not appear on the **Services** drop-down list in the Sky or Advanced UIs after it has been disabled.
- Disabled web applications will not get classified into the service group to which the service belonged.
- If a class already exists for a disabled web application, the class will still appear in the traffic tree but will no longer get any class hits. It will *not* have a configuration error.
- It is not necessary to remove plug-ins for web services that have been disabled.

class web-app enable

Re-enables a web application (such as Facebook or YouTube) in the current configuration after it has been disabled. You would use this command if you change your mind about disabling a service and want to start classifying the service separately again.

```
class web-app enable <service>
```

where <service> is the name of the web application to enable. To see a list of web services that are currently disabled, use the [class web-app show](#) command.

Additional Information

- If the PolicyCenter management console is open when you enable/disable the service, you will need to refresh the browser window to load the configuration change.
- To verify that the service has been re-enabled, use the [class web-app show](#) command. The service should *not* be listed in the output of the **show** command.
- After you have re-enabled a web service, you can manually create classes based on this service or let the PacketShaper auto-discover the class. If the class already existed in the tree, it will start getting class hits once the service is re-enabled.

class web-app show

Lists the services in the current configuration that have been disabled with the [class web-app disable](#) command.

```
class web-app show
```

Example:

```
class web-app show
```

```
Classification of the following Web applications is disabled:
```

```
Facebook  
YouTube
```

config backup

Make a backup copy of a PolicyCenter configuration. After you issue the **config backup** command, you will be prompted to confirm that you want to create a backup of the specified configuration. Enter the word **Yes**, or press the **Enter** key. Backup configurations will appear in the PolicyCenter configuration tree with a "-backup" after the configuration name.

```
config backup [<cfg_path>]
```

Restore a backup copy of a PolicyCenter configuration with the [config restore](#) command.

config clear

For PolicyCenter / PacketShapers in Shared Configuration Mode

Clears all non-default configuration values from the named configuration. If none is named, it clears the current configuration. Clearing a child configuration means that the child will derive its sharable attributes and settings from its parent configuration. If you clear a parent configuration, its child configurations will no longer inherit any values from its parent.

```
config clear [<cfg_path>]
```

config cp

Copies an existing configuration to a new or existing configuration. Include the `-r` (recursive) option to include the selected configuration's child configurations in the copy operation. Note that if the configuration to be copied and the destination configuration both have a child configuration with the same name, the destination configuration's child will be overwritten. If the `<source cfg_path>` argument is omitted, it copies the current active configuration.

This command does not allow a parent configuration to be copied to its child configuration with the `"-r"` option. You also may not copy to a draft configuration, or to any configuration that has a draft anywhere in its configuration hierarchy. The individual serial-number configuration of a PacketShaper is unique to that unit, and cannot be copied to another location in the configuration tree unless you also rename the new copy of the unit configuration as a part of the copy operation.

```
config cp [-r] [<source cfg_path>] <dest cfg_path>
```

Where the `<source cfg_path>` is the source configuration to be copied, and the `<dest cfg_path>` is the destination for the new copy of that configuration. Specify a slash (`/`) for the `<dest cfg_path>` value to copy the source configuration to the root of the configuration tree.

See also [config mv](#) for details on moving PolicyCenter configurations

config dump

For PolicyCenter/ PacketShapers in Shared Configuration Mode

This command prints out the current effective configuration objects' formats and attributes in something like LDAP data interchange format. Useful mainly for development and diagnostic purposes.

```
config dump
```

config edit

Locks the current configuration, creates a draft copy of that configuration if a draft does not exist, and opens the draft configuration for display and modification. If a draft copy of that configuration already exists, this command only opens the draft configuration for display, but does not create a new draft.

```
config edit <cfg_path>
```

Draft configurations impose limitations not present in other configurations. Once you have created a draft copy of a configuration, neither the original configuration or any of its parent or child configurations can be modified until the draft configuration is permanently committed or deleted.

If, for example, you had a PolicyCenter configuration tree with the following configurations

PolicyCenter CLI Commands

- `/parent_cfg`
- `/parent_cfg/child1`
- `/parent_cfg/child1/grandchild1`
- `/parent_cfg/child2`
- `/parent_cfg/child2/grandchild2`

the command `config edit parent_cfg/child1` would lock the configurations `/parent_cfg/parent_cfg/child1` and `/parent_cfg/child1/grandchild`, and would create a new draft configuration called `parent_cfg/child1-draft`. The configuration tree would then appear as follows:

- `/parent_cfg` (locked)
- `/parent_cfg/child1` (locked)
- `/parent_cfg/child1-draft` (locked)
- `/parent_cfg/child1/grandchild1` (locked)
- `/parent_cfg/child2`
- `/parent_cfg/child2/grandchild2`

A draft configuration can only be edited by one PolicyCenter user at a time--no other user can modify a draft until the first user logs out of PolicyCenter or sets the focus of his PolicyCenter session on another configuration (for example, by using the [config view](#) or [config edit](#) commands and specifying another configuration). However, while one user is modifying a draft, other users are allowed to view (but not change) the draft.

Once you have made the required modifications to a draft configuration, you can test that configuration on one or more PacketShapers with the command [draft try](#), or permanently commit the changes using the command [draft commit](#).

config errors

Display configuration errors for the unit. When issued from PolicyCenter, this command displays errors for the PolicyCenterconfiguration currently being edited.

```
config errors
```



Configuration errors are also shown in the output of the [banner show](#) command.

config load

Load saved configuration files (such as `config.ldi` and `config.cmd`). Sharable settings are saved in files with the `.ldi` file extension, while nonsharable settings are saved in the `.cmd` file.

This command can load the traffic tree, partitions, policies, host lists, events, agents, basic settings (such as shaping and traffic discovery), security settings (such as passwords and login access protocols), SNMP, SNTP, and Syslog settings, DNS server, and gateway addresses, domain names, time zones, and network interface settings.



Use the [setup show](#) command to see a list of sharable and nonsharable settings that are stored in the configuration files.

```
config load <file>[<cfg_path>] [complete]
```

<file>	<p>The location and name of a saved configuration file. Include the .ldi file extension to load just an .ldi file, or omit the file extension and include the complete parameter to load both a .ldi and a .cmd file with the specified filename.</p> <p>By default, this command loads files from the PolicyCenter directory. To load saved files from a different folder, specify the entire path.</p> <p>For example, to load the configuration files test.ldi and test.cmd from the 9.258/ folder, type:</p> <pre>config load 9.258/test complete</pre>
<cfg-path>	Include the path of the PolicyCenter configuration to which you want to load the files.
[complete]	Include the complete parameter to load the saved .ldi and .cmd files. If this parameter is omitted, the command will load only the sharable settings in the .ldi file.

The **config load** command discards the current configuration and institutes the loaded configuration; it does not merge the loaded configuration with the pre-existing one. The new configuration settings are then stored in 9.256/CFG/config.ldi.

Keep in mind that the .ldi file includes the unit's password, and if you load the configuration on another unit, you will change its password. If you want to load a traffic configuration on another unit without changing the password, use the [class load](#) command instead of the **config load** command.

config mode

For PolicyCenter/ PacketShapers in Shared Configuration Mode

Indicates whether a unit is in local or shared mode.

```
config mode
```



This command does not enable or disable the LDAP client, which is normally initialized with `config setup` and disabled with `config unset`.

config mv

Moves a configuration to another location within the PolicyCenter configuration tree. This command copies the specified source configuration to the destination configuration name, switches any assigned units from their source sharable configuration to the new destination configuration, and deletes the source configuration. Note that you cannot move the

PolicyCenter CLI Commands

/default configuration or the individual unit configurations of PacketShapers that have not been assigned to a sharable configuration.

If the configuration is a parent configuration with child configurations, the selected configuration's child configurations will be included in the move operation.



You may not move a configuration under a draft configuration, or to any configuration that has a draft anywhere in its configuration hierarchy.

If the `config mv` command on the previous page the source configuration name is omitted, this command will assume the current active configuration is the configuration to be moved. You must, however, specify the destination configuration path.

```
config mv [<source_cfg_path>] <dest_cfg_path>
```

Where the `<source_cfg_path>` is the source configuration to be moved, and the `<dest_cfg_path>` is the destination for that configuration. If the first `<cfg_path>` value is omitted, PolicyCenter will move the current active configuration. Specify a slash (/) for the `<dest_cfg_path>` value to move the source configuration to the root of the configuration tree.

config new

Creates a new, empty configuration with the given name. You can use this command to create a new configuration at the top of the configuration tree, or to add a new child configuration under an existing parent.

```
config new <cfg_path> legacy|s-series
```

Examples:

```
config new 7500 legacy
```

```
config new /PS-S-configs/S200 s-series
```

config owner set

Assign a configuration to a specified organization. Include the `-r` (recursive) option to assign the selected configuration and all its child configurations to the same organization.

A child configuration can only be assigned to a different organization than its parent if the parent configuration is assigned to *PC*, the default PolicyCenter organization. If the parent configuration is assigned to any other organization, all of its child configurations must be assigned to that same organization.

For example, if the parent configuration `/parent` is assigned to the *PC* organization, its child `/parent/child` can be assigned to *PC* or any other existing organization. However, if the parent config `/parent` is assigned to any other organization besides *PC*, such as *New_York*, then the child configuration `/parent/child` must also be assigned to that *New_York* organization.

You must be logged in as a PolicyCenter administrator to issue this command. You may not change the organization on a configuration that has a draft anywhere in its configuration hierarchy.

```
config owner set [-r] </cfg-path> <organization>
```

Examples:

```
config owner set -r /TriStateConfig New_York
```

```
config owner set /PacificNorth/Corvallis Oregon
```

config owner show

Lists PolicyCenter configurations and the organization to which those configurations are assigned. Include the `</cfg-path>` parameter to view the organization for that single configuration, or omit the parameter to view the assigned organization for all PolicyCenter configurations. You must have touch access to *PC*, the default PolicyCenter organization, in order to issue this command.

```
config owner show [</cfg-path>]
```

Example:

```
config owner show
```

Configuration	Owner Organization
901-20000132	PC
default	PC
branch_west	California
los_angeles	California
portland	California
san_francisco	California
branch_east	PC
new_york	East_Sales
raleigh	East_Sales
washington_dc	East_Sales
branch_central	Manufacturing

config publish

This command publishes a child configuration to its parent, replacing classes and settings in the parent configuration with classes and settings in the child configuration. The child configuration is then cleared, so it will inherit its entire configuration from the new settings of parent.

Use this command to publish discovered traffic classes to a parent configuration, or to publish a prototype configuration that should be inherited by all child configurations under the same parent. If the `<cfg_path>` argument is omitted, this command publishes the current active configuration.

```
config publish [<cfg_path>]
```



PolicyCenter cannot publish traffic classes from or to a draft configuration. This command will not work if either the parent or child configuration is a draft configuration.

config reset

For PolicyCenter/ PacketShapers in Shared Configuration Mode

PolicyCenter CLI Commands

When you issue this command from the PolicyCenter command-line interface, communication between PolicyCenter and the directory server will be disabled. With this connection disabled, PolicyCenter will no longer be able to contact PacketShapers in shared mode. To restore the connection between PolicyCenter and the directory server, use [config setup](#).

When issued from the command-line interface of an individual PacketShaper, this command disables the unit's connection to the PolicyCenter directory server, returning the unit to local mode and setting the unit's sharable attributes to their factory-default state. The **config reset** command will not remove a unit entry from the PolicyCenter directory server, and the unit's non-sharable settings (IP address, DNS and management port settings, etc.) will not be changed. To completely remove the unit entry from PolicyCenter, use [unit clean](#).

```
config reset
```



If you want to return a unit to local mode *without* clearing the unit's sharable attributes, use [config unset](#), instead. You may restore a unit's previous PolicyCenter configuration at any time by resetting its connection to the directory server with the [config setup](#) command.

config restore

Restore a [backup copy of a PolicyCenter configuration](#). (Backup configurations appear in the PolicyCenter configuration tree with a "-backup" after the configuration name.) The **config restore** command does not delete a backup configuration after it copies it to its original configuration, so you can restore a single backup configuration as often as desired.

```
config restore [<cfg_path>]
```



When you issue the **config restore** command, specify the original configuration you want restored, and not the backup configuration.

For example:

```
config restore Florida/Miami
```

config rm

Removes a configuration or group of configurations from PolicyCenter. If the configuration name is omitted, this command will assume the current active unit configuration is the configuration to be deleted.

```
config rm [-r] [<cfg_path>]
```

This command cannot delete a configuration if it or any of its child configurations have units assigned to them. Before you delete a configuration that has a unit assigned to it, be sure to reassign the units to another configuration. Include the **-r** (recursive) argument to delete both the selected configuration and all its child configurations. Omit the **-r** argument to delete a configuration with no children.



The **default** configuration can't be removed.

config save

Save the current configuration's sharable settings in an .ldi file and its nonsharable settings in a .cmd file.

```
config save [<cfg_path>] <file> [unit]
```

<cfg-path>	<p>The name of the configuration you want to save, including its complete path. For example, to save a configuration named <i>legacy</i> that has a parent configuration named <i>master</i>, type <i>/master/legacy</i> for the <cfg_path>.</p> <p>If you don't specify a <cfg_path>, PolicyCenter saves the current configuration.</p>
<file>	<p>Specify a filename. The .ldi and .cmd extensions are automatically added to the configuration file name.</p> <p>By default, this command saves the files to the PolicyCenter directory. To save the files to a different folder, specify the entire path.</p> <p>For example, to save a configuration in files named test.ldi and test.cmd on the 9.258/ folder, type:</p> <pre>config save 9.258/test</pre>
[unit]	<p>Saves a unit's local sharable and nonsharable settings. If the unit parameter is omitted, the config save command will save a configuration's inherited and local settings.</p>

This command can save the traffic tree, partitions, policies, host lists, basic settings (such as shaping and traffic discovery), security settings (such as passwords and login access protocols), SNMP, SNTP, and Syslog settings, DNS server, and gateway addresses, domain names, time zones, and network interface settings. Use the [setup show](#) command to see a list of sharable and nonsharable settings that are stored in the configuration files.

The **config save** and **config load** commands are useful for experimenting with different configuration settings. For example, you can save your current settings, make changes to the configuration (such as create new partitions or policies), and then return to the original configuration if you prefer it. You can create as many configurations as you like.

This feature can also be used to share configurations with other units. You can download the configuration files to a local workstation, upload them to another PacketShaper unit, and then activate the configuration with the [config load](#) command.



Keep in mind that the .ldi file includes the unit's password, and if you load the configuration on another unit, you will change its password. If you want to load a configuration on another unit without changing the password, use the [class load](#) command instead of the **config load** command.

config secure

Issue this command to enable or disable Secure LDAP communication between PacketShapers assigned to this configuration and PolicyCenter.

```
config secure [<cfg_path>] on|off
```

config setup

Configures the unit to access shared configurations in Lightweight Directory Access Protocol (LDAP). Initializes the LDAP client to communicate with the directory server and establish the default unit configuration name. A unit's initial PolicyCenter configuration is based on its DNS name (if known) or IP address. When this command is complete, the unit

PolicyCenter CLI Commands

will obtain its configuration from the directory server, replacing any previous local setup, policy, or other sharable configuration values. If you add the optional `convert` option, the configuration of the unit is preserved.

```
config setup <ldap_host>[:<port>] [secure | unsecure] [<directory_server_password>]
[convert]
```

Where:

<code><ldap-host></code>	DNS name or IP address of a PolicyCenter directory server
<code><port></code>	TCP port number to connect to on the directory server
<code>secure/</code> <code>nonsecure</code>	Specify secure to establish a secure LDAP connection between the PacketShaper and the PolicyCenter directory server, or specify nonsecure for a standard LDAP connection.
<code><directory_</code> <code>server_</code> <code>password></code>	Password for the PolicyCenter directory server.
<code>[convert]</code>	Specify the convert option to convert the unit's existing configuration into a new PolicyCenter configuration with the same attributes and values. Because the unit's new PolicyCenter configuration will be based upon its previous configuration, the unit will continue to operate the same in PolicyCenter as it did in local mode. If you do <i>not</i> select the convert option, the unit's new PolicyCenter configuration is cleared, and will have default settings only.

If you previously issued the command [config unset](#) to disable communication between PolicyCenter and the directory server, you can issue the command `config setup <ldap_host>[:<port>] [secure | unsecure] [<directory_server_password>]` from the *PolicyCenter* configuration (the configuration for the PolicyCenter software) to restore communications between PolicyCenter and the directory server. Note that this use of the **config setup** command doesn't support the *convert* option.

config show

Lists available PolicyCenter configurations. Depending on the subcommand, shows the available configurations and unit status information. Useful for monitoring units, verifying the PolicyCenter configuration hierarchy, or determining software image versions.

```
config show all|units|versions|{details <unit name>|<unit serial number>}
```

<code>all</code>	Displays a table of all units subscribing to the directory server, the configuration they are assigned to, IP address, and status. If a unit has not recently updated its status entry, the time since last update is noted as its 'Out of Contact' time. The status column reports whether a unit has found any errors in its configuration.
<code>units</code>	Displays a table of all units that are posting status to the directory server, with serial number, group/unit name, model, and domain name.

versions	Displays a table of all units that are posting status to the directory server, with serial number, IP address, and image version.
details <unit name>/<unit serial number>	Shows all status information reported by the unit to its status entry in the directory server. You can designate the unit by its unit configuration name (e.g. '/default/austin') or its serial number (e.g. '100-10000105').

The example output below shows a configuration tree with fourteen configurations, including the configuration for the PolicyCenter appliance itself, configuration *901-20000132*. The other configurations at the top of the configuration tree are *default*, *branch_west*, *branch_east* and *branch_central*.

The *branch_west*, *branch_east* and *branch_central* configurations each have three child configurations with an assigned unit. The names of each of these child configurations are indented in the **Configuration Name** column, to show that they are child configurations under another parent. Information on the individual PacketShapers, such as unit name, IP address, Out of Contact time, and the status of the unit is displayed beside the unit's assigned configuration.

```
/025-10001808# config show
```

Configuration Name	Unit Name	IP Address	Out Of Contact	Status
901-20000132	901-20000132	172.21.7.50		OK
default				
branch_west	main_site	172.21.29.129		OK
los_angeles	shaper_1	172.21.29.130		OK
portland	shaper_2	172.21.29.135		OK
san_francisco	shaper_3	172.21.29.139		OK
branch_east				
new_york	shaper_4	172.21.18.75		OK
raleigh	shaper_5	172.21.18.45		OK
washington_dc	shaper_6	172.21.18.99		OK
branch_central				
denver	shaper_7	172.21.25.160		OK
madison	shaper_8	172.21.25.170		OK
oklahoma_city	shaper_9	172.21.27.203		OK

config unset

For PolicyCenter/ PacketShapers in Shared Configuration Mode

This command disables directory server access for a unit, and returns the unit to local mode. The **config unset** command removes a unit entry from the PolicyCenter directory server, so the PacketShaper no longer appears on the PolicyCenter Configurations tab, but allows the unit to retain its last PolicyCenter configuration after it returns to local mode. To set the unit to local mode and return its configuration to a factory-default state, use [config reset](#).

PolicyCenter CLI Commands

```
config unset
```

When you issue this command from the PolicyCenter command-line interface, PolicyCenter will disable communication between PolicyCenter and the directory server. With this connection disabled, PolicyCenter will no longer be able to contact PacketShapers in shared mode. To restore the connection between PolicyCenter and the directory server, use [config setup](#).

config view

Closes the current configuration and opens the named configuration for display.

```
config view <cfg_path>
```

This command allows you to view, but not modify, configuration settings. To open a draft copy of a configuration for editing, use the command [config edit](#).

After you enter the command, the CLI prompt will display the name of the current configuration.

Example:

```
/default # config view /master1/legacy  
/master1/legacy #
```

date

View or set the date and/or time. When initially setting the date and time, use **setup timezone**.

```
date [<yyyymmddhhmm>[<.ss>]]
```

Note that this command has the same functionality as the **setup date** command.



You should always do a [system reset](#) immediately after changing the date so that the underlying time-sensitive scheduled operations of PolicyCenter can be correctly initialized.

dns lookup

List the IP address(es) associated with a domain name. PolicyCenter keeps the mapping data up to date so that when a site changes an IP address, the matching rule knows about the change.

```
dns lookup <hostname>
```

If the name that you enter is different from the canonical or official name, the canonical name record (CNAME) is displayed at the end of the address list. A canonical name record defines an alias for the official host name, facilitating the transition from an old name to a new name.

Some sites return multiple addresses to a lookup query. The PacketShaper classification process compares the traffic flows to the address lists when looking for a match.

dns names

List all domain names and addresses that are configured in PolicyCenter.

dns names

Domain Name	IP Address	TTL	Age	RQSNCRP	Error
(luna-corp.bluecoat.com)...	192.168.0.33	3600	647	Q	
(m10-pat-corp.bluecoat.com)...	192.168.0.207	3600	427	Q	
percy.xyz.com	(204.202.49.73)	86400	12512	Q	

The resolved values are shown in parentheses. The other columns in the output are described below.

TTL: The time interval that the DNS entry may be cached before the source of the information should again be consulted.

Age: The time, in seconds, since PolicyCenter received the last name refresh.

R: If a name server cannot be reached, the entry's retry count is incremented. This is a high-level retry, and each one may include multiple queries to each name server. If the retry value is greater than 9, an asterisk is displayed in this column. If the retry value is zero, nothing is displayed in the column.

Q: Displays a Q if PolicyCenter sent a query and received a response for the name.

S: Displays an S if PolicyCenter learned the name's address (or vice versa) by spying on DNS traffic instead of making a query.

N: The number of successful responses received since the one containing this address. If the value is 0, nothing is displayed in the column.

C: The number of responses received before getting one without any new addresses. This is the length of a round-robin cycle. If the value is 1, nothing is displayed in the column.

R: The number of matching rules that refer to this name. It will be incremented by one while a name is being resolved. If the value is 1, nothing is displayed in this column.

P: Displays a P if PolicyCenter is currently resolving this name.

Error: Shows the problem (if any) encountered by the last refresh attempt. Some possible errors are:

name not found: The authoritative server for this domain has no such name.

server offline: The resolver could not reach the authoritative name server, either directly or indirectly through the locally-configured name servers.

rqst refused: The name server knows (or might know) but won't tell you.

no data record: The name exists, but does not have an address (or vice versa).

internal error: The name server is not functioning.

dns refresh

Clear the resolved DNS values—that is, names and IP addresses—in the names database. The entries then are repopulated at the next ten-second polling interval.

```
dns refresh
```

Immediately after executing **dns refresh**, if you use the [dns names](#) command, the resolved values will be listed as <unknown> in the output. These entries are repopulated at the next polling interval.

dns rlookup

Find the host name associated with an IP address.

```
dns rlookup <ipaddress>
```

dns servers

List the DNS servers, their online/offline status, and the time since the servers either timed out or responded to a DNS request.

```
dns servers
```

Address	Status	Idle
192.168.0.33	on line	4
192.168.0.22	unknown	

draft commit

Merge changes made to a draft copy of a configuration into the original target configuration. After merging the changes, PolicyCenter reassigns any PacketShapers using the draft configuration back to their original target configuration, then deletes the draft. Once a draft has been committed, PolicyCenter removes the configuration locks on the draft's parent and sibling configurations, so other PolicyCenter users may edit them.

```
draft commit <config-draft>
```

Example:

```
draft commit legacy-draft
```

draft discard

Discard a draft copy of configuration without merging any of the changes into the original target configuration. If any PacketShapers were assigned to this draft configuration with the [draft try](#) command, you will not be able to discard the draft configuration until the units are assigned back to their original target configuration with the [draft revert](#) command. This command also removes the configuration locks on the draft's parent and sibling configurations, so other PolicyCenter users may edit them.

```
draft discard <config-draft>
```

Example:

```
draft discard legacy-draft
```

draft revert

Reassign any PacketShapers using a draft configuration back to their original target configuration. The changes made to the draft configuration are retained, and the draft's parent and sibling configurations remain locked.

```
draft revert <config-draft>
```

Example:

```
draft revert legacy-draft
```

draft try

Applies a modified draft configuration to one or more selected PacketShapers, allowing you to test the draft configuration before you apply it to a larger group of units. You re-issue this command to assign additional PacketShapers to a draft, though the draft may not be modified while any PacketShaper is trying it.

```
draft try [<cfg_path>] [all | <unit_name>|<unit_sn> <unit_name>|<unit_sn> ....>]
```

If you don't like the result, you can revert the PacketShapers running the draft configuration back to their original target configuration with the command [draft revert](#). If the test goes well and you would like to make the draft changes permanent, you can commit the draft to the original configuration with the command [draft commit](#). Once a draft configuration has been committed, all PacketShaper running or inheriting from the target configuration will get the draft changes.

draft view

Change the focus of your PolicyCenter session to the selected draft configuration, but only with read access. (You will only be allowed to view, but not modify, the draft configuration.) You can also use this command to release your session's lock on a draft configuration you are finished editing, so another PolicyCenter user can access and edit the draft.

```
draft view [<cfg_path>]
```



To edit and modify a draft configuration use the **config edit** command.

event delete

Delete an event and all its registrations in the current configuration.

```
event delete <name>
```

event email

Add or delete an email recipient for event notifications in the current configuration.

PolicyCenter CLI Commands

```
event email add [<recipient> ... <recipient>]
```

```
event email delete [<recipient> ... <recipient>]|all
```

Separate recipients with a space. You can add up to four recipient addresses.

To use the command-prompt mode, use:

```
event email add
```

event log reset

Delete current and archived event log files in the current configuration.

```
event log reset
```

event log status

Display information about current and archived event log files, such as their location, current capacities, and limitations.

```
event log status
```

event new

Define a new event in the current configuration. When you define an event, you specify a measurement variable in an expression— that is, the condition for which you want to be notified. In addition, you can define a default event-checking interval. The maximum number of events that can be defined is 32. Defined events are not active until registered.

To initiate the command-prompting mode use:

```
event new
```

You may exit 'event new' at any time by typing 'exit'

Name of the event: **WebQoS**

Type of object to be tested: Link, Partition, or traffic Class: (class):

Measurement Engine variable to be tested: **tcp-conn-aborts%**

Default checking interval [1m,1h] (1m):

Enter a relational operator. When you register this event later, you will supply a threshold on 'tcp-conn-aborts%' that triggers the event.

The event can be triggered when 'tcp-conn-aborts%'

becomes >, >=, <, or <= the threshold.

Relational operator (>, >=, <, or <=) (>):

As an alternative to the prompting mode, you can use a single command line to create an event, as follows:

```
event new <name> <expression> [<default checking interval>]
```

<i><name></i>	Event names must begin with an alphabetic character and contain only alphabetic characters, numbers, and underscores, up to a maximum of 32 characters. Note that you cannot use hyphens in event names.
<i><expression></i>	<p>The expression specifies the condition that will be checked and requires adherence to the following syntax:</p> <p><i><variable>[. <object type>] <relational operator> <constant></i></p> <p>Where:</p> <p><i><variable>.<object type></i> is one of the PacketShaper measurement variables with an appended object type that is, a link, partition, or class. This object type is required for most variables those that are common to link, partition, and class objects. Some variables are unique to the object type. For example, peak-excess-bps is relevant only to partitions, so it does not need the object-type qualifier in this syntax. Later, when you register the event, you will supply a specific name for the object type. For a list of measurement variables, use the measure show command.</p> <p><i><relational operator></i> is one of the following: <, <=, =, >=, or >.</p> <p><i><constant></i> is a placeholder for the threshold value. Use the \$n syntax for example, \$1 or \$2. When you register an event, you supply a value that is substituted for this constant in the expression.</p> <p>Example of an expression: tcp-rtx-pkts% > 30</p>
<i>[<default checking interval>]</i>	The default frequency that PacketShaper will use to check for this event. When you register this event, you can substitute a different interval. For standard PacketShaper units, you can specify 1m (one minute) or 1h (one hour).

Examples:

```
event new NetworkInefficiency tcp-efficiency%.link<$1 1m
```

```
event new WebQos tcp-conn-aborts%.class>$1 1h
```

event override

Override the inherited user event by creating a local copy of the event in the current configuration.

```
event override <event_name>
```

You must make a local copy of an inherited user event before you can change the user event on the child configuration.

event register

Initiate event-checking and notification for an event in the current configuration. The maximum number of events that can be registered at one time is 32. To use the command-prompting mode, simply use event register, otherwise use the following command syntax.

```
event register <event name>(<object>,<threshold>,<re-arm>) [<checking interval>]
[email] [trap] [syslog] [limit=<n>]
```

<i><event name></i>	An existing predefined or user-defined event
<i><object></i>	The name of a link, partition, or class that is relevant to the event definition
<i><threshold></i>	The value used to trigger event notification. The value is substituted in the event's expression, which you defined with the event new command. If the condition in the expression occurs, it triggers the event notification that is registered for the event.
<i><re-arm></i>	The value that tells PacketShaper that it's okay to once again send event notifications. After the initial notification occurs for the threshold crossing, additional event messages traps, email, or syslog will not be sent until the re-arm condition occurs. The purpose of the re-arm value is to prevent excessive event notification.
<i>[<checking interval>]</i>	The frequency at which this condition should be checked. For standard PacketShaper units, you can specify 1m (one minute) or 1h (one hour).
<i>[email] [trap] [syslog]</i>	The notification mechanism for this event email, trap, or Syslog.
<i>[limit=<n>]</i>	The number of notifications to be sent within the 24-hour period from midnight to midnight. If you omit this option, the number of notifications is limitless.

Example:

```
event new WebQos tcp-conn-aborts%.class>$1 1h
```

```
event register WebQos(inbound/outside/http,70,50) 1m email limit=20
```

Note that in the above example, the event was defined with a default interval of one hour. When the event was registered, the specific class was identified with a threshold of 70%, a re-arm level of 50%, a 1-minute interval, and a limit of 20 notifications within a 24-hour period.

When an event exceeds the predefined threshold value, the event is in violation and the PacketShaper will automatically send out notification. PacketShaper will also send a notification when the re-arm level is crossed, allowing you to be alerted automatically when the event has been cleared.

event reset

Reset the user events system in the current configuration. This command removes all user-defined events and unregisters all events (user-defined and predefined).

```
event reset
```



Issuing the **event reset** command from the PolicyCenter command line interface can incorrectly trigger an error message stating that the operation failed, even if the operation executed correctly.

event show

Display email notification recipients, available events (both user-defined and predefined), registered events, and their status in the current configuration.

```
event show
```

event unregister

Stop checking an event in the current configuration.

```
event unregister <registration-id>|all
```

Use [event show](#) to display the registration IDs.

highav add

Define an access router for the access-link monitoring (high availability) feature. This feature allows legacy PacketShapers to deal with imperfect load-balancing and has the ability to respond to the occurrence of WAN link failure. When high availability is enabled, PacketWise can adjust partitions appropriately to prevent overloading any given WAN link and to account for lost available capacity due to router or link failure. High availability has two modes: [basic](#) and [advanced](#). *This command is applicable to legacy PacketShaper configurations only.*

```
highav add <address> <community>
```

where

<address>	IP address of the router
<community>	SNMP community string (password) for the router

Example:

```
highav add 10.10.10.10 pAss4WoRD
```

highav community

Change the community string of a high availability router. Use this command when the community string changes after you have already defined the router with the [highav add](#) command. *This command is applicable to legacy PacketShaper configurations only.*

```
highav community <address/sysname> <community>
```

where

<address/sysname>	The router's IP address or system name
<community>	New SNMP community string (password) for the router

highav delete

Remove an existing router from the high availability configuration. *This command is applicable to legacy PacketShaper configurations only.*

```
highav delete <address/sysname>
```

highav disable

Disable link monitoring ([basic](#) mode) as well as link overload protection ([advanced](#) mode, if enabled). *This command is applicable to legacy PacketShaper configurations only.*

```
highav disable
```

highav enable advanced

Enable two high availability features: link monitoring/resizing (as in [basic](#) mode) and link overload protection. With the link monitoring/resizing feature, a legacy PacketShaper polls the configured router(s) every 30 seconds to assess the status (link up or link down) of the WAN link interfaces. If a link goes down, PacketWise will automatically adjust the total available capacity by subtracting out the capacity of the down link. With link overload protection, PacketWise can help prevent the overloading of an interface. PacketWise will use SNMP polling to access the actual throughput of each configured WAN link interface. If an interface approaches its configured capacity, PacketWise will pace the traffic sent through that interface to prevent overloading the link and reduce the number of retransmissions. This is accomplished by adjusting the size of the Inbound and Outbound partitions. *This command is applicable to legacy PacketShaper configurations only.*

```
highav enable advanced
```

To turn off the advanced mode of high availability, use the [highav disable](#) command.

highav enable basic

Enable the link monitoring/resizing high availability feature. When this feature is enabled, a legacy PacketShaper polls the configured router(s) every 30 seconds to assess the status (link up or link down) of the WAN link interfaces. If a link goes down, PacketWise will automatically adjust the total available capacity by subtracting out the capacity of the down link. *This command is applicable to legacy PacketShaper configurations only.*

```
highav enable basic
```

Suppose you have two routers, A and B. Router A has two 200K interfaces and Router B has one 100K interface. The total available capacity is 500K (unless you have set up an override see [highav override](#)). Now suppose one of Router A's 200K links goes down. With basic high availability enabled, PacketWise will not only detect the down link, it will also automatically reduce the total available capacity by the capacity of the down link (500K minus 200K = 300K).

