DDoS Managed Service Customer Guide

Version 7.0

Contents

Preface

Introduction

This guide includes instructions and information for managed services users of the MTS DDoS Protection 7.0 Web user interface.

Audience

This guide is intended for managed services users. The scope of the audience for this guide includes network security teams, system administrators, technical project managers, and anyone who uses Managed Services to access data and reports in the Arbor system.

Support

MTS is your primary point of contact for all service and technical assistance issues.

In this section

This section contains the following topics:

Conventions Used in this Guide

Introduction

This guide uses typographic conventions to make the information in procedures, commands, and expressions easier to recognize.

Conventions for procedures

The following conventions represent the elements that you select, press, and type while following procedures.

Typographic conventions for procedures

Convention	Description	Examples
Bold	An element on the graphical user interface.	Type the computer's *address* in the **IP Address** box. Select the **Print** check box. Click **OK**.
SMALL CAPS	A key on the keyboard.	Press ENTER. To interrupt long outputs, press CTRL + C.
`Monospaced`	A file name, folder name, or path name. Also represents computer output.	Open the `readme.txt` file. Expand the `Addresses` folder. Navigate to the `C:\Users\Default\Favorites` folder.
`Monospaced bold`	Information that you must type exactly as shown.	Type **`https://`** followed by the *`IP address`*.
`Monospaced italic`	A file name, folder name, path name, or other information that you must supply.	Type the server's *`IP address`* or *`hostname`*.
>	A navigation path or sequence of commands.	In the Web UI, select **Mitigation > Threat Management**. Navigate to the Alerts Ongoing page (**Alerts > Ongoing**).

Conventions for commands and expressions

The following conventions show the syntax of commands and expressions. Do not type the brackets, braces, or vertical bar in commands or expressions.

Typographic conventions for commands and expressions

Convention	Description
`Monospaced bold`	Information that you must type exactly as shown.
`Monospaced italic`	A variable for which you must supply a value.
[ ] (square brackets)	A set of choices for options or variables, one of which is required. For example: [*`option1`* | *`option2`*].
{ } (braces)	A set of choices for options or variables, any of which is optional. For example: {*`variable1`* | *`variable2`*}.
| (vertical bar)	Separates the mutually exclusive options or variables.

Chapter 1:

About the Managed Services Web UI

Introduction

This section describes how to log on and navigate in the Managed Services Web user interface (UI). It also describes how to commit configuration changes and how to use selection and FCAP wizards.

In this section

This section contains the following topics:

Logging On and Off

Introduction

Follow the procedures in this topic to log on and off.

Initial login steps

To log on, follow the steps below, based on your user group:

User group	Steps
administrator	1. Log on using the administrator name and password that your service provider gave you. 2. Change your password for security purposes. See "Editing Your User Account" on page 27. 3. Create user accounts.
user	1. Log on using the user name and password that your administrator gave you. 2. Change your password for security purposes. See "Editing Your User Account" on page 27.

Before you begin

Before you log on to Peakflow SP, set your browser preferences to allow pop-ups and accept cookies from Peakflow SP.

Logging On:

Important: You must use a secure connection to access the User Interface.

To log on:

1. Open your Web browser.
2. Type **`https://`** ddos.mts.ca
3. If applicable, select the appropriate option for accepting the site's certificate, and then click **OK**.
4. Type your *`user name`* and *`password`*.
5. Click **Login**.

Logging Off:

To log off:

- In the upper-right corner of any page in the Web UI, click **Log Out**.

Navigating the Peakflow SP Web UI

Introduction

You can navigate the DDos Web UI menus and pages using a variety of navigation controls.

About the Web UI menu bar

The Web UI menu bar displays the current date and time, indicates which menu is active, and allows you to navigate the Web UI menus and pages.

The Web UI is divided into the following menus:

Menu	Description
Status	Displays summary information about the state of Peakflow SP.
Alerts	Allows you to view the alerts detected by Peakflow SP.
Traffic	Allows you to view pre-defined reports about traffic data from different perspectives.
Mitigation	Allows you to view and configure mitigations in Peakflow SP.
Administration	Allows you to configure and maintain the Peakflow SP system.

You can hover the mouse pointer over a menu item to view that item's submenus.

Note: The menus that are available depend on a user's account group. Your service provider configures the account groups.

About the Arbor Smart Bar

The Arbor Smart Bar is a collection of icons that can appear to the left of the Help button. The number of icons that appear depends on the page that you are on.

The following icons can appear on the Arbor Smart Bar:

Icon	Description
	Hover over this icon to display the download options. The download options that appear vary depending on the page that you are on and can include the following: ■  PDF - Click to download the page in PDF format. ■  XML- Click to download the page in XML format. ■  CSV - Click to download a page in CSV format. The download can be a CSV text file or CSV zip archive file. When the download is a zip archive file, “zip archive” is appended to CSV. ■  Excel-XML - Click to download a page in Excel-XML format. The  icon appears only if there are multiple options for downloading a page.
	Click to download a page in PDF format. This icon appears only when the PDF format is the only download option for a page.
	Click to download and email a page as a PDF.

Sorting data tables

You can sort most tables by certain columns. The system displays column headings as links (underlined text) to enable table sorting by column. You can recognize the way in which a column is sorted by the up or down arrow that appears next to the column header. Peakflow SP sorts columns by default in the Web UI as follows:

- Columns that contain alphabetical lists are initially sorted in alphabetical order, from A-Z. Click an alphabetical column header to re-sort the table by that column in reverse order (Z-A).
- Columns that contain numerical lists are initially sorted in ascending order. Click a numerical column header to re-sort the table by that specific column in reverse (descending) order.

Note: By default, the Importance column on alert pages is sorted first by the severity level and then by the severity percent.

See “About the layout of the alert listing pages” on page 61.

Navigating multiple pages

Data is often displayed in tables that continue on multiple pages. In these cases, Peakflow SP displays at the bottom of the page the current page number in a text box followed by the total number of pages. You can use the following links to help you navigate among multiple pages:

Link	Description
One arrow pointing left (<)	Displays the previous page.
Two arrows pointing left (<<)	Displays the first page.
One arrow pointing right (>)	Displays the next page.
Two arrows pointing right (>>)	Displays last page.

To navigate directly to another page, you can type its page number in the text box and then press ENTER.

Resizing frames

To resize a frame of network objects on an administrative page so that it fits your browser window:

- Click (maximize) below the frame.

Frames are maximized by default in Peakflow SP.

Viewing status messages

Peakflow SP displays status messages in a box at the top of the Web UI page.

Select one of the following steps:

- To view the status message, click (expand) or **EXPAND**.
- To hide the status message, click (collapse) or **COLLAPSE**.

Committing Configuration Changes

Introduction

When you make a configuration change, you must “commit” it in order for the change to go into effect. You can commit configuration changes on the Configuration Commit page (**Administration > Commit Configuration**) or from any page in the Web UI.

User access

Only managed services administrators can commit configuration changes.

Committing configuration changes

To commit configuration changes:

1. Do one of the following:
   - Click the **Config Commit** button in the upper-right corner of the Web UI page.
   - Navigate to the Configuration Commit page (**Administration > Commit Configuration**).
2. (Optional) Type a *`log message`* to describe the changes.
3. Click **Commit**.

Using Selection Wizards

Introduction

Throughout the Peakflow SP Web UI are various selection wizards that you can use to select objects. In general, all wizards function similarly.

Using a selection wizard

To select an object using a selection wizard:

1. (Optional) From the **Group** list, select an option.
2. (Optional) In the **Name Regexp** box, type a *regular expression* and then click **Filter**.
3. Choose one of the following steps, and then click **Select**:
   - To add an object, select it in the Available Choices pane, and then click the down arrow to move it to the Selected pane.
   - To delete an object, select it in the Selected pane, and then click the up arrow to move it to the Available Choices pane.

Using the FCAP Wizard

Introduction

The fingerprint expression language is an extended version of the standard fingerprint expression language used by programs, such as tcpdump, to describe layer 2/3 traffic information. The FCAP Wizard helps you to add filtering criteria to a fingerprint expression.

An Open FCAP Wizard button appears whenever you can use the wizard to configure a fingerprint expression.

Using the FCAP wizard to configure a fingerprint expression

To configure a fingerprint expression using the FCAP Wizard:

1. Click **Open FCAP Wizard**.
2. Configure the settings in the FCAP Wizard window.

   Note: The settings that appear in the FCAP Wizard depend on the object you are configuring.
3. Click **Add** or **Add to Fingerprint**.
4. Click **Close**.
5. To add additional fingerprint expressions, repeat this procedure.

   For details about the settings, see "FCAP Wizard settings" below.

FCAP Wizard settings

The FCAP Wizard contains the following settings.

Note: The settings that appear in the FCAP Wizard depend on the object you are configuring.

Setting	Description
Source addresses box	Type one or more *source CIDR addresses*.
Source ports box	Type one or more *source TCP port numbers*.
Destination addresses box	Type one or more *destination CIDR addresses*.
Destination ports box	Type one or more *destination TCP ports numbers*.
Protocols box	Type one or more *protocol names* or *protocol numbers*.
Types of service box	Type one or more *types of service bits*. The ToS bits are as follows: ■ D - Minimizes delay ■ T - Maximizes throughput ■ R - Maximizes reliability ■ M - Minimizes monetary cost In some router implementations, this bit is labeled *C*, for cost.

Setting	Description
Average packet lengths box	Type one or more *packet lengths* or ranges of lengths.
TCP Flags boxes	For each type of TCP flag, select **on** or **off**. If you do not make a selection for a TCP flag, Peakflow SP ignores it.
Router list	Select a router to add to the fingerprint.
Input Interfaces (SNMP ID) box	To select input interfaces, click **Select Interfaces** and then, in the Router Interfaces window, click the name links of one or more interfaces. The interfaces that are available depend on your selection in the Router list. The selected interfaces appear in the Input Interfaces (SNMP ID) box.
Output Interfaces (SNMP ID) box	To select output interfaces, click **Select Interfaces** and then, in the Router Interfaces window, click the name links of one or more interfaces. The interfaces that are available depend on your selection in the Router list. The selected interfaces appear in the Output Interfaces (SNMP ID) box.
ICMP Type list, **ICMP Type** box	Select an ICMP type from the **ICMP Type** list or type an *ICMP type* in the **ICMP Type** box.
ICMP Code box	Type an *ICMP code number*.

Chapter 2:

Configuring User Accounts

Introduction

This section describes how to configure user accounts.

In this section

This section contains the following topics:

About the User Accounts Page

Introduction

You can use the User Accounts page to create, edit, delete, and view detailed user account information. You can view all of the users on your network and configure user accounts on the User Accounts page.

You can access the User Accounts page at **Administration > User Accounts**.

For information about configuring user accounts, see “Configuring User Accounts” on page 24.

For information about the last login attempt of users, see “About the User Account Login Records Page” on page 29.

User access

Only managed services administrators can view this page.

User Accounts page

The User Accounts page displays the following information:

Icon	Description
	Hover over this icon to display the download options. The download options that appear vary depending on the page that you are on and can include the following: ■ PDF - Click to download the page in PDF format. ■ XML- Click to download the page in XML format. ■ CSV - Click to download a page in CSV format. The download can be a CSV text file or CSV zip archive file. When the download is a zip archive file, “zip archive” is appended to CSV. ■ Excel-XML - Click to download a page in Excel-XML format. The icon appears only if there are multiple options for downloading a page.
	Click to download a page in PDF format. This icon appears only when the PDF format is the only download option for a page.
	Click to download and email a page as a PDF.

Column	Description

About searching on the User Accounts page

You can use the **Search** box to search for user accounts. The following are some guidelines for using the **Search** box:

- You can enter search values with or without keywords.
- Keywords allow you to search on a specific column.

See "Acceptable search keywords and values for user accounts" below.

- When you enter a keyword followed by a value, do not put a space between the colon and the value that you enter.
- A space between search values creates an AND statement.
- A comma between search values creates an OR statement.
- Search values are case-insensitive.
- A match occurs when a search value matches any part of a text string.

Acceptable search keywords and values for user accounts

The following table lists the columns on the User Accounts page and the keywords and values that you can use to search on that column in the **Search** box:

Column to search on	Acceptable keywords and values	Examples
Name	■ **name:***user name*	■ name:account_group1
Real Name	■ **realname:***user real name*	■ realname:John Doe
Account Group	■ **group_name:***account group name* ■ **account_group:***account group name* ■ **account_group_name:***account group name*	■ group_name:system_admin ■ account_group:system_user ■ account_group_name:system_operator
Capability Level	■ **capability_level:** [**administrator** | **user**]	■ capability_level:user
Email	■ **email:***email address*	■ email:user@example.com
Device	■ **device_name:***appliance name*	■ device_name:global
UI Menu	■ **ui_menu:***UI menu*	■ ui_menu:default
Status	■ **disabled:disabled** ■ **status:disabled**	■ disabled:disabled ■ status:disabled

Configuring User Accounts

Introduction

You can configure user accounts on the User Accounts page.

You can access the User Accounts page at **Administration > User Accounts**.

User access

Only managed services administrators can view this page.

Choosing a secure and acceptable password

Passwords must meet the following criteria:

- Must contain at least 7 characters.
- Must not exceed the maximum length if an administrator has configured a maximum length
- Can include special characters, spaces, and quotation marks
- Cannot be all digits
- Cannot be all uppercase letters
- Cannot be all lowercase letters
- Cannot be only letters followed by only digits (for example, abcd123)
- Cannot be only digits followed by only letters (for example, 123abcd)
- Cannot consist of alternating letter-digit combinations (for example, 1a3A4c1 or a2B4c1d)

Adding and editing user accounts

To edit a user account:

1. Navigate to the User Accounts page (**Administration > User Accounts**).
2. Choose one of the following steps:
   - To edit a user, click the user name link on the User Accounts page.
3. On the **Account Configuration** tab, configure the user account settings.
   See “User account configuration settings” on the facing page.
4. Click **Save**, and then commit your changes.

For information on how users change their passwords, see "Editing Your User Account" on page 27.

User account configuration settings

Use the following table to configure the settings on the Account Configuration tab:

Setting	Description
Username box	Type a *unique name*. The user name must meet the following criteria: ■ Must contain from 1 to 31 characters, digits, or any combination of both ■ Can begin with and include uppercase and lowercase letters, digits, a period (.), an underscore (_), and an @ ■ Cannot begin with a hyphen but can include a hyphen ■ Cannot contain spaces You cannot change the user name in an existing account.
Real Name box	Type the user's *full name*.
Email Address box	Type the user's *email address* as a fully qualified domain name. For example, user@example.com.
Old Password for <user> box	(Existing account only) Type the *current password* for this user.
New Password box **Confirm New Password** box	Type a *new password*, and then re-type it to confirm it. For information about password criteria, see "Choosing a secure and acceptable password" on the previous page.
Account Group list	Select the account group to assign to this user. To filter the list, type any part of the account group name that does not include a space. The account group determines the user's level of system access. Only your service provider can configure user account groups.
Capability Level list	Select the capability level to assign to this user. The capability level is either user or administrator.

Setting	Description
Timezone list	Select the time zone in which the appliance is located.
UI Menu list	Select the Web UI menu skin to be displayed for this user. The menu skin you select determines what a user can see in the Web UI. Your service provider can configure the menus.

Disabling user accounts

To disable a user account:

1. Navigate to the User Accounts page (**Administration > User Accounts**).
2. Select the check boxes for the accounts that you want to disable, and then click **Disable**.

Deleting user accounts

To delete a user account:

1. Navigate to the User Accounts page (**Administration > User Accounts**).
2. Select the check boxes for the users that you want to delete, and then click **Delete**.

You cannot delete your own user account.

Editing Your User Account

Introduction

The Edit My Account page allows you to edit your user account settings.

You can edit the following settings in your user account:

- Password
- Real name
- Email address
- Timezone
- UI menu (administrators only)

Note: A local administrator can edit all of the settings on this page except the Appliance list, which can be configured only by your service provider.

Editing your user account

To edit your user account:

1. Navigate to the Edit My Account page (**Administration > User Accounts**).
2. On the **Account Configuration** tab, configure the user account settings.
   See "Your user account configuration settings" below.
3. Click **Save**, and then commit your changes.
   See "Committing Configuration Changes" on page 16.

Your user account configuration settings

Use the following table to configure your user account configuration settings:

Setting	Description
Username box	Displays your user name. A local user cannot edit this setting.
Real Name box	Type your *full name*.
Email Address box	Type your *email address* as a fully qualified domain name. For example, user@example.com.
Old Password for <user> box	To change your password, type your *old password*.
New Password box **Confirm New Password** box	Type a *new password*, and then re-type it to confirm it. For information about password criteria, see "Choosing a secure and acceptable password" on page 24.
Appliance list	Displays the appliances to which you are assigned. You can be assigned to a single appliance or to all appliances.
Account Group list	Displays the account group that is assigned to you. The account group determines your level of system access.

Setting	Description
Capability Level list	Displays the capability group that is assigned to you. The capability group determines the features in Peakflow SP that you can access.
Timezone list	Select your time zone.
UI Menu list	Displays the Web UI menu skin that is assigned to you. The menu skin determines which UI menu items are displayed. Only administrators can select UI menu skins.

About the User Account Login Records Page

Introduction

The User Account Login Records page displays information about the last login attempt by users. This list includes users that have been deleted. You can access the User Account Login Records at **Administration > Login Records**.

User Account Login Records page

The User Account Login Records page displays the following information:

Column	Description
Username	The user name of an account.
Last Login Location	The IP address from which a user last attempted to connect to Peakflow SP.
Last Login Time	The time at which a user last attempted to connect to Peakflow SP.
Login Failures	The number of times that a user last tried to log on but was unsuccessful. The number reverts to zero when a user successfully logs on.

About searching on the User Account Login Records page

You can use the Search box to search on the User Account Login Records page. The following are some guidelines for using the Search box:

- You can enter search values with or without keywords.
- Keywords allow you to search on a specific column.

  See "Acceptable search keywords and values for user account login records" on the next page.
- When you enter a keyword followed by a value, do not put a space between the colon and the value that you enter.
- A space between search values creates an AND statement.
- A comma between search values creates an OR statement.
- Search values are case-insensitive.
- A match occurs when a search value matches any part of a text string.

Acceptable search keywords and values for user account login records

The following table lists the columns on the User Account Login Records page and the keywords and values that you can use to search on that column in the Search box:

Column to search on	Acceptable keywords and values	Examples
Name	■ **name:***user name* ■ **username:***user name*	■ name:admin ■ username:user1
Last Login Location	■ **location:***IP address* ■ **login_location:***IP address*	■ location:10.0.0.1 ■ login_location:10.0.0.2
Login Failures	■ **count_last:***number of failures* ■ **login_failures:***number of failures* ■ **num_fails:***number of failures*	■ count_last:1 ■ login_failures:2 ■ num_fails:3

Chapter 3:
Configuring Profiles

Introduction

This section describes how to configure profiles and VPNs in the Managed Services Web UI.

User access

Only managed services administrators can configure profiles and only managed services VPN administrators can configure a VPN.

In this section

This section contains the following topics:

About Profiles

Introduction

Profiles are managed objects. Profile managed objects are administrator-configured network resources that Peakflow SP uses to sort, filter, and store traffic and flow data. Profile managed objects define what Peakflow SP protects.

A profile managed object is an arbitrary subset of your network or of another network. For example, you might create a profile managed object to monitor your DNS servers or a data center within your network. You can also create a profile managed object to monitor external services or providers, such as YouTube or an upstream ASN.

See “Configuring Profile Managed Objects” on page 35.

User access

Only managed services administrators can configure profile managed objects.

About naming profile managed objects

A profile managed object name can include up to 64 characters. Use the standard printable ASCII characters, except for the following characters:

- backslash (\)
- exclamation point (!)
- quotation mark (“)

About the Configure Profiles Page

Introduction

The Configure Profiles page (**Administration > Profiles**) list the names of profile managed objects and their match values. You can search for profile managed objects on this page by using the following tools:

- Search box
- Search wizard

See "About Profiles" on the previous page.

User access

Only managed services administrators can configure profile managed objects .

About searching on a Configure Profiles page

You can use the **Search** box to search for profile managed objects on the Configure Profiles page. The following guidelines describe how to use the **Search** box effectively:

- You can enter search values with or without keywords.

  See "Acceptable search keywords and values for profile managed objects" on the next page.
- Keywords allow you to search on a specific column.
- When you enter a keyword followed by a value, do not put a space between the colon and the value that you enter.
- A space between search values creates an AND statement. If a keyword is followed by more than one value, only the first value is associated with the keyword. For any additional values, the search looks for those values in the name, description, or tag fields of the profile managed objects. For example, if you type **`name:XYZ 123`**, then the search returns all occurrences of profile managed objects that have XYZ in their name and 123 in their name, description, or tag fields.
- A comma between search values creates an OR statement.
- A match occurs when a search value matches any part of a text string.
- You can use quotation marks (") to match a phrase. For example, to search for a profile managed object with "This is the Chicago office," you can type **`description:"Chicago office"`**.

Acceptable search keywords and values for profile managed objects

The following table lists the acceptable keywords and values that you can use to search in the **Search** box for profile managed objects:

Attribute to search by	Acceptable keywords and values	Examples
name	■ **name:***profile managed object name*	■ name:customer1
description	■ **description:***profile managed object description*	■ description:”chicago office”
match	■ **match:***profile managed object match value*	■ match:1.1.0.0/16

Configuring Profile Managed Objects

Introduction

Managed services administrators can define profile managed objects on the Configure Profiles page (**Administration > Profiles**). Profile managed objects can represent entities such as network servers, data centers, and upstream ASNs.

See "About Profiles" on page 32 and "About the Configure Profiles Page" on page 33.

User access

Only managed services administrators can configure profile managed objects.

Adding and editing a profile managed object

To add or edit a profile managed object:

1. Navigate to the Configure Profiles page (**Administration > Profiles**).
2. Choose one of the following steps:
   - To add a profile managed object, click **Add: Profile**.
   - To edit a profile managed object, expand its parent profile, and then click its name link.
3. On the Add Profiles page or Edit Profiles page, on the **Description** tab, configure the following settings:

Setting	Description
Name box	Type the *name* of the profile managed object. See "About naming profile managed objects" on page 32.
Description box	Type a *description* of the profile managed object.
Parent list	Select the profile managed object to be the parent of this profile managed object.

4. Click the following tabs and add or edit their the settings:

Setting	Description
Match tab	Allows you to configure the match settings for a profile managed object. See "Configuring match settings for a profile managed object" on page 37.
Profiled Router Detection tab	Allows you to configure profiled router detection settings for a profile managed object. See "Configuring profiled router detection settings" on page 42.
Host Detection tab	Allows you to configure host detection settings for a profile managed object. See "Configuring Host Detection for Profile Managed Objects" on page 50.

Setting	Description
Profiled Network Detection tab	Allows you to configure profiled network detection settings for a profile managed object. See "Configuring Profiled Network Detection for Profile Managed Objects" on page 56.

5. Click **Save**, and then commit your changes.
   See "Committing Configuration Changes" on page 16.

Deleting profile managed objects

To delete a profile managed object:

1. Navigate to the Configure Profiles page (**Administration > Profiles**).
2. Expand the object's parent profile.
3. Select the check boxes next to the profile managed objects to delete.
   Caution: Peakflow SP does not prompt you for confirmation before it deletes a profile managed object.
4. Click **Delete**.

Configuring Match Settings for Profile Managed Objects

Introduction

Match settings are used to define how Peakflow SP should associate traffic with profile managed objects.

You can use the **Match** tab to add or edit the match settings when you configure a profile managed object. See "Configuring Profile Managed Objects" on page 35.

Configuring match settings for a profile managed object

To configure match settings for a profile managed object:

1. Navigate to the **Match** tab of the profile managed object.

   See "Adding and editing a profile managed object" on page 35.
2. From the **Match 1** list, select a match type that defines the profile managed object.

   Peakflow SP displays the match settings that you can configure for the profile managed object. For more information about each match type, see "About match types" on the next page.
3. Complete the next steps based on the match settings that you want to configure:

Match Type	Procedure
None	Go to Step 4.
CIDR Blocks	In the **Match Values** box, type one or more *CIDR block prefixes*.>
CIDR Groups	a. Click **Edit CIDR Groups**. b. Do one of the following: • Type the *CIDR groups* in the CIDR Groups Wizard. • Browse to your file that contains a list of CIDR groups, and then click **Upload**. c. Click **Select**. **Note:** To open or save a file of the CIDR groups that are listed in the CIDR Groups Wizard, click **Download CIDR Groups**.
CIDR IPv6 Blocks	In the **Match Values** box, type one or more *CIDR blocks*.

4. Click **Save**, and then commit your changes.

About match types

Peakflow SP defines objects by name, match type, match values, and optional protocol and port filters. Peakflow SP supports the following match types:

Match Type	Description
CIDR Blocks	One or more IPv4 CIDR block prefixes with the form A.B.C.D/N. To separate multiple prefixes, use spaces. Peakflow SP treats all CIDRs in aggregate for traffic reports and DoS alert detection.
CIDR Groups	One or more CIDR block prefixes with the form A.B.C.D/N with the name you assign to the group and a semicolon (;). To separate multiple prefixes, use spaces. Peakflow SP performs the DoS profiled router detection independently for each CIDR group but reports the traffic data for all CIDRs as a whole.
CIDR IPv6 Blocks	One or more IPv6 CIDR blocks. To separate multiple blocks, use a comma (,) followed by a space. Peakflow SP treats all CIDRs in aggregate for traffic reports and DoS alert deletion. **Example:** 2001:DB8:FF00::/40, 2001:DB8:0000::/48

About Profiled Router Detection

Introduction

Profiled router detection identifies traffic rates on a router that exceed expected levels for a profile managed object. The traffic rate that Peakflow SP expects for a profile managed object is referred to as the baseline. A baseline is the learned traffic rate for a profile managed object. When Peakflow SP detects a profiled router anomaly, it gathers details about the anomalous traffic on the affected routers. When the traffic significantly exceeds the baseline for a sustained period of time, Peakflow SP triggers an alert.

For information about configuring profiled router detection, see “Configuring Profiled Router Detection for Profile Managed Objects” on page 42.

The severity thresholds for profiled router detection are applied on a per router basis for profiled router protocol alerts and on a per interface basis for profiled router bandwidth alerts.

This type of detection generates alerts for IPv4 traffic.

Profiled router detection terminology

An understanding of the following terminology is needed to configure profiled router detection:

- Baseline

  The expected or normal rate of traffic.

- Sensitivity threshold

  How far traffic must be above the baseline before it is considered anomalous.

- Profiled router latency period

  The length of time that traffic must remain above the sensitivity threshold before an alert is generated. It is also used to determine when an alert has ended. This value is a global setting that is configured by your service provider.

- Severity duration

  The length of time that traffic must exceed a given threshold before Peakflow SP escalates its severity level.

- Severity threshold

  A threshold that Peakflow SP uses to differentiate between medium and high alert severity. If traffic exceeds the severity threshold for the severity duration, then the alert is classified as high. If traffic exceeds this threshold but does not stay there for the severity duration, then the alert is classified as medium.

- Middle line

  A calculated value that is approximately 50% of the way between the sensitivity threshold and the severity threshold. It is used to differentiate between low and medium levels of severity.

- Ignore rate

  A traffic rate that must be exceeded before Peakflow SP generates an alert. The ignore rate is not affected by the baseline.

- Forced alert threshold

  A threshold that causes Peakflow SP to generate an alert when traffic exceeds it for the profiled router latency period. This threshold is manually configured.

About profiled router detection baselines

Baselines are learned traffic rates of normal traffic for a profile managed object. Each collector keeps a separate set of baselines for each profile managed object. Each collector compares real-time flow information with its stored baselines. A profiled router alert is generated when traffic is significantly above the baseline for a sustained period of time. The sensitivity threshold defines how far traffic must be above the baseline before it is considered anomalous.

For each profile managed object, data is collected per interface on the total traffic (bps and pps) and per router on the traffic for each IP protocol (bps and pps). From this data, baseline traffic rates are calculated using the average traffic rate from each of the following 30 minute periods:

- Previous 30 minutes
- Equivalent 30 minute period 24 hours ago
- Equivalent 30 minute period 7 days ago

When computing the baseline, the older information is weighted more heavily in order to reduce the effect of recent changes.

How Peakflow SP creates and classifies profiled router detection alerts

A profiled router detection alert is generated when traffic exceeds the baseline or a forced alert threshold for a sustained period of time.

See "About profiled router detection baselines" above and "About the use of forced alert thresholds" on the facing page.

Peakflow SP creates profiled router detection alerts and assigns their severity level based on the following conditions:

Severity	Conditions
low	An alert has a low severity level if the following conditions are true: ■ Traffic goes above the ignore rate. ■ Traffic goes above a forced alert threshold or the baseline plus the sensitivity threshold and stays there for longer than the profiled router latency period. ■ Traffic does not stay above the middle line for the severity duration. ■ Traffic never goes above a severity threshold.
medium	An alert has a medium severity level if the following conditions are true: ■ Traffic goes above the ignore rate. ■ Traffic goes above a forced alert threshold or the baseline plus the sensitivity threshold and stays there for longer than the profiled router latency period. ■ Traffic goes above the middle line and stays there for the severity duration or traffic goes above the severity threshold but does not stay there for the severity duration.
high	An alert has a high severity rate if the following conditions are true: ■ Traffic goes above the ignore rate. ■ Traffic goes above a forced alert threshold or the baseline plus the sensitivity threshold and stays there for longer than the profiled router latency period. ■ Traffic goes above a severity threshold and stays there for the severity duration.

About the use of forced alert thresholds

A forced alert threshold is a manually configured threshold for profiled router detection. If traffic exceeds this threshold for the profiled router latency period, then Peakflow SP generates an alert.

You can use a forced alert threshold to generate alerts instead of using the baseline. For example, with a profile managed object that has a fairly constant rate of traffic, you don't really need to use baselines to trigger alerts. You can then configure profiled router detection so that the forced alert thresholds trigger the alerts instead of the baselines. If you set the alert ignore rates to the same value as the forced alert thresholds, then alerts are generated only when the forced alert thresholds are exceeded.

You can also use forced alert thresholds with baselines to ensure that alerts are generated when traffic rates exceed certain thresholds. With a baseline, the rate of traffic that is required to generate an alert can increase over time. If you configure forced alert thresholds, then an alert is generated when a forced alert threshold is exceeded even when an alert would not be generated because of the baseline.

About automatic rate calculation for profiled router detection

When you enable profiled router detection for a profile managed object, you can configure the ignore rates and severity thresholds manually or you can enable automatic rate calculation. It is recommended that you use the automatic rate calculation whenever possible, because the calculated rates automatically adjust to changes in traffic patterns.

For more information about configuring automatic rate detection, see "Profiled router detection configuration settings" on page 43.

Automatic rate calculation is based on rate settings that you configure and the last 30 days of actual traffic of a profile managed object. When the calculated ignore rates and severity thresholds become available, they override your configured rates. Peakflow SP calculates rates every day at 00:35 GMT, 08:35 GMT, and 16:35 GMT.

Allow Peakflow SP to monitor the traffic of a profile managed object for at least 24 hours before you enable automatic rate calculation for that object. Peakflow SP can calculate rates in less time, but gathering a larger sample size of data ensures better accuracy.

When automatic rate calculation is enabled, the automatic rate calculation results for a profile managed object appear in a graph and tables in the Profiled Router Detection Configuration window. The graph displays traffic rates for the past 30 days and the current calculated ignore rates and severity thresholds.

Configuring Profiled Router Detection for Profile Managed Objects

Introduction

On the **Profiled Router Detection** tab, you can enable profiled router detection. After you enable profiled router detection, you can access the Profiled Router Detection Configuration window to configure settings that determine when an alert is triggered and the severity level that it is assigned. Separate incoming and outgoing traffic settings are provided because the rate of traffic in one direction might be significantly different than the rate of traffic in the other direction. For additional information about profiled router detection, see "About Profiled Router Detection" on page 39.

You can also enable and configure automatic rate calculations. It is recommended that you use the automatic rate calculations whenever possible. For more information about the calculations and their settings, see "About automatic rate calculation for profiled router detection" on the previous page.

Configuring profiled router detection settings

To configure profiled router detection settings:

1. Navigate to the **Profiled Router Detection** tab.

   See "Adding and editing a profile managed object" on page 35.
2. To enable profiled router detection, select the **Enable Profiled Router Detection** check box.
3. Click **Edit Profiled Router Configuration**.
4. Configure the settings in the Profiled Router Detection Configuration window.

   See "Profiled router detection configuration settings" on the facing page.
5. From the **Outgoing Detection** list, select one of the following settings:
   - **Default (Use Global Setting)**

     This setting uses the global setting for profiled router outgoing detection that is configured on the Configure Global Detection Settings page (**Administration > Detection > DDoS**). The global profiled router outgoing detection setting is set to **Disabled** by default.
   - **Always Enabled** or **Always Disabled**

     These settings allow you to enable or disable profiled router outgoing detection on a per profile managed object basis.
6. Click **Save**, and then commit your changes.

   See "Committing Configuration Changes" on page 16.

Profiled router detection configuration settings

Use the following table to configure the settings in the Profiled Router Detection Configuration window:

Setting	Procedure
Severity Duration box	Type the *number of seconds* that traffic must exceed a given threshold before Peakflow SP escalates its severity. For more information about how the severity duration is used to classify an alert's severity, see "How Peakflow SP creates and classifies profiled router detection alerts" on page 40.
Incoming Severity Thresholds and **Outgoing Severity Thresholds** boxes	Type the *severity thresholds* (in bps and pps). The severity thresholds are applied on a per router basis for profiled router protocol alerts and on a per interface basis for profiled router bandwidth alerts. For more information about how the severity thresholds are used to classify an alert's severity, see "How Peakflow SP creates and classifies profiled router detection alerts" on page 40.
Enable SNMP Link Rate Severity Calculation check box	Select if you want Peakflow SP to use the SNMP link rate of an interface as a severity threshold. Peakflow SP calculates the severity threshold based on the lower of the auto-configured or manually configured high severity rate and the SNMP link rate of the router interface on which the traffic was detected.
Incoming Forced Alert Thresholds and **Outgoing Forced Alert Thresholds** boxes	Type the *forced alert thresholds* (in bps and pps). For information on the use of forced alert thresholds, see "About the use of forced alert thresholds" on page 41. If traffic exceeds a forced alert threshold for the profiled router latency period, Peakflow SP generates an alert. The severity of the alert is then determined by the severity duration, the severity thresholds, and other factors. For more information about the classification of an alert's severity, see "How Peakflow SP creates and classifies profiled router detection alerts" on page 40. The forced alert thresholds are applied on a per router basis for profiled router protocol alerts and on a per interface basis for profiled router bandwidth alerts.
Incoming Alert Ignore Rates and **Outgoing Alert Ignore Rates** boxes	Type the *alert ignore rates* (in bps and pps). Traffic must exceed an ignore rate for an alert to be generated. If the ignore rates are the same as the forced alert thresholds, then the baselines are ignored when generating alerts. **Note:** Forced alert thresholds supersede ignore rates.

Setting	Procedure
Enable Automatic Rate Calculation check box	a. Select if you want to enable automatic rate calculation. b. Configure the Automatic Rate Calculation settings. See "Automatic rate calculation settings" below. c. Finish configuring the profiled router detection settings described in this table.
Interface Bandwidth Alerts, **Interface Packets Alerts**, and **All Protocols Alerts** lists	In the Detection Sensitivity Thresholds section, select the sensitivity thresholds for the different types of alerts. A low number results in more alerts and a high number results in fewer alerts. It is recommended that you select 3 as a starting point in a production environment. You can then adjust this setting to reduce or to increase the number of alerts that you receive in your deployment.

Automatic rate calculation settings

Use the following table to configure the automatic rate calculation settings for profiled router detection:

Setting	Procedure
Severity Percentile box	Type the *percentage* of normal traffic that you want Peakflow SP to use as a base value to calculate incoming and outgoing severity rates. Typical percentile values range from 95 to 98.
Severity Multiplier box	Type the *number* that you want to multiply with the severity percentile to calculate the high severity rate. **Example:** If the 95th percentile value for incoming traffic is 100 Mbps and the multiplier is 1.1, then the high severity threshold for that profile managed object becomes 110 Mbps.
Ignore Percentile box	Type an ignore *percentile* to calculate the ignore rate. The default value is 40. This means that 60% of the data points over the last 30 days are greater than the calculated trigger rate. It is recommended that you enter a value between 40 and 50.
Severity Rate Floor settings	Type the lowest *values* for which you want Peakflow SP to generate a severity rate, and then select the corresponding traffic units from the lists.
Ignore Rate Floor settings	Type the lowest *values* for which you want Peakflow SP to generate an ignore rate, and then select the corresponding traffic units from the lists.

See "About automatic rate calculation for profiled router detection" on page 41.

To finish configuring profiled router detection, see "Configuring profiled router detection settings" on page 42.

About Host Detection

Introduction

Host detection monitors the IPv4 and IPv6 traffic to a host on all monitored routers. Host detection can be configured to monitor the traffic of profile managed objects. It can also be configured to monitor traffic that is not associated with a profile managed object.

Host detection can trigger a standard DoS Host alert or a Fast Flood DoS Host alert. A standard DoS Host alert is triggered when the traffic on a monitored router towards a single host exceeds the configured threshold of an enabled misuse type for a specified time period. A Fast Flood DoS Host alert is triggered when large amounts of traffic toward a single host are detected for an enabled misuse type.

If excessive traffic is detected for multiple misuse types that are enabled, then a single alert is created instead of separate alerts for each misuse type. The alert includes each misuse type that had excessive traffic. See “Host detection misuse types” on page 48.

For information about the other types of detection, see “About Profiled Network Detection” on page 53 and “About Profiled Router Detection” on page 39.

Note: Withe Peakflow SP 7.0, host detection replaces misuse detection.

Host detection terminology

An understanding of the following terminology is needed to configure host detection:

- Trigger rate

  A traffic rate that must be exceeded before Peakflow SP generates an alert.

  The trigger rate is applied on a per router basis. The trigger rate accounts for all interfaces on the router.

- Host detection start latency period

  Defines how long traffic must be above the trigger rate before a host alert is generated. This value is a global setting that is configured by your service provider.

- Severity duration

  The length of time that traffic must exceed a given rate before Peakflow SP escalates the alert’s severity level. If the traffic exceeds 75% of the high severity rate for the severity duration, then the alert is classified with a severity of Medium. If the traffic exceeds the high severity rate for the severity duration, then the alert is classified with a severity of High.

- High severity rate

  A traffic rate that Peakflow SP uses to differentiate between medium and high alert severity. If traffic exceeds the high severity rate for the severity duration, then the alert severity is set to high. If traffic exceeds this rate but does not stay there for the severity duration, then the alert severity is set to medium.

  The high severity rate is based on the highest rate of traffic at the managed object boundary, network boundary, or an individual router.

  If traffic exceeds the high severity rate for at least one minute during the start latency period, then the alert is classified with a medium severity when it starts instead of a low severity.

- Host detection end latency period

  Defines how long traffic must be below the trigger rate before a host alert is ended. This value is a global setting that is configured by your service provider.

- Fast flood detection

  An option that can be enabled to trigger a host alert much faster when large amounts of traffic toward a single host are detected.

About host detection with profile managed objects

Host detection detects excessive rates of traffic toward a single host that matches a profile managed object. It detects the excessive rates of traffic for the host misuse types that are enabled. When you enable host detection for a profile managed object, you can configure the host detection settings or you can use the global host default settings. The global host default settings are configured by your service provider. See "Configuring Host Detection for Profile Managed Objects" on page 50.

