QUANTUM

22 January 2026

PERFORMANCE TUNING

R82

Administration Guide

Check Point Copyright Notice

Important Information

Latest Software

We recommend that you install the most recent software release to stay up-to-date with the latest functional improvements, stability fixes, security enhancements and protection against new and evolving attacks.

Certifications

For third party independent certification of Check Point products, see the Check Point Certifications page.

Check Point R82

For more about this release, see the R82 home page.

Latest Version of this Document in English

Open the latest version of this document in a Web browser.
Download the latest version of this document in PDF format.

Feedback

Check Point is engaged in a continuous effort to improve its documentation.
Please help us by sending your comments.

Revision History

Date	Description
14 January 2026	Updated: ■ *"SecureXL in User Mode (UPPAK)" on page 22*
14 January 2026	Updated: ■ *"fwaccel stats" on page 147* ■ *"Example Outputs of the "fwaccel stats" Commands" on page 158*
25 December 2025	Updated: ■ *"Multi-Queue Basic Configuration" on page 445*
07 December 2025	Updated: ■ *"Multi-Queue Requirements and Limitations" on page 437*
12 November 2025	Updated: ■ *"Rate Limiting for DoS Mitigation" on page 36*
29 January 2025	Updated: ■ *"Working with Kernel Parameters" on page 479* ■ *"Kernel Debug" on page 480*
04 December 2024	Updated: ■ *"Rate Limiting for DoS Mitigation" on page 36*
25 November 2024	Added: ■ *"fw ctl multik queues" on page 386* ■ *"fw ctl multik snd_dist" on page 388* Updated: ■ *"fw ctl multik print_heavy_conn" on page 383*
21 October 2024	First release of this document

Table of Contents

Introduction to Performance Tuning

There features improve the performance of Check Point Security Gateway / ClusterXL / Scalable Platform Security Group:

- **SecureXL** - accelerates traffic (see *"SecureXL" on page 13*)
- **CoreXL** - runs multiple Firewall instances at the same time (see *"CoreXL" on page 314*)
- **HyperFlow** - handles elephant connections on more than one CPU core in parallel (see *"HyperFlow" on page 463*)
- **Multi-Queue** - configures multiple traffic queues for each network interface (see *"Multi-Queue" on page 436*)

SecureXL

This feature accelerates traffic that passes through a Security Gateway / each Cluster Member / Scalable Platform Security Group.

Accelerated Features

- Access Control
- Encryption
- NAT
- Software Blades
  - Firewall
  - IPS features
  - Application Control
  - URL Filtering
  - Anti-Virus
  - Anti-Bot
  - Identity Awareness (SecureXL does not create templates for traffic from Identity Agents)
  - VPN Site-to-Site
  - HTTPS Inspection
  - QoS
- Policy installation
- Accounting and logging
- Connection/session rate
- General security checks
- ClusterXL High Availability and Load Sharing
- TCP Sequence Verification
- Dynamic VPN
- Passive streaming
- Active streaming

SecureXL Packet Flow in Kernel Mode (KPPAK)

For more information about SecureXL modes, see *"Configuring SecureXL" on page 20*.

This is the general description of the packet flow through a Security Appliance without a NVIDIA ConnectX 100G Card when SecureXL works in Kernel Mode (KPPAK):