To turn off the basic mode of high availability, use [highav disable](#).

highav interface add

Define the WAN link interface used on a previously-defined access router. *This command is applicable to legacy PacketShaper configurations only.*

```
highav interface add <address> <interface number/name> <inbound-bps> <outbound-bps>
```

where

<address>	The router's IP address or sysname
<interface number/name>	The name (ifname) or index number (ifindex) that identifies the interface. Examples of interface names are <i>ethernet 3/1</i> and <i>serial 0/1</i> . It is recommended that you identify the interface by name, not index, because ifnames are unique and persistent while index numbers can change dynamically. If you are using Cisco IOS v12.1 or above and have configured the router to make the ifindex persistent, you can safely identify the interface by index number. Note that ifname was not available in Cisco IOS before v11.1.
<inbound-bps>	Maximum inbound throughput that is expected to pass through the interface. Rates may be specified as integer bits per second, followed by a k (thousands), M (millions), or G (billions).
<outbound-bps>	Maximum outbound throughput that is expected to pass through the interface

Adding an interface will increase the router's available bandwidth unless you have set [override](#) values. The lowest value (override versus sum of interfaces) takes precedence. For example, suppose a router has two 400K interfaces and you have set an override of 600K. If you add another 200K interface, the override will take precedence (in other words, the router's available bandwidth will still be 600K). Make sure that you adjust

highav interface delete

Delete a previously-defined interface from the high availability configuration. *This command is applicable to legacy PacketShaper configurations only.*

```
highav interface delete <address> <interface number/name>
```

where

<address>	The router's IP address or sysname
<interface number/name>	The name or SNMP index number of the interface you want to remove

Deleting an interface may reduce the router's available bandwidth, depending on the override value. For example, suppose a router has two 400K interfaces and you have set an override of 600K. If you then delete an interface, the router's available bandwidth would be reduced to 400K; the override would be ignored since it's greater than the sum of the router's interfaces.

highav interface modify

Modify the settings for a previously-defined WAN link interface. *This command is applicable to legacy PacketShaper configurations only.*

```
highav interface modify <address> <interface number/name> <inbound-bps> <outbound-bps>
```

where

<address>	The router's IP address or sysname
<interface number/name>	The name or SNMP index number whose settings you want to modify
<inbound-bps>	Maximum inbound throughput that is expected to pass through the interface. Rates may be specified as integer bits per second, followed by a k (thousands), M (millions), or G (billions).
<outbound-bps>	Maximum outbound throughput that is expected to pass through the interface

highav override

Configure the inbound and outbound speed of the router. When an override is set, PacketWise uses this speed for calculating the WAN link capacity for the router, as opposed to using the sum of the interfaces. *This command is applicable to legacy PacketShaper configurations only.*

```
highav override <address> {<inbound-bps> <outbound-bps>} | none
```

where

<address>	The router's IP address or sysname
<inbound-bps> <outbound-bps>	Maximum inbound and outbound throughput that is expected to pass through the router. Rates may be specified as integer bits per second, followed by a k (thousands), M (millions), G (billions).
none	To remove the override, use none .

This optional approach might be used in a situation with multiple WAN access line interfaces on a router. If you don't expect to get perfect load balancing between the interfaces, you can configure a smaller value for the router than for the sum of the interfaces. If both interfaces are up, PacketWise would use the override value for the router when calculating the WAN access line capacity available for the router. If one of the interfaces goes down, PacketWise would use the capacity configured for the active interface (the values configured with the [highav interface add](#) command).

highav show

Show current high availability configuration and status. The output indicates the overall high availability capacity as well as the settings of each interface and router. *This command is applicable to legacy PacketShaper configurations only.*

```
highav show
```

High Availability: Mode = Basic

Access Set: In 500k Out 500k

Total Available Capacity: In 500k Out 500k

Router Address: 192.168.176.5

Active: yes

SysName: testnetrouter.bluecoat.com

Override Capacity: No Override Set

Interface: + ET0(1) speed: 10.0M

Interface Capacity: In 200k Out 200k

Router Address: 192.168.176.2

Active: yes

SysName: router1

Override Capacity: In 300k Out 300k

Interface: + ET0(1) speed: 10.0M

Interface Capacity: In 200k Out 200k

Interface: + ET2(3) speed: 1.5M

Interface Capacity: In 200k Out 200k

The table below describes the output.

High Availability	The current high availability mode (basic, advanced, or disabled)
Access Set	<p>Inbound and Outbound access link speed. In basic mode, these values are the same as Total Available Capacity.</p> <p>With link overload protection (a feature of advanced mode), these values are based on actual throughput observed through SNMP polling. More detailed information about observed minimum values are also listed.</p>

PolicyCenter CLI Commands

Total Available Capacity	<p>The total bps available based on the values configured for the interfaces and routers. It is the sum of the routers' capacities. A router's capacity is determined by the values set with the highav override command or by summing all the interfaces' capacities (if no override has been set).</p> <p>When high availability is enabled and a link becomes inactive, the Total Available Capacity will reflect this reduction of available bandwidth (that is, the inactive link's capacity will be subtracted out, assuming it is less than the override value).</p>
Router	<p>The router's IP address and sysname that were configured with the highav add command, the router status (active vs. inactive), and the override capacity (if one was set with the highav override command).</p>
Interface	<p>Interface name, SNMP index number, and the inbound and outbound capacities that were configured with the highav interface add command. If you see "Unknown" for the interface name, your router's OS may not support the ifname variable. For example, ifname was not available in Cisco IOS before v11.1.</p> <p>A "+" indicates the interface is active; a "-" indicate the interface is inactive.</p> <p>If advanced mode is enabled, the actual bps throughput (based on SNMP polling) is listed.</p>

history

The **history** command displays the last 20 commands that were entered into the command line interface; each command is prefixed by a number. Any command on the history list can be executed by using the `!<n>` command, where **<n>** is the number next to the command on the history list.

```
history
```

For example:

```
history
```

```
31: setup show
32: help me dump
33: help class rule
34: help setup secure
35: traffic flow -tuO
36: traffic flow -tIPc /inbound/default
37: setup shaping on
38: setup discovery on
39: traffic tree
40: link show
41: class show /inbound/default
42: traffic bandwidth /inbound
43: help class new
44: hostdb show
45: sys info
46: traffic bandwidth
47: cat 9.256/log/bootlog
48: ls 9.258/diag
```

49: setup shaping off
 50: setup discovery off

Typing **!40** would repeat the **links show** command.

hl add

Add entries to an existing host list in the current configuration. When specifying multiple names and/or addresses, separate each with a space.

```
hl add <hostlist> <host> [<host> ...]
```

where *<hostlist>* is an existing host list name, and *<host>* can be specified in any of the following ways:

Type of <i><host></i>	Example
Host IP address	192.168.1.10
Range of IP addresses Use a dash – with no spaces – between the low and high address in the range.	192.168.1.100- 192.168.1.200
Address of the subnet; the CIDR number specifies the number of constant bits in the address range	192.168.10.0/24
Range of subnet addresses; the CIDR number specifies the number of constant bits in the address range Use a dash between the low and high address in the range. Spaces are not allowed before or after the dash or slash characters.	192.168.10.0- 192.168.20.0/24
DNS name Note: Do not use domain names if you will be using the host list with the host sidedness feature.	www.yourcompany.com

Example:

```
hl add competitors yourcompany.com 192.168.1.00-192.168.1.200
```

hl delete

Remove one or more items from an existing host list in the current configuration.

```
hl delete <hostlist> <host> [<host> ...]
```

where *<hostlist>* is an existing host list name, and *<host>* can be specified in any of the following ways:

Type of <i><host></i>	Example
Host IP address	192.168.1.10
Range of IP addresses Use a dash – with no spaces – between the low and high address in the range.	192.168.1.100- 192.168.1.200
Address of the subnet; the CIDR number specifies the number of constant bits in the address range	192.168.10.0/24
Range of subnet addresses; the CIDR number specifies the number of constant bits in the address range Use a dash between the low and high address in the range. Spaces are not allowed before or after the dash or slash characters.	192.168.10.0- 192.168.20.0/24
DNS name	www.yourcompany.com



You can only remove hosts the way they were originally added to the host list. For instance, suppose you add a host to the host list by specifying a single IP address. The only way to remove the host is by specifying the single address. You cannot remove this host by entering a range of addresses, a subnet, or a range of subnet addresses.

hl new

Create a host list of DNS names, IP addresses, and/or subnets in the current configuration. You can combine names and addresses in the same list. When specifying multiple names and/or addresses, separate each with a space.

```
hl new <hostlist> [<host> [<host> ...]]
```

where *<hostlist>* is a descriptive name, up to 127 characters; the slash (/) and backslash (\) characters may not be used.

The *<host>* can be specified in any of the following ways:

Type of <i><host></i>	Example
Host IP address	192.168.1.10
Range of IP addresses Use a dash – with no spaces – between the low and high address in the range.	192.168.1.100- 192.168.1.200
Address of the subnet; the CIDR number specifies the number of constant bits in the address range	192.168.10.0/24

Type of <host>	Example
Range of subnet addresses; the CIDR number specifies the number of constant bits in the address range Use a dash between the low and high address in the range. Spaces are not allowed before or after the dash or slash characters.	192.168.10.0-192.168.20.0/24
DNS name Note: Do not use domain names if you will be using the host list with the host sidedness feature.	www.yourcompany.com

Host lists are useful when creating classes based on hosts, retrieving host accounting data, and defining exception lists for adaptive response host agents.

The **hl new** command accepts any addresses and/or names that are syntactically correct. It does not validate the existence of the entries.

To add entries to the host list after it's created, use the [hl add](#) command.

Examples:

```
hl new BigGifs www.yourcompany.com 192.168.0.116
```

```
hl new insidelist
```

hl override

For PolicyCenter / Units in shared mode only

Override an inherited host list by creating a local copy of the list in the current configuration.

```
hl override <list-name>
```

You must make a local copy of an inherited host list before you can change the host list on the child configuration.

hl refresh

Update the host lists in the current configuration with the latest data from the DNS server.

```
hl refresh
```

hl resolve

Display the addresses that are mapped to a particular host list name in the current configuration.

```
hl resolve <hostlist>
```

Example:

```
hl resolve BigGifs
```

PolicyCenter CLI Commands

ldap:///biggifs,ou=hostlists,ou=m10-pat,ou=pscfg,o=bluecoat.com: 198.3.99.199,
192.168.0.116, 204.71.177.35

hl rm

Remove a host list from the current configuration.

```
hl rm <hostlist>
```

Host lists cannot be removed if they are currently being used (for example, in a class matching rule or a host side list).

hl show

Display a list of all host lists in the current configuration or show the details of a specific host list.

```
hl show [<list_name>]
```

To show all host lists and all host values:

```
hl show *
```

Host values are listed alphabetically or by top-level domain order.

host accounting categories

Create the names of categories to be used in the host accounting feature in the current configuration. For example, an ISP can create categories such as premium, standard, or free, and then assign traffic classes to the appropriate category. This would allow the service provider to charge different rates for premium and standard traffic and discount traffic to free services (such as the ISP support computers).

```
host accounting categories none|<category> [<category> ...]
```

You must create all your categories at once, separating each name with a space. The total string of category names must be 200 characters or less. Forty four is the maximum number of categories you can create, assuming sufficient resources on your PacketShaper.

Note that you cannot issue this command while the measurement engine is in the process of starting or resetting.

After creating categories, use the [class category](#) command to assign traffic classes to the categories.

You cannot selectively add or delete categories. If you later want to modify your category list, you must use the **host accounting categories** command and specify the complete list of categories you want to have. You will then need to reset the unit to create the new categories in the measurement database. Note that the host accounting measurement data will be cleared when you create categories, so make sure you retrieve your measurement data before adding categories.

Each category name is actually the name of a measurement variable. If you give the **measure show host accounting** command, you will see the host accounting categories listed as variables.

To remove all categories, use:

```
host accounting categories none
```

host accounting enable

Enable or disable recording of per-host accounting data in the measurement engine database in the current configuration. The host accounting feature allows you to store and retrieve measurement data per IP address, without having to create a traffic class for each user. This feature is especially useful with dynamic partitions, which create subpartitions for each IP address or subnet. You can use host accounting to see total bytes sent and received per IP address.

```
host accounting enable <mode> [<interval-minutes> [<max-samples>]]
```

<mode>	Specify the host location for which you want to record data: inside record data for inside hosts outside record data for outside hosts both record data for inside and outside hosts none turn off host accounting
<interval-minutes>	Number of minutes between each recorded sample (the default is 10, the minimum is 1, and the maximum is 1440 minutes)
<max-samples>	Maximum number of samples that can be stored in the host database (default is 1,000,000). This value needs to be greater than the concurrent host limit on your unit (this limit varies by model). <max-samples> may require some experimentation. Try a large number (such as 3,000,000) and see if the unit stores host data for a sufficient length of time. If it stores only three weeks of data (and you need it to store a month's worth), you'll need to increase the <max-samples> value. Bear in mind that the larger number of samples you store, the more disk space you will need.

Note that you cannot issue this command while the measurement engine is in the process of starting or resetting.

The building of the host accounting measurement data file can take awhile; the more categories and samples you have, the longer it takes to build the file. While the file is building, you will not be able to issue any commands in the current remote login session. If you open another session while you are waiting, you can issue any command except for the **measure show** command.

Note that any pre-existing host accounting data will be cleared when the data file is built, so make sure you retrieve your measurement data before enabling host accounting.

After you enable host accounting host accounting measurement data will not begin recording again until the next full interval. For example, assume <interval-minutes> is 2 and you reset the unit. When you give the **measure show** command, the message indicates "Measurement engine is waiting until 15:03 to start." At 15:03 all other measurement groups will begin recording, but host accounting will not begin recording until 15:04 (the next interval).

To disable recording, use:

```
host accounting enable none
```

host accounting retrieve

Retrieve host accounting data in the current configuration. You can display the data on the screen or save it into a comma separated value (.CSV) file. You can select hosts by IP address, CIDR subnet, or membership in a host list.

PolicyCenter CLI Commands

```
host accounting retrieve [dns] <ip-addr>|<subnet>/<cidr>|<hostlist>|all from <start-date-time> to <end-date-time> [into <file>|to <file>]
```

[dns]	Specify dns to include DNS names, when available, in the output (instead of IP addresses)
<ip-addr>	Designate the hosts you want to retrieve data for, using one of the following specifications:
<subnet>	<ip-addr>— host IP address
/<cidr>	<subnet>/<cidr>— the name of the subnet; the CIDR number specifies the number of constant bits in the address range
<hostlist>	<hostlist> — the name of a host list file
all	all — all entries in the time range
from <start-date-time>	Specify the starting and ending date and/or time to retrieve data for. The <start-date-time> and <end-date-time> are required parameters you must specify a date, a time, or the date and time. If a time is omitted, midnight is assumed; if a date is omitted, today's date is assumed. Dates and times can be entered in the following formats:
to <end-date-time>	M/D (for example, 5/3 midnight of May 3) M/D HH:MM (for example, 5/3 13:15 1:15pm on May 9) HH:MM (for example, 9:00 9am today) You can also specify a relative date for example, -7 for 7 days ago. Note: If the end date is after today's date, PacketShaper assumes the date you meant was last year's date. If the end date is before the current date and after the start date, PacketShaper will display the requested data (if any exists).
into <file> / to <file>	The into literal dumps the records to the file named <file>. If <file> already exists, the records will be appended to the existing file. The to literal also dumps the records to the file named <file>, but it overwrites the contents of <file> if it already exists. If no path is specified, the file is stored in the current folder (9.256/, by default). If into <file> or to <file> is omitted, the records appear on the screen.

The output contains, for each specified host, a comma-separated-values list of the total bytes recorded over the time period as well as the total for each category.

```
host accounting retrieve dns all from 14:12 to 14:14
```

```
# 26-Jul-2001 14:12:00 to 26-Jul-2001 14:14:00
```

```
host , bytes , web , overhead
```

```
r2.us.rmi.yahoo.com ,9085,9085,0
```

```
ck101.rmi.yahoo.com ,377,377,0
```

```
store.yahoo.com ,1684,1684,0
```

```
...
```

```
...
```

```
10.7.6.62,42,0,0
pal.ads.vip.sc5.yahoo.com,1289,1289,0
```

Or, to see the total usage for the month of March for all of the hosts in a certain subnet:

```
host accounting retrieve 192.168.1.0/24 from 3/1 00:00 to 3/31 23:59
```



You can also use the **measure dump** command to retrieve host accounting data.

host accounting show

List the recording parameters for host accounting that were set with the [host accounting enable](#) command. It also lists all the categories that were created with the [host accounting categories](#) command, along with the traffic classes assigned to each category (with the [class category](#) command).

```
host accounting show
```

Recording Mode: both

Sample Interval: 5 minutes

Allocated Samples: 1000000

Category	Traffic Class

web	
	/Inbound/http
	/Outbound/http
overhead	
	/Inbound/ICMP
	/Outbound/ICMP

hostdb side auto

Enable automatic side detection in the current configuration. In this mode (the default), PacketShaper automatically determines whether a host is inside or outside, relative to the PacketShaper. This is appropriate for many network topologies; however, for complex topologies, you can manually override the automatic host side detection and force the placement of certain hosts or subnets on the appropriate side (see [hostdb side manual](#)).

```
hostdb side auto
```

After you turn on automatic side detection, the current settings display. For example:

```
Mode: Automatic
```

```
Inside:
```

```
Host list: none
```

```
Outside:
```

```
Host list: none
```

hostdb side default

Set host side detection to its default mode (auto mode) in the current configuration. If the unit is subscribed to PolicyCenter, the **default** option tells PolicyCenter to delete the setting in the local configuration and inherit from the parent configuration.

```
hostdb side default
```

hostdb side manual

Enable manual side mode in the current configuration; this is necessary when PacketShaper isn't able to automatically detect the correct side for certain hosts. When this mode is enabled, you can force the placement of certain hosts on the appropriate side (inside or outside). This is done by using the [hostdb side set](#) command to assign hosts or host lists to the inside or outside. The side lists are not actually used until manual side mode is enabled.

```
hostdb side manual
```

After you turn on manual mode, the current settings display:

```
Mode: Manual
Inside:
  Host list: none
  192.21.18.175
  192.21.18.177
  192.21.18.178-192.21.18.180

Outside:
  Host list: none
```

For any host that isn't assigned to a specific side when manual side mode is enabled, PacketShaper will use its normal mechanism for determining and setting a side. In other words, the sides of all other hosts are detected automatically.

hostdb side reset

Clear the side settings for a particular host or all hosts in the current configuration. The next time PacketShaper sees a flow from that address it will again try to figure out whether the host is inside or outside. This might be necessary if a particular host is seen on the wrong side – you can add the host to the proper side list (inside or outside) and then reset the host so that PacketShaper will rediscover the host and place it on the correct side. To see which side PacketShaper considers a host to be on, use the `hostdb show` command.

```
hostdb side reset all|<ip-addr>
```

Example:

To clear the side setting for 224.0.1.3:

```
hostdb side reset 224.0.1.3
```

If you immediately issue the **hostdb show** command, you'll see that the entry in the Side column for this particular host has been cleared.

IP Address	Side	Speed/Effective	TCP/UDP	Idle Time	Cache I O	Compress I R S NC
10.10.1.52	in	10.0M/4.5M	3/0	50s	? ?	n y n n
10.1.1.81	out	0/44756	1/0	0s	I ?	y n n n
224.0.1.3		0/0	0/0	38s	N O	n n n n
10.1.5.1	in	0/0	0/0	3s	? O	n n n n
10.7.12.20	out	0/0	0/0	24s	I ?	n n n n
10.7.10.10	out	0/0	0/0	0s	? ?	n n y n
6 entries						

hostdb side rm

Remove a host from the manually configured side list in the current configuration. Use this command if you no longer want a particular host assigned to a side. After you remove a host, PacketShaper will determine and assign a side to the host, using its normal mechanism.

```
hostdb side rm list:<hostlist>|<ip-addr>|<subnet>/<cidr>|hosts|all
```

list: <hostlist>	Name of the host list file to be removed from the side list Note: This does not delete the list – use the hl rm command if you want to delete the list.
<ip-addr>	Host IP address or a range of IP addresses to be removed from the side list To specify a range, use a dash – with no spaces – between the low and high address in the range (for example, <i>192.168.1.100-192.168.1.200</i>).
<subnet> /<cidr>	The address of the subnet or a range of subnet addresses to be removed from the side list; the CIDR number specifies the number of constant bits in the address range To specify a subnet range, use a dash between the low and high address in the range (for example, <i>192.168.10.0-192.168.20.0/24</i>). Spaces are not allowed before or after the dash or slash characters.
hosts	Removes all individually defined IP addresses, ranges, and subnets (but not host lists)
all	Removes all hosts from the inside and outside lists, including host lists



To verify the host was removed, use the [hostdb side show](#) command.

Examples:

```
hostdb side rm 172.17.72.0-172.17.75.0/22
```

```
hostdb side rm list:outside_list
```

hostdb side set

Assign hosts to the inside or outside in the current configuration. This is necessary when PacketShaper isn't able to automatically detect the correct side for certain hosts.

```
hostdb side set inside|outside list:<hostlist>|<ip-addr>|<subnet>/<cidr>
```

Designate inside or outside hosts, using one of the following specifications:

list: <hostlist>	<p>The name of a host list created with the hl new command; only one host list per side is allowed.</p> <p>Note: Host lists are the recommended method of specifying hosts. However, you cannot use a host list that contains domain names.</p>
<ip-addr>	<p>Host IP address or a range of IP addresses</p> <p>To specify a range, use a dash – with no spaces – between the low and high address in the range (for example, <i>192.168.1.100-192.168.1.200</i>).</p>
<subnet> /<cidr>	<p>The address of the subnet or a range of subnet addresses; the CIDR number specifies the number of constant bits in the address range</p> <p>To specify a subnet range, use a dash between the low and high address in the range (for example, <i>192.168.10.0-192.168.20.0/24</i>). Spaces are not allowed before or after the dash or slash characters.</p>

Additional Information

- After you assign hosts to sides, you will need to enable manual side mode with the [hostdb side manual](#) command.
- For any host that isn't assigned to a specific side when manual side mode is enabled, PacketShaper will use its normal mechanism for determining and assigning a side.
- To remove a host after you have assigned it to a side, use the [hostdb side rm](#) command.
- To view a list of hosts assigned to each side, use the [hostdb side show](#) command.
- A maximum of 32 entries can be assigned to the inside and outside. An entry can be a single IP address, a range of IP addresses, a subnet, a subnet range, or a host list. Only one host list can be assigned to a side.

Examples:

In this example, host lists named *inside_list* and *outside_list* were created with the **hl new** command. *inside_list* contains a list of hosts and subnets that are known to be on the inside of PacketShaper and *outside_list* contains the hosts known to be on the outside of PacketShaper. To assign each of the hosts in *inside_list* an inside designation, use this command:

```
hostdb side set inside list:inside_list
```

To assign each of the hosts in *outside_list* an outside designation, use this command:

```
hostdb side set outside list:outside_list
```

And then enable manual mode:

```
hostdb side manual
```


Each time you assign IP addresses or subnets with the **hostdb side set** command, the specified hosts are added to the appropriate side – you do not overwrite previous settings. For example, the following two commands will assign two hosts to the outside:

```
hostdb side set outside 192.21.18.172
hostdb side set outside 192.15.17.45
```

However, this rule does not apply to host lists since only one host list is allowed per side. If you assign a host list to a side that already has a host list defined, this list will override the one that was previously defined.

hostdb side show

Display host side settings in the current configuration. The current mode (auto vs. manual) is displayed, along with the list of hosts assigned to each side.

```
hostdb side show
```

Mode: Manual

Inside:

Host list: inside_list

192.21.18.175

192.21.18.177

192.21.18.178-192.21.18.180>

Outside:

Host list: outside_list

hostdb topusers

Determine which hosts or users are consuming the most bandwidth in the current configuration. You can configure PacketShaper to track the Top Talkers (hosts which initiate the most traffic) and Top Listeners (hosts which receive the most traffic).

To display statistics for the top 20 bandwidth users per traffic class either receivers or senders use the following commands:

```
hostdb topusers start <tclass> [talk|listen]
```

```
hostdb topusers stop <tclass> [talk|listen]
```

```
hostdb topusers reset <tclass> [talk|listen]
```

```
hostdb topusers show [<tclass>] [talk|listen]
```

start	Starts tracking top hosts (talkers or listeners) for a traffic class
stop	Stops tracking top hosts for a traffic class

PolicyCenter CLI Commands

reset	Clears the list of top hosts and restarts the host-tracking process
show	<p>Displays the hosts or users that have used the highest percentage of bandwidth in the class since tracking was started. The list is cleared with the hostdb topusers reset <aclass> command or when you reset the unit.</p> <p>A host stays on the top-20 list until another host uses more bandwidth, at which point the host may drop off the list entirely or move further down the list. For example, suppose top talkers is turned on for the Inbound/HTTP class, and cnn.com is the top consumer with 22%. If another host, yahoo.com, later consumes more bandwidth than cnn.com, yahoo.com might go to the top of the list and cnn.com would drop lower on the list.</p>

Additional Information

- A total of 32 Top Talkers and Top Listeners (combined) can be enabled at one time. Thus, you can enable both Top Talkers and Top Listeners for up to 16 different classes. Or, if you enable one or the other (just Top Talkers or just Top Listeners), you can track top hosts on 32 different classes.
- For non-IP traffic, PacketShaper does not track sessions or hosts. Therefore, traffic for non-IP protocols (IPX, AppleTalk, NetBEUI, DECnet, FNA, and SNA) will not appear in the Top Talker or Top Listener lists.
- The Group Name column in the output shows one of the user group names to which the user belongs. If a top user belongs to more than one group, an ellipses will appear. To see all the group names for a user, use the [setup bcaaa server-test](#) command. Note that the only group names that display are ones for which a user group class has been created. The user may belong to other Active Directory groups, but they will not be listed for the user unless a class exists for that group.

Examples:

To start top talker tracking on the Inbound/HTTP class:

```
hostdb topusers start inbound/http talk
```

To see a list of top talkers in the Inbound/HTTP class:

```
# hostdb topusers show inbound/http talk
```

Top talker analysis for inbound class HTTP.

Duration: 02:09:18

7 active entries.

User Name	Group Name	DNS Name
Percent	IP Address	
N/A	N/A	a184-84-222-
35.deploy.akamaitechnologies.com	52	184.84.222.35
N/A	N/A	a184-84-222-
107.deploy.akamaitechnologies.com	24	184.84.222.107
john.smith	group-sales	No such name
7	10.200.10.129	
N/A	N/A	20052lpweb01.redcrossblood.org
4	174.120.176.2	
N/A	N/A	75.126.14.205-

```
static.reverse.softlayer.com          4          75.126.14.205
N/A                                   N/A          server-205-251-203-
169.lax3.r.cloudfront.net            2          205.251.203.169
N/A                                   N/A          a184-84-222-
115.deploy.akamaitechnologies.com     1          184.84.222.115
```

```
# hostdb topusers show
```

```
2 active top user sessions.
```

Direction	Class	T/L	Duration
inbound	HTTP	talker	02:13:13
outbound	HTTP	listener	02:13:13

ipfilter clear

Removes all IP filters (default) in the current configuration or the IP filter that you specify.

```
ipfilter clear [<id>]
```

ipfilter clear	Removes all IP filters. To remove all IP filters from a PolicyCenter configuration, access that configuration via the config view command, then issue the CLI command ipfilter clear ; no draft configuration is necessary.
[<id>]	Removes the IP filter identified by [<id>]. To delete a single IP filter from a PolicyCenter sharable configuration, create a draft version of that configuration and then issue the command ipfilter clear <id> from the draft. The deleted filter will be permanently deleted when you commit the draft. To see a list of all configured IP filters and their identifiers, use the ipfilter show command.

Examples

This example removes all configured IP filters from the PacketShaper.

```
ipfilter clear
```

This example removes only the IP filter with the identifier "DC73DA16".

```
ipfilter clear DC73DA16
```

ipfilter iponly

Configures a configuration to relay only IP traffic.

```
ipfilter iponly on|off
```

on	Creates an IP filter that relays only IP traffic (applies to all interfaces).
off	(Default). Configures the PacketShaper to relay both IP and non-IP traffic.

ipfilter show

Shows all configured IP filters in the current configuration.

```
ipfilter show
```

Returned data include the status of the [ipfilter iponly](#) setting (whether or not the PacketShaper will relay only IP or both IP and non-IP traffic), the unique identifier of the IP filter, number of hits (instances where the IP filter rule matched a packet), and the configuration of the IP filter.



You cannot configure IP filters for a PolicyCenter configuration, although PolicyCenter will show you the filters that have been configured directly on a unit. To add filters, log in to the unit and issue the desired **ipfilter discard|onlyaccept|passthrough** command.

Example

In this example, the PacketShaper is configured to relay both IP and non-IP traffic, and one passthrough IP filter has been configured on the outside interface:

```
ipfilter show
Relay all traffic.
Exclude Filters: total 1
  [DC73DA16] hits 0 Outside
             src 172.21.1.44 (ffffffff) --> passthru
Include filters: total 0

ipfilters MIB:
[ 0] outs                0 [ 1] onlyAccepts                0
[ 2] onlyExcludes        0 [ 3] nonIpDiscarded            0
[ 4] atalkDiscarded      0 [ 5] ipxDiscarded              0
[ 6] netbiosDiscarded    0 [ 7] snaDiscarded              0
[ 8] fnaDiscarded        0 [ 9] decDiscarded              0
```

license birthcert

Reinstalls the hardware birth certificate. This operation may be necessary if the birth certificate got corrupted or accidentally deleted. This command requires that the PolicyCenter S-Series appliance have Internet access.

```
license birthcert
```

license load

License keys, as well as licenses for add-on components, can be downloaded from the Symantec Licensing Portal, uploaded to the PolicyCenter S-Series appliance or a web server, and then activated with the **license load** command.

```
license load file://localhost/<file path> | http://<url>
```

where *<file path>* is the name and location of the license key file on the PolicyCenter S-Series appliance and *<url>* is the URL of the file on a web server. Filenames are case sensitive so make sure to use the proper upper/lower case.

Examples:

```
license load file://localhost/9.256/mylicense.bcl
```

```
license load http://mywebserver.com/licenses/mylicense.bcl
```

Additional Information

- After loading the license, you will need to [reset](#) the PolicyCenter S-Series appliance to activate the license.
- The license activation and expiration dates appear in the [version verbose](#) output and next to **Licensing Status** on the Info tab.
- If a license will be expiring within the next 90 days, a message (such as *PolicyCenter license will expire in 15 days*) appears on the Info tab.

license update

Connects to the Symantec licensing server and downloads the latest license keys. Use this command to immediately retrieve the license after initially configuring PolicyCenter or after a licensing renewal/purchase. This command requires that the PolicyCenter S-Series appliance have Internet access.

```
license update
```

Additional Information

- After updating the license keys, you will need to [reset](#) the PolicyCenter S-Series appliance to activate the license.
- The license activation and expiration dates appear in the [version verbose](#) output and next to **Licensing Status** on the Info tab..
- If a license will be expiring within the next 90 days, a message (such as *PolicyCenter license will expire in 15 days*) appears on the Info tab.

ls

Display file listings of a specific folder.

```
ls [<directory> | <file>]...
```

more

Display the named file, showing a single page and pausing before displaying the next page. More than one filename can be specified.

```
more [-<number>] <filename>
```

Providing an optional number will display the specified number of lines on one page.

PolicyCenter CLI Commands

net ip

View IP MIB statistics. These statistics are accumulated since the last PolicyCenter reset.

```
net ip
```

net nic

Show or clear network statistics, such as number of packets sent and received, dropped packets, and packets with errors. These statistics are accumulated since the last PolicyCenter reset.

```
net nic [clear]
```

net stat

Show network protocol statistics for TCP, UDP, and IP.

```
net stat
```

Sample output:

Tcp:

```
11997 active connection openings
1596 passive connection openings
11411 failed connection attempts
568 connection resets received
9 connections established
226806 segments received
227518 segments sent out
427 segments retransmitted
0 bad segments received.
12001 resets sent
```

Udp:

```
239029 packets received
0 packets to unknown port received.
0 packet receive errors
582639 packets sent
```

Ip:

```
741833 total packets received
0 with invalid headers
0 with invalid addresses
0 forwarded
0 with unknown protocol
0 incoming packets discarded
```

```

741815 incoming packets delivered
0 outgoing packets dropped
28 dropped because of missing route
0 fragments dropped after timeout
0 reassemblies required
0 packets reassembled ok
0 packet reassemblies failed
0 fragments received ok
0 fragments failed
0 fragments created

```

partition apply

Create a static partition for a traffic class in the current configuration.

```
partition apply <aclass><minBps>|<minPct>|uncommitted [<maxBps>|<maxPct>|none|fixed]
```

<aclass>	This traffic class and all of its children are partitioned together (those that are not already separately partitioned), so this traffic class becomes the root of a partitioned subtree of traffic classes.
<minBps> <minPct> uncommitted	<p>The minimum size of the new partition, specified in bits per second (<i>minBps</i>) or as a percentage of the parent partition's minimum size (<i>minPct</i>). If <i><minPct></i> is used, you must include the percent sign (for example, 10%). The minimum partition size is 1024 bps.</p> <p>The sum of all the partitions within either Inbound or Outbound can exceed the link size, allowing you to oversubscribe the link.</p> <p>Use the literal uncommitted to indicate that the guaranteed minimum allocation is whatever is not committed to other partitions. Normally uncommitted is used only by the default Inbound and Outbound partitions.</p>
[<maxBps> <maxPct> none fixed]	<p>Limit the maximum bandwidth used by a burstable partition. The maximum can be specified in bits per second (<i>maxBps</i>) or as a percentage of the parent maximum (<i>maxPct</i>). If <i><maxPct></i> is used, you must include the percent sign (for example, 10%). The maximum must be greater than the minimum. The maximum partition size varies by model.</p> <p>Specify none to allow the partition to use any available bandwidth. Specify fixed to prevent a partition from exceeding the <i><minBps></i> or <i><minPct></i> size. If you do not specify burstable or fixed, the partition defaults to burstable.</p>

Additional Information

- In order for partitions to take effect, traffic shaping must be enabled. See [setup shaping](#).
- When creating partitions, make sure you don't allocate bandwidth in such a way that Inbound/Default and Outbound/Default get "starved," that is, there is no bandwidth available for these classes. If this happens, traffic classification and policies may not work as expected.

Examples:

Create an inbound burstable partition (no maximum limit specified) of 10000 bps:

PolicyCenter CLI Commands

```
partition apply inbound/outside/http 10k
```

Create an inbound burstable partition of 20000 bps with the ability to borrow additional bandwidth from other partitions, if it is available, up to 30000 bps:

```
partition apply inbound/outside/http 20k 30k
```


Create a burstable partition for SAP that is 30% of the link (Inbound partition) size, with a maximum size of 40%:

```
partition apply inbound/sap 30% 40%
```

In the above example, if the link size is 1.5 Mbps, the SAP partition would get a minimum of 450 Kbps and a maximum of 600 Kbps.

partition dynamic apply

Create a dynamic per-user partition for a traffic class in the current configuration. Can be specified by IP address or by subnet.