When host detection is enabled for a profile managed object, and the host detection triggers an alert, the profile managed object is associated with the alert. A host alert has only one profile managed object associated with it, but it can have multiple misuse types associated with it.

About host global detection

Host global detection detects excessive rates of traffic toward a single host that does not match a profile managed object. It can also detects excessive rates of traffic for dark IP addresses that are configured by your service provider. It detects the excessive rates of traffic for the host misuse types that are enabled. Host global detection must be configured by your service provider.

Note: If a profile managed object matches a host, then host global detection does not monitor the traffic to that host even if host detection is disabled for the profile managed object.

When host global detection triggers an alert, a profile managed object with the name Global Detection is associated with the alert. The name Global Detection appears in the alert wherever the name of a profile managed object would appear for an alert triggered by host detection that is configured for a profile managed object.

About host detection with fast flood detection enabled

When you enable host detection for a profile managed object, you can also enable fast flood detection. When fast flood detection is enabled, host detection is able to detect large amounts of traffic toward a single host for the misuse types that are enabled. Fast flood detection can then trigger an alert in as little as 1 second if the traffic is high enough.

You can use fast flood detection with auto-mitigation to protect a target against a flood of traffic that lasts just a few minutes. You do this by enabling fast flood detection and then having your service provider configure auto-mitigation for the profile managed object. A sudden flood of traffic can then be mitigated very quickly before it is able to take down the target. Peakflow SP can also trigger fast flood host alerts for traffic that is not monitored by a profile managed object if your service provider has enabled global host fast flood detection. See "Configuring Host Detection for Profile Managed Objects" on page 50.

With fast flood detection, a host alert is triggered when Peakflow SP detects that the amount of traffic seen exceeds the amount of traffic that would be received in 60 seconds at the high severity rate. The following graphs illustrate the difference between standard host detection and fast flood detection:

Standard Host Detection

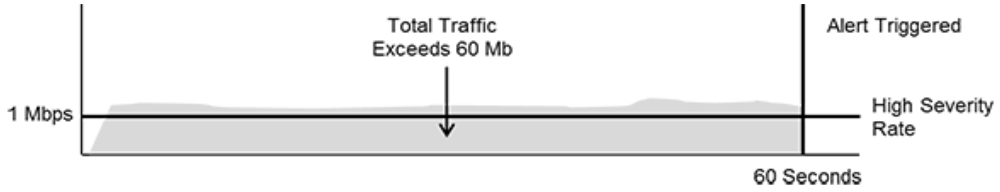

Fast Flood Detection

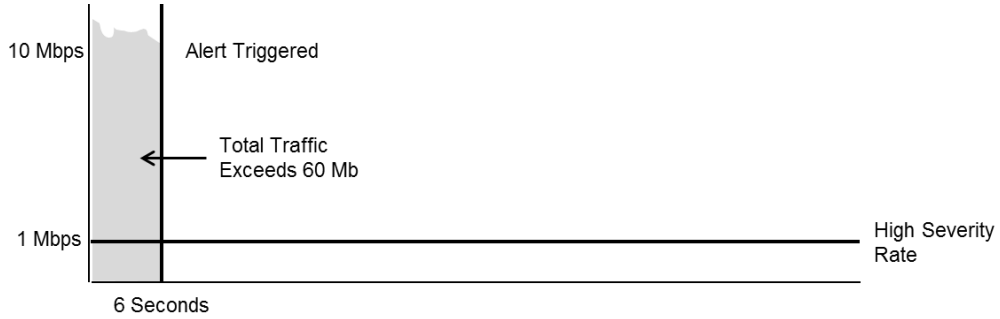

With standard host detection, Peakflow SP can trigger an alert only after 60 seconds of high traffic. With fast flood detection, if there is a large amount of traffic, an alert can be triggered after 1 or more seconds. If fast flood detection is enabled, but the amount of traffic seen does not exceed the amount of traffic that would be received in 60 seconds at the high severity rate, then Peakflow SP uses the standard host detection settings to determine if an alert should be triggered.

For example, the following table displays the results for different traffic rates when fast flood detection is enabled and the high severity rate is set at 1 Mbps :

Traffic Rate	Result
60 Mbps	Peakflow SP would trigger a fast flood host alert after 1 second.
10 Mbps	Peakflow SP would trigger a fast flood host alert after 6 seconds.
1 Mbps	Peakflow SP would not trigger a fast flood host alert, but it might trigger a standard host alert after 60 seconds.

When a fast flood host alert is triggered, the alert has a severity level of high and the severity level is followed by **Fast Flood**. See "About DoS Host alerts" on page 66.

Before you enable fast flood detection, you should be aware of the following fast flood detection limitations:

- It can trigger alerts when you have spikes in your legitimate traffic.
- It uses more system resources than standard host detection. If it is used extensively, it might impact your system performance.

- It classifies all of the alerts that it triggers with a severity of High when these same alerts might have a lower severity with standard host detection.

Host detection misuse types

Peakflow SP uses the following misuse types with host detection:

Misuse Type	Type of Traffic	Can Help Detect
DNS	Domain Name Server traffic (TCP and UDP port 53 traffic to a host)	This misuse type can help detect DNS amplification and reflection attacks and other DNS misuse.
ICMP	IPv4 Internet Control Message Protocol traffic	This misuse type can help detect smurf attacks.
ICMPv6	IPv6 Internet Control Message Protocol traffic	This misuse type can help detect modern variants of smurf attacks.
IP Fragment	Traffic with the IP fragment flag	This misuse type can help detect IP fragmentation attacks, such as the "ping of death" attack.
IP Null	Traffic with the protocol number set to 0	This misuse type can help determine how null routing is impacting an attack.
IP Private	Traffic for private IP address space	This misuse type can help detect spoofed IP addresses that are used in attacks.
TCP Null	TCP traffic that contains a sequence number but no flags	This misuse type can help detect “shrew” attacks.
TCP RST	TCP traffic with the reset flag set	This misuse type can help detect TCP reset attacks.
TCP SYN	TCP traffic with the synchronize flag set	This misuse type can help detect common TCP SYN flood attacks.
UDP	UDP traffic	This misuse type can help detect UDP attacks.
Total Traffic	The total traffic for a given host, for this profile managed object	This misuse type can help detect host attacks that do not follow a known attack pattern.

How Peakflow SP creates and classifies standard host alerts

A standard host alert occurs when the traffic that is sent to a host for a configured misuse type exceeds the configured trigger rate. For information about host alerts that are triggered with fast flood detection enabled, see "About host detection with fast flood detection enabled" on page 46.

Peakflow SP creates host alerts and assigns their severity level based on the following conditions:

Severity	Conditions
low	The alert severity is low if the traffic: ■ Exceeds the trigger rate for longer than the host detection start latency period. ■ Does not exceed 75% of the high severity rate for the severity duration. ■ Never exceeds the high severity rate.
medium	The alert severity is medium if the traffic: ■ Exceeds the trigger rate for longer than the host detection start latency period. ■ Exceeds the high severity rate for one minute or exceeds 75% of the high severity rate for the severity duration ■ Does not exceed the high severity rate for the severity duration.
high	The alert severity is high if the traffic: ■ Exceeds the trigger rate for longer than the host detection start latency period. ■ Exceeds the high severity rate and stays there for the severity duration.

The following are important things to know about host alert classification:

- When Peakflow SP initially classifies a host alert, the severity is based on traffic data from the router that has the highest rate of alert-triggering traffic.
- After a host alert is triggered, the traffic data that is used to classify the severity of the alert also includes data from the boundaries that are configured for the profile managed object.
- The severity of a host alert can increase, but it can never decrease.
- The traffic rate used for severity classification is gathered once a minute and is the average rate per second for the minute.
- If the severity duration is greater than 60 seconds, a host alert cannot have an initial severity of high because an alert will be triggered before the end of the severity duration.

For more information about the rates and time periods that control the host alerts, see "Host detection terminology" on page 45.

Configuring Host Detection for Profile Managed Objects

Introduction

Host detection monitors the IPv4 and IPv6 traffic to a host on all monitored routers. Host detection can trigger a standard host alert or a fast flood host alert. A standard host alert is triggered when the traffic on a monitored router towards a single host exceeds the configured threshold of an enabled misuse type for a specified time period. A fast flood host alert is triggered when large amounts of traffic towards a single host are detected for an enabled misuse type. See “About Host Detection” on page 45..

If excessive traffic is detected for multiple misuse types that are enabled, then a single alert is created instead of separate alerts for each misuse type. The alert identifies each misuse type that had excessive traffic. See “Host detection misuse types” on the facing page.

Configuring host detection for a profile managed object

You configure the host detection settings when you add or edit a profile managed object. The settings you configure determine when an alert is generated and the severity of the alert. See “How Peakflow SP creates and classifies standard host alerts” on the previous page.

To configure host detection for a profile managed object:

1. Navigate to the **Host Detection** tab of the profile managed object.

   See “Adding and editing a profile managed object” on page 35.
2. From the **Enable Host Detection** list, select one of the following:
   - **Use Default Host Settings** to use the default settings for host detection, and then go to Step 6.

     The default settings for host detection are configured by your service provider.
   - **Always Disabled** to disable host detection, and then go to Step 6.
   - **Always Enabled** to enable host detection, and then go to Step 3.
3. In the **Severity Duration** box, type the *`number of seconds`* that you want Peakflow SP to wait before it escalates the severity level of an alert.

   If the traffic exceeds 75% of the high severity rate for the severity duration, then the alert is classified with a severity of Medium. If the traffic exceeds the high severity rate for the severity duration, then the alert is classified with a severity of High.

   This setting is used only by standard host detection and is not used by fast flood host detection.
4. From the **Fast Flood Detection** option, select **Enabled** or **Disabled**. The default select is **Disabled**.

   Host fast flood detection triggers an alert much faster than standard host detection when large amounts of traffic toward a host are detected. See “About host detection with fast flood detection enabled” on page 46.

   Note: If you want a host alert that is triggered by fast flood detection to start an auto-mitigation, then you must have your service provider configure auto-mitigation for this profile managed object.

5. In the Misuse Types section, configure the following settings:

Setting	Procedure
Type column	Select the check boxes of the misuse types for which you want to configure a threshold for detection. See "Host detection misuse types" below.
Trigger Rate column	For each misuse type that you enable for detection, type the *trigger rate* in the **Trigger Rate** box, and then select the appropriate units from the **Trigger Rate** list. See "Host detection terminology" on page 45.
High Severity Rate column,	For each misuse type that you enable for detection, type the *high severity rate* in the **High Severity** box, and then select the appropriate unit from the **High Severity** list. See "Host detection terminology" on page 45. The high severity rate is applied on a per router basis. The high severity rate accounts for all profile managed object boundary interfaces.
Set to Default button	Click to reset the trigger rates and high severity rates to their default settings.

6. Click **Save**, and then commit your changes.

Host detection misuse types

Peakflow SP uses the following misuse types with host detection:

Misuse Type	Type of Traffic	Can Help Detect
DNS	Domain Name Server traffic (TCP and UDP port 53 traffic to a host)	This misuse type can help detect DNS amplification and reflection attacks and other DNS misuse.
ICMP	IPv4 Internet Control Message Protocol traffic	This misuse type can help detect smurf attacks.
ICMPv6	IPv6 Internet Control Message Protocol traffic	This misuse type can help detect modern variants of smurf attacks.
IP Fragment	Traffic with the IP fragment flag	This misuse type can help detect IP fragmentation attacks, such as the "ping of death" attack.
IP Null	Traffic with the protocol number set to 0	This misuse type can help determine how null routing is impacting an attack.
IP Private	Traffic for private IP address space	This misuse type can help detect spoofed IP addresses that are used in attacks.

Misuse Type	Type of Traffic	Can Help Detect
TCP Null	TCP traffic that contains a sequence number but no flags	This misuse type can help detect “shrew” attacks.
TCP RST	TCP traffic with the reset flag set	This misuse type can help detect TCP reset attacks.
TCP SYN	TCP traffic with the synchronize flag set	This misuse type can help detect common TCP SYN flood attacks.
UDP	UDP traffic	This misuse type can help detect UDP attacks.
Total Traffic	The total traffic for a given host, for this profile managed object	This misuse type can help detect host attacks that do not follow a known attack pattern.

About Profiled Network Detection

Introduction

Profiled network detection identifies excessive rates of network-wide traffic based on baselines that Peakflow SP has calculated for your network. Peakflow SP generates a profiled network alert if the rate of the traffic at a profile managed object boundary for one or more hosts exceeds the baseline by the detection percentage for a sustained period of time. When Peakflow SP generates a profiled network detection alert, it classifies the severity of the alert as low, medium, or high.

When Peakflow SP detects a profiled network alert, it gathers details about the alert traffic from across the entire network. The alert traffic details that Peakflow SP gathers are broader than the alert traffic details for profiled router detection. It combines all protocols for which attacks have been detected on the same profile managed object into one alert. It also provides the source ASN.

Profiled network detection terminology

An understanding of the following terminology is needed to configure profiled network detection:

- Baseline

  The expected or normal rate of traffic.

  See "About profiled network detection baselines" on the next page.

- Detection percentage

  The percentage above the baseline that the rate of traffic must reach before Peakflow SP can generate an alert. The traffic must maintain this rate for the profiled network start latency period before an alert is generated.

- Trigger rate

  A traffic rate that must be exceeded before Peakflow SP generates an alert. This rate is the baseline plus the detection percentage.

- Profiled network start latency period

  The length of time that the rate of traffic must exceed the trigger rate before Peakflow SP generates an alert. This value is a global setting that is configured by your service provider.

  Note: If the rate of traffic exceeds the baseline by the high severity percentage for at least a minute, an alert is generated even if the profiled network start latency period has not elapsed.

- High severity duration

  The length of time that the rate of traffic must exceed the baseline by a specified percentage before Peakflow SP classifies an alert as medium or high. An alert is classified as medium severity if the rate of traffic exceeds the baseline by at least 75 percent of the high severity percentage for the high severity duration. An alert is classified as high severity if the rate of traffic exceeds the baseline by the high severity percentage for the high severity duration.

- High severity percentage

  The percentage above the baseline that the rate of traffic must reach before Peakflow SP can classify an alert as medium or high. If the rate of traffic exceeds the baseline by the high severity percentage for at least one minute but for less than the high severity duration, the

alert is classified as medium. If the rate of traffic exceeds the baseline by the high severity percentage for the high severity duration, then the alert is classified as high.

- Ignore rate

  A traffic rate that must be exceeded before Peakflow SP generates an alert. The ignore rate is not affected by the baseline.

- Profiled network end latency period

  The length of time that the rate of traffic must remain below the trigger rate before Peakflow SP ends an alert. The profiled network end latency period is a global setting that is configured by your service provider.

  For more information about the ending of profiled network alerts, see "How Peakflow SP determines if a DoS alert should be ended or ongoing" on page 71.

About profiled network detection baselines

Baselines are learned traffic rates of normal traffic. Peakflow SP generates an alert for a profile managed object when the rate of traffic exceeds the baseline by a specified percentage (detection percentage) for a sustained period of time.

Peakflow SP starts collecting baseline data as soon as profiled network detection is turned on for that profile managed object. However, baseline data does not appear in the reports for about 24 hours. Baselines are updated every 30 minutes at 15 and 45 minutes past the hour.

How Peakflow SP creates and classifies profiled network alerts

Peakflow SP creates profiled network alerts and assigns their severity level based on the following conditions:

Severity	Conditions
low	An alert has a low severity level if the following conditions are true: ■ Traffic exceeds the ignore rate. ■ Traffic exceeds the baseline by the detection percentage and stays there for the profiled network start latency period. ■ Traffic exceeds the baseline by 75 percent of the high severity percentage and has a duration that is less than the high severity duration or the traffic exceeds the baseline by the high severity percentage and has a duration that is less than a minute and less than the high severity duration.
medium	An alert has a medium severity level if the following conditions are true: ■ Traffic exceeds the ignore rate. ■ Traffic exceeds the baseline by the detection percentage and stays there for the profiled network start latency period, or traffic exceeds the baseline by the high severity percentage for at least a minute. ■ Traffic exceeds the baseline by at least 75 percent of the high severity percentage for the severity duration. ■ Traffic does not exceed the baseline by the high severity percentage for the high severity duration.

Severity	Conditions
high	An alert has a high severity rate if the following conditions are true: ■ Traffic exceeds the ignore rate. ■ Traffic exceeds the baseline by the detection percentage and stays there for the profiled network start latency period. ■ Traffic exceeds the baseline by the high severity percentage and stays there for the high severity duration.

Configuring Profiled Network Detection for Profile Managed Objects

Introduction

Profiled network detection identifies excessive rates of network-wide IPv4 and IPv6 traffic based on baselines that Peakflow SP has calculated for your network. Peakflow SP generates a profiled network alert if the rate of the traffic at a profile managed object boundary for one or more hosts exceeds the baseline by the detection percentage for a sustained period of time. When Peakflow SP generates a profiled network detection alert, it also classifies the severity of the alert as low, medium, or high. See "About Profiled Network Detection" on page 53.

Configuring profiled network detection for a profile managed object

To configure profiled network detection for a profile managed object:

1. Add or edit a profile managed object.

   See "Adding and editing a profile managed object" on page 35.
2. Click the **Profiled Network Detection** tab.
3. Select the **Enable Profiled Network Detection** check box to enable profiled network detection.
4. Use the following table to configure the profiled network detection settings:

Setting	Procedure
Enable Profiled Country Detection check box	Select if you want to enable profiled country detection. If enabled, Peakflow SP generates alerts when the traffic from a country exceeds the baseline values for that country.
Incoming Detection Percent and **Outgoing Detection Percent** box	Type the *percentage* above the baseline that either incoming or outgoing traffic must be before Peakflow SP triggers the alert.
Severity Duration box	Type the *number* of minutes that an alert must exceed the severity threshold before Peakflow SP sets the alert to high severity. The severity rates are applied on a network wide basis.
Incoming Severity Percent and **Outgoing Severity Percent** boxes	Type the *percentage* above the baseline that either incoming or outgoing traffic must be before Peakflow SP sets the alert to high severity.
Incoming Ignore Rates and **Outgoing Ignore Rates** boxes	Type the traffic *rates* (in bps and pps) below which you do not want Peakflow SP to generate alerts.

5. Click **Save**.

Chapter 4:
About DoS Alerts

Introduction

This section describes how to investigate DoS alerts. Peakflow SP tracks the activity based on user-configured thresholds and can alert you to any anomalous activity in your network.

User access

Managed services administrators and non-administrative users have access to these features.

In this section

This section contains the following topics:

How Alerts Work

Introduction

You can use the alerts pages to detect DoS attacks on your network. Peakflow SP triggers DoS alerts based on detection settings. See "About Host Detection" on page 45, "About Profiled Network Detection" on page 53, and "About Profiled Router Detection" on page 39.

For information about the different alert pages, see "About the Alert Listing Pages" on the facing page.

For information about navigating the alerts pages, see "Navigating the Peakflow SP Web UI" on page 13.

How Peakflow SP uses samples to collect data for alerts

Peakflow SP aggregates data into statistically significant groupings, such as subnets and port ranges. Peakflow SP uses one-minute samples to collect the data for alerts. Flow records that match the alert are gathered from all Peakflow SP systems every 60 seconds. These flow records are parsed network-wide for the following information:

- ingress and egress interfaces
- protocols

Note: With a Fast Flood DoS Host alert, Peakflow SP bypasses this method of collecting data in order to trigger the alert more quickly.

Alert levels of importance

Peakflow SP assigns each alert one of the following levels of importance, based on its severity:

Importance	Color	Recommended action
High	Red	Address the alert immediately.
Medium	Orange	Analyze the alert to determine whether it is an attack.
Low	Green	Ignore if it is not worth your time to address them.

For a DoS alert, Peakflow SP uses the default thresholds or the thresholds that you set to determine the levels of importance. For additional information, see the following:

- "How Peakflow SP creates and classifies standard host alerts" on page 49
- "How Peakflow SP creates and classifies profiled network alerts" on page 54
- "How Peakflow SP creates and classifies profiled router detection alerts" on page 40

About the Alert Listing Pages

Introduction

The alert listing pages display information about the alerts that are triggered by Peakflow SP. To search for alerts on the alert pages, you can use the **Search** box and the Alert Search Wizard.

The All Alerts page (**Alerts > All Alerts**) displays all current and past alert activity. You can also use the Alerts Ongoing page (**Alerts > Ongoing**) to view the same information for ongoing alerts.

For additional information, see the following:

- "How Alerts Work" on the previous page
- "Introduction to DoS Alerts" on page 68
- "Navigating multiple pages" on page 15
- "About searching for alerts on the alert listing pages" on the next page

About the layout of the alert listing pages

The alert listing pages contain the following information:

Information	Description
Search box	Use to search for alerts, with or without keywords. See "About searching for alerts on the alert listing pages" on the next page.
Wizard button	Click to search for alerts by using the Alert Search Wizard. See "About searching for alerts on the alert listing pages" on the next page.
ID	The unique number that is assigned to each alert. If the ID is a link, you can click the link to navigate to the alert's detail page.
Graph	For traffic alerts, a thumbnail traffic graph that is a visual depiction of an alert's ongoing activity. You can click the graph to navigate to the alert's detail page.
Importance	The alert's severity level (high, medium, or low), severity percentage, and impact. For DoS alerts, the severity percentage is the percentage by which traffic in a DoS alert exceeded the configured pps or bps threshold for a managed object. Impact indicates the bandwidth that an alert consumes in your network. See "Why severity percent, impact, and max values might not match" on page 71. Peakflow SP sorts alerts in the Importance column first by severity level and then by severity percent.
Alert	The type of alert with key information about the alert. This information includes the resource associated with the alert.

Information	Description
Start Time	The time at which the alert activity was first detected, followed by the duration of the alert in days, hours, and minutes (DD d, HH:MM). If the alert has not ended, Peakflow SP displays **Ongoing**.
page navigation links	Click to navigate to other Alerts pages. See "Navigating multiple pages" on page 15.

About searching for alerts on the alert listing pages

You can use the **Search** box to search on the alert listing pages. The following are some guidelines for using the **Search** box:

- You can enter search values with or without keywords.

  See "Acceptable search keywords and values for alerts" on the facing page.
- You can enter the search values "fast," "flood," or "fast flood" to search for all DoS Host alerts that are triggered by fast flood detection.
- Keywords allow you to search on a specific attribute.

  See "Acceptable search keywords and values for alerts" on the facing page.
- When you enter a keyword followed by a value, do not put a space between the colon and the value that you enter.
- If you do not enter a keyword, then Peakflow SP tries to match your search entry to specific elements in the list of alerts. These elements include the alert ID (if you entered a positive integer), alert type, severity level, status, and resource.

  A resource is a profile managed object.
- A space between search values creates an AND statement.
- A comma between search values creates an OR statement.
- You can use quotation marks (") to match a phrase.
- A match occurs when a search value matches any part of a text string.

If you want to add time search criteria to your search, use the Alert Search Wizard. See "Using the Alert Search Wizard" on page 64.

Acceptable search keywords and values for alerts

The following table describes the acceptable keywords and values that you can use to search for alerts in the **Search** box:

Attribute to search by	Acceptable keywords and values	Examples
resource (a profile managed object)	■ **resource:** *profile managed object name* ■ **mo:***profile managed object name*	■ resource:object3,example_service ■ mo:object1 ■ service:new_serv1 The “resource” keyword searches for alerts that involve profiles. This search is case-insensitive, and Peakflow SP matches on partial resources.
alert ID	■ *ID number* ■ **alert_id:***ID number*	■ 12345 ■ alert_id:23456
alert class	■ **ac:***alert class* ■ **alert_class:***alert class*	■ ac:TMS ■ alert_class:TMS
severity level	■ *severity* ■ **sev:***severity* ■ **severity:** *severity*	■ low ■ sev:low ■ severity:high,low
alert type	■ *alert type* ■ **at:***alert type* ■ **alert_type:***alert type*	■ “BGP Trap” ■ at:“BGP Trap” ■ alert_type:“BGP Trap” This search is case-insensitive, and Peakflow SP matches on partial alert types. For example, if you type **at:udp** in the **Search** box, Peakflow SP returns all the alerts that have UDP in the Alert column.
alert status	■ *alert status* ■ **sts:***alert status* ■ **status:***alert status*	■ ongoing ■ sts:recent ■ status:all You can type **all**, **ongoing**, **recent**, **ended**, **stopped**, **done**, or **completed** for the alert status.
prefix	■ **prefix:***CIDR block*	■ prefix:10.0.0.0/8 **Note:** If an alert is very short-lived, you might not be able to find it by using the prefix keyword.

About the search results

By default, the search returns the top 100 results in order of relevance. You can override the default setting for specific searches by using the Alert Search Wizard. See "Using the Alert Search Wizard" below.

Using the Alert Search Wizard

To search for alerts from the Alert Search Wizard:

1. From the **Alerts** menu, navigate to any alert listing page, and then click **Wizard**.
2. In the Alert Search Wizard, configure the search settings.

   See "Settings in the Alert Search Wizard" below.

   When you configure multiple settings, Peakflow SP combines them using AND statements.
3. Click **Search**.
4. (Optional) If you searched by the start or stop time, you can configure the time controls that appear, and then click **Search**.

   These time controls further refine the search, based on the initial search results.
5. (Optional) If you do not click away from the page, then you can repeat these steps to add or change the search criteria.

Settings in the Alert Search Wizard

Use the following table to configure the Alert Search Wizard settings:

Setting	Procedure
Severity level check boxes	Select the check boxes for the severity levels by which to search. The severity levels are High, Medium, and Low.
Alert Class list	Select the alert class by which to search.
Alert Type list	Select the alert type by which to search.
Search Limit box	Type the maximum *number* of results to return.
Items per Page box	Type the *number* of results to view per page.
Status check boxes	Select the check boxes for the alert statuses to include in the search. The statuses are Ongoing and Recent.
Start and **Stop** settings	Configure the start and stop times by which to search.
Impact boxes	Type the bps and pps values by which to search for impact data that falls above, below, or within a range of traffic thresholds. This is measured by the highest single-minute sum of alert-traffic rates that occurred at a single router or at the boundary interfaces of the managed object.
Severity Percent boxes	Type the severity percentage range (from lowest to highest) by which you want to search for alerts. Severity Percent is the percentage by which traffic in an alert exceeded the configured pps or bps threshold for a managed object.

About deleting alerts

To manage the alert pages, you can delete alerts manually or schedule Peakflow SP to delete alerts automatically after a specified number of days.

See “Deleting Alerts” on page 92.

About the DoS Alert Pages

Introduction

The DoS alert pages can display the following types of alerts:

- DoS Host

  A DoS Host alert is triggered when the traffic on a monitored router exceeds the configured threshold of an enabled misuse type for a specified time period.

- DoS Profiled Router

  A DoS Profiled Router alert is triggered when traffic on a router significantly exceeds the expected levels for a sustained period of time.

- DoS Profiled Network

  A DoS Profiled Network alert is triggered when network-wide traffic for one or more hosts significantly exceeds the expected baseline for a sustained period of time.

A DoS alert provides details about a DoS attack and how it affects your network. It displays breakdowns of what triggered an alert. For additional information about DoS alerts, see the following topics:

- “Introduction to DoS Alerts” on page 68
- “About the Summary Tab on a DoS Alert Page” on page 72
- “About the Traffic Details Tab on a DoS Alert Page” on page 79

Navigating to a DoS alert page

To navigate to a DoS alert page:

1. Navigate to the All Alerts page (**Alerts > All Alerts**).
2. In the **Search** box, type one of the following, depending on the type of DoS alert you want to view, and then click **Search**.
   - `host`
   - `profiled router`
   - `profiled network`
3. Click the graph or ID link for the DoS alert.

About DoS Host alerts

DoS Host alerts are triggered by host detection. Host detection can trigger a standard DoS Host alert or a Fast Flood DoS Host alert. A standard DoS Host alert is triggered when the traffic on a monitored router towards a single host exceeds the configured threshold of an enabled misuse type for a specified time period. See “About Host Detection” on page 45.

A Fast Flood DoS Host alert is triggered when large amounts of traffic toward a single host are detected for an enabled misuse type. A Fast Flood DoS Host alert always has a severity of High, and the severity is always followed by **Fast Flood**. See “About host detection with fast flood detection enabled” on page 46.

Note: You can enter the search values “fast,” “flood,” or “fast flood” in the **Search** box on an alert listing page to search for Fast Flood DoS Host alerts.

If excessive traffic is detected for multiple misuse types that are enabled, then a single DoS Host alert is created instead of separate alerts for each misuse type. The alert includes each misuse type that had excessive traffic. See “Host detection misuse types” on page 48.

About DoS Profiled Router alerts

DoS Profiled Router alerts are triggered by profiled router detection. A DoS Profiled Router alert is triggered for a profiled managed object when traffic on a router significantly exceeds the expected levels for a sustained period of time. The traffic rate that Peakflow SP expects for a profile managed object is referred to as the baseline. A baseline is the learned traffic rate for a profile managed object. See “About Profiled Router Detection” on page 39.

About DoS Profiled Network alerts

DoS Profiled Network alerts are triggered by profiled network detection. Peakflow SP triggers a profiled network alert if the rate of the traffic at a profile managed object boundary for one or more hosts exceeds the baseline by the detection percentage for a sustained period of time. Profiled network detection identifies excessive rates of network-wide traffic based on baselines that Peakflow SP has calculated for your network. See “About Profiled Network Detection” on page 53.See “About Profiled Network Detection” on page 53.

Introduction to DoS Alerts

Introduction

A DoS alert provides details about a possible DoS attack and how it affects your network. It displays breakdowns of what triggered an alert.

You can access a DoS alert to perform the following tasks:

- determine if an alert represents an attack
- determine how to mitigate an attack
- add traffic data to an Alert Scratchpad for use in a mitigation
- initiate a mitigation

You can use the icons on the Arbor Smart Bar to download or email the information in a DoS alert. See "About the Arbor Smart Bar " on page 13.

For additional information about DoS alerts, see the following topics:

- "About the DoS Alert Pages" on page 66
- "About the Summary Tab on a DoS Alert Page" on page 72
- "About the Traffic Details Tab on a DoS Alert Page" on page 79
- "About the Top Traffic Patterns Table" on page 82
- "Recognizing a Potential DoS Attack" on page 89

About the information in the header of a DoS alert

The header above the tabs of a DoS alert displays the following information:

Information Type	Description
Alert type and alert ID	The page title includes the alert type and the alert ID. **Example:** DoS Host Alert 35803
Alert timeframe	The alert timeframe appears below the page title. The timeframe includes the start time, the end time (or "Ongoing" if the alert is still active), and the duration. **Examples:** ■ Apr 22 04:07 - Apr 23 21:42 (1d, 17:35) ■ Apr 9 20:08-21:06 (0:58) See "How Peakflow SP determines if a DoS alert should be ended or ongoing" on page 71.

Information Type	Description
Mitigations	For each type of mitigation that has been applied to a DoS alert, a mitigation type link appears below the **Mitigate Alert** button in the upper-right corner of a DoS alert. The mitigation type link includes the number of mitigations of that type. If you click a mitigation type link, a list of the mitigations of that type appears. The name of each mitigation in the list is a link that opens that mitigation. For each TMS mitigation, a summary of the traffic that is passed or dropped by that mitigation is displayed. For more information about initiating a mitigation from a DoS alert, see "Initiating a Mitigation from a DoS Alert" on page 114. **Note:** You must have the proper user privileges for the mitigations information to appear.

About the Alert Scratchpad of a DoS alert

You can add traffic data from a DoS alert to an Alert Scratchpad and then copy the data from the scratchpad and paste it into a mitigation that is associated with the alert. The Alert Scratchpad opens when you click the **View Scratchpad** button at the top of a DoS alert page. The number of items that you have added to the Alert Scratchpad is in parentheses on the **View Scratchpad** button. For information about using the Alert Scratchpad, see "About the Alert Scratchpad" on page 85.

About initiating a mitigation from a DoS alert

You can click the **Mitigate Alert** button to initiate a mitigation from a DoS alert page. See "Initiating a Mitigation from a DoS Alert" on page 114.

Note: You must have the proper user privileges for the **Mitigate Alert** button to appear.

About the traffic data displayed for a DoS alert

You can use the **Period** and **Units** to control the traffic data that is displayed on the **Summary** and **Traffic Details** tabs of a DoS alert . After you make changes to either of these lists, click **Update** to update the display of the traffic data. When you make any changes to these lists on one tab, the same changes are made on the other tab.

The **Period** list allows you to look at the alert's traffic data for a selected period of time. You might look at a subset of the timeframe of an alert for purposes of forensics. If you select **Other** from this list, you can then specify a start and end time. You can type the time in the **Start** and **End** boxes or you cab click the calendar icon to select the data and time. You can also type entries like "2 weeks ago," "100 hours ago," "last Monday," or "5 May" in the **Start** and **End** boxes.

By default, the displayed timeframe of a DoS alert is set to the duration of an alert. If you change the timeframe of a DoS alert, you can select **Alert Timeframe** from the **Period** list to redisplay the data for the alert's duration.

About the aggregation of IP addresses and ports in a DoS alert

Peakflow SP aggregates IP addresses and ports in a DoS alert to help identify attack traffic.

About the aggregation of IP addresses in a DoS alert

Peakflow SP aggregates IP addresses to consolidate the data and make it more useful. These aggregated IP addresses can help you identify the source and destination IP addresses of potential attack traffic.

During each minute of a DoS alert, Peakflow SP collects data on the source and destination IP addresses of the alert traffic and aggregates them as follows:

- Aggregates the IP addresses until it identifies an IP prefix that represents at least 10% of the alert traffic.
- Continues to aggregate IP addresses until it identifies an IP prefix that represents at least 10% of the alert traffic in addition to the traffic of the previously identified prefix.
- Continues this process of aggregation as long as it can identify IP prefixes that represent at least 10% of the alert traffic in addition to the traffic of previously identified prefixes.

After Peakflow SP aggregates the source and destination IP addresses of a DoS alert, it can display these aggregated IP addresses in the following locations of a DoS alert page:

- Top Traffic Patterns (last 5 minutes) table and the CSV file of all traffic patterns

  See "About the Top Traffic Patterns Table" on page 82.
- Alert Characterization table

  See "About the Alert Characterization table on the Summary tab" on page 77.
- Source IP Addresses and Destination IP Addresses tables and the View More Details window for source and destination IP addresses

  See "About the traffic statistics tables on the Traffic Details tab" on page 79.

  Note: The IP Addresses tables and the View More Details window for IP addresses can also display top individual IP addresses even if they do not represent at least 10% of the alert traffic during any minute of the alert. Peakflow SP displays these individual addresses if it identifies them as top IP addresses for the alert for the selected timeframe.

When Peakflow SP displays aggregated IPv4 addresses, it can use any CIDR block from "/32" to "/8." When Peakflow SP displays aggregated IPv6 addresses, it uses every fourth CIDR block from "/128" to "/8" (for example: "/124," "/120," and "/116"). If Peakflow aggregates IP addresses above "/8," it displays Highly Distributed for the name of the aggregated IP prefix, which represents any IP address.

When Peakflow SP displays the percentage of the traffic that an aggregated IP address represents, it is the percentage of the overall traffic of the alert for the selected timeframe. Consequently, an aggregated IP address that represents at least 10% of the traffic during any minute of an alert may only represent 2% of the overall traffic of the alert.

Note: Although the aggregation of IP addresses is enabled by default, you can use the CLI to disable it. .

About the aggregation of ports in a DoS alert

Peakflow SP gathers the source and destination ports of the TCP and UDP traffic of a DoS alert. These ports can help determine if the traffic is normal traffic or attack traffic.

Peakflow SP displays data on individual ports that represent at least 10% of the alert traffic during any minute of the alert. Peakflow SP also aggregates system and dynamic ports and displays the port range with the name of the range, as follows:

- 1-1023 (System)
- 1024-65535 (Dynamic)

Note: In the traffic statistic tables for ports on the **Traffic Details** tab, the name of the port range is in its own column.

Peakflow SP can display individual and aggregated ports in the following tables of a DoS alert:

- Top Traffic Patterns (last 5 minutes) table

  See "About the Top Traffic Patterns Table" on page 82.

- Alert Characterization table

  See "About the Alert Characterization table on the Summary tab" on page 77.

- TCP and UDP source and destination port tables and the View More Details window for ports

  See "About the traffic statistics tables on the Traffic Details tab" on page 79.

Note: The source and destination port tables and the View More Details window for ports can also display top individual ports even if they do not represent at least 10% of the alert traffic during any minute of the alert. Peakflow SP displays these individual ports if it identifies them as top ports for the alert for the selected timeframe.

How Peakflow SP determines if a DoS alert should be ended or ongoing

After a DoS alert is triggered, Peakflow SP continues to monitor all of the sources of traffic associated with the alert. If the traffic at the source that triggered the alert indicates that the alert should be ended, but another source indicates that the alert should be ongoing, then the alert remains ongoing.

For example, with a DoS Host alert, if an attack is triggered at the managed object boundary and the attack traffic is being mitigated, then the traffic at the managed object boundary might indicate that the alert should be ended. However, if the attack is still in progress, then the data from the network boundary or from an individual router would indicate that the alert should remain ongoing. Because the data from the network boundary or an individual router most accurately reflects the state of the attack, the alert remains ongoing.

Why severity percent, impact, and max values might not match

The severity percent, impact, and affected router max values will not always match. Non-matching values can usually be attributed to differences between when the measurements are taken for each of the values. The measurements for the severity percent and impact values can be taken both during the latency period and after an alert has been generated. The measurements for the affected router max values are only taken after an alert has been generated. If the highest rate of traffic for an alert occurs during the latency period, then the severity percent and impact values can be higher than the affected router max values.

With DoS Host alerts, the severity percent and impact can be based on the traffic of different misuse types which can cause these values not to match. The impact value is always based on the total traffic, while the severity percent is based on the traffic of the misuse type that triggered the alert. If the misuse type that triggered the alert is total traffic, then the impact value and severity percent are both based on total traffic. However, if the alert was not triggered by total traffic, then the impact value would be based on total traffic, while the severity percent would be based on the traffic of another misuse type. For more information about severity percent and impact, see "About key alert information on the Summary tab" on the next page.

About the Summary Tab on a DoS Alert Page

Introduction

The **Summary** tab on a DoS alert page displays a summary of the information concerning a DoS alert for the selected timeframe.

For general information about a DoS alert page including how to control the traffic data that is displayed, see "Introduction to DoS Alerts" on page 68.

About key alert information on the Summary tab

The following information is displayed above the traffic graph on the **Summary** tab of a DoS alert page:

Icon	Description
	Hover over this icon to display the download options. The download options that appear vary depending on the page that you are on and can include the following: ■ PDF - Click to download the page in PDF format. ■ XML- Click to download the page in XML format. ■ CSV - Click to download a page in CSV format. The download can be a CSV text file or CSV zip archive file. When the download is a zip archive file, "zip archive" is appended to CSV. ■ Excel-XML - Click to download a page in Excel-XML format. The icon appears only if there are multiple options for downloading a page.
	Click to download a page in PDF format. This icon appears only when the PDF format is the only download option for a page.
	Click to download and email a page as a PDF.