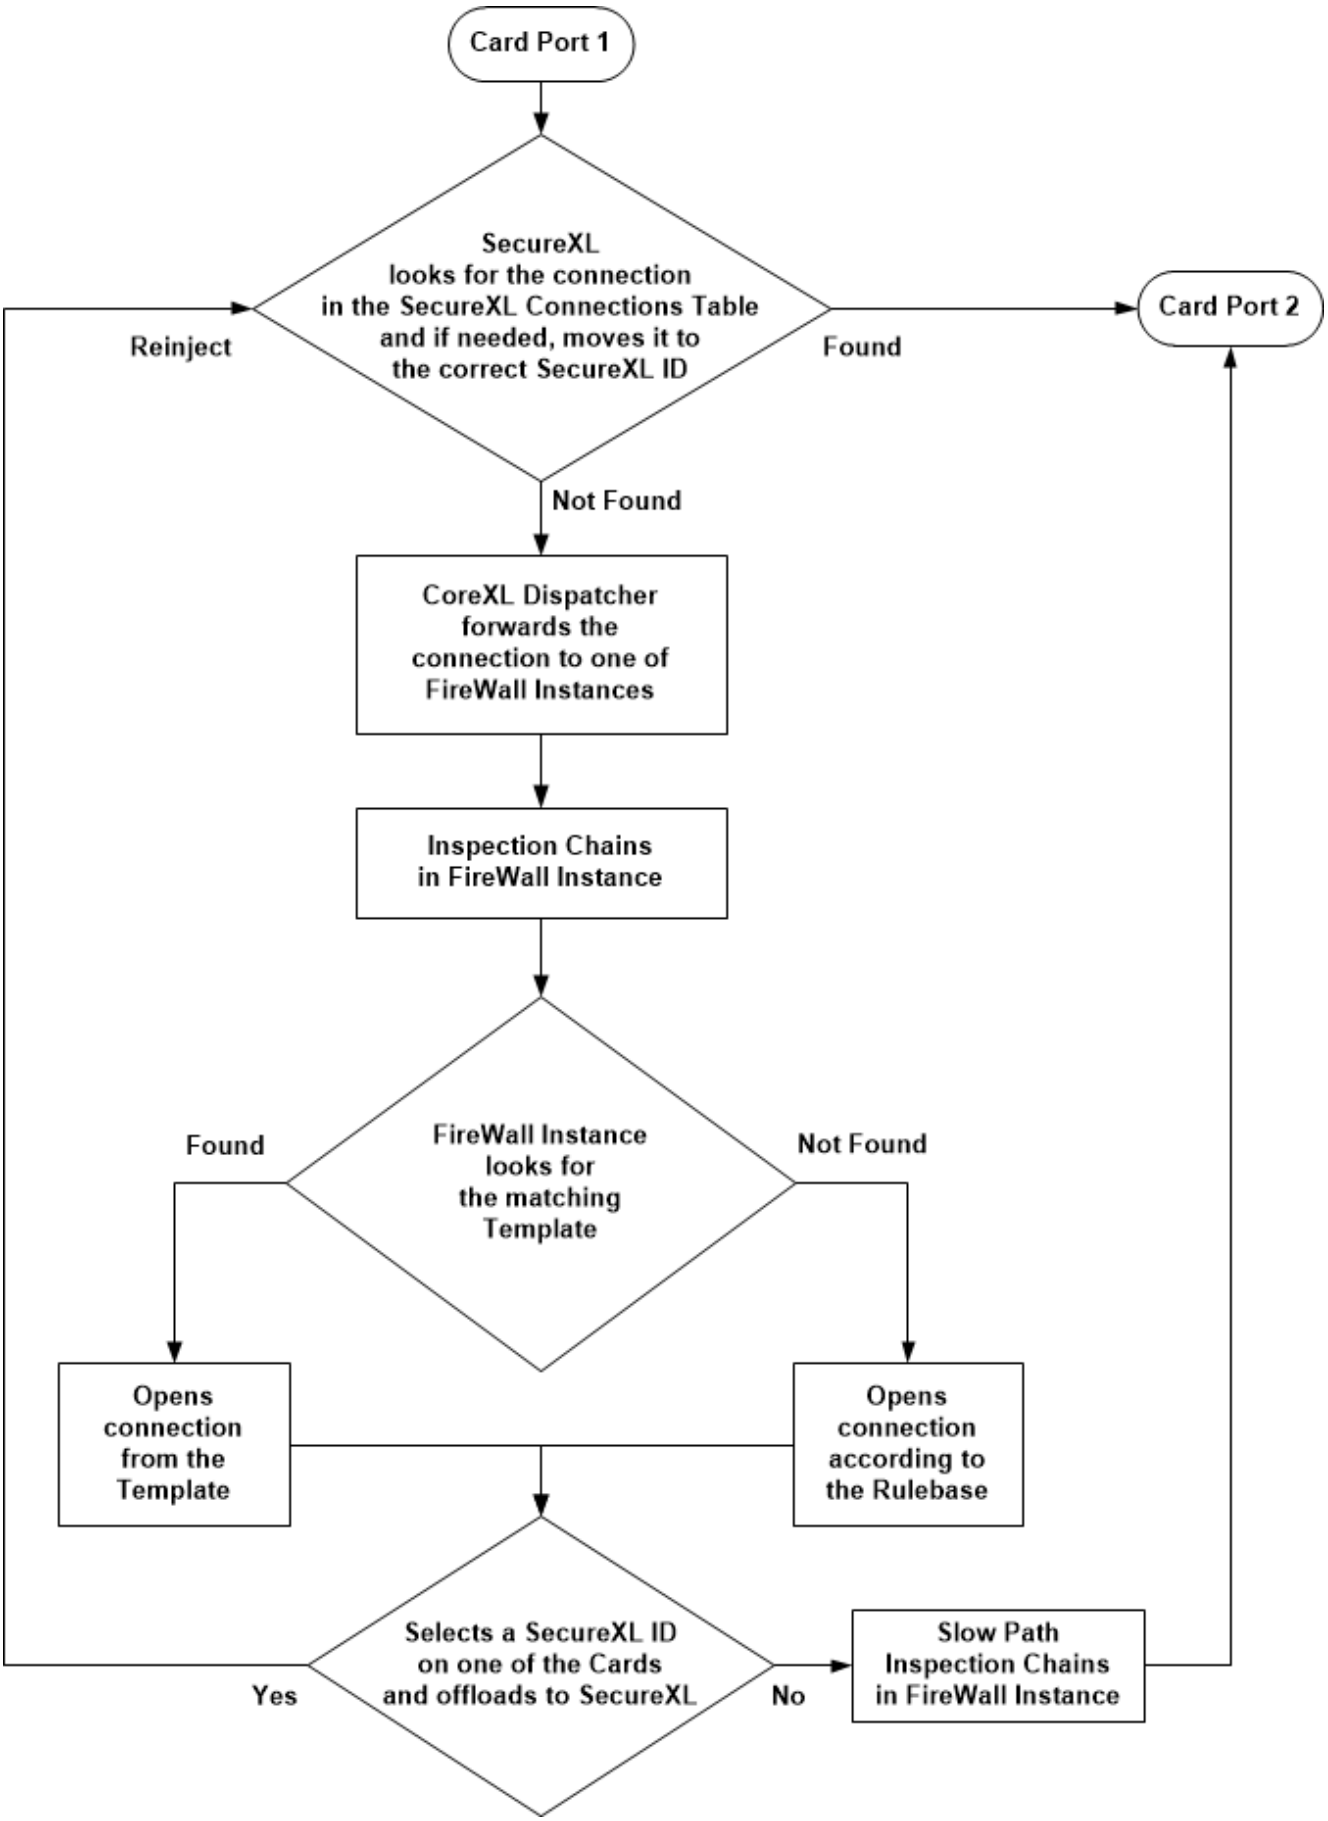
Card Port 1
SecureXL looks for the connection in the SecureXL Connections Table and if needed, moves it to the correct SecureXL ID
Reinject
Found
Card Port 2
Not Found
CoreXL Dispatcher forwards the connection to one of FireWall Instances
Inspection Chains in FireWall Instance
FireWall Instance looks for the matching Template
Found
Not Found
Opens connection from the Template
Opens connection according to the Rulebase
Selects a SecureXL ID on one of the Cards and offloads to SecureXL
Yes
No
Slow Path Inspection Chains in FireWall Instance

Connection Templates

The Connection Templates feature accelerates the speed, at which new connections from the same source IP address to the same destination IP address and to the same destination port are established.