Before you can create dynamic per-user partitions, you must create a static partition for the class using the [partition apply](#) command.

To create a dynamic partition which handles traffic for an IP address:

```
partition dynamic apply <tclass> per-address <side> <minBps>|<minPct>|uncommitted  
<maxBps>|<maxPct>|none|fixed
```

<tclass>	Name of the traffic class having a static partition you would like to subdivide for each user
<side>	Side (inside or outside) of the PacketShaper on which the user is located

<code><minBps> <minPct> uncommitted</code>	<p>Minimum amount of bandwidth to be assigned to each user, specified in bits per second (<i>minBps</i>) or as a percentage of the parent partition's size (<i>minPct</i>). If <i><minPct></i> is used, you must include the percent sign (for example, 10%).</p> <p>Use the literal uncommitted to indicate that the guaranteed minimum allocation is whatever is not committed to other partitions.</p> <p>Set this field to zero (0) to have PacketShaper allocate bandwidth equitably to each subpartition, so that the total of all subpartitions equals the static partition's size.</p> <p>Note: Minimum subpartition size is usually best handled by setting this field to zero and setting a maximum number of subpartitions (using the partition dynamic cap command). However, you must use a non-zero size if you want to implement per-session guaranteed rates within rate policies for this same traffic.</p>
<code><maxBps> <maxPct> none fixed</code>	<p>Maximum amount of bandwidth to be assigned to each subpartition, specified in bits per second (<i>maxBps</i>) or as a percentage of the parent partition's size (<i>maxPct</i>). If <i><maxPct></i> is used, you must include the percent sign (for example, 10%).</p> <p>Specify a <i><maxBps></i> value if you want to enforce a cap on each user or subnet even if more bandwidth is available. Managed bandwidth service providers are most frequently in this position, needing to cut off usage at agreed-upon, paid-for limits. If you don't want a maximum, specify none. Specify fixed to prevent a subpartition from exceeding the <i><minBps></i> or <i><minPct></i> size.</p> <p>Even if this field is left blank, the limit on the static, parent partition still restricts the total bandwidth for the aggregate of all subpartitions.</p>

To create a dynamic partition which handles traffic for a subnet:

```
partition dynamic apply <tclass> per-subnet /<cidr> <side><minBps>|<minPct>
<maxBps>|<maxPct>|none
```

<code><tclass></code>	Name of the traffic class having a static partition you would like to subdivide for each user
<code>/<cidr></code>	CIDR number specifying the number of constant bits in the address range
<code><side></code>	Side (inside or outside) of the PacketShaper on which the user is located

PolicyCenter CLI Commands

<code><minBps> / <minPct> uncommitted</code>	<p>Minimum amount of bandwidth to be assigned to each user, specified in bits per second (<i>minBps</i>) or as a percentage of the parent partition's size (<i>minPct</i>). If <i><minPct></i> is used, you must include the percent sign (for example, 10%).</p> <p>Use the literal uncommitted to indicate that the guaranteed minimum allocation is whatever is not committed to other partitions.</p> <p>Set this field to zero (0) to have PacketShaper allocate bandwidth equitably to each subpartition, so that the total of all subpartitions equals the static partition's size.</p> <p>Note: Minimum subpartition size is usually best handled by setting this field to zero and setting a maximum number of subpartitions (using the partition dynamic cap command). However, you must use a non-zero size if you want to implement per-session guaranteed rates within rate policies for this same traffic.</p>
<code><maxBps> / <maxPct> none fixed</code>	<p>Maximum amount of bandwidth to be assigned to each subpartition, specified in bits per second (<i>maxBps</i>) or as a percentage of the parent partition's size (<i>maxPct</i>). If <i><maxPct></i> is used, you must include the percent sign (for example, 10%).</p> <p>Specify a maximum value if you want to enforce a cap on each user or subnet even if more bandwidth is available. Managed bandwidth service providers are most frequently in this position, needing to cut off usage at agreed-upon, paid-for limits. If you don't want a maximum, specify none. Specify fixed to prevent a subpartition from exceeding the <i><minBps></i> or <i><minPct></i> size.</p> <p>Even if this field is left blank, the limit on the static, parent partition still restricts the total bandwidth for the aggregate of all subpartitions.</p>

After the dynamic partition is set up, whenever a new user begins generating flows in that class, a subpartition will be created for the user on the fly. The per-user partition remains in existence until it's re-used for new flows by the same user or needed by another user. A subpartition may be given to another user if there have not been any recent flows in the partition. To be more precise, a subpartition may be given to another user if 30 seconds have passed without any flows or if it's been five minutes since an established flow has sent any packets.

partition dynamic cap

Set the maximum number of active users allowed in the dynamic partition in the current configuration, and, optionally, create an overflow partition for users to tap into if the cap is exceeded.

```
partition dynamic cap <tclass><maxusers> [<overflowMinBps>|<overflowMinPct>|uncommitted  
<overflowMaxBps>|<overflowMaxPct>|none|fixed]
```

<code><tclass></code>	Name of the traffic class having a dynamic partition for which you would like to set a cap
<code><maxusers></code>	Maximum number of per-user partitions that can be created in this traffic class

<code><overflowMinBps></code> <code>/ <overflowMinPct></code> <code> uncommitted</code>	<p>Minimum amount of bandwidth in the overflow partition, specified in bits per second (<i>overflowMinBps</i>) or as a percentage of the parent partition's size (<i>overflowMinPct</i>). If <i><overflowMinPct></i> is used, you must include the percent sign (for example, 10%).</p> <p>Use the literal uncommitted to indicate that the guaranteed minimum allocation is whatever is not committed to other partitions.</p>
<code><overflowMaxBps></code> <code>/</code> <code><overflowMaxPct></code> <code> none fixed</code>	<p>Maximum amount of bandwidth in the overflow partition, specified in bits per second (<i>overflowMaxBps</i>) or as a percentage of the parent partition's size (<i>overflowMaxPct</i>). If <i><overflowMaxPct></i> is used, you must include the percent sign (for example, 10%). When a value is specified, the overflow partition can use available excess bandwidth if needed.</p> <p>Specify none to allow the overflow partition to use any available bandwidth. Specify fixed to prevent the partition from exceeding the <i><overflowMinBps></i> or <i><overflowMinPct></i> specification.</p> <p>If you don't specify a value, the overflow partition has a fixed size; when it's not using its reserved bandwidth, that bandwidth is available to other traffic.</p>

To remove the cap on a dynamic partition:

```
partition dynamic cap <tclass> none
```

partition dynamic remove

Remove a dynamic partition in the current configuration. The partition reverts to being static.

```
partition dynamic remove <tclass>
```

partition dynamic summary

Show all the configured dynamic partitions in the current configuration and the number of users currently using each partition. Users that are not active can be replaced by new users.

```
partition dynamic summary
```

Example:

```
partition dynamic summary
```

Partition Name	--- Users ---		--- Current User Details ---			
	Current	Cap	Active	Idle	Gone	LongGone
Inbound	7	none	3	0	3	1
http						

After a dynamic partition is set up, whenever a new user begins generating flows in that class, a subpartition will be created for the user on the fly. The per-user partition remains in existence until it's re-used for new flows by the same user or needed by another user. A subpartition may be given to another user if there have not been any recent flows in the partition.

PolicyCenter CLI Commands

A subpartition is considered **Idle** if it has not been active for 300 seconds (5 minutes). Idle subpartitions still have flows which are sending packets. A subpartition is considered **Gone** if the flows associated with it have been gone 30 seconds or less, or **LongGone** if they have been gone more than 30 seconds. When the dynamic partition cap has been reached, new subpartitions are created from LongGone and Gone partitions.

In other words, a subpartition may be given to another user if 30 seconds have passed without any flows or if it's been five minutes since an established flow has sent any packets.

partition remove

Remove a static partition from a traffic class in the current configuration. The bandwidth allocated to this traffic class is returned to the parent partition.

```
partition remove <tclass>
```

partition show

Display current partition usage for static, dynamic, or both types of partitions in the current configuration.

```
partition show [<tclass>] [static|dynamic {<ip-addr>|<subnet>/<cidr>|<ip-addr-range>}|clear|config]
```

where:

<tclass>	Displays partition statistics for the specified traffic class
static	Lists only static partitions
dynamic	Lists only dynamic partitions
dynamic {<ip-addr> <subnet>/<cidr> <ip-addr-range>}	Lists only dynamic partitions for the specified IP address, subnet, CIDR, or range of IP addresses
clear	Resets the displayed partition statistics and then displays a partition list. Note that the statistics will not necessarily show as zero after this reset, because traffic activity could be recorded instantaneously. Note: The clear option is not applicable in conjunction with the <tclass> option.
config	Displays the minimum and maximum usage. The usage maximum is a partition's burst limit.

Examples:

To display partition statistics for all partitions, omit all parameters.

```
PacketShaper# partition show
```

Partition name	Size	Grntd	Prior	Curr	1-Min
----------------	------	-------	-------	------	-------

Peak	Min /	Max	Excess		Usage		Avg	

--								
/Inbound	1.5M	1.5M	0	0	1024	1024	2048	3687
537k								
/Inbound/MPEG-Audio	500k	1.5M*	0	0	0	0	0	0
0								
/Inbound/WinMedia	0	1.5M*	0	0	0	0	0	0
0								

--								
/Outbound	1.5M	1.5M	0	0	0	0	1143	405
6986								
/Outbound	1.5M	1.5M	0	0	0	0	1143	405
6986								
/Outbound/157	0	1.5M	0	0	0	0	0	0
0								
/Outbound/157	0	1.5M	0	0	0	0	0	0
0								
/Outbound/157/74.125.77.99	0	100k	0	0	0	0	0	0
0								
/Outbound/157/74.125.77.102	0	100k	0	0	0	0	0	0
0								
/Outbound/157/10.2.2.100	0	100k	0	0	0	0	0	0
0								
/Outbound/157/91.189.90.41	0	100k	0	0	0	0	0	0
0								
/Outbound/157/74.125.77.104	0	100k	0	0	0	0	0	0
0								
/Outbound/157/74.125.79.102	0	100k	0	0	0	0	0	0
0								
/Outbound/157/91.189.90.40	0	100k	0	0	0	0	0	0
0								

This output lists both minimum and maximum partition size settings. It also lists the rate of priority traffic. In addition, it displays an asterisk (*) next to any minimum or maximum value that isn't "pure" that is, if the programmed value was adjusted due to (1) oversubscription or (2) the use of the strings **fixed** or **none**. The adjusted values, not the programmed values, are listed, followed by an asterisk.

PolicyCenter CLI Commands

The Usage field represents the current bandwidth assigned to the partition, including guaranteed rate and excess rate for classes with rate policies, and any bandwidth currently allocated to classes with priority policies. Current rate and one-minute averages are bits-per-second rates.

To list only dynamic partitions:

PacketShaper# **partition show dynamic**

Partition name	Size		Grntd	Excess	Prior	Usage	Curr	1-Min	
	Min /	Max						Avg	Peak

-									
/Inbound/10.2.12.171	0	1.0M	0	0	4608	4608	2804	944	3325
/Inbound/10.9.50.93	0	1.0M	0	0	0	0	0	0	0
/Inbound/10.9.50.47	0	1.0M	0	0	0	0	4	13	53
/Inbound/10.9.50.27	0	1.0M	0	0	0	0	0	8	214
/Inbound/152.86.13.0	0	1.0M	0	0	0	0	0	0	0
/Inbound/10.9.50.60	0	1.0M	0	0	0	0	0	0	0
/Inbound/10.9.50.48	0	1.0M	0	0	0	0	0	0	0
/Inbound/10.9.50.44	0	1.0M	0	0	0	0	0	0	0
/Inbound/8.81.20.0	0	1.0M	0	0	0	0	0	0	0
/Inbound/10.9.50.92	0	1.0M	0	0	0	0	0	0	0
/Inbound/10.9.51.3	0	1.0M	0	0	0	0	0	9	33
/Inbound/10.9.50.75	0	1.0M	0	0	0	0	0	39	1943
/Inbound/114.185.22.0	0	1.0M	0	0	0	0	56	33	302
/Inbound/10.9.50.1	0	1.0M	0	0	0	0	61	39	62
/Inbound/10.9.50.109	0	1.0M	0	0	0	0	276	122	655
/Inbound/10.9.51.13	0	1.0M	0	0	0	0	18	38	238
/Inbound/114.86.25.0	0	1.0M	0	0	1280	1280	265	178	338

-									
/Outbound/157/224.0.0.251	0	100k	0	0	0	0	0	0	0

To list dynamic partitions in a subnet (CIDR=24):

PacketShaper# **partition show dynamic 10.9.50.0/24**

Partition name	Size		Grntd	Prior	Curr	1-Min			
	Min /	Max				Excess	Usage	Avg	Peak

-									
/Inbound/10.9.50.49	0	1.0M	0	0	0	0	0	0	0
/Inbound/10.9.50.93	0	1.0M	0	0	0	0	0	0	0
/Inbound/10.9.50.47	0	1.0M	0	0	0	0	0	2	53
/Inbound/10.9.50.27	0	1.0M	0	0	0	0	0	2	214
/Inbound/10.9.50.60	0	1.0M	0	0	0	0	0	0	0
/Inbound/10.9.50.48	0	1.0M	0	0	0	0	0	0	0
/Inbound/10.9.50.44	0	1.0M	0	0	0	0	5	7	27

```

/Inbound/10.9.50.92      0      1.0M      0      0      0      0      0      0
0
/Inbound/10.9.50.75      0      1.0M      0      0      0      0      0      10
1943
/Inbound/10.9.50.1       0      1.0M      0      0      0      0      9      30
233
/Inbound/10.9.50.109     0      1.0M      0      0      0      0      535     127
655

```

pc password

Modify the password for the directory server or touch user (serial console access only). A password can be up to 19 characters long and include all printable characters, including spaces, periods, underscores, and dashes.

```
pc password [touch|DS] [default]
```

After you enter the command, the system will prompt you for the current password, new password, and have you retype the new password.

If you forget the directory server password, you can restore the default password (admin) by logging in from the serial console as the touch user and then issuing the following command:

```
pc password ds default
```

If you forget the touch user password, you need to log in as the recovery user.

pc portal library

Show the current portfolios of customer portal files available for distribution from PolicyCenter to individual PacketShapers. This command can only be issued by network administrators with touch-role access to the *PC* organization.

```
pc portal library [verbose]
```

The **pc portal library** command shows the name of the available portfolios only. Use **pcportal library verbose** to view the names of all the customer portfolio files within each portfolio.

pc radius acct

Set up the configuration of the RADIUS accounting service for use with PolicyCenter. This feature allows you to have an audit trail for user logins. This command can only be issued by network administrators with touch-role access to the *PC* organization. Note that PolicyCenter does not allow a RADIUS user to log in with the user name **admin**.

```
pc radius acct default | off | on | [primary {<host> <shared_secret> [<port>]}|delete]
| [secondary {<host> <shared_secret> [<port>]}|delete]
```

default	Return RADIUS accounting to its default off setting
off	Disable RADIUS accounting

PolicyCenter CLI Commands

on	Enable RADIUS accounting
<host>	IP address or DNS of the RADIUS server
<shared_secret>	Specify the designated secret (password)
<port>	To access the RADIUS server with a specific port, specify a port number. Otherwise, the default port will be used.
delete	Delete this RADIUS accounting server.

Example:

```
pc radius acct primary 172.21.8.50 secretpwd
```

pc radius auth

Set up the configuration of the RADIUS authentication service so that users can log in to PolicyCenter using RADIUS credentials. RADIUS authentication is an optional method for users to log into the PolicyCenter UI. Using third-party RADIUS servers enables you to have central configuration of user accounts. This command can only be issued by network administrators with touch-role access to the *PC* organization. Note that PolicyCenter does not allow a RADIUS user to log in with the user name **admin**.

```
pc radius auth default | off | on | [primary {<host> <shared_secret> [<port>]}|delete]
| [secondary {<host> <shared_secret> [<port>]}|delete]
```

default	Return RADIUS authentication to its default off setting
off	Disable RADIUS authentication
on	Enable RADIUS authentication
<host>	IP address or DNS of the RADIUS server
<shared_secret>	Specify the designated secret (password)
<port>	To access the RADIUS server with a specific port, specify a port number. Otherwise, the default port will be used.
delete	Delete this RADIUS authentication server.

Example:

```
pc radius auth primary 172.21.8.55 secretpwd
```


pc radius interval

Adjust the RADIUS retry interval. By default, the RADIUS client waits five seconds before retrying a login when the RADIUS server fails to respond. You can select a value between 1 and 30 seconds. This command can only be issued by network administrators with touch-role access to the *PC* organization.

```
pc radius interval <seconds>|default
```

pc radius limit

Adjust the RADIUS retry limit. By default, if the RADIUS server fails to respond, the RADIUS client will try to log onto the server three times before reporting a server failure. You can select a value between 1 and 10. If you have specified a secondary authentication host, the RADIUS client will alternate attempts to log onto each server. This command can only be issued by network administrators with touch-role access to the *PC* organization.

```
pc radius limit <attempts>|default
```

pc radius method

Select the RADIUS authentication method: PAP, CHAP, or MSCHAP.

```
pc radius method pap|chap|mschap|default
```

PAP (<i>Password Authentication Protocol</i>)	With PAP, the user name and password are transmitted in clear, unencrypted text. ASCII or PAP authentication is required for RADIUS configurations that require access to clear text passwords (for example, when passwords are stored and maintained in a database external to the RADIUS server).
CHAP (<i>Challenge Handshake Authentication Protocol</i>)	In some environments, CHAP may be preferred for greater security. The RADIUS server sends a challenge that consists of a session ID and an arbitrary challenge string, and the user name and password are encrypted before they are sent back to the server. CHAP is the default authentication method.
MS-CHAP (<i>Microsoft Challenge Handshake Authentication Protocol</i>)	This protocol is similar to CHAP, but with MS-CHAP authentication, the RADIUS server can store an encrypted version of a user password to validate the challenge response. Standard CHAP authentication requires that the server stores unencrypted passwords. Note: MS-CHAP v1 and v2 are supported. PacketShaper attempts authentication with MS-CHAP v2 first. If the remote server doesn't support v2 or if authentication is denied, PacketShaper re-attempts authentication with MS-CHAP v1.
default	Return to default protocol.

pc radius show

View current settings for RADIUS authentication and accounting. This command can only be issued by network administrators with access to the *PC* organization.

```
pc radius show
```

Example output:

```
pc radius show
```

Setup values:

```
Radius Method      :CHAP
  Authentication    :off
  Accounting         :off
  Retry limit       :3
  Retry interval    :5
```

Service records:

Type	Host	Port	Secret
acct1	172.21.18.170	1813	secretpwd
acct2	radius.mycompany.com	1813	secretpwd

pc setup date

View or set the date and/or time for the PolicyCenter appliance. This command can only be issued by network administrators with touch-role access to the *PC* organization.

```
pc setup date [[yyyymmdd]hhmm[.ss]]
```

To define a timezone so PolicyCenter can change its local time automatically at the start and end of daylight savings time, use the command [pc setup timezone](#) .

pc setup disablehttp

Disable nonsecure HTTP access to the PolicyCenter management console and force all PolicyCenter users to log in via secure HTTPS only. Note that disabling HTTP access will immediately end all current HTTP sessions

```
pc setup disablehttp on|off|show
```

If you want to use the PolicyCenter file distribution server to distribute files to PacketShapers in shared mode, you must also configure the file distribution server for HTTPS (applicable to distributing files to legacy PacketShapers only).

pc setup dns

Configure one or more DNS servers for PolicyCenter to access.

```
pc setup dns none|<ipaddress> ...
```

Specify up to eight IP addresses, separating each with a space, or use **none** to clear previously set addresses.

pc setup domain

Define a default domain name that PolicyCenter can append to domain name lookups that are not fully qualified.

```
pc setup domain none|<domain_name>
```

pc setup images active

Select the image you want to run. After rebooting, this image will be the one running on the system.

```
pc setup images active <release_id>
```

Use the [pc setup images show](#) command to see the release ID number to specify. Use the [version](#) command to see the release ID of the current active image.

You will need to [reset](#) the PolicyCenter S-Series appliance to complete the activation process.



If connected to the serial console at boot up time, you can press the space bar and select the desired image to boot.

pc setup images add

Add an image to the PolicyCenter S-Series appliance. Use this command so that you can upgrade to a new release. PolicyCenter can manage up to six images. If your PolicyCenter S-Series appliance already has six images installed and you add another image, the oldest unlocked image will be replaced with the new image.

```
pc setup images add <URL to image> | file://<filename>
```

where:

<URL to image> is the path to an image on a web server that the PolicyCenter S-Series appliance has access to. Before initiating the **setup images add** command, you need to go the Symantec Support site, and download the new software image to a web server PolicyCenter can access.

file://<filename> is the filename of an image on the PolicyCenter S-Series appliance. Before initiating the **setup images add** command, you need to go to the Symantec Support site, download the new software image to a local workstation or server, and then upload it to the 9.258/image folder on the PolicyCenter S-Series appliance.

Examples:

```
pc setup images add http://webserver.mycompany.com/images/542386.bcs
```

```
pc setup images add file://542386.bcs
```

Additional Information

- When you add an image, the image manager assigns it a unique release ID (such as 532537); you must refer to this ID when [activating](#), [removing](#), [locking](#), and [unlocking](#) images.
- The **pc setup images add** command does not support https.

pc setup images detail

Display details about an image, such as the software version, whether the image is locked, and when the image was last booted.

```
pc setup images detail <release_id>
```

Use the [pc setup images show](#) command to see the release ID numbers.

Example:

```
#pc setup images detail 133755
```

```
-----  
-  
Get Image Detail: 133755  
-----  
-  
Release Id       : 133755  
DisplayName      : PolicyCenter 1.1.1.1, Release ID: 133755  
Product Name    : PolicyCenter  
Version         : 1.1.1.1  
Default Image   : YES  
Image Locked    : NO  
Last Boot Failed : NO  
Image Booted    : YES  
Boot Timestamp  : 2015-09-15T18:45:43
```

pc setup images lock

Lock an image to protect it from accidental deletion.

```
pc setup images lock <release_id>
```

Use the [pc setup images show](#) command to see the release ID numbers.

If you later decide you want to delete a locked image, you must use the [pc setup images unlock](#) command first.

pc setup images remove

Remove an image from the PolicyCenter S-Series appliance.

```
pc setup images remove <release_id>
```

Use the [pc setup images show](#) command to see the release ID numbers.

You cannot remove a locked image or the active image. If an image is locked, use the [pc setup images unlock](#) command before issuing the **pc setup images remove** command.

pc setup images show

Display a list of installed images and their version and release ID numbers. You will need to refer to the release ID when [activating](#), [removing](#), [locking](#), and [unlocking](#) images.

```
pc setup images show
```

Example:

```
pc setup images show
```

```
-----
--
List of Installed Images
-----
--
PolicyCenter 1.1.1.1      Release Id: 161344
PolicyCenter 1.1.1.1      Release Id: 165561
PolicyCenter 1.1.1.1      Release Id: 167309
PolicyCenter 1.1.1.1      Release Id: 169150
PolicyCenter 1.1.1.1      Release Id: 169795
PolicyCenter 1.1.1.1      Release Id: 170386  (A)

Total 6 images. (A) Active image
-----
--
```

The (A) marker indicates which image will be loaded when the PacketShaper is next rebooted. This may or may not be the currently running image; to see which image is currently active, use the [version](#) command.

pc setup images unlock

Unlock an image that you no longer want to protect from deletion. You have to unlock a locked image before you can remove it.

```
pc setup images unlock <release_id>
```

Use the [pc setup images show](#) command to see the release ID numbers.

pc setup message

Configure a message that will display before logging into PolicyCenter. The message displays before users login via the browser login page. This feature is useful for informing users about the company's access policies and consequences for unauthorized use. This command can only be issued by network administrators with touch-role access to the PC organization.

PolicyCenter CLI Commands

```
pc setup message {set <message>}|show|default
```

set <message>	Defines the message text. The text should be enclosed in quotation marks and can be up to 511 characters long.
show	Displays the content of the login message
default	Clears the message text

Examples

```
pc setup message set "Access to this system is restricted to authorized users only."
```

```
Message set to: "Access to this system is restricted to authorized users only..."
```

```
pc setup message show
```

```
Configured Message:
```

```
Access to this system is restricted to authorized users only.
```

Notes

- Quotation marks indicate the beginning and end of the login message. You cannot use a quotation mark within the body of the login message.
- To configure a login message for PacketShapers managed by PolicyCenter sharable configurations, see [setup message](#)

pc setup nic

Set the PolicyCenter S-Series's speed and duplex state.

```
pc setup nic <device> auto|(10bt|100bt) (full|half)|(1000b|10000b) full
```

where <device> is the interface name.

Specify **auto** (auto-negotiate) to automatically configure the unit for the appropriate mode. If you do not specify a state, it defaults to **auto**.

Additional Information

- Whenever you wish to change Network Interface Card (NIC) settings, always select auto-negotiate first, then select a different value if desired. Do not change from one non-auto setting to another non-auto setting directly; re-negotiation may fail.
- Although you can specify different fixed speeds on the Inside and Outside interfaces, such a configuration will result in a network interruption if the PolicyCenter is turned off because the end devices will not be able to negotiate the correct speed for the link.
- PolicyCenter does not support a 1000b or 10000b half-duplex interface.

pc setup show

Display the basic configuration for your PolicyCenter software. This command can only be issued by network administrators with access to the *PC* organization.

```
pc setup show
```

General Settings:

```
IP address:172.16.16.16 Subnet mask: 255.255.0.0
Gateway:172.16.16.1
DNS server(s):172.16.64.10
Default domain:mycompany.com
Date, time, timezone:Thu Dec 9 17:38:52 2006 PST (LosAngeles)
SNTP Client:off
SNTP Primary Server:time.nist.gov
SNTP Secondary Server:time-a.nist.gov
SNTP Poll Seconds:300
Syslog:off
```

RADIUS Setup values:

```
Radius Method      :CHAP
Authentication     :on
Accounting         :off
Retry limit        :3
Retry interval     :5
```

RADIUS Service records:

Type	Host	Port	Secret
auth1	server.mycompany.com	1812	mysecret

```
Directory Server password:myDSpassword
```

pc setup sntp

Set or display the Simple Network Time Protocol (SNTP) configuration for your PolicyCenter software. SNTP is used to synchronize the time in PolicyCenter to a server configured to propagate highly accurate time information through the Internet. This command can only be issued by network administrators with touch-role access to the *PC* organization.

```
setup sntp on|off|servers {<primary> [<secondary>]|none}|poll|reset|sync
```

To define a primary and secondary SNTP server, enter a standard dotted-decimal IP address for *<primary>* or *<secondary>*. To view current settings, issue the command [pc setup show](#).

pc setup ssl

Controls the strength of ciphers that PolicyCenter allows.

```
pc setup ssl strong|weak|show
```

When this option is set to **strong** (the default), PolicyCenter does not allow ciphers that don't have authentication or encryption, nor does it allow ciphers that don't have at least a 56-bit encryption key. When this option is set to **weak**, PolicyCenter allows ciphers of all strengths, as well as ciphers with no encryption or authentication.

pc setup timezone

When you configure a time zone, PolicyCenter can change its local time automatically at the start and end of daylight savings time. It also can retrieve time updates from time servers. This command can only be issued by network administrators with touch-role access to the *PC* organization.

```
pc setup timezone [<name>|custom <tz_spec>]
```

Each time zone has a unique name—usually the name of the best-known city in that zone. The default time zone is Los Angeles, CA. To display the valid time zones, use `setup timezone help`.

<tz_spec> is a string defined by POSIX.1 as:

```
<std><offset>[<dst>[<offset>],<date>[/<time>],<date>[/<time>]]
```

Where:

<std> and <dst>	3 or more characters specifying the standard and daylight saving time (DST) zone names
<offset>	[-]hh:[mm[:ss]] specifies the offset west of UTC. The default DST offset is one hour ahead of standard time
<date> [/ <time>]	Specifies the beginning and end of DST. If this is absent, the system applies US DST rules (first Sunday of April at 2:00 AM to last Sunday of October at 2:00 AM)
<time>	hh:[mm[:ss]] with a default of 02:00
<date>	One of the following forms: Jn (1<=n<=365): origin-1 day number, not counting February 29 n (0<=n<=365): origin-0 day number, counting February 29, if present Mm.n.d (0[Sunday]<=d<=6[Saturday], 1<=n<=5, 1<=m<=12): for the dth day of week n of month m of the year, where week 1 is the first week in which day d appears, and 5 stands for the last week in which day d appears (which may be either the 4th or 5th week)

For example, you could configure a time zone for Cairo, Egypt with the command:

```
pc setup timezone custom EET-2EEST,M4.5.5/01:00,M9.5.5/03:00
```

Current time zone:

Time zone name: Custom

Time zone desc: Custom time spec in POSIX format

Time zone spec: EET-2EEST,M4.5.5/01:00,M9.5.5/03:00

Time zone offset: GMT+02:00

DST offset: 60 minutes

DST starts: Last Friday of April at 01:00 AM

DST ends: Last Friday of September at 03:00 AM

In this example, the standard time, known as EET, is two hours ahead of GMT and daylight savings time, known as EEST, is the default 60 minutes ahead of EET. Rather than using US default rules, EEST begins on the last Friday of April at 1:00 AM and ends on the last Friday of September at 3:00 AM.

pc setup variable

Change a default variable setting for the PolicyCenter software configuration.

```
pc setup variable [<variable> <value>|default] | [-reset|-nd]
```

where *<variable>* is one of the variables listed below and *<value>* is the value you want to set the variable to. The default, minimum, and maximum values for each *<variable>* are listed in the table.

To reset all system variables to their defaults, use the **pc setup variable -reset** command. To reset a specific variable to its default, use the **pc setup variable <variable> default** command. To see a list of all variables that have non-default settings, use the **pc setup variable -nd** command.



Although additional variables are available, only the variables described below affect the PolicyCenter software configuration. The other variables have no effect on PolicyCenter.

Variable/ Description	Default Value	Min. Value	Max. Value
sessionTimeout Number of seconds before a user session times out after a period of inactivity. This variable applies to browser and console connections to PolicyCenter. After the session times out, the user will need to re-enter login credentials to continue PolicyCenter management.	180	5	86400
userSessionIdleTimeout This variable controls how many seconds it takes for an unauthenticated login session to get purged from the system. You might need to increase this value if the session times out before PolicyCenter can authenticate a login password, for example, when there is latency on the network or they are using a RADIUS or TACACS implementation. Note that this variable does not apply to idle sessions that have already been authenticated—just new sessions that have not yet been authenticated.	30	30	360

pc tacacs acct

Set up the configuration of the TACACS+ accounting service for use with PolicyCenter. This feature allows you to have an audit trail for user logins. This command can only be issued by network administrators with touch-role access to the *PC*

PolicyCenter CLI Commands

organization.

To define the TACACS+ accounting service to work with PolicyCenter, use:

```
pc tacacs acct primary|secondary {<host> <shared_secret> [<port>]}|delete|override
```

primary secondary	Enter the literal primary or secondary to indicate which server you are defining. (The secondary server is used when the primary server isn't accessible.)
<host>	The IP address or DNS name of the TACACS+ accounting server
<shared_secret>	The designated secret for the server; quotes are not required
[<port>]	The port number to access the server; if omitted, the default port 49 is used.
delete	Deletes the configuration of the primary or secondary server (whichever is specified)
override	This option is not supported by the pc tacacs command

To turn the service on or off, or to return the service to its default **off** value, use:

```
pc setup tacacs acct on|of|default
```

Example:

```
pc tacacs acct primary 10.10.10.10 P4assw0rd1
```

```
pc tacacs acct secondary 10.10.20.10 Paa55w0rd2
```

```
pc tacacs acct on
```

This example defines a primary accounting server at 10.10.10.10 which has a shared secret of P4ssw0rd, as well as a secondary server at 10.10.20.10. The third command line enables the TACACS+ accounting service. Once this service is configured and enabled, PolicyCenter will send a PW_STATUS_START accounting message to the accounting server when a user logs in and a PW_STATUS_STOP message when a user logs off or is disconnected.

pc tacacs auth

Set up the configuration of the TACACS+ authentication service so that users can log in to PolicyCenter using TACACS+ credentials. Using third-party TACACS+ servers enables you to have central configuration of user accounts.

```
pc tacacs auth primary|secondary {<host> <shared_secret> [<port>]}|delete
```

primary secondary	Enter the literal primary or secondary to indicate which server you are defining. (Note: The TACACS+ client uses the secondary server when the primary server isn't accessible or authentication failed.)
<host>	The IP address or DNS name of the TACACS+ authentication server
<shared_secret>	The designated secret for the server; quotes are not required

[<port>]	The port number to access the server; if omitted, the default port 49 is used
delete	Deletes the configuration of the primary or secondary server (whichever is specified)
override	This option is not supported by the pc tacacs command

To turn the service on or off, or to return the service to its default **off** value, use:

```
pc tacacs auth on|off|default
```

Example:

```
pc tacacs auth primary 10.10.10.10 CupServ44
```

```
pc tacacs auth on
```

This example first defines a primary authentication server at 10.10.10.10 which has a shared secret of CupServ44. The second command line enables TACACS+ authentication service. Once this is configured and enabled, PolicyCenter will prompt users for user name and password when they log in.

pc tacacs method

Select the TACACS+ authentication method for your PolicyCenter server:

- **ASCII** (*American Standard Code for Information Interchange*): With ASCII, the username and password are transmitted in clear, unencrypted text.
- **PAP** (*Password Authentication Protocol*). With PAP, the username and password are transmitted in clear, unencrypted text. ASCII or PAP authentication is required for TACACS+ configurations that require access to clear text passwords (for example, when passwords are stored and maintained in a database external to the TACACS+ server)
- **CHAP** (*Challenge Handshake Authentication Protocol*). In other environments, CHAP may be preferred for greater security. The TACACS server sends a challenge that consists of a session ID and an arbitrary challenge string, and the username and password are encrypted before they are sent back to the server.
- **MS-CHAP** (*Microsoft Challenge Handshake Authentication Protocol*): This protocol is very similar to CHAP, but with MS-CHAP authentication, the TACACS+ server can store an encrypted version of a user password to validate the challenge response. Standard CHAP authentication requires that the server stores unencrypted passwords.

```
pc tacacs method ascii|pap|chap|mschap|default
```

The default authentication method is **ascii**.

pc tacacs show

View current settings for TACACS+ authentication and accounting. This command can only be issued by network administrators with access to the *PC* organization.

```
pc tacacs show
```

PolicyCenter CLI Commands

Example output:

```
pc tacacs show
```

Setup values:

```
TACACS Method      :CHAP
Authentication      :off
Accounting           :off
Retry limit         :3
Retry interval      :5
```

Service records:

Type	Host	Port	Secret
acct1	172.21.18.170	1813	secretpwd
acct2	tacacs.mycompany.com	1813	secretpwd

pc tacacs timeout

Set the amount of time for TACACS+ to wait for a response from a server. By default, the TACACS+ client waits 20 seconds before retrying a login when the TACACS+ server fails to respond.

```
pc tacacs timeout <seconds>|default
```

where <seconds> is a value between 1 and 180 seconds. For example:

```
pc tacacs interval 20
```

In this example, the timeout interval is 25 seconds; this interval applies to any configured TACACS+ server.