Column	Description
	Select this icon to delete a user account.
Username	A user name as a link to the Edit Existing Account page.
Real Name	A user's full name.
Account Group	The account group to which a user belongs.
Capability Level	A user's capability level, which is either an administrator or a user.
Email	A user's email address.
Device	The SP appliance with which a user is associated. The SP appliance is either a specific appliance name or *global*, which associates a user

Information Type	Description
Impact	The Impact column displays the bandwidth that an alert consumes in your network. For a DoS Profiled Network alert, the Impact is based on the same highest single-minute of traffic that is used to calculate the Severity Percent. For a DoS Profiled Router alert, the Impact is based on the highest single-minute of traffic at the network boundary, managed object boundary, or an individual router. For a DoS Host alert, the impact is based on the highest single-minute of the “total traffic” misuse type, while the Severity Percent is based on the misuse type that triggered the alert. The misuse type that triggered the alert can be “total traffic” or it can be one of the other misuse types. With a DoS Host alert and a DoS Profiled Network alert, this column also displays where the impact data was recorded. With a DoS Host alert, the impact data can be recorded at the managed object boundary, the network boundary, or an individual router. With a DoS Profiled Network alert, the impact data can be recorded at the managed object boundary or the network boundary.
Direction	The Direction column displays the direction of the alert traffic in the local network (incoming or outgoing). **Note:** With a DoS Host alert, the direction is always incoming.
Misuse Types	The Misuse Types column appears only with a DoS Host alert. It displays the misuse types that had traffic that exceeded the configured trigger rate threshold for that type of traffic.
Type	The Type column appears only with a DoS Profiled Router alert. It displays the alert’s type. The type can be Bandwidth, Multi-Protocol, or an individual protocol. The type includes IPv4 or IPv6.

Icon	Description
	Hover over this icon to display the download options. The download options that appear vary depending on the page that you are on and can include the following: ■ PDF - Click to download the page in PDF format. ■ XML- Click to download the page in XML format. ■ CSV - Click to download a page in CSV format. The download can be a CSV text file or CSV zip archive file. When the download is a zip archive file, "zip archive" is appended to CSV. ■ Excel-XML - Click to download a page in Excel-XML format. The icon appears only if there are multiple options for downloading a page.
	Click to download a page in PDF format. This icon appears only when the PDF format is the only download option for a page.
	Click to download and email a page as a PDF.

Column	Description
	Select this icon to delete a user account.
Username	A user name as a link to the Edit Existing Account page.
Real Name	A user's full name.
Account Group	The account group to which a user belongs.
Capability Level	A user's capability level, which is either an administrator or a user.
Email	A user's email address.
Device	The SP appliance with which a user is associated. The SP appliance is either a specific appliance name or *global*, which associates a user with all appliances. For more information about associating a user with appliances, see "About user-appliance association" on page 24.
UI Menu	The UI menu that is assigned to a user. The UI menu determines what menu choices are available to a user.
Status	"Disabled" appears in this column for a user account that is disabled; otherwise, this column is blank.

About the Alert Traffic graph on the Summary tab

The Alert Traffic graph displays different data for each of the different types of DoS alerts.

About the Alert Traffic graph on the Summary tab of a DoS Host Alert page

The **Summary** tab on a DoS Host Alert page has an Alert Traffic graph that can display the following information about the traffic of an alert for the selected timeframe:

- Traffic for each misuse type that is part of the alert, including the total traffic misuse type
- The traffic and trigger rate of a misuse type that has exceeded the trigger rate
- Traffic that is dropped by any TMS mitigations associated with the alert

If you move your mouse across a line on an Alert Traffic graph, the amount of traffic at that point on the line is displayed along with the time.

Colored selectors appear above the graph for the different types of alert traffic. A black threshold selector can also appear that allows you to display the trigger rate threshold for a misuse type.

The following table describes the different selectors that appear above the graph:

Selector	Description
Misuse Types	These selectors hide or display lines or areas on the Alert Traffic graph for each misuse type of alert traffic including total traffic. For information about the different misuse types, see "Host detection misuse types" on page 51. When the Total Traffic selector is selected, a gray background represents all of the alert traffic if the graph does not include traffic that is being dropped by a TMS mitigation. If the graph includes traffic that is being dropped by a TMS mitigation, then the gray background represents traffic that is passed, and a red background represents traffic that is dropped.
Trigger Rate	This selector appears only when **Router** is selected in the **View** list and only when the traffic of a single misuse type is displayed. With total traffic, the trigger rate selector appears for both bps and pps traffic. For the other misuse types, the trigger rate selector appears only for pps traffic. You must click the trigger rate selector to display the trigger rate threshold on the Alert Traffic graph.
Mitigated Traffic	This selector appears only when **Network Boundary** is selected in the **View** list and only when a TMS mitigation associated with the alert is dropping traffic. This selector hides or displays the traffic that is being dropped by any TMS mitigations. A red background represents the traffic that is being dropped.

A selector is a solid-colored circle when what it represents is displayed in the graph, and it appears as an empty circle when the traffic it represents is hidden. When the traffic for a misuse type exceeds its configured trigger rate, then the text of its misuse type selector is red and is followed by an asterisk (*).

You can click a selector to hide or display its traffic or the trigger rate threshold. You can double-click a misuse type selector to display just the traffic for that misuse type. If only one type of traffic is selected, you can click its selector to display all the types of traffic that are associated with the alert.

About the Alert Traffic graph on the Summary tab of a DoS Profiled Router Alert page

The **Summary** tab on a DoS Profiled Router Alert page has an Alert Traffic graph that can display the following types of traffic for the selected timeframe:

- Total traffic for the alert
- Traffic that is dropped by any TMS mitigations associated with the alert

The following table describes the different selectors that appear above the graph:

Selector	Description
Total Traffic	This selector always appears above the graph, but it only functions as a selector when the **Mitigated Traffic** selector is also present. When the Total Traffic selector is selected, a gray background represents all of the alert traffic if the graph does not include traffic that is being dropped by a TMS mitigation. If the graph includes traffic that is being dropped by a TMS mitigation, then the gray background represents traffic that is passed, and a red background represents traffic that is dropped.
Mitigated Traffic	This selector appears only when **Network Boundary** is selected in the **View** list and only when a TMS mitigation associated with the alert is dropping traffic. This selector hides or displays the traffic that is being dropped by any TMS mitigations. A red background represents the traffic that is being dropped.

A selector is a solid-colored circle when what it represents is displayed in the graph, and an empty circle when the traffic it represents is hidden. You can click a selector to hide or display the traffic that the selector represents. You can double-click a selector to display just that type of traffic. If only one type of traffic is selected, you can click its selector to display both types of traffic that are associated with the alert.

About the Alert Traffic graph on the Summary tab of a DoS Profiled Network Alert page

The **Summary** tab on a DoS Profiled Network Alert page has an Alert Traffic graph that displays the following information about the alert for the selected timeframe:

- Traffic

  The total incoming or outgoing traffic observed for the alert for the selected timeframe. If profiled country detection is enabled for the managed object associated with the alert, then a tab appears for each of the top 5 countries for which traffic is detected. Each of these tabs displays the traffic observed for that country. If traffic is detected for more than one country, then a Stacked Countries tab appears that displays the traffic for each of the top five countries. For information about enabling country detection, see “Configuring Profiled Network Detection for Profile Managed Objects” on page 56.

  Note: If both the incoming and outgoing traffic of a managed object trigger an alert, then two separate alerts are triggered.

- Baseline

  The learned traffic rate for normal traffic. See “About profiled network detection baselines” on page 54.

- Detection threshold

  The threshold that traffic must exceed before a DoS Profiled Network alert can be triggered. The detection threshold is determined by adding a percentage of the baseline to the baseline. The percentage that is added is configured when the managed object is

configured. See "Configuring Profiled Network Detection for Profile Managed Objects" on page 56.

The following vertical colored lines can appear on the Alert Traffic graph of a DoS Profiled Network alert:

Line Color	Description
gray	Indicates when an annotation was applied to an alert.
yellow	Indicates when an alert was changed to medium importance.
red	Indicates when an alert was changed to high importance.
green	Indicates when an alert started.
black.	Indicates when an alert stopped.

You can drill down into the Alert Traffic graph of a DoS Profiled Network alert to see a more detailed view of the traffic. To drill down into the graph, click and drag across the graph to select the timeframe that you want to view.

About the Top Traffic Patterns (last 5 minutes) table on the Summary tab

The Top Traffic Patterns table appears on the **Summary** tab and **Traffic Details** tab of a DoS Host alert or a DoS Profiled Router alert. Peakflow SP looks at the traffic in an alert and aggregates the src/dst CIDRs and the src/dst port ranges to identify groups of flows that have the same 5-tuple traffic pattern (src/dst IP, src/dst port, and protocol). Peakflow SP then populates this table with traffic patterns for the alert that represent at least 10% of the traffic during the last 5 minutes of the selected timeframe. See "About the Top Traffic Patterns Table" on page 82.See "About the Top Traffic Patterns Table" on page 82.

About the Alert Characterization table on the Summary tab

The Alert Characterization table lists different elements associated with the alert. For each element, it lists the items that contributed at least 25% of the traffic of the alert. For each element that appears in this table, a data table appears on the Traffic Details page.

A (context menu) icon appears to the left of each element in the Alert Characterization table. When you click , the options that you can select depend on the traffic item. The following options can appear:

- **Add Item to Alert Scratchpad**

  See "Adding traffic items to an Alert Scratchpad" on page 86.
- **Lookup IP Address (Whois)** (IP addresses only)

  See "Performing a Whois Lookup for an IP Address on a DoS Alert Page" on page 88.

The data for each item is the total for that item from all of the traffic of the alert. For example, if an alert has the following amounts and types of traffic:

- 40% to TCP port 80
- 40% to UDP port 53
- 20% to UDP port 80

then the Alert Characterization table would include the following data:

Protocols	udp	60%
Protocols	tcp	40%
Destination TCP Ports	80	40%
Destination UDP Ports	53	40%

The 20% to UDP port 80 would not appear in the table because it was less than 25% of the alert traffic.

Note: By default, Peakflow SP aggregates IP prefixes. For information about how Peakflow SP aggregates IP prefixes, see "About the aggregation of IP addresses and ports in a DoS alert" on page 69.

Note: The Alert Characterization table displays data for individual items of the traffic of an alert, while the Top Traffic Patterns (last 5 minute) table displays traffic that shares a 5-tuple pattern. See "About the Top Traffic Patterns Table" on page 82.

About the Packet Size Distribution graph on the Summary tab

The **Summary** tab on a DoS alert page has a histogram that displays the distribution of the packet sizes for the alert for the selected timeframe. The left side of the graph lists groups of packet size ranges of 150 bytes each. Each horizontal bar shows the number of packets within that 150-byte range. A jumbo frames bar appears at the bottom of the graph for packets that are larger than 1500 bytes.

The Packet Size Distribution graph can often help you determine if an alert represents an attack. You can use the graph to identify whether packet sizes are spread out or concentrated. If the packet sizes are concentrated, you can then use the graph to determine if the areas of concentration are what would be expected for that type of traffic.

For example, if you receive a UDP flood alert for packets sourced from port 123 (NTP), and the majority of the packets are large (400 bytes or larger), you are probably looking at a reflection attack because these NTP packets would normally be much smaller.

The Packet Size Distribution graph can also be used for post-attack forensic analysis to identify patterns in packet size distribution for different types of attacks. You can then use this information to help you identify future attacks.

About the Traffic Details Tab on a DoS Alert Page

Introduction

The **Traffic Details** tab on the DoS Alert pages displays data graphs and tables about the most significant elements that contributed to an alert during a selected timeframe. The element whose data appears in the tab's main graph is highlighted with a blue background.

The **Traffic Details** tab also includes the Top Traffic Patterns (5-tuple) section. See "About the Top Traffic Patterns Table" on page 82.

A (context menu) icon appears next to most traffic items on the **Traffic Details** tab. You can click to add the item to the Alert Scratchpad or to perform a whois lookup for source and destination IP addresses. For additional information, see:

- "Adding traffic items to an Alert Scratchpad" on page 86
- "Adding a 5-tuple traffic pattern to an Alert Scratchpad" on page 85
- "Performing a Whois Lookup for an IP Address on a DoS Alert Page" on page 88

For general information about a DoS alert page including how to control the traffic data that is displayed, see "Introduction to DoS Alerts" on page 68.

Displaying and viewing data on the Traffic Details tab

The following are different options for displaying and viewing data on the **Traffic Details** tab:

- Change the traffic data displayed on the **Traffic Details** tab

  You can use **Period**, **Units**, and **View** lists at the top of the **Traffic Details** tab to change the traffic data displayed on this tab. See "About the traffic data displayed for a DoS alert" on page 69.

  In addition to using the **Period** list to change the timeframe for the alert traffic, you can also click and drag across the graph to select the timeframe that you want to view.

- Display the data of a statistics table in the traffic graph

  You can click any table or **View Graph** below a table to display that table's data the traffic graph.

- View additional data that does not appear in a traffic statistics table

  If a contributing element has more than 5 entries, you can click the **View More** link below the statistics table to view more of the entries for that contributing element. For source and destination IP addresses, the View More Details window that appears displays up to 100 aggregated IP addresses.

- Download all the source IP addresses

  If an alert has more than 5 source IP address, you can view all of its source IP addresses. When you click **View More** below a source IP address table, the View More Details window that appears has a **Download All** button. When you click the **Download All** button, Peakflow SP downloads a CSV file of all the source IP addresses that are associated with the alert.

About the traffic statistics tables on the Traffic Details tab

The **Traffic Details** tab displays traffic statistics tables and allows you to investigate the traffic to determine if is malicious. For information about viewing more information than is displayed in these tables, see "Displaying and viewing data on the Traffic Details tab" above.

For information about the Top Traffic Patterns (last 5 minutes) table, see "About the Top Traffic Patterns Table" on page 82.

Peakflow SP gathers the data that is displayed in these tables every minute. Peakflow updates the data in the tables whenever the DoS alert is manually updated. Each table displays the rate and percentage of the traffic for the items listed in the table. The rate and percentage are based on the overall traffic of the alert for the selected timeframe.

For the IP address tables and the port tables, Peakflow SP aggregates the IP addresses and the ports. For information on how Peakflow SP aggregates IP addresses, see "About the aggregation of IP addresses and ports in a DoS alert" on page 69.

The **Traffic Details** page contains the following data tables:

Table	Description
Source IP Addresses	Displays the top 5 aggregated source IP addresses of the alert traffic. By default, Peakflow SP aggregates IP addresses. **Note:** Attackers can forge source IP addresses. Do not rely on these statistics to identify the actual source of traffic.
Destination IP Addresses	Displays the top 5 aggregated destination IP addresses of the alert traffic. By default, Peakflow SP aggregates IP addresses. This table can help you determine the destination of potential attack traffic and the volume of this traffic.
Source TCP Ports	Displays the top 5 source ports or aggregated ports for the TCP packets. A port is followed by the service name, and an aggregated port is followed by the aggregated port name. This table can sometimes help you determine if the traffic is normal traffic or attack traffic.
Destination TCP Ports	Displays the top 5 destination ports or aggregated ports for the TCP packets. A port is followed by the service name, and an aggregated port is followed by the aggregated port name. This table can help you determine the type of ports that are likely to be affected by this traffic. **Example:** If most of the traffic has a destination (DST) port of 80 and the protocol is TCP, then the HTTP service on one or more hosts is the target service for most of the traffic. If the destination ports are listed as 0-65535, this is most likely an attack against all services on the destination host. Consult the destination address table to determine which hosts might be the target.
Source UDP Ports	Displays the top 5 source ports or aggregated ports for the UDP packets. A port is followed by the service name, and an aggregated port is followed by the aggregated port name. This table can sometimes help you determine if the traffic is normal traffic or attack traffic.

Table	Description
Destination UDP Ports	Displays the top 5 destination ports or aggregated ports for the UDP packets. A port is followed by the service name, and an aggregated port is followed by the aggregated port name. This table can help you determine the type of ports that are likely to be affected by this traffic.
Source Countries	Displays the top 5 source countries for the alert traffic. The country's flag is followed by its name.
Source ASNs	Displays the top 5 ASNs for the alert traffic. The ASN number is followed by the ASN name.
Protocols	Displays the top 5 protocols for the alert traffic. The most common protocols are as follows: ■ TCP — usually legitimate traffic ■ UDP — usually legitimate traffic ■ ICMP or IPv6-ICMP— large amounts of ICMP traffic usually indicates a problem Other protocols can appear here if your network uses or provides service for customers who are using GRE tunnels, IPSec tunnels, OSPF, or other facilities that use their own protocol number.
TCP Flags	Displays the top 5 sets of TCP flags for the alert traffic. It displays the TCP flags that have been set to 1. The letter or letters for the TCP flag are followed by the name or names of the TCP flag. TCP uses the TCP flags to signal the beginning and end of connections and other conditions. The individual letters that appear for the flags indicate which flags are set to 1 in the associated flows. Not all packets associated with the listed flag sets have all of the flags set to 1. Instead, it indicates that at least one packet in each flow has the associated flag set to 1.
ICMP Types	Displays counts and rates for ICMP packets with the specified ICMP type. If ICMP packets were not seen during the sampling process, then this table might be empty. You can use this information to determine the ratio of ICMP Echo Request packets to ICMP Destination Unreachable packets.
Misuse Types	Displays the top 5 misuse types for the alert traffic.

About the Top Traffic Patterns Table

Introduction

The Top Traffic Patterns table appears on the **Summary** tab and **Traffic Details** tab of a DoS Host alert or a DoS Profiled Router alert. This table displays the top traffic patterns identified in the traffic of an alert.

How Peakflow SP populates the Top Traffic Patterns table

Peakflow SP looks at the traffic in an alert and aggregates the src/dst CIDRs and the src/dst port ranges to identify groups of flows that have the same 5-tuple traffic pattern (src/dst IP, src/dst port, and protocol). Peakflow SP then populates this table with traffic patterns for the alert that represent at least 10% of the traffic during the last 5 minutes of the selected timeframe.

Note: The Top Traffic Patterns table displays traffic that shares a 5-tuple pattern, while the Alert Characterization table displays data for individual items of the traffic of an alert. See "About the Alert Characterization table on the Summary tab" on page 77.

How to use the top traffic patterns

The traffic patterns can help you determine if an alert represents an attack. If you determine that a traffic pattern represents an attack, you can then add data from the traffic pattern to an Alert Scratchpad, and then copy the data into a mitigation.

Top traffic patterns represent interesting traffic, but not necessarily bad traffic. You should look for top traffic patterns that stand out because they are in some way abnormal for your network. An abnormal traffic pattern could have a protocol that you normally do not see or a high volume of traffic when that pattern usually has a low volume of traffic.

For example, if you have a pattern for TCP traffic that has a source IP address of Widely Varied, a Dynamic source port, and all the traffic targeting an NTP server (port 123), then this traffic pattern might indicate an NTP reflection attack. If the Packet Size Distribution table has an unusually high number of small packets, it would confirm that the attack is an NTP reflection attack. For information about the names of aggregated IP address and ports, see "About the aggregation of IP addresses and ports in a DoS alert" on page 69.

About the display of top traffic patterns

When a DoS alert is initially triggered, Peakflow SP does not display any top traffic patterns. It can only display top traffic patterns 2 to 3 minutes after an alert is triggered.

Peakflow SP can display up to 10 top traffic patterns in the Top Traffic Patterns table. You can also click the **Download All Patterns** button to view all of the traffic patterns that are associated with an alert. When you click **Download All Patterns**, a CSV file is generated that lists the traffic patterns.

If an alert is displaying data for a specific router, then the Top Traffic Patterns table displays only patterns for the alert traffic associated with that router. If an alert is displaying data for a network boundary or a managed object boundary, then this table displays patterns for the alert traffic of all the routers associated with the alert. See "About the traffic data displayed for a DoS alert" on page 69.

Peakflow SP updates the top traffic pattern data every minute, but it does not automatically update the Top Traffic Patterns table. To update this table, you must refresh the DoS alert page or click **Update**.

Peakflow SP does not display top traffic patterns if any of the following occur:

- A Peakflow SP appliance does not have enough system resources to generate the traffic patterns for every alert.
- The alert had no traffic in the last 5 minutes.
- The alert is a DoS Profiled Network alert.
- The alert is triggered on one of the following:
  - An appliance with a serial number that begins with AZLR
  - A virtual machine with less than 4 processors
  - A virtual machine running Xen

  Note: Because the generation of traffic patterns is very processor intensive, it is disabled on appliances or VMs where it would significantly disrupt the normal operation of Peakflow SP.

Information that appears in a traffic pattern

For each traffic pattern, the following information is listed:

Information	Decription
Source	The source IP addresses of the alert traffic in this traffic pattern. By default, Peakflow SP aggregates IP prefixes. For information about how Peakflow SP aggregates IP prefixes, see "About the aggregation of IP addresses and ports in a DoS alert" on page 69. If the name of an IP address is truncated, you can hover your mouse over the name to display the full name.
Protocol	The protocol of the alert traffic in this traffic pattern.
Flags	For the TCP protocol, the TCP flags of the alert traffic in this traffic pattern.
Src Port	The source port or aggregated source port of the alert traffic in this traffic pattern. A port is followed by the service name, and an aggregated port is followed by the aggregated port name. If no ports are displayed, then the ports have been aggregated to include all ports. For information about how Peakflow SP aggregates ports, see "About the aggregation of IP addresses and ports in a DoS alert" on page 69.
Destination	The destination IP addresses of the alert traffic in this traffic pattern. By default, Peakflow SP aggregates IP prefixes. For information about how Peakflow SP aggregates IP addresses, see "About the aggregation of IP addresses and ports in a DoS alert" on page 69. If the name of an IP address is truncated, you can hover your mouse over the name to display the full name.

Icon	Description
	Hover over this icon to display the download options. The download options that appear vary depending on the page that you are on and can include the following: ■ PDF - Click to download the page in PDF format. ■ XML- Click to download the page in XML format. ■ CSV - Click to download a page in CSV format. The download can be a CSV text file or CSV zip archive file. When the download is a zip archive file, “zip archive” is appended to CSV. ■ Excel-XML - Click to download a page in Excel-XML format.

About the context menu icons in a top traffic pattern

A (context menu) icon appears to the left of each traffic pattern, and additional icons appear for individual traffic items when you hover your mouse pointer to the right of the item (except for items in the Alert Traffic column). When you click , the options that you can select depend on the traffic item. The following options can appear:

- **Add Pattern to Alert Scratchpad** (traffic patterns only)

  Adds a traffic pattern to the Alert Scratchpad. See “Adding a 5-tuple traffic pattern to an Alert Scratchpad” on the facing page.

- **Add Item to Alert Scratchpad**

  Adds a traffic item to the Alert Scratchpad. See “Adding traffic items to an Alert Scratchpad” on page 86.

- **Lookup IP Address (Whois)** (IP addresses only)

  Performs a whois lookup on a source or destination IP address or aggregated IP address. See “Performing a Whois Lookup for an IP Address on a DoS Alert Page” on page 88.

About the Alert Scratchpad

Introduction

You can add traffic data from a DoS alert page to an Alert Scratchpad and then copy the data from the scratchpad and paste it into a mitigation that is associated with the DoS alert.

Important things to know about the Alert Scratchpad

The following are important things that you should know about the Alert Scratchpad:

- Each alert has its own Alert Scratchpad.
- Alert Scratchpads are user-specific.

  Any traffic data that you add to the Alert Scratchpad of a DoS alert does not appear in the Alert Scratchpad when another user accesses that DoS alert.
- The (context menu) icon next to traffic data on a DoS Alert page is used to add traffic data to the Alert Scratchpad for that alert.
- The Alert Scratchpad for a DoS alert opens when you click the **View Scratchpad** button at the top of a DoS alert page or a mitigation page.
- An Alert Scratchpad can be moved anywhere on a DoS alert or mitigation page.
- The number of traffic items that you have added to an Alert Scratchpad appears in parentheses on the **View Scratchpad** button and at the end of the title of the Alert Scratchpad window.
- The traffic data that you add to an Alert Scratchpad remains in the scratchpad when you log off.
- A downloaded PDF of a DoS alert includes any traffic data that you have added to the Alert Scratchpad for that alert.

Traffic data that you can add to an Alert Scratchpad

You can add the following types of traffic data to an Alert Scratchpad from a DoS alert page:

- Traffic patterns

  You can add a 5-tuple traffic pattern (src/dst IP, src/dst port, and protocol) to an Alert Scratchpad from the Top Traffic Patterns (last 5 minutes) table on the **Summary** tab and the **Traffic Details** tab of a DoS alert page. See “Adding a 5-tuple traffic pattern to an Alert Scratchpad” below.
- Traffic elements

  Traffic elements include any traffic items that appear on the DoS alert pages. See “Adding traffic items to an Alert Scratchpad” on the next page..

Traffic patterns are added to the Traffic Patterns section of an Alert Scratchpad and traffic elements are added to the Traffic Elements section. When traffic items are added to the Traffic Elements section, they are arranged in the same order that they appear on the **Traffic Details** tab.

Adding a 5-tuple traffic pattern to an Alert Scratchpad

You can add a 5-tuple traffic pattern (src/dst IP, src/dst port, and protocol) from the Top Traffic Patterns (last 5 minutes) table to an Alert Scratchpad. A traffic pattern is added to an Alert Scratchpad as an FCAP expression so that it is in a format that can be pasted into a mitigation.

To add a specific traffic item from the Top Traffic Patterns (last 5 minutes) table to an Alert Scratchpad, see "Adding traffic items to an Alert Scratchpad" below.

To add a 5-tuple traffic pattern to an Alert Scratchpad:

1. Navigate to the DoS alert page **(Alerts > All Alerts > alert ID link)**.
2. Click the **Summary** tab or the **Traffic Details** tab.
3. In the Top Traffic Patterns (last 5 minutes) table, hover your mouse pointer over the (context menu) icon to the left of a traffic pattern.

   When you hover your mouse pointer over the (context menu) icon , it becomes more visible.
4. Click (context menu), and then click **Add Pattern to Alert Scratchpad**.

Adding traffic items to an Alert Scratchpad

In addition to traffic patterns,you can also add most of the other traffic items that appear on a DoS alert page to an Alert Scratchpad.

You can use the (context menu) icon to add traffic items to an Alert Scratchpad from the following locations on a DoS alert page:

- Alert Characterization table on the **Summary** tab

  Note: Any information that you can add to an Alert Scratchpad from the Alert Characterization can also be added from the data tables on the **Traffic Details** tab.
- Data tables on the **Traffic Details** tab
- View More Details window on the **Traffic Details** tab

  You access this window by clicking **View all** below the data table. If a data table has 5 or less items, **View all** does not appear.
- Top Traffic Patterns (5-tuple) table on the **Summary** tab or the **Traffic Details** tab

  You can add items to an Alert Scratchpad from every column in the table except the % Alert Traffic column.

  Important: The (context menu) icon for individual items in the Top Traffic Patterns (5-tuple) is to the right of the item and does not appear until you hover you mouse pointer over it.

Deleting traffic data from an Alert Scratchpad

If you see items in an Alert Scratchpad that no longer apply to your mitigation strategy, you can delete them. You can delete items from an Alert Scratchpad on a DoS alert page or on a mitigation page. You can delete individual items or all of the items in the Traffic Patterns or the Traffic Elements sections of an Alert Scratchpad. When you delete all of the items from an Alert Scratchpad on a mitigation page, the Alert Scratchpad and the **View Scratchpad** button disappear from the mitigation page.

To delete traffic data from an Alert Scratchpad:

1. Navigate to a DoS alert page **(Alerts > All Alerts > alert ID link)** or the mitigation page for the DoS alert.

   When you initiate a mitigation from a DoS alert page, you are taken to the mitigation page. For a mitigation that already exists, you can access the mitigation from the mitigation type link that appears in the upper right corner of the DoS alert. For information about the mitigation type link, see "About the information in the header of a DoS alert " on page 68.

2. At the top of the page, click **View Scratchpad** to open the Alert Scratchpad.

   On the TMS Mitigation Status page, **View Scratchpad** is in the header of the Countermeasures pane.
3. To delete a single item, click the X to the left of the item.
4. To delete all of the items in the Traffic Patterns section or the Traffic Elements section, click **Clear All** in the heading of that section.

Copying traffic data from an Alert Scratchpad into a mitigation

When a mitigation is associated with a DoS alert for which you have added traffic data to its Alert Scratchpad, you can access the Alert Scratchpad and its traffic data from the mitigation page. You can then copy and paste the traffic data into the mitigation settings. See "Initiating a Mitigation from a DoS Alert" on page 114.

You can copy traffic data from the Alert Scratchpad into the mitigation settings for the following types of mitigations:

- TMS
- Flow Specification
- Blackhole
- Fingerprint

To copy traffic data from an Alert Scratchpad into the mitigation settings:

1. Navigate to the mitigation page for the DoS alert.

   When you initiate a mitigation from a DoS alert page, you are taken to the mitigation page. For a mitigation that already exists, access it from the **Mitigation** menu.
2. At the top of the page, click **View Scratchpad** to open the Alert Scratchpad.

   On the TMS Mitigation Status page, **View Scratchpad** is in the header of the Countermeasures pane.

   If **View Scratchpad** does not appear on the mitigation page, then you have not added any traffic data to the Alert Scratchpad for the alert associated with the mitigation.
3. In the Alert Scratchpad, copy the traffic data that you want to use in the mitigation.
4. On the mitigation page, navigate to where you want to insert the data and paste it into the appropriate settings box.

Performing a Whois Lookup for an IP Address on a DoS Alert Page

Introduction

You can perform a whois lookup to view ASN information about any source or destination IP address on a DoS alert page. The information that a whois lookup provides includes the company name, contact information, and AS data. You use the (context menu) icon of an IP address to perform a whois lookup.

Performing a whois lookup

You can perform a whois lookup on a source or destination IP address in any of the following tables or windows on a DoS alert page:

- Top Traffic Patterns (5-tuple) table on the **Summary** tab or the **Traffic Details** tab
- Alert Characterization table on the **Summary** tab
- Source and destination addresses tables on the **Traffic Details** tab
- View More Details window for source or destination addresses

  To access the View More Details window, click **View More** below the source and destination addresses tables on the **Traffic Details** tab.

To perform a whois lookup:

1. Hover your mouse pointer over the (context menu) icon of the address for which you want to do a whois lookup.

   In the Top Traffic Patterns (5-tuple) table, the (context menu) icon is to right of the address and appears when you hover your mouse pointer over it. For all the other addresses, the (context menu) icon is to the left of the IP address and becomes more visible when you hover your mouse pointer over it.
2. Click (context menu), and then click **Look Up IP Address (Whois)**.
3. In the Whois Lookup window, you can use the ARIN, RIPE, or APNIC registries to find information about the IP address.

   If you do not select the correct registry, the whois lookup will indicate that you need to use one of the other registries.

Recognizing a Potential DoS Attack

Introduction

The following example workflow describes some of the key traffic data that you can use to determine if a DoS Host alert represents an attack. Most of this same traffic data can also be used to determine if a DoS Profiled Router alert or DoS Profiled Network alert represents an attack.

Example workflow for recognizing an attack from a DoS Host alert

The following example describes how to recognize an attack from a DoS Host alert:

1. Do the following to navigate to the DoS Host alert:
   - On the Alerts Ongoing page (**Alerts > Ongoing**), type `host` in the **Search** box, and then click **Search**.
   - Look for a DoS Host alert with an importance level of High that has been ongoing for more than 5 minutes.

     This type of alert is alarming because of its high importance level and the duration of the attack.
   - Click the ID link of the alert to access the DoS Host Alert page to view more information about the traffic of the alert and to determine if it represents an attack.
2. In the key alert information that is above the Alert Traffic graph on the **Summary** tab, look at the severity percent and the impact data, and do the following.
   - From the severity percent data, make sure the alert is using a reasonable threshold.

     If the threshold is too low, then the alert might represent traffic that does not need your attention.
   - Use the severity percent and impact data combined with your understanding of your network to determine if the alert deserves further attention.

     The severity percent displays the highest single-minute ratio of the rate of the alert traffic to the high severity rate over the lifetime of the alert. The impact displays the bandwidth that an alert consumes in your network and where this impact data was recorded.

   See “About key alert information on the Summary tab” on page 72.
3. In the Alert Traffic graph on the **Summary** tab, look for anything that is unusual about the traffic displayed.

   Base your analysis on your knowledge of normal peaks in your network as well as known events that could cause spikes in the traffic.
4. On the Alert Traffic graph on the **Summary** tab, display the trigger rate for each misuse type that exceeded the trigger rate.

   This trigger rate can then help you determine if the rate of traffic can be explained by known events or if it represents attack traffic. To display the trigger rates, do the following.

   a. Select **Router** from the **View** list, and then click **Update**.

   b. Double-click the selector of a misuse type that exceeds the trigger rate so that only the traffic of that misuse type is displayed.

   c. Click the trigger rate selector for that misuse type to display the trigger rate.

   See “About the Alert Traffic graph on the Summary tab of a DoS Host Alert page” on page 74.

5. Use the following tables on the **Summary** tab to look for additional traffic data that can help you determine if the alert traffic represents an attack, as follows:

Icon	Description
	Hover over this icon to display the download options. The download options that appear vary depending on the page that you are on and c include the following: ■ PDF - Click to download the page in PDF format. ■ XML- Click to download the page in XML format. ■ CSV - Click to download a page in CSV format. The download can be a CSV text file or CSV zip archive file. When the download a zip archive file, “zip archive” is appended to CSV. ■ Excel-XML - Click to download a page in Excel-XML format. The icon appears only if there are multiple options for downloadin page.
	Click to download a page in PDF format. This icon appears only when the PDF format is the only download opti for a page.
	Click to download and email a page as a PDF.

Column	Description
	Select this icon to delete a user account.
Username	A user name as a link to the Edit Existing Account page.

6. If the top traffic patterns do not help you identify a potential attack , then see if you can correlate any unusual data in the tables on the **Traffic Details** tab to identify a potential attack. The following table lists some of the things you can look for in these tables:

Table	Traffic Data to Look For
Source IP Addresses	Unexpected high traffic volumes from a source IP address.
Destination IP Addresses	The IP address where the traffic is going.
Source TCP Ports or Source UDP Ports	Whether the source ports represent normal traffic. **Example:** Normal traffic might be 1-1023 (System) and attack traffic might be 1024-65535 (Dynamic).

Table	Traffic Data to Look For
Destination TCP Ports or DestinationUDP Ports	Whether the levels of the traffic that are sent to these destination ports are normal. **Example:** If you see that the packets that were sent to port TCP 80 were 40 bytes each, then this is an attack and not normal traffic.
Source Countries or Source ASNs	High volumes of traffic from unexpected sources.
TCP Flags	A flag that can help you determine the type of the attack.

Note: You can click on any of these tables to display its data in traffic graph at the top of the **Traffic Details** tab.

See "About the traffic statistics tables on the Traffic Details tab" on page 79.

If you determine that the traffic of an alert represents an attack, you can then add the traffic data to the Alert Scratchpad to use when you configure a mitigation for the attack. See "About the Alert Scratchpad" on page 85.

Deleting Alerts

Introduction

You can use the following methods to delete alerts:

- manually delete alerts on an individual basis on the Delete Alerts page (**Administration > DoS Management > Delete Alerts**)
- schedule Peakflow SP to delete alerts automatically on the Schedule Auto-Deletion of Alerts page (**Administration > DoS Management > Schedule Alert Deletion**)

Important: When you delete an alert, any record of it is removed and the alert cannot be restored.

User access

Only managed services administrators can delete alerts.

Deleting alerts manually

To delete alerts manually:

1. Navigate to the Delete Alerts page (**Administration > DoS Management > Delete Alerts**).
2. Choose one of the following steps:
   - To delete all alerts, select **All Alerts**, and then go to Step 4.
   - To delete specific alerts, select **Alerts Matching**.
3. Choose your next steps based on the matching criteria on which you want to base the deletions:

Matching criteria	Steps
alert ID	Select the **ID** check box, and then type the *alert IDs*.
alert class	Select the **Class** check box, and then in the **Class** list, select the class of alerts to delete.
alert type	Select the type of alerts that you want to delete from the **Type** list. The options that are available depend on the alert class that you select.
alert age	Select the **Age** check box, and then configure the timeframe options for the age of alerts to delete.
alert duration	Select the **Duration** check box, and then configure the timeframe options for the duration of alerts to delete.
alert importance	Select the **Importance** check box, and then select the check boxes for the importance levels of alerts to delete.
affected resource	Select the **Resource** check box, and then type the *affected resource* in the box.

4. Click **Delete**, and then click **Delete** again on the Confirm Delete Alerts page.

Deleting alerts automatically

To schedule the automatic deletion of alerts:

1. Navigate to the Schedule Auto-Deletion of Alerts page (**Administration > DoS Management > Schedule Alert Deletion**).
2. Select the check boxes for the importance levels of alerts that you want to delete (low, medium, or high).
3. For the **older than** settings, follow these steps:
   - In the boxes, type the *number* of days, weeks, or months at which you want the alerts to be deleted.
   - From the lists, select the corresponding timeframes.
4. Click **Save**.

These settings are automatically constrained for the system-wide configuration for managed services users. Peakflow SP might delete alerts that are more recent than what you set here if the system-wide deletion time is less than your setting. Settings that are followed by an asterisk (*) are overridden by the current system-wide configuration.

Chapter 5:
Introduction to TMS Mitigations

Introduction

This section describes general information about using TMS to protect your network against attacks.

For general information about mitigating attacks without using TMS appliances, see Chapter 9: “Other Ways to Mitigate Attacks” on page 191.

User access

Managed services administrators and non-administrative users have access to the mitigation views.

In this section

This section contains the following topics:

About TMS Mitigations

Introduction

You can create a mitigation to filter malicious traffic and permit expected traffic through intelligent filtering devices. Because Peakflow SP provides a robust filtering language and real-time traffic reports, you can precisely define filters and observe their effect. Peakflow SP uses the system's reporting capabilities to monitor the removal of unwanted traffic. This ability, combined with the DoS detection functionality of Peakflow SP, protects your network from attacks.

For a description of all the ways you can mitigate attacks with Peakflow SP, see "Mitigating Attacks Using Peakflow SP" on page 192.

About the TMS Mitigations page

The TMS Mitigations page (**Mitigation > Threat Management**) allows you to do the following:

- Configure and delete TMS mitigations.

  See "Configuring and Deleting TMS Mitigations" on page 110.
- Search for TMS mitigations.

  See "Searching for Mitigations" on page 194.
- View information about TMS mitigations.

  The TMS Mitigations page displays the same information as the Mitigations Ongoing and Mitigations Recent pages, but also includes mitigations that have not started.

  See "About the Mitigations Pages" on page 193.

  Note: Traffic graphs of hardware mitigations might show some traffic (approximately 100-200 bps) even if offramp routes or filters are not active. This is the result of various broadcast packets (ARP, STP, etc.) from the routers and the TMS appliance and is not an issue.

  For information about navigating through multiple pages of TMS Mitigations, see "Navigating multiple pages" on page 15.
- Start or stop TMS mitigations.

  See "Starting and Stopping TMS Mitigations" on page 107.
- Download or email a TMS mitigation report by clicking an icon on the Arbor Smart Bar.

  See "About the Arbor Smart Bar " on page 13.

For more information see the following:

- "About the Mitigations Pages" on page 193
- "About the TMS Mitigation Status Page" on page 100

About TMS Mitigation Countermeasures

Introduction

Countermeasures are defense mechanisms that you can use to target and remove attack traffic so that your network can continue to operate. Different countermeasures are designed to stop different types of attack traffic.