To achieve the maximum acceleration enhancement, only the Firewall on the Security Appliance creates these Connection Templates from active connections according to the Rule Base.

Important - For the list of restrictions that apply to the Connection Templates, see sk32578.

Policy Installation Acceleration

Acceleration is enabled during policy installation.

SecureXL continues to run and stay enabled during a policy installation.

This decreases the load on the Security Gateway's CPU.

Scalable Performance

R80.20 and higher versions include improved SecureXL scalability during high session rate.

As a result, there are no longer limitations on the number of CoreXL SND cores (see *"CoreXL" on page 314*).

Configuring SecureXL

The Gaia First Time Configuration Wizard automatically installs, enables, and configures SecureXL on your Security Gateway (Scalable Platform Security Group). No additional configuration is required.

SecureXL can work in these modes:

SecureXL Mode	Description
User Mode (UPPAK)	SecureXL runs as processes in the user space (UPPAK - "User Space Performance Pack"). This mode increases performance and unlocks more advanced features in SecureXL. This is the default mode on the supported Check Point appliances. **Important** - For the list of supported Check Point appliances and Known Limitations, see the *LightSpeed 10/25/40/100G QSFP28 Ports Administration Guide*.
Kernel Mode (KPPAK)	SecureXL runs as a kernel module in the kernel space (KPPAK - "Kernel Space Performance Pack").

See:

- *"Viewing the Current SecureXL Mode" on page 23*.
- *"Changing the Current SecureXL Mode" on page 24*.

SecureXL in Kernel Mode (KPPAK)

SecureXL runs as a kernel module in the kernel space (KPPAK - "Kernel Space Performance Pack").

Description

On Security Gateways that do not support SecureXL in User Mode (UPPAK), it works in Kernel Mode.

SecureXL kernel modules:

SecureXL in Kernel Mode uses these kernel modules (see Introduction to Kernel Parameters):

- `$PPKDIR/boot/modules/sim_kern_64_3_10_64.o`
- `$PPKDIR/boot/modules/sim_kern_64_3_10_64_v6.o`

SecureXL configuration file:

SecureXL in Kernel Mode uses this configuration file for its parameters (see SecureXL Kernel Parameters):

- `$PPKDIR/conf/simkern.conf`

SecureXL in User Mode (UPPAK)

SecureXL runs as processes in the user space (UPPAK - "User Space Performance Pack").

Important:
This feature is available on the supported Check Point appliances. For the list of supported Check Point appliances and Known Limitations, see the *LightSpeed 10/25/40/100G QSFP28 Ports Administration Guide*.
SecureXL in UPPAK mode is not supported on ClusterXL in Active-Active mode.
Refer to sk32578.

Description

SecureXL user space processes:

SecureXL in User Mode uses these processes and log files:

Process	Log File	Description
`$PPKDIR/bin/usim_x86`	`/var/log/usim_x86.elg`	The main SecureXL process.
`$PPKDIR/bin/usim_wd_agent`	N / A	The Watch Dog process that monitors the main SecureXL process "`usim_x86`". If the main process crashes, this Watch Dog process starts it again.
`$PPKDIR/bin/usim_launcher`	N / A	Starts the main SecureXL process "`usim_x86`" during boot.

SecureXL configuration file:

SecureXL in User Mode uses this configuration file for its parameters (see SecureXL Kernel Parameters):

- `$PPKDIR/conf/simkern.conf`

SecureXL core dump files:

SecureXL in User Mode creates these files when its user space processes crash:

- `/var/log/dump/usermode/usim_x86.<PID>.core`
- `/var/log/dump/usermode/lcore-worker<ID>.core`
- `/var/log/dump/usermode/fwk_snd<ID>.core`
- `/var/log/usim_crash/crash_list`

Viewing the Current SecureXL Mode

Procedure

Step	Instructions
1	Connect to the command line on your Security Gateway / each Cluster Member. **Note** - On Scalable Platforms, you must connect to the Gaia Portal of the applicable Security Group.
2	Log in to Gaia Clish, or Expert mode. **Note** - On Scalable Platforms, you must run the applicable commands in Gaia gClish or the Expert mode of the applicable Security Group.
3	Examine the SecureXL status and mode (see *"fwaccel stat" on page 145*). ■ On a Security Gateway (each Cluster Member), run in Gaia Clish or the Expert mode: `fwaccel stat` ■ On a Scalable Platform Security Group, run in Gaia gClish: `fwaccel stat` ■ On a Scalable Platform Security Group, run in the Expert mode: `g_fwaccel stat`
4	Examine the column **Name**: ■ `KPPAK` - Kernel Mode ■ `UPPAK` - User Mode

Changing the Current SecureXL Mode

Procedure

Notes:

- During a clean installation, the Gaia First Time Configuration Wizard automatically configures SecureXL to work in User Mode (UPPAK) on Security Gateways that meet the requirements.
- After an upgrade from R81.10 and lower versions, SecureXL works in Kernel Mode (KPPAK).
- SecureXL does **not** change its mode automatically if the hardware specifications of a Security Gateway change, and now they meet the requirements.
  You must change the SecureXL mode manually.
  If the Security Gateway now supports USFW, then enable it, but do not reboot yet. Change the SecureXL mode to User Mode (UPPAK) and then reboot only one time.
- When SecureXL works in User Mode (UPPAK), it does not allow you to enable features that UPPAK does not support.

You can change the current SecureXL mode between Kernel Mode (KPPAK) and User Mode (UPPAK).