To return to the default timeout interval, use:

```
pc tacacs timeout default
```

ping

Generate pings to test connectivity with another device on the network. If the device answers the pings from the PacketShaper, the message "x.x.x.x is alive" or "x packets transmitted, x packets received" will appear. If PacketShaper is unable to connect with the device, the message "no answer from x.x.x.x" or "0 packets received" will display.

```
ping <host> [<timeout>]
ping [-s] <host> [<count>]
```

<host>	IP address or DNS name
<timeout>	Number of seconds to transmit packets; if you don't specify a <timeout> value, PacketShaper will ping the host for up to 10 seconds

<code>[-s]</code>	Send a continuous ping
<code><count></code>	Number of pings to transmit; if you don't specify a <code><count></code> value, PacketShaper will ping the host 10 times

Examples of Successful Pings

```
PacketShaper# ping 172.21.1.26
ping (172.21.1.26): 56 data bytes
172.21.1.26 is alive
```

```
PacketShaper# ping 172.21.1.26 10
ping (172.21.1.26): 56 data bytes
172.21.1.26 is alive
```

```
PacketShaper# ping -s 172.21.1.26 5
ping (172.21.1.26): 56 data bytes
64 bytes from 172.21.1.26: icmp_seq=0
64 bytes from 172.21.1.26: icmp_seq=1
64 bytes from 172.21.1.26: icmp_seq=2
64 bytes from 172.21.1.26: icmp_seq=3
64 bytes from 172.21.1.26: icmp_seq=4
5 packets transmitted, 5 packets received
```

Examples of Unsuccessful Pings

```
PacketShaper# ping 192.168.0.1
ping (192.168.0.1): 56 data bytes
no answer from 192.168.0.1
```

```
PacketShaper# ping 192.168.0.1 30
ping (192.168.0.1): 56 data bytes
no answer from 192.168.0.1
```

```
PacketShaper# ping -s 192.168.0.1
ping (192.168.0.1): 56 data bytes
10 packets transmitted, 0 packets received
```

```
PacketShaper# ping -s 192.168.0.1 5
ping (192.168.0.1): 56 data bytes
5 packets transmitted, 0 packets received
```

plugin library

Show the current library of plug-in files available for distribution from PolicyCenter to individual PacketShapers. *This command is applicable to legacy PacketShaper configurations only.*

```
plugin library
```

The **plugin library** command shows the version name and type, version number and description for available plug-in files.

Example output of this command:

```
plugin library
```

PolicyCenter CLI Commands

Name	Type	Version	Description
ntpplug	bt03	1.0.0.0	Network News Transport Protocol
rogue	bt03	1.0.0.0	FileRogue - File Sharing Application
sms	bt03	1.0.0.0	Microsoft SMS pre Windows Service Pack 2

plugin prescribe

Prescribe plug-in files for a PolicyCenter configuration by filename. Use the **plugin library** command to determine the names of available files. *This command is applicable to legacy PacketShaper configurations only.*

```
plugin prescribe [<filename> <filename> ...] default|none|show
```

<filename>	The filename of the plug-in file you wish to prescribe to a PolicyCenter configuration.
default none show	<p>Specify default if the configuration should inherit its plug-ins from a parent configuration, or specify none if the configuration should not inherit its plug-ins. The show option shows the configuration's current plug-in files.</p> <p>Note: Issuing the plugin prescribe default command on a configuration with an inherited a plug-in prescription may incorrectly indicate that there are no inherited plug-ins. Use the command plugin prescribe show to correctly show all plug-ins prescribed for that configuration.</p>

plugin subscribe

Configure when and how often PacketShapers assigned to a PolicyCenter configuration update plug-in files. *This command is applicable to legacy PacketShaper configurations only.*

```
plugin subscribe asap|scheduled|default
```

The **plugin subscribe** command has the following options:

asap	PacketShapers assigned to the configuration will automatically update their plug-in files as soon as they are prescribed.
scheduled	PacketShapers assigned to the configuration will wait for the plugin sync command before downloading prescribed files.
default	If set to default , the PolicyCenter configuration inherits its plug-in subscription behavior from its parent configuration.

plugin sync

For units in shared mode only

Issue this command from an individual PacketShaper to immediately download plug-in files prescribed for the unit's PolicyCenter configuration. *This command is applicable to legacy PacketShaper configurations only.*

```
plugin sync <seconds>
```

The `<seconds>` variable allows you to specify in seconds the duration of the synchronization process. This command is only required when the PolicyCenter configuration prescription mode has been set to **scheduled** with the **plugin subscribe** command.



It is not necessary to issue this command if the prescription mode has been set to **asap** with the **plugin subscribe** command. If you issue the `plugin sync <seconds>` command from PolicyCenter with a specified number of seconds, PolicyCenter may display an error message that incorrectly states that there are no prescribed plug-in files for that configuration.

policy admit

Set the admission-control mechanism for a policy in the current configuration.

```
policy admit <tclass> squeeze|refuse|"<redirect-URL>" [nontcp|nonweb|web]
```

<code><tclass></code>	The traffic class whose policy is to be changed
<code>squeeze refuse "<redirect-URL>"</code>	This admission-control mechanism determines what happens when there isn't enough bandwidth to satisfy a guaranteed rate allocation. When the mechanism is squeeze , new connections will get at most 1024 bps. When the mechanism is refuse , the connection is refused. For web traffic <i>only</i> , when the mechanism is <code>"<redirect-URL>"</code> , the connection will be redirected to the specified URL.
<code>[nontcp nonweb web all]</code>	The traffic type

The policy admit command supports these combinations:

Admission Control Mechanism	Traffic Types
squeeze	nontcp, nonweb, web
tcp_refuse	nonweb tcp
http_refuse	web
http_redirect	web

policy apply discard

Toss all packets for a class in the current configuration.

```
policy apply discard <tclass>
```

policy apply ignore

Exempt a traffic class in the current configuration from bandwidth allocation and treat the traffic type as "pass-through" traffic.

```
policy apply ignore <tc<div data-bbox="63 200 937 262" data-label="Text">

By default, any traffic that you haven't explicitly classified is classified as Inbound/Default or Outbound/Default, and is factored into the bandwidth allocation scheme. When you apply an ignore policy to a traffic class, that traffic type will not be considered at all by the bandwidth allocation process. That is, it won't be counted as part of the virtual link traffic under management.


```

policy apply never-admit

Force admission control to occur on every use of a policy in the current configuration.

```
policy apply never-admit <tc<div data-bbox="63 384 937 447" data-label="Text">

The never-admit policy invokes the appropriate admission-control mechanism at the beginning of each session. For TCP and web traffic, use a never-admit policy to notify users that a service is unavailable. Admission-control mechanisms are configured using the policy admit command. For non-TCP traffic, use the policy apply discard command. For TCP non-web traffic, you can only use the policy admit refuse mechanism with a never-admit policy.


```

A never-admit policy must be applied to classes on the *requesting* flow. If a never-admit policy is applied to a class representing the response flow, PacketShaper responds as if the policy were a discard policy.

The never-admit policy has proven particularly effective in controlling certain viruses. This type of policy can also be used to redirect certain users to alternate URLs. For example, you might redirect a competitor to a URL that presents a customized message with a competitive analysis. The redirect option works only on the *response side* of the HTTP flow, not the request side.

policy apply priority

Apply a priority-based policy of a traffic class in the current configuration.

```
policy apply priority <tc<div data-bbox="75 679 904 742" data-label="Table">

|                                                                                                                                                                                                                                                                                |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <tc<div data-bbox="63 753 937 784" data-label="Text"> <p>Priority-based policies are used to establish a priority for traffic without specifying a particular rate. Use priority policies for traffic that does not burst, or whenever rate is not your primary objective.</p> |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|


```

policy apply rate

Apply a rate-based policy to a traffic class in the current configuration.


```
policy apply rate <tclass> <guar_lo_bps> <guar_hi_bps> [<priority>[[automatic|<excess_lo_bps><excess_hi_bps>] [<excess_limit_bps>]]]
```

<tclass>	The traffic class to which to apply the policy
<guar_lo_bps> <guar_hi_bps>	The guaranteed rate for this class' low- and high-speed connections (both parameters are required, even if you specify the same value). Rates may be specified as integer bits per second, followed by a K (thousands), M (millions), or G (billions).

For example, to guarantee 10k to Inbound/HTTP, use the following command:

```
policy apply rate inbound/http 10k 10k
```

To allow a policy to use excess rate, specify the following additional parameters:

<priority>	The excess rate priority for this traffic class, ranging from 0 (lowest) to 7 (highest)
[automatic]	Adjusts scaling automatically at run time
[<excess_lo_bps> <excess_hi_bps>]	The excess rate for this class' low- and high-speed connections (if you don't specify automatic). If you choose to use this option, both speeds must be specified. The minimum value allowed for <excess_lo_bps> is 1024.
[<excess_limit_bps>]	The maximum excess rate that can be used by this class

Guaranteed rate represents the minimum acceptable service level and thus the minimum acceptable rate to allocate. Low- and high-speed rate specifications are used to scale rate allocation to the user's access speed.

For example, to guarantee 10k to Inbound/HTTP burstable up to 48K at priority 3, use the following command:

```
policy apply rate inbound/http 10k 10k 3 automatic 38k
```



Excess rate is expressed differently in the CLI command than in the management console. In the management console, you specify 48k for the limit, but in the CLI you specify 38k for the amount of excess (the 48k limit minus the guaranteed rate of 10k).

To change the guaranteed rate later, use the [policy guaranteed](#) command. To adjust the excess rate, use the [policy excess](#) command.

policy default

Apply the PacketShaper-recommended policy to a traffic class in the current configuration.


```
policy default <tclass>
```

policy delaybound

Set the delay bound for a policy assigned to a traffic class in the current configuration to perform non-TCP rate control. PacketShaper uses a *UDP latency control mechanism* to rate-control individual UDP traffic flows and minimize packet loss. PacketShaper accumulates incoming UDP packets on a flow-by-flow basis when they are not scheduled for immediate transfer. With the UDP latency control mechanism, you define a *delay bound* how long the packets can remain buffered before they become too old to be useful. If UDP flows don't get sent immediately (because of link congestion, for example), they are placed in a buffer or *queue*. UDP flows stay in the queue until they are sent or until the delay bound time is exceeded, in which case the packets are dropped.

```
policy delaybound <tcldass> [<bound_in_millliseconds>]|default
```

<tcldass>	The traffic class whose policy is to be changed
[<bound_in_millliseconds>]	The new delay bound, from 1 to 10,000 milliseconds. The default delay bound is set to 200 milliseconds.

 Unless you have specific requirements for buffering non-TCP traffic, it is recommended that you do not change the delay bound size, as it has been optimized for most network environments.

policy dscp

Substitute a value into the Differentiated Services Code Point (DSCP) field in each packet for the specified traffic class in the current configuration. As defined in the Differentiated Services specification (RFC 2474), the DSCP field is the first six bits of the Type of Service (TOS) field in the IP header. This field is used by routers to make prioritized routing decisions.

```
policy dscp <tcldass> unchanged|<dscp>
```

Valid <dscp> values are 0-63, inclusive.

policy excess

Modify a rate-based policy's excess rate allocation for a traffic class in the current configuration.

```
policy excess <tcldass> <priority> [automatic|<lo_speed_bps> <hi_speed_bps>] [<excess_limit_bps>]
```

<tcldass>	The traffic class whose policy is to be changed
<priority>	The new highest priority for excess rate allocation
Optional rate allocations can be specified:	
automatic	Automatically scale the low-speed and high-speed rates

[<lo_speed_bps><hi_speed_bps>]	The new low- and high-speed rates. If you choose to use this option, both speeds must be specified.
[<excess_limit_bps>]	The maximum excess rate that can be used by this class

For example, the following command sets the excess rate limits for the FTP traffic class Inbound/Outside/ftp. It is assigned a priority of 4, and assigns both high- and low-speed users an excess rate of 50,000 bps with a total excess rate limit of 200,000 bps:

```
policy excess /inbound/outside/ftp 4 50k 50k 200k
```

policy failover

Configure a policy to react to failover mode, replacing the policy's guaranteed rate with a rate that is appropriate for the loss of a router link. Use this command if the unit has been configured to go into failover mode when it detects a problem with a site router link. *This command is applicable to legacy PacketShaper configurations only.*

```
policy failover <tclass> none|<speed_bps>
```

<tclass>	Traffic class with the policy that is to be changed
none	Remove the failover guaranteed rate from this class
<speed_bps>	Guaranteed rate to apply to the class when failover is active. Rates may be specified as integer bits per second, followed by a k (thousands), M (millions), or G (billions). The guaranteed rate must be a minimum of 256 bps (1024 bps on PacketShaper 7500, 10000, and 12000 models).

For example, the following commands set the guaranteed rate for the test class for normal link conditions. The **policy failover** command sets the guaranteed rate for the test class when the router link fails and a backup link with less bandwidth is used:

```
policy apply rate test 100k 100k 5 10k 10k
```

```
policy failover test 25k
```

policy flowlimit

Limit the rate of new flows to or from a unique host. This command can be used to detect and control a SYN Flood or similar denial-of-service attack directed at a particular host or if the attack is from a specific IP address. Flows exceeding the rate are blocked from passing through the unit. The limits are set to default values of 10,000 flows per minute on client hosts and 100,000 flows per minute on servers; depending on your network, you may need to change these defaults for effective control of SYN floods. Flow limits are automatically set on any classes that have a rate or priority policy assigned to them; if the **PolicyFlowLimitForAllClasses** [system variable](#) is enabled, PacketShaper will automatically block any flows that exceed these limits. (This variable is disabled by default. If you want to enforce flow limit policies, you need to enable the **PolicyFlowLimitForAllClasses** variable.)



You cannot set a flow limit on a class unless it already has a rate or priority policy assigned to it.

If you want to set or adjust the default limits on a particular class in the current configuration, use:

PolicyCenter CLI Commands

```
policy flowlimit <tclass> none|<client-fpm> <server-fpm>
```

<tclass>	Traffic class where the policy is located
none	Remove the flow limit
<client-fpm>	Maximum number of flows per minute to allow from each individual host; valid values are 0 600000
<server-fpm>	Maximum number of flows per minute to allow to each individual host; valid values are 0 600000

Note that the <client-fpm> and <server-fpm> rates include new flows of all types from an individual client or to an individual server (not just flows of the type of traffic matching this specific traffic class or policy).

PacketShaper offers measurement variables to track the number of flows that were blocked due to a server (flow destination) or a client (flow initiator) exceeding the flow limit rate specified in the **policy flowlimit** command: server-flood-block and client-flood-block.

If you don't want flow limits to be set automatically for newly created classes, enter the following commands:

```
policy flowlimit inbound/default none
policy flowlimit outbound/default none
```

policy guaranteed

Modify a rate policy's guaranteed rate allocations for a traffic class in the current configuration.

```
policy guaranteed <tclass> <lo_speed_bps> <hi_speed_bps>
```

<tclass>	Traffic class whose policy is to be changed
<lo_speed_bps> <hi_speed_bps>	New low-speed and high-speed guaranteed rates. Rates may be specified as integer bits per second, followed by a k (thousands), M (millions), or G (billions).

For example, the following command sets the low-speed and high-speed rates (10000 bps for low-speed users and 100000 bps for high-speed users) for a class named inbound/jup_202_http:

```
policy guaranteed inbound/jup_202_http 10000 100000
```

policy mpls-exp

Add or change the experimental (EXP) bits field of the MPLS (Multi-Protocol Label Switching) label on a packet. This field can be used in different ways for example, some routers use the EXP field to set class of service. *This command is applicable to legacy PacketShaper configurations only.*

```
policy mpls-exp <tclass> swap|delete <exp>
```

<code><tc></code>	Traffic class for which you want to modify the experimental bits field of the MPLS label
<code>swap <exp></code>	Marks the EXP field of an MPLS packet with the specified <code><exp></code> value (0 7)
<code>delete <exp></code>	Deletes the mpls-exp policy on the class

For example, to mark /outbound/http packets with an `<exp>` value of 7, use this command:

```
policy mpls-exp /outbound/http swap 7
```

To remove an mpls-exp policy that has an `<exp>` value of 7, use this command:

```
policy mpls-exp /outbound/http delete 7
```

The mpls-exp policy can be applied only to a class that already has a rate or priority policy. Note that the mpls-exp policy is applicable only if the packet has an MPLS label. If the packet doesn't have a label, the mpls-exp policy will simply be ignored. If packets don't already have MPLS labelling, you can use the [policy mplslabel](#) command to create an MPLS-tagging policy.

If more than one MPLS label exists in the stack, only the outermost packet's EXP field can be marked.

policy mplslabel

Add or change an MPLS (Multi-Protocol Label Switching) label on a packet. It can be applied only to a class that already has a rate or priority policy defined. *This command is applicable to legacy PacketShaper configurations only.*

```
policy mplslabel <tc> push|swap <mplslabel> | pop <times> | delete <operation>
```

<code><tc></code>	Traffic class for which you want to modify the MPLS label
<code>push <mplslabel></code>	Puts an MPLS label in a packet (and creates the MPLS stack if it doesn't exist); the <code><mplslabel></code> is the value of the label to be pushed (0 1048575)
<code>swap <mplslabel></code>	Swaps the topmost label of the MPLS stack with the specified <code><mplslabel></code> (0 1048575)
<code>pop <times></code>	Pops off the topmost label of the MPLS stack in the packet the specified number of <code><times></code>
<code>delete <operation></code>	Deletes the specified <code><operation></code> (pop, swap, or push) from the policy

Note that MPLS policies will work only on IP traffic.


A class can have a combination of push, swap, and pop operations in its MPLS policy; the pop operation can be specified multiple times (up to 8). If more than one operation type is specified for a given class, they are executed in the following order: pop, swap, push. For example, a class might have a policy that specifies a swap, three pops, and a push. In this case, the three pops occur first, then the swap, and then the push.

policy precedence

Substitute a precedence value for IP-based traffic classes in the current configuration.

```
policy precedence <tc> unchanged|<precedence>
```

<tc>	The traffic class for which you want to change precedence.
unchanged <precedence>	Use unchanged to turn off precedence substitution, restoring precedence to its default value. Or, enter a precedence value 0-7, where 7 is the highest priority.

 The **policy precedence** command supplements rate and priority policies—that is, a traffic class must have a policy already applied to it before you use the **policy precedence** command to substitute a precedence value.

policy remove

Remove a policy from a traffic class in the current configuration.

```
policy remove <tc>
```

policy route

Divert specific traffic to an alternative route by sending the class' traffic to a secondary gateway or router.

Set the MAC address routing for a traffic class.

```
policy route <tc> none|<macaddr>
```

PacketShaper substitutes the MAC address and transmits the packet accordingly.

policy show

Display policy information for a traffic class in the current configuration.

```
policy show <tc> [clear]
```

<tc>	Explicit traffic class name whose policy is to be displayed - for example, Inbound/Outside/http
[clear]	Reset the associated traffic class and policy hit counts

policy substitute

PacketShaper can detect the speed of a web connection at the first HTTP get request. You can use the **policy substitute** command to re-map the requested URL by substituting a URL that's more appropriate for the speed of the connection.

```
policy substitute <tclass> none
```

```
policy substitute <tclass> above|below <speed> "<pattern>" "<newpattern>"
```

<tclass>	The traffic class to which you are applying the policy
above below	Specify above or below a connection speed to indicate when the URL should be substituted.
<speed>	The connection speed that, in conjunction with above or below , triggers the content substitution
"<pattern>"	Specify in quotes the current URL pattern, which will be substituted with a new pattern. Wildcard patterns are not supported. This URL string is compared with the pattern in the /directory/file portion of a URL. PacketShaper ignores the http://computer-name portion of a URL when performing matching or substitution.
"<newpattern>"	The URL that you specify for substitution <i>must be the same length as the original URL</i> . The formatting rules are the same as those listed for the <pattern> parameter.

For example, to better serve a low-speed user, you could substitute a text-based web page for the regular home page:

```
policy substitute inbound/outside/web-in below 28800 "home-1.htm" "home-2.htm"
```

policy test

Test a policy on a traffic class in the current configuration to determine what rate will be allocated.

```
policy test <tclass> <rate_bps>
```

<tclass>	The traffic class whose policy is to be tested
<rate_bps>	The access speed to use to determine rate allocation

Example:

Assume the class inbound/http has the following policy settings: 10k guaranteed, burstable at priority 5, limit of 100k. To see how excess rate is allocated when there is 150 Kbps of demand, use the following command:

policy test inbound/http 150k

Policy Settings

Guaranteed rate lo 10k hi 10k

Excess rate default priority 5 CAP 90k

Allocation for flow at rate 150000

Guaranteed rate 10000

Excess rate at priority 5 -> 25088

Excess rate total 90000

Excess rate demand 0 0 15k 25k 25k 25k 0 0

PolicyCenter CLI Commands

This output shows how PacketShaper would allocate bandwidth when traffic class inbound/http generates 150 Kbps of demand. The top part of the display summarizes the policy settings. The excess rate (90k, next to CAP) is calculated by subtracting the guaranteed rate from the limit (100k-10k=90k).

The lower portion of the output lets you see how the excess rate is allocated between priority levels, 0-7. The sum of the rates allocated at each priority level equals the total excess rate (90k, in this example).

policy tos

Set a specific type of service (TOS) for IP traffic flows in a traffic class in the current configuration. It can be applied only to a class that already has a rate control policy defined.

```
policy tos <tc> unchanged|<tos>
```

<tc>	Explicit traffic class name for which you want to change the type of service
unchanged <tos>	<p>Use unchanged to turn off TOS substitution. Enter a <tos> value according to the following standard:</p> <ul style="list-style-type: none">8 = minimize delay4 = maximize throughput2 = maximize reliability1 = minimize monetary cost0 = normal service <p>Values can be combined to define broader results. For example, a value of 3 indicates "maximize reliability and minimize monetary cost."</p>

policy vlan

Add or change a VLAN identification (802.1Q) or priority (802.1p) on a packet in a class in the current configuration. It can be applied only to a class that already has a rate or priority policy defined.

VLAN Priority (802.1p)

To change the priority tag on an 802.1p class:

```
policy vlan type:8021p <tc> swap <priority>
```

<tc>	Traffic class for which you want to modify the VLAN priority tag
swap <priority>	Swaps the topmost priority level on the VLAN stack with the specified <priority>, 0 to 7.

For example, to change the VLAN priority to 6:

```
policy vlan type:8021p vlantestclass swap 6
```

VLAN Identification (802.1Q)

To modify the identification tag on an 802.1Q class:


```
policy vlan type:8021q <tclass> push|swap <vlanid> | pop <times> | delete <operation>
```

<tclass>	Traffic class for which you want to modify the VLAN ID
push <vlanid>	Puts an ID entry in a packet (and creates the stack if it doesn't exist); the <vlanid> is the value of the label to be pushed (0-4095)
swap <vlanid>	Swaps the topmost ID of the VLAN stack with the specified <vlanid> (0-4095)
pop <times>	Pops off the topmost label of the VLAN stack in the packet the specified number of <times>
delete <operation>	Deletes the specified <operation> (pop, swap, or push) from the policy

Examples:

```
policy vlan type:8021q testclass pop 2
policy vlan type:8021q testclass push 1
policy vlan type:8021q testclass swap 6
policy vlan type:8021q testclass delete pop
```

A class can have a combination of push, swap, and pop operations in its VLAN policy; the pop operation can be specified multiple times (up to 8). If more than one operation type is specified for a given class, they are executed in the following order: pop, swap, push. For example, a class might have a policy that specifies a swap, three pops, and a push. In this case, the three pops occur first, then the swap, and then the push.



A VLAN ID swap policy will automatically zero out the existing VLAN priority. To keep an existing non-zero priority value or to set a priority, be sure to specify a VLAN priority swap policy as well.

portal delete

Delete a customer portal account in the current configuration. *This command is applicable to legacy PacketShaper configurations only.*

```
portal delete <name>|all
```

portal library

Show the current portfolios of customer portal files available for distribution from PolicyCenter to individual PacketShapers. *This command is applicable to legacy PacketShaper configurations only.*

```
portal library [verbose]
```

The **portal library** command shows the name of the available portfolios only. Use **portal library verbose** to view the names of all the customer portfolio files within each portfolio.

portal modify

Modify customer account information in the current configuration. *This command is applicable to legacy PacketShaper configurations only.*

```
portal modify <name> <password> <directory> <message-of-the-day>
```

If RADIUS or TACACS+ authentication is enabled, passwords are not used (they are entered at the RADIUS/TACACS+ server). Thus, the syntax when RADIUS or TACACS+ is enabled is:

```
portal modify <name> <directory> <message-of-the-day>
```

Parameter	Description
<name>	The existing customer login name
<password>	The password for the customer account. If you are using RADIUS or TACACS+ authentication, you do not specify a password here—the customer portal will use the password specified for this user in the RADIUS/TACACS+ server.
<directory>	The new name of the customer's home directory (up to 8 characters); this directory will be created on the unit's data disk under 9.258/customer (optional)
<message-of-the-day>	The new custom message-of-the-day (optional)

If you don't specify the parameters, Packetshaper will prompt you for the information:

```
portal modify
```

Enter the name of the customer : **mycust**

Enter the new password :

Confirm new password :

Enter the new home directory name, 8 characters or less : **newdir**

Enter the new custom message-of-the-day (optional) : **All network resources online**

Customer mycust was modified



You will not be prompted for a password if RADIUS authentication is enabled.

After this is executed, mycust's home directory will be 9.258/customer/newdir.



You must explicitly type each entry when you use prompted mode. If, for example, you press Enter at the password prompt, the new password value becomes (none).

portal new

Create a new customer portal account in the current configuration. *This command is applicable to legacy PacketShaper configurations only.*

```
portal new <name> <password> <directory> <message-of-the-day>
```

If RADIUS or TACACS+ authentication is enabled, passwords are not used (they are entered at the RADIUS/TACACS+ server). Thus, the syntax when RADIUS or TACACS+ is enabled is:

```
portal new <name> <directory> <message-of-the-day>
```

Parameter	Description
<name>	The login name the customer will use; up to 32 characters long, use numbers, letters and underscores— spaces are not allowed. If you are using RADIUS or TACACS+ authentication, this name must match the user name entered in the RADIUS/TACACS+ server. Note: If the directory name is not specified, then the login name is used for the directory name. In this case, the login name is limited to 8 characters because the directory name is limited to 8 characters.
<password>	The password for the customer account. If you are using RADIUS or TACACS+ authentication, you do not specify a password here— the customer portal will use the password specified for this user in the RADIUS/TACACS+ server.
<directory>	The name of the customer's home directory (up to 8 characters); this directory will be created on the unit's data disk under 9.258/customer (optional)
<message-of-the-day>	A text string of 128 characters or less, intended to carry simple messages such as <i>System will be down from 5:00 am to 6:00 am tomorrow</i> (optional)

You must use empty quotes ("") if you don't want to enter a value for a parameter. For example, to create a user MyCust with a directory named cust01 (no password, no message of the day), use:

```
portal new MyCust "" cust01 ""
```

If you don't specify any parameters with the **portal new** command, PolicyCenter will prompt you for the values.

This is an example of prompted mode:

```
portal new
```

Enter the customer login name, password, home directory name (8 characters or less) and an optional custom message-of-the-day (128 characters or less).

Enter the customer's login name, e.g. 'marysmith' : **mycust**

Enter the password :

Confirm the password :

Enter the customer's home directory name, e.g. 8 characters or less : **mycust**


Enter a custom message-of-the-day (optional): **No network outages**

Customer mycust was added.



You will not be prompted for a password if RADIUS authentication is enabled.

After this is executed, a directory 9.258/customer/mycust will exist. The service provider must FTP an INDEX.HTM file to it before the mycust customer can use it effectively.

 You must explicitly type each entry when you use prompted mode. If, for example, you press Enter at the password prompt, the new password value becomes (none).

portal prescribe

Prescribe a group of customer portal files by portfolio name. Use the **portal library** command to determine available customer portal portfolios. *This command is applicable to legacy PacketShaper configurations only.*

```
portal prescribe <portfolio> default|none|show
```

<portfolio>	A portfolio is any sub-folder of <i>publish/portal</i> that contains a group of portal files.
default none show	Specify default if the configuration should inherit its portfolio of customer portal files from a parent configuration, or specify none if the configuration should not inherit its portfolio. The show option shows the configuration's current prescribed portfolio of customer portal files.

portal show

Display the current customer portal settings in the current configuration. *This command is applicable to legacy PacketShaper configurations only.*

```
portal show
```

Customer Name	Password	Directory	Message

Farnsworths	*****	books	Inventory starts Friday!
Sigma_Air	*****	air	No scheduled network outages

The Password column does not appear if RADIUS authentication is enabled.

portal subscribe

Configure when and how often PacketShapers assigned to a PolicyCenter configuration update their portfolio of customer portal files. *This command is applicable to legacy PacketShaper configurations only.*

```
portal subscribe asap|scheduled|default
```

The **portal subscribe** command has the following options:

<i>asap</i>	PacketShapers assigned to the configuration will automatically update their customer portal portfolio as soon as it is prescribed.
<i>scheduled</i>	PacketShapers assigned to the configuration will wait for the portal sync command before downloading the prescribed portfolio of files.
<i>default</i>	If set to default , the PolicyCenter configuration inherits its portal subscription behavior from its parent configuration.

portal sync

For units in shared mode only

Issue this command from an individual PacketShaper to immediately download customer portal files prescribed for the unit's PolicyCenter configuration. This command is only required when the PolicyCenter configuration prescription mode has been set to **scheduled** with the **portal subscribe** command. *This command is applicable to legacy PacketShaper configurations only.*



It is not necessary to issue this command if the prescription mode is currently in its default state, or has been set to **asap** with the **portal subscribe** command.

```
portal sync
```

pwd

Show the working directory.

```
pwd
```

radius acct

Test and debug the setup of your RADIUS accounting server in the current configuration. This command sends test accounting messages to the server.

```
radius acct start|stop|on|off
```

Specify **start** to send a test message that tells the accounting server that someone logged in and **stop** to send a log-off test message. The administrator can then verify that these messages are in the accounting server log. They will appear in the log under the name *RadiusAccountingTestUser*.

You can use the **on** and **off** parameters to send a message to the server that the RADIUS accounting service is on or off. Note that this command does not affect the setup of the accounting service; if the service was enabled with the **setup radius acct on** command, it will remain enabled (even if you used the **radius acct off** to send a test message that the accounting service is off).

radius chapcookie

Generate and display a single-use CHAP cookie in the current configuration.

```
radius chapcookie
```

radius chaplogin

Send a test CHAP login request to the RADIUS authentication server in the current configuration. This command is useful for testing and debugging the setup of your RADIUS authentication server, when Challenge-Handshake Authentication Protocol (CHAP) is used.

```
radius chaplogin <username> <password>
```

PolicyCenter CLI Commands

For example:

```
radius chaplogin bob 12567
```

chap ID = 0x1

challenge = 37a9aa04189c7ac5c826fde6a52c988f

password = 12567

response = 7610c93540dc90422fb4b077d23dd63a

"bob" RADIUS Authentication OK

Vendor-Specific: access=touch

The above output indicates that the authentication of the user Bob was successful. If authentication fails, you will see one of the following messages:

Message	What it means	What you should do
Authentication turned off	You need to enable the authentication service in the configuration.	Use the setup radius auth on command to enable authentication, and then send another test login request.
No server configured	The RADIUS authentication service is turned on in the configuration, but the server is not configured.	Use the setup radius auth primary command and specify the authentication server's IP address, port number, and shared secret. Then send another test login request.
Access rejected by server	The user name and/or password is invalid.	Contact your RADIUS administrator to verify that you are using the correct user name and password.