Types of countermeasures

Peakflow SP uses the following types of countermeasures:

Type	Description
Per-packet	This type of countermeasure is applied to every packet that matches the prefix associated with a mitigation. Per-packet countermeasures are processed before event-driven countermeasures. See "Configuring Per-Packet Countermeasures" on page 125.
Event-driven	This type of countermeasure is divided into the following groups: ■ Application-specific stream-based — Peakflow SP identifies the traffic stream with an application ID before it applies the countermeasure. ■ Time-based — Timers detect specific events. For example, the TCP Connection Reset countermeasure drops traffic when a connection remains idle for too long. See "Configuring Event-Driven Countermeasures" on page 157.

About dynamic blacklisting

Blacklisting countermeasures identify source hosts that violate standards of proper behavior and add those IP addresses or flows to a dynamic blacklist. For most countermeasures, all traffic from a blacklisted source is dropped for a default period of one minute. After a source is removed from the blacklist, it remains in the blacklist cache until the cache entry is overwritten with a new blacklist source. If the same source is blacklisted again while it is in the blacklist cache, it is moved to a repeat-offender blacklist for five minutes. Some blacklist timers may be adjusted in advanced settings

Important: Dynamic blacklisting is different from the Black/White Lists countermeasure.

Hardware blacklisting

With a TMS 4000 appliance, the top source hosts that are repeatedly blacklisted as attackers are blocked by hardware blacklisting. Up to 2,044 source hosts can be blocked in this way. This method of blocking individual hosts provides a substantial boost to the TMS 4000 performance because the packets are not passed to a mitigation processor.

Countermeasure processing order

To enforce the configured countermeasures, the TMS appliance collects and views all raw packets. It applies the filters and countermeasures to all of the traffic that the Black/White Lists countermeasure did not explicitly pass or drop.

An IP packet received on a TMS mitigation interface is evaluated for per-packet mitigation countermeasures in the following order:

1. Dynamic Blacklist (set automatically by other countermeasures)

   Note: The TMS appliance does not apply this countermeasure to packets that are truncated or corrupted.
2. Invalid Packets

   This non-configurable countermeasure drops invalid TCP/IP packets. The criteria used to validate packets are listed under this countermeasure on the TMS Mitigation Status page.
3. IP Address Filter Lists

   See "Configuring the IP Address Filter Lists Countermeasure" on page 131.
4. Black/White Lists
   a. Inline Filter
   b. Black/White Filter Lists
   c. Blacklist Fingerprints
5. IP Location Filter Lists
6. Zombie Detection
7. Per Connection Flood Protection
8. TCP SYN Authentication (includes HTTP Authentication)
9. DNS Authentication (except in active mode with DNS Scoping)
10. TCP Connection Limiting
11. TCP Connection Reset (traffic detection only, also happens later)
12. Payload Regular Expression
13. Source /24 Baselines
14. Protocol Baselines
15. DNS Malformed (missing payload check only)
16. SIP Malformed (missing payload check only)
17. Shaping
18. IP Location Policing
19. HTTP Malformed
20. HTTP Scoping
21. HTTP Rate Limiting
22. AIF and HTTP/URL Regular Expression
23. SSL Negotiation
24. SIP Request Limiting

Countermeasures supported for IPv6

The following countermeasures are supported for IPv6:

- Black/White Lists
- Invalid Packets
- IP Address Filter Lists
- TCP SYN Authentication

- Payload Regular Expression
- Zombie Detection

About configuring countermeasure settings

You can configure the countermeasure settings in the following ways:

- When you configure mitigations

  See "Configuring and Deleting TMS Mitigations" on page 110.

- When you view a mitigation on the TMS Mitigation Status page

  When you configure countermeasure settings on the TMS Mitigation Status page, you can see the results in real time. This allows you to refine the mitigations to make them more effective.

  See "About the TMS Mitigation Status Page" on the next page.

About the TMS Mitigation Status Page

Introduction

The TMS Mitigation Status page displays detailed statistics about a mitigation and allows you to edit the countermeasures being applied to a mitigation. The name of the mitigation is appended to the title of the page.

Navigating to the TMS Mitigation Status page

To navigate to the TMS Mitigation Status page:

1. Navigate to one of the following pages:
   - Mitigations Ongoing page (**Mitigation > Ongoing**)
   - Mitigations Recent page (**Mitigation > Recent**)
   - TMS Mitigations page (**Mitigation > Threat Management**)
2. Click the ID link for a TMS mitigation.

About the Summary tab

The Summary tab displays important information about a mitigation. It includes a **Start** button that allows you to start a mitigation and a **Stop** button that allows you to stop an ongoing mitigation. You can also download lists of blocked hosts and add or view annotations.

For additional information see:

- "Mitigation information displayed on the Summary tab" below
- "Editing mitigation settings on the Summary tab" on the facing page
- "Viewing data on the Summary tab traffic graph" on page 102
- "Downloading blocked hosts on the Summary tab" on page 102
- "Adding or viewing comments on the Summary tab" on page 103

Mitigation information displayed on the Summary tab

The Summary tab displays the following information about a TMS mitigation:

Information	Description
Status	The start and end time of the mitigation. If the mitigation is currently active, the end time is replaced with "Ongoing."
Alert	The alert associated with a mitigation, if applicable. The alert's link opens the page that displays information about the alert.
Template	The mitigation template that is used in a mitigation, if applicable. The template's link opens the page where you can edit the template.
Managed Object	The managed object associated with a mitigation. The managed object's link opens the page where you can edit the managed object.
Learning Dataset	The learning dataset that is selected for the mitigation.

Information	Description
TMS Group	The group to which a TMS appliance belongs. The TMS group's link opens the page where you can edit the TMS group.
Offramp Prefixes	The offramp prefixes that are protected by this mitigation.
Flow Specification settings	Flow specification settings appear if they were configured for a mitigation. You can configure flow specification settings only when you use flow specification to offramp traffic. For more information about the flow specification settings, see "Flow specification filter settings" on page 118.
Traffic graph	A graph of the traffic that is involved in a mitigation, based on impact data. See "Viewing data on the Summary tab traffic graph" on the next page.
Traffic data table	A table of mitigation traffic data. Dropped traffic is displayed in red, and passed traffic is displayed in green.
Comment / Annotation list	A list of the three most recent comments (annotations) that are applied to a mitigation. You can click the Show All button to view additional comments (annotations) that are applied to a mitigation. See "Adding Annotations to a Mitigation" on page 197.

Editing mitigation settings on the Summary tab

To edit mitigation settings on the Summary tab:

1. Navigate to the TMS Mitigation Status page.
   See "Navigating to the TMS Mitigation Status page" on the previous page.
2. In the upper-right corner of the Summary tab, click **Edit**.
3. Use the following table to edit the mitigation settings:

Setting	Procedure
Learning Dataset list	Select the learning dataset to apply to a mitigation. Peakflow SP displays only the learning datasets for the managed object that is selected in the mitigation.
Offramp Prefixes box	Type the *prefixes*, in CIDR notation, to specify one or more address ranges to be offramped. The match criteria for the managed object that is selected for the mitigation limit the prefixes that you can offramp.
Flow Specification Filters boxes	(IPv4 only) Configure any flow specification filters that you want to use in a mitigation. You can use flow specification filters only when you use flow specification to offramp traffic. See "Flow specification filter settings" on page 118.

4. To edit additional mitigation settings, click the **Edit Full Configuration** link in the lower-right corner of the Summary tab.

   Note: If you have unsaved changes on the Summary tab, they are lost when you click **Edit Full Configuration**.
5. To undo your unsaved edits, click **Cancel** in the upper-right corner of the Summary tab.
6. To save your edits, click **Save** in the lower right corner of the Summary tab.

Viewing data on the Summary tab traffic graph

To view data on the Summary tab traffic graph:

1. Navigate to the TMS Mitigation Status page.

   See "Navigating to the TMS Mitigation Status page" on page 100.
2. On the Summary tab, click the following tabs on the top left of the graph for the breakdown of data that you want to view:

Tab	Description
Total tab	This tab displays dropped and passed traffic totals for all of the TMS appliances that are involved in the mitigation.
Per TMS tab	This tab displays the traffic that each TMS appliance drops. The graph on the **Per TMS** tab can display up to ten TMS appliances. Each TMS appliance that is involved in the mitigation is also listed below the graph. In this list, each TMS appliance is preceded by a small graph for that TMS appliance's traffic. Hover your mouse pointer over the smaller graph to display a larger version of the graph. Each of the first ten graphs in this list has a different color that corresponds to the colors in the graph above. If the list has more than ten TMS appliances, those additional TMS appliances have grayscale graphs.
Per Countermeasure tab	This tab displays the traffic that each countermeasure drops.

3. Click the tab on the top right of the graph for the unit of measure by which you want to view traffic data.
4. Click the tab on the bottom left of the graph for the timeframe of data that you want to view.

Downloading blocked hosts on the Summary tab

The following countermeasures can add blocked hosts to a mitigation:

- AIF and HTTP/URL Regular Expression
- DNS NXDomain Rate Limiting
- DNS Rate Limiting
- HTTP Malformed
- HTTP Rate Limiting
- Payload Regular Expression

- SIP Malformed
- SIP Request Limiting
- SSL Negotiation
- TCP Connection Limiting
- TCP Connection Reset
- Zombie Detection

To download blocked hosts, do one of the following:

- To download and save a list of the hosts that a mitigation is currently blocking, click **Download Blocked Hosts**.
- To download and save a list of the hosts that the mitigation blocked the most often, click **Download Top Blocked Hosts**.

Adding or viewing comments on the Summary tab

To add a comment (annotation) to a mitigation on the Summary tab:

- Click **Add Comment**.

To view all of the comments (annotations) that are applied to a mitigation on the Summary tab:

- Click **Show All**.

About the Countermeasures tab

The TMS Mitigation Status page Countermeasures tab displays the following information:

Icon	Description
	Hover over this icon to display the download options. The download options that appear vary depending on the page that you are on and can include the following: ■ PDF - Click to download the page in PDF format. ■ XML- Click to download the page in XML format. ■ CSV - Click to download a page in CSV format. The download can be a CSV text file or CSV zip archive file. When the download is a zip archive file, "zip archive" is appended to CSV. ■ Excel-XML - Click to download a page in Excel-XML format. The icon appears only if there are multiple options for downloading a page.
	Click to download a page in PDF format. This icon appears only when the PDF format is the only download option for a page.
	Click to download and email a page as a PDF.

Column	Description
	Select this icon to delete a user account.
Username	A user name as a link to the Edit Existing Account page.
Real Name	A user's full name.
Account Group	The account group to which a user belongs.
Capability Level	A user's capability level, which is either an administrator or a user.
Email	A user's email address.
Device	The SP appliance with which a user is associated. The SP appliance is

For more information about countermeasures and how to configure them, see "About TMS Mitigation Countermeasures" on page 97.

Countermeasures on the Countermeasures tab

The Countermeasures tab on the TMS Mitigation Status page has the following countermeasures:

- AIF and HTTP/URL Regular Expression

  See "Configuring the AIF and HTTP/URL Regular Expression Countermeasure" on page 158.
- Black/White Lists

  See "Configuring the Black/White Lists Countermeasure" on page 126.
- DNS Authentication

See “Configuring the DNS Authentication Countermeasure” on page 129.

- DNS Malformed

  See “Configuring the DNS Malformed Countermeasure” on page 163.
- DNS NXDomain Rate Limiting

  See “Configuring the DNS NXDomain Rate Limiting Countermeasure” on page 164.
- DNS Rate Limiting

  See “Configuring the DNS Rate Limiting Countermeasure” on page 166.
- DNS Regular Expression

  See “Configuring the DNS Regular Expression Countermeasure” on page 168.
- DNS Scoping

  See “Configuring Advanced Settings for TMS Mitigations” on page 121.
- HTTP Malformed

  See “Configuring the HTTP Malformed Countermeasure” on page 174.
- HTTP Rate Limiting

  See “Configuring the HTTP Rate Limiting Countermeasure” on page 176.
- HTTP Scoping

  See “Configuring Advanced Settings for TMS Mitigations” on page 121.
- Invalid Packets
- IP Address Filter Lists

  See “Configuring the IP Address Filter Lists Countermeasure” on page 131.
- IP Location Filter Lists

  See “Configuring the IP Location Filter Lists Countermeasure” on page 133.
- IP Location Policing

  See “Configuring the IP Location Policing Countermeasure ” on page 135.
- Payload Regular Expression

  See “Configuring the Payload Regular Expression Countermeasure” on page 138.
- Per Connection Flood Protection

  See “Configuring the Per Connection Flood Protection Countermeasure ” on page 141.
- Protocol Baselines

  See “Configuring the Protocol Baselines Countermeasure” on page 144.
- Source /24 Baselines

  See “Configuring the Source /24 Baselines Countermeasure” on page 148.
- Shaping

  See “Configuring the Shaping Countermeasure” on page 146.
- SIP Malformed

  See “Configuring the SIP Malformed Countermeasure” on page 178.
- SIP Request Limiting

  See “Configuring the SIP Request Limiting Countermeasure” on page 180.
- SSL Negotiation

  See “Configuring the SSL Negotiation Countermeasure” on page 182.

- TCP Connection Limiting
  See “Configuring the TCP Connection Limiting Countermeasure” on page 185.
- TCP Connection Reset
  See “Configuring the TCP Connection Reset Countermeasure” on page 188.
- TCP SYN Authentication
  See “Configuring the TCP SYN Authentication Countermeasure” on page 150.
- Zombie Detection
  See “Configuring the Zombie Detection Countermeasure ” on page 154.

Starting and Stopping TMS Mitigations

Introduction

TMS mitigations filter malicious traffic and permit expected traffic through a TMS appliance. To enable a specific mitigation, you must start it. You can disable a mitigation by stopping it.

Depending on the failure settings on the TMS appliance, a mitigation can be suspended due to a system failure. When a mitigation is suspended, then an attack might go unmitigated until you manually restart the mitigation.

See "About TMS Mitigations" on page 96.

Starting and restarting TMS mitigations

You can start or restart a TMS mitigation from any of the following pages:

Page	Procedure
TMS Mitigations page (**Mitigation > Threat Management**)	Select the check boxes next to the TMS mitigations to start, and then click **Start** or **Restart**.
Create TMS Mitigation page (**Mitigation > Threat Management > Add Mitigation**) or Edit TMS Mitigation page (**Mitigation > Threat Management > mitigation name link**)	After you created or edited the mitigation, click **Save and Start**.
TMS Mitigation Status page (**Mitigation > Threat Management >** *mitigation name link*)	On the Summary tab, click **Start**.

Stopping TMS mitigations

To stop a TMS mitigation:

1. Navigate to the Threat Management Mitigations page (**Mitigation > Threat Management**).
2. Choose one of the following steps:
   - Click the name of the TMS mitigation to stop, and then click **Stop**.
   - Select the check boxes next to the TMS mitigations that you want to stop, and then click **Stop**.

Chapter 6:
Configuring TMS Mitigations

Introduction

This section describes how to configure TMS mitigations to filter malicious traffic and allow legitimate traffic.

For an introduction to TMS mitigations, see Chapter 5: “Introduction to TMS Mitigations” on page 95.

User access

Only managed services administrators can configure these settings.

In this section

This section contains the following topics:

Configuring and Deleting TMS Mitigations

Introduction

You can configure TMS mitigations to filter malicious traffic and to allow legitimate traffic through TMS appliances.

You can configure mitigations on the TMS Mitigations page (**Mitigation > Threat Management**). See "About TMS Mitigations" on page 96.

You can also initiate a TMS Mitigation from a DoS Alert. See "Initiating a Mitigation from a DoS Alert" on page 114.

Adding and editing a TMS mitigation

To add or edit a TMS mitigation

1. Navigate to the TMS Mitigations page (**Mitigation > Threat Management**).
2. Do one of the following:
   - To add a mitigation, click **Add Mitigation**, and then click **IPv4** or **IPv6** for the IP version of the traffic that you want to mitigate.
   - To edit a mitigation, click (edit).
3. Configure the settings on the tabs of the Create TMS Mitigation page or the Edit TMS Mitigation page.

   See "Tabs on the Create TMS Mitigation page or Edit TMS Mitigation page" below.
4. Choose one of the following steps:
   - To save the mitigation without starting it, click **Save**.
   - To save and start the mitigation, click **Save And Start**.

Tabs on the Create TMS Mitigation page or Edit TMS Mitigation page

The following tabs appear on the Create TMS Mitigation page and the Edit TMS Mitigation page:

Tab	Description
Mitigation	Allows you to configure the basic description settings for a mitigation. See "Configuring Basic Identification Settings for TMS Mitigations" on page 115.
Protect	Allows you to specify the managed object and address ranges that you want to protect in a mitigation. See "Configuring Protect Settings for TMS Mitigations" on page 117.
TMS Appliances	Allows you to configure the TMS appliances to use to mitigate alert traffic. See "Configuring TMS Appliance Settings for TMS Mitigations" on page 120.

Tab	Description
Black/White Lists	Allows you to configure the following lists: ■ A black list that uses custom or system-defined fingerprints to designate which traffic to drop. ■ A white list that uses custom fingerprints to designate which traffic to pass. See "Configuring the Black/White Lists Countermeasure" on page 126.
IP Based Filter List	Allows you to create IP-based filter lists that specify which IP addresses to drop or pass in a mitigation. See "Configuring the IP Address Filter Lists Countermeasure" on page 131. See "Configuring the IP Location Filter Lists Countermeasure" on page 133..
Payload	Allows you to configure the TMS appliance to drop TCP or UDP traffic that either matches or does not match a Payload regex (regular expression), HTTP Header or request regex, or a DNS request regex. The regular expressions use PCRE syntax. See "Configuring the AIF and HTTP/URL Regular Expression Countermeasure" on page 158. See "Configuring the DNS Regular Expression Countermeasure" on page 168. See "Configuring the Payload Regular Expression Countermeasure" on page 138.

Tab	Description
Countermeasures	Allows you to configure the filters and other settings that allow Peakflow TMS to block the effects of malicious traffic. For additional information, see: ■ "Configuring the DNS Authentication Countermeasure" on page 129 ■ "Configuring the DNS Malformed Countermeasure" on page 163 ■ "Configuring the DNS NXDomain Rate Limiting Countermeasure" on page 164 ■ "Configuring the DNS Rate Limiting Countermeasure" on page 166 ■ "Configuring the HTTP Malformed Countermeasure" on page 174 ■ "Configuring the HTTP Rate Limiting Countermeasure" on page 176 ■ "Configuring the IP Location Policing Countermeasure " on page 135 ■ See "Configuring the Per Connection Flood Protection Countermeasure " on page 141. ■ "Configuring the SIP Malformed Countermeasure" on page 178 ■ "Configuring the SIP Request Limiting Countermeasure" on page 180 ■ "Configuring the SSL Negotiation Countermeasure" on page 182 ■ "Configuring the TCP Connection Limiting Countermeasure" on page 185 ■ "Configuring the TCP Connection Reset Countermeasure" on page 188 ■ "Configuring the TCP SYN Authentication Countermeasure" on page 150 ■ "Configuring the Zombie Detection Countermeasure " on page 154 ■ "Configuring the Source /24 Baselines Countermeasure" on page 148 ■ "Configuring the Protocol Baselines Countermeasure" on page 144
Shaping	(IPv4 mitigations only) Allows you to set filters for any IPv4 traffic that remains after all other countermeasures and filters are applied. See "Configuring the Shaping Countermeasure" on page 146.
Advanced	Allows you to configure proxy list threshold exceptions. See "Configuring Advanced Settings for TMS Mitigations" on page 121.

Deleting TMS mitigations

When you delete a mitigation configuration, Peakflow SP also deletes all of its associated mitigation events.

To delete a TMS mitigation:

1. If the mitigation that you want to delete is running, then stop it.

   See "Starting and Stopping TMS Mitigations" on page 107.
2. To delete a mitigation, navigate to the Threat Management Mitigations page (**Mitigation > Threat Management**).
3. Select the check boxes next to the TMS mitigations to delete, and then click **Delete**.

Initiating a Mitigation from a DoS Alert

Introduction

After Peakflow SP generates a DoS alert, you can analyze the traffic data in the alert, and then, if needed, initiate a mitigation to reduce or stop the impact of an attack. You can initiate a mitigation from a DoS Host Alert page, a DoS Profiled Router Alert page, and a DoS Profiled Network Alert page.

For general mitigation information, see "Mitigating Attacks Using Peakflow SP" on page 192.

Initiating a Mitigation from a DoS alert

To initiate a mitigation from a DoS alert:

1. Navigate to the alert for which you want to start a mitigation, as follows:
   - DoS Host Alert (**Alerts > All Alerts >** *DoS Host alert ID link*)
   - DoS Profiled Router Alert (**Alerts > All Alerts >** *DoS Profiled Router alert ID link*)
   - DoS Profiled Network Alert (**Alerts > All Alerts >** *DoS Profiled Network alert ID link*)
2. Click **Mitigate Alert**, and then click the type of mitigation that you want to perform.
3. On the configuration page that appears, configure the settings for the mitigation.

   Based on the type of mitigation that you select, see one of the following topics:
   - Threat Management — "Configuring and Deleting TMS Mitigations" on page 110
   - Generate Filter — "Mitigating Using ACL Filters" on page 199
   - Blackhole — "Mitigating Using Blackhole Routing" on page 201

Configuring Basic Identification Settings for TMS Mitigations

Introduction

When you configure a TMS mitigation, you can use the **Mitigation** tab to define the basic identification settings.

For information about adding and editing mitigations, see "Configuring and Deleting TMS Mitigations" on page 110.

Configuring basic identification settings for TMS mitigations

To configure basic identification settings for TMS mitigations:

1. Do one of the following:
   - Navigate to the **Mitigation** tab.
     See "Adding and editing a TMS mitigation" on page 110.
   - Navigate to the TMS Mitigation Status page. On the Summary tab, click **Edit**.
     See "Navigating to the TMS Mitigation Status page" on page 100.
2. Configure the basic identification settings.
   See "Identification settings on the Mitigation tab" below.
3. Click **Save**.

Identification settings on the Mitigation tab

Use the following table to configure the settings on the **Mitigation** tab:

Setting	Procedure
Name box	Type a *unique name* for the mitigation.
Source Alert ID (Optional) box	Type the *ID number for the alert* from which the mitigation was created. Peakflow SP populates this box if the mitigation is generated from a DoS alert. If this mitigation is not associated with an existing DoS alert, then leave this box empty.
Description box	Type a *description* that can help to easily identify this mitigation in a list.
Learning Dataset list	(Optional) Select the learning dataset to apply to a mitigation. Peakflow SP displays only the learning datasets for the managed object that is selected in the mitigation. **Note:** You can also edit this setting on the Summary tab of the TMS Mitigation Status page.
Enable CDN Proxy Support check box	Select to prevent the blacklisting of a content delivery network (CDN) proxy. This setting is a global setting that applies to all countermeasures in a mitigation that can blacklist a source IP address. See "About CDN proxy support" on the next page.

About CDN proxy support

Peakflow SP countermeasures can blacklist an attacker's IP address. When traffic is routed through a CDN proxy, the source IP address of that traffic is the IP address of the last CDN proxy device. That source IP address is shared by all of the users whose traffic passes through that device. Therefore, the countermeasure settings that blacklist an attacker's IP address might blacklist all traffic from the CDN proxy.

When you enable CDN Proxy Support, you can prevent the blacklisting of a CDN proxy. Peakflow SP then uses the countermeasures of the mitigation to block just the malicious traffic from a CDN proxy.

The following countermeasures modify blacklist behavior for detected proxy hosts:

- HTTP Malformed
- SIP Malformed
- SSL Negotiation
- DNS Regular Expression
  - UDP DNS flows are dropped but not blacklisted.
  - TCP DNS flows are blacklisted regardless of whether the countermeasure's Blacklist on Blocked setting is selected.
- HTTP Regular Expressions
  - Includes AIF filters.
  - Flows are blacklisted regardless of whether the countermeasure's Blacklist on Blocked setting is selected.

Source IP addresses of detected proxies are exempted from the following rate-based countermeasures:

- DNS Rate Limiting
- DNS NXDomain Rate Limiting
- HTTP Rate Limiting
- Protocol Baselines
- SIP Request Limiting
- Source /24 Baselines
- TCP Connection Limiting
- TCP Connection Reset
- Zombie Detection

Note: Proxy hosts are not exempted from rate-based countermeasures such as IP Location Policing and Shaping that are designed to regulate summary traffic rather than individual sources.

In cases where Peakflow TMS cannot automatically detect proxies with CDN Proxy Support, Proxy List Threshold Exceptions may be useful. For information about Proxy List Threshold Exceptions, see "Configuring Advanced Settings for TMS Mitigations" on page 121.

About locked settings

When  (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring Protect Settings for TMS Mitigations

Introduction

You can configure the protect settings to define the managed object and address ranges that you want to protect in mitigations. With an IPv4 TMS mitigation that uses flow specification to offramp traffic, you can also use the protect settings to define additional flow specification filter settings.

You can use the **Protect** tab to define the protect settings when you configure a TMS mitigation.

For information about adding and editing mitigations, see “Adding and editing a TMS mitigation” on page 110.

Configuring protect settings

To configure protect settings for TMS mitigations:

1. Do one of the following:
   - Navigate to the **Protect** tab.
     See “Adding and editing a TMS mitigation” on page 110.
   - Navigate to the TMS Mitigation Status page. On the Summary tab, click **Edit**.
     See “Navigating to the TMS Mitigation Status page” on page 100.
2. Configure the protect settings:
   See “Protect settings” on the next page.
3. Click **Save**.

Protect settings

Use the following table to configure the protect settings:

Setting	Procedure
Managed Object box	Click **Select Managed Object** and in the Select a Managed Object window, select a managed object for one of the following tasks: ■ Use the managed object to obtain baseline information if you plan to enable Protocol or Source /24 countermeasures ■ Control user access to a mitigation if the managed object involves a resource group's managed object The match criteria for the managed object that you select limit the prefixes that you can type in the Offramp Prefixes box below. You can also search for a managed object by using the Search options.
Offramp Prefixes box	Type the *prefixes*, in CIDR notation, to specify one or more address ranges to be offramped. The match criteria for the managed object that you selected limit the prefixes that you can offramp. **Note:** You can also edit this setting on the Summary tab of the TMS Mitigation Status page.
Timeout box	Type the *number of seconds* that a TMS mitigation should run before it stops automatically.
Flow Specification Filters boxes	(Mitigations and IPv4 only) Configure any flow specification filters that you want to use in a mitigation. You can use flow specification filters only when you use flow specification to offramp traffic. See "Flow specification filter settings" below.

Flow specification filter settings

When you use flow specification to offramp IPv4 traffic into a VPN that is tied to a TMS infrastructure, you can also configure several flow specification filter settings.

You can edit the flow specification filter settings on the **Protect** tab of a mitigation or on the Summary tab of the TMS Mitigation Status page.

Use the following table to configure the flow specification filter settings:

Setting	Procedure
Protocol Numbers box	To filter on packets using protocol numbers, type the *protocol numbers* or *ranges* to match. For example, 6 or 10-20.
Source Prefix box	To filter on packets using the source prefix, type the *source CIDR block* to match.

Setting	Procedure
Match any specified source ports AND any specified destination ports option	If you select this option, then the **Source Ports** and **Destination Ports** boxes appear. You can then configure a flow specification filter that uses source ports and destination ports. **Note:** You do not have to specify both source ports and destination ports. For example, if you leave the **Source Ports** box blank, then only the destination ports will be considered as part of the match.
Match any specified ports option	If you select this option, then the **Destination OR Source Ports** box appears. You can then configure a flow specification filter that uses either source ports or destination ports.
Source Ports box	To filter on packets using the source port of the packets, type the *source port numbers* or *ranges* to match. For example, 80 or 24-30.
Destination Ports box	To filter on packets using the destination port of the packets, type the *destination port numbers* or *ranges* to match. For example, 80 or 24-30.
Destination OR Source Ports box	To filter on packets using either the source or destination port of the packets, type the *destination* or *source port numbers* or *ranges* to match.

About locked settings

When (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring TMS Appliance Settings for TMS Mitigations

Introduction

You can use the **TMS Appliances** tab to designate which TMS appliances to use for mitigation when you configure a TMS mitigation. You can include all TMS appliance groups or a specific group.

For information about adding and editing mitigations, see "Configuring and Deleting TMS Mitigations" on page 110.

Configuring the TMS appliance settings for TMS mitigations

To configure the TMS appliances settings for TMS mitigations:

1. Do one of the following:
   - Navigate to the **TMS Appliances** tab.
     See "Adding and editing a TMS mitigation" on page 110.
   - Navigate to the TMS Mitigation Status page. On the Summary tab, click **Edit**.
     See "Navigating to the TMS Mitigation Status page" on page 100.
2. Use the following table to configure the TMS appliances settings:

Setting	Procedure
TMS Group list	Select one of the following options: ■ **All** to include all TMS groups ■ A group name to configure a specific group
Announce Route check box	Select to allow Peakflow SP to announce BGP or flow specification routes to peering routers when a mitigation starts. See "About the Announce Route setting" below.

3. Click **Save**.

About the Announce Route setting

The TMS appliances that are involved in the mitigation determine the default setting of the Announce Route check box, as follows:

- If any appliances are used in an offramp deployment, then Announce Route is selected.
  In an offramp deployment scenario, Peakflow SP redirects anomalous traffic to the TMS appliance using a BGP or flow specification route announcement.
- If any appliances are used in an inline or portspan deployment, Announce Route is cleared.
  In an inline deployment scenario, the traffic typically flows through the TMS appliance before the mitigation starts so that the TMS appliance can begin filtering the traffic.

The default is set when you select or change the appliance group or the appliances that are selected for the Other group.

About locked settings

When  (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring Advanced Settings for TMS Mitigations

Introduction

The advanced settings allow you to configure proxy list threshold scaling and DNS and HTTP scoping for a mitigation. The proxy list threshold scaling applies to the Zombie Detection, DNS Rate Limiting, HTTP Rate Limiting, and SIP Request Limiting countermeasures.

You configure the advanced settings on the **Advanced** tab when you configure a TMS mitigation. You can also configure the DNS scoping settings and the HTTP scoping settings on the Countermeasures tab on the TMS Mitigation Status page.

For information about adding and editing mitigations, see "Configuring and Deleting TMS Mitigations" on page 110.

For information about adjusting these and other mitigation settings on the TMS Mitigation Status page, see "About the TMS Mitigation Status Page" on page 100.

About the advanced settings

The **Advanced Settings** tab allows you to configure the following TMS mitigation settings:

Feature	Description
Proxy List Threshold Exceptions	Uses a specified scaling factor to scale the countermeasure threshold rates for IPv4 or IPv6 traffic that is sourced from proxies. These scaled threshold rates are applied in place of the configured traffic rates. This feature allows you to manage known proxies, network address translation (NAT) locations, and sources that consistently maintain a much higher traffic level than that of a typical /32.
DNS Scoping	Limits the application of DNS countermeasures to specific IPv4 domains.
HTTP Scoping	Limits the application of HTTP countermeasures to specific HTTP IPv4 requests. For example, you can use scoping to apply HTTP countermeasures to a virtual server that resides on a shared resource.

Configuring advanced TMS mitigation settings

To configure the proxy list threshold exceptions settings for IPv4 or IPv6 TMS mitigations:

1. Do one of the following:
   - Navigate to the **Advanced** tab of the mitigation.

     See "Adding and editing a TMS mitigation" on page 110.
   - Navigate to the TMS Mitigation Status page. On the Countermeasures tab, click (expand) for the DNS Scoping and the HTTP Scoping countermeasures.

     See "Navigating to the TMS Mitigation Status page" on page 100.

     Note: On the TMS Mitigation Status page, you can only edit countermeasure settings that are not locked. You must also be in an account group that has been assigned the capability to edit mitigations.
2. Use the following table to configure the proxy list threshold exceptions settings for IPv4 or

IPv6 TMS mitigations:

Setting	Procedure
Enable Proxy List Threshold Exceptions check box	Select to enable these settings.
Proxy Source CIDRs box	Type a list of IPv4 or IPv6 *CIDR blocks* to scope. The address version of the CIDR blocks that you type must correspond to the address version of the mitigation that you are configuring. For example, if you are configuring an IPv6 mitigation, then you can only type IPv6 CIDRs.
Scaling Factor box	Type the estimated *number* of hosts that are behind the proxy.

3. Use the following table to configure the DNS scoping settings for IPv4 TMS mitigations:

Setting	Procedure
Enable DNS Scoping check box	Select to enable these settings. If a traffic flow is determined to be within scope, then Peakflow SP applies countermeasures normally. If a traffic flow is out of scope, then Peakflow SP does not apply the countermeasures.
DNS Scoping Regular Expressions boxes	Type up to five *regular expressions* (in PCRE format and single-line mode) for the domains to which you want to limit the DNS countermeasures. **Note:** DNS scoping regular expressions are case-insensitive by default. To perform case-sensitive matching, preface the expression with "(?-i)".
DNS Scoping Action list	Select whether to apply the countermeasure to matched or unmatched traffic

4. Use the following table to configure the HTTP scoping settings for IPv4 TMS mitigations:

Setting	Procedure
Enable HTTP Scoping check box	Select to enable these settings. If a traffic flow is determined to be within scope, then Peakflow SP applies countermeasures normally. If a traffic flow is out of scope, then Peakflow SP does not apply the countermeasures.
HTTP Scoping Regular Expressions boxes	Type up to five *regular expressions* (in PCRE format and single-line mode) for the domains to which you want to limit the HTTP countermeasures. **Note:** HTTP scoping regular expressions are case-sensitive by default. To perform case-insensitive matching, preface the expression with "(?i)".
HTTP Scoping Action list	Select whether to apply the countermeasure to matched or unmatched traffic.

5. Click **Save**.

About locked settings

When  (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Chapter 7:

Configuring Per-Packet Countermeasures

Introduction

This section describes how to configure the per-packet countermeasures for TMS mitigations.

For information about the types of countermeasures, see “Types of countermeasures” on page 97.

User access

Only managed services administrators can configure these settings.

In this section

This section contains the following topics:

Configuring the Black/White Lists Countermeasure

Introduction

The Black/White Lists countermeasure uses a configurable list of fingerprint expression filters to drop or pass traffic without additional scrutiny. This countermeasure can also be configured to blacklist every source host whose traffic is dropped by this countermeasure.

This countermeasure is applied near the beginning of the mitigation process, and it excludes traffic flows from additional mitigation processing. For information about the order in which the countermeasures are processed, see "Countermeasure processing order" on page 97.

When you configure a TMS mitigation, you can use the **Black/White Lists** tab to configure the Black/White Lists countermeasure to mitigate IPv4 or IPv6 traffic. You can also configure the Black/White Lists countermeasure on the Countermeasures tab on the TMS Mitigation Status page.

For information about adding and editing mitigations, see "Configuring and Deleting TMS Mitigations" on page 110.

For information about adjusting these and other mitigation settings on the TMS Mitigation Status page, see "About the TMS Mitigation Status Page" on page 100.

About the Black/White Lists countermeasure

The Black/White Lists countermeasure evaluates and processes the traffic as follows:

- Any traffic that matches a drop statement is dropped immediately. No additional statements are checked, and no additional countermeasures are applied.
- Any traffic that matches a pass statement is passed immediately. No additional statements are checked, and no additional countermeasures are applied.
- Any traffic that does not match either a drop statement or a pass statement is passed to the remaining countermeasures.

Use the Black/White Lists countermeasure to mitigate based on specific situations. For example, if the mitigation protects a server group that obtains content from other sources, then add the connections to those other sources to a pass rule. Because you already know that those connections are legitimate, you can exempt them from mitigation countermeasures.

Alternatively, if bandwidth is consumed by legitimate-appearing traffic of a type that is not used by the protected hosts, add those traffic characteristics to a drop rule. For example, you might drop any DNS traffic that is directed at Web servers.

You can also use drop rules along with the **Blacklist Matching Addresses** option to consistently drop traffic from the hosts or networks that have been identified as chronic offenders. This eliminates the need to continually re-evaluate whether to allow that traffic.

Example: IPv4 Black/White Lists settings

If you want to block all TCP/22 SSH traffic on your network except for a select block of addresses, you can type the following fingerprint expression filter:

```
pass port 22 and src 10.0.1.0/24
drop port 22
```

All port 22 traffic from 10.0.1.0/24 is automatically whitelisted, and all other port 22 traffic is automatically dropped. To exempt that net block from the countermeasures that you enabled, type the following fingerprint expression filter:

```
pass 10.0.1.0/24
```

No additional filtering or countermeasures are applied to traffic going to or from that block.

Example: IPv6 Black/White Lists settings

If you want to block all TCP/22 SSH traffic on your network except for a select block of addresses, you can type the following fingerprint expression filter:

```
pass port 22 and src 2001:DB8:FF00::/40
drop port 22
```

All port 22 traffic from 2001:DB8:FF00::/40 is automatically whitelisted, and all other port 22 traffic is automatically dropped. To exempt that net block from the countermeasures that you enabled, type the following fingerprint expression filter:

```
pass 2001:DB8:FF00::/40
```

No additional filtering or countermeasures are applied to traffic going to or from that block.

Configuring the Black/White Lists countermeasure

To configure the Black/White Lists countermeasure for TMS mitigations:

1. Do one of the following:
   - Navigate to the **Black/White Lists** tab of the mitigation.

     See “Adding and editing a TMS mitigation” on page 110.
   - Navigate to the TMS Mitigation Status page. On the Countermeasures tab, click (expand) for the Black/White Lists countermeasure.

     See “Navigating to the TMS Mitigation Status page” on page 100.

     Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations.

2. Use the following table to configure the Black/White Lists settings:

Setting	Procedure
Inline Filters box	To add a custom Black/White list FCAP filter for IPv4 or IPv6 traffic, choose one of the following steps: ■ Type a *`fingerprint expression`* that corresponds to the traffic that you want to match. ■ (IPv4 only) Click **Open FCAP Wizard** to add a fingerprint expression by using the FCAP Wizard. See "Using the FCAP Wizard" on page 18.
Black/White Filter Lists box	To add a custom Black/White list filter for IPv4 or IPv6 traffic: a. Click **Select Filter List**. b. Select the filter lists to apply to the mitigation, and then click **OK**. c. In the **Black/White Filter Lists** box, drag the lists to arrange them in the order in which they should be applied.
Blacklist Fingerprints box (IPv4 only)	To add a configured fingerprint to blacklist IPv4 traffic: a. Click **Select Fingerprint**. b. Select the fingerprints to blacklist, and then click **OK**. c. In the **Blacklist Fingerprints** box, drag the fingerprints to arrange them in the order in which they should be applied.
Blacklist Sources check box	Select to blacklist any source host that sends traffic that is dropped because it matches the Black/White Filter Lists. Peakflow SP then drops all the traffic from these source hosts including traffic that matches a pass rule in one of the mitigation's filters. For information about how blacklisting works, see "About dynamic blacklisting" on page 97.
Test Filter button (TMS Mitigation Status page only)	Click to test the effectiveness of an inline filter in mitigating the attack associated with this alert.
View Rates button (TMS Mitigation Status page only)	Click to view all filter list and blacklist fingerprint rate information.

3. Click **Save**.

Peakflow SP's integrated filtering feature supports filtering on source and destination addresses and ports, protocol, TCP flags, ICMP type and code, ToS, and TTL. You can configure the settings for the filters when you create or edit a mitigation.