Step	Instructions
1	Connect to the command line on your Security Gateway / each Cluster Member. **Note** - On Scalable Platforms, you must connect to the Gaia Portal of the applicable Security Group.
2	Log in to Gaia Clish, or Expert mode. **Note** - On Scalable Platforms, you must run the applicable commands in Gaia gClish or the Expert mode of the applicable Security Group.
3	Run: `cpconfig`
4	Enter the number of the **Check Point SecureXL** option.
5	The menu shows the current SecureXL mode.
6	Enter the number of the **Change SecureXL Mode** option.
7	Enter **y** to confirm the change.
8	Exit from the `cpconfig` menu.

Step	Instructions
9	Reboot. **Important** - In cluster, this can cause a failover.
10	Examine the SecureXL status and mode: ■ On a Security Gateway (each Cluster Member), run in Gaia Clish or the Expert mode: `fwaccel stat` ■ On a Scalable Platform Security Group, run in Gaia gClish: `fwaccel stat` ■ On a Scalable Platform Security Group, run in the Expert mode: `g_fwaccel stat`

Disabling SecureXL

It is not supported to disable SecureXL.
You can disable SecureXL only if Check Point Support explicitly instructs you to do so for debug purposes.

Explanation and Procedure

Starting from R80.20, you can disable the SecureXL only *temporarily*.

The SecureXL starts automatically when you start Check Point services (with the `cpstart` command), or reboot the Security Gateway (Scalable Platform Security Group Member).

Important:

- If you disable the SecureXL, this change does **not** survive reboot.
  SecureXL remains disabled until you enable it again on-the-fly, or reboot the Security Gateway (Scalable Platform Security Group Member).
- If you disable the SecureXL, this change applies only to new connections that arrive after you disabled the acceleration.
  SecureXL continues to accelerate the connections that are already accelerated.
  Other non-connection oriented processing continues to function (for example, virtual defragmentation and VPN decrypt).
- In a Cluster, you must configure all the Cluster Members in the same way.

To disable SecureXL for IPv4 temporarily

Step	Instructions
1	Connect to the command line on your Security Gateway / each Cluster Member. **Note** - On Scalable Platforms, you must connect to the Gaia Portal of the applicable Security Group.
2	Log in to Gaia Clish, or Expert mode. **Note** - On Scalable Platforms, you must run the applicable commands in Gaia gClish or the Expert mode of the applicable Security Group.
3	Examine the SecureXL status. ■ On a Security Gateway (each Cluster Member), run in Gaia Clish or the Expert mode: `fwaccel stat` ■ On a Scalable Platform Security Group, run in Gaia gClish: `fwaccel stat` ■ On a Scalable Platform Security Group, run in the Expert mode: `g_fwaccel stat`
4	Disable the SecureXL. ■ On a Security Gateway (each Cluster Member), run in Gaia Clish or the Expert mode: `fwaccel off [-a]` ■ On a Scalable Platform Security Group, run in Gaia gClish: `fwaccel off [-a]` ■ On a Scalable Platform Security Group, run in the Expert mode: `g_fwaccel off [-a]`

Step	Instructions
5	Examine the SecureXL status. ■ On a Security Gateway (each Cluster Member), run in Gaia Clish or the Expert mode: `fwaccel stat` ■ On a Scalable Platform Security Group, run in Gaia gClish: `fwaccel stat` ■ On a Scalable Platform Security Group, run in the Expert mode: `g_fwaccel stat`

To disable SecureXL for IPv6 temporarily

Step	Instructions
1	Connect to the command line on your Security Gateway / each Cluster Member. **Note** - On Scalable Platforms, you must connect to the Gaia Portal of the applicable Security Group.
2	Log in to Gaia Clish, or Expert mode. **Note** - On Scalable Platforms, you must run the applicable commands in Gaia gClish or the Expert mode of the applicable Security Group.
3	Examine the SecureXL status. ■ On a Security Gateway (each Cluster Member), run in Gaia Clish or the Expert mode: `fwaccel6 stat` ■ On a Scalable Platform Security Group, run in Gaia gClish: `fwaccel6 stat` ■ On a Scalable Platform Security Group, run in the Expert mode: `g_fwaccel6 stat`
4	Disable the SecureXL. ■ On a Security Gateway (each Cluster Member), run in Gaia Clish or the Expert mode: `fwaccel6 off [-a]` ■ On a Scalable Platform Security Group, run in Gaia gClish: `fwaccel6 off [-a]` ■ On a Scalable Platform Security Group, run in the Expert mode: `g_fwaccel6 off [-a]`

Step	Instructions
5	Examine the SecureXL status. ■ On a Security Gateway (each Cluster Member), run in Gaia Clish or the Expert mode: `fwaccel6 stat` ■ On a Scalable Platform Security Group, run in Gaia gClish: `fwaccel6 stat` ■ On a Scalable Platform Security Group, run in the Expert mode: `g_fwaccel6 stat`

To enable SecureXL again for IPv4

Step	Instructions
1	Connect to the command line on your Security Gateway / each Cluster Member. **Note** - On Scalable Platforms, you must connect to the Gaia Portal of the applicable Security Group.
2	Log in to Gaia Clish, or Expert mode. **Note** - On Scalable Platforms, you must run the applicable commands in Gaia gClish or the Expert mode of the applicable Security Group.
3	Examine the SecureXL status. ■ On a Security Gateway (each Cluster Member), run in Gaia Clish or the Expert mode: `fwaccel stat` ■ On a Scalable Platform Security Group, run in Gaia gClish: `fwaccel stat` ■ On a Scalable Platform Security Group, run in the Expert mode: `g_fwaccel stat`
4	Enable the SecureXL. ■ On a Security Gateway (each Cluster Member), run in Gaia Clish or the Expert mode: `fwaccel on [-a]` ■ On a Scalable Platform Security Group, run in Gaia gClish: `fwaccel on [-a]` ■ On a Scalable Platform Security Group, run in the Expert mode: `g_fwaccel on [-a]`

Step	Instructions
5	Examine the SecureXL status. ■ On a Security Gateway (each Cluster Member), run in Gaia Clish or the Expert mode: `fwaccel stat` ■ On a Scalable Platform Security Group, run in Gaia gClish: `fwaccel stat` ■ On a Scalable Platform Security Group, run in the Expert mode: `g_fwaccel stat`