Message	What it means	What you should do
Timeout: Unable to obtain a response from server	The RADIUS authentication service is turned on in the configuration and the server is configured. This message could be caused by any of the following situations:	
	Server could be down.	Contact your RADIUS administrator to check the status of the RADIUS authentication server. It's a good idea to configure a secondary server to have a backup in case the primary server fails.
	Incorrect IP address for the server.	Contact your RADIUS administrator to verify the host name or IP address of the authentication server.
	The authentication service may not be enabled on the RADIUS server side.	Contact your RADIUS administrator to verify that the authentication service is enabled on the RADIUS server.
	The server may not be configured to work as a PacketShaper client.	For information on configuring the RADIUS server with PacketShaper-specific attributes, see Configure RADIUS Servers in PacketGuide .
	The LAN may be busy or down.	Check the status of the network.

radius clear

Clear the accounting drop count and remove the drop-notice banner in the current configuration. When an accounting request is dropped because the accounting server was not configured correctly or was unreachable for some reason, PacketShaper keeps track of these dropped accounting requests and displays a banner alerting you that requests have been dropped. You can use the **radius clear** command to clear this banner.

```
radius clear
```

radius login

Send a test PAP login request to the RADIUS authentication server in the current configuration. This command is useful for testing and debugging the setup of your RADIUS authentication server, when Password Authentication Protocol (PAP) is used.

```
radius login <username> <password>
```

For example:

```
radius login bob 12567
```

PolicyCenter CLI Commands

"bob" RADIUS Authentication OK
Vendor-Specific: access=touch

The above output indicates that the authentication of the user Bob was successful. If authentication fails, you will see one of the following messages:

Message	What it means	What you should do
Authentication turned off	You need to enable the authentication service in the configuration.	Use the setup radius auth on command to enable authentication, and then send another test login request.
No server configured	The RADIUS authentication service is turned on in the configuration, but the server is not configured.	Use the setup radius auth primary command and specify the authentication server's IP address, port number, and shared secret. Then send another test login request.
Access rejected by server	The user name and/or password is invalid.	Contact your RADIUS administrator to verify that you are using the correct user name and password.
Timeout: Unable to obtain a response from server	The RADIUS authentication service is turned on in the configuration and the server is configured. This message could be caused by any of the following situations:	
	Server could be down.	Contact your RADIUS administrator to check the status of the RADIUS authentication server. It's a good idea to configure a secondary server to have a backup in case the primary server fails.
	Incorrect IP address for the server.	Contact your RADIUS administrator to verify the host name or IP address of the authentication server.
	The authentication service may not be enabled on the RADIUS server side.	Contact your RADIUS administrator to verify that the authentication service is enabled on the RADIUS server.
	The server may not be configured to work as a PacketShaper client.	For information on configuring the RADIUS server with PacketShaper-specific attributes, see Configure RADIUS Servers in PacketGuide .

Message	What it means	What you should do
	The LAN may be busy or down.	Check the status of the network.
Error: Reply didn't contain an access level attribute	The user name and password are valid, but the user wasn't configured with an access level attribute.	Configure the RADIUS server with an access level attribute for this user.

radius mschaplogin

Send a test MSCHAP login request to the RADIUS authentication server in the current configuration. This command is useful for testing and debugging the setup of your RADIUS authentication server, when Microsoft Challenge-Handshake Authentication Protocol (MSCHAP) is used.

```
radius mschaplogin <username> <password>
```

For example:

```
radius mschaplogin bob 12567
```

```
chap ID = 0x1
challenge = 37a9aa04189c7ac5c826fde6a52c988f
password = 12567
response = 7610c93540dc90422fb4b077d23dd63a
"bob" RADIUS Authentication OK
Vendor-Specific: access=touch
```

The above output indicates that the authentication of the user Bob was successful. If authentication fails, you will see one of the following messages:

Message	What it means	What you should do
Authentication turned off	You need to enable the authentication service in the configuration.	Use the setup radius auth on command to enable authentication, and then send another test login request.
No server configured	The RADIUS authentication service is turned on in the configuration, but the server is not configured.	Use the setup radius auth primary command and specify the authentication server's IP address, port number, and shared secret. Then send another test login request.
Access rejected by server	The user name and/or password is invalid.	Contact your RADIUS administrator to verify that you are using the correct user name and password.

Message	What it means	What you should do
Timeout: Unable to obtain a response from server	The RADIUS authentication service is turned on in the configuration and the server is configured. This message could be caused by any of the following situations:	
	Server could be down.	Contact your RADIUS administrator to check the status of the RADIUS authentication server. It's a good idea to configure a secondary server to have a backup in case the primary server fails.
	Incorrect IP address for the server.	Contact your RADIUS administrator to verify the host name or IP address of the authentication server.
	The authentication service may not be enabled on the RADIUS server side.	Contact your RADIUS administrator to verify that the authentication service is enabled on the RADIUS server.
	The server may not be configured to work as a PacketShaper client.	For information on configuring the RADIUS server with PacketShaper-specific attributes, see Configure RADIUS Servers in PacketGuide.
	The LAN may be busy or down.	Check the status of the network.

radius sessions

Show a list of current user sessions (RADIUS, TACACS, ds, local) and detailed information about each session in the current configuration.

```
radius sessions
```

ID	Status	Age	Idle	Limit	Type	Access	User Name
3b61a571	logged in	40 mins	0 secs	60 mins	CLI	touch	bob
3b61afb6	logged in	91 secs	91 secs	60 mins	WUI	touch	john
3b61af6a	logged in	167 secs	17 secs	60 mins	WUI	touch	george
3b61a571	logged in	45 mins	0 secs	60 mins	CLI	touch	bob

Column	Description
ID	Identification given to the user session
Stat	The status of the session: logged in– the session is active timed out– the session timed out logged off– the session is inactive
Age	Length of time the session has been active– that is, the amount of time since the user logged in
Idle	Amount of time since the user gave a command; whenever a user gives a command, the idle value is reset to zero
Limit	Amount of time a session is idle before the user will be timed out and logged off; for example, if the limit is 60 minutes, a user will get logged off when no commands are given for a 60-minute period. Note: The PacketShaper default session life limit is 60 minutes. However, the RADIUS or TACACS+ server can be independently configured with different limits for different users and these limits override PacketShaper's.
Type	Type of interface used: CLI (command-line interface), WUI (web user interface)
Access	Type of access: Touch, Look
User Name	Name of the user who logged into the session

radius show

Display RADIUS client settings in the current configuration. Use this command to verify that RADIUS authentication and accounting are enabled, to see the current settings for the retry limit and retry interval, and to view the configuration settings for the primary and secondary authentication and accounting servers.

```
radius show
```

```
Radius method is CHAP
Radius Authentication is ON
Radius Accounting is OFF
Retry limit: 3
Retry interval: 5
```

	auth1	auth2	acct1	acct2
Server	172.23.225.203	172.23.225.213	–	–
Secret	packet	packet	–	–
Port	1812	1812	0	0
Status	Up	Unknown	Unknown	Unknown

PolicyCenter CLI Commands

Attempts	1	0	0	0
Success	1	0	0	0
Timeout	0	0	0	0

Auth1 last accessed: Wed Jul 11 14:16:48 2007

Auth2 was never accessed!

The output also indicates the number of attempts made to connect to each server, the number of successful connections, and the number of connections that timed out.

reset

Reset the PolicyCenter software or PolicyCenter S-Series appliance. A reset performs a proper reboot, without losing power to the system.

```
reset [soft]
```

To reboot the appliance, use **reset**. To reset just the software, use **reset soft**.

To reset configurations, use the [setup reset](#) command.

rtm accept

Set an acceptable service level threshold percentage for response time measurement (RTM) in the current configuration. The default is 100%.

```
rtm accept <tc> <percent>
```

where *<tc>* is the traffic class to be defined.

rtm hosts

Enable or disable worst client/server tracking for a class in the current configuration.

```
rtm hosts <tc> enable|disable
```

where *<tc>* is the traffic class to be tracked.

rtm show

Display a summary of the RTM statistics for all traffic classes with response-time data in the current configuration.

```
rtm show
```

The display has one row per traffic class with the following information:

Traffic Class	The name of the traffic class
Goodness	The number of good transactions (those within the Total Delay Threshold) divided by the transaction count, multiplied by 100. In other words, the percentage of good transactions.
Response Time: Total, Network, Server, Normal	<p>The average number of milliseconds required by the class' transactions.</p> <p>The value in the Normal column is the component of the transaction time that is directly related to the transaction size. An increase or decrease in that number does not indicate any change in network or server performance and requires no user intervention. This value is not tracked by the measurement engine.</p>

rtm threshold

Differentiate between acceptable and unacceptable response by supplying a threshold that defines good performance for a traffic class in the current configuration. PacketShaper uses the threshold when evaluating each transaction's total delay figure. If the transaction completes within the time indicated with the threshold, the transaction is considered "good."

```
rtm threshold <tclass> <delay>|none [total|network|server]
```

Specify the delay threshold in milliseconds or remove the threshold using the **none** literal. The threshold maximum is 99 seconds.

If you set a network or server delay threshold in the CLI, the setting will not appear in the management console, as these types of thresholds are not supported in the browser version of PacketShaper. (Only total delay threshold can be set in the management console.) Note that using the management console to make any changes to a class' RTM settings will clear the network or server delay you set in the CLI.

run

Run a command file or script. This command runs in the context of your current folder. The filename must have a **cmd** suffix. An output file, <filename>.out, contains the results of command-file execution. To view the contents of this output file, use the **more** command.



Diagnostic commands (dns, ping, net, sys, uptime) and utility commands (cd, date, history, ls, more, pwd) cannot be executed from a command file or used with the **schedule** command.

schedule delete

Delete a scheduled command execution in the current configuration.

```
schedule delete <item_id>|all
```

The scheduled item is removed from the list, but the remaining items have the same IDs.

This command produces output similar to the following:

PolicyCenter CLI Commands

Id	Time Range	Issued	Date	Command
00000001:	08:00:00-08:00:40	1	weekday	"p2pday.cmd"*
2DA2B03E:	18:00:00-18:00:40	1	weekday	"p2peve.cmd"*
5545F94E:	06:00:00-06:00:40	0	daily	"test.cmd"*

schedule disable

Disable a scheduled item in the current configuration so that it won't be executed. If you want to permanently remove the scheduled item, use the [schedule delete](#) command.

```
schedule disable <item_id> | all
```

where <item_id> is the ID of the item displayed in the [schedule show](#) output. Use the **all** parameter to disable all entries.

After you have disabled an item, it will still be listed in the **schedule show** output, but [disabled] will appear at the end of the line.

Id	Time Range	Issued	Date	Command
00000001:	08:00:00-08:00:40	1	weekday	"p2pday.cmd"*
2DA2B03E:	18:00:00-18:00:40	1	weekday	"p2peve.cmd"*
7EC07402:	06:00:00-06:00:40	1	weekend	"p2psat.cmd"*
5545F94E:	06:00:00-06:00:40	0	daily	"test.cmd"*

To enable the item later, use the [schedule enable](#) command.

schedule enable

Enable a scheduled item in the current configuration that has been disabled with the [schedule disable](#) command.

```
schedule enable <item_id> | all
```

where <item_id> is the ID of the item displayed in the [schedule show](#) output. Use the **all** parameter to enable all entries.

schedule new

Schedule a command in the current configuration to execute at a specific time and date. When using the scheduling feature, it is important that your unit has the correct date and time. Use the [date](#) command to check the date and time. If you need to correct the date or time, use the setup date command.

Scheduling is limited to 64 scheduled commands. Scheduled commands that are no longer needed (for example, expired commands scheduled to run only once) should be removed from the list via the [schedule delete](#) command, so they do not continue to consume available resources.

Note for units in shared mode: To permanently delete a command scheduled via PolicyCenter, you must remove the command from the unit's sharable PolicyCenter configuration, and not just from the unit itself. If PolicyCenter detects that a scheduled command is on the unit's PolicyCenter configuration but no longer on the unit itself, PolicyCenter will synchronize the unit's settings with its PolicyCenter configuration to restore that command to the unit. If the time range for the scheduled command is not over or the PolicyCenter time zone is not correctly configured, the command may run again.

When setting the time for a scheduled command, keep biannual time changes into consideration. For example, if you set a command to execute at 2:30am in the United States, the command will not be executed when the clock changes ahead one hour for Daylight Saving Time. You can ensure a command will be executed during a clock change by specifying a time range (such as 02:00-04:00).

```
schedule new [<day_option>] <time_range>[utc] <cmd>|{-f <cmd_file>} [mail:<address>]  
[id:<item_id>] [disable]
```

PolicyCenter CLI Commands


<p>[<day option>]</p>	<p>Specifies the day(s) the schedule should run. If you don't specify the <day option>, the scheduled item will run every day.</p> <p>Specify one of the following for <day option>:</p> <p>now</p> <p>today + <n></p> <p>[once:]<date>[, <date>]</p> <p>[once:]weekday weekend <dow>[, <dow>]</p> <p>[once:]<dom>[, <dom>]</p> <p>+<n> is <n> days from today. For example, +1 is tomorrow.</p> <p><date> is a specific date in the format <i>mm/dd</i>. The date is assumed to be a future date, within the next twelve months. For example, if today is 5/30/02 and you specify the date 5/29, the item will be scheduled for execution on May 29, 2003. You can specify up to 10 dates separated by commas.</p> <p>The once: option, that optionally precedes the <date> and the following options, specifies that the item should be executed once for each of the specified dates. If you don't specify once:, the item will be executed on an ongoing basis, according to the date(s) you specified.</p> <p>weekday executes the item on weekdays only (Monday through Friday). weekend executes the item on weekends (Saturday and Sunday). These two options are useful for setting different policies for weekdays and weekends. For example, you might want music file sharing to have less bandwidth during the week than on the weekend.</p> <p><dow> is the day of the week, specified with the first three letters of the day (mon, tue, wed, thu, fri, sat, sun). If you specify more than one day, each day is separated by a comma and no space, for example, <i>mon,wed</i>. You can specify up to seven days of the week.</p> <p><dom> is a specific day of the month, for example, <i>15</i> for the fifteenth of the month. You can specify up to 31 days, separated by commas.</p>
-----------------------------	---

<code><time_range></code> <code>[utc]</code>	<p>Specifies the time at which the command or command file should be executed. The syntax is:</p> <pre>hh:mm[.ss] [-hh:mm[.ss]]</pre> <p>where <i>hh</i> is the hour from 0 to 23, and <i>mm</i> and <i>ss</i> are minutes and seconds from 0 to 59. For example, to specify the time 5pm, enter <i>17:00</i>.</p> <p>If a range is not specified, PacketShaper will attempt to execute the command within a 40-second window. If you want to allow more time for the command(s) to be executed, you can specify a range (for example, 08:00-08:02).</p> <p>The legacy syntax <i>hhmm.ss</i> is supported for backward compatibility.</p> <p>Use the optional suffix utc to specify a coordinated universal time (UTC). Specifying times in UTC (similar to Greenwich mean time) is useful when managing units in different time zones. For example, 1800Z is 1pm in Eastern standard time and 4pm in Pacific standard time.</p>
<code><cmd> </code> <code>-f <cmd_file></code>	<p>You can specify one of the following:</p> <ul style="list-style-type: none"> • The CLI command <code><cmd></code> to be executed. The command should be enclosed in quotation marks. • The name of the file (<code>-f <cmd_file></code>) that contains a list of CLI commands. Specify a path to the <code><cmd_file></code> unless the file is in the default folder (9.256/cmd). The filename should have a .CMD extension.
<code>[mail:</code> <code><address></code> <code>]</code>	<p>Sends the output of the command or command file execution to the specified email address(es), allowing you to confirm that the command executed at the specified date and time. You can specify up to four email addresses, separated by commas.</p> <p>In order to use this feature, you must configure a mail server. See setup email.</p>
<code>[id: <item_id>]</code>	<p>Assigns the specified ID to the scheduled item. <code><item_id></code> can be up to eight characters long and can contain the numbers 0-9 and the letters A-F and a-f. This parameter is primarily used to override inherited entries when using shared mode (PolicyCenter). The ID is shown in the list of scheduled entries via the schedule show command. If you don't specify an ID, a random number is assigned.</p> <p>Note: It's recommended that you allow PolicyCenter to automatically create the ID rather than manually assign the ID with the <code>id</code> option. If you do manually assign an ID, make sure you follow the guidelines for ID names, as described above.</p>
<code>[disable]</code>	<p>Disables the scheduled item so that it won't be executed. If you want to enable or disable the item after it is created, use the schedule enable or schedule disable command.</p>

You will typically want to create scheduled items in complementary pairs. For example, you can create one scheduled item for a policy that is applicable during work hours and another schedule for a policy that is applicable after hours.

```
schedule new weekday 08:00 policy apply rate /inbound/gnutella 4800 9600 2 4800 4800
schedule new weekday 18:00 policy apply rate /inbound/gnutella 128k 256k 4 256k 256k
```

If you use the `mail: <address>` parameter, an email message containing the command output will be sent to the specified `<address>` shortly after the schedule is executed.

 Diagnostic commands (dns, ping, net, sys, uptime) and utility commands (cd, date, history, ls, more, pwd) cannot be executed from a command file or used with the **schedule** command.

schedule reset-counter

Zero out the counters for all command schedules in the current configuration. The counters keep track of the number of times each schedule has been executed since it was created. Use this command to reset the counters to zero (for example, after resetting the appliance).

```
schedule reset-counter
```

Use the [schedule show](#) command to see the values of the counters for all configured schedules; the Issued column displays the counter values.

schedule show

List the currently scheduled commands in the current configuration.

```
schedule show [-time] [-utc]
```

[-time]	Sorts schedules by time, with the earliest start time listed first. Without the -time switch, schedules are sorted by ID.
[-utc]	Lists schedules in their original time input. If time was entered in coordinated universal time (UTC) format, the UTC time will be displayed with a Z after the time (for example, 00:00:00-00:00:40Z). If the time was specified in local time, the local time will be displayed. Without the -utc switch, all times are listed in local time; in other words, any UTC times are converted to local time on the display. An L displays after the converted UTC time (for example, 08:00:00-08:00:40L).

This command produces output similar to the following:

Id	Time Range	Issued	Date	Command
00000001:	08:00:00-08:00:40	1	weekday	"p2pday.cmd"*
2DA2B03E:	18:00:00-18:00:40	1	weekday	"p2peve.cmd"*
7EC07402:	06:00:00-06:00:40	1	weekend	"p2psat.cmd"*
5545F94E:	06:00:00-06:00:40	0	daily	"test.cmd"*

Each scheduled event has a unique ID, which can be used to delete items from a schedule. The number in the "Issued" column indicates how many times the command has executed. An asterisk (*) flags command-file items. If [disabled] appears, the item has been disabled with the [schedule disable](#) command and will not be executed.

setup access default

When issued from the command-line interface of an individual PacketShaper unit, this command returns the security access settings for the unit to its default value, allowing users to access the unit via all available secure and non-secure protocols.

```
setup access default
```

If this command is issued from PolicyCenter for a *child configuration*, the selected configuration will discard its existing security access settings and inherit all security access settings from its parent configuration. If the parent configuration has one or more access protocols disabled, the child configuration will disable those protocols as well.

If this command is issued from PolicyCenter for a *root-level configuration*, the selected configuration will return its security access settings to its default value, enabling all available secure and non-secure protocols for accessing the unit.



To enable or disable a single protocol for accessing the unit, use instead the commands [setup access enable](#) or [setup access disable](#).

setup access disable

Allow access to a PacketShaper's browser and command-line interfaces (CLI) only via specific access protocols or with a secure connection. By default, the secure access protocols (HTTPS, SSH) are enabled and the non-secure protocols (HTTP, SNMP) are disabled.



Changing this setting might cause the session to be dropped.

For S-Series configurations:

```
setup access disable all|[https|ssh|http|snmp]
```

For legacy configurations:

```
setup access disable all|[https|ssh|http|snmp|ftp|telnet|echo]
```

where:

all	Enable all secure and nonsecure protocols for accessing your unit.
https	HTTP over Secure Sockets Layer protocol (SSL). The management console uses the SSL protocol to securely access the unit.
ssh	Secure Shell remote login protocol (SSH). The CLI uses SSH to securely access the unit.
http	Hypertext Transport Protocol. This is a non-secure protocol.
snmp	Simple Network Time Protocol. This is a non-secure protocol.
ftp	File Transfer Protocol. This is a non-secure protocol available for legacy configurations only.
telnet	Network termination protocol. This is a non-secure protocol available for legacy configurations only.
echo	Echo protocol. This is a non-secure protocol available for legacy configurations only.

setup access enable

Issue this command to re-enable access to the PacketShaper unit via a service protocol that was previously disabled with the [setup access disable](#) command. By default, all services are enabled, allowing you to access the unit by all available secure and non-secure protocols.

For S-Series configurations:

```
setup access enable all | [https|ssh|http|snmp
```

For legacy configurations:

```
setup access enable all | [https|ssh|http|snmp|ftp|telnet|echo
```

Where:

all	Enable all secure and nonsecure protocols for accessing your unit.
https	HTTP over Secure Sockets Layer protocol (SSL). The management console uses the SSL protocol to securely access the unit.
ssh	Secure Shell remote login protocol (SSH). The CLI uses SSH to securely access the unit.
http	Hypertext Transport Protocol. This is a non-secure protocol.
snmp	Simple Network Time Protocol. This is a non-secure protocol.
ftp	File Transfer Protocol. This is a non-secure protocol available for legacy configurations only.
telnet	Network termination protocol. This is a non-secure protocol available for legacy configurations only.
echo	Echo protocol. This is a non-secure protocol available for legacy configurations only.

setup access show

Display access security settings for the current configuration. The output lists all service protocols available for accessing the unit, and indicates whether each protocol is enabled or disabled.

```
setup access show
```

To disable or re-enable a protocol for accessing the unit, use the command [setup access disable](#) or [setup access enable](#).

For example:

```
Unit Access:
Access Method      Disabled
-----
HTTPS              No
SSH                No
HTTP               Yes
SNMP               No
```

setup adaptiveresponse

Turns *all* configured adaptive response agents in the current configuration on or off, or returns all agents to their default state.



To enable or disable a single agent, use instead the commands [agent on](#) or [agent off](#).

```
setup adaptiveresponse on|off|default
```

setup bcaaa force-register

Force synchronization of PacketShaper's groups of interest table with BCAA; use this command if BCAA was down or the PacketShaper was not connected to BCAA while user group classes were being added or removed on PacketShaper.

```
setup bcaaa force-register
```

setup bcaaa server-test

Verify BCAA servers are configured properly in the current configuration. After you have configured the BCAA server(s), you can verify the server configuration by testing whether an IP address returns the correct user name and group names listed in Active Directory.

```
setup bcaaa server-test <ip-address>
```

where *<ip-address>* is a single IPv4 address that you know is assigned to an Active Directory user.

Examples:

If the BCAA servers are configured properly and the IP address is associated with a user name in Active Directory, the **test** command will return a user name for the IP address that is specified, as well as the names of the groups the user belongs to.



The only group names that display are ones for which a user group class has been created. The user may belong to other Active Directory groups, but they will not be listed for the user unless a class exists for that group. If the user isn't a member of any groups for which group classes have been created, NOT AVAILABLE appears for the group name.

```
# setup bcaaa server-test 10.9.116.215
```

```
user name: BCAA\ADMINISTRATOR
group name(s): group-sanjose, group-engineering
```



If your BCAA server was configured incorrectly (such as an invalid IP address), it may take several minutes for the *Cannot establish a connection to the BCAA server* message to return. It may appear that the session has hung, but the prompt will appear after a few minutes.

If there is a problem with the configuration or BCAA cannot locate the IP address in Active Directory, you will see one of the following error messages in response to the **server-test** command:

Message	Action to Take
BCAAA must be enabled in order to test the configuration.	Enable BCAAA: setup bcaaa service on
There is no primary server configured. Please configure a primary server and try again.	Configure BCAAA server: setup bcaaa service primary <host> [<port>]
Can't establish a connection to the BCAAA server. Please make sure you have configured the primary or secondary server properly.	Show BCAAA settings and make sure they are correct: setup bcaaa show
The user name could not be determined for this IP: x.x.x.x	Find a valid IP address and issue setup bcaaa test <ip-address> again.

setup bcaaa service

Configure PacketShaper as a Blue Coat Authentication & Authorization Agent (BCAAA) client in the current configuration; this feature enables the PacketShaper to classify and report on user names and group names in Active Directory.

To enable or disable the BCAAA service, use:

```
setup bcaaa service on|off|default
```

To configure the settings of the BCAAA servers, use:

```
setup bcaaa service primary|secondary {<host> [<port>]}|delete
```

primary secondary	Enter the literal primary or secondary to indicate which server you are defining. (The secondary server is used when the primary server isn't accessible.)
<host>	The IPv4 address or DNS name of the BCAAA server
[<port>]	The port number to access the server; if omitted, the default port 16101 is used.
delete	Deletes the configuration of the primary or secondary server (whichever is specified)
override	For a PolicyCenter child configuration, create local BCAAA settings that override the settings it inherits from its parent configuration. Remove these override settings at any time with the command setup bcaaa primary secondary delete .

Example:

This example defines a primary BCAAA server at 10.10.10.10 using the default port and a secondary server at 10.10.20.10 using port 903. The third command line enables the BCAAA service.

```
setup bcaaa service primary 10.10.10.10
setup bcaaa service secondary 10.10.20.10 903
setup bcaaa service on
```

setup bcaaa show

Display current Blue Coat Authentication & Authorization Agent (BCAAA) settings and the status of each BCAAA server in the current configuration.

```
setup bcaaa show
```

Example output:

BCAAA Setup values:

```
BCAAA Service      : on
Timeout            : 10 seconds
```

BCAAA Servers:

Type	Host	Port	Status
Primary	10.9.112.240	16101	In use
Secondary	10.9.112.245	16101	Available

The Status column has the following possible values:

Status	Description
Attempting to connect	PacketShaper is in the process of trying to establish a connection to this BCAAA server.
In use	PacketShaper is connected to this BCAAA server and is using it to look up users.
Available	PacketShaper is able to connect to this server but is using another working server (one with an "In use" status) for user lookups.
Failed connection	PacketShaper tried to connect to this BCAAA server but the connection failed. Possible reasons for a failed connection include: <ul style="list-style-type: none"> • An incorrect IP address is specified. • BCAAA is configured to use a different port than PacketShaper. • BCAAA server is down. • Link is down. • No route exists between the networks.

Issue the **setup bcaaa show** command to refresh the status.

setup bcaaa timeout

Specify the number of seconds that PacketShaper will wait for a response from the BCAAA server when looking up a user name.

```
setup bcaaa timeout <seconds>|default
```

where *<seconds>* is a value between 2-30 seconds. The default timeout is 10 seconds.

Example:

```
setup bcaaa timeout 20
```

setup capture

Capture the current configuration. The output can be created in a portable form (without specific IP addresses) or complete form (IP addresses included). The output file includes commands to recreate the configuration.

 This command is not intended to be a substitute for backing up PacketShaper configurations.

```
setup capture [[portable|complete] [<filename>]]
```

[portable complete]	Indicate the format of the output. Portable is the default, if no output format is specified. The portable file "comments out" the unit-specific details, such as the setup commands for the IP addresses. The complete file contains address information.
<filename>	The filename is limited to an eight-character name with a three-character suffix. If no filename is specified, the file 9.256/cmd/config.cmd is created. If you specify a filename, you must also specify the format type: portable or complete. If you specify a filename without a suffix, the .cmd suffix is appended to the filename. The file is placed in the 9.256/cmd folder, unless you specify an explicit pathname.

Additional Information

- To restore the captured configuration, use the [run](#) command (for example, **run config.cmd**).
- Auto-discovered port classes are not recreated when you run the CMD file created with the **setup capture complete** command.
- For security reasons, the **setup capture** command does not store the look and touch passwords.
- The **setup capture** command will not capture SNMPv3 authentication and privacy passwords modified via SNMP Set requests unless the user is defined with the *localizedKey* option.

setup discover

Turn traffic discovery on or off for the default inbound and outbound classes in the current configuration.

When you turn on traffic discovery, PacketShaper monitors the traffic going through the unit and classifies the traffic by service type. The traffic discovery process inserts classes into your traffic tree.

```
setup discover off|on
```

Use [class services](#) to list supported protocols and services.

Use [class discover](#) to enable/disable traffic discovery within a specific class.

setup email

Configure email settings for use with the scheduled command, adaptive response action files, and/or user event features in the current configuration.

```
setup email <address>[:<port>] [<sender name>] | none | default
```


<code><address></code>	Specify the email server using either its DNS name or IP address.
<code>[:<port>]</code>	The default SMTP port is port 25. To specify a non-standard port for email messages, enter number.
<code>[<sender name>]</code>	<p>The <code><sender name></code> will appear in the From line of any email message that the user event or scheduled command feature sends out. Specify a complete mail address, including the domain name for example, <i>john_doe@example.com</i></p> <p>Note: In the prompting mode for this command, full form can also be used for the sender that is, a quoted name followed by the explicit email address.</p> <p>For example: "Bob" <bob@examplecompany.com></p>

After PacketShaper has been configured to send email using the **setup email** command, use [setup show](#) to view the configuration.

To clear the email settings:

```
setup email none
```

Use the **default** option to remove the local override. This command allows the child configuration to inherit the parent's email setting:

```
setup email default
```

setup failover

Bandwidth on a WAN link can change, causing the router to fail over to a secondary link. PacketShaper can be configured to detect this failover condition and enforce the new, lower link speed. This feature applies only to site routers that have been configured for failover.

PacketShaper polls the site router every two seconds to determine the status of the links. If the site router has two links sharing the load and one of the links goes down, PacketShaper uses the failover settings to adjust its link speed. If the site router has a primary and a backup link configured, and the primary link fails, PacketShaper also handles the failover condition.

```
setup failover none|show
```

```
setup failover none|show|(<primary ifIndex> <secondary ifIndex> [<backup speed>|either])
```

<code>none</code>	Turn off failover mode.
<code>show</code>	Display current failover statistics.
<code><primary ifIndex></code>	Specify the SNMP index number of the first router interface.
<code><secondary ifIndex></code>	Specify the SNMP index number of the secondary router interface.

<code><backup speed></code>	Set the speed to be used if failover is activated. Rates may be specified as integer bits per second, followed by a k (thousands), M (millions), or G (billions).
<code>either</code>	Use either when the two interfaces are being used for load balancing, not as primary and backup links.

setup flowrecords engineID/engineType

Assign an identifying number to a PacketShaper, when using the NetFlow-5 flow detail record (FDR) format. EngineID and EngineType are two of the fields in NetFlow-5 headers; you can use either, or a combination, of these fields to identify the PacketShaper that is emitting records. These fields are not relevant for the Packeteer-1 and Packeteer-2 formats.

```
setup flowrecords engineID|engineType <value>|none|default
```

<code><value></code>	An integer (0-255) that identifies the PacketShaper. The default value for engineID and engineType is 0.
<code>none</code>	Clears the EngineID or EngineType field
<code>default</code>	Removes the local settings for the EngineID or EngineType field so that the unit inherits the settings of the parent configuration. If the parent configuration doesn't have any EngineID or EngineType settings, the local settings will be cleared so that the unit can inherit any future settings that are set.

Example

To assign an ID of 12 to the current PacketShaper:

```
setup flowrecords engineID 12
```

All records emitted from this PacketShaper to all defined NetFlow-5 collectors will have a value of 12 in the header's engineID field. Thus, the source PacketShaper is easily identifiable when interpreting the flow detail records.

setup flowrecords filters

Defines an include or exclude list to specify whether or not flow detail records (FDR) are emitted for traffic that matches the specified classes, services, and/or subnets. You can use this command in either of two ways:

- To define an **include** list so that flow detail records are *always* emitted for traffic that matches the specified classes, services, and/or subnets. This approach is recommended when you want flow detail records for only a few specific classes, services, and/or subnets. Flow detail records will not be emitted for traffic that does not match the class, service, or subnets specified on the include list.

or

- To define an **exclude** list so that flow detail records are *never* emitted for traffic that matches the specified classes, services, and/or subnets. This approach is recommended when you want flow detail records for all traffic except that which matches the classes, services, and/or subnets on the exclude list.

```
setup flowrecords filters [add|remove|show] [class|service|subnet] include|exclude
[[<class name>|<class id>]| [<service name>]| [<ip:netmask>|<ip/netmask>|<ip>]]
```

[add remove show]	add adds the specified class, service, or subnet to the include or exclude list remove removes the specified class, service, or subnet from the include or exclude list show displays the classes, services, and subnets on both the include and exclude lists
[class service subnet]	Indicates that the FDR filter applies to a class , service , or subnet
include exclude	include specifies that flow detail records will always be emitted for traffic that matches the class, service, and/or subnet exclude specifies that flow detail records will never be emitted for traffic that matches the class, service, and/or subnet
<class name> <class id>	The name of the class or the class id
<service name>	The name of the service
<ip:netmask>	The IP address and subnet mask subject to the IP filter, where <i>netmask</i> is the subnet mask in decimal notation
<ip/netmask>	The IP address and subnet mask subject to the IP filter, where <i>netmask</i> is an integer (the CIDR value) that specifies the number of binary 1s in a mask
<ip>	The IP address subject to the IP filter

Examples

When you add classes to the include list, FDRs are emitted only for those classes. To add a class to the include list:

```
setup flowrecords filters add class include /Inbound/SNMP
```

To add the FTP service to the include list:

```
setup flowrecords filters add service include ftp
```

To add subnets specified by IP address in decimal notation to the include list:

```
setup flowrecords filters add subnet include 10.10.10.01:255.255.255.255
```

To add subnets specified by IP address and CIDR value to the include list:

```
setup flowrecords filters add subnet include 10.10.10.01/32
```

To add an IP address and all of its subnets to the include list:

```
setup flowrecords filters add subnet include 10.10.10.01
```

PolicyCenter CLI Commands

When you add classes, services, and subnets to the exclude list, FDRs are emitted for all classes, services, and subnets except for those specified on the exclude list. For example, to add a class to the exclude list:

```
setup flowrecords filters add class exclude /Inbound/SNMP
```

To remove a class, service, or subnet from the include or exclude list, use the remove keyword. For example:

```
setup flowrecords filters remove class include /Inbound/SNMP
setup flowrecords filters remove class exclude /Inbound/SNMP
setup flowrecords filters remove service include ftp
setup flowrecords filters remove subnet include 10.10.10.01
```

To show all FDR filters:

```
setup flowrecords filters show
```

setup flowrecords id

Define the settings for a flow detail record (FDR) collector in the current configuration. Up to four collectors can be defined.

```
setup flowrecords id [<ID> <collectorDefinition>|off|on|none|default]
```

<ID>	Identifying number of the collector (1, 2, 3, or 4)
<collectorDefinition>	where <collectorDefinition> is <recordType> <ipaddr> [<port> on off] <recordType> is the type of record format to be emitted (netflow-5, packeteer-1, or packeteer-2) <ipaddr> is the IP address of the collector <port> is the UDP port number of the collector (default = 9800)

off on none default	<p>on enables the collector. When a collector is enabled, PacketShaper will emit flow detail records to the collector.</p> <p>off disables the collector; flow detail records will not be emitted.</p> <p>none clears the collector settings; the row will be empty in the setup flowrecords show output.</p> <p>default removes the local settings for the ID so that the unit inherits the collector settings of the parent configuration. If the parent configuration doesn't have any settings for this ID, the local settings will be cleared so that the unit can inherit any future collector settings that are set. This command is only applicable to shared mode with PolicyCenter.</p> <p>You can enable/disable a collector when you are defining it:</p> <pre>setup flowrecords id 1 netflow-5 10.10.10.10 9800 on</pre> <p>or, after a collector has been defined:</p> <pre>setup flowrecords id 1 off</pre>
---------------------	---

A collector is defined by its record type (NetFlow-5, Packeteer-1, or Packeteer-2) and its location (IP address and UDP port number). You can define collectors with the same IP address but different record types, or with the same record type but different IP address. For example, you can create two collectors with the same IP address (but different ports), with one collector collecting NetFlow data records and the other collecting Packeteer-2 data records.