About locked settings

When  (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring the DNS Authentication Countermeasure

Introduction

The DNS Authentication countermeasure authenticates DNS requests before they reach the DNS server and drops the requests that cannot be authenticated within a specified time. The DNS Authentication countermeasure mitigates IPv4 attack traffic.

You can configure the DNS Authentication countermeasure when you create or edit a mitigation and when you edit a mitigation on the TMS Mitigation Status page. For information about configuring mitigations, see "Configuring and Deleting TMS Mitigations" on page 110. For information about the TMS Mitigation Status page, see "About the TMS Mitigation Status Page" on page 100.

About the DNS Authentication countermeasure

This countermeasure can protect your network against spoof attacks, which occur when an attacker spoofs multiple source addresses in an attempt to overload a DNS server with queries. This countermeasure filters traffic at the packet level.

Configuring the DNS Authentication countermeasure when adding or editing a mitigation

To configure the DNS Authentication countermeasure when adding or editing a mitigation:

1. Navigate to the **Countermeasures** tab of the mitigation.
   See "Adding and editing a TMS mitigation" on page 110.
2. Scroll to the DNS Mitigation section.
3. Configure the settings for the DNS Authentication countermeasure.
   See "DNS Authentication countermeasure settings" on the next page.
4. Click **Save**.

Configuring the DNS Authentication countermeasure on the TMS Mitigation Status page

To configure the DNS Authentication countermeasure on the TMS Mitigations Status page:

1. Navigate to the TMS Mitigation Status page.
   See "Navigating to the TMS Mitigation Status page" on page 100.
2. On the Countermeasures tab, click ⊞ (expand) for the DNS Authentication countermeasure.
   Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations.
3. Configure the settings for the DNS Authentication countermeasure.
   See "DNS Authentication countermeasure settings" on the next page.
4. Click **Save**.

DNS Authentication countermeasure settings

Use the following table to configure the DNS Authentication countermeasure settings:

Setting	Procedure
Enable DNS Authentication check box	Select to enable this countermeasure.
Protection Mode list	Select one of the following modes in which to enable DNS authentication: ■ **Passive** The TMS appliance forces any new UDP DNS queries from a host on port 53 to authenticate within the timeframe configured in the **DNS Authentication Timeout** box. For a period of time after a host is authenticated, subsequent valid DNS queries from that host are passed through unhindered. Passive mode can protect any type of DNS server. ■ **Active UDP** The TMS appliance intercepts DNS queries before they reach an authoritative DNS server. Then the TMS appliance issues a challenge to the client to verify that those queries are valid before it passes the traffic to the original server. Active UDP mode can only protect authoritative DNS servers. ■ **Active TCP** The TMS appliance forces any clients sending DNS queries to respond with a TCP DNS request within the timeframe configured in the **DNS Authentication Timeout** box. This change to a TCP DNS request validates that the original request came from a legitimate client, and the TCP DNS request is then forwarded to its destination. Active TCP mode can protect any type of DNS server. When you enable Active TCP mode, the TCP SYN Authentication countermeasure is disabled for port 53.
DNS Authentication Timeout box	If you selected **Passive** or **Active TCP** from the **Protection Mode** list, then type the *number of seconds* after which a DNS request is considered to have failed authentication.

About locked settings

When  (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring the IP Address Filter Lists Countermeasure

Introduction

The IP Address Filter Lists countermeasure contains user-chosen lists of IP addresses that are configured in Peakflow SP for use in TMS mitigations. These filter lists are the first configurable filters that can be used in a mitigation. The source IP address that are known in advance to be undesirable can be dropped immediately, and the source IP addresses that are known in advance to be desirable can be forwarded immediately. No further countermeasure evaluation is done on any traffic matching selected IP address filter lists. Because filter list processing is efficient, skillful use of filter lists can make better use of mitigation countermeasure processing capacity for those sources that are suspect or unknown.

When you configure a TMS mitigation, you can use the **IP Based Filter Lists** tab to configure the IP Address Filter Lists countermeasure to mitigate IPv4 or IPv6 traffic. You can also configure this countermeasure on the TMS Mitigation Status page.

You can configure one mitigation for either IPv4 or IPv6 filtering.

For information about adding and editing mitigations, see “Configuring and Deleting TMS Mitigations” on page 110.

For information about adjusting these and other mitigation settings on the TMS Mitigation Status page, see “About the TMS Mitigation Status Page” on page 100.

Configuring the IP Address Filter Lists countermeasure

To configure the IP Address Filter Lists countermeasure for TMS mitigations:

1. Do one of the following:
   - Navigate to the **IP Based Filter Lists** tab of the mitigation.

     See “Adding and editing a TMS mitigation” on page 110.
   - Navigate to the TMS Mitigation Status page. On the Countermeasures tab, click (expand) for the IP Address Filter Lists.

     See “Navigating to the TMS Mitigation Status page” on page 100.

     Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations.

2. Use the following table to configure the IP Address Filter Lists settings:

Setting	Procedure
IPv4 or **IPv6 Address Drop Filter Lists** box	To apply a drop rule to a list of IPv4 or IPv6 addresses: a. Under the **IP Address Drop Filter Lists** box, click **Select Filter List**. b. Select the IP address lists to include in the rule, and then click **OK**.
IPv4 or **IPv6 Address Pass Filter Lists** box	To apply a pass rule to a list of IPv4 or IPv6 addresses: a. Under the **IP Address Pass Filter Lists** box, click **Select Filter List**. b. Select the IP address lists to include in the rule, and then click **OK**.
Blacklist Sources check box	Select to blacklist any source host that sends traffic that is dropped because the source host's IP address is in the **IP Address Drop Filter Lists**. This check box is selected by default. Blacklisting of source hosts makes it more efficient for a Threat Management System appliance to drop traffic. For information about how blacklisting works, see "About dynamic blacklisting" on page 97.
View All Filter List Rates button	Click to view all filter list rate information.

3. Click **Save**.

About locked settings

When  (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring the IP Location Filter Lists Countermeasure

Introduction

IP Location filter lists are assembled from one or more geographic country lists that are loaded in Peakflow SP software for use in TMS mitigations. Each IP Location country object is internally defined as a large list of IP addresses that is not visible or configurable in Peakflow SP. Each IP Location filter list is configured in Peakflow SP as a selection of any number of IP Location countries. Default IP Location filter lists for several continental regions are installed by the software.

When you configure a TMS mitigation, you can use the **IP Based Filter Lists** tab to configure the IP Location Filter Lists (IPv4) countermeasure. You can also configure this countermeasure on the Countermeasures tab on the TMS Mitigation Status page.

For information about adding and editing mitigations, see "Configuring and Deleting TMS Mitigations" on page 110.

For information about adjusting these and other mitigation settings on the TMS Mitigation Status page, see "About the TMS Mitigation Status Page" on page 100.

Configuring the IP Location Filter Lists countermeasure

To configure the IP Location Filter Lists countermeasure for TMS mitigations:

1. Do one of the following:
   - Navigate to the **IP Based Filter Lists** tab of the mitigation.

     See "Adding and editing a TMS mitigation" on page 110.
   - Navigate to the TMS Mitigation Status page. On the Countermeasures tab, click ⊞ (expand) for the IP Location Filter Lists countermeasures.

     See "Navigating to the TMS Mitigation Status page" on page 100.

     Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations.
2. Use the following table to configure the IP Location Filter Lists settings:

Setting	Procedure
IP Location Filter Lists box	(IPv4 only) To drop traffic based on a geographic IPv4 address list: a. Under the **IP Location Filter Lists** box, click **Select Filter List**. b. Select the IP Location filter lists to include in the rule, and then click **OK**. c. From the list below **Select Filter List**, select whether to drop matched or unmatched traffic.
View All button	Click to view all country filter rate information.

3. Click **Save**.

About locked settings

When 🔒 (lock) appears beside a setting, you cannot edit that setting because it has been

locked by your service provider.

Configuring the IP Location Policing Countermeasure

Introduction

IP Location Policing is a per-packet countermeasure that uses individual settings of IP location data to mitigate attacks. Each selected country can be configured with an action to drop, pass, or rate shape the matching traffic. The IP Location Policing countermeasure mitigates IPv4 attack traffic.

For Peakflow SP to create per-country traffic rate suggestions, the Generate IP Location Policing Rate Suggestions setting must be selected. These rate suggestions can be transferred into rate limit settings for a country. This setting must be configured by your service provider.

You can configure the IP Location Policing countermeasure when you create or edit a mitigation and when you edit a mitigation on the TMS Mitigation Status page. For information about configuring mitigations, see "Configuring and Deleting TMS Mitigations" on page 110.For information about the TMS Mitigation Status page, see "About the TMS Mitigation Status Page" on page 100.

About the IP Location Policing countermeasure

This countermeasure allows you to mitigate traffic from either specified or unspecified (Other) countries by doing the following:

- Allowing all traffic to enter your network from either specified or unspecified (Other) countries

  All "allowed" traffic is not necessarily passed. Some allowed traffic might ultimately be dropped as the result of other enabled countermeasures.
- Blocking all traffic from entering your network from either specified or unspecified (Other) countries
- Limiting (rate shaping) the rate of traffic that enters your network from either specified or unspecified (Other) countries

Configuring the IP Location Policing countermeasure when adding or editing a mitigation

To configure the IP Location Policing countermeasure when adding or editing a mitigation:

1. Navigate to the **Countermeasures** tab of the mitigation.

   See "Adding and editing a TMS mitigation" on page 110.
2. Scroll to the IP Location Policing section.
3. Configure the settings for the IP Location Policing countermeasure.

   See "IP Location Policing countermeasure settings" on the next page.
4. Click **Save**.

Configuring the IP Location Policing countermeasure on the TMS Mitigation Status page

To configure the IP Location Policing countermeasure on the TMS Mitigations Status page:

1. Navigate to the TMS Mitigation Status page.

   See "Navigating to the TMS Mitigation Status page" on page 100.
2. On the Countermeasures tab, click (expand) for the IP Location Policing

countermeasure.

Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations.

3. Configure the settings for the IP Location Policing countermeasure.

   See "IP Location Policing countermeasure settings" below.

4. Click **Save**.

IP Location Policing countermeasure settings

Use the following table to configure the IP Location Policing countermeasure settings:

Setting	Procedure
Enable IP Location Policing check box	Select to enable this countermeasure.
Add Country button	Click to specify a country whose traffic should be policed, by completing the following steps: a. Click **Add Country**. b. In the Add IP Location Policing Country window, click **Select Country**. c. In the Select a Country window, select a country, and then click **OK**. d. From the **Action** list, select whether to drop all, allow all, or rate shape traffic from the selected country. e. If you selected **rate shape**, then type the maximum *amount* of traffic to allow and select their corresponding unit of measure from the **bps** and **pps** lists. f. Click **OK**.
Load All Countries (TMS Mitigation Status page only)	Click to add all of the countries' traffic for which Peakflow SP has data.
Load Rates (TMS Mitigation Status page only)	Click to load the generated rates for all countries whose configured actions are "rate shape." For Peakflow SP to load per-country traffic rate suggestions, the **Generate IP Location Policing Rate Suggestions** setting must be selected on the **Mitigation** tab for the protected managed object. This setting must be configured by your service provider.
Load All Countries and Rates on Mitigation Start check box	Select if, when a mitigation starts, you want to load traffic rates automatically for countries that you did not specify but for which Peakflow SP has managed object-specific data. To load rates, the **Generate IP Location Policing Rate Suggestions** setting must be selected on the **Mitigation** tab for the protected managed object. This setting must be configured by your service provider.

Icon	Description
	Hover over this icon to display the download options. The download options that appear vary depending on the page that you are on and can include the following:

About locked settings

When  (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring the Payload Regular Expression Countermeasure

Introduction

The Payload Regular Expression countermeasure allows you to configure a mitigation to drop malicious TCP or UDP traffic or to blacklist the hosts of malicious TCP or UDP traffic. The payload of a packet consists of the data after the TCP and UDP headers. If the packet matches one of the specified destination ports, the payload regular expression is applied to the packet payload. If the packet payload either matches or does not match a payload regular expression (depending on your settings), the packet is dropped or the host is blacklisted. The payload regular expression can be applied to the packet header in addition to the packet payload.

The regular expressions are applied to individual packets only. Any match that requires spanning multiple packets is not detected.

When you configure a TMS mitigation, you can use the **Payload** tab to configure the Payload Regular Expression countermeasure for IPv4 or IPv6 mitigations. You can also configure this countermeasure on the Countermeasures tab on the TMS Mitigation Status page.

For information about adding and editing mitigations, see “Configuring and Deleting TMS Mitigations” on page 110.

For information about adjusting these and other mitigation settings on the TMS Mitigation Status page, see “About the TMS Mitigation Status Page” on page 100.

Configuring the Payload Regular Expression countermeasure

To configure the Payload Regular Expression countermeasure for TMS mitigations:

1. Do one of the following:
   - Navigate to the **Payload** tab of the mitigation.

     See “Adding and editing a TMS mitigation” on page 110.
   - Navigate to the TMS Mitigation Status page. On the Countermeasures tab, click (expand) for the Payload Regular Expression countermeasure.

     See “Navigating to the TMS Mitigation Status page” on page 100.

     Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations. See “About locked settings” on page 140.
2. Configure the Payload Regular Expression countermeasure settings.

   See “ Payload Regular Expression countermeasure settings” on the facing page.
3. Click **Save**.

Payload Regular Expression countermeasure settings

Use the following table to configure the Payload Regular Expression countermeasure settings:

Setting	Procedure
Payload Regular Expression TCP Ports box	Type the destination *port numbers* for TCP traffic that needs to be checked for the payload regular expression. Use spaces or commas to separate multiple port numbers. The payload regular expression is applied to the TCP packets with these destination ports only. If you do not specify ports, then the payload regular expression is not applied to the traffic.
Payload Regular Expression UDP Ports box	Type the destination *port numbers* for UDP traffic that needs to be checked for the payload regular expression. Use spaces or commas to separate multiple port numbers. The payload regular expression is applied to the UDP packets with these destination ports only. If you do not specify ports, then the payload regular expression is not applied to the traffic.
Payload Regular Expression box	Type the *regular expression* (in PCRE format and single-line mode) to apply to the payload traffic that matches the appropriate ports. **Note:** Payload regular expressions are case-sensitive by default. To perform case-insensitive matching, preface the expression with "(?i)".
Action to Apply to Offending Hosts options	Click **Blacklist Hosts** or **Drop Traffic**. The **Blacklist Hosts** option is selected by default. If you click **Blacklist Hosts**, then all of the traffic from the offending source hosts is dropped. If you click **Drop Traffic**, then only the offending traffic from these hosts is dropped. Blacklisting of source hosts is a more efficient way for a Threat Management System appliance to drop traffic. For information about how blacklisting works, see "About dynamic blacklisting" on page 97.
Apply Action to options	Click **Matched Traffic** or **Unmatched traffic**. If you click **Matched Traffic**, then the traffic that matches the payload regular expression is either dropped or the host is blacklisted. If you click **Unmatched traffic**, then the traffic that does not match the payload regular expression is either dropped or the host is blacklisted.
Apply Regular Expression to Packet Headers check box	Select this check box to apply the regular expression to the packet header in addition to the packet payload. This option allows you to block attacks based on specific patterns in the packet header.

Setting	Procedure
Download Blocked Hosts button (TMS Mitigation Status page only)	Click to download a .txt file that contains a list of the hosts that this countermeasure currently blocks.You can use this information to refine your configuration of the countermeasure.
Download Top Blocked Hosts button (TMS Mitigation Status page only)	Click to download a txt. file that contains a list of the hosts that this countermeasure has blocked the most. You can use this information to refine your configuration of the countermeasure.
Test Regular Expression button (TMS Mitigation Status page only)	Click to test the effectiveness of a regular expression in mitigating the attack associated with this alert.

About locked settings

When  (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring the Per Connection Flood Protection Countermeasure

Introduction

The Per Connection Flood Protection countermeasure monitors IPv4 traffic on a per-connection basis (5-tuple) rather than on a per-source basis. When the IPv4 traffic of any connection exceeds the maximum configured rates for bps or pps, then the countermeasure can block all of the traffic of that connection or limit the rate of the traffic of that connection.

You can use the Per Connection Flood Protection countermeasure when blacklisting the source of the offending traffic is not a good option. For example, if the attacker is behind a NAT, you can use this countermeasure to block or rate limit the traffic of an attacker's connection without blacklisting legitimate users who are also behind the same NAT.

Important: Per Connection Flood Protection should be used only with applications that send traffic at a consistent rate. With applications such as HTTP that send bursts of traffic, this countermeasure may produce inconsistent results.

You can configure the Per Connection Flood Protection countermeasure when you add or edit a mitigation and when you edit a mitigation on the TMS Mitigation Status page. For information about configuring mitigations, see "Configuring and Deleting TMS Mitigations" on page 110.For information about the TMS Mitigation Status page, see "About the TMS Mitigation Status Page" on page 100.

Configuring the Per Connection Flood Protection countermeasure when adding or editing a mitigation

To configure the Per Connection Flood Protection countermeasure when adding or editing a mitigation:

1. Navigate to the **Countermeasures** tab of the mitigation.
   See "Adding and editing a TMS mitigation" on page 110.
2. Scroll to the Per Connection Flood Protection section.
3. Configure the settings for the Per Connection Flood Protection countermeasure.
   See "Per Connection Flood Protection countermeasure settings" on the next page.
4. Click **Save**.

Configuring the Per Connection Flood Protection countermeasure on the TMS Mitigation Status page

To configure the Per Connection Flood Protection countermeasure on the TMS Mitigations Status page:

1. Navigate to the TMS Mitigation Status page.
   See "Navigating to the TMS Mitigation Status page" on page 100.
2. On the Countermeasures tab, click ⊞ (expand) for the Per Connection Flood Protection countermeasure.
   Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations. See "About locked settings" on page 143.
3. Configure the settings for the Per Connection Flood Protection countermeasure.
   See "Per Connection Flood Protection countermeasure settings" on the next page.

4. Click **Save**.
5. To evaluate the impact of the settings that you selected, view the following statistics below the settings of the countermeasure:
   - Connection Rate

     The number of connections per second to the configured TCP and UDP ports.
   - Enforcement Rate

     The number of connections per second that the countermeasure is blocking or rate limiting.
   - Packets Ignore Rate

     The number of packets per second that the countermeasure ignores. The countermeasure ignores the packets of traffic that are going to TCP or UDP ports that are not configured in the countermeasure. If you type **`all`** in **TCP Ports** and **UDP Ports**, then no packets are ignored.

Per Connection Flood Protection countermeasure settings

Use the following table to configure the Per Connection Flood Protection countermeasure settings:

Setting	Procedure
Enable Per Connection Flood Protection check box	Select to enable this countermeasure.
TCP Ports box	Type the *`destination port numbers`* for the TCP traffic that you want this countermeasure to monitor. Use spaces or commas to separate multiple port numbers. To monitor a range of ports, separate the first and last number in the range with a hyphen (for example: 21-26). To monitor all TCP traffic, type **`all`**. This countermeasure monitors only the TCP packets that have these destination ports. If you do not specify TCP ports, then this countermeasure is not applied to TCP traffic. **Note:** You must configure at least 1 TCP or 1 UDP port to be able to use this countermeasure.
UDP Ports box	Type the *`destination port numbers`* for the UDP traffic that you want this countermeasure to monitor. Use spaces or commas to separate multiple port numbers. To monitor a range of ports, separate the first and last number in the range with a hyphen (for example: 67-69). To monitor all UDP traffic, type **`all`**. This countermeasure monitors only the UDP packets that have these destination ports. If you do not specify UDP ports, then this countermeasure is not applied to UDP traffic. **Note:** You must configure at least 1 TCP or 1 UDP port to be able to use this countermeasure.

Setting	Procedure
Enforcement options	Click **Block** or **Rate Limit**. The default setting is **Block**. If you click **Block**, then Peakflow drops all of the traffic of a connection when that traffic exceeds the maximum configured pps or bps. If you click **Rate Limit**, then Peakflow drops packets from the traffic of a connection to keep the traffic within the maximum configured pps or bps. Use the **Rate Limit** option if you do not want to drop all of the traffic of a connection. You configure the maximum pps or bps in the **Maximum Packets/sec per Connection** and **Maximum Bits/sec per Connection** boxes.
Maximum Packets/sec per Connection box	Type the *maximum number of packets* that each connection is allowed per second before its traffic is blocked or rate limited, and then select the unit of measure. If you type **0** or leave the box blank, then the packets per second rate is unlimited.
Maximum Bits/sec per Connection box	Type the *maximum number of bits* that each connection is allowed per second before its traffic is blocked or rate limited, and then select the unit of measure. If you type **0** or leave the box blank, then the bits per second rate is unlimited.

About locked settings

When a (lock) icon appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring the Protocol Baselines Countermeasure

Introduction

The Protocol Baselines countermeasure helps protect your network from uncharacteristic surges in traffic volume. For this countermeasure, Peakflow SP collects historical traffic data from the configured managed object. If traffic rates exceed a calculated baseline threshold, then the TMS appliance dynamically blacklists the source host of the traffic. The Protocol Baselines countermeasure mitigates IPv4 attack traffic.

For Peakflow SP to collect the historical data for a configured managed object, the Enable Enforced Baseline Protection setting must be selected at least two days in advance of a mitigation. This setting must be configured by your service provider.

When a mitigation starts that has this countermeasure enabled, Peakflow TMS downloads from Peakflow SP the observed traffic data for the time period 48 to 24 hours before the start of the mitigation. The downloaded data is assumed normal and becomes the baseline data that is used by the countermeasure for comparison with real-time traffic statistics. This baseline data is not updated for the duration of the mitigation.

You can configure the Protocol Baselines countermeasure when you create or edit a mitigation and when you edit a mitigation on the TMS Mitigation Status page. For information about configuring mitigations, see “Configuring and Deleting TMS Mitigations” on page 110.For information about the TMS Mitigation Status page, see “About the TMS Mitigation Status Page” on page 100.

Configuring the Protocol Baselines countermeasure when adding or editing a mitigation

To configure the Protocol Baselines countermeasure when adding or editing a mitigation:

1. Navigate to the **Countermeasures** tab of the mitigation.
   See “Adding and editing a TMS mitigation” on page 110.
2. Scroll to the Baseline Enforcement section.
3. To enable this countermeasure, select the **Enable Protocol Baseline Enforcement** check box.
4. If you are configuring a mitigation, verify that you selected a managed object on the **Protect** tab and that it appears in the **Baselines from Managed Object** box.
   See “Configuring Protect Settings for TMS Mitigations” on page 117.
5. Click **Save**.

Configuring the Protocol Baselines countermeasure on the TMS Mitigation Status page

To configure the Protocol Baselines countermeasure on the TMS Mitigations Status page:

1. Navigate to the TMS Mitigation Status page.
   See “Navigating to the TMS Mitigation Status page” on page 100.
2. On the Countermeasures tab, click ⊞ (expand) for the Protocol Baselines countermeasure.
   Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations.
3. To enable this countermeasure, select the **Enable Protocol Baseline Enforcement**

check box.

4. To download information about blocked protocols, use the following buttons:

Button	Procedure
Download Protocol Block Data	Click to download a .txt file containing a list of the protocols blocked by this countermeasure, including the number of bytes and packets blocked for each protocol. You can use this information to refine your mitigation.
Download Top Protocol Block Data	Click to download a .txt file containing a list of the protocols most frequently blocked by this countermeasure with the number of bytes and packets blocked for each protocol. You can use this information to refine your mitigation.

5. Click **Save**.

About locked settings

When  (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring the Shaping Countermeasure

Introduction

The TMS appliance can shape any IPv4 traffic that remains after the preceding countermeasures and filters are applied. Shaping allows you to control the level of bps / pps traffic that reaches the customer to ensure that your links do not become overwhelmed. A mitigation can have up to ten traffic shaping filters. You can also configure the traffic shaping settings on the Countermeasures tab on the TMS Mitigation Status page.

When you configure a TMS mitigation, you can use the **Shaping** tab to configure traffic shaping settings. You can also configure this countermeasure on the Countermeasures tab on the TMS Mitigation Status page.

For information about adding and editing mitigations, see "Configuring and Deleting TMS Mitigations" on page 110.

For information about adjusting these and other mitigation settings on the TMS Mitigation Status page, see "About the TMS Mitigation Status Page" on page 100.

Example: Shaping countermeasure with two filters

Consider a Shaping countermeasure that has two traffic shaping filters that have different purposes. Each traffic shaping filter then has its own FCAP filter expression and maximum level settings. For example, the filters might have settings similar to those in the following table:

Purpose	FCAP Filter Expression	Maximum Levels
limit the rate of traffic to an application server	`dst net 10.0.0.1 and dst port 26 and proto UDP`	1 Mbps 500 pps
protect a Web server from resource exhaustion	`dst net 10.0.0.2 and proto TCP and (proto tcp and tflags A/A)`	2 Mbps 1 Kpps

Configuring the Shaping countermeasure

To configure the Shaping countermeasure for TMS mitigations:

1. Do one of the following:
   - Navigate to the **Shaping** tab of the mitigation.

     See "Adding and editing a TMS mitigation" on page 110.
   - Navigate to the TMS Mitigation Status page. On the Countermeasures tab, click (expand) for the Shaping countermeasure.

     See "Navigating to the TMS Mitigation Status page" on page 100.

     Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations.
2. Select the **Enable Shaping** check box, to enable traffic shaping.

3. Use the following table to configure the traffic shaping settings for the first filter:

Setting	Procedure
Flow Capture Filter Expression box	Choose one of the following steps: ■ Type a *`fingerprint expression`* that corresponds to the data that you want to match. ■ Click **Open FCAP Wizard** to use the FCAP Wizard to add a fingerprint expression. See "Using the FCAP Wizard" on page 18.
Maximum Levels boxes	Type the *`maximum amount`* of bps and pps traffic to allow, and then select the traffic rate unit of measure for each.

4. To add settings for another traffic shaping filter, click **Add Shaping Queue** and configure the settings according to Step 3.
5. Repeat Step 4 for each additional traffic shaping filter that you want to add.

   A mitigation can have up to ten traffic shaping filters.
6. To remove a traffic shaping filter, click **Remove** within the section that contains that filter's settings.
7. Click **Save**.

   On the TMS Mitigation Status page, each filter is assigned a color that is displayed in a square above its filter expression. The traffic graph at the bottom of the countermeasure displays the traffic that each filter drops, using the colors assigned to each filter.

About locked settings

When  (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring the Source /24 Baselines Countermeasure

Introduction

The Source /24 Baselines countermeasure has been deprecated. This countermeasure helps protect your network from uncharacteristic surges in traffic volume. For this countermeasure, Peakflow SP collects historical traffic data from the configured managed object. If traffic rates exceed a calculated baseline threshold, then the TMS appliance dynamically blacklists the traffic. The Source /24 Baselines countermeasure mitigates IPv4 attack traffic.

For Peakflow SP to collect the historical data for a configured managed object, the Enable Enforced Baseline Protection setting must be selected at least two days in advance of a mitigation. This setting must be configured by your service provider.

When a mitigation starts that has this countermeasure enabled, Peakflow TMS downloads from Peakflow SP the observed traffic data for the time period 48 to 24 hours before the start of the mitigation. The downloaded data is assumed normal and becomes the baseline data that is used by the countermeasure for comparison with real-time traffic statistics. This baseline data is not updated for the duration of the mitigation.

You can configure the Source /24 Baselines countermeasure when you create or edit a mitigation and when you edit a mitigation on the TMS Mitigation Status page. For information about configuring mitigations, see "Configuring and Deleting TMS Mitigations" on page 110.For information about the TMS Mitigation Status page, see "About the TMS Mitigation Status Page" on page 100.

Configuring the Source /24 Baselines countermeasure when adding or editing a mitigation

To configure the Source /24 Baselines countermeasure when adding or editing a mitigation:

1. Navigate to the **Countermeasures** tab of the mitigation.
   See "Adding and editing a TMS mitigation" on page 110.
2. Scroll to the Baseline Enforcement section.
3. If you are configuring a mitigation, verify that you selected a managed object on the **Protect** tab and that it appears in the **Baselines from Managed Object** box.
   See "Configuring Protect Settings for TMS Mitigations" on page 117.
4. To enable this countermeasure, select the **Enable Source /24 Baselines Enforcement** check box.
5. Click **Save**.

Configuring the Source /24 Baselines countermeasure on the TMS Mitigation Status page

To configure the Source /24 Baselines countermeasure on the TMS Mitigations Status page:

1. Navigate to the TMS Mitigation Status page.
   See "Navigating to the TMS Mitigation Status page" on page 100.
2. On the Countermeasures tab, click ⊞ (expand) for the Protocol Baselines countermeasure.
   Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations.

3. To enable this countermeasure, select the **Enable Source /24 Baseline Enforcement** check box.
4. To download information about blocked source /24 data, use the following buttons:

Button	Procedure
Download Source /24 Block Data	Click to download a .txt file containing a list of the source /24 data blocked by this countermeasure. You can use this information to refine your mitigation.
Download Top Source /24 Block Data	Click to download a .txt file containing a list of the source /24 data most frequently blocked by this countermeasure. You can use this information to refine your mitigation.

5. Click **Save**.

About locked settings

When  (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring the TCP SYN Authentication Countermeasure

Introduction

TCP SYN Authentication is a per-packet countermeasure that intercepts and authenticates all inbound IPv4 and IPv6 TCP connections to the protected hosts. It can protect against TCP SYN flood attacks and any TCP flag attack, such as ACK floods or illegal TCP flag combinations. In these attacks, the TCP protocol is misused to consume a target's resources. The TCP SYN Authentication countermeasure mitigates IPv4 and IPv6 attack traffic.

You can configure the TCP SYN Authentication countermeasure when you create or edit a mitigation and when you edit a mitigation on the TMS Mitigation Status page. For information about configuring mitigations, see "Configuring and Deleting TMS Mitigations" on page 110.For information about the TMS Mitigation Status page, see "About the TMS Mitigation Status Page" on page 100.

About the TCP SYN Authentication countermeasure

In the TCP SYN Authentication countermeasure, the TMS appliance acts as a proxy for the protected hosts to verify that the source host completes a three-way SYN/ACK handshake. If the source host is authenticated, then the TMS appliance approves that host and allows it to connect to the protected hosts. The source host remains approved until it does not send a TCP packet within the configured timeout period.

If the source host is not authenticated, then it is assumed to be malicious, and the connection is not allowed. A host that fails TCP SYN authentication is not blacklisted; any subsequent TCP connection attempt can be used to authenticate that host.

If the first received packet of a TCP connection is not a SYN packet, then the TCP SYN Authentication countermeasure assumes that it has intercepted a connection already in progress. Packets from the source host will continue to be dropped until the TMS detects a retransmission of the data in the dropped packet. The retransmission packet is forwarded to the destination, and the source host is approved to continue sending TCP packets directly to the protected hosts.

Note: TCP SYN Authentication ignores TCP port 53 traffic, and the traffic is forwarded untouched.

Configuring the TCP SYN Authentication countermeasure when adding or editing a mitigation

To configure the TCP SYN Authentication countermeasure when adding or editing a mitigation:

1. Navigate to the **Countermeasures** tab of the mitigation.
   See "Adding and editing a TMS mitigation" on page 110.
2. Scroll to the TCP SYN Authentication section.
3. Configure the settings for the TCP SYN Authentication countermeasure.
   See "TCP SYN Authentication countermeasure settings" on the facing page.
4. Click **Save**.

Configuring the TCP SYN Authentication countermeasure on the TMS Mitigation Status page

To configure the TCP SYN Authentication countermeasure on the TMS Mitigations Status page:

1. Navigate to the TMS Mitigation Status page.
   See "Navigating to the TMS Mitigation Status page" on page 100.
2. On the Countermeasures tab, click (expand) for the TCP SYN Authentication countermeasure.
   Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations.
3. Configure the settings for the TCP SYN Authentication countermeasure.
   See "TCP SYN Authentication countermeasure settings" below.
4. Click **Save**.

TCP SYN Authentication countermeasure settings

Use the following table to configure the TCP SYN Authentication settings:

Setting	Procedure
Enable TCP SYN Authentication check box	Select to enable this countermeasure. When TCP SYN authentication is enabled and a legitimate source host completes the TCP handshake, the source host has only a TCP connection with the TMS appliance. The TMS appliance then normally sends a TCP reset to the source host. This TCP reset usually results in an error from the application that is visible to the user and that can require the user to refresh their Web browser manually. To resolve this problem and to make the connection to the real server transparent to the user, it is recommended that you also select **Enable Out-of-sequence Authentication**.
Ignore Source Ports box	Type the *source ports* that the countermeasure should ignore.
Ignore Destination Ports box	Type the *destination ports* that the countermeasure should ignore.
TCP SYN Authentication Idle Timeout box	Type the *number of seconds* that an authenticated TCP SYN connection can be idle before it is dropped. After a connection times out, the next TCP connection attempt by that host must be authenticated.

Setting	Procedure
Enable Out-of-sequence Authentication check box	Select to enable the TMS appliance to use out-of-sequence authentication instead of TCP SYN authentication. This authentication method allows the TMS appliance to transparently authenticate all applications without displaying error messages to the user or requiring them to refresh their Web browsers manually. It is recommended that you use this authentication method in most instances. If for some reason out-of-sequence authentication fails, then the TMS appliance reverts to TCP SYN authentication. If out-of-sequence authentication causes problems with clients, then it is recommended that you select **Enable Application Reset** instead.
Enable Application Reset check box	Select to enable the TMS appliance to use a simple HTTP redirect to the source host to make it open a new connection to the real server. The user should then not see an error message and should not have to refresh the Web browser manually. Application reset only supports HTTP. It is recommended that you select this setting if you cannot use out-of-sequence authentication. For example, if you have clients where out-of-sequence authentication causes problems, then you should select this setting. You can also select this setting when you select out-of-sequence authentication. If out-of-sequence authentication fails for a client, then application reset will still make the connection to the real server transparent to the user.
Enable HTTP Authentication check box	(IPv4 mitigations only) Select to apply additional authentication steps to specific HTTP ports. While TCP SYN authentication can identify spoofed SYN floods, HTTP authentication can identify attacks by botnets or malicious users that are not spoofed. HTTP authentication makes sure that the source host is a valid HTTP client. It does this by making sure that the source host correctly responds to an HTTP redirect that the TMS sends. If the source host correctly responds to the redirect, then it is allowed to connect to the protected host. It is recommended that you enable this setting only when an attack has multiple components that include both a spoofed SYN flood and an HTTP request flood. **Note:** If you select this check box and the **Require JavaScript for HTTP Authentication** check box, then this check box has no effect and the ability to run JavaScript is used for HTTP authentication,

Setting	Procedure
HTTP Authentication Ports box	(IPv4 mitigations only) If you enabled HTTP Authentication, type the *HTTP ports* on which to look for authenticated HTTP traffic. **Note:** The TMS appliance does not perform TCP SYN authentication on any of the ports where HTTP Authentication is performed.
Require JavaScript for HTTP Authentication check box	(IPv4 mitigations only) Select this check box to require the browser of the source host to be able to run JavaScript for HTTP authentication. If the browser of the source host is able to run JavaScript, then the source host is allowed to connect to the protected host. If the browser of the source host cannot run JavaScript, then the TMS appliance does not allow the source host to connect to the protected hosts. **Note:** This check box is disabled if the **Enable HTTP Authentication** check box is not selected. **Note:** If you select this check box and have legitimate clients that cannot run JavaScript, then they will not be allowed to connect to the protected host.

About locked settings

When (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring the Zombie Detection Countermeasure

Introduction

The Zombie Detection countermeasure uses configured threshold values to identify and block hosts ("zombies") that send excessive amounts of IPv4 or IPv6 traffic to protected hosts or networks. This packet-based countermeasure can protect against common attacks including flood, TCP SYN, and protocol attacks. The Zombie Detection countermeasure mitigates IPv4 and IPv6 attack traffic.

You can configure the Zombie Detection countermeasure when you create or edit a mitigation and when you edit a mitigation on the TMS Mitigation Status page. For information about configuring mitigations, see "Configuring and Deleting TMS Mitigations" on page 110.For information about the TMS Mitigation Status page, see "About the TMS Mitigation Status Page" on page 100.

About the Zombie Detection countermeasure

You can configure the following Zombie Types:

Zombie Type	Description
All Hosts	Peakflow SP checks the configured bit and packet rates from all hosts. If the traffic from a host exceeds any of the configured thresholds, then the host is blacklisted for one minute for the first offense and for five minutes for a repeat offense. Typically, you should set the All Hosts thresholds to rates that are higher than any legitimate host is expected to send on a regular basis. These rates can vary, depending on the services that a host offers. For example, if the protected hosts are content servers and the source hosts are clients that send only requests and acknowledgments, then the expected traffic rates are low.
Flexible 1 Flexible 2 Flexible 3 Flexible 4 Flexible 5	You can specify up to five Flexible Zombie configurations to handle specific types of attacks. The Flexible Zombie configurations allow you to specify bps and pps thresholds and a filter (SYN flag, packet size, etc.). Only packets matching the filter are counted to determine if a host should be blocked. If the rate is exceeded per source address, then the host is blacklisted for one minute for the first offense and for five minutes for a repeat offense. See "Example: Flexible Zombie mitigation" on page 156.

Configuring the Zombie Detection countermeasure when adding or editing a mitigation

To configure the Zombie Detection countermeasure when adding or editing a mitigation:

1. Navigate to the **Countermeasures** tab of the mitigation.
   See "Adding and editing a TMS mitigation" on page 110.
2. Scroll to the Zombie Removal section.
3. To enable this countermeasure, select the **Enable Zombie Detection** check box.

4. Configure the following settings for the different types of zombies:

Setting	Procedure
Zombie Filter box (Flexible Zombie configurations only)	Do one of the following: ■ Type a fingerprint expression that corresponds to the data that you want to match. ■ (IPv4 mitigations only) Click **Open FCAP Wizard** to use the FCAP Wizard to add a fingerprint expression. See "Using the FCAP Wizard" on page 18. **Note:** A filter is required for a Flexible Zombie configuration.
Zombie Thresholds boxes	In each box, type a *rate*, and then select the corresponding unit of measure (bps or pps). These values determine the level of traffic that a host can send before it is considered a zombie. **Note:** A threshold setting is required for All Hosts and for each Flexible Zombie. You can configure a bps setting, a pps setting, or both.

5. Click **Save**.

Configuring the Zombie Detection countermeasure on the TMS Mitigation Status page

To configure the Zombie Detection countermeasure on the TMS Mitigations Status page:

1. Navigate to the TMS Mitigation Status page.
   See "Navigating to the TMS Mitigation Status page" on page 100.
2. On the Countermeasures tab, click ⊞ (expand) for the Zombie Detection countermeasure.
   Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations.
3. To enable this countermeasure, select the **Enable Zombie Detection** check box.
4. To see configuration details for a specific Zombie Type, in the **Zombie Type** box, select the Zombie Type.
   When a Zombie Type is selected, the following information is displayed:
   - threshold configurations
   - filter information (Flexible Zombie configurations only)
   - graphs showing the number of hosts that are blocked based on the traffic thresholds that you configure
     These graphs include sliders (💧) that you can use to adjust the configured traffic thresholds.
   - graph showing the dropped traffic

   Note: If you click the **Save** button in an individual Zombie Type section, Peakflow SP saves only your configuration changes for that Zombie Type. If you click the **Save** button in the overview section, Peakflow SP will save all configuration changes.