To enable SecureXL again for IPv6

Step	Instructions
1	Connect to the command line on your Security Gateway / each Cluster Member. **Note** - On Scalable Platforms, you must connect to the Gaia Portal of the applicable Security Group.
2	Log in to Gaia Clish, or Expert mode. **Note** - On Scalable Platforms, you must run the applicable commands in Gaia gClish or the Expert mode of the applicable Security Group.
3	Examine the SecureXL status. ■ On a Security Gateway (each Cluster Member), run in Gaia Clish or the Expert mode: `fwaccel6 stat` ■ On a Scalable Platform Security Group, run in Gaia gClish: `fwaccel6 stat` ■ On a Scalable Platform Security Group, run in the Expert mode: `g_fwaccel6 stat`
4	Enable the SecureXL. ■ On a Security Gateway (each Cluster Member), run in Gaia Clish or the Expert mode: `fwaccel6 on [-a]` ■ On a Scalable Platform Security Group, run in Gaia gClish: `fwaccel6 on [-a]` ■ On a Scalable Platform Security Group, run in the Expert mode: `g_fwaccel6 on [-a]`

Step	Instructions
5	Examine the SecureXL status. ■ On a Security Gateway (each Cluster Member), run in Gaia Clish or the Expert mode: `fwaccel6 stat` ■ On a Scalable Platform Security Group, run in Gaia gClish: `fwaccel6 stat` ■ On a Scalable Platform Security Group, run in the Expert mode: `g_fwaccel6 stat`

For more information on the "`fwaccel`" commands, see:

- *"fwaccel stat" on page 145*
- *"fwaccel off" on page 128*
- *"fwaccel on" on page 133*

Analyzing the Accelerated Traffic

To capture and analyze the accelerated traffic, use the *"fw monitor" on page 196* command.

Rate Limiting for DoS Mitigation

Introduction

DoS / Rate Limiting is a defense against DoS (Denial-of-Service) attacks.

DoS / Rate Limiting includes these features:

- Rate Limiting Rules
- IP Deny List
- Block IP Fragments
- Block IP Options
- Penalty Box

In general, these features solve separate problems and are managed / configured separately. However, be aware that there are some global settings that will affect the behavior of multiple features simultaneously.

To maximize performance, most of the DoS / Rate Limiting policy is enforced as early as possible in the packet flow. For most features this means it is enforced in SecureXL. Connection-based policy is the single exception. This policy is enforced by the Firewall Software Blade, because this is where the related connection state is stored and managed.

Important:

- To configure Rate Limiting for DoS Mitigation, follow: sk182350 - How to configure Rate Limiting rules for DoS Mitigation.
- The DoS / Rate Limiting policy is enforced even if SecureXL is disabled with the "`fwaccel off`" command.

  Warning - Do **not** disable SecureXL, unless Check Point R&D explicitly instructed you to do so.
- By design, the DoS / Rate Limiting policy is **not** applied to packets after VPN decryption.
- During the installation of the Security Policy on the Security Gateway / ClusterXL, the DoS / Rate Limiting policy is **not** enforced.
- You cannot use the Rate Limiting feature for specific URLs. This feature applies to all traffic.
- SecureXL Rate Limiting rules for DoS Mitigation supports these parameters with automatic IP range updating enabled by default starting R82 Jumbo Hotfix Accumulator Take 44:
  - `cc:<COUNTRY_CODE>`
  - `asn:<AUTONOMOUS_SYSTEM_NUMBER>`

Monitoring Events Related to DoS Mitigation on a Security Gateway / ClusterXL

To see some information related to DoS Mitigation, run these commands:

Command in Gaia Clish or the Expert mode	Description
`fwaccel stats` `fwaccel6 stats`	Shows all SecureXL statistics (for IPv4 and IPv6 kernel modules). See: ■ *"fwaccel stats" on page 147* ■ *"The /proc/ppk/ and /proc/ppk6/ entries" on page 259*
`fwaccel stats -d` `or` `cat /proc/ppk/drop_statistics` `fwaccel6 stats -d` `or` `cat /proc/ppk6/drop_statistics`	Shows SecureXL drop statistics only (for IPv4 and IPv6 kernel modules). See: ■ *"fwaccel stats" on page 147* ■ *"The /proc/ppk/ and /proc/ppk6/ entries" on page 259* ■ *"fw sam_policy" on page 231*
`fw samp get -l |\` `grep '^<[0-9a-f,]*>$' \|\` `xargs fwaccel dos rate get` `fw samp get -l \|\` `grep '^<[0-9a-f,]*>$' |` `xargs fwaccel6 dos rate get`	Shows details of active policy rules in long format (for IPv4 and IPv6 kernel modules). See *"fw sam_policy get" on page 253*.
`cat /proc/ppk/rlc`	Shows: ■ Total drop packets ■ Total drop bytes See *"The /proc/ppk/ and /proc/ppk6/ entries" on page 259*.

Monitoring Events Related to DoS Mitigation on Scalable Platforms

Note - On Scalable Platforms, you must connect to the Gaia Portal of the applicable Security Group.