To view your collector settings, use the [setup flowrecords show](#) command.

Examples

To define a collector that collects Packeteer-2 flow detail records:

```
setup flowrecords id 1 packeteer-2 10.10.10.1 9800 on
```

Because 9800 is the default port and "on" is the default, you can use the following alternative command:

```
setup flowrecords id 1 packeteer-2 10.10.10.1
```

To turn off collector 1 (assuming collector 1 has been previously defined):

```
setup flowrecords id 1 off
```

With the above command, PacketShaper will stop emitting flow detail records to collector 1, but will retain the collector settings. To start emitting records again, use this command:

```
setup flowrecords id 1 on
```

To clear the settings for collector 3:

```
setup flowrecords id 3 none
```

setup flowrecords show

Display the flow detail record (FDR) collector settings in the current configuration. Use this command to see the collectors that have been configured and check which ones have been enabled. This command is also useful to look up the ID number associated with a collector (the ID is needed for defining and clearing collector settings).

```
setup flowrecords show
```

ID	RecordType	CollectorIP	Port	Enabled
1	packeteer-2	10.10.10.1	9800	on
2	netflow-5	10.10.10.2	9800	off
3				
4				

In the above sample output, two collectors have been defined. The first collector (ID of 1) collects Packeteer-2 flow detail records and is currently enabled. The second collector (ID of 2) collects NetFlow-5 records but is currently disabled. Collector IDs 3 and 4 have not been defined.

setup heartbeat

Configure PacketShapers in the current configuration to emit messages (heartbeats) to the Symantec heartbeat server. Using the information contained in the heartbeat messages, Symantec is able to compile statistics on the stability of various software releases and hardware products. The heartbeats can also be used to identify and resolve defects.

```
setup heartbeat on|off|default|show
```

Heartbeat emission is enabled by default. Symantec recommends that you *not* disable the feature. Be assured that the messages are encrypted and sent securely via HTTPS. The size of the daily heartbeat message is negligible (30-40K) and has virtually no impact on PacketShaper performance.

Examples:

To disable the heartbeat feature:

```
setup heartbeat off
```

To find out when/if the last heartbeat message was sent:

```
setup heartbeat show
```

```
Heartbeat                               : On
Most Recent Attempt Status : Heartbeat sent successfully, Jun 17 2010
10:31:06
```

If PacketShaper isn't able to send the heartbeat, the **setup heartbeat show** output will indicate the reason for the failure:

- DNS resolution errors
- Transport protocol error
- File I/O errors
- Error running certain commands
- Insufficient free hard disk space



Depending on how long it takes for the configuration system to initialize DNS server entries, the **setup heartbeat show** output may show a DNS resolution error if the PacketShaper cannot resolve the heartbeat server name via DNS quickly enough.

setup link

Configure the access link capacity in the current configuration. To effectively manage the traffic on the link, PacketShaper must know the capacity it is managing.



PacketShaper will enforce the link size that you set.

```
setup link inbound|outbound|default [<size_bps>|default]
```

Specify a rate as either a bits-per-second value or a symbolic name, as shown in the following list of valid link sizes.

<n>	Size in bits per second
<n>k	Size in kilobits per second
<n>m	Size in megabits per second
<n>g	Size in gigabits per second
T1	1.5 Mbps
E1	2 Mbps
T3	45 Mbps

Examples:

```
setup link inbound 1500000
```

```
setup link outbound 1.5m
```

```
setup link inbound T1
```

Considerations

- For full-duplex Ethernet, enter the total link speed for the inbound and outbound rates. Because full-duplex has wires that can simultaneously communicate in both inbound and outbound directions, you should enter the same rate for Inbound Rate and Outbound Rate. For example, if you have two T1 lines (3 Mbps), you should enter 3M for Inbound Rate and 3M for Outbound Rate.
- For half-duplex Ethernet, split the rate between the inbound and outbound links. For example, if you are managing 10 Mbps Ethernet, you could configure 5 Mbps for the inbound rate and 5 Mbps for the outbound rate.

PolicyCenter CLI Commands

- If your appliance is using NICs to manage different WAN links and you don't want to control each NIC separately, the rate should be the size of the smallest NIC. For example, if you have two 100 Mbps NICs managing two links, you should specify 100M for the rate.

On the other hand, if you want to control each link separately, the rate should be the sum of the link speeds on all devices. For example, if the built-in device is controlling a T1 line (1.5 Mbps) and a NIC option card is managing two T1 lines (3.0 Mbps), you should specify 4.5M for the rate. To control traffic across each link separately, you can create a class for each device (for example, *Slot1* and *Slot3*) and assign partitions that match the link size (1.5M for the Slot1 class and 3.0M for the Slot3 class).

If your appliance is using two NICs to manage a single WAN link, specify the WAN link speed for the rate. Although the Info page will give you an error message (such as Link speed of 155 Mbps exceeds outside NIC speed of 100 Mbps) in the latter situation, it is still appropriate to specify the actual size of the link for the rate.

- Software configuration determines maximum shaping capacity.



10BaseT links rarely reach the 10 Mbps limit. Keep Ethernet's practical limits in mind when configuring rates.

setup message

Configure a message that will display before logging into PacketShapers in the current c onfiguration. The message displays before you log in via the browser login page, before logging in using a remote login utility (such as putty), and when you first console connect to the unit. This feature is useful for informing users about the company's access policies and consequences for unauthorized use.

```
setup message {set <message>}|show|default
```

where

set <message>	Defines the message text. The text should be enclosed in quotation marks and can be up to 511 characters long.
show	Displays the content of the login message
default	Clears the message text. In PolicyCenter's shared mode, the unit will then be able to inherit the message of the parent configuration.

Examples

```
setup message set "Access to this system is restricted to authorized users only."  
Message set to: "Access to this system is restricted to authorized users  
onl...
```

```
setup message show
```

```
Configured Message:  
Access to this system is restricted to authorized users only.
```

Notes

- Quotation marks indicate the beginning and end of the login message. You cannot use a quotation mark within the body of the login message.
- If you want to display a message that is longer than 511 characters, you can create a text file that contains your message text. Name the file login.txt and upload it to the 9.256/ folder. The first 2048 characters of the text file will display after any message that is configured with the **setup message set** command. Thus, the text file is appended to the message text, allowing the message to have a total approximate length of 2500 characters. Note that quotation marks are allowed in the login.txt file.
- The **setup message show** command does not display the content of the login.txt file.

setup mirrorlink

Enable/disable link state mirroring. With link state mirroring, PacketShaper will bring down the second port of a NIC pair if the first goes down. This feature allows each PacketShaper to sit between a WAN router and a switch without blocking detection of switch outages by the router.

```
setup mirrorlink on|off
```

This CLI command is the equivalent of the [setup variable mirrorlinks](#) command; these two commands are interchangeable.

setup password

Configure a touch (read/write) or look (read-only) password for PacketShapers in the current configuration. The default password for the touch user is *touch* and the default password for the look user is *look*.

```
setup password look|touch
```

Passwords can be up to nineteen characters long and are case sensitive. They can consist of a combination of letters, numbers, and all special characters on the U.S. keyboard. Passwords cannot be blank. Passwords are set to their defaults when the **setup reset all** or **setup reset clear** commands are executed.

You will be prompted to enter the old password, type a new password, and retype the password to confirm. For example:

```
setup password touch
```

```
Old touch password: (none)
New touch password:
Confirm touch password:
Changed the touch password
```

To abort this command and return to the command prompt, press Ctrl+D.

setup radius acct

Set up or change the settings of the RADIUS accounting service in the current configuration. This feature allows you to have an audit trail for user logins.

PolicyCenter CLI Commands

To define the RADIUS accounting service, use:

```
setup radius acct primary|secondary <host> <shared_secret> [<port>] |delete|override
```

primary secondary	Enter the literal primary or secondary to indicate which server you are defining. (The secondary server is used when the primary server isn't accessible.)
<host>	The IP address or DNS name of the RADIUS accounting server
<shared_secret>	The designated secret for the server; quotes are not required
[<port>]	The port number to access the server; if omitted, the default port is used
delete	Deletes the configuration of the primary or secondary server (whichever is specified)
override	Inherits setup of the primary or secondary server from PolicyCenter

To turn the service on or off, or to return the service to its default **off** value, use:

```
setup radius acct on|off|default
```

Example:

```
setup radius acct primary 10.10.10.10 bobolink
```

```
setup radius acct secondary 10.10.20.10 parrot
```

```
setup radius acct on
```

This example defines a primary accounting server at 10.10.10.10 which has a shared secret of bobolink, as well as a secondary server at 10.10.20.10. The third command line enables RADIUS accounting service. Once this service is configured and enabled, PacketShaper will send a PW_STATUS_START accounting message to the accounting server when a user logs in and a PW_STATUS_STOP message when a user logs off or is disconnected.

setup radius auth

Set up or change the settings of the RADIUS authentication service in the current configuration. RADIUS authentication is an optional method for users to log into the PacketShaper browser interface, command-line interface, or customer portal or when FTPing to the unit. Using third-party RADIUS servers enables you to have central configuration of user accounts.

```
setup radius auth primary|secondary <host> <shared_secret> [<port>] |delete|override
```

primary secondary	Enter the literal primary or secondary to indicate which server you are defining. (Note: The RADIUS client uses the secondary server when the primary server isn't accessible or authentication failed.)
<host>	The IP address or DNS name of the RADIUS authentication server
<shared_secret>	The designated secret for the server; quotes are not required

[<port>]	The port number to access the server; if omitted, the default port is used
delete	Deletes the configuration of the primary or secondary server (whichever is specified)
override	Inherits setup of the primary or secondary server from PolicyCenter

To turn the service on or off, or to return the service to its default **off** value, use:

```
setup radius auth on|off|default
```

Example:

```
setup radius auth primary 10.10.10.10 bobolink
```

```
setup radius auth on
```

This example first defines a primary authentication server at 10.10.10.10 which has a shared secret of *bobolink*. The second command line enables RADIUS authentication service. Once this is configured and enabled, PacketShaper will prompt users for user name and password when they log into PacketShaper.

setup radius interval

Set the amount of time for RADIUS to wait for a response from a server. By default, the RADIUS client waits 5 seconds before retrying a login when the RADIUS server fails to respond.

```
setup radius interval <seconds>|default
```

where <seconds> is a value between 1 and 30 seconds. For example:

```
setup radius interval 20
```

In this example, the retry interval is 20 seconds; this interval applies to any configured RADIUS server.

To return to the default retry interval, use:

```
setup radius interval default
```

setup radius limit

Set the number of retry attempts the RADIUS client will make to a server before cancelling the login. By default, if the RADIUS server fails to respond, the RADIUS client will try to log onto the server three times before reporting a server failure. If you have specified a secondary server, the RADIUS client will alternate attempts to log onto each server.

```
setup radius limit <n>|default
```

where <n> is a value between 1 and 10. For example:

```
setup radius limit 6
```

In this example, the RADIUS client will try to log onto the server six times.

PolicyCenter CLI Commands

To return to the default retry limit, use:

```
setup radius limit default
```

setup radius method

Select the RADIUS authentication method for the current configuration: PAP, CHAP, or MSCHAP.

```
setup radius method PAP|CHAP|MSCHAP|default
```

PAP (<i>Password Authentication Protocol</i>)	With PAP, the user name and password are transmitted in clear, unencrypted text. ASCII or PAP authentication is required for RADIUS configurations that require access to clear text passwords (for example, when passwords are stored and maintained in a database external to the RADIUS server).
CHAP (<i>Challenge Handshake Authentication Protocol</i>)	In some environments, CHAP may be preferred for greater security. The RADIUS server sends a challenge that consists of a session ID and an arbitrary challenge string, and the user name and password are encrypted before they are sent back to the server. CHAP is the default authentication method.
MS-CHAP (<i>Microsoft Challenge Handshake Authentication Protocol</i>)	This protocol is similar to CHAP, but with MS-CHAP authentication, the RADIUS server can store an encrypted version of a user password to validate the challenge response. Standard CHAP authentication requires that the server stores unencrypted passwords. Note: MS-CHAP v1 and v2 are supported. PacketShaper attempts authentication with MS-CHAP v2 first. If the remote server doesn't support v2 or if authentication is denied, PacketShaper re-attempts authentication with MS-CHAP v1.

setup radius show

Display the current RADIUS settings in the current configuration. Use this command to verify that RADIUS authentication and accounting are enabled, to see the current settings for the retry limit and retry interval, and to view configuration settings on each of the RADIUS servers.

```
setup radius show
```

Authentication :on

Accounting :on

Retry limit :3

Retry interval :5

Service records:

Type	Host	Port	Secret
auth1	10.7.55.1	1812	testing123

```
acct1 10.7.55.1 1813 testing123
```

setup reset

Clears PolicyCenter configurations and settings.

```
setup reset [all|clear]
```

setup reset (issued for a unit configuration)	Restores the unit configuration to its default settings so that it inherits settings from its parent configuration
setup reset all	Resets the PolicyCenter configuration settings
setup reset clear	Resets the PolicyCenter configuration and <i>all</i> shared configurations in PolicyCenter

setup secure

Limit management access from the inside or outside interface.

```
setup secure inside|outside|mgmt on|off|default|list <addr>[:<mask>]...
```

Use the **setup secure outside on** command to secure the outside interface, that is, the Internet. For example, when the outside interface is set to secure, Telnet, HTTP, FTP and ping requests from external sources will not be permitted. By default, the inside and outside interfaces are not secured.

The **list** parameter enables access to up to 16 IP addresses, separated by spaces. This is an exception list—the interface is secured except for the IP addresses on the list. To specify a subnet, use the format: *ipaddress:subnet_mask* or *ipaddress/CIDR*. The **list** option accepts IP addresses only, not host names. To find the IP address associated with a host name, use the [dns lookup](#) CLI command.

Notes:

- If you secure the interfaces, you will be able to access the unit *only* via a console connection. The browser interface will be disabled because you will not have management access over the network. Another way to secure the interface is to specify a list of IP addresses that can access the unit. For example, **setup secure outside list 10.1.1.100 10.1.12.1** would allow access from only two IP addresses.
- Keep in mind that securing an interface means that queries such as DNS and SNTP cannot be made via the secured interface. Consider using the **list** option and including these servers and your gateway in the list.
- The WebPulse classification features require access to a number of outside web servers. Therefore, do not completely secure the outside interface. Instead, use the **list** option and add the IP addresses of the following servers to the exception list: WebPulse service points (use the **setup webpulse show service** command to find the IP addresses of the one or two fastest servers), the WebPulse map update server (sitereview.bluecoat.com), the support update server (updates.bluecoat.com), the heartbeat server (hb.bluecoat.com), and the traffic information reporting server (cda.bluecoat.com). If you are using a web proxy, you will also need to add this server's IP address to the list if it will be accessible via the outside interface.

PolicyCenter CLI Commands

- If you plan on using direct standby, do *not* set the outside interface to **secure**. For standby to work, each device must be able to communicate with the other device. If you set the outside interface to **list**, you must add both the partner's and the unit's IP addresses to the Outside security list.
- The PacketShaper will not be able to process local ARP requests via a secured interface.
- If you secure the outside interface and your gateway is on the outside, a "gateway not found" message will be displayed in the login banner or on the info page. In this state, tasks such as upgrading the software image from a non-local address will be disabled.
- The MGMT port (available on certain models) is considered an outside port. Therefore, securing the outside interface will secure the MGMT port as well.

setup shaping

When shaping is turned on, traffic is classified and measured, and control policies are enforced for all units assigned to the configuration. When shaping is off, traffic is classified and measured but not managed.

```
setup shaping on|off|bypass|passthru|watch
```

Where:

on	Turns traffic shaping on
off	Turns off shaping mode (traffic, bypass, passthru, watch)
bypass	Sets the unit to pure bypass mode. Bypass mode prevents both packet shaping and further network management access; it is as if the unit were removed, and cables connected around it.
passthru	Turns off all shaping, classification, and measurement
watch	Sets the unit into a non-inline, monitor-only mode.

setup show

Display the basic configuration.

```
setup show
```

The output is divided into non-sharable (local) and sharable settings. The sharable settings are part of the configuration file (config.idi). If a configuration is loaded on another unit, the sharable settings will be copied to the other unit (see [config save](#) and [config load](#)).

Non-sharable (local) settings:

```
IP address:          10.9.28.170 Subnet mask: 255.255.252.0
Gateway:             10.9.28.1
DNS server(s):       10.2.2.100 10.2.2.11 10.2.2.10
Default domain:      (none)
```

```

Date, time, timezone:   Thu Nov  7 10:12:49 2013 PST (LosAngeles)
Category Map:          v9 downloaded on Wed Nov  6 10:28:38 2013
Web Applications Map:   downloaded on Wed Nov  6 10:28:38 2013
Slot2_inside_0 nic speed: 1Gbps (Twisted pair) Full-Duplex(auto-negotiate)
Slot2_outside_1 nic speed: 1Gbps (Twisted pair) Full-Duplex(auto-negotiate)
Slot1_inside_0 nic speed: No link (manual)
Slot1_outside_1 nic speed: No link (manual)
Mgmt_inside_0 nic speed: 1Gbps Full-Duplex(auto-negotiate)
Installed Keys:        demo                               on (Expires: DEC 31, 2013)

```

```

Sharable settings:
  Inside interfaces:    unsecure
  Outside interfaces:   unsecure
  Look password:        *****
  Touch password:       *****
  Link speed:           1G (1000BT))
  Packet shaping:       on
  Traffic discovery:    on
  WebPulse:             on
  Web Proxy:            off
  SNTP Client:          off
  SNTP Primary Server:  time.nist.gov
  SNTP Secondary Server: time-a.nist.gov
  SNTP Poll Seconds:    300
  BCAAA:                on
  BCAAA Primary Server: 10.9.30.140:16101
  BCAAA Secondary Server: 10.9.30.140:16101
  BCAAA Timeout Seconds: 10
  Syslog:               off

```

setup snmp accessgroup

Each SNMP access group is defined by a group name, a security model (and level), and a set of views that specifies which types of MIB data that access group can read or write.

```

setup snmp accessgroup new|modify|override <groupname>
[noAuthNoPriv|authNoPriv|authPriv] [read <viewname>] [write <viewname>]}

```

where

PolicyCenter CLI Commands

<code><groupname></code>	The name of the user group you are creating or modifying in the configuration. An access group name can be up to 32 characters; hyphens, underscores, and periods are acceptable. If the group name contains spaces, it must be enclosed in quotations marks, for example "admin group."
noAuthNoPriv authNoPriv authPriv	SNMPv3 access groups support the following security levels: <ul style="list-style-type: none"> • noAuthNoPriv: Identifies a user for access control, but does not provide authentication. • authNoPriv: Identifies a user for access control, and authenticates the user's password. • authPriv: Identifies a user for access control, authenticates the user's password, and provides encryption. <p>If you do not specify a usm security model, the group will use the default noAuthnoPriv.</p>
read <code><viewname></code>	Access groups have <i>read</i> (look) access to the information specified by the read view. To give the group read access to all MIB data, specify the predefined view name isoAll for the <code><viewname></code> parameter. To block all read access, specify isoNone . To limit a group's read access to a subset of available MIB data, enter the name of a user-defined view created with the setup snmp view command. If you do not specify a read view, the group will apply the default isoAll setting.
write <code><viewname></code>	Access groups have <i>write</i> (touch) access to the information specified by by the write view. To give the group write access to all MIB data, specify the predefined view name isoAll for the <code><viewname></code> parameter. To block all write access, specify isoNone . To limit a group's write access to a subset of available MIB data, enter the name of a user-defined view created with the setup snmp view command. If you do not specify a write view, the group will apply the default isoAll setting.

Examples:

```
setup snmp accessgroup new engineering usm authpriv read isoall write isoall
setup snmp accessgroup new admin usm authpriv read snmpTraps write isoNone
```



Since each individual SNMP setting change requires a full configuration reload, Blue Coat recommends that you temporarily disable SNMP (or wait to enable it) before modifying multiple SNMP settings. With this approach, the configuration only needs to reload once (upon enabling SNMP) which avoids significant overhead.

Delete an Access Group

To delete an access group, use:

```
setup snmp accessgroup delete <groupname>
```

where *<groupname>* name of the group you want to delete. Note that you will not be able to delete a group that currently has users assigned to it.

Example:

```
setup snmp accessgroup delete marketing
```

View Access Group Settings

To view current SNMP access group settings, issue the command:

```
setup snmp accessgroup show
```

Example output:

AccessGroupName	Model	Level	ReadViewName	WriteViewName		
Refs Status						
admin	usm	authPriv	all_mib	isoNone		
10 ok						
engineering	usm	authPriv	all_mib	Trap	23	
ok						
v2	v2c		isoAll	isoAll	1	ok
v1_users	v1		isoAll	isoAll	0	ok

setup snmp configmode

Specify which mode of SNMP you will use to access your PacketShaper. Omit the `[simple|complex|default]` parameters to view the currently configured setting. *Although you can set the parent configuration to complex mode for configuring SNMPv3, the specific complex mode commands (such as add and modify) can only be set on legacy PacketShaper appliances.*



Complex mode configuration is only recommended for *advanced users with previous experience working with SNMPv3*, as this mode does not display error messages for incorrectly configured settings that can prevent SNMP from working correctly. Complex mode should only be used in PolicyCenter to set SNMPv3 values for an individual unit configuration. Any complex mode SNMPv3 values set on a PolicyCenter sharable configuration will not be inherited by units assigned to that configuration.

```
setup snmp configmode [ simple|complex|default]
```

where

PolicyCenter CLI Commands

Parameter	Description
simple	Simple SNMPv1 configuration relies on IP address-based access lists and community strings for authentication
complex	Complex SNMPv3 configuration allows access to the SNMP configuration tables and provides security features for authentication, privacy, and access control
default	<p>When the setup snmp configmode default command is issued for an individual PacketShaper or a PolicyCenter configuration at the top of the PolicyCenter configuration tree, the default parameter returns the SNMP mode to the default simple (SNMPv1) setting.</p> <p>When issued for a PolicyCenter child configuration, that child configuration will clear its local SNMP version, and inherit its SNMP version from its parent.</p>

Examples:

```
setup snmp configmode complex
```

or

```
setup snmp configmode
```

```
SNMP configmode: complex
```

setup snmp destinations

The IP address(es) of SNMPv1 trap destination(s) to where PacketShapers in the current configuration will send traps. For *<ipaddress>*, you can specify up to eight IP addresses in dotted-decimal notation, separated by spaces. Note that DNS names are not supported. If the **default** parameter is specified, all IP address destinations are cleared.



If the **setup snmp destinations default** command is issued from a PolicyCenter child configuration, that configuration will clear its local destination settings, but will immediately reinherit SNMP destinations defined in its parent configuration. Use the command **setup snmp destinations none** to clear all local SNMP destination settings from a PolicyCenter child configuration without inheriting additional destinations from the parent configuration.

```
setup snmp destinations <ipaddress>...|default|none
```

Example:

```
setup snmp destinations 172.22.20.156 172.23.21.19
```

setup snmp disable

Disable SNMP functionality on the PacketShapers assigned to the current configuration. When SNMP is disabled, the PacketShaper is not able to accept incoming SNMP requests or send SNMP traps. SNMP is disabled by default.

```
setup snmp disable
```



Since each individual SNMP setting change requires a full configuration reload, Blue Coat recommends that you temporarily disable SNMP (or wait to enable it) before modifying multiple SNMP settings. With this approach, the configuration only needs to reload once (upon enabling SNMP) which avoids significant overhead.

setup snmp enable

Enable SNMP functionality on the PacketShapers assigned to the current configuration. (SNMP is disabled by default.) When SNMP is enabled, the PacketShaper is able to accept incoming SNMP requests and send SNMP traps.

```
setup snmp enable
```

For full SNMPv1 or v2 functionality on PacketShaper, you also need to configure the look or touch community strings and set trap destinations . For SNMPv3, you need to configure notify targets.

setup snmp look|touch

Set SNMP look or touch community strings (passwords) for the current configuration. SNMP is not functional on PacketShaper until you configure the look community string and enable it with the "setup snmp enable" above command.

```
setup snmp look|touch <string>|default|none
```

Where

Parameter	Description	With default setting
look	The context-sensitive SNMP look (read) password	Look password is cleared (set to public)
touch	The context-sensitive SNMP touch (read/write) password	Touch password is cleared (set to public)

The community string can contain alphanumeric characters, hyphens, underscores, and periods; all other special characters should be avoided.

To see the current settings for the look or touch community string values, use the [setup show](#) command.

Examples:

```
setup snmp look lookpwd1
setup snmp touch touchpwd2
```

```
setup snmp show
```

```
SNMP config mode: simple
SNMP look community: lookpwd1
SNMP touch community: touchpwd2
SNMP Trap destinations: (none)
```

ViewName	SubtreeOID	Type	Refs	Status
B isoAll	iso	included	0	ok
B isoNone	iso	excluded	0	ok

setup snmp remoteuser

An SNMP remote user defines a user or a management system that receives notification of SNMPv3 traps and informs. Unlike a local SNMP user, a remote user is not associated with an access group and therefore has only a notify view, rather than a read or write view.

```
setup snmp remoteuser new|modify|override <username> [<engine-id>] [auth {md5|sha}
<auth-pw>] [priv {des|3des|aes128|aes192|aes256} <priv-pw>]
```

where

<username>	Name of the user you are creating or modifying. Remote user names can have up to 32 characters; hyphens, underscores, and periods are acceptable. If the name includes a space, it must be enclosed within quotation marks, for example "John Doe." Each SNMP remote user name must be unique.
<EngineID>	<p>An SNMP Engine ID identifies an SNMP engine that will receive trap and inform notifications. The default Engine ID for a remote SNMP user is <i>LocalSnmpld</i>, the SNMP agent's own SNMP Engine. If you omit this parameter, the remote user will use this default <i>LocalSnmpld</i> Engine ID.</p> <p>To specify a <i>different</i> remote SNMP engine with which this user can communicate, specify the 24-digit hexadecimal Engine ID of a remote SNMP engine.</p>
auth {md5 sha}	If the remote user requires authentication, specify either the MD5 or SHA authorization protocol and enter an authentication password for the user. If the remote user does not require authentication, this parameter can be omitted.
<auth-pw>	Authorization password for the user. Passwords can have up to 32 characters; hyphens, underscores, and periods are acceptable.
priv {des 3des aes128 aes192 aes256}	<p>Specify one of the following privacy protection protocols if the remote user requires privacy protection. Otherwise, this parameter is not required.</p> <ul style="list-style-type: none"> • des: CBC-DES Symmetric Encryption Protocol • 3des: 3DES-EDE Symmetric Encryption Protocol • aes128: 128-bit AES (Advanced Encryption Standard) • aes192: 192-bit AES • aes256: 256-bit AES
<priv-pw>	Privacy password for the user. Passwords can have up to 32 characters; hyphens, underscores, and periods are acceptable.

Examples:

```
setup snmp remoteuser new "Jane Killick" auth md5 authpwd12$ priv aes245 privpwd12!
0000091E000000A1AC1512AC
setup snmp remoteuser new "Nonsecure user"
```

Delete a Remote User

To delete a remote user, use:

```
setup snmp remoteuser delete <username>
```

Example:

```
setup snmp remote user delete "Sean Wood "
```

View Remote User Settings

To view current remote user settings, issue the command:

```
setup snmp remoteuser show
```

Example output:

```
setup snmp remoteuser show
```

RemoteUserName	AuthProt	PrivProt	RemoteEngineID	Refs	Status
IT_remote	md5	aes256	0000091E000000A1AC1512AA	4	ok
sys admin	none	none		3	ok
Todd Gray	md5	des	0000091E000000A1AC1512A0	1	ok



Since each individual SNMP setting change requires a full configuration reload, Blue Coat recommends that you temporarily disable SNMP (or wait to enable it) before modifying multiple SNMP settings. With this approach, the configuration only needs to reload once (upon enabling SNMP) which avoids significant overhead.

setup snmp show

Display the basic SNMP settings in the current configuration.

```
setup snmp show
```

The output is divided into SNMPv1/SNMPv2c and SNMPv3 settings. The first two sections of the output display settings for an SNMPv1 or SNMPv2 configuration. If you have configured SNMPv3 views, access groups, users and targets, the information for each table entry appears in the *ViewName*, *AccessGroupName*, *UserName*, *RemoteUserName* and *TargetName* sections.

The following is an example of the output of the **setup snmp show** command.

```
setup snmp show
```

```
SNMP config mode:      simple
SNMP look community:   lookpwd1
SNMP touch community:  touchpwd2
SNMP Trap destinations: 172.21.18.166 172.21.18.167

syslocation: Northwest Corner of Building 4
syscontact:  Jill Smith
sysname:     PKTR_9500_42
localSnmpID: 0000091E000000A1AC1512AA
```

PolicyCenter CLI Commands

ViewName	SubtreeOID	Type	Refs
Status			
all_mib	1.3.6.1.6.3.1	included	2
ok			
isoAll	iso	included	
ok			
isoNone	iso	excluded	1
ok			
Traps	snmpTraps	included	1
ok			

AccessGroupName	Model	Level	ReadViewName	WriteViewName
Refs Status				
admin	usm	authPriv	all_mib	isoNone
engineering	usm	authPriv	all_mib	Traps
test_1	usm	noAuthNoPriv	isoAll	isoAll
vlonly	v1		isoAll	isoAll
v2	v2c		isoAll	isoAll

UserName	GroupName	AuthProt	PrivProt
Amit	engineering	md5	des
IT	engineering	md5	des
Marcia	admin	md5	des
RemoteUserName	AuthProt	PrivProt	RemoteEngineID
Refs Status			
IT_remote	md5	aes256	0000091E000000A1AC1512AA
sys admin	none	none	
Todd Gray	md5	des	0000091E000000A1AC1512A0

TargetName	RemoteHost	RemoteUserName	ViewName	Type	Ver
M1 Status					
System admin	10.10.14.55	IT_remote	isoAll	trap	v3
um ok					
Target_it	172.21.18.170	IT_remote	isoAll	trap	v3
um ok					
Vltraps	172.21.18.160	public	isoAll	trap	v1
v1 ok					

setup snmp syscontact

The name of the person managing the units in the current configuration; *<string>* can be up to 256 characters long and must be enclosed in quotation marks if spaces are used. If the **default** parameter is specified, any existing syscontact value is cleared.

```
setup snmp syscontact <string>|default
```

Example:

```
setup snmp syscontact "Gail Jellison"
```

setup snmp syslocation

Issue this command to specify the physical location of the PacketShapers (room, floor, building) in the current configuration.

```
setup snmp syslocation <string>|default
```

where *<string>* can be up to 256 characters long and must be enclosed in quotation marks if spaces are used. If desired, specify the **default** parameter to clear the existing syslocation variable.

Example:

```
setup snmp sysname "4th floor"
```

setup snmp sysname

Issue this command to configure the PacketShaper's fully-qualified domain name for SNMPv1, where *<string>* can be up to 256 characters long and must be enclosed in quotation marks if spaces are used. If the **default** parameter is specified, the unit's IP address is used for the sysName variable.

```
setup snmp sysname <string>|default
```

Example:

```
setup snmp sysname example.com
```

If the PacketShaper is using its IP address as a default sysName and the unit's IP address changes, the sysName will update to the new IP address once the unit resets or an SNMP MIB walk is performed on the unit.

setup snmp target

Create new SNMP targets or modify existing SNMP targets to determine where SNMPv3 notifications from units in the configuration should be sent.

There are two different commands to modify target settings. Modify target settings on a PacketShaper in local mode or a top-level PolicyCenter configuration with the command **setup snmp target modify**. Use the command **setup snmp target override** on a PacketShaper in shared mode or a PolicyCenter child configuration to create a local copy of a SNMP target that overrides the inherited SNMP target.

```
setup snmp target modify|new|override <targetname> <targethost> <remoteuser> [port <port>] [version v1|v2c|v3] [model v1|v2c|usm] [type trap|inform] [view <notifyView>] [timeout <seconds>] [retry <n>]
```

where

<i><targetname></i>	<p>Name of the target you are creating or modifying. Target names can have up to 32 characters; hyphens, underscores, and periods are acceptable. If the target name includes a space, it must be enclosed within quotation marks, for example "target four."</p> <p>Note: A target name can be any text string, and does not have to be related to the target system or remote user name.</p>
<i><targethost></i>	IP address of a remote IP host, in dotted-decimal format.