5. Configure the following settings for the different types of zombies:

Setting	Procedure
Zombie Filter box (Flexible Zombie configurations only)	Do one of the following: ■ Type a fingerprint expression that corresponds to the data that you want to match. ■ (IPv4 mitigations only) Click **Open FCAP Wizard** to use the FCAP Wizard to add a fingerprint expression. See "Using the FCAP Wizard" on page 18. **Note:** A filter is required for a Flexible Zombie configuration.
Zombie Thresholds boxes	In each box, type a *rate*, and then select the corresponding unit of measure (bps or pps). These values determine the level of traffic that a host can send before it is considered a zombie. **Note:** A threshold setting is required for All Hosts and for each Flexible Zombie. You can configure a bps setting, a pps setting, or both.

6. To download information about blocked hosts, use the following buttons:

Button	Procedure
Download Blocked Hosts	Click to download a .txt file containing a list of the hosts blocked by this countermeasure. You can use this information to refine other countermeasure settings in the mitigation.
Download Top Blocked Hosts	Click to download a .txt file containing a list of the most frequently blacklisted host. You can use this information to refine other countermeasure settings in the mitigation.

7. Click **Save**.

Example: Flexible Zombie mitigation

A botnet attacks a server by connecting over TCP, completing the 3-way TCP handshake, but then disconnecting a short time after the handshake completes without sending any data. You cannot use the TCP SYN Authentication countermeasure to block this traffic because the hosts are completing the 3-way handshake.

However, you can use Flexible Zombie configurations to mitigate this attack. You specify an FCAP expression of "proto tcp and tflags S/S" with a pps rate that is lower than what the attacking sources are sending but that is high enough to allow legitimate users to connect. The Flexible Zombie mitigation then blacklists the attacking hosts that send SYN packets at too great a rate even though they completed the 3-way handshake.

About locked settings

When (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Chapter 8:

Configuring Event-Driven Countermeasures

Introduction

This section describes how to configure the event-driven countermeasures for TMS mitigations.

For information about the types of countermeasures, see “Types of countermeasures” on page 97.

User access

Only managed services administrators can configure these settings.

In this section

This section contains the following topics:

Configuring the AIF and HTTP/URL Regular Expression Countermeasure

Introduction

The AIF and HTTP/URL Regular Expression countermeasure allows you to configure a mitigation to use the regular expressions downloaded by the AIF feed to drop traffic associated with malware families. It also allows you to use HTTP header regular expressions and/or URL filter lists to match traffic in a mitigation.

When you configure a TMS mitigation, you can use the **Payload** tab to configure the AIF and HTTP/URL Regular Expression countermeasure to mitigate IPv4 traffic. You can also configure this countermeasure on the Countermeasures tab on the TMS Mitigation Status page.

For information about configuring mitigations, see "Configuring and Deleting TMS Mitigations" on page 110.

For information about adjusting these and other mitigation settings on the TMS Mitigation Status page, see "About the TMS Mitigation Status Page" on page 100.

About the AIF and HTTP/URL Regular Expression Countermeasure

If the AIF feed regular expressions are used with the HTTP header regular expressions or URL filter lists, then the HTTP header regular expressions or URL filter lists will block matching traffic. If the HTTP header regular expressions and/or URL filter lists are used without the AIF feed regular expressions, these regular expression and/or filter lists block traffic that matches or does not match these regular expressions and/or filter lists.

The AIF and HTTP/URL Regular Expression countermeasure blacklists the source host of dropped traffic by default. This countermeasure scans across packets.

Configuring the AIF and HTTP/URL Regular Expression countermeasure

To configure the AIF and HTTP/URL Regular Expression countermeasure for TMS mitigations:

1. Do one of the following:
   - Navigate to the **Payload** tab of the mitigation.

     See "Adding and editing a TMS mitigation" on page 110.
   - Navigate to the TMS Mitigation Status page. On the Countermeasures tab, click ⊞ (expand) for the AIF and HTTP/URL Regular Expression countermeasure.

     See "Navigating to the TMS Mitigation Status page" on page 100.

     Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations.
2. Configure the AIF and HTTP/URL Regular Expression countermeasure settings.

   See "AIF and HTTP/URL Regular Expression countermeasure settings " on the facing page.
3. Click **Save**.

AIF and HTTP/URL Regular Expression countermeasure settings

Use the following table to configure the AIF and HTTP/URL Regular Expression countermeasure settings:

Icon	Description
	Hover over this icon to display the download options. The download options that appear vary depending on the page that you are on and can include the following: ■ PDF - Click to download the page in PDF format. ■ XML- Click to download the page in XML format. ■ CSV - Click to download a page in CSV format. The download can be a CSV text file or CSV zip archive file. When the download is a zip archive file, "zip archive" is appended to CSV. ■ Excel-XML - Click to download a page in Excel-XML format. The icon appears only if there are multiple options for downloading a page.
	Click to download a page in PDF format. This icon appears only when the PDF format is the only download option for a page.
	Click to download and email a page as a PDF.

Column	Description
	Select this icon to delete a user account.
Username	A user name as a link to the Edit Existing Account page.
Real Name	A user's full name.
Account Group	The account group to which a user belongs.
Capability Level	A user's capability level, which is either an administrator or a user.
Email	A user's email address.
Device	The SP appliance with which a user is associated. The SP appliance is either a specific appliance name or *global*, which associates a user with all appliances. For more information about associating a user with appliances, see "About user-appliance association" on page 24.
UI Menu	The UI menu that is assigned to a user. The UI menu determines what menu choices are available to a user.

Setting	Procedure
URL Filter Lists box	To add a URL filter list to the mitigation: 1. Click **Select Filter List**. 2. Select the filter lists to add to the mitigation, and then click **OK**. 3. From **Logical Connective**, select the **AND Filter Lists with Expressions** or the **OR Filter Lists with Expressions** operator. If the **AND Filter Lists with Expressions** operator is selected, the payload of a packet must match the regular expressions and the filter lists. If the **OR Filter Lists with Expressions** operator is selected, the payload of a packet only has to match the regular expressions or the filter lists. If you selected **Enable AIF Malware Family Blocking**, the **OR Filter Lists with Expressions** operator is selected and disabled. If the **AND Filter Lists with Expressions** operator is selected, then the option to enable the blocking of AIF malware families is disabled.
Graph Dataset list	Select one of the following options: ■ **Total** to display a graph of all the traffic that the countermeasure drops with its current configuration. ■ **AIF Low**, **AIF Medium**, or **AIF High** to display a graph of the traffic that each of the AIF enforcement levels drops. If you select an enforcement level that is higher than the configured AIF enforcement level, the graph displays how much traffic is currently matching the regular expressions that were downloaded for that level. If you select an enforcement level that is equal to or lower than the configured enforcement level, the graph displays how much traffic is being dropped for hosts whose traffic matches the regular expressions that were downloaded for that level.
Download Top URLs button (TMS Mitigation Status page only)	Click to download a .txt file containing a list of the most frequently accessed URLs. You can use this information to help you refine your regular expression.
Download Top User Agents button (TMS Mitigation Status page only)	Click to download a .txt file containing a list of the top user agents. You can use this information to help you refine your regular expression.
Download Blocked Hosts button (TMS Mitigation Status page only)	Click to download a .txt file containing a list of the hosts blocked by this countermeasure. You can use this information to help you refine your regular expression.

Setting	Procedure
Download Top Blocked Hosts button (TMS Mitigation Status page only)	Click to download a .txt file containing a list of the hosts that are most frequently blocked by this countermeasure. You can use this information to help you refine your regular expression.
Test Regular Expression button (TMS Mitigation Status page only)	Click to test the effectiveness of a regular expression in mitigating the attack associated with this alert.
View All Filter List Matched Rates button (TMS Mitigation Status page only)	Click to view all URL filter list rate information.

About the malware family list

The Enable AIF Malware Family Blocking setting includes the (information) icon. When you click , a window displays the malware families for which the AIF feed has downloaded regular expressions. For each malware family, the AIF feed downloads one or more regular expressions. The regular expressions for malware families in the Low list are “conservative.” The regular expressions for malware families in the Medium list are “moderate” unless the malware family inherits the regular expression used in the Low list. The regular expressions used in the High list are “aggressive” unless the malware family inherits the regular expression from the Low or Medium list.

In the Malware families window, you can view a list of all of the malware families or a list of malware families for each AIF enforcement level. You can also search for specific malware families.

If the AIF feed downloads more than one regular expression for a malware family, then a descriptor is appended to the malware family in the lists where it appears. The descriptors are “conservative,” “moderate,” and “aggressive.” If a malware family in a higher list inherits the regular expression that is used in a lower list, then the same descriptor is appended to the malware family in both lists. For example, if a malware family uses a conservative regular expression in the Low and Medium lists and an aggressive regular expression in the High list, then “conservative” is appended to the malware family in the Low and Medium lists and “aggressive” is appended to the malware family in the High list.

If the AIF feed downloads only one regular expression for a malware family, then a descriptor is not appended to the malware family in any of the lists where it appears. The regular expression used for that malware family in the lowest enforcement level is then used in any higher enforcement levels.

For information about the Enable AIF Malware Family Blocking setting, see “AIF and HTTP/URL Regular Expression countermeasure settings ” on page 159.

Note: If the AIF Malware Family Blocking setting is locked or if you have read-only access to TMS mitigations, you can view only the malware families for the selected AIF enforcement level.

The following table describes the malware family lists that you can view in the Malware families window:

Malware Family List	Description
Low	This list displays all of the malware families for which the AIF feed provides protection when the Low AIF enforcement level is selected. The regular expressions that the AIF feed downloads for these malware families have a low or conservative risk of dropping legitimate traffic.
Medium	This list displays all of the malware families for which the AIF feed provides protection when the Medium AIF enforcement level is selected. This list includes malware families that inherit the conservative regular expression used for that family in the Low list. It also includes malware families for which the AIF feed downloads regular expressions that have a moderate risk of dropping legitimate traffic.
High	This list displays all of the malware families for which the AIF feed provides protection when the High AIF enforcement level is selected. This list includes malware families that inherit the regular expression used for that family in the Low or Medium list. It also includes malware families for which the AIF feed downloads regular expressions that have a high risk of dropping legitimate traffic and are considered aggressive.
All	This list displays all of the malware families for which the AIF feed provides protection. For each malware family, it displays the lists in which it occurs.

AIF setting only drops matched traffic (IPv4 only)

When you select the **Enable AIF Malware Family Blocking** setting, the **Action** list is disabled and set to **Drop matched traffic** because the AIF feed is designed to drop bad HTTP traffic signatures. Because the AIF setting works in conjunction with any other settings in this countermeasure, those settings will also drop matched traffic. See “AIF and HTTP/URL Regular Expression countermeasure settings ” on page 159.

If you edit an existing mitigation that has **Drop unmatched traffic** selected from the **Action** list or **AND Filter Lists with Expressions** selected from the **Logical Connective** list, then the AIF setting is disabled.

About locked settings

When  (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring the DNS Malformed Countermeasure

Introduction

The DNS Malformed countermeasure filters DNS requests that do not conform to RFC standards. This countermeasure protects against attacks that send invalid or blank DNS messages to a server to exhaust resources or to exploit vulnerabilities. The DNS Malformed countermeasure mitigates IPv4 attack traffic.

You can configure the DNS Malformed countermeasure when you create or edit a mitigation and when you edit a mitigation on the TMS Mitigation Status page. For information about configuring mitigations, see "Configuring and Deleting TMS Mitigations" on page 110.For information about the TMS Mitigation Status page, see "About the TMS Mitigation Status Page" on page 100.

About the DNS Malformed countermeasure

This countermeasure checks DNS requests in the following ways:

- If a packet is sent to destination port 53, then the packet is checked for a payload that could be part of a valid DNS message. If the payload is missing, then the packet is dropped.
- If the packet is a valid DNS message, then the message is checked for RFC conformance. If the message does not conform to RFC standards, then the packet is dropped.

The source host is not blacklisted in this countermeasure.

Configuring the DNS Malformed countermeasure when adding or editing a mitigation

To configure the DNS Malformed countermeasure when adding or editing a mitigation:

1. Navigate to the **Countermeasures** tab of the mitigation.
   See "Adding and editing a TMS mitigation" on page 110.
2. Scroll to the DNS Mitigation section.
3. Select the **Enable Malformed DNS Filtering** check box.
4. Click **Save**.

Configuring the DNS Malformed countermeasure on the TMS Mitigation Status page

To configure the DNS Malformed countermeasure on the TMS Mitigations Status page:

1. Navigate to the TMS Mitigation Status page.
   See "Navigating to the TMS Mitigation Status page" on page 100.
2. On the Countermeasures tab, click (expand) for the DNS Malformed countermeasure.
   Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations.
3. Select the **Enable Malformed DNS Filtering** check box.
4. Click **Save**.

About locked settings

When (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring the DNS NXDomain Rate Limiting Countermeasure

Introduction

The DNS NXDomain Rate Limiting countermeasure monitors response packets for hosts that send requests that might cause non-existent domain (NXDomain) responses to be generated. This countermeasure protects against DNS cache poisoning and dictionary attacks. Any host that generates more consecutive failed DNS requests than the configured limit is blacklisted. The DNS NXDomain Rate Limiting countermeasure mitigates IPv4 attack traffic.

You can configure the DNS NXDomain Rate Limiting countermeasure when you create or edit a mitigation and when you edit a mitigation on the TMS Mitigation Status page. For information about configuring mitigations, see "Configuring and Deleting TMS Mitigations" on page 110.For information about the TMS Mitigation Status page, see "About the TMS Mitigation Status Page" on page 100.

About the DNS NXDomain Rate Limiting countermeasure

This countermeasure requires that the TMS appliance be able to receive requests and responses so that it can detect and correlate the domain-specific relationship. With an offramp deployment, a TMS port must be configured to listen to DNS NXDomain responses from a network SPAN port. With an inline deployment, the appliance's mitigation capability option must be enabled on both the input and output interfaces. The TMS port can be configured on the Patch Panel tab of the TMS appliance.

Configuring the DNS NXDomain Rate Limiting countermeasure when adding or editing a mitigation

To configure the DNS NXDomain Rate Limiting countermeasure when adding or editing a mitigation:

1. Navigate to the **Countermeasures** tab of the mitigation.
   See "Adding and editing a TMS mitigation" on page 110.
2. Scroll to the DNS Mitigations section.
3. Configure the following settings for the DNS NXDomain Rate Limiting countermeasure:

Setting	Procedure
Enable DNS NXDomain Rate Limiting check box	Select to enable this countermeasure.
DNS NXDomain Rate Limit box	Type the *number of failed queries per second* to allow.

4. Click **Save**.

Configuring the DNS NXDomain Rate Limiting countermeasure on the TMS Mitigation Status page

To configure the DNS NXDomain Rate Limiting countermeasure on the TMS Mitigations Status page:

1. Navigate to the TMS Mitigation Status page.
   See "Navigating to the TMS Mitigation Status page" on page 100.

2. On the Countermeasures tab, click (expand) for the DNS NXDomain Rate Limiting countermeasure.

   Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations.

3. Configure the following settings for the DNS NXDomain Rate Limiting countermeasure:

Setting	Procedure
Enable DNS NXDomain Rate Limiting check box	Select to enable this countermeasure.
DNS NXDomain Rate Limit box	Type the *number of failed queries per second* to allow.

4. To download information about blocked hosts, use the following buttons:

Button	Procedure
Download Blocked Hosts	Click to download a .txt file containing a list of the hosts blocked by this countermeasure. You can use this information to refine other countermeasure settings in the mitigation.
Download Top Blocked Hosts	Click to download a .txt file containing a list of the most frequently blacklisted host. You can use this information to refine other countermeasure settings in the mitigation.

5. Click **Save**.

About locked settings

When (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring the DNS Rate Limiting Countermeasure

Introduction

The DNS Rate Limiting countermeasure limits the number of DNS queries that a host can send per second. This countermeasure prevents attacks from legitimate hosts who misuse DNS requests to flood DNS servers. After a host is authenticated, this countermeasure monitors the DNS queries from the source IP address. Any traffic that exceeds the configured rate limit is dropped and the source is blacklisted. The DNS Rate Limiting countermeasure mitigates IPv4 attack traffic.

You can configure the DNS Rate Limiting countermeasure when you create or edit a mitigation and when you edit a mitigation on the TMS Mitigation Status page. For information about configuring mitigations, see "Configuring and Deleting TMS Mitigations" on page 110.For information about the TMS Mitigation Status page, see "About the TMS Mitigation Status Page" on page 100.

Configuring the DNS Rate Limiting countermeasure when adding or editing a mitigation

To configure the DNS Rate Limiting countermeasure when adding or editing a mitigation:

1. Navigate to the **Countermeasures** tab of the mitigation.
   See "Adding and editing a TMS mitigation" on page 110.
2. Scroll to the DNS Mitigations section.
3. Configure the following settings for the DNS Rate Limiting countermeasure:

Setting	Procedure
Enable DNS Query Rate Limiting check box	Select to enable this countermeasure.
DNS Query Rate Limit box	Type the *number of queries* per second to allow.

4. Click **Save**.

Configuring the DNS Rate Limiting countermeasure on the TMS Mitigation Status page

To configure the DNS Rate Limiting countermeasure on the TMS Mitigations Status page:

1. Navigate to the TMS Mitigation Status page.
   See "Navigating to the TMS Mitigation Status page" on page 100.
2. On the Countermeasures tab, click ⊞ (expand) for the DNS Rate Limiting countermeasure.
   Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations.
3. Configure the following settings for the DNS Rate Limiting countermeasure:

Setting	Procedure
Enable DNS Query Rate Limiting check box	Select to enable this countermeasure.
DNS Query Rate Limit box	Type the *number of queries* per second to allow.

4. To download information about blocked hosts, use the following buttons:

Button	Procedure
Download Blocked Hosts	Click to download a .txt file containing a list of the hosts blocked by this countermeasure. You can use this information to refine other countermeasure settings in the mitigation.
Download Top Blocked Hosts	Click to download a .txt file containing a list of the most frequently blacklisted host. You can use this information to refine other countermeasure settings in the mitigation.

5. Click **Save**.

About locked settings

When (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring the DNS Regular Expression Countermeasure

Introduction

You can use the DNS Regular Expression countermeasure to drop malicious inbound DNS message packets based on regular expression matching and other filter settings you configure. This countermeasure can also blacklist hosts that send packets that are dropped. You can apply this countermeasure to your choice of inbound DNS queries, inbound DNS replies, or both. The selected DNS message types are compared against settings in DNS filters and/or regular expression entries in DNS filter lists.

Note: This countermeasure does not blacklist hosts by default. You must choose to do so.

For inline TMS appliance deployments only: If this countermeasure is configured to drop inbound replies, it only drops replies destined for an attack target within the protected network. It does not drop outbound replies from DNS servers in the network.

About configuring this countermeasure in a TMS mitigation

When you configure a TMS mitigation, you can use the Payload tab to configure the DNS Regular Expression countermeasure to mitigate IPv4 or IPv6 traffic. You can also configure this countermeasure in the Countermeasures section on the TMS Mitigation Status page.

For information about configuring mitigations, see “Configuring and Deleting TMS Mitigations” on page 110.

For information about adjusting these and other mitigation settings on the TMS Mitigation Status page, see “About the TMS Mitigation Status Page” on page 100.

About domain name regular expression matching

This countermeasure can compare domain name fields in a DNS message to one or more regular expressions. It can detect regular expression matches within single DNS packets.

Note: This countermeasure cannot detect regular expression matches that span multiple packets.

The regular expressions to match in this countermeasure are specified in DNS filters, DNS filter lists, or both.

For more information, see “About configuring DNS filters” below and “About configuring DNS filter lists” on the facing page.

About configuring DNS filters

A DNS filter is a group of settings. This countermeasure compares the settings in each DNS filter to fields in a DNS message.

You can add up to five DNS filters to the DNS Regular Expression countermeasure inline. You can also edit DNS filter settings and remove DNS filters inline.

Note: In order to add, edit, or remove DNS filters, your account must allow editing and the DNS filter section must be unlocked. (See “Configuring the DNS Regular Expression Countermeasure” above.)

A DNS filter is enabled for matching when one or more of its settings are specified (not blank). If all settings in a DNS filter are blank, the filter is disabled.

You can remove a DNS filter without clearing its settings first. All DNS filters can be removed except DNS Filter 1. However, you can clear the settings in DNS Filter 1 to disable it.

For more information about matching, see "About DNS filter matching" below. For descriptions of the settings in a DNS filter, see "DNS Filter Settings" on page 172.

About DNS filter matching

The TMS mitigation uses this countermeasure to compare the following fields in a DNS message to their corresponding settings in a DNS filter:

DNS Message Field	Description	DNS Filter Setting	Value
Resource Record (RR)	A numeric DNS RR type value, such as "2" for a name server (NS) record.	**Resource Record Types**	DNS RR types, such as "NS" and "PTR," and/or RR type numeric values.
Recursion Desired (RD) flag	The state of the RD flag, "1" (set) or "0" (unset).	**Recursion Desired Flag**	Set, Unset, or Ignored
QNAME	The queried domain name in the Question section of a DNS message.	**Domain Regular Expression**	A domain name regular expression, such as ".+\.example\.com"
NAME	A domain name in the Answer section, Authority section, or Additional section of a DNS message.		

If each message field matches its corresponding setting in a filter, the TMS mitigation classifies the DNS message as a match for that filter. When matching, null message fields and blank filter settings are ignored.

When multiple DNS filters are configured, the match results for individual filters are combined using an OR operation. So, if any DNS filter is a match, the TMS mitigation classifies the DNS message as a match.

For more information, see "DNS Filter Settings" on page 172.

About configuring DNS filter lists

DNS filter lists are lists of domain regular expression entries that are configured outside the mitigation. You can select up to 32 DNS filter lists to include in this countermeasure. You can also remove DNS filter lists from this countermeasure.

To create or edit DNS filter lists, contact your service provider.

Note: In order to select and remove DNS filter lists, your account must allow editing and the DNS Filter Lists section must be unlocked. (See "Configuring the DNS Regular Expression Countermeasure" on page 168.) Removing a filter list from this countermeasure does not delete the list itself.

If you include multiple DNS filter lists, when matching, this countermeasure combines them using an OR operation. This means that the TMS mitigation yields a match when a DNS message matches any regular expression entry in any single DNS filter list that you included in this countermeasure.

For more about configuring DNS filter list settings in this countermeasure, see "Configuring the DNS Regular Expression countermeasure" below and "DNS Regular Expression countermeasure settings" on the facing page.

About combined matching between DNS filters and DNS filter lists

To determine if a DNS message is a match, this countermeasure first matches the message separately against DNS filters and DNS filter lists. Then, it combines the results using either an AND or an OR operation (your choice). If you choose AND, a match occurs when both DNS filters and DNS filter lists yield a match. If you choose OR, a match occurs when either DNS filters yield a match or DNS filter lists yield a match.

For more information about combined matching, see "DNS Filter List Settings" on page 173.

Configuring the DNS Regular Expression countermeasure

Note: You can only configure this countermeasure to mitigate IPv4 DNS traffic.

To configure the DNS Regular Expression countermeasure for TMS mitigations:

1. Do one of the following:
   - Navigate to the **Payload** tab of the mitigation.

     See "Adding and editing a TMS mitigation" on page 110.
   - Navigate to the TMS Mitigation Status page. In the Countermeasures section, click  (expand) for the DNS Regular Expression countermeasure to show its settings.

     See "Navigating to the TMS Mitigation Status page" on page 100.

     Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations.

     For more information about locked settings, see "About locked settings" on page 173.
2. Configure the DNS Regular Expression countermeasure settings. These include the settings for inline DNS filters and DNS filter lists. They also include settings that are only available on the TMS Mitigation Status page.

   See "DNS Regular Expression countermeasure settings" on the facing page.
3. Click **Save**.

DNS Regular Expression countermeasure settings

Use the following table to configure the DNS regular expression settings.

Setting	Procedure
Message Types to Filter options	Click **Inbound Queries** (default), **Inbound Replies**, or **Both** to select the type(s) of inbound DNS messages to match against the settings in DNS filters and/or regular expression entries in DNS filter lists. **Note:** If you select an option and then later select a different option, the graph that appears on the TMS Mitigation Status page displays all the traffic that was dropped regardless of the option that is currently selected.
Action to Apply options	Click **Blacklist Hosts** (default) or **Drop Traffic** to select the action to apply to matched or unmatched DNS messages. ■ **Blacklist Hosts**—Performs the **Drop Traffic** action (described below) and also blacklists hosts that send DNS messages that are dropped. ■ **Drop Traffic**—Drops the selected DNS Message types that either match, or do not match, a DNS filter and/or DNS filter list. Use **Apply Action to** (described below) to apply the selected action to matched traffic or unmatched traffic.
Apply Action to options	Click **Matched Traffic** (default) or **Unmatched Traffic** to apply the action **Blacklist Hosts** or **Drop Traffic** (described above) to matched or unmatched DNS messages. The **Matched Traffic** and **Unmatched Traffic** options specify DNS messages with the following characteristics: ■ **Matched Traffic**—DNS message types that match the settings in at least one DNS filter and/or at least one regular expression in a DNS filter list. ■ **Unmatched Traffic**—DNS message types that do not match the settings in any DNS filter and/or any regular expression in any DNS filter list. The **Message Types to Filter** setting (described above) specifies the type of matched or unmatched DNS messages that the action will be applied to: inbound queries, responses, or both. Use the **AND** or **OR** options under **Combine DNS Filters with DNS Filter Lists using** to specify how this countermeasure combines the separate DNS filter and DNS filter list results to classify a DNS message as matched or unmatched. (See "About combined matching between DNS filters and DNS filter lists" on the previous page.)

Setting	Procedure
DNS Filter Settings	
Add Filter button	Click **Add Filter** to add a new group of inline DNS filter settings to this countermeasure. Initially, all settings in the new filter are blank. You must specify at least one setting to enable the filter.
Resource Record Types selector box	Click in the **Resource Record Types** box. In the list, click a DNS resource record (RR) type to add it to the selection in the box. You can add multiple RR types to the selection. Each entry in the list is an RR type name and its numeric value. **Tip:** To quickly find the RR type you want to add, start typing its name or numeric value in the selector box. To delete an RR type from the selection, click the "x" in the RR type label. You can also click in the box and then press the `DELETE` key multiple times to delete RR type labels from the selection. To disable matching by RR type, clear the **Resource Record Types** box. **Note:** Your selection can include numeric values for RR types that are not on the list. To add the numeric value for an unlisted RR type to the selection, in the selector box, type the numeric value. Press `ENTER` or click the highlighted value to add it to the selection. For a list of DNS resource record types along with their values and meanings, see section 3.2.2 of RFC 1035 on the IETF.org Web site (https://www.ietf.org/rfc/rfc1035.txt).
Recursion Desired Flag options	Click **Ignore** (default), **Set**, or **Unset** to match on the value of the Recursion Desired flag bit in a DNS message. ■ **Ignored**—RD flag bit is ignored when matching. ■ **Set**—Match if RD flag bit is set (1). ■ **Unset**—Match if RD flag bit is unset (0) In a DNS query from a DNS client to a DNS name server, the RD flag bit can be set or unset. If the RD flag bit is set, and the DNS name server cannot resolve the query, it forwards the query to successive upstream name servers until it receives a response that contains a fully resolved domain name. The RD flag bit value is copied into every response to the query.
Domain Regular Expression box	Type a regular expression (in PCRE format and single-line mode) to match on the specified pattern of characters in the domain name (QNAME or NAME) field in a DNS message. **Note:** DNS regular expressions are case-insensitive by default. To perform case-sensitive matching, preface the expression with "(?-i)".
Remove button	Click **Remove** in a DNS filter to remove that DNS filter from this countermeasure. **Tip:** You cannot remove DNS Filter 1, however, you can disable it by clearing all of its settings. To clear all DNS Filter 1 settings at once, click **Remove** in the DNS Filter 1 group.

Setting	Procedure
DNS Filter List Settings	
DNS Filter Lists box	Shows all DNS filter lists that are currently included in this mitigation.
Select Filter List button	Click **Select Filter List** to add one or more DNS filter lists to the mitigation. In the Select one or more DNS Filter Lists window, select all the DNS filter lists to add to the mitigation, and then click **OK**. If you want to configure more DNS filter lists, contact your service provider.
Combine DNS Filters with DNS Filter Lists using options	Click **OR** (default) or **AND** to specify the condition under which this countermeasure returns a match for DNS filters and DNS filter lists combined. ■ **OR**—Returns a match if either DNS filters or DNS filter lists yield a match. ■ **AND**—Returns a match if both DNS filters and DNS filter lists yield a match. See "About combined matching between DNS filters and DNS filter lists" on page 170.
Settings on the TMS Mitigation Status Page (only)	
Download Top FQDNs button	Click to download a .txt file containing a list of the most frequently accessed FQDNs. You can use this information to help you refine your regular expression.
Download Top RDNs button	Click to download a .txt file containing a list of the most frequently accessed RDNs. You can use this information to help you refine your regular expression.
Test Regular Expression button	Click to test the effectiveness of a regular expression in mitigating the attack associated with this alert.
View All Filter List Matched Rates button	Click to view all filter list matched rate information.
Save	Click to save any changes you made to the settings for this countermeasure on the TMS Mitigation Status page.

About locked settings

When (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring the HTTP Malformed Countermeasure

Introduction

The HTTP Malformed countermeasure can filter HTTP traffic that does not conform to RFC standards and HTTP traffic that behaves abnormally. The HTTP Malformed countermeasure mitigates IPv4 attack traffic.

You can configure the HTTP Malformed countermeasure when you create or edit a mitigation and when you edit a mitigation on the TMS Mitigation Status page. For information about configuring mitigations, see "Configuring and Deleting TMS Mitigations" on page 110.For information about the TMS Mitigation Status page, see "About the TMS Mitigation Status Page" on page 100.

About the HTTP Malformed countermeasure

By default, the HTTP Malformed countermeasure filters HTTP traffic that does not conform to RFC standards for valid request headers. This countermeasure protects against attacks that send invalid or blank HTTP requests to a server to exhaust resources or to exploit vulnerabilities. Each request is checked for compliance with RFC standards. If a request does not conform to standards, then the packet is dropped and the source host is blacklisted.

The HTTP Malformed countermeasure can also filter HTTP traffic that conforms to RFC standards for valid request headers but has other abnormal HTTP behavior. To filter this type of HTTP traffic, you must change the default enforcement level of the countermeasure. If a request does not conform to these higher enforcement standards, the packet is dropped and the source is blacklisted. For example, these higher enforcement levels would block many botnet attacks.

Configuring the HTTP Malformed countermeasure when adding or editing a mitigation

To configure the HTTP Malformed countermeasure when adding or editing a mitigation:

1. Navigate to the **Countermeasures** tab of the mitigation.

   See "Adding and editing a TMS mitigation" on page 110.
2. Scroll to the HTTP Mitigations section.
3. Select the **Enable Malformed HTTP Filtering** check box.
4. Click (low), (medium), or (high) to set the enforcement level for the countermeasure.

   Low is selected by default when this setting is enabled. When Low is selected, the countermeasure filters traffic that does not conform to RFC standards for valid request headers.

   A medium or high enforcement level blocks traffic that conforms to RFC standards for valid request headers but has other abnormal HTTP behavior. As you increase the enforcement level, more malicious HTTP traffic is dropped, but the likelihood of dropping legitimate traffic also increases.

   Note: If the **Enable Malformed HTTP Filtering** setting is locked, only the locked enforcement level appears.
5. Click **Save**.

Configuring the HTTP Malformed countermeasure on the TMS Mitigation Status page

To configure the HTTP Malformed countermeasure on the TMS Mitigations Status page:

1. Navigate to the TMS Mitigation Status page.

   See "Navigating to the TMS Mitigation Status page" on page 100.

2. On the Countermeasures tab, click (expand) for the HTTP Malformed countermeasure.

   Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations.

3. Select the **Enable Malformed HTTP Filtering** check box.

4. Click (low), (medium), or (high) to set the enforcement level for the countermeasure.

   Low is selected by default when this setting is enabled. When Low is selected, the countermeasure filters traffic that does not conform to RFC standards for valid request headers.

   A medium or high enforcement level blocks traffic that conforms to RFC standards for valid request headers but has other abnormal HTTP behavior. As you increase the enforcement level, more malicious HTTP traffic is dropped, but the likelihood of dropping legitimate traffic also increases.

   Note: If the **Enable Malformed HTTP Filtering** setting is locked, only the locked enforcement level appears.

5. To download information about blocked hosts, use the following buttons:

Button	Procedure
Download Blocked Hosts	Click to download a .txt file containing a list of the hosts blocked by this countermeasure. You can use this information to refine other countermeasure settings in the mitigation.
Download Top Blocked Hosts	Click to download a .txt file containing a list of the most frequently blacklisted host. You can use this information to refine other countermeasure settings in the mitigation.

6. Click **Save**.

About locked settings

When (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring the HTTP Rate Limiting Countermeasure

Introduction

The HTTP Rate Limiting countermeasure limits the rates at which a host can send HTTP requests. This countermeasure prevents a host from overwhelming the resources of a Web server, either by sending too many requests or by requesting too many unique objects. This countermeasure monitors the HTTP requests from the source IP address. Any traffic that exceeds either of the configured rate limits is dropped and the source host is blacklisted. The HTTP Rate Limiting countermeasure mitigates IPv4 attack traffic.

You can configure the HTTP Rate Limiting countermeasure when you create or edit a mitigation and when you edit a mitigation on the TMS Mitigation Status page. For information about configuring mitigations, see "Configuring and Deleting TMS Mitigations" on page 110.For information about the TMS Mitigation Status page, see "About the TMS Mitigation Status Page" on page 100.

About the HTTP Rate Limiting countermeasure

The default HTTP rate limits are usually acceptable for typical users. Because a Web server can be heavily loaded by a small number of HTTP requests, do not increase the limits by large amounts without careful consideration. If you must make an exception for a content mirror server, you can add it to a pass rule in the Black/White Lists countermeasure.

See "Configuring the Black/White Lists Countermeasure" on page 126.

Configuring the HTTP Rate Limiting countermeasure when adding or editing a mitigation

To configure the HTTP Rate Limiting countermeasure when adding or editing a mitigation:

1. Navigate to the **Countermeasures** tab of the mitigation.
   See "Adding and editing a TMS mitigation" on page 110.
2. Scroll to the HTTP Mitigations section.
3. Configure the settings for the HTTP Rate Limiting countermeasure.
   See "HTTP Rate Limiting countermeasure settings" on the facing page.
4. Click **Save**.

Configuring the HTTP Rate Limiting countermeasure on the TMS Mitigation Status page

To configure the HTTP Rate Limiting countermeasure on the TMS Mitigations Status page:

1. Navigate to the TMS Mitigation Status page.
   See "Navigating to the TMS Mitigation Status page" on page 100.
2. On the Countermeasures tab, click (expand) for the HTTP Rate Limiting countermeasure.
   Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations.
3. Configure the settings for the HTTP Rate Limiting countermeasure.
   See "HTTP Rate Limiting countermeasure settings" on the facing page.

4. To download information about blocked hosts, use the following buttons:

Button	Procedure
Download Blocked Hosts	Click to download a .txt file containing a list of the hosts blocked by this countermeasure. You can use this information to refine other countermeasure settings in the mitigation.
Download Top Blocked Hosts	Click to download a .txt file containing a list of the most frequently blacklisted host. You can use this information to refine other countermeasure settings in the mitigation.

5. Click **Save**.

HTTP Rate Limiting countermeasure settings

Use the following table to configure the HTTP Rate Limiting countermeasure settings:

Setting	Procedure
Enable HTTP Object Limiting check box	Select to mitigate attack traffic when HTTP objects exceed the limit that is configured in the **HTTP Object Limit** box. An HTTP object is a GET request for a particular URL.
HTTP Object Limit box	Type the *number of objects per second* to allow. **Example:** An object might be www.abc.net/an_image.png. When HTTP Object rate limiting is enabled, the number of requests that can be made for that URL are limited to the rate specified in **HTTP Object Limit**.
Enable HTTP Request Limiting check box	Select to mitigate attack traffic when HTTP requests exceed the limit that is configured in the **HTTP Request Limit** box. An HTTP request is any type of request like GET, POST, HEAD, or OPTIONS.
HTTP Request Limit box	Type the *number of requests per second* to allow.

About locked settings

When (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring the SIP Malformed Countermeasure

Introduction

The SIP Malformed countermeasure filters the SIP traffic that does not conform to the RFC standards for valid request headers. This countermeasure protects against attacks that disrupt VoIP service by sending invalid or blank SIP requests. Each request is checked for compliance with RFC standards. If a request does not conform to standards, then the packet is dropped and the source host is blacklisted. The SIP Malformed countermeasure mitigates IPv4 attack traffic.

You can configure the SIP Malformed countermeasure when you create or edit a mitigation and when you edit a mitigation on the TMS Mitigation Status page. For information about configuring mitigations, see "Configuring and Deleting TMS Mitigations" on page 110.For information about the TMS Mitigation Status page, see "About the TMS Mitigation Status Page" on page 100.

Configuring the SIP Malformed countermeasure when adding or editing a mitigation

To configure the SIP Malformed countermeasure when adding or editing a mitigation:

1. Navigate to the **Countermeasures** tab of the mitigation.
   See "Adding and editing a TMS mitigation" on page 110.
2. Scroll to the VoIP / SIP section.
3. Select the **Enable Malformed SIP Filtering** check box.
4. Click **Save**.

Configuring the SIP Malformed countermeasure on the TMS Mitigation Status page

To configure the SIP Malformed countermeasure on the TMS Mitigations Status page:

1. Navigate to the TMS Mitigation Status page.
   See "Navigating to the TMS Mitigation Status page" on page 100.
2. On the Countermeasures tab, click (expand) for the SIP Malformed countermeasure.
   Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations.
3. Select the **Enable Malformed SIP Filtering** check box.
4. To download information about blocked hosts, use the following buttons:

Button	Procedure
Download Blocked Hosts	Click to download a .txt file containing a list of the hosts blocked by this countermeasure. You can use this information to refine other countermeasure settings in the mitigation.
Download Top Blocked Hosts	Click to download a .txt file containing a list of the most frequently blacklisted host. You can use this information to refine other countermeasure settings in the mitigation.

5. Click **Save**.

About locked settings

When  (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring the SIP Request Limiting Countermeasure

Introduction

The SIP Request Limiting countermeasure limits the number of SIP requests that a host can send per second. This countermeasure monitors SIP requests by source IP address, to prevent attacks that disrupt VoIP service by flooding the network with excessive SIP requests. Any traffic that exceeds the configured rate limit is dropped, and the source host is blacklisted. The SIP Request Limiting countermeasure mitigates IPv4 attack traffic.

You can configure the SIP Request Limiting countermeasure when you create or edit a mitigation and when you edit a mitigation on the TMS Mitigation Status page. For information about configuring mitigations, see "Configuring and Deleting TMS Mitigations" on page 110.For information about the TMS Mitigation Status page, see "About the TMS Mitigation Status Page" on page 100.

About the SIP Request Limiting countermeasure

Because Web servers can send a large amount of data in a single request, communication between SIP servers can greatly exceed the source limit. You can protect those servers by adding them to a pass rule in the Black/White Lists countermeasure.

See "Configuring the Black/White Lists Countermeasure" on page 126.