To see some information related to DoS Mitigation, run these commands:

Command in Gaia gClish	Command in the Expert mode	Instructions
`fwaccel stats` `fwaccel6 stats`	`g_fwaccel stats` `g_fwaccel6 stats`	Shows all SecureXL statistics (for IPv4 and IPv6 kernel modules). See: ■ *"fwaccel stats" on page 147* ■ *"The /proc/ppk/ and /proc/ppk6/ entries" on page 259*
`fwaccel stats -d` `fwaccel6 stats -d`	`g_fwaccel stats -d` or `cat /proc/ppk/drop_statistics` `g_fwaccel6 stats -d` or `cat /proc/ppk6/drop_statistics`	Shows SecureXL drop statistics only (for IPv4 and IPv6 kernel modules). See: ■ *"fwaccel stats" on page 147* ■ *"The /proc/ppk/ and /proc/ppk6/ entries" on page 259* ■ *"fw sam_policy" on page 231*

Command in Gaia gClish	Command in the Expert mode	Instructions
`fw samp get -l |\ grep '^<[0-9a-f,]*>$' |\ xargs fwaccel dos rate get`	`g_fw samp get -l |\ grep '^<[0-9a-f,]*>$' |\ xargs fwaccel dos rate get`	Shows details of active policy rules in long format (for IPv4 and IPv6 kernel modules). See *"fw sam_policy get" on page 253*.
`fw samp get -l |\ grep '^<[0-9a-f,]*>$' | xargs fwaccel6 dos rate get`	`g_fw samp get -l |\ grep '^<[0-9a-f,]*>$' | xargs fwaccel6 dos rate get`	
N / A	`cat /proc/ppk/rlc`	Shows: ■ Total drop packets ■ Total drop bytes See *"The /proc/ppk/ and /proc/ppk6/ entries" on page 259*.

Note - In addition, see *"SecureXL Debug" on page 282*.

Accelerated SYN Defender

Introduction

A TCP SYN Flood attack occurs when a host, typically with a forged IP address, sends a flood of TCP [SYN] packets. Each of these TCP [SYN] packets is handled as a connection request, which causes the server to create a half-open (unestablished) TCP connection. This occurs because the server sends a TCP [SYN+ACK] packet, and waits for a response TCP packet that does not arrive.

These half-open TCP connections eventually exceed the maximum available TCP connections. This causes a denial of service condition.

The Check Point Accelerated SYN Defender protects the Security Gateway (Scalable Platform Security Group) by preventing excessive TCP connections from being created.

The Accelerated SYN Defender uses TCP [SYN] Cookies (particular choices of initial TCP sequence numbers) when under a suspected TCP SYN Flood attack. Using TCP [SYN] Cookies can reduce the load on Security Gateway and on computers behind the Security Gateway (Scalable Platform Security Group). The Accelerated SYN Defender acts as proxy for TCP connections and adjusts TCP {SEQ} and TCP {ACK} values in TCP packets.

This is a sample TCP timeline diagram that shows a TCP connection through the Security Gateway (Scalable Platform Security Group) with the enabled Accelerated SYN Defender:

Note - In this example, we assume that there no TCP retransmissions and no early data.

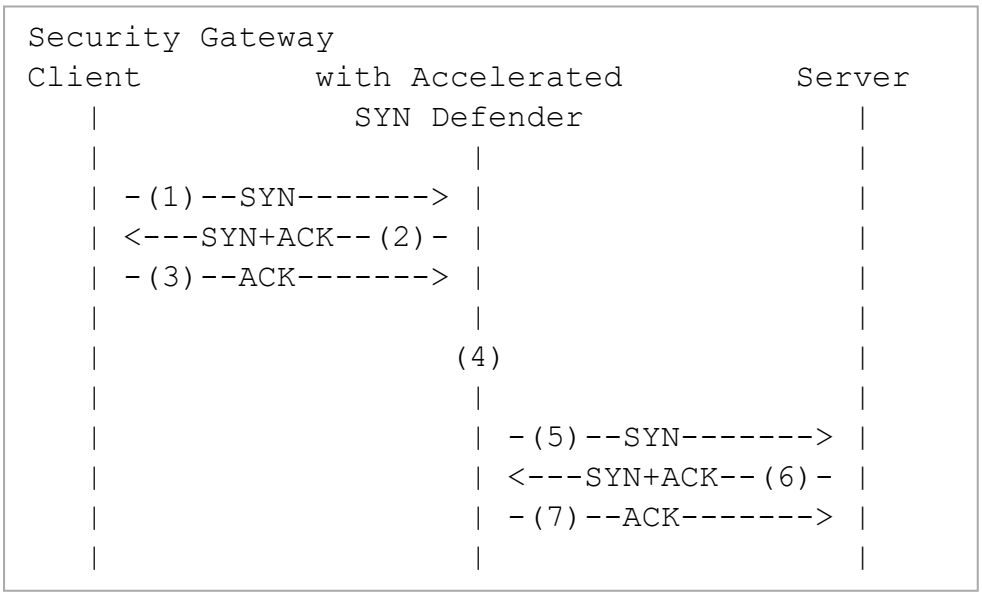

1. A Client sends a TCP [SYN] packet to a Server.
2. The Accelerated SYN Defender replies to the Client with a TCP [SYN+ACK] packet that contains a special cookie in the `Seq` field.

The Security Gateway (Scalable Platform Security Group) does not maintain the connection state at this time.

3. The Client sends a reply TCP [ACK] packet. This completes the Client-side of the TCP connection.
4. The Accelerated SYN Defender checks if the SYN cookie in the Client's TCP [ACK] packet is legitimate.
5. If the SYN cookie in the Client's TCP [ACK] packet is legitimate, the Accelerated SYN Defender sends a TCP [SYN] packet to the Server to begin the Server-side of the TCP connection.
6. The Server replies with a TCP [SYN+ACK] packet.
7. The Accelerated SYN Defender sends a TCP [ACK] packet to complete the Server-size of the TCP 3-way handshake.
8. The Accelerated SYN Defender marks the TCP connection as established and records the TCP sequence adjustment between the two sides.

SecureXL handles the TCP [SYN] packets. The Security Gateway (Scalable Platform Security Group) handles the rest of the TCP connection setup.

For each TCP connection the Accelerated SYN Defender establishes, the Security Gateway (Scalable Platform Security Group) adjusts the TCP sequence number for the life of that TCP connection.

Command Line Interface

Use the *"fwaccel synatk" on page 166* commands to configure the Accelerated SYN Defender.