PolicyCenter CLI Commands

<code><remoteuser></code>	<p>If the new target will use both the v3 protocol version <i>and</i> the <i>usm</i> security model, you must also specify a remote user.</p> <p>To associate an existing remote user with this target, specify a user already defined by the command setup snmp remoteuser. To create a <i>new</i> remote user for this target with a localSnmpIpd and no authorization or privacy protection, specify the name for the new remote user. Note that remote users created with this command will not appear in the <i>Remote Users</i> table.</p>
<code><port></code>	Port number on the remote host to which the notifications will be sent.
version <code>v1/v2/v3</code>	Specify v1 , v2 or v3 to indicate which SNMP version of notifications the user will receive. The default SNMP version is v3 .
model <code>v1/v2/usm</code>	Select the security model for this notification by specifying v1 (for SNMPv1), v2 (for SNMPv2c), or usm (for SNMPv3). The default security models for the different versions of SNMP are as follows:
	<ul style="list-style-type: none">• SNMPv1: defaults to <i>v1</i>• SNMPv2c: defaults to <i>v2</i>• SNMPv3: defaults to <i>usm</i>
type trap inform	<p>Specify whether the user should receive trap notifications or just informs. If no parameter is specified, the default setting will be trap.</p> <p>Note: SNMPv1 supports trap notifications only.</p>
view <code><notifyview></code>	To allow the remote user to receive all types of MIB notifications, specify the predefined view name isoAll for the <code><notifyview></code> parameter. To limit the user's access to a subset of available MIB notifications, enter the name of a user-defined view created with the setup snmp view command. If you do not specify a notify view, the group will apply the default isoAll setting.
timeout <code><seconds></code>	<p>Maximum round trip time for communications between the PacketShaper and the SNMP target address, in seconds. Valid timeout values 1-60 , and the default value is 10.</p> <p>If an inform message is sent to this address but a response is not received within this specified time frame, the PacketShaper will assume that there will be no response.</p>
retry <code><n></code>	Number of times the PacketShaper should attempt to retransmit an inform message when it does not receive a response. Valid retry values are 1-10, and the default value is 3 retries.

Example:

```
setup snmp target new targ1 10.1.2.3 trapuser
```

Delete a Target

To delete a target, use:

```
setup snmp target delete <target>
```

Example:

```
setup snmp target delete "admin_target"
```

View SNMP Target Settings

To view current SNMP target settings, issue the command:

```
setup snmp target show
```

Example output:

TargetName	RemoteHost	RemoteUserName	ViewName	Type	Ver
M1 Status					
System admin	10.10.14.55	IT_remote	isoAll	trap	v3
usm ok					
Target_it	172.21.18.170	IT_remote	isoAll	trap	v3
usm ok					
V1traps	172.21.18.160	public	isoAll	trap	v1
v1 ok					



Since each individual SNMP setting change requires a full configuration reload, Blue Coat recommends that you temporarily disable SNMP (or wait to enable it) before modifying multiple SNMP settings. With this approach, the configuration only needs to reload once (upon enabling SNMP) which avoids significant overhead.

setup snmp user

Each SNMP user entry defines a user (login) name, an association with an existing access group, and authentication and privacy keys that a management system can use to access the PacketShaper. This user name is not related to any other user names such as those defined for RADIUS or PolicyCenter access.

There are two different commands to modify user settings. Modify user settings on a PacketShaper in local mode or a top-level PolicyCenter configuration with the command **setup snmp user modify**. Use the command **setup snmp user override** on a PacketShaper in shared mode or a PolicyCenter child configuration to create a local copy of a SNMP user that overrides the inherited SNMP user.



If you have not yet defined access groups for your SNMP users, use the CLI command [setup snmp accessgroup](#) to create one or more access groups before you add users to these groups.

```
setup snmp user new|modify|override <username> <groupname> [auth {md5|sha} <auth-pw>]
[priv {des|3des|aes128|aes192|aes256} <priv-pw>]
```

where

<username>	Name of the user you are creating or modifying. A user name can have up to 32 characters; hyphens, underscores, and periods are acceptable. If the name includes a space, it must be enclosed within quotation marks, for example "Jane Doe." Each SNMP user name must be unique.
<groupname>	User's access group

PolicyCenter CLI Commands

auth {md5 sha}	<p>If the user's access group uses the <i>usm</i> (SNMPv3) security model with the <i>authNoPriv</i> or <i>authpriv</i> security levels, Specify either the MD5 or SHA authorization protocol and enter an authentication password for the user.</p> <p>If the user's access group uses the <i>v1</i> (SNMPv1) or <i>v2</i> (SNMPv2c) security model or an <i>noAuthNoPriv</i> security level, this parameter is not required.</p>
<auth-pw>	Authorization password for the user. Passwords can have up to 32 characters; hyphens, underscores, and periods are acceptable.
priv {des 3des aes128 aes192 aes256}	<p>Specify one of the following privacy protection protocols only if the user's access group uses the <i>authpriv</i> security level. Otherwise, this parameter is not required.</p> <ul style="list-style-type: none">• des: CBC-DES Symmetric Encryption Protocol• 3des: 3DES-EDE Symmetric Encryption Protocol• aes128: 128- bit AES (Advanced Encryption Standard)• aes192: 192- bit AES• aes256: 256-bit AES
<priv-pw>	Privacy password for the user. Passwords can have up to 32 characters; hyphens, underscores, and periods are acceptable.

Examples:

```
setup snmp user new "Kim Johnson" snmpv3Eng auth md5 authpwd123 priv aes245 privpwd123
setup snmp user new "v1_user" snmpv1Group
setup snmp user modify "Kim Johnson" snmpv3Eng auth md5 new_pwd1 priv aes245 new_pwd2
```



Since each individual SNMP setting change requires a full configuration reload, Blue Coat recommends that you temporarily disable SNMP (or wait to enable it) before modifying multiple SNMP settings. With this approach, the configuration only needs to reload once (upon enabling SNMP) which avoids significant overhead.

Delete a User

To delete a user, use:

```
setup snmp user delete <username>
```

Example:

```
setup snmp user delete "Ken Traum"
```

View Users' Group and Security Settings

To view current SNMP user settings, issue the command:

```
setup snmp user show
```

Example output:

```

setup snmp user show
  UserName      GroupName      AuthProt  PrivProt  Status
  Amit          engineering   md5       des       ok
  Example_v1    vlonly       none      none      ok
  IT            engineering   md5       des       ok
  Jane Doe      engineering   md5       des       ok
  Kim Johnson   admin        md5       des       ok
  Tom Jones     authnopriv   md5       none      ok
  VP_Marcia     admin        md5       des       ok
  Wendy Ho      engineering   md5       des       ok

```

setup snmp view

An SNMP view filters objects from the entire MIB and defines a subset of MIB objects. Every SNMP access group has views for read and write access which either allow or limit that group's access to MIB objects. There are two predefined views; *isoAll* and *isoNone*. The *isoAll* view gives a group access to all MIB information, and the *isoNone* view blocks all access.

If you want your SNMP groups to have either complete access or no access to all MIB information, your groups only need to use the built-in *isoAll* or *isoNone* views. If, however, you want a group to access just a subset of MIB information, you will have to create a new view that describes those MIB object identifiers (OIDs) that should be included or excluded.

There are two different commands to modify view settings. Modify view settings on a PacketShaper in local mode or a top-level PolicyCenter configuration with the command **setup snmp view modify**. Use the command **setup snmp view override** on a PacketShaper in shared mode or a PolicyCenter child configuration to create a local copy of a SNMP view that overrides the inherited SNMP view.

```
setup snmp view add|modify|new|override <viewName> <OID> [included|excluded]
```

where

<viewname>	Name of the view you are creating or modifying in the current configuration. A view name can be up to 32 characters; hyphens, underscores, and periods are acceptable. If the name includes a space, it must be enclosed within quotation marks, for example "admin view." Each SNMP view name must be unique.
<OID>	The <OID> parameter may be an OID name, number, or an initial OID name and a number, e.g., <i>packeteerMibs</i> , <i>1.3.6.1.4.1.2334.2</i> or <i>packeteerMibs.1.4</i> This parameter also supports the use of asterisks as wildcards for OID numbers, for example, <i>interfaces.*.*.1</i>
included excluded	Select included if the OID subtree should be included in this view, or excluded if it is explicitly not accessible. If you do not specify an include or exclude parameter, the OID will automatically be included.

The `setup snmp view new` command only allows you to specify a single OID. To include or exclude an additional OID in the view, use the command `setup snmp view add`.

Examples:

PolicyCenter CLI Commands

```
setup snmp view new sysadmin 1.3.6.1.6.3.18 exclude
setup snmp view add sysadmin 1.3.6.1.6.3.15.1.2.2 exclude
```

When you modify snmp view settings, all OIDs defined for that view are removed and replaced with the one OID specified in the **setup snmp view modify** or **setup snmp view override** command. To add additional OIDs to the modified view, use **setup snmp view add**.

Examples:

```
setup snmp view modify sysadmin 1.3.6.1.6.3.19 exclude
setup snmp view add sysadmin 1.3.6.1.6.3.15.1.2.2 exclude
```



Since each individual SNMP setting change requires a full configuration reload, Blue Coat recommends that you temporarily disable SNMP (or wait to enable it) before modifying multiple SNMP settings. With this approach, the configuration only needs to reload once (upon enabling SNMP) which avoids significant overhead.

Delete a View

To delete a view, use:

```
setup snmp view delete <viewName>
```

Example:

```
setup snmp vew delete IT_view
```

Display View Settings

To display current SNMP view settings, issue the command:

```
setup snmp view show
```

Example output is shown below. The letter **B** before the *isoAll* and *isoNone* view names indicate that these views are built-in, and cannot be modified or deleted.

```
setup snmp view show
```

ViewName	SubtreeOID	Type	Refs	Status
adminview	1.3.6.1.6.3.1	included	8	ok
iface_view	1.3.6.1.2.1.2	included	12	ok
B isoAll	iso	included	24	ok
B isoNone	iso	excluded	2	ok
sysadmin	1.3.6.1.6.3.15.1.2.2	excluded	3	ok
	1.3.6.1.6.3.19	excluded		
Traps	snmpTraps	included	5	ok

setup sntp

Set or display the Simple Network Time Protocol (SNTP) settings for the current configuration. SNTP is used to synchronize the time in PacketShaper to a server configured to propagate highly accurate time information through the Internet.

```
setup sntp on|off|servers {<primary> [<secondary>]|none}|poll|reset|sync
```

Enter a standard dotted-decimal IP address.

setup ssh keygen

Generate new key pairs for accessing the PacketShaper command-line interface (CLI) with a secure connection. If you believe the key's security was compromised, you can use this command to generate new keys. *This command is applicable to legacy PacketShaper configurations only.*

```
setup ssh keygen [<size>]
```

where <size> can be 512, 768, 1024, or 2048 bits; 2048 is the default size starting in PacketWise 9.2.6; 1024 is the default for earlier versions. If you are using SSHv1, you should choose 512 or 1024. If you are using SSHv2, specify either 768 or higher.

You can use the [setup ssh show](#) command to see the fingerprints that were generated.

setup ssh port

Change the SSH (Secure Shell) listening port in the current configuration. PacketShaper is automatically configured to run SSH on port 22; use this command to select a different port. *This command is applicable to legacy PacketShaper configurations only.*



Secure Shell is a program and protocol that provides strong authentication and secure communications for logging onto a remote computer. For secure connections to the PacketShaper command-line interface, you can choose any SSH client, such as SecureCRT for Windows or OpenSSH for UNIX operating systems.

```
setup ssh port <port_number>|default
```

where <port_number> is the new SSH port number and **default** uses the default SSH port, 22.



If your unit is configured in shared mode with PolicyCenter, the default is the SSH port number of the parent group, which may or may not be 22.

Examples:

To use SSH on port 25:

```
setup ssh port 25
```

SSH service will be restarted on port 25. It may take up to 10 seconds for the new value to take effect. Please use "setup ssh show" to verify the service status.

Or, to use SSH on the default port:


```
setup ssh port default
```

PolicyCenter CLI Commands

The SSH service will start on the designated port in less than 10 seconds. If the configured port was already in use, PacketShaper automatically uses the last valid port number specified, or the default value (22). Use the [setup ssh show](#) command to verify that the port number was accepted.

setup ssh show

Display the status of SSH (Secure Shell) in the current configuration. The output indicates whether the SSH service is running and on which port. In addition, the output lists the RSA1, DSA, and RSA fingerprints. (The fingerprinting mechanism makes sure that you are contacting the intended remote host.) The fingerprint appears as a sequence of 16 octets in hexadecimal, separated by colons. *This command is applicable to legacy PacketShaper configurations only.*

 Secure Shell is a program and protocol that provides strong authentication and secure communications for logging onto a remote computer. For secure connections to the PacketShaper command-line interface, you can choose any SSH client, such as SecureCRT for Windows or OpenSSH for UNIX operating systems.

```
setup ssh show
```

Example output (when the configured port is the same as the port that is actually being used):

SSH service is listening on port #: 22 (default)

DSA key fingerprint is:

fc:b9:01:88:cf:02:74:50:5e:e1:c0:f7:ab:e9:62:92

RSA key fingerprint is:


2f:77:8a:d6:ff:72:2f:f6:7d:b3:87:53:60:80:a3:ec

RSA1 key fingerprint is:

9d:fd:1f:8f:bc:29:16:29:f8:a7:b6:a0:6b:c2:5e:a7

You can use the **setup ssh show** command to verify that your port number was accepted. If the configured port (specified with the [setup ssh port](#) command) was already in use, PacketShaper automatically uses the last valid port number specified, or the default value (22). In this situation, the **setup ssh show** output will display a Fail binding to port message and indicate the port number that is being used instead. In addition, a notification will appear in the system banner when you log on. For example:

Attention: SSH service failed to start on the port configured in the configuration file. Port 22 is used instead.

 The output of the **setup show** command lists the SSH port number that was configured, which is not necessarily the port number currently in effect. Use **setup ssh show** to see the SSH port number that is in effect.

setup ssl cipherstrength

Controls the strength of ciphers that PacketShapers in the current configuration allows. *This command is applicable to legacy PacketShaper configurations only.*

```
setup ssl cipherstrength weak|strong|show
```

When this option is set to **strong**, PacketShaper does not allow ciphers that don't have authentication or encryption, nor does it allow ciphers that don't have at least a 56-bit encryption key. When this option is set to **weak** (the default), PacketShaper allows ciphers of all strengths, as well as ciphers with no encryption or authentication.

setup ssl secureReneg

Configure secure renegotiation for SSL/TLS, specifying whether to allow PacketShapers in the current configuration to communicate with patched and unpatched SSL clients and servers. Note that patched SSL clients and servers have fixed the SSL/TLS vulnerability and use OpenSSL v0.9.8m or higher; unpatched clients/servers use earlier OpenSSL versions. *This command is applicable to legacy PacketShaper configurations only.*

```
setup ssl secureReneg off|serverOnly|serverAndClient|default
```

off	Preserves legacy behavior: allows PacketShaper to communicate with all unpatched and patched SSL clients and servers. This is the default setting.
serverOnly	The PacketShaper, as an SSL server, supports renegotiation from patched clients only.
serverAndClient	Allows patched and unpatched SSL clients (browsers) to communicate with PacketShaper, and supports renegotiation from patched clients and servers only. Note: Renegotiation attempts from unpatched clients using SSLv3 protocol will fail, and PacketShaper will terminate the connection. Renegotiation attempts from unpatched clients using TLSv1 protocol will also fail, but PacketShaper will not terminate the connection; it will send the client a "no renegotiation" alert message.
default	Sets secure renegotiation back to its default setting so the configuration can inherit the setting from the parent configuration.

Notes:

- **ServerOnly** mode allows PacketShaper (SSL client) to communicate with SSL servers, such as Symantec's heartbeat and update servers, regardless of whether the servers are patched.
- Use the [setup ssl show](#) command to view the current secure renegotiation setting.

setup ssl show

Display the secure renegotiation setting for SSL/TLS in the current configuration. *This command is applicable to legacy PacketShaper configurations only.*

```
setup ssl show
```

setup syslog add

Add a Syslog server to the current configuration. The logging feature gives administrators a way to centrally log and analyze user events and system warning messages. For example, if you are using RADIUS authentication, each failed login attempt will be sent to the defined Syslog server.

Adaptive response action files and user events can be configured to send messages to a Syslog server. For example, when you register an event, you will be asked if you want to send events to Syslog; you can define and register an event that sends a message to a Syslog server when retransmissions rise to 30 percent of your network activity.

You can add up to four servers.

PolicyCenter CLI Commands

```
setup syslog add host:<ipaddress> [output:<facility>,<level>] [port:<portnum>]
[datetime]
```

host: <ipaddress>	The Syslog server IP address for example, <i>host:10.7.38.100</i>
output: <facility> ,<level>	<p>The facility and severity level for example, <i>output:local1,6</i></p> <p>Up to three outputs can be specified. The default facility is local4 and the default level is 7. PacketShaper user events are at severity level 6; if you want to capture them with Syslog, you must set the level to 6 or 7.</p> <p>See Facility Types and Severity Levels for lists of the valid facility types and levels.</p>
port: <portnum>	The port number of the Syslog server; if the port isn't specified, port 514 is used
datetime	Include the date and time in the message; the date and time are not included unless you specify the datetime parameter

For example:

```
setup syslog add host:10.7.38.100 output:local1,3 datetime
```

If you need to modify any of the settings later, you need to remove the server and then add it again (see [setup syslog remove](#)).

Messages are not sent until you enable the logging feature. See [setup syslog state on](#). If you want a PacketShaper event to be recorded in a Syslog, you need to specify this option when registering the event (see [event register](#)).

Facility Types

You can enter the keyword or value specified in the following table.

Description	Keyword	Value
Kernel	kern	0
User Processes	user	1
Electronic Mail	mail	2
Background System Processes	sysd	3
Authorization	auth	4
System Logging	sysl	5
Printing	lpr	6
Usenet News	news	7
Unix-to-Unix Copy Program	uucp	8
Clock Daemon	ckld	9
Security	sec2	10

Description	Keyword	Value
NTP Subsystem	ntp	12
Log Audit	audit	13
Log Alert	alert	14
Clock Daemon	clkd2	15
For Local Use	local0 local7	16-23

Severity Levels

You can enter the keyword or value specified in the following table. Set the level to specify which messages to suppress to the Syslog server. For example, setting the severity level to 3 allows messages with levels 0 3 and suppresses messages with levels 4 7. If you don't specify a severity level, 7 is used. With the default severity level, messages of all levels will get sent to the Syslog server.

Description	Keyword	Value
System unusable	emerg	0
Take immediate action	alert	1
Critical condition	crit	2
Error message	err	3
Warning message	warn	4
Normal but significant condition	notice	5
Informational	info	6
Debug message	debug	7

At the "warn" level, Blue Coat sends the following types of messages to the Syslog server:

- Login failed
- Hard drive status
- Measurement Engine status

User events that are configured to send a syslog message when a threshold is crossed are sent at the info severity level (6). See [event register](#) for more information on configuring an event to send a syslog message.

Adaptive response action files that include the **send syslog** command can designate the severity level at which the message is sent to the Syslog server; any level can be specified.

setup syslog rate

Set the maximum number of syslog messages that will be sent per second.

```
setup syslog rate <number>
```

The default rate is 20 messages per second and the valid range is 1-200. You may want to increase the rate if you are experiencing a problem with your unit.

setup syslog remove

Remove a syslog server from the current configuration. If you need to modify the settings of a server you have added, you will need to remove the server first.

```
setup syslog remove <ipaddress>
```

setup syslog show

Display the settings for Syslog servers defined in the current configuration.

```
setup syslog show [<ipaddress>]
```

If no <ipaddress> is specified, the setup of all Syslog servers is displayed. For example:

```
setup syslog show
```

Status: On		
Max Rate: 35		
Total Sent: 5		
Total Lost: 0		
Server Addr	Facility	Level

10.7.38.200	local4, 20	warn, 4
10.7.38.100	local4, 20	warn, 4

If you specify an <ipaddress>, the settings for a single Syslog server are displayed. For example:

```
setup syslog show 10.7.38.200
```

Server Addr: 10.7.38.100	
UDP Port: 514	
DateTime Option: Not Enabled	

Facility	Level

local4, 20	warn, 4

Message Format

When viewing the messages at the Syslog server, you will see the format of a Syslog message is as follows:

```
ReceiveDateTime address SendDateTime module-severity-MNEMONIC: description
```

ReceiveDateTime	The date and time the message was received by the Syslog server (may not be included, depending on the setup of the Syslog server)
address	The PacketShaper unit s IP address

SendDateTime	The date and time the message was sent to the Syslog server (if the datetime parameter was specified when defining the syslog server)
module	A four-byte string that identifies the type of message. For example, USRE is a user event and SYSW is a system warning.
severity	A single digit code (0 7) that reflects the severity of the condition; see Severity Levels
MNEMONIC	A code that uniquely identifies the error message for example, BAD_WR (bad write) or INSERT_F (insert into a list fails)
description	A text string describing the condition

Example message:

```
Aug 6 17:06:27 10.7.38.5 SYSW-4-LOG_WARN: Hard drive is down.
```

Or, if the datetime parameter was specified:

```
Aug 6 17:07:25 10.7.38.5 Mon Aug 6 17:05:01 2001 BST (London) SYSW-4-LOG_WARN: Hard drive is down.
```

setup syslog state

Enable or disable the logging feature so that messages will be sent to the syslog server(s) defined in the current configuration. If this command is issued from the command-line interface of an individual PacketShaper or a PolicyCenter sharable configuration, it will enable or disable the syslog servers defined for the selected configuration or unit.

```
setup syslog state on|off|default
```

Select the **default** option to set the logging feature on a PacketShaper to its default *off* state. To check whether the logging feature is on or off, use the [setup syslog show](#) command.



If the **setup syslog state default** command is issued from PolicyCenter for a *child configuration*, the selected configuration will discard its existing syslog state and inherit its syslog on/off setting from its parent configuration. If the **setup syslog state default** command is issued from PolicyCenter for a *root-level configuration*, the syslog state for that configuration will be returned to the default *off* setting.

setup tacacs acct

Set up or change the settings of the TACACS+ accounting service in the current configuration. This feature allows you to have an audit trail for user logins.

To define the TACACS+ accounting service, use:

```
setup tacacs acct primary|secondary {<host> <shared_secret> [<port>]}|delete|override
```

PolicyCenter CLI Commands

primary secondary	Enter the literal primary or secondary to indicate which server you are defining. (The secondary server is used when the primary server isn't accessible.)
<host>	The IP address or DNS name of the TACACS+ accounting server
<shared_secret>	The designated secret for the server; quotes are not required
[<port>]	The port number to access the server; if omitted, the default port 49 is used.
delete	Deletes the configuration of the primary or secondary server (whichever is specified)
override	For a PolicyCenter child configuration, create local TACACS+ settings that override the settings it inherits from its parent configuration. Remove these override settings at any time with the command setup tacacs acct primary secondary delete .

To turn the service on or off, or to return the service to its default **off** value, use:

```
setup tacacs acct on|off|default
```

Example:

```
setup tacacs acct primary 10.10.10.10 P4assw0rd1
setup tacacs acct secondary 10.10.20.10 Paa55w0rd2
setup tacacs acct on
```

This example defines a primary accounting server at 10.10.10.10 which has a shared secret of P4ssw0rd1, as well as a secondary server at 10.10.20.10. The third command line enables the TACACS+ accounting service. Once this service is configured and enabled, PacketShaper will send a PW_STATUS_START accounting message to the accounting server when a user logs in and a PW_STATUS_STOP message when a user logs off or is disconnected.

setup tacacs auth

Set up or change the configuration of the TACACS+ authentication service in the current configuration. TACACS+ authentication is an optional method for users to log into the PacketShaper browser interface, command-line interface, or customer portal or when FTPing to the unit. Using third-party TACACS+ servers enables you to have central configuration of user accounts.

```
setup tacacs auth primary|secondary {<host> <shared_secret> [<port>]}|delete|override
```

primary secondary	Enter the literal primary or secondary to indicate which server you are defining. (Note: The TACACS+ client uses the secondary server when the primary server isn't accessible or authentication failed.)
<host>	The IP address or DNS name of the TACACS+ authentication server
<shared_secret>	The designated secret for the server; quotes are not required

[<port>]	The port number to access the server; if omitted, the default port 49 is used
delete	Deletes the configuration of the primary or secondary server (whichever is specified)
override	For a PolicyCenter child configuration, create local TACACS+ settings that override settings inherited from its parent configuration. Remove these override settings at any time with the command setup tacacs primary secondary delete .

To turn the service on or off, or to return the service to its default **off** value, use:

```
setup tacacs auth on|off|default
```

Example:

```
setup tacacs auth primary 10.10.10.10 CupServ44
setup tacacs auth on
```

This example first defines a primary authentication server at 10.10.10.10 which has a shared secret of CupServ44. The second command line enables TACACS+ authentication service. Once this is configured and enabled, PacketShaper will prompt users for user name and password when they log into PacketShaper.

setup tacacs method

Select the TACACS+ authentication method to be used in the current configuration.

```
setup tacacs method ascii|pap|chap|mschap|default
```

- **ASCII** (*American Standard Code for Information Interchange*): With ASCII, the username and password are transmitted in clear, unencrypted text.
- **PAP** (*Password Authentication Protocol*): With PAP, the username and password are transmitted in clear, unencrypted text. ASCII or PAP authentication is required for TACACS+ configurations that require access to clear text passwords (for example, when passwords are stored and maintained in a database external to the TACACS+ server)
- **CHAP** (*Challenge Handshake Authentication Protocol*): In other environments, CHAP may be preferred for greater security. The TACACS server sends a challenge that consists of a session ID and an arbitrary challenge string, and the username and password are encrypted before they are sent back to the server.
- **MS-CHAP** (*Microsoft Challenge Handshake Authentication Protocol*): This protocol is very similar to CHAP, but with MS-CHAP authentication, the TACACS+ server can store an encrypted version of a user password to validate the challenge response. Standard CHAP authentication requires that the server stores unencrypted passwords.

The default authentication method is **ascii**.



MS-CHAP v1 and v2 are supported. PacketShaper attempts authentication with MS-CHAP v2 first. If the remote server doesn't support v2 or if authentication is denied, PacketShaper re-attempts authentication with MS-CHAP v1.

setup tacacs show

Display the TACACS+ settings in the current configuration. Use this command to verify that TACACS+ authentication and accounting are enabled, to see the timeout setting, and to view configuration settings on each of the TACACS+ servers.

setup tacacs show

TACACS Setup values:

```
Tacacs Method    : ASCII
Authentication   : on
Accounting       : off
Timeout          : 10
```

TACACS Service records:

Type	Host	Port	Secret
auth1	192.21.18.190	49	test

setup tacacs timeout

Set the amount of time for PacketShaper to wait for a response from the TACACS+ server defined in the current configuration. If the server doesn't send a reply within the timeout period, the PacketShaper will disconnect and the authorization attempt will fail. The default timeout period is 10 seconds.

```
setup tacacs timeout <seconds>|default
```

where *<seconds>* is a value between 1 and 60 seconds.

In the example below, the timeout interval is 25 seconds; this interval applies to any configured TACACS+ server.

```
setup tacacs timeout 25
```

To return to the default timeout interval, use:

```
setup tacacs timeout default
```

setup variable

Change a default variable setting in the current configuration.

```
setup variable [<variable> <value>|default] | [-reset|-nd]
```

where *<variable>* is one of the variables listed below and *<value>* is the value you want to set the variable to. The default, minimum, and maximum values for each *<variable>* are listed in the table.



After changing a variable's setting, many variables require that you [reset](#) the unit in order for the change to take effect.

To reset all system variables to their defaults, use the **setup variable -reset** command. To reset a specific variable to its default, use the **setup variable <variable> default** command. To see a list of all variables that have non-default settings, use the **setup variable -nd** command.

Variable/ Description	Default Value	Min. Value	Max. Value
autoCreateSameSide When this variable is enabled, the SameSide class is created automatically. When disabled, the SameSide class will not be auto-created. You may want to disable this variable if traffic is being misclassified into the SameSide class.	1 (on)	0 (off)	1 (on)
Bridge PassThru With Bridge PassThru enabled, the PacketShaper forwards packets that have a source and destination MAC address on the same side of the unit. When Bridge PassThru is disabled and traffic shaping is enabled, the PacketShaper drops packets that have source and destination MAC addresses on the same side.	1 (on)	0 (off)	1 (on)
DiffservClassSortPref Controls the sort order of the traffic tree, with respect to Diffserv classes (those with DSCP marks). Three settings are available: 0 Diffserv classes are sorted below IP-address-based classes, but above port-based classes (the default). 1 Diffserv classes are sorted above IP-address-based classes 2 Legacy sort order (Diffserv classes are sorted after IP-address-based classes, port-based classes, and auto-discovered classes) Note: The new sort order doesn't take effect until the unit is rebooted.	0	0	2
discoveryThresholdDynamicPort The number of new connections of an identifiable service to a port greater than 1024 that must be identified within a one-minute timeframe before PacketShaper creates a class	2	1	1000000
discoveryThresholdNonIP The number of new non-IP connections of a given type that must be identified within a one-minute timeframe before PacketShaper creates a class	2	1	1000000
discoveryThresholdNormal The number of new connections of an identifiable service to a port less than or equal to 1024 that must be identified within a one-minute timeframe before PacketShaper creates a class	1	1	1000000
discoveryThresholdPort The number of new connections to a particular port within a one-minute timeframe before PacketShaper creates a Port_#### class in the DiscoveredPorts folder It may be necessary to increase this value on Internet link deployments to prevent excessive number of DiscoveredPorts classes being created. If you don't want any Port_#### classes discovered, set this variable to its maximum value.	100	1	1000000

Variable/ Description	Default Value	Min. Value	Max. Value
discoveryThresholdUrlCategories The number of new flows belonging to a particular URL category that must be identified within a one-minute time frame before PacketShaper creates a class for the category	1	1	1000000
dynPtnActiveReuseSeconds The number of seconds a dynamic partition will be retained after an established flow has sent packets Note: If no other user needs a dynamic partition, the partition will be retained indefinitely.	300 (5 min)	10	7200 (2 hrs)
dynPtnIdleReuseSeconds The number of seconds a dynamic partition will be retained after an established flow has not sent or received packets Note: If no other user needs a dynamic partition, the partition will be retained indefinitely.	30	10	7200 (2 hrs)
dynPtnSequestrationCount The number of partitions reserved for static partitions; all other partitions can be used for dynamic or static partitions (applicable to PacketShaper 1200 and 1500 only)	3	0	99
enableCongestion Enable/disable the calculation of packet exchange time. When this variable is disabled, the Pkt Exch column on the Monitor Traffic page will not appear, RTM will not be available, and the packet exchange time and RTM measurement variables will always have a value of 0. After disabling the enableCongestion variable, you should reset the unit.	1 (on)	0 (off)	1 (on)
enableLatency Enable/disable the calculation of VoIP metrics. When this variable is enabled, PacketShaper collects data that measure packet loss, jitter, and latency for VoIP flows. Additional Information <ul style="list-style-type: none"> VoIP metrics can only be measured between PacketShapers with the VoIP metrics feature enabled. The VoIP metrics feature can measure traffic only from VoIP applications whose data is classified as RTP-I. For instance, latency metrics are not provided for DialPad, iChat, Vonage, and Skype. 	0 (off)	0 (off)	1 (on)

Variable/ Description	Default Value	Min. Value	Max. Value
enableSTUNclassification Enable/disable classification of the STUN (Session Traversal Utilities for NAT) protocol. Audio/video flows are first classified as STUN and then as RTP/RTCP. When enabled, PacketShaper will auto-discover the STUN class. If your network has a lot of STUN traffic, you may want to disable STUN classification to improve performance. Note: RTP/RTCP will still be classified even when STUN classification is disabled.	1 (on)	0 (off)	1 (on)
enableVoIPUseragentAutoDiscovery RTP auto-discovery is based on the VoIP user-agent attribute when this variable is enabled. When this variable is disabled (as it is by default), RTP auto-discovery is based on the RTP-I encoding attribute. When this variable is enabled and RTP-I is auto-discovered, it will auto-discover child classes based on VoIP user agent traffic (such as RTP-I-Motorola_VT1000 and RTP-I-Google_Talk).	0 (off)	0 (off)	1 (on)
enableWinnyClassification Enable/disable classification of the Winny service. For optimal performance, enable only when management of Winny traffic is required. Note: The Winny peer-to-peer application is used primarily in Japan.	0 (off)	0 (off)	1 (on)
flowRecordsIntermediateTimeout Number of milliseconds between generation and sending of intermediate flow detail records when traffic is present	1500	1000	36000
flowRecordsPktr0Timeout Number of seconds between generation and sending of Packeteer-0 flow records.	3600	10	5000
flowRecordsPktrPTimeout Number of seconds between generation and sending of Packeteer-P flow records.	60	10	5000
flowRecordsResetCounters Controls whether or not the counter fields in FDR packets are reset with each intermediate FDR sent Note: This variable only affects Packeteer-1 and Packeteer-2 format FDRs: counter fields are always reset in the NetFlow-5 format.	1 (on)	0 (off)	1 (on)
flowRecordsSendIntermediate Enable/disable the intermediate flow detail records feature. When this variable is enabled, PacketShaper emits intermediate FDRs at the interval specified by the flowRecordsIntermediateTimeout variable. Note: Enable the intermediate flow detail records feature only when using a suitably-instrumented collector, such as Cisco-based Netflow-5 collectors. IntelligenceCenter does not support intermediate FDRs.	0 (off)	0 (off)	1 (on)

Variable/ Description	Default Value	Min. Value	Max. Value
flowRecordsSendPktrP Enable/disable emission of Packeteer-P packets to Packeteer-1 and Packeteer-2 flow detail record collectors. Packeteer-P packets contain statistics that are not related to particular flows, but rather provide information about utilization on the PacketShaper at the time flows are recorded. If this variable is enabled, Packeteer-P records are sent after each UDP flow record packet is sent to Packeteer-1 or Packeteer-2 collectors (not more than once per minute).	0 (off)	0 (off)	1 (on)
flowRecordsSendPktr0 Enable/disable emission of Packeteer-0 packets to Packeteer-1 and Packeteer-2 flow detail record collectors. Packeteer-0 packets are mapping messages that allow collectors to decipher PacketShaper-related information in the FDRs they receive. For example, in the FDR's ClassID field, a value identifies the traffic class. In order for the collector to understand what class is actually associated with the ID, it uses the class map a list that contains each traffic class on the unit along with the identifying number assigned to each class. If this variable is enabled, Packeteer-0 mappings are sent out approximately once each hour. Note that this variable needs to be enabled only if the collector does not know this information through other means.	0 (off)	0 (off)	1 (on)
graphTimeoutSeconds The maximum number of seconds a graph can take to generate in the browser interface; if the graph takes longer to generate than this value, a system timeout error message will appear. Note: Increasing this setting can make the browser interface appear to "freeze" while PacketShaper is generating some of the more complex graphs. Sometimes the browser will not display the page until all of the graphs are generated.	60	1	600 (10 min)
hostTspecCacheInside Enable/disable caching of IP address-based classes on the inside. Change this setting to outside (0) to increase performance of classification if the majority of IP addresses in manually created classes are on the outside, rather than the inside. To disable the caching of inside IP address-based classes, use the setup variable hostTspecCacheInside 0 command. After you reset the PacketShaper, IP address-based classes will be cacheable on the outside. To re-enable caching for inside classes, use the setup variable hostTspecCacheInside 1 command.	1 (inside)	0 (outside)	1 (inside)