Configuring the SIP Request Limiting countermeasure when adding or editing a mitigation

To configure the SIP Request Limiting countermeasure when adding or editing a mitigation:

1. Navigate to the **Countermeasures** tab of the mitigation.
   See "Adding and editing a TMS mitigation" on page 110.
2. Scroll to the VoIP / SIP section.
3. Use the following table to configure the SIP Request Limiting countermeasure settings:

Setting	Procedure
Enable SIP Source Limiting check box	Select to enable this countermeasure.
SIP Source Limit box	Type the maximum *number of SIP request per second* to allow.

4. Click **Save**.

Configuring the SIP Request Limiting countermeasure on the TMS Mitigation Status page

To configure the SIP Request Limiting countermeasure on the TMS Mitigations Status page:

1. Navigate to the TMS Mitigation Status page.
   See "Navigating to the TMS Mitigation Status page" on page 100.
2. On the Countermeasures tab, click ⊞ (expand) for the SIP Request Limiting countermeasure.

Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations.

3. Use the following table to configure the SIP Request Limiting countermeasure settings:

Setting	Procedure
Enable SIP Source Limiting check box	Select to enable this countermeasure.
SIP Source Limit box	Type the maximum *number of SIP request per second* to allow.

4. To download information about blocked hosts, use the following buttons:

Button	Procedure
Download Blocked Hosts	Click to download a .txt file containing a list of the hosts blocked by this countermeasure. You can use this information to refine other countermeasure settings in the mitigation.
Download Top Blocked Hosts	Click to download a .txt file containing a list of the most frequently blacklisted host. You can use this information to refine other countermeasure settings in the mitigation.

5. Click **Save**.

About locked settings

When (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring the SSL Negotiation Countermeasure

Introduction

The SSL Negotiation countermeasure is designed to protect arbitrary services from attacks that target the SSL and TLS protocols. The SSL Negotiation countermeasure mitigates IPv4 attack traffic.

You can configure the SSL Negotiation countermeasure when you create or edit a mitigation and when you edit a mitigation on the TMS Mitigation Status page. For information about configuring mitigations, see "Configuring and Deleting TMS Mitigations" on page 110.For information about the TMS Mitigation Status page, see "About the TMS Mitigation Status Page" on page 100.

About the SSL Negotiation countermeasure

The SSL Negotiation countermeasure uses protocol events to create a blacklist of source addresses and a whitelist of source addresses. This countermeasure passes packets that have a source address on the whitelist and drops packets that have a source address on the blacklist.

When the source address of a packet is not on the whitelist or blacklist, the countermeasure associates the packet with a pending connection. The countermeasure then performs various checks to determine if the packet should be passed or if its source address should be added to the whitelist or blacklist.

The SSL Negotiation countermeasure also blacklists hosts that take too long to complete a handshake.

Configuring the SSL Negotiation countermeasure when adding or editing a mitigation

To configure the SSL Negotiation countermeasure when adding or editing a mitigation:

1. Navigate to the **Countermeasures** tab of the mitigation.
   See "Adding and editing a TMS mitigation" on page 110.
2. Scroll to the SSL Negotiation section.
3. Configure the following settings for the SSL Negotiation countermeasure.

Icon	Description
	Hover over this icon to display the download options. The download options that appear vary depending on the page that you are on and c include the following: ■ PDF - Click to download the page in PDF format. ■ XML- Click to download the page in XML format. ■ CSV - Click to download a page in CSV format. The download

4. Click **Save**.

Configuring the SSL Negotiation countermeasure on the TMS Mitigation Status page

To configure the SSL Negotiation countermeasure on the TMS Mitigations Status page:

1. Navigate to the TMS Mitigation Status page.
   See "Navigating to the TMS Mitigation Status page" on page 100.
2. On the Countermeasures tab, click ⊞ (expand) for the SSL Negotiation countermeasure.
   Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations.
3. Configure the following settings for the SSL Negotiation countermeasure.

Icon	Description
	Hover over this icon to display the download options. The download options that appear vary depending on the page that you are on and c include the following: ■ PDF - Click to download the page in PDF format. ■ XML- Click to download the page in XML format.

4. To download information about blocked hosts, use the following buttons:

Button	Procedure
Download Blocked Hosts	Click to download a .txt file containing a list of the hosts blocked by this countermeasure. You can use this information to refine other countermeasure settings in the mitigation.
Download Top Blocked Hosts	Click to download a .txt file containing a list of the most frequently blacklisted host. You can use this information to refine other countermeasure settings in the mitigation.

5. Click **Save**.

SSL Negotiation countermeasure advanced settings

The SSL Negotiation countermeasure advanced settings appear when you click ▶ (expand) next to Advanced Settings.

Caution: The default values are based on extensive profiling of normal SSL client behaviors. It is recommended that you do not change the default advanced setting values unless required to avoid blocking specific SSL client traffic in your network.

Use the following table to configure the SSL Negotiation countermeasure advanced settings:

Setting	Procedure
Maximum cipher suites	Type the maximum number of cipher suites for which a client is allowed to indicate support. When this value is exceeded, the countermeasure blacklists the client, and drops the packet. The default value is 100. While this default value significantly exceeds the norm, it is small enough to reduce the amount of time that a server spends searching the list of ciphers to look for a supported option.
Maximum client extensions	Type the maximum number of extensions that a client is allowed to include. When this value is exceeded, the countermeasure blacklists the client, and drops the packet. The default value is 10 to reduce the impact of malicious clients on the server.
Maximum open uncompleted connections	Type the maximum number of open uncompleted connections. This is the maximum number of times that a client can open and close a connection without completing the SSL handshake and sending encrypted data. When this value is exceeded, the countermeasure blacklists the client and drops the packet. The default value is 25. This value allows valid clients to open multiple parallel connections and only use a few of them, while still blocking attacks.
Maximum seconds before application data	Type the maximum number of seconds that a client is allowed between opening a connection and completing the SSL handshake and sending the first bytes of encrypted application data. If no application data is sent in this amount of time, the countermeasure blacklists the client, and drops the packet. The default value is 30 seconds, which allows for multiple round-trip times and multiple packet retransmits.
Minimum seconds connection tracked	Type the number of seconds for the minimum amount of time that it will take for a valid client to complete the SSL handshake and start sending encrypted application data. The default value is 15 seconds, which allows sufficient time for a valid SSL connection to complete during an attack.
Allow client alert messages	Select to allow client alert messages. This setting is selected by default because Arbor is not currently aware of the use of client alert messages in attacks.

About locked settings

When  (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring the TCP Connection Limiting Countermeasure

Introduction

The TCP Connection Limiting countermeasure limits the number of concurrent TCP connections that can originate from a single host. This countermeasure prevents attacks that overwhelm the victim's connection resources with an excessive number of TCP connections. The TCP Connection Limiting countermeasure mitigates IPv4 attack traffic.

For example, some botnets open hundreds of active or inactive TCP connections. A sufficiently large number of connections can consume all of the resources of a server and prevent the server from accepting legitimate traffic.

You can configure the TCP Connection Limiting countermeasure when you create or edit a mitigation and when you edit a mitigation on the TMS Mitigation Status page. For information about configuring mitigations, see “Configuring and Deleting TMS Mitigations” on page 110.For information about the TMS Mitigation Status page, see “About the TMS Mitigation Status Page” on page 100.

About the TCP Connection Limiting countermeasure

The TCP Connection Limiting countermeasure monitors the TCP requests from a source host and checks for a SYN followed by an ACK for the same 4-tuple (src/dst IP and src/dst port combination). This countermeasure does not require replies from the server to function because it is designed to work in an asymmetric environment.

When the number of concurrent connections from a single host exceeds the connection limit configured in this countermeasure, then one of the following happens depending on how this countermeasure is configured:

- The host is blacklisted.
- The host's connections that exceed the connection limit are dropped and the connections are reset.
- The host's idle connections are ignored and not counted to keep the host within the connection limit.

Configuring the TCP Connection Limiting countermeasure when adding or editing a mitigation

To configure the TCP Connection Limiting countermeasure when adding or editing a mitigation:

1. Navigate to the **Countermeasures** tab of the mitigation.
   See “Adding and editing a TMS mitigation” on page 110.
2. Scroll to the TCP Connection Limiting section.
3. Configure the settings for the TCP Connection Limiting countermeasure.
   See “TCP Connection Limiting countermeasure settings” on the next page.
4. Click **Save**.

Configuring the TCP Connection Limiting countermeasure on the TMS Mitigation Status page

To configure the TCP Connection Limiting countermeasure on the TMS Mitigations Status page:

1. Navigate to the TMS Mitigation Status page.

   See "Navigating to the TMS Mitigation Status page" on page 100.
2. On the Countermeasures tab, click (expand) for the TCP Connection Limiting countermeasure.

   Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations.
3. Configure the settings for the TCP Connection Limiting countermeasure.

   See "TCP Connection Limiting countermeasure settings" below.
4. To evaluate the impact of the settings that you selected, view the following statistics:
   - Connections Blocked Rate

     If you selected the **Blacklist** action, then the Connections Blocked Rate displays the number of hosts that are being added to the blacklist per second. If you selected the **Drop** action, then the Connections Blocked Rate displays the number of connections that are dropped per second.
   - Successful Connections Rate

     Successful Connections Rate displays the number of successful connections per second.
5. To download a .txt file containing a list of all the hosts currently blacklisted by this countermeasure, click the **Download Blocked Hosts** button.

   You can use this information to refine other countermeasure settings in the mitigation.
6. Click **Save**.

TCP Connection Limiting countermeasure settings

Use the following table to configure the TCP Connection Limiting countermeasure settings:

Setting	Procedure
Enable TCP Connection Limiting check box	Select to enable this countermeasure. This countermeasure is disabled by default.
Action to Apply to Offending Host options	Click **Blacklist** or **Drop** to select the action to apply to the offending host. The **Blacklist** option is selected by default, and it blacklists the offending host for one minute. For additional information about blacklisting, see "About dynamic blacklisting" on page 97. The **Drop** option drops the connections from the offending host that exceed the number of connections specified in the **Open Connection Limit Per Host** box. When a connection is dropped, it is also reset on the server to free up server state.

Icon	Description
	Hover over this icon to display the download options. The download options that appear vary depending on the page that you are on and can include the following: ■ PDF - Click to download the page in PDF format. ■ XML - Click to download the page in XML format.

TCP Connection Limiting countermeasure advanced settings

The TCP Connection Limiting countermeasure advanced settings appear when you click ▶ (expand) next to Advanced Settings.

Important: You should only edit these advanced settings if the default values do not work well in your environment.

Use the following table to configure the TCP Connection Limiting countermeasure advanced settings:

Setting	Procedure
Ignore Idle Connections options	Click **Enabled** to ignore idle connections. Click **Disabled** to not ignore idle connections. **Enabled** is selected by default. A connection must be inactive for the number of seconds specified in the **Idle Timeout Value** box before it is considered to be an idle connection. With **Enabled** selected, if a host exceeds the connection limit and the host has idle connections, then the idle connections are ignored and are not counted towards the connection limit. If ignoring the host's idle connections keeps the host from exceeding the connection limit, then no action is taken to blacklist the host or to drop any of the host's active connections. You can use this countermeasure in conjunction with the TCP Connection Reset countermeasure. You can use the TCP Connection Reset countermeasure to identify and drop idle connections that are part of an attack that opens a large number of idle TCP connections. You can then use this countermeasure to ignore idle connections that are legitimate. With **Disabled** selected, idle connections are not ignored to keep a host from exceeding the connection limit.
Idle Timeout Value box	Type the number of seconds that a connection must not send any data before it can be considered idle. The default value is 60 seconds.

About locked settings

When (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Configuring the TCP Connection Reset Countermeasure

Introduction

The TCP Connection Reset countermeasure tracks established TCP connections and drops the traffic when a connection remains idle for too long. This countermeasure can prevent idle TCP connections from filling server connection tables. This countermeasure also allows you to blacklist hosts that send extremely slow requests. The TCP Connection Reset countermeasure mitigates IPv4 attack traffic.

Although TCP Connection Reset is primarily event-driven, it includes per-packet monitoring of TCP packets so that TCP packet fragments are detected both to reset idle timers and to detect highly fragmented slow application requests.

You can configure the TCP Connection Reset countermeasure when you create or edit a mitigation and when you edit a mitigation on the TMS Mitigation Status page.For information about configuring mitigations, see “Configuring and Deleting TMS Mitigations” on page 110.For information about the TMS Mitigation Status page, see “About the TMS Mitigation Status Page” on page 100.

About the TCP Connection Reset countermeasure

When a TCP connection is first detected, the source host must send a specified amount of payload data (Initial Timeout Required Data) within a certain amount of time (TCP Connection Initial Timeout). After the required amount of data is sent, the source host must only send a TCP packet more frequently than the TCP Connection Idle Timeout setting. A source host that does not send the specified amount of data is blacklisted.

Peakflow SP applies the TCP Connection Reset countermeasure to the following ports:

- 80—HTTP traffic (Web traffic)
- 443—HTTPS traffic (Web traffic)
- 25—SMTP traffic (mail)

You cannot manually configure ports for this countermeasure.

Configuring the TCP Connection Reset countermeasure when adding or editing a mitigation

To configure the TCP Connection Reset countermeasure when adding or editing a mitigation:

1. Navigate to the **Countermeasures** tab of the mitigation.
   See “Adding and editing a TMS mitigation” on page 110.
2. Scroll to the TCP Connection Reset section.
3. Configure the settings for the TCP Connection Reset countermeasure.
   See “TCP Connection Reset countermeasure settings” on the facing page.
4. Click **Save**.

Configuring the TCP Connection Reset countermeasure on the TMS Mitigation Status page

To configure the TCP Connection Reset countermeasure on the TMS Mitigation Status page:

1. Navigate to the TMS Mitigation Status page.
   See “Navigating to the TMS Mitigation Status page” on page 100.

2. On the Countermeasures tab, click (expand) for the TCP Connection Reset countermeasure.

   Note: You can edit only the countermeasure settings that are not locked. You must also be in an account group that is assigned the capability to edit mitigations.

3. Configure the settings for the TCP Connection Reset countermeasure.

   See "TCP Connection Reset countermeasure settings" below.

4. To download information about blocked hosts, use the following buttons:

Button	Procedure
Download Blocked Hosts	Click to download a.txt file containing a list of the hosts blocked by this countermeasure. You can use this information to refine other countermeasure settings in the mitigation.
Download Top Blocked Hosts	Click to download a .txt file containing a list of the most frequently blacklisted host. You can use this information to refine other countermeasure settings in the mitigation.

5. Click **Save**.

TCP Connection Reset countermeasure settings

Use the following table to configure the TCP Connection Reset countermeasure settings

Setting	Procedure
Enable TCP Connection Reset check box	Select to enable this countermeasure.
TCP Connection Idle Timeout box	Type the *number of seconds* before a connection is filtered. Any connection that is idle for this amount of time is reset.
TCP Connection Initial Timeout box	Type the *number of seconds* that a connection can be idle after it is first established. The default value is 10 seconds.
Initial Timeout Required Data box	Type the *number of bytes* that a host must send within the initial timeout period for the timeout to be canceled. The default value is 40 bytes. With the default values, the connection has 10 seconds in which to send 40 bytes of data. If that amount of data is not sent, then the timeout period in the TCP Connection Idle Timeout box begins.
Track Connections After Initial State check box	Select to track a connection after it leaves the initial state.

Icon	Description
	Hover over this icon to display the download options. The download options that appear vary depending on the page that you are on and can include the following: ■ PDF - Click to download the page in PDF format. ■ XML- Click to download the page in XML format. ■ CSV - Click to download a page in CSV format. The download can be a CSV text file or CSV zip archive file. When the download is a zip archive file, “zip archive” is appended to CSV. ■ Excel-XML - Click to download a page in Excel-XML format. The icon appears only if there are multiple options for downloading a page.
	Click to download a page in PDF format.

Application Slow Request advanced settings

The Application Slow Request advanced settings provide additional configurable options.

Caution: The default values are based on extensive profiling of the behavior of attacks that use extremely slow HTTP requests. It is recommended that you do not change the default advanced setting values unless required to avoid blocking specific HTTP requests in your network.

Use the following table to configure the Application Slow Request advanced settings:

Setting	Procedure
Minimum Request Bit Rate box	Type the minimum bit rate that clients must maintain to avoid being blacklisted. The default value is 200.
Time Period for Computing the Minimum Rate box	Type the number of seconds allowed for computing the minimum rate. The default value is 60 seconds.
Minimum Time to Allow for Header Transmission box	Type the number of seconds allowed for header transmission. The default value is 60 seconds.

About locked settings

When (lock) appears beside a setting, you cannot edit that setting because it has been locked by your service provider.

Chapter 9:
Other Ways to Mitigate Attacks

Introduction

This section describes ways to mitigate attacks without using TMS.

To mitigate attacks using TMS, see "About TMS Mitigations" on page 96.

User access

Only managed services administrators can configure these settings. Managed services administrators and non-administrative users can view mitigations.

In this section

This section contains the following topics:

Mitigating Attacks Using Peakflow SP

Introduction

Peakflow SP provides a variety of mitigation options that you can use to stop or prevent network attacks. You can initiate a mitigation from a DoS alert or from the configuration page of a specific type of mitigation. For additional information about initiating a mitigation, see "Initiating a Mitigation from a DoS Alert" on page 114 and "Mitigation types" below.

For information about navigating the mitigations pages, see "Navigating the Peakflow SP Web UI" on page 13.

Mitigation types

The following table describes the Peakflow SP mitigation types and includes references to information about configuring them:

Type	Description	Reference
Threat Management	Offramps network traffic to a TMS. This mitigation type is useful for attacks on critical resources that use main service ports. This mitigation type provides detailed mitigation statistics.	"About TMS Mitigations" on page 96
Generate Filter	Mitigates an attack with unique characteristics that can be defined using layer 3-4 access control list (ACL) filters. You can use this mitigation type to mitigate a DDoS attack if the results of the attack are not critical to your network operations.	"Mitigating Using ACL Filters" on page 199
Blackhole (null-routing using BGP)	Temporarily blackholes network traffic by redirecting it elsewhere in the network. This mitigation can also offramp network traffic at the peering edge of the network without redirecting it. This mitigation type uses a BGP announcement with a new nexthop to redirect the traffic to the filter device.	"Mitigating Using Blackhole Routing" on page 201

About the Mitigations Pages

The Mitigations Ongoing page (**Mitigation > Ongoing**) lists all of the active traffic mitigations in your network. The Mitigations Recent page (**Mitigation > Recent**) lists inactive traffic mitigations in your network.

For information about searching on this page, see "Searching for Mitigations" on the next page.

About the Mitigations Ongoing and Mitigations Recent pages

The Mitigations Ongoing and Mitigations Recent pages display the following information:

Information	Description
Search box	Use to search for mitigations by keyword. See "Searching for Mitigations" on the next page.
Wizard button	Click to search for mitigations using the Mitigation Search Wizard. See "Searching for Mitigations" on the next page.
Graph	A representation of the relevant traffic data (if available) that is involved in a mitigation. The graph is also a link to the Mitigation Status page.
Name	The unique name of a mitigation. The name is also a link to the Mitigation Status page.
Prefixes	The prefixes that are involved in the mitigation.
Duration	The amount of time (in days, hours, and minutes) that a mitigation was active. This column also displays the status of the mitigation as "Ongoing" or "Ended."
Start Time	The time and date when a mitigation was initiated. By default the "Ongoing" mitigations are listed first followed by the "Ended" mitigations. In each category, the mitigations are listed from the most recent to the least recent. When you click **Start Time**, the order of the list is reversed.
User	The user who initiated a mitigation.
Type	The mitigation type.
Annotations	The annotations (comments) that are applied to a mitigation. You can click the icon to apply an annotation to a mitigation. See "Adding Annotations to a Mitigation" on page 197.

Searching for Mitigations

Introduction

You can search for mitigations by using the **Search** box and the Mitigation Search Wizard:

You can search for mitigations on the Mitigations Ongoing page (**Mitigation > Ongoing**), the Mitigations Recent page (**Mitigation > Recent**), and the TMS Mitigations page (**Mitigation > Threat Management**):

See "About the Mitigations Pages" on the previous page and "About the TMS Mitigations page" on page 96.

About searching for mitigations on the mitigations pages

You can use the **Search** box to search on the mitigations pages. The following are some guidelines for using the **Search** box:

- You can enter search values with or without keywords.

  See "Acceptable search keywords and values for mitigations" below.
- Keywords allow you to search on a specific attribute.
- When you enter a keyword followed by a value, do not put a space between the colon and the value that you enter.
- A space between search values creates an AND statement.
- A comma between search values creates an OR statement.
- You can use quotation marks (") to match a phrase. For example, to search for a mitigation that has "mitigation stopped" in the annotation, you can type **`ann:"mitigation stopped"`**.
- A match occurs when a search value matches any part of a text string.

Acceptable search keywords and values for mitigations

The following table describes the keywords and values that you can use to search in the **Search** box on the mitigations pages:

Attribute to search by	Acceptable keywords and values	Examples
Mitigation name	▪ **`mitigation_name:`***`name`* ▪ **`name:`***`name`*	▪ mitigation_name:test_mit ▪ name:mitigation_test This search is case-insensitive, and Peakflow SP matches on partial mitigation names.
User name	▪ **`mitigation_user:`***`user name`* ▪ **`user:`***`user name`*	▪ mitigation_user:admin, user:John This search is case-sensitive, and Peakflow SP matches on the exact names of users who initiated mitigations.

Attribute to search by	Acceptable keywords and values	Examples
Mitigation type	■ **mitigation_type:***type* ■ **type:***type*	■ mitigation_type:TMS ■ type:blackhole This search is case-insensitive, and Peakflow SP matches on partial mitigation types. You do not need to use quotation marks (“) to use phrases in type searches.
Mitigation status	■ *mitigation status* ■ **sts:** *mitigation status* ■ **status:** *mitigation status*	■ ongoing ■ sts:recent ■ status:all You can type all, ongoing, recent, ended, stopped, done, or completed for mitigation status.
Annotation	■ **mitigation_annotation:** *annotation* ■ **ann:** *annotation* ■ **annotation:** *annotation* ■ **comment:** *annotation*	■ ann:Stop ■ annotation:Critical ■ comment:”TMS down”

About the search results

By default, the search returns the top 100 results in order of relevance. You can override the default setting for specific searches by using the Alert Search Wizard.

See “Using the Mitigation Search Wizard” on the next page.

Using the Search box

You can use the **Search** box to further refine the results of a previous search.

To search for mitigations from the **Search** box:

1. Navigate to one of the following pages:
   - Mitigations Ongoing page (**Mitigation > Ongoing**)
   - Mitigations Recent page (**Mitigation > Recent**)
   - TMS Mitigations page (**Mitigation > Threat Management**)
2. In the **Search** box, type *keywords*, *values*, or both.
3. Click **Search**.

Using the Mitigation Search Wizard

To search for mitigations with the Mitigation Search Wizard:

1. Navigate to one of the following pages:
   - Mitigations Ongoing page (**Mitigation > Ongoing**)
   - Mitigations Recent page (**Mitigation > Recent**)
   - TMS Mitigations page (**Mitigation > Threat Management**)
2. On the mitigation page, click **Wizard**.
3. In the Mitigation Search Wizard, configure the following settings:

Setting	Procedure
Status check boxes	(All Mitigations page only) Select the check boxes next to the mitigation statuses to include in the search.
Search Limit box	Type the maximum *number* of results to return.
Items per Page box	Type the maximum *number* of items to include per page.
IP Version check boxes	Select the IP version type (IPv4 and/or IPv6) that you want to include in the search.
Mitigation Type check boxes	(All Mitigations page only) Select the mitigation type by which to search. See "Mitigation types" on page 192.

   When you search by multiple attributes, Peakflow SP combines them using AND operators.
4. Click **Search**.
5. (Optional) If you do not click away from the page, then you can repeat these steps to add or change the search criteria.

Adding Annotations to a Mitigation

Introduction

You can add annotations (comments) to a mitigation to help you track the history of the actions that are taken on it. You can add annotations to any mitigation on a mitigation listing page and to TMS mitigations on the TMS Mitigation Status page.

Adding an annotation to a mitigation on a mitigation listing page

To add an annotation to a mitigation on a mitigation listing page:

1. Navigate to one of the following pages:
   - Mitigations Ongoing (**Mitigation > Ongoing**)
   - Mitigations Recent (**Mitigation > Recent**)
   - TMS Mitigations (**Mitigations > Threat Management**)
2. Click  (annotation) in the Annotations column for the mitigation that you want to annotate.
3. In the first Annotations window, click **Add Comment**.
4. In the second Annotations window, configure the following settings:

Setting	Procedure
box	Type your *annotation*.
Customer called, **Crippling attack**, and **Escalated** check boxes	(Optional) Select one or more of these check boxes to indicate why you added the annotation.

5. Click **Save**.

Adding an annotation to a TMS mitigation on the TMS Mitigation Status page

To add an annotation to a TMS mitigation on the TMS Mitigation Status page:

1. Navigate to one of the following pages:
   - Mitigations Ongoing (**Mitigation > Ongoing**)
   - Mitigations Recent (**Mitigation > Recent**)
   - TMS Mitigations (**Mitigations > Threat Management**)
2. Click the name link of the mitigation.
3. On the TMS Mitigation Status page, in the Summary pane, click **Add Comment**.

4. In the Add a Comment window, configure the following settings.

Setting	Procedure
box	Type your *annotation*.
Customer called, **Crippling attack**, and **Escalated** check boxes	(Optional) Select one or more of these check boxes to indicate why you added the annotation.

5. Click **Save**.

Mitigating Using ACL Filters

Introduction

You can use ACL (Access Control List) filters to mitigate a DDoS attack when the DDoS attack does not critically impede network operation. ACL filters specify who or what can access an object and which operations can be performed on the object. You can use ACL filters to filter traffic on the following routers:

- Alaxala
- Cisco
- Foundry
- Juniper

Mitigating using an ACL filter

To mitigate an attack using an ACL filter:

1. Navigate to the DoS alert page.
   See "Navigating to a DoS alert page" on page 66.
2. Click **Mitigate Alert**, and then click **Generate Filter**.
3. On the Generate Filter page, configure the following settings:

Setting	Procedure
Name box	Type a *name* if you want to change the existing name. The system automatically names the ACL filter by the alert number assigned (for example, alert-29). Some routers require that the name include a valid access-list number. The valid ranges of numbers for Cisco and Foundry routers are 100-199 and 2000-2699. Before you choose a number, check the configuration of the router to verify that the number is unique.
Vendor list	Select the vendor for which you want to generate an ACL or rate limiter.
Rate Limit (optional) box	Type a *rate limit* that is accepted by the router that you selected. The rate limit depends on the router, as follows: ■ Alaxala routers — The rate limit must fall between 1 kbps and 10 Gbps. ■ Cisco and Foundry routers — The rate limit must fall between a minimum of 8,000 bps and 10 Gbps. Because these routers only accept rates as multiples of 8,000, Peakflow SP rounds the rate limit to the next lowest multiple of 8,000. ■ Juniper routers — The rate limit must fall between 30,520 bps and 4.29 Gbps.

4. Click **Generate**.

5. After the ACL text appears on the Generate Filter page, copy and paste the text into your router configuration.

Mitigating using shun commands for Cisco PIX, ASA, or FWSM

Peakflow SP generates shun commands that allow you to mitigate using the Cisco PIX, ASA, and FWSM firewalls. The shun commands block traffic based on a specific source address.

To mitigate using a Cisco Pix/ASA/FWSM Shun or Cisco Firewall Shun:

1. Navigate to the DoS Alert <ID number> page (**Alerts > All Alerts> DoS alert ID link**).
2. On the alert page, on the Affected Routers tab, click **Details** for the device to filter.
3. On the DoS Alert Traffic Details page, select the **Filter** check boxes for the components that you want to block.

   Select only one item for each type (source address, destination address, source port, destination port, or protocol).
4. Click **Mitigate**.
5. In the DoS Alert Mitigation window, from the list, select **Generate Filter**, and then click **Mitigate**.
6. On the Generate Filter page, select **Cisco PIX Shun**, and then click **Generate**.
7. After the text appears on the Generate Filter page, copy and paste the command text into your router configuration.

Mitigating Using Blackhole Routing

Introduction

You can configure blackhole mitigations on the Blackholes page (**Mitigation > Blackhole**). Blackhole mitigations route traffic to specific IP prefixes that can be injected into the network using BGP with characteristics that inform routers how to treat traffic going toward a prefix.

You can also initiate a blackhole mitigation from a DoS Alert. See "Initiating a Mitigation from a DoS Alert" on page 114.

The following are the most common use cases of blackhole routing:

- null routing–where routers drop all traffic towards the injected prefix.
- offramp routing–where all traffic going towards the injected prefix is diverted to an alternate location.

Using blackhole routing, a network administrator can precisely define filters and observe their ability to protect your network from DoS and DDoS attacks.

You can use the Blackhole Mitigation Status page to view the status of and add annotations to blackhole mitigations.

See "About the Blackhole Mitigation Status Page" on page 205.

Mitigating an attack with a BGP blackhole

Blackhole mitigations allow you to drop or null-route all of the traffic that enters your network and is destined to one of your host addresses. This mitigation is typically used when the amount of attack traffic that enters the network overwhelms routes and devices.

Example

You assign a customer a /24 CIDR block, and a /32 host within that CIDR block is under attack. The amount of attack traffic that reaches the customer completely fills up their link to you. You create a blackhole mitigation that drops all traffic to that /32 host as it enters the network. This prevents the attack traffic to that host from reaching the customer link, allowing the remaining traffic to the customer's /24 network to still arrive.

This mitigation ends the DoS attack, but the customer is unable to offer the services that were associated with that /32 host until you remove the blackhole mitigation.

About the Blackholes page

Use the Blackholes page to temporarily blackhole network traffic or to offramp network traffic without redirecting it. You can use the Blackholes page to view, add, edit, and delete traffic blackhole filters. You also can start blackholes or stop ongoing blackholes.

If you create blackholes as part of the TMS mitigations or FlowSpec router mitigations, then they do not appear on the Blackholes page.

The Blackhole page displays the following information:

Icon	Description
	Hover over this icon to display the download options. The download options that appear vary depending on the page that you are on and can include the following: ■ PDF - Click to download the page in PDF format. ■ XML- Click to download the page in XML format. ■ CSV - Click to download a page in CSV format. The download can be a CSV text file or CSV zip archive file. When the download is a zip archive file, “zip archive” is appended to CSV. ■ Excel-XML - Click to download a page in Excel-XML format. The icon appears only if there are multiple options for downloading a

Adding and editing blackhole mitigations

To add or edit a blackhole mitigation:

1. Navigate to the Blackholes page (**Mitigation > Blackhole**).
2. Choose one of the following steps:
   - To add a blackhole mitigation, click **Add Blackhole**, and then click **IPv4** or **IPv6** for the IP version of the traffic that you want to mitigate.
   - To edit a blackhole mitigation, click its name link.
3. On the Add Blackhole page or the Edit Blackhole page, configure the blackhole settings.
   See “Settings for blackhole mitigations” below.
4. Click **Save**, and then commit your changes.

Settings for blackhole mitigations

Use the following table to configure the settings on the Add Blackhole page or the Edit Blackhole page:

Setting	Description
Name box	Type a *unique name* for the blackhole mitigation.
Source Alert ID box	Type the number of the DoS alert with which to associate this mitigation. This number is pre-populated if you created this mitigation from a DoS alert.
Select Managed Object button	Click and then select a managed object if you want users in an account group with access to the managed object to be able to configure this mitigation. This button does not appear when you edit a blackhole mitigation.
Internet Protocol Version options	Displays the IP version that was selected for this blackhole mitigation.

Icon	Description
	Hover over this icon to display the download options. The download options that appear vary depending on the page that you are on and can include the following: ■ PDF - Click to download the page in PDF format. ■ XML- Click to download the page in XML format. ■ CSV - Click to download a page in CSV format. The download can be a CSV text file or CSV zip archive file. When the download is a zip archive file, "zip archive" is appended to CSV. ■ Excel-XML - Click to download a page in Excel-XML format. The icon appears only if there are multiple options for downloading a page.
	Click to download a page in PDF format. This icon appears only when the PDF format is the only download option for a page.
	Click to download and email a page as a PDF.

Column	Description
	Select this icon to delete a user account.
Username	A user name as a link to the Edit Existing Account page.
Real Name	A user's full name.
Account Group	The account group to which a user belongs.
Capability Level	A user's capability level, which is either an administrator or a user.
Email	A user's email address.
Device	The SP appliance with which a user is associated. The SP appliance is either a specific appliance name or *global*, which associates a user with all appliances. For more information about associating a user with appliances, see "About user-appliance association" on page 24.
UI Menu	The UI menu that is assigned to a user. The UI menu determines what menu choices are available to a user.
Status	"Disabled" appears in this column for a user account that is disabled; otherwise, this column is blank.

Information Type	Description

Starting blackhole mitigations

To start a blackhole mitigation:

1. Navigate to the Blackholes page (**Mitigation > Blackhole**).
2. Choose one of the following steps:
   - Click **Start** in the row of the mitigation to start.
   - Select the check boxes next to the mitigations to start, and then click **Start**.

Stopping blackhole mitigations

To stop a blackhole mitigation:

1. Navigate to the Blackholes page (**Mitigation > Blackhole**).
2. Do one of the following:
   - Click **Stop** for the mitigation that you want to stop.
   - Select the check boxes for the mitigations that you want to stop, and then click **Stop**.

You can also stop a blackhole mitigation on the Blackhole Mitigation Status page.

See "About the Blackhole Mitigation Status Page" on the facing page.

Deleting blackhole mitigations

When you delete a blackhole mitigation, all of its associated mitigation events are deleted also.

To delete a blackhole mitigation:

1. Navigate to the Blackholes page (**Mitigation > Blackhole**).
2. Select the check boxes next to the mitigations to delete, and then click **Delete**.

About the Blackhole Mitigation Status Page

Introduction

The Blackhole Mitigation Status page allows you to view the status of and add annotations to blackhole mitigations.

Navigating to the Blackhole Mitigation Status page

To navigate to the Blackhole Mitigation Status page:

1. Choose one of the following steps:
   - Navigate to the Ongoing Mitigation Events page (**Mitigation > Ongoing**).
   - Navigate to the Recent Mitigation Events page (**Mitigation > Recent**).
2. Click the name link for the blackhole mitigation that you want to view.

About the blackhole Mitigation Summary section

The blackhole Mitigation Summary section contains the following information:

Information	Description
Name	The name of the alert that is associated with a mitigation.
Duration	The length of time that a mitigation ran and whether it ended or is ongoing.
Start Time	The date and time when a mitigation started.
Source Alert ID	The ID number of the alert that is associated with a mitigation.
Offramp Prefix	The CIDR prefix of the address to which the traffic is offramped.
Nexthop	The nexthop IP address to which Peakflow SP sends mitigated traffic.
Community	The mitigation communities.
Description	The description of a blackhole mitigation that you configured.
Done button	Click to navigate to the Blackholes page.
Stop button	Click to stop a mitigation.
Edit button	Click to edit a mitigation.

References

See the following sections for more information about this page:

- “Adding Annotations to a Mitigation” on page 197
- “Stopping blackhole mitigations” on the previous page
- “Adding and editing blackhole mitigations” on page 202

Chapter 10:
Traffic Reports

Introduction

This section describes how to use the various reports and configurations in Peakflow SP to monitor your network traffic.

In this section

This section contains the following topics:

Introduction to Traffic Reports

Introduction

You can use reports to monitor the traffic patterns and usage across your network. Traffic reports can assist you in managing your devices, routers, links, and interfaces.

For more information about how to use report controls, see "About the Traffic Report Pages" on page 211.

About the structure of traffic reports

The structure of traffic reports allows you to view traffic data based on different time periods and perspectives.

Managed services administrators and non-administrative users can view traffic reports from the following perspectives:

- Network
- Profiles

How Peakflow SP bins traffic data for reports

Peakflow SP stores all of its time-series data for traffic reports in a round-robin database. Initially, Peakflow SP represents all traffic data by five-minute samples. To save disk space over time, the samples are aggregated to report historical traffic. For example, the week, month, and year views are available in most traffic reports.

When you query long time periods, an empty graph might appear. An empty graph means that the queried object might not have been in place long enough for Peakflow SP to have binned traffic over that period. To avoid this issue, you can decrease the queried time period until data appears.

Data granularity for reports

Peakflow SP stores all of its time-series data for traffic reports in a round-robin database. The following table shows how Peakflow SP returns data based on the time period of a report and how long Peakflow SP stores the data:

Time period	Granularity	Maximum age of data
Today	5 minutes	14 days
Yesterday	5 minutes	14 days
2 Days Ago	5 minutes	14 days
1 Week	30 minutes	8 weeks
4 Weeks	120 minutes	6 months

Time period	Granularity	Maximum age of data
52 Weeks	24 hours	3 years
Other	varies	varies

Example: If you select **Today** for the time period of any default report, the report includes data for the previous 24 hours and returns samples with a five-minute granularity. If you select a start time of **10 days ago** and a stop time of **now**, Peakflow SP returns samples with a 30-minute granularity because the report covers more than two days but less than two weeks.

Correcting aggregation data

To change your sample from aggregated data to a five-minute sample:

1. From the **Period** list in the query window, select **Other**.
2. In the **Start** box, type a new query *`start time`*.
3. In the **End** box, type a new query *`stop time`*, and then click **Update**.

Change the query start and stop times so that the query spans fewer than two days GMT.

You can also run two separate queries to divide the data.

How time zones affect data granularity

Peakflow SP stores data in terms of Greenwich Mean Time (GMT). When Peakflow SP renders reports, users in different time zones have their data adjusted automatically. Some time zones do not correspond well to GMT over certain data reporting bins, which can lead to non-intuitive sample granularities. For example, when a query spans more than two days GMT, the system pushes the query results to the next aggregation level (one week).

Depending on the queried time period and time zone, Peakflow SP might aggregate the data as expected or it might aggregate it differently than expected. For example, Peakflow SP might return the data in half-hour bins instead of five minute bins.

Example: Central European Time (CET) users see that the queried data does not align with the specified time period, and the query allows more data than they want. For example, the time period is two days and the query returns three days worth of data.

Configuring traffic reports

To configure a traffic report:

1. Navigate to the Web UI page for the report that you want to configure.
2. Where applicable, configure the following settings:

Setting	Description
Period list	Select the time period for which you want data.
Units list	Select the unit in which to measure traffic. If you select **Bytes** or **Packets** from the list, then Peakflow SP replaces the report calculation types with “Total.” See “Report data calculation options” on page 213.

Setting	Description
Graph Type list	Select the type of graph in which you want to display the data.
Class list	From the **Class** list, select the set of data that you want to display. You can select **In**, **Out**, or **Total**. The **Class** list only appears when you select **Pie** or **Bar** from the **Graph Type** list and does not appear for all reports.
Display list	Select to display all in and out traffic or just off-net traffic. This setting appears on only a few reports.
Select <object> button	Click this to select the object for which you want to view data. This button does not appear on all of the reports.
Select All button	Click this to select all of the items in the following table. This button does not appear on all of the reports.
Clear All button	Click this to clear all of the items selected in the following table. This button does not appear on all of the reports.

3. Click **Update**.

For more information about updating graphs and how Peakflow SP displays report data, see "About the Traffic Report Pages" on the facing page.

About the Traffic Report Pages

Introduction

Peakflow SP pre-defined reports allow you to view data about Internet traffic that traverses your network. These reports are available in the **Reports** menu. For Peakflow SP pre-defined reports, you can:

- Obtain a brief report description by hovering the mouse pointer over the (information) icon.
- (For some reports) drill down to more detail by hovering the mouse pointer over the (drill down) icon. When you hover your mouse pointer over , Peakflow SP displays a list of related reports that you can click to drill down into more detailed data
- Download and email the report by using the icons that appear on the Arbor Smart Bar. See "About the Arbor Smart Bar " on page 13.