Configuring the IPS 'SYN Attack' protection in SmartConsole

Important - Scalable Platform Security Group does not support the configuration of the IPS 'SYN Attack' protection in SmartConsole(Known Limitation MBS-5415).

The IPS 'SYN Attack' protection is intended to mitigate SYN Flood attacks.

Step	Instructions
1	Connect with SmartConsole to the Management Server.
2	From the left navigation panel, click **Security Policies**.
3	In the **Shared Policies** section, click **Inspection Settings**.
4	In the top field, search for **SYN Attack**.
5	Double-click on the **SYN Attack** protection.
6	Edit the applicable Inspection profile.
7	Configure the applicable settings in the profile: ■ On the **General Properties** page: If you select **Override with Action** and then **Accept** or **Drop**, it overrides the settings you make on the Security Gateway with the *"fwaccel synatk" on page 166* commands. ■ On the **Advanced** page: The option you select in the **Activation Settings** (**Protect all interfaces** or **Protect external interfaces only**) overrides the settings you make on the Security Gateway with the *"fwaccel synatk" on page 166* commands.
9	Install the Access Control Policy.

For more information about the **SYN Attack** protection in SmartConsole, see sk120476.

SecureXL Commands and Debug

This section describes:

- SecureXL CLI commands
- SecureXL CLI Debug

Syntax Legend for CLI Commands

Whenever possible, this guide lists commands, parameters and options in the alphabetical order.

This guide uses this convention in the Command Line Interface (CLI) syntax:

Character	Description
TAB	Shows the available nested subcommands: `main command` `→ nested subcommand 1` `→ → nested subsubcommand 1-1` `→ → nested subsubcommand 1-2` `→ nested subcommand 2` Example: `cpwd_admin` `config` `-a <options>` `-d <options>` `-p` `-r` `del <options>` Meaning, you can run only **one** of these commands: ■ This command: `cpwd_admin config -a <options>` ■ Or this command: `cpwd_admin config -d <options>` ■ Or this command: `cpwd_admin config -p` ■ Or this command: `cpwd_admin config -r` ■ Or this command: `cpwd_admin del <options>`
Curly brackets or braces **{ }**	Enclose a list of available commands or parameters, separated by the vertical bar **|**. User can enter only one of the available commands or parameters.

Character	Description
Angle brackets **< >**	Enclose a variable. User must explicitly specify a supported value.
Square brackets or brackets **[ ]**	Enclose an optional command or parameter, which user can also enter.

cpview

Overview of CPView

Description

CPView is a text based built-in utility on a Check Point computer.

CPView Utility shows statistical data that contain both general system information (CPU, Memory, Disk space) and information for different Software Blades (only on a Security Gateway / ClusterXL / Scalable Platform Security Group).

The CPView continuously updates the data in easy to access views.

On a Security Gateway / ClusterXL / Scalable Platform Security Group, you can use this statistical data to monitor the performance.

For more information, see sk101878.

Syntax

```
cpview --help
```

CPView User Interface

The CPView user interface has three sections:

Section	Description
Header	This view shows the time the statistics in the third view are collected. It updates when you refresh the statistics.
Navigation	This menu bar is interactive. Move between menus with the arrow keys and mouse. A menu can have sub-menus and they show under the menu bar.
View	This view shows the statistics collected in that view. These statistics update at the refresh rate.

Using CPView

Use these keys to navigate the CPView:

Key	Description
Arrow keys	Moves between menus and views. Scrolls in a view.
Home	Returns to the **Overview** view.
Enter	Changes to the **View Mode**. On a menu with sub-menus, the **Enter** key moves you to the lowest level sub-menu.
Esc	Returns to the **Menu Mode**.
Q	Quits CPView.

Use these keys to change CPView interface options:

Key	Description
R	Opens a window where you can change the refresh rate. The default refresh rate is 2 seconds.
W	Changes between wide and normal display modes. In wide mode, CPView fits the screen horizontally.
S	Manually sets the number of rows or columns.
M	Switches on/off the mouse.
P	Pauses and resumes the collection of statistics.

Use these keys to save statistics, show help, and refresh statistics:

Key	Description
C	Saves the current page to a file. The file name format is: `cpview_<ID of the cpview process>.cap<Number of the capture>`
H	Shows a tooltip with CPView options.
Space bar	Immediately refreshes the statistics.

'fwaccel' and 'fwaccel6'

Description

The "`fwaccel`" commands control the acceleration for IPv4 traffic.

The "`fwaccel6`" commands control the acceleration for IPv6 traffic.

Important:

- In a Cluster, you must configure all the Cluster Members in the same way.
- On Scalable Platforms, you must connect to the applicable Security Group.

Syntax for IPv4

Syntax for IPv4 on a Security Gateway / ClusterXL - in Gaia Clish and the Expert mode

```
fwaccel help
```

```
fwaccel
      cfg <options>
      conns <options>
      dbg <options>
      dos <options>
      feature <options>
      if <options>
      ip_mr_cache
      nonaccel <options>
      off <options>
      on <options>
      ranges <options>
      stat <options>
      stats <options>
      synatk <options>
      tab <options>
      templates <options>
      ver
```

Syntax for IPv4 on a Scalable Platform Security Group - in Gaia gClish

```
fwaccel help
```

```
fwaccel
      cfg <options>
      conns <options>
      dbg <options>
      dos <options>
      feature <options>
      if <options>
      ip_mr_cache
      nonaccel <options>
      off <options>
      on <options>
      ranges <options>
      stat <options>
      stats <options>
      synatk <options>
      tab <options>
      templates <options>
      ver
```

Syntax for IPv4 on a Scalable Platform Security Group - in the Expert mode

```
g_fwaccel help
```

```
g_fwaccel
      cfg <options>
      conns <options>
      dbg <options>
      dos <options>
      feature <options>
      if <options>
      ip_mr_cache
      nonaccel <options>
      off <options>
      on <options>
      ranges <options>
      stat <options>
      stats <options>
      synatk <options>
      tab <options>
      templates <options>
      ver
```