Variable/ Description	Default Value	Min. Value	Max. Value
httpStealth503 Control the display of the 503 - Service unavailable server error message when a connection is refused because of admission control (such as a never-admit policy). 0 The 503 - Service unavailable message will be customized with the text This message is sent by Blue Coat PacketShaper. 1 The PacketShaper text is not displayed with the 503 - Service unavailable message. 2 PacketShaper performs a TCP reset and drops the HTTP request; the error message will likely be The attempt to load <i>http://...</i> failed.	0	0	2
ipUserCacheNegativeTTL The number of seconds an IP will be stored in the PacketShaper cache, when the IP lookup does not result in a user.	1800	300	86400
ipUserCachePositiveTTL The number of seconds an IP-user name mapping will be stored in the PacketShaper cache. By default, the user mappings are stored for one hour. Because querying the cache is faster than querying the BCAA server, you can accelerate user name look ups by increasing the cache timeout. However, the tradeoff is that stale mappings could cause incorrect user name identification.	3600	300	86400
latencyProbeDiscard Allows the PacketShaper to be configured to discard VoIP latency probes after responding. If VoIP devices located on the Inside of the PacketShaper are sensitive to VoIP latency probes, enabling this variable will prevent potential VoIP call drops.	0	0	1
LFNSupport When enabled, this setting improves performance on Long Fat Networks (LFN) which require larger TCP window sizes. An LFN is a long distance network with large bandwidth and long delay; for example, high-capacity satellite channels are LFNs.	0 (off)	0 (off)	1 (on)
linkOverheadBytes Number of bytes that are added to each packet to account for WAN protocol header overhead	0	0	256
linkOverheadPpt Number of parts per thousand* by which packet sizes are increased to account for link overhead. This adjustment is useful for links that do bit stuffing. (<i>Bit stuffing</i> is the practice of adding bits to a stream of data. Bit stuffing is required by many network and communications protocols, for example to prevent data from being interpreted as control information.) * to be more precise, it s actually parts per 1024	35 (3.5%)	0	1024

PolicyCenter CLI Commands

Variable/ Description	Default Value	Min. Value	Max. Value
mirrorLinks Enable/disable link state mirroring. With link state mirroring, PacketShaper will bring down the second port of a NIC pair if the first goes down. This feature allows each PacketShaper to sit between a WAN router and a switch without blocking detection of switch outages by the router. This command is equivalent to the setup mirrorlink CLI command; these two commands are interchangeable.	0 (off)	0 (off)	1 (on)
PolicyFlowLimitForAllClasses Enables/disables the policy flow limit feature. When enabled, PacketShaper will enforce all policy flow limits that have been set on traffic classes. When disabled, all policy flow limits will be ignored. Disabled is the appropriate setting for PacketShapers deployed in proxy or NAT environments. For additional information, see policy flowlimit .	0 (off)	0 (off)	1 (on)
probeIntervalSeconds Number of seconds between the issuance of VoIP latency probes that measure VoIP metrics, enabled by the enableLatency variable.	5	1	60
rtolnboundClampMsecs Number of milliseconds delay for clamping early retransmission timeout on Inbound packets. Puts a maximum on retransmit time.	1600	0 (disable)	3000 (3 sec)
rtoutboundClampMsecs Number of milliseconds delay for clamping early retransmission timeout on Outbound packets.	1600	0 (disable)	3000 (3 sec)
syntheticReadTimeoutSeconds Number of seconds after which a synthetic transaction will end when the response received is incomplete	5	1	1000
syntheticWriteTimeoutSeconds Number of seconds after which a synthetic transaction will be canceled if the server fails to respond to a request	60	10	5000
tcpClipInitialWindow When tcpClipInitialWindow is enabled, the PacketShaper will always reduce the initial TCP window size to 1x MSS (maximum segment size). When this variable is disabled, new flows will ramp up faster but enforcement of small rate policies and/or partitions may not work at the beginning of flows.	1 (on)	0 (off)	1 (on)
tcpMssInbound Maximum segment size of TCP packets on Inbound flows. This setting can help avoid packet fragmentation when using VPN and not being able to support 1500-byte packets (the default size) through the VPN tunnel.	1460 bytes	0	65535
tcpMssOutbound Maximum segment size of TCP packets on Outbound flows	1460 bytes	0	65535

Variable/ Description	Default Value	Min. Value	Max. Value
tcpSmallMssLinkSpeed Link speeds slower than this value will force the use of smaller MSS (maximum segment size). Prevents PacketShaper from changing the MSS on large WAN links.	384000 bps	0	512000
trafficIsAsymmetric By turning on this setting, PacketShaper will automatically assume all flows are asymmetric and stop TCP Rate Control. In topologies where there are a large percentage of asymmetric flows, this may be more efficient than attempting to apply regular rate control. In addition to disabling rate control, turning on this setting disables all layer 7 classification activities (PacketShaper must see traffic in both directions in order to classify layer 7).	0 (off)	0 (off)	1 (on)
uiDefaultSky The user interface that appears after logging in to the browser interface: Blue Coat Sky or the original (legacy) user interface.	1 (Sky)	0 (legacy)	1 (Sky)
userEventExtSnmpVersion Enable/disable the extended SNMP trap for user events. When this variable is turned on, there will be an additional field in the trap that indicates the type of situation that triggered the trap. The field indicates <i>violated</i> (when the threshold was exceeded) or <i>rearm</i> (when the re-arm value was crossed).	0 (off)	0 (off)	1 (on)
userEventMaxDefinitions The maximum number of events that can be user-defined	32	32	128
userEventMaxRegistrations The maximum number of events that can be registered	32	32	128
userSessionIdleTimeout This variable controls how many seconds it takes for an unauthenticated login session to get purged from the system. You might need to increase this value if the session times out before PacketShaper can authenticate a login password, for example, when there is latency on the network or they are using a RADIUS or TACACS implementation. Note that this variable does not apply to idle sessions that have already been authenticated—just new sessions that have not yet been authenticated.	30	30	360 (6 min)

setup web-proxy

Enable/disable the web proxy feature in the current configuration. When this feature is enabled, the [configured web-proxy server](#) handles the WebPulse requests, category map downloads, heartbeat emissions, support status updates, and image updates. Because some PacketShaper features (such as URL categories) access external servers on the Internet, you must configure an explicit web proxy for the PacketShaper if your company's security policy requires that all outbound traffic go through a proxy. The topology should look like this:

LAN Switch-->PacketShaper-->Proxy--->Router-->Internet

```
setup web-proxy on|off
```

setup web-proxy default

Returns web proxy settings to their defaults so that they can be inherited from the parent configuration.

```
setup web-proxy default
```

setup web-proxy server

Configure the IP address and port of the web proxy server used in the current configuration. All PacketShaper features that access external servers on the Internet will go through the proxy server. This server handles WebPulse requests, category map downloads, heartbeat emissions, support status updates, and image updates.

```
setup web-proxy server <ip-address>:<port>
```

The *<ip-address>* of the web proxy server must be reachable by the PacketShaper.

Example:

```
setup web-proxy server 10.9.66.12:8000
```



You may need to add this server's IP address to the firewall's white list, if required.

setup web-proxy show

Display the web proxy settings in the current configuration.

```
setup web-proxy show
```

Sample output:

```
Web Proxy Settings
    Status : on
    Server IP : 10.9.66.12
    Server port : 8000
```

setup webpulse

Enables/disables the WebPulse enhanced classification feature in the current configuration. When WebPulse is enabled, the PacketShaper can classify by URL categories, web applications, and web operations. It is enabled by default.

```
setup webpulse on|off
```

When you enable WebPulse, you will have the ability to:

- create classes based on specific URL categories, and PacketShaper will then classify web traffic that corresponds to each of these categories into the appropriate class
- auto-discover category classes (you must first create an "all categories" class and enable class discovery)
- create classes based on specific web applications and operations
- auto-discover web application classes

When WebPulse is turned on, PacketShaper will send URL queries to the WebPulse cloud service; the service will look up the URL in its extensive database to find the categories, web applications, and/or actions associated with the URL, and send the response back to PacketShaper for classification. If WebPulse is disabled, the PacketShaper will look up the URL in its cache; if the URL is not in the cache, PacketShaper won't be able to classify the web traffic into a category.

The PacketShaper S500 has the option of using WebPulse Express, a feature that downloads the URL database to the appliance. When WebPulse is enabled and Express is configured, URL queries don't need to be sent to the WebPulse cloud service: they are looked up on the appliance, so classification is much faster. See [setup webpulse express](#).

setup webpulse cache-test

Looks up the specified URL in the PacketShaper's local cache and returns the category, web application, and operation names for the URL. Use this command to find out if a URL is in the cache or to see how a URL is being categorized.

```
setup webpulse cache-test <url>
```

where *<url>* is the URL to be looked up. Follow these guidelines when specifying the URL:

- The "www" is optional unless the server can host more than just Web, or if there is another subdomain you want to look up.
- Enter the URL exactly as it appears in your browser; you can copy the URL from the browser and paste it into the command window.
- Include https:// to do a lookup for a URL hosted on port 443.

```
# setup webpulse cache-test facebook.com
```

```
Category: Social Networking
```

```
Web application: Facebook
```

```
Operation: Unknown
```

```
# setup webpulse cache-testgroupon.com
```

```
No categorization forgroupon.com found.
```



If the command responds with *No categorization for <url> found*, the specified *<url>* is not in the cache. You can use the [setup webpulse cache-update](#) command to query WebPulse for the category.

setup webpulse cache-update

Sends a query to the WebPulse cloud service to look up the category, web application, and operation of the URL. The URL and its category ID, web application ID, and operation ID are then placed into the URL cache, overwriting any existing entry. If the [setup webpulse cache-test](#) command failed because the URL was not in the cache, you can use the **cache-update** command to find the category, web application, or operation. This command is also useful if you suspect a change in categorization, such as a recent malware notification.

PolicyCenter CLI Commands

```
setup webpulse cache-update <url>
```

where *<url>* is the URL to be looked up. Follow these guidelines when specifying the URL:

- The "www" is optional unless the server can host more than just Web, or if there is another subdomain you want to look up.
- Enter the URL exactly as it appears in your browser; you can copy the URL from the browser and paste it into the command window.
- Include https:// to do a lookup for a URL hosted on port 443.

If the command responds with *<<Unrated>>*, the specified *<url>* is not in the WebPulse database. Check your spelling and try again.

Examples:

```
# setup webpulse cache-update craigslist.org
```

```
Category: Shopping  
Web application: Craigslist  
Operation: Unknown
```

```
# setup webpulse cache-update youtube.com
```

```
Category: Open/Mixed Content, Audio/Video Clips  
Web application: YouTube  
Operation: Unknown
```

setup webpulse discovery

Enables/disables discovery of specific URL categories, or all categories in the current configuration. By default, discovery is enabled for each category. You may want to disable discovery of the categories that you aren't interested in monitoring or controlling. By doing so, you will not clutter your traffic tree with classes you don't need to track.

```
setup webpulse discovery on|off|inherit <category_name>|all
```

where *<category_name>* is the exact name of the category. To see a list of category names, use the [setup webpulse show categories](#) command. If the category name has a space, enclose the name in quotes.

If the unit is in shared mode, you can use the **inherit** option to inherit the category discovery settings from the parent configuration.

Example

To disable discovery of the Society/Daily Living category:

```
setup webpulse discovery off "Society/Daily Living"
```

To reenable discovery of all categories:

```
setup webpulse discovery on all
```


setup webpulse express

WebPulse Express initiates the download of the entire WebPulse database to a PS-S500 appliance, and regularly downloads incremental updates to the database to keep the local content current. By having a local copy of the database on the appliance, the process of URL categorization and web application identification is significantly faster.

```
setup webpulse express set {username | password | interval} <value>
```

Contact your Symantec sales representative for the **username** and **password** to use with WebPulse Express.

The **interval** defines how frequently the PacketShaper checks for updates to the WebPulse database. The interval value is a positive integer, defined in seconds; the default value is 60 seconds.

After WebPulse verifies your credentials, the download of the 500MB database will begin. Use the **setup webpulse express show** command to monitor the status of the download. During the download process, the output will say *Loading the Database*. When download is complete, it will say *Database is loaded*, indicating the date and time of the download. If there was a problem loading the database (for example, the user name and password weren't configured properly), the output will say *Database is not loaded*.

The WebPulse Express username, password, and interval must be specified in separate commands, as shown in the examples below.

Examples

```
setup webpulse express set username bcwfname
```

```
setup webpulse express set password *****
```

```
setup webpulse express set interval 120
```

Additional Information:

- The WebPulse Express feature is available on PS-S500 appliances with 64 GB of RAM. Some models may have sufficient memory already; others may require the purchase and installation of a memory upgrade kit. To see if your appliance has sufficient memory, issue the "version" on page 205 command.
- WebPulse Express is not available on PS-S200 and PS-S400 models.
- Symantec will email you a username and password for WebPulse Express. Until WebPulse Express is configured properly with the correct user name and password, WebPulse classification will still function but will use the database in the cloud.
- Use the [setup webpulse on](#) command to enable URL categorization and web application classification.

To see the status of WebPulse Express, use the following command:

```
setup webpulse express show
```

```
Database is loaded at 2015-01-16 14:33:02.
```

```
Current database version is 350160400.
```

```
Database will update every 60 seconds.
```

```
Username: BCWF-DEC2814
```

```
Password: *****
```

setup webpulse express-test

Looks up the specified URL in WebPulse Express database and returns the category, web application, and operation names for the URL. Use this command to find out if a URL is in the database or to see how a URL is being categorized.

```
setup webpulse express-test <url>
```

where <url> is the URL to be looked up. Follow these guidelines when specifying the URL:

- The "www" is optional unless the server can host more than just Web, or if there is another subdomain you want to look up.
- Enter the URL exactly as it appears in your browser; you can copy the URL from the browser and paste it into the command window.
- Include https:// to do a lookup for a URL hosted on port 443.

```
# setup webpulse express-test facebook.com
```

Category: Social Networking

Web application: Facebook

Operation: Unknown

```
# setup webpulse express-testgroupon.com
```

No categorization forgroupon.com found.



If the command responds with *No categorization for <url> found*, the specified <url> is not in the database. You can use the [setup webpulse cache-update](#) command to query WebPulse for the category.

setup webpulse map-download

Initiates a download of the latest WebPulse mapping files for identifying URL category names, web application names, and web operation names.

```
setup webpulse map-download
```

The WebPulse maps support the following features:

- **URL categorization:** The category map associates URL category names with numeric IDs. PacketShaper refers to the map to look up the category name after WebPulse has assigned a category ID to a flow.
- **Web application identification:** The web application map associates web application names with numeric IDs, and lists the operation IDs applicable to the application. PacketShaper refers to the map to look up the application name after WebPulse has assigned an application ID to a flow. In addition, PacketShaper uses the map to determine which operations are applicable to an application, for example, to create a class for a specific operation for a web application (such as Facebook-Post_Messages, Gmail-Upload_Attachment).

Additional Information

- When you enable WebPulse, the latest map files are automatically downloaded to your PacketShaper. Thereafter, PacketShaper automatically downloads the maps every day and after a device reset, regardless of whether the files have changed or not.
- You will see a *Success* message after the download completes. The map files are downloaded, replacing the current files (even if the file hasn't changed). The version number is also indicated.
- The map version is incremented when a category, web application, or operation has been added, renamed, or deleted.

setup webpulse reset

Clears the URL cache on the PacketShaper, deletes the cache backup, and returns all WebPulse settings to their defaults. If WebPulse Express is in use, the reset command removes the on-box database and clears the Express username and password. This command returns the feature to its factory default settings: WebPulse is enabled and discovery is enabled for all URL categories.

```
setup webpulse reset
```



Because this command deletes the cache as well as its backup copy in the PacketShaper 9.1026/urlcat/cache folder, use this command with caution. You will have an opportunity to confirm the reset after you enter the command.

setup webpulse show

Displays additional information about WebPulse statistics, caches, service points, and so forth.

```
setup webpulse show categories|statistics|cache|service|applications|operations
```

Option	Description
categories	Display URL category names, IDs, discovery state, and hits
statistics	Display WebPulse statistics: how many flows weren't processed due to load or because the flow ended, number and speed of queries to the WebPulse database, number and speed of queries that used the Dynamic Real-Time Rating (DRTR) service.
cache	<p>The URL caches contain URLs and their category, application, and operation IDs. This command displays statistics about the URL caches: number of hits and entries in the domain, directory, and filename caches. Also indicates when the cache was last backed up to the 9.1026/urlcat/cache folder.</p> <p>Examples of URLs that go in each type of cache:</p> <p>Domain cache: bluecoat.com Directory cache: nps.gov/yose Filename cache: cnn.com/forum/viewforum.php?f=1</p>

PolicyCenter CLI Commands

Option	Description
service	Display WebPulse service point information. For each of eight service points, the output lists the IP address, number of hits and speed of the WebPulse Rating Service (RS) database and Dynamic Real-Time Rating (DRTR) requests. Also indicates the health of the service points. Tip: The fastest servers appear at the top of the list.
applications	Display name and supported operations for each web application
operations	List all supported web operation names

Examples:

setup webpulse show statistics

Unprocessed Flows

```
Current Daily Unprocessed Due to Load: 0
Average Daily Unprocessed Due to Load: 0
Current Daily Unprocessed Due to Flow Ended: 89
Average Daily Unprocessed Due to Flow Ended: 102
```

Service Points

```
Current Daily Queries           : 161
Current Daily Query RTT         : 48 ms
Average Daily Queries           : 598
Average Daily Query RTT         : 51 ms
Current Daily DRTR Queries      : 3
Current Daily DRTR Query RTT    : 966 ms
Average Daily DRTR Queries      : 0
Average Daily DRTR Query RTT    : 0 ms
```

setup webpulse show cache

Cache

```
Daily Queries           : 1100
Average Daily Queries    : 2417
Cache Hits               : 926
Average Daily Cache Hits : 1736
Average Cache Efficiency : 71.83%
```

Domain Cache

```
Cache Hits           : 872
Average Daily Cache Hits : 1575
Number of Entries     : 575
Average Hourly Size    : 516
```

Directory Cache

```

Cache Hits           : 54
Average Daily Cache Hits: 161
Number of Entries     : 930
Average Hourly Size   : 863
Filename Cache
Cache Hits           : 0
Average Daily Cache Hits: 0
Number of Entries     : 0
Average Hourly Size   : 0
URL Cache Backup Status:
Cache backup successfully updated on Tue Aug  3 14:56:20 2010

```



To show details about WebPulse Express, use the [setup webpulse express show](#) command.

shutdown

Shut down the operating system and file system and power off the PolicyCenter S-Series. This command is used to prepare the appliance for transport.

```
shutdown
```

synthetic add

Create a new synthetic transaction in the current configuration. Using synthetic transactions allows PacketShaper to initiate ICMP, web, or other TCP transactions at periodic intervals to verify the availability of critical hosts.

```
synthetic add <interval>[,<repeat>] [<id>] <url>
```

<interval>	Number of minutes between issuance of the transaction (the maximum interval is 1440)
[,<repeat>]	Number of times to issue the request on the established TCP connection (default is 1; the maximum is 100)

PolicyCenter CLI Commands

<code><id></code>	String that identifies the synthetic transaction; if omitted, PacketShaper will automatically create a unique eight-character ASCII ID for each transaction. Note: Do not specify a transaction ID for a synthetic transaction on a PolicyCenter sharable configuration, as PolicyCenter requires that each synthetic transaction has a unique auto-generated transaction ID.
<code><url></code>	Type of transaction to issue, in the following format: <code><type>://<host>[:<port>][/<path>]</code> where: <code><type></code> is http , https , icmp , pop3 , smtp , ftp , echo , or custom Note: The <code><type></code> must be entered in lowercase. <code><host></code> is the IP address or DNS name of the host <code><port></code> is the TCP port number to connect to; the default varies by type (for example, for http the default is port 80) <code><path></code> is additional information necessary for the request (such as a folder or file name)

Additional information about each type:

- The **http** type will issue a GET request for the file specified by the `<path>` parameter. (The default port is 80.)
- The **https** type does an SSL handshake and issues a GET request for the file specified by the `<path>` parameter. (The default port is 443.)
- The **icmp** type sends a ping request to the designated server, using the ICMP-ECHO Protocol (RFC 792). Example:
`synthetic add 17,2 "icmp://www.bluecoat.com"`
- The **smtp** and **pop3** types also do not send or receive mail; they issue a single command over the channel to elicit a response. (The default port is 25.)
- The **ftp** type will issue a single retrieve command (RETR) for the file specified in the `<path>` parameter but doesn't do any user authentication. (The default port is 21.)
- The **echo** type sends a string to the designated host and the host echos it back. TCP echo requires that the target host have an echo server process running and listening on port 7. The optional `<path>` argument has the format `<length>[<fill>]` where `<length>` is the number of bytes to send on each request (the default is 512) and `<fill>` is a string to fill the request buffer. The `<fill>` string can be up to 511 bytes.

For example:

```
echo://test.domain.com/10/xyz
```

The above example sends requests containing xyzxyzxyzx (10 bytes).

- The **custom** type allows you to specify a series of requests to be sent alternatively for as many messages as requested by the `<repeat>` parameter. The request strings are separated by the `|` character. For example:

```
custom://my.test.com:25/HELO|MAIL FROM:<bob>|RCPT TO:<brett>|DATA|hey|.
```

The above example sends a simple message to a mail server on port 25 (the default port for SMTP).

synthetic delete

Delete a synthetic transaction in the current configuration.

```
synthetic delete <id>
```

where *<id>* is the identifying name specified in the **synthetic add** command. To view IDs of all synthetic transactions, use the [synthetic show](#) command.

When a synthetic transaction is deleted, its corresponding *Inbound/SyntheticTransactions* or *Outbound/SyntheticTransactions* traffic class is not deleted, so that measurement data can still be retrieved from that traffic class. Even after the synthetic transaction is deleted, network traffic may still be classified in that traffic class until the class is also manually removed.

synthetic options

Create traffic classes in the current configuration for the hosts specified in synthetic transactions. The classes will be created in the *SyntheticTransactions* class.

```
synthetic options create-classes show|on|off|default
```

The default value is **on**. If you have already created traffic classes for your critical hosts and you want synthetic transaction measurement data to be recorded in these classes, set this option to **off**.



If you use the CLI command **synthetic options create-classes on** to create traffic classes for synthetic transactions hosts and then later issue the **command synthetic options create-classes off** to turn off this option, any traffic classes already created for previous synthetic transactions will remain a part of the configuration's traffic tree.

When the **synthetic option create-classes default** command is issued for an individual PacketShaper or a PolicyCenter parent configuration at the top of the PolicyCenter configuration tree, this command restores the default **on** value. When this command is issued for a PolicyCenter child configuration, the child configuration clears its local setting and inherits the **synthetic option create-classes on|off** value from its parent configuration.

synthetic show

Display information about synthetic transactions in the current configuration.

```
synthetic show
```

The output displays all the active synthetic transactions, when they are next scheduled to run, and a count of how many TCP connections have been attempted and were accepted.

Transaction ID	URL	Interval	Repeat	Next Scheduled	Attempts	Connections
shop1	http://www.cdnnow.com	5	1	26-Jul-2001 14:30:38	0	0

PolicyCenter CLI Commands

st4	echo://10.10.10.10/monsters					
	1	5	26-Jul-2001	14:28:34	0	0
st2	http://www.lucent.com/minds/innovating/index.html					
	3	3	26-Jul-2001	14:30:00	6	5
st3	custom://my.test.com:80/HEY YOU					
	4	1	26-Jul-2001	14:31:17	5	4
st1	http://www.amazon.com					
	2	2	26-Jul-2001	14:30:04	9	9

ul add

Add entries to an existing user list in the current configuration. When specifying multiple names, separate each with a space.

```
ul add <user_list> u:<user> | g:<group> [u:<user> | g:<group> ...]
```

where *<user_list>* is an existing user list name, and *<user>* and *<group>* must be specified in the *<domain-name>\<user or group name>* format. To differentiate between user and group names, you must precede each name with **u:** or **g:**.

Example:

```
ul add list1 g:cal\group-marketing u:cal\john.smith u:cal\peter.hanson
```

ul delete

Remove one or more users or groups from an existing user list in the current configuration.

```
ul delete <user_list> u:<user> | g:<group> [u:<user> | g:<group> ...]
```

where *<user_list>* is an existing user list name, and *<user>* and *<group>* must be specified in the *<domain-name>\<user or group name>* format.

ul new

Create a user list in the current configuration by defining a unique name and specifying the user and group names that should be included in the list. When specifying multiple names, separate each with a space.

```
ul new <user_list> u:<user> | g:<group> [u:<user> | g:<group> ...]
```

where *<user_list>* is a descriptive name, up to 127 characters; the slash (/) and backslash (\) characters may not be used. The *<user>* and *<group>* must be specified in the *<domain-name>\<user or group name>* format. To differentiate between user and group names, you must precede each name with **u:** or **g:**.

To add entries to the user list after it's created, use the [ul add](#) command.

Examples:


```
ul new list2 u:ny\pharrison u:ny\speters g:ny\group-sales
ul new list3
```



The **ul new** command does not validate the existence of the user and group names.

ul override

For PolicyCenter/ Units in shared mode only

Override an inherited user list by creating a local copy of the list.

```
ul override <list-name>
```

where **<user_list>** is an existing user list name. You must make a local copy of an inherited user list before you can change the user list on the child configuration.

ul publish

For PolicyCenter/ Units in shared mode only

Publish a local or overridden user list.

```
ul publish <list_name>
```

where **<user_list>** is an existing user list name.

ul reformat

This command ensures v9.2.1 user lists are formatted properly after upgrading to v9.2.2 or higher.

```
ul reformat upgrade
```

The **upgrade** option adds **u:** before each user name in the list because a prefix is necessary to distinguish user names from group names in v9.2.2+. It performs the operation on all user lists at once.



If you fail to use the upgrade command, user list classification will still work properly. However, you will see an “invalid entries in the list” error if you ever try to modify the list, and you will need to manually add the **u:** prefix on each user name (or run the **ul reformat upgrade** command to fix all user lists).

ul rm

Remove a user list from the current configuration.

```
ul rm <user_list>
```

User lists cannot be removed if they are currently being used in a class matching rule.

ul show

Display a list of all user lists defined in the current configuration or show the details of a specific user list.

```
ul show [<list_name>]
```

To show all user lists, as well as the user and group names in each list:

```
ul show *
```

User and group names are listed alphabetically.


unit assign

Assign a unit to a different PolicyCenter configuration.

When a PacketShaper subscribes to PolicyCenter, a unique configuration for that unit is created in the configuration tree, then assigns the unit to that configuration. When you assign a unit to a sharable configuration, the unit remains attached to its individual unique unit configuration, so that unit configuration will appear in the configuration tree below the sharable parent configuration to which it is assigned. Because the unit is not directly assigned to a sharable configuration, changes made to the individual unit configuration will not affect its sharable parent configuration. The unit will, however, continue to inherit the settings from its sharable parent.

```
unit assign <unit_name>|<unit_sn> [cfg_path]
```

<unit_name>	The name of the unit
<unit_sn>	The serial number of the unit
<cfg_path>	The path of the unit's new configuration. If you omit this parameter, the unit will be assigned to the current active configuration.

 You cannot assign a unit to a draft configuration. To try a draft configuration on a unit, use the command [draft try](#).

unit clean

Deletes old unit status entries from the directory server, so that they will no longer clutter the **config show** command's output. You can specify the minimum age of entries to be deleted, where age is the number of seconds since the corresponding unit has reported status to the directory server. Does not delete any configurations.

```
unit clean [<age in seconds>]
```

unit details

Display detailed information about the selected unit, such as model number, serial number, IP address, image version, PolicyCenter information, and the unit's banner messages.

```
unit details <unit_name>|<unit_sn>
```

<unit_name>	The name of the unit
<unit_sn>	The serial number of the unit

```
unit details 085-10000215
```

```

Serial number      175-10010178
Unit name          7500-2
Model              7500
Directory Server   10.9.64.232
Status update age  189
Uptime             17 days
IP address         10.78.52.21
HTTPS port         443
Domain name        (unknown)
Configuration name /Legacy/175-10010178
Image Version Currently Running PacketWise v9.2.10g1 2015-02-20
Description        (none)

```

Installed Feature Keys:

```

acceleration      on
classes           N/A
compatibility      1
compression       on
control           on (off)
isp               2000
linksize          200m (1.5M/1.5M)

```

The *status update age* is the number of seconds that has elapsed since the unit confirmed its connection to PolicyCenter, while the *uptime* is the number of hours that the unit has reported a consistent connection. (After resetting the unit, the *uptime* will be 0.)

unit rename

Give a new name or assign a different name to a PacketShaper. You can identify the unit by unit name or serial number.

```
unit rename <unit_name>|<unit_sn> <unit_name>
```

<unit_name>	The name of the PacketShaper; the name can be up to 20 characters long and may include alphanumeric characters, dashes (-), underlines (_), and periods (.) and may not contain spaces.
<unit_sn>	The serial number of the PacketShaper

The following example identifies a unit by its serial number, and gives it the name *ShaperOne*.

```
unit rename 025-10000215 ShaperOne
```

unit show

Display information for all PacketShapers assigned to PolicyCenter, such as serial number, unit name, the name of its assigned PolicyCenter configuration, the unit's domain name and IP address.

```
unit show
```



Although unit names can be up to 20 characters long, the **unit show** output displays only the first 16 characters of the name.

unit sync

Show unit sync status. The output displays the amount of time since PolicyCenter last contacted each subscribed unit.

```
unit sync
```

Example output:

Serial Number	Configuration Name	P	C	Status	Age	Sync	Status...
014-10000000	/master1/legacy/014-10000000	P	C	290 secs	OK		
0310000000	/PolicyCenter	P	C	126 secs	na		
109-10000000	/109-10000000	P		35 days	mismatch	on	
/109-10000399							
135-10000000	/master1/legacy/135-10000000	P		7 days	mismatch	on	
/135-10005388							
175-10000000	/master1/legacy/175-10000000	P		7 days	mismatch	on	
/175-10010178							
175-10000000	/master1/legacy/175-10000000	P		7 days	mismatch	on	
/175-10017460							

Legend:

P: Configuration digests have been posted for the unit

C: Configuration digests have been confirmed by the unit

unit versions

Identify the software version of each unit assigned to PolicyCenter.

```
unit versions
```

For example:

```
unit versions
```

IP Address	Type	Version
10.9.64.193	STD	v9.2.11g1 2015-05-05
10.9.64.180	S400	11.5.2.1 build 169789

```

10.9.64.160      S200 11.5.2.1 build 169789
10.9.64.164      S200 11.5.2.1 build 169789
10.9.64.235      STD  v9.2.10g1 2015-02-03
10.9.64.110      S500 11.5.2.1 build 169789
10.9.64.208      S500 11.5.2.1 build 169789
10.9.64.151      S400 11.5.2.1 build 169789
10.9.64.150      S400 11.5.2.1 build 169789
10.9.64.209      S500 11.5.2.1 build 169789
10.9.64.183      S400 11.5.2.1 build 169789
10.9.64.184      S400 11.5.2.1 build 169789
10.9.64.182      S400 11.5.2.1 build 169789

```

uptime

Determine how long the PolicyCenter appliance has been up and running. It measures the time since the unit was booted, either from a power-up or a software reset.

```
uptime
```

version

Display the PolicyCenter software version that is currently running, build number (release ID of the active image), model, and serial number. Use the **verbose** option to list the part number, the inside and outside MAC addresses, and installed license keys.

```
version [verbose]
```

Example:

```

# version
    Version:          PolicyCenter 1.1.1.1 build 174115
    Product:          PolicyCenter S200
    Serial Number:    000000000
    Memory:           5.8GB RAM, 4GB System Disk total, 3.2GB System Disk
available
Copyright (c) 1996-2015, Blue Coat Systems, Inc. All rights reserved.

```

watch add

Add a router to the watch mode router list (the routers whose traffic is being monitored) in the current configuration.

```
watch add <name> <MAC address>
```

where

PolicyCenter CLI Commands

<name>	Description of router; up to 32 characters (no spaces are allowed, the only special characters allowed are colon, dash, underline, and period.)
<MAC address>	MAC address of the router, for example 08:00:20:c0:56:a6

Example:

```
watch add router1 08:00:20:c0:56:a6
```

A PacketShaper in watch mode can monitor traffic from up to 256 routers.

watch delete

Delete a router from the watch mode router list in the current configuration.

```
watch remove <name>
```

where <name> is the name that was defined when the router was added. To see a list of defined router names, use the [watch show](#) command.

Example:

```
watch remove router1
```

watch show

Display the watch mode settings in the current configuration. The output lists the current management port and configured watch mode routers.

```
watch show
```

Example:

```
watch show
```

Watch Mode Status: Enabled
Management Port: MGMT

Name	MAC Address
router6	00:10:7b:3c:30:39
router5	08:00:20:c0:56:a6
router4	00:03:e3:6b:46:c2
router3	00:03:e3:6b:46:c2
router2	01:02:03:04:05:06
router10	00:60:fb:60:1f:16

The following information is displayed in this screen:

Field	Description
Watch Mode Status	Indicates whether watch mode is currently enabled or disabled
Management Port	Indicates which port PacketShaper has determined will be used to manage the unit. The management port is not user-definable. PacketShaper decides which port to use for management access by checking which ports are connected.
Name MAC Address	The list of configured routers. Use the watch add command to configure routers and the watch delete command to remove them from the list.