About report graphs

Most of the pre-defined report pages in Peakflow SP provide configuration options for displaying the data in graphs. Most reports offer the following types of graphs:

- stacked
- bar
- pie

For stacked graphs, note the following information:

- The data above the center line represents outgoing traffic, and the data below the center line represents incoming traffic.
- Peakflow SP converts all data to the configured time zone that is selected in your user profile.
- The Total row, the last row of the data table, displays the total traffic of the target object. These totals are not the sum of the values of each column.

To remove the gray background that appears in some line graphs, remove the Total traffic from the graph by deselecting the last row of the data table.

To drill down to view graph details:

- Click and drag across the timeframe of graph data that you want to view in detail. A new graph loads with a finer level of detail for the selected timeframe.

To change a graph to a different type:

- From the **Graph Type** list in the report, select the type of graph that you want to view, and then click **Update**.

About report data tables

Most reports display a data table. You can use the data table to perform the following tasks:

- Select which rows of data to include in the graph and the "Sum of selected items" row. By default, Peakflow SP selects the top five items to include in the graph.
- View the top 100 items that match a report's selection criteria
- Re-sort the data in a table by clicking any of the underlined column headings

The “Sum of selected items” row in tables displays the data sum of all of the selected rows in a table, except for the Total row.

To change the table data that is displayed in the graph and the “Sum of selected items” row, choose one of the following steps, and then click **Update**:

- To add data to a graph and “Sum of selected items,” select the check box for a data row.
- To remove data from a graph and “Sum of selected items,” clear the check box for a data row.
- To add all data rows to a graph and “Sum of selected items," click **Select All**.
- To remove all data rows that you added to a graph and “Sum of selected items,” click **Clear All**. When you use the **Clear All** function, Peakflow SP clears the rows that you added and reselects the top five items.

By default, reports show the top 100 items that match a report’s selection criteria. You can modify this default setting for most reports in the CLI. See “Overriding the Default Number of Items Listed in a Report Data Table” in the *Peakflow SP and Threat Management System (TMS) Advanced Configuration Guide*.

About traffic types in reports

Report graphs commonly display the following types of traffic:

Traffic type	Description
In	The total traffic entering the network through a selected object.
Out	The total traffic leaving the network through a selected object.
Backbone	The total traffic that passes through the backbone and does not leave the network.
Multicast	The total multicast traffic that enters your network. Multicast traffic is sent from one source address to one destination address that many people share, called a multicast address. It allows traffic to be sent from one host to many hosts simultaneously. Multicast traffic potentially uses less bandwidth.
In %	The percentage of incoming traffic that this traffic consumed.
Out %	The percentage of outgoing traffic that this traffic consumed.

Report data calculation options

Most report tables in Peakflow SP allow you to choose from the following options for calculating the data to display (also known as CAMP values):

Calculation type	Description
Current	Displays the values of the most recent five-minute sample. This calculation is available for the “Today” time period only.
Average	Displays the average of all samples for the selected time period. If you select Bytes or Packets from a **Units** list, the “average” calculation type is selected and all other types are disabled.
Max	Displays the maximum of all samples for the selected time period.
PCT95	Displays the 95th percentile of all samples for the selected time period.

The available calculation options appear in the lower-right corner below the data table.

To change the calculation option for the displayed table data click **Current**, **Average**, **Max**, or **PCT95**.

About Summary Reports

Introduction

Summary reports (**Traffic > Summary >** report) display network-wide traffic data. This data includes several breakdowns of overall network traffic, such as per protocol and per port.

User access

Only managed services administrators and non-administrative users can view these reports.

About the Summary reports

The following table describes the Summary reports and how to navigate to them:

Report	Navigation Path	Description
Applications Summary	Traffic > Summary > Applications	Displays all of the network's incoming and outgoing traffic, organized by the top 100 applications observed.
TCP Applications Summary	Traffic > Summary > Ports > TCP	Displays all of the network's incoming and outgoing traffic for the top TCP applications observed, organized by application port. The report lists the application names and ports in a traffic table.
UDP Applications Summary	Traffic > Summary > Ports > UDP	Displays all of the network's incoming and outgoing traffic for the top UDP applications observed, organized by application port. The report lists the application names and ports in a traffic table.
Protocols Summary	Traffic > Summary > Protocols	Displays all of the network's incoming and outgoing traffic, organized by IP protocol. The report maps protocols to names (when known) or displays the protocol number.
Top Talker Summary	Traffic > Summary > Top Talkers	Displays the top prefixes within the network that consume the most bandwidth. For each prefix, the report displays the peak traffic rate that is consumed by that prefix. Click the Details button to see more information about the prefix. The Details button is disabled if your service provider has not given you the capability to use it.

About Profile Reports

Introduction

You can use the Profile reports (**Traffic > Profiles**) to monitor your profile managed object traffic across your network.

User access

Only managed services administrators and non-administrative users can view these reports.

About the Profile reports

The following table describes the Profile reports and how to navigate to them:

Profile Report	Navigation Path	Description
Customer Compare	Traffic > Profiles > Compare Profiles	Displays the data that compares the top profiles from a selected profile group. Because configured profiles can overlap, the graph might display the network's total in and out traffic as less than the combined managed object traffic. This is normal and is not an indication of data inaccuracy or double counting.
Customer Summary	Traffic > Profiles > Profile Detail	Displays detailed traffic information for a selected profile.
Customer All Applications	Traffic > Profiles > Applications	Displays the traffic into and out of a profile for the top 100 applications observed.
Customer TCP Applications	Traffic > Profiles > Ports > TCP	Displays the traffic into and out of a profile for the top TCP applications observed, organized by application port. The report lists the application names and ports in a traffic table.
Customer UDP Applications	Traffic > Profiles > Ports > UDP	Displays the traffic into and out of a profile for the top UDP applications observed, organized by application port. The report lists the application names and ports in a traffic table.

Profile Report	Navigation Path	Description
Customer Protocols	Traffic > Profiles > Protocols	Displays the traffic into and out of a profile, organized by IP protocol. The report maps protocols to names (when known) or displays the protocol number.
Profile Top Talkers	Traffic > Profiles > Top Talkers	Displays the top prefixes that consume the most bandwidth for a profile. For each prefix, the report displays the peak traffic rate that is consumed by that prefix. Click the Details button to see more information about the prefix. The Details button is disabled if your service provider has not given you the capability to use it. The report displays subnets and individual hosts that belong to the selected profile and that send a significant proportion of the profile's traffic. The default address 0.0.0.0/0 collects all of the traffic that was not vital enough to list as a more specific host or subnet.

Note about the ICMP Application report

The Name column of the data table in the ICMP Application report displays the name of the ICMP type, followed by the ICMP code in parentheses. For example, ICMP traffic that had the type set to 0 and the code set to 11 would be displayed as Echo Reply (code=11).

About Cisco Standard DSCP values

The following table describes Cisco's suggestions for standard DSCP values:

DSCP value	Purpose
0	Best effort
26	Voice control (SIP, H.323)
46	Voice data (RTP, RTSP)
18	Better effort data
10	Streaming video
48	Network-layer protocol (OSPF, RIP, EIGRP)

Non-Country Entries in a List of Countries

Introduction

With the Customer Countries report, the list of countries can include entries that are not countries.

For more information about the Customer Countries report, see “Additional Pre-defined Customer Reports” on the previous page.

About non-country entries in a list of countries

Non-country entries that appear in a list of countries include the following:

- Anonymous Proxy (A1)

  This entry represents IP addresses that are part of anonymous proxies. An attacker can use an anonymous proxy to hide its IP address or its geographical location.

- Satellite Provider (A2)

  This entry represents ISPs that use satellites to provide Internet access to various countries. For these ISPs, the county of the end user is often unknown.

- Europe (EU) or Asia-Pacific (AP)

  These entries appear when the end user's location within these regions is unclear. For example, a corporate proxy that is located in Paris, France, could be listed as Europe if the actual end users connect from different parts of Europe. Because the traffic originates from various places in Europe, “Europe” is used for the country and not France.

  When you block Europe or Asia-Pacific traffic, you are not blocking all of the traffic that comes from Europe or Asia-Pacific. Instead, you are blocking only the traffic that could not be clearly defined as coming from a specific country within these regions. If you want to block all of the traffic from Europe or Asia-Pacific, you must block the traffic for each of the countries within these regions.

Chapter 11:
Monitoring the System Status

Introduction

This section describes how to monitor the Peakflow SP system status.

In this section

This section contains the following topics:

About the Security Status Page

Introduction

The Security Status page (**Status** menu) allows you to view a summary of alerts, ongoing mitigations, and the general health of your Peakflow SP system.

User access

Only managed services administrators and non-administrative users can view this page.

About the Network Summary tab

The Network Summary tab includes a graph that displays network activity over the last 24 hours. Time is graphed on the X-axis and in and out traffic is graphed on the Y-axis.

About the Alerts tab

The Alerts tab displays the number of ongoing alerts, recent alerts, and alerts in the last 24 hours, organized by importance (severity) level. You can click the number links on this tab to navigate to the corresponding Alerts Ongoing page or Alerts Recent page. The page only includes alerts for the importance level that the link was in.

About the Ongoing Alerts and Ongoing Mitigations tabs

The **Ongoing Alerts** tab and the **Ongoing Mitigations** tab contain the following information:

Tab	Description
Ongoing Alerts	Displays the five most severe ongoing alerts that Peakflow SP has detected and are still active. The information on this tab is the same as the information that is displayed on the Alerts Ongoing page. See “About the Alert Listing Pages” on page 61.
Ongoing Mitigations	Displays the five most recent mitigations. The information on this tab is the same as the information that is displayed on the Mitigations Ongoing page (**Mitigation > Ongoing**). See “About the Mitigations Pages” on page 193.

Glossary

a

AAA (Authentication, Authorization, & Accounting) — This is an acronym used to describe the process of authorizing access to a system, authenticating the identity of users, and logging their behaviors.

ACL (Access Control List) — A list composed of rules and filters stored in a router to allow, deny, or otherwise regulate network traffic based on network parameters such as IP addresses, protocol types, and port numbers.

active route — A network route installed in a routing table.

address — A coded representation that uniquely identifies a particular network identity.

AES (Advanced Encryption Standard) — A commonly used encryption block cipher adopted as the standard of the U.S. government.

AIF (ATLAS Intelligence Feed) — Real-time threat information from ATLAS. Peakflow SP regularly downloads this information and uses it to detect and block emerging botnet attacks and application-layer attacks.

anomaly — An event or condition in the network that is identified as an abnormality when compared to a predefined illegal traffic pattern.

anonymous statistic sharing — A service whereby service providers and enterprise businesses share anonymized statistics on ongoing attacks in order to provide a Internet-wide view of ongoing attacks.

API (Application Programming Interface) — A well-defined set of function calls providing high-level controls for underlying services.

appliance — A Peakflow server that gathers network statistics from adjacent routers via either packet capture or flow and performs first-order traffic analysis. Anomalous activities are compressed into alert messages that are periodically sent to the listening leader.

ARP (Address Resolution Protocol) — A protocol for mapping an IP address to a physical machine address.

AS (Autonomous System) — A collection of IP networks and routers under the control of one entity and assigned a single ASN for purposes of BGP routing.

ASCII (American Standard Code for Information Interchange — A coded representation for standard alphabetic, numeric, and punctuation characters, also referred to as “plain text.”

ASN (Autonomous System Number) — A unique number assigned to an autonomous system for purposes of BGP routing.

AS Path (Autonomous System Path) — The ASNs that comprise a packet's path through the internet using BGP.

ATF (Active Threat Feed) — An Arbor maintained database of security threats and signatures that automatically updates each minute to include the most recent Internet-wide information and sends it to your Peakflow deployment.

ATLAS (Active Threat Level Analysis System) — A globally scoped threat analysis network that analyzes data from darknets and the Internet's core backbone to provide information to participating customers about malware, exploits, phishing, and botnets.

authentication — An identity verification process.

b

backbone router — An OSPF router with all operational interfaces within 0.0.0.0.

baseline — A description of typical traffic patterns over a period of time. Baselines are generated by reducing collections of fine-grained profiles into a more monolithic data representation that includes a chronological component.

BGP (Border Gateway Protocol) — The core routing protocol of the Internet.

binning — Grouping data into chunks or "bins" usually defined by time periods, for example traffic for the last 24 hours.

blackhole routing — A technique to route traffic to null interfaces that can never forward the traffic.

bogon — An IP packet that claims to originate from "dark" IP space.

border router — A router at the border of an AS or network.

bps — Bits per second.

c

CA (Certificate Authority) — A third party which issues digital certificates for use by other parties. CAs are characteristic of many public key infrastructure (PKI) schemes.

CAR (Committed Access Rate) — A tool for managing bandwidth that provides the same control as ACL with the additional property that traffic can be regulated based on bandwidth usage rates in bits per second.

CIDR (Classless Inter-Domain Routing) — Method for classifying and grouping Internet addresses.

CIDR Group — CIDR addresses grouped together to share common manged object configuration. The equivalent of Peakflow DoS "detection groups."

cflowd — Developed to collect and analyze the information available from NetFlow. It allows the user to store the information and enables several views of the data. It produces port matrices, AS matrices, network matrices, and pure flow structures.

challenge packets — Information sent by a TMS model to an unknown host in response to a request from the unknown host. The unknown host must provide a valid response to the challenge packets. If it does not, the TMS model refuses the request and adds the unknown host to the blacklist. Several TMS countermeasures use challenge packets to authenticate unknown hosts.

CLI (Command Line Interface) — A user interface that uses a command line, such as a terminal or console (as opposed to a graphical user interface).

client — The component of client/server computing that uses a service offered by a server.

Collector — An appliance that gathers network information from adjacent routers through flow and performs first-order traffic analysis. Anomalous events are compressed into event messages that are then sent to the listening Controller.

commit — The process of saving a configuration change so that the changes take affect on the Peakflow SP system.

customer — A managed object that defines traffic for a business or organization who purchases internet service from an Internet service provider. Note, this type of managed object should be used to define most managed services clients.

customer edge router — A router within a customer's network connected to an ISP's customer peering edge.

d

Dark IP — Regions of the IP address space that are reserved or known to be unused.

DDoS (Distributed Denial of Service) — An interruption of network availability typically caused by many, distributed malicious sources.

designated router — The router designated by other routers (via the OSPF protocol) as the sender of link state advertisements.

DHCP (Dynamic Host Configuration Protocol) — A protocol used to distribute IP addresses to host machines, which has a list of available addresses.

DNS (Domain Name System) — A system that translates numeric IP addresses into meaningful, human-consumable names and vice-versa.

DoS (Denial of Service) — An interruption of network availability typically caused by malicious sources.

DoS alert — A notification indicating an event or condition in the network that is identified as a statistical abnormality when compared to typical traffic patterns gleaned from previously collected profiles and baselines or that matches a predefined illegal traffic pattern.

e

encryption — The process by which plain text is scrambled in such a way as to hide its content.

ESP (Encapsulating Security Payload) — An IPSec protocol for establishing secure tunnels. See "IPSec (Internet Protocol Security)" on the facing page.

Ethernet — A series of technologies used for communication on local area networks.

exploit — Tools intended to take advantage of security holes or inherent flaws in the design of network applications, devices, or infrastructures.

f

failover — Configuring two appliances so that if one appliance fails, the second appliance takes over the duties of the first, ensuring continued service.

fate sharing — Putting a mitigation out of service when a part of the mitigation's deployment fails or becomes unreachable. Fate sharing can occur when a dependent interface loses link, a nexthop becomes unreachable, a BGP peer is down, a GRE tunnel is down, one or more TMS appliances or TMS clusters are out of service, or the leader appliance becomes unreachable, . For example, if nexthop fate sharing is configured for a TMS appliance and the nexthop used by a mitigation becomes unreachable, then the mitigation is put out of service.

FCAP — A fingerprint expression language that describes and matches traffic information.

Fibre Channel — Gigabit-speed network technology primarily used for storage networking.

firewall — A security measure that monitors and controls the types of packets allowed in and out of a network, based on a set of configured rules and filters.

flow — Flow is a characterization of the network traffic. It defines the traffic that is seen. It provides Peakflow SP with information from layers 1, 3, and 4 for the traffic that traverses a network.

FlowSpec — A BGP based IETF standard for exchanging flexible firewall and ACL rules implemented by Juniper routers utilizing JunOS 7.3 or later.

fps — Traffic flows per second (NetFlow, ArborFlow, SFlow, etc.).

FQDN (Fully Qualified Domain Name) — A complete domain name, including both the registered domain name and any preceding node information.

— A TCP/IP protocol for transferring files across a network.

g

GMT (Greenwich Mean Time) — A deprecated world time standard, replaced by UTC.

GRE (Generic Routing Encapsulation) — A tunneling protocol commonly used to build VPNs.

h

host — A networked computer (client or server); in contrast to a router or switch.

HTTP (HyperText Transfer Protocol) — A protocol used to transfer or convey information on the World Wide Web. Its original purpose was to provide a way to publish and retrieve HTML pages.

HTTPS (HyperText Transfer Protocol over SSL) — The combination of a normal HTTP interaction over an encrypted Secure Sockets Layer (SSL) or Transport Layer Security (TLS) transport mechanism.

I

IANA (Internet Assigned Numbers Authority) — An entity that oversees global IP address allocation, DNS root zone management, and other Internet protocol assignments. It is operated by ICANN.

ICMP (Internet Control Message Protocol) — An IP protocol that delivers error and control messages between TCP/IP enabled network devices, for example, ping packets.

IETF (Internet Engineering Task Force) — An Internet standards organization that develops draft documents and RFC documents defining protocols for the Internet.

IGMP (Internet Group Management Protocol) — A communications protocol used to manage the membership of Internet Protocol multicast groups.

intelligent filtering — A feature that adds the ability to work with an integrated filtering device to automatically filter traffic.

IMAP (Internet Message Access Protocol) — An application layer Internet protocol that allows a local client to access email on a remote server. (Also known as Internet Mail Access Protocol, Interactive Mail Access Protocol, and Interim Mail Access Protocol.)

interface — An interconnection between routers, switches, or hosts.

IP (Internet Protocol) — A connectionless network layer protocol used for packet delivery between hosts and devices on a TCP/IP network.

IP Address — A unique identifier for a host or device on a TCP/IP network.

IPS (Intrusion Prevention System) — A computer security device that exercises access control to protect computers from exploitation.

IPSec (Internet Protocol Security) — A suite of protocols for securing Internet Protocol (IP) communications by authenticating and/or encrypting each IP packet in a data stream.

ISP (Internet Service Provider) — A business or organization that provides to consumers access to the Internet and related services.

L

LAN (Local Area Network) — A typically small network that is confined to a small geographic space.

leader — A designated Peakflow appliance that accepts alert messages from one or more normal devices and performs second-order traffic analysis in order to identify and visualize potential attacks. (These were referred to as "Controllers" in previous Peakflow products.)

m

MAC (Media Access Control) Address — A unique hardware number associated with a networking device.

managed object — User-defined network objects used to classify logical portions of your network or network traffic. Managed objects can be customers, peers, profiles, VPNs, or VPN sites.

MD5 (Message Digest algorithm 5) — A widely used cryptographic hash function.

MDI (Media Dependent Interface) — An Ethernet port connection that allows network hubs or switches to connect to other hubs or switches without a null-modem, or Ethernet crossover cable.

MIB (Management Information Base) — A database used by the SNMP protocol to manage devices in a network. Your SNMP polling device uses this to understand Peakflow SNMP traps.

MPLS label — An identifying string for packets using the MPLS protocol.

mitigation — The process of using recommendations from Peakflow SP to apply policies to your network to reduce the affects of a worm or DoS attack.

mitigation device — A device that filters network traffic passing through it based upon a ruleset provided by Peakflow SP. This can be either a dedicated network device (Peakflow TMS appliance or Flowspec capable router) or a Peakflow SP appliance with software mitigation enabled.

MPLS (Multiprotocol Label Switching) — A packet-switching protocol developed by the Internet Engineering Task Force (IETF) initially to improve switching speeds, but other benefits are now seen as being more important.

MS (Managed Services) — A Peakflow SP appliance that has the ability to provide a Web UI to allow customers a special, restricted access to the Peakflow SP system.

MTU (Maximum Transmission Unit) — The size (in bytes) of the largest packet that a given layer of a communications protocol can efficiently forward.

multicast — Protocols that address multiple IP addresses with a single packet (as opposed to unicast and broadcast protocols).

n

NAT (Network Address Translation) — Rewriting the source and destination addresses of IP packets as they pass through a router or firewall.

NetFlow — A technology developed by Cisco Systems, Inc. that allows routers and other network devices to periodically export information about current network conditions and traffic volumes.

netmask — A dotted quad notation number used by routers determine which part of the address is the network address and which part is the host address.

network object — Network objects are portions of your network or network traffic and include both managed objects (customers, peers, profiles, VPNs, or VPN sites) and physical network objects (routers and interfaces).

NIC (Network Interface Card) — A hardware component that maintains a network interface connection.

NTP (Network Time Protocol) — A protocol that is used to synchronize clock times in a network of computers.

o

OC-3 — A fiber optic network line with transmission speeds of up to 155.52 Mbit/s.

OC-12 — A fiber optic network line with transmission speeds of up to 622.08 Mbit/s.

offnet — Traffic that leaves the network through a BGP boundary and is not destined for a configured customer entity.

p

packet — A unit of data transmitted across the network that includes control information along with actual content.

password — A secret code used to gain access to a computer system.

PCC (Packet Capture Collector) — Packet capture is a method of passively monitoring network traffic to create flow information. The packet capture mode on a Peakflow appliance can be used in cases where flow from routers is unavailable or unwanted.

PE (Provider Edge) Router — A router in a service provider's network that is connected to a customer edge router.

peer — A managed object that describes other networks that are peering with yours.

peer to peer — (Sometimes abbreviated P2P) a computer network that relies primarily on the computing power of the clients in the network rather than concentrating it in a relatively low number of servers. P2P networks are typically used for connecting nodes via largely ad hoc connections.

pps — Packets per second.

ping — An ICMP request to determine if a host is responsive.

POP (Post Office Protocol) — A TCP/IP email protocol for retrieving messages from a remote server.

PoP (Point of Presence) — A physical connection between telecommunications networks.

port — A field in TCP and UDP protocol, packet headers that corresponds to an application level service (for example TCP port 80 corresponds to HTTP).

profile — A managed object that defines an arbitrary subset of network traffic that does not fit any of the other managed object types.

protocol — A well-defined language used by networking entities to communicate with one another.

q

QoS (Quality of Service) — A method of providing different priority to different traffic, or guaranteeing a certain level of performance to a data flow for a particular traffic type.

r

RADIUS (Remote Authentication Dial In User Service) — A client/server protocol that enables remote access servers to communicate with a central server to authenticate dial-in users and authorize their access to the requested system or service.

RDN (Registered Domain Name) — A domain name as registered, without any preceding node information (for example, “arbor.net” instead of www.arbor.net).

refinement — The process of continually gathering information about anomalous activity seen.

remediation — The process of minimizing attack damage by taking the recommendations from Peakflow SP and applying reasonable changes to the network.

remote BGP routeviews — External route servers maintained by Arbor Networks which provide information on route availability with remote ASNs.

report — An informational page presenting data about a traffic type or event.

RFC (Request For Comments) — An IETF document that defines a protocol or other standard for Internet communications.

route — A path a packet takes through a network.

route target — A VPN identifier. A VPN might require more than one route target.

router — A device that connects one network to another. Packets are forwarded from one router to another until they reach their ultimate destination.

s

scoping — The container managed object within which a managed services customer's traffic view is restricted.

secret key — A secret shared only between a sender and receiver of data.

SFlow — A standard similar to NetFlow which describes a mechanism to capture traffic data in switched or routed networks.

shrew attack — A DoS attack that exploits a weakness in TCP's retransmission timeout mechanism to disrupt TCP connections.

skins — Sets of UI parameters, including menus ,used to facilitate different Peakflow SP workflows.

SMTP - (Simple Mail Transfer Protocol) — The de facto standard protocol for email transmissions across the Internet.

smurf attack — A DDoS attack that exploits misconfigured network devices to broadcast large numbers of ICMP packets to all the computer hosts on a network.

SNMP (Simple Network Management Protocol) — A standard protocol that allows routers and other network devices to export information about their routing tables and other state information.

spoofing — A situation in which one person or program successfully masquerades as another by falsifying data (usually the IP address) and thereby gains an illegitimate advantage.

SSH (Secure Shell) — A command line interface and protocol for securely getting access to a remote computer. SSH is also known as Secure Socket Shell.

SSL (Secure Sockets Layer) — A protocol for secure communications on the Internet for such things as web browsing, email, instant messaging and other data transfers.

t

TACACS+ (Terminal Access Controller Access Control System +) — An authentication protocol common to UNIX networks that allows a remote access server to forward a user's logon password to an authentication server to determine whether that user is allowed to access a given system.

target — A victim host or network of a worm or other malicious denial of service (DoS) attacks.

TCP (Transmission Control Protocol) — A connection-based, transport protocol that provides reliable delivery of packets across the Internet.

TCP/IP — A suite of protocols that controls the delivery of messages across the Internet.

Telnet — A TCP protocol used primarily for unencrypted CLI communications (usually deprecated and replaced by SSH).

TMS (Threat Management System) — A Peakflow SP appliance designed for intelligent traffic filtering and DNS monitoring in conjunction with a Peakflow SP deployment.

tunnel — A method of communication where one protocol is encapsulated within another.

u

UDP (User Datagram Protocol) — An unreliable, connectionless, communication protocol.

UNC (Universal Naming Convention) — A standard which originated from the UNIX for identifying servers, printers, and other resources in a network.

uptime — The time elapsed since a given host or server was last rebooted.

URI (Uniform Resource Identifier) — A protocol, login, host, port, path, etc. in a standard format used to reference a network resource, (for example http://arbor.net/).

URL (Uniform Resource Locator) — Usually a synonym for URI.

UTC (Universal Time Coordinated) — The time zone at zero degrees longitude which replaced GMT as the world time standard.

V

VLAN (Virtual Local Area Network) — Hosts connected in an infrastructure that simulates a local area network, when the hosts are remotely located, or to segment a physical local network into smaller, virtual pieces.

VoIP (Voice over Internet Protocol) — Routing voice communications (such as phone calls) through an IP network.

VPN (Virtual Private Network) — A private communications network often used within a company, or by several companies or organizations, to communicate confidentially over a public network using encrypted tunnels.

vulnerability — A security weakness that could potentially be exploited.

W

WAN (Wide Area Network) — A computer network that covers a broad area. (Also, Wireless Area Network meaning a wireless network.)

WEP (Wired Equivalent Privacy) — A security scheme for wireless networks intended to provide comparable confidentiality to a traditional wired network (in particular it does not protect users of the network from each other).

worm — A self propagating program, usually used to spread a malicious payload across networked computers.

X

XML (eXtensible Markup Language) — A metalanguage written in Standard Generalized Markup Language (SGML) that allows one to design a markup language for easy interchange of documents on the World Wide Web.

Index

a

b

c

d

f

h

i

l

m

p

r

S

t

u

v

w

z

Software License Agreement

ARBOR NETWORKS, INC., IF YOUR PRINCIPAL PLACE OF BUSINESS IS IN THE UNITED STATES, OR ARBOR NETWORKS UK LTD., IF YOUR PRINCIPAL PLACE OF BUSINESS IS OUTSIDE OF THE UNITED STATES ("ARBOR") LICENSES THE PRODUCT AND/OR USE OF ARBOR'S CLOUD SERVICE ("CLOUD SERVICE") AND DOCUMENTATION (TOGETHER, THE "SOFTWARE") TO YOU ("YOU" OR "YOUR") PROVIDED YOU ACCEPT ALL OF THE TERMS OF THIS LICENSE AND CLOUD SERVICE AGREEMENT (the "AGREEMENT"). IF YOU'VE PURCHASED THE CLOUD SERVICE, YOU ALSO AGREE TO THE ADDITIONAL TERMS AND CONDITIONS LOCATED AT www.arbornetworks.com/cloud-suppterms. BY SIGNING THE ATTACHED FORM, OPENING THIS PACKAGE, BREAKING THE SEAL, CONNECTING PRODUCT TO YOUR NETWORK, OR ACCESSING THE CLOUD SERVICE, YOU AGREE TO THE TERMS AND CONDITIONS OF THIS AGREEMENT. IF YOU DO NOT AGREE TO THESE TERMS AND CONDITIONS, RETURN THE UNUSED PRODUCT WITHIN TEN (10) DAYS OF RECEIPT AND, WHERE APPLICABLE, YOU'LL BE DISCONNECTED FROM THE CLOUD SERVICE FOR A REFUND OF FEES PAID.

1. License to Use. Arbor grants You a limited, revocable non-exclusive, non-transferable license (the "License") to: a) use Arbor's software in machine-readable form that is shipped to You and/or identified on the attached form ("Form") and accompanying documentation (collectively "Product") on the machines on which the software has been installed or authorized by Arbor; and/or b) access and use the Cloud Services as described herein. The term of the license shall be as stated on the Form. Your affiliate (s), purchasing agents, and outsourcing vendors ("Affiliates") may on your behalf purchase or use Product and/or Cloud Services hereunder so long as each is bound to terms as in this Agreement and You indemnify Arbor for their breach of this Agreement. Any future trial or purchase of Product and services and future trials or purchases of Cloud Services is governed exclusively by this Agreement and may be effected by You or Your Affiliates providing a purchase order or trial request. Trial term licenses for Cloud Services shall be as stated on the Form. Trial term licenses for Product shall be for the longer of thirty (30) days from date of Product's delivery to You or as stated on the Form supplied by Arbor. Any feed, release, revision or enhancement to the Software that Arbor may furnish to You becomes a part of Product or Cloud Service and is governed by this Agreement. Specifically for Product, if You have not purchased a license by the end of a Product trial term or You breach this Agreement, You agree to return Product and any machine provided by Arbor to Arbor in its original condition less normal wear and tear in original packaging or equivalent and in accordance with Arbor's RMA process within 10 days. You agree to pay for any damage to Product occurring prior to receipt by Arbor. If You purchase a license to Product, this Agreement will control that purchase and title to machines (where applicable) provided hereunder vests in You.

2. Proprietary Rights and Restrictions. Arbor and/or its licensors and outsourcing vendors (together, "Vendors") retain all right, title, and interest in the Software and in all copies thereof, and no title to the Software or any intellectual property or other rights therein, are transferred to You other than as specified herein. No right, title or interest to any trademarks, service marks or trade names of Arbor or its Vendors is granted by this Agreement. Software is copyrighted and contains proprietary information and trade secrets belonging to Arbor and/or its Vendors. You will only use Software for Your own internal business purposes. You may not make copies of the Software, other than a single copy in machine-readable format for back-up or archival purposes. You may make copies of the associated documentation for Your internal use only. You shall ensure that all proprietary rights notices on Software are reproduced and applied to any copies. Licenses are limited to use in accordance with the "Description" on the Form and user documentation. You agree not to cause or permit the reverse engineering or decompilation of the Software or to derive source code therefrom. You may not create derivative works based upon all or part of Software. You may not transfer, lend, lease, assign, sublicense, and/or make available through timesharing, Software, in whole or in part. If you are purchasing spare Product, You're only licensed to use such spare during such time as another Product is removed from service for repair.

3. Confidentiality. When disclosing information under this Agreement, the disclosing party will be the "Disclosing Party" and the receiving party will be the "Receiving Party." The term "Confidential Information" includes: (a) a party's technical, financial, commercial or other proprietary information including without limitation product roadmaps, pricing, software code and documentation, Software, techniques or systems and (b) information or data that is confidential and proprietary to a third party and is in the possession or control of a party. The Receiving Party will not disclose any of the Disclosing Party's Confidential Information to any third party except to the extent such disclosure is necessary for performance of the Agreement or it can be documented that any such Confidential Information is in the public domain and generally available to the general public without any restriction or license, or is required to be disclosed by any authority having jurisdiction so long as Disclosing Party is provided advance notice of such disclosure by the Receiving Party. Each party's respective Confidential Information shall remain its own property. Notwithstanding the foregoing, Arbor may use anonymized data from the Product or Cloud Service for its business purposes provided that Arbor shall not identify You to any third party as the source of such data.

4. Product Warranty, Indemnification. Arbor warrants, for sixty (60) days from shipment, that Product will perform in compliance with user manuals accompanying Product. If, within sixty (60) days of shipment, You report to Arbor that Product is not performing as described above, and Arbor is unable to correct it within sixty (60) days of the date You report it, You may return the non-performing Product at Arbor's expense, and Arbor will refund amounts paid for such Product. The foregoing is Your sole and exclusive remedy. Arbor agrees to defend You from and against any third party claim or action based on any alleged infringement of any U.S. patent or copyright arising from use of the Product or Cloud Service according to the terms and conditions of this Agreement ("Claim"), and Arbor agrees to indemnify You from damages awarded against You in any such Claim or settlement thereof, provided that (i) Arbor is promptly notified in writing of such Claim, (ii) You grant Arbor sole control of the defense and any related settlement negotiations, and (iii) You cooperate with Arbor in defense of such Claim. Notwithstanding the foregoing, Arbor shall have no liability to You if the infringement results from (a) use of the Product or Cloud Service in combination with software not provided by Arbor; (b) modifications to the Product or Cloud Service not made by Arbor; (c) use of the Product or Cloud Service other than in accordance with the Documentation or this Agreement; or (d) failure to use an updated, non-infringing version of the applicable Product or Cloud Service. The foregoing states the entire liability of Arbor with respect to infringement.

5. Limitations. EXCEPT AS OTHERWISE PROVIDED HEREIN, ARBOR AND ITS THIRD PARTY VENDORS MAKE NO OTHER WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. ARBOR'S AGGREGATE LIABILITY FOR ANY AND ALL CLAIMS ARISING OUT OF OR IN CONNECTION WITH THIS AGREEMENT, THE PERFORMANCE OF PRODUCT PROVIDED HEREUNDER, AND/OR ARBOR'S PERFORMANCE OF SERVICES (INCLUDING, WITHOUT LIMITATION, THE CLOUD SERVICE), SHALL NOT EXCEED THE AMOUNT PAID UNDER THIS AGREEMENT FOR PRODUCT AND/OR CLOUD SERVICES WITHIN THE TWELVE (12) MONTH PERIOD IMMEDIATELY PRECEDING THE CLAIM, WHETHER A CLAIM IS BASED ON CONTRACT OR TORT, INCLUDING NEGLIGENCE. IN NO EVENT SHALL ARBOR OR ITS VENDORS BE LIABLE FOR ANY INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, INCLUDING, WITHOUT LIMITATION, DAMAGES RESULTING FROM LOSS OF PROFITS, DATA, OR BUSINESS ARISING OUT OF OR IN CONNECTION WITH THIS AGREEMENT, EVEN IF ARBOR HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL ARBOR BE LIABLE FOR ANY UNAUTHORIZED ACCESS TO, ALTERATION OF, OR THE DELETION, DESTRUCTION DAMAGE, LOSS OR FAILURE TO STORE ANY OF YOUR CONTENT OR OTHER DATA. YOUR SOLE RECOURSE HEREUNDER SHALL BE AGAINST ARBOR AND YOU SHALL HOLD THIRD PARTY VENDORS HARMLESS.

6. Product Installation and Support. Installation purchased directly from Arbor with Product is governed by this Agreement, but Arbor shall not be required to continue any installation for longer than 90 days following receipt of Product. If a perpetual license is granted hereunder, You agree to purchase support ("Support") for at least the initial year from shipment. Thereafter, Arbor will invoice approximately sixty (60) days prior to the end of the Support term for additional one-year periods so long as Product is covered by Support. Failure to pay such invoice will result in a lapse of Your Support. If Support lapses, upon renewal of Support a 10% reinstatement fee will be assessed and you shall pay all Support fees back to the date Support lapsed. Each annual renewal service price shall be no less than the previous service price. With Support, Arbor will provide You (i) telephone and email based technical support in accordance with the level purchased and (ii) all new maintenance releases to Product when and if available during Your participation in Support. Arbor shall not be required to provide Support on any Product (i) for more than twelve months after its general release, or (ii) more than one release behind the currently shipping release. Arbor shall be permitted to subcontract any or all of its services or Support obligations under this Agreement to an affiliated company including, without limitation, Arbor Networks, Inc. in the United States.

7. Export Regulation and Government Rights. You agree to comply strictly with all U.S. export control laws, including the U.S. Export Administration Act and Export Administration Regulations ("EAR"). Product is prohibited for export or re-export to the list of terrorist supporting countries or to any person or entity on the U.S. Department of Commerce Denied Persons List or on the U.S. Department of Treasury's lists of Specially Designated Nationals, Specially Designated Narcotics Traffickers or Specially Designated Terrorists. If Product is being shipped by Arbor, then it is exported from the U.S. in accordance with the EAR. Diversion contrary to U.S. law is prohibited. If You are licensing Product or its accompanying documentation on behalf of the U.S. Government, it is classified as "Commercial Computer Product" and "Commercial Computer Documentation" developed at private expense, contains confidential information and trade secrets of Arbor and its licensors, and is subject to "Restricted Rights" as that term is defined in the Federal Acquisition Regulations ("FARs"). Contractor/Manufacturer is: Arbor Networks, Inc., and its subsidiaries, Burlington, Massachusetts, USA.

8. Modifications to the Agreement. Notwithstanding anything to the contrary in this Agreement, Arbor may modify Sections 1-3 and 6-8 of this Agreement (including any referenced policies or terms) as they relate to the Cloud Service at any time by posting a revised version at www.arbornetworks.com/cloud-suppterms and any successor site designated by Arbor. The modified terms will become effective upon posting. By continuing to use the Cloud Service after the effective date of any modification to this Agreement, you agree to be bound by the modified terms. It is Your responsibility to check the referenced site regularly for modifications to this Agreement.

9. General. This Agreement is made under the laws of the Commonwealth of Massachusetts, USA, excluding the choice of law and conflict of law provisions. You consent to the federal and state courts of Massachusetts as sole jurisdiction and venue for any litigation arising from or relating to this Agreement. This Agreement is the entire agreement between You and Arbor relating to Product and Cloud Service and supersedes all prior, contemporaneous and future communications, proposals and understandings with respect to its subject matter, as well as without limitation terms and conditions of any past, present or future purchase order. No modification to this Agreement is binding unless in writing and signed by a duly authorized representative of each party. The waiver or failure of either party to exercise any right provided for herein shall not be deemed a waiver of any further right hereunder. If any provision of this Agreement is held invalid, all other provisions shall continue in full force and effect. All licenses and rights granted hereunder shall terminate upon expiration of the term or Your breach of this Agreement. Neither party shall be liable for the failure to perform its obligations under this Agreement due to events beyond such party's reasonable control including, but not limited to, strikes, riots, wars, fire, acts of God or acts in compliance with any applicable law, regulation or order of any court or governmental body. Neither party may assign its rights, duties or obligations under this Agreement without the prior written consent of the other party and any attempt to do so shall be void; except to a successor by merger, acquisition or restructuring that assumes the rights and duties of this Agreement. The following sections survive termination or expiration of this Agreement: Proprietary Rights and Restrictions, Confidentiality, Limitations, Export and Government Rights, and General. All Product shipments are FCA Shipping Point and title to machines shall pass upon shipment. (09-16-14)