Syntax for IPv6

Syntax for IPv6 on a Security Gateway / ClusterXL - in Gaia Clish and the Expert mode

```
fwaccel6 help
```

```
fwaccel6
      conns <options>
      dbg <options>
      dos <options>
      feature <options>
      if <options>
      ip_mr_cache
      nonaccel <options>
      off <options>
      on <options>
      ranges <options>
      stat <options>
      stats <options>
      synatk <options>
      tab <options>
      templates <options>
      ver
```

Syntax for IPv6 on a Scalable Platform Security Group - in Gaia gClish

```
fwaccel6 help
```

```
fwaccel6
      conns <options>
      dbg <options>
      dos <options>
      feature <options>
      if <options>
      ip_mr_cache
      nonaccel <options>
      off <options>
      on <options>
      ranges <options>
      stat <options>
      stats <options>
      synatk <options>
      tab <options>
      templates <options>
      ver
```

Syntax for IPv6 on a Scalable Platform Security Group - in the Expert mode

```
g_fwaccel6 help
```

```
g_fwaccel6
      conns <options>
      dbg <options>
      dos <options>
      feature <options>
      if <options>
      ip_mr_cache
      nonaccel <options>
      off <options>
      on <options>
      ranges <options>
      stat <options>
      stats <options>
      synatk <options>
      tab <options>
      templates <options>
      ver
```

Parameters and Options

Parameter and Options	Description
`help`	Shows the built-in help.
`cfg <options>`	Controls the SecureXL acceleration parameters (for IPv4 only). See *"fwaccel cfg" on page 53*.
`conns <options>`	Shows all connections that pass through SecureXL. See *"fwaccel conns" on page 57*.
`dbg <options>`	Controls the *"SecureXL Debug" on page 282*. See *"fwaccel dbg" on page 283*.
`dos <options>`	Controls the Rate Limiting for DoS Mitigation in SecureXL. See *"fwaccel dos" on page 75*.
`feature <options>`	Controls the specified SecureXL features. See *"fwaccel feature" on page 124*.
`if <options>`	Shows information about interfaces in SecureXL. See *fwaccel if*.

Parameter and Options	Description
`ip_mr_cache`	Shows the IPv4 multicast routing cache when SecureXL works in the User Mode (UPPAK). See *"fwaccel ip_mr_cache" on page 127*.
`nonaccel <options>`	Enables or disables SecureXL acceleration for the specified interfaces. See *fwaccel nonaccel*.
`off <options>`	Stops the acceleration on-the-fly. This does **not** survive reboot. See *"fwaccel off" on page 128*.
`on <options>`	Starts the acceleration on-the-fly, if it was previously stopped. See *"fwaccel on" on page 133*.
`ranges <options>`	Shows the loaded ranges. See *"fwaccel ranges" on page 138*.
`stat <options>`	Shows the SecureXL status. See *"fwaccel stat" on page 145*.
`stats <options>`	Shows the acceleration statistics. See *"fwaccel stats" on page 147*.
`synatk <options>`	Controls the Accelerated SYN Defender. See *"fwaccel synatk" on page 166*.
`tab <options>`	Shows the contents of the specified SecureXL table. See *"fwaccel tab" on page 187*.
`templates <options>`	Shows the SecureXL templates. See *"fwaccel templates" on page 190*.
`ver`	Shows the SecureXL and FireWall version. See *"fwaccel ver" on page 195*.

fwaccel cfg

Description

The "`fwaccel cfg`" command controls the SecureXL acceleration parameters (for IPv4 only).

Important:

- In a Cluster, you must configure all the Cluster Members in the same way.
- On Scalable Platforms, you must connect to the applicable Security Group.
- In the VSNext mode / Traditional VSX mode, you must go to the context of an applicable Virtual Gateway / Virtual System.
  - In Gaia Clish, run: `set virtual-system <VSID>`
  - In the Expert mode, run: `vsenv <VSID>`
- These commands do not provide output. You cannot see the currently configured values.
- Changes made with these commands do **not** survive reboot.

Syntax

Syntax on a Security Gateway / ClusterXL - in Gaia Clish and the Expert mode

```
fwaccel cfg
      -h
      -a {<Number of Interface> | <Name of Interface> | reset}
      -b {on | off}
      -c <Number>
      -d <Number>
      -e <Number>
      -i {on | off}
      -l <Number>
      -m <Seconds>
      -p {on | off}
      -r <Number>
      -v <Seconds>
      -w {on | off}
```

Syntax on a Scalable Platform Security Group - in Gaia gClish

```
fwaccel cfg
      -h
      -a {<Number of Interface> | <Name of Interface> | reset}
      -b {on | off}
      -c <Number>
      -d <Number>
      -e <Number>
      -i {on | off}
      -l <Number>
      -m <Seconds>
      -p {on | off}
      -r <Number>
      -v <Seconds>
      -w {on | off}
```

Syntax on a Scalable Platform Security Group - in the Expert mode

```
g_fwaccel cfg
      -h
      -a {<Number of Interface> | <Name of Interface> | reset}
      -b {on | off}
      -c <Number>
      -d <Number>
      -e <Number>
      -i {on | off}
      -l <Number>
      -m <Seconds>
      -p {on | off}
      -r <Number>
      -v <Seconds>
      -w {on | off}
```

Parameters

Parameter	Description
`-h`	Shows the applicable built-in help.

Parameter	Description
`-a <Number of Interface>` `-a <Name of Interface>` `-a reset`	■ `-a <Number of Interface>` Configures the SecureXL not to accelerate traffic on the interface specified by its internal number in Check Point kernel. ■ `-a <Name of Interface>` Configures the SecureXL not to accelerate traffic on the interface specified by its name. ■ `-a reset` Configures the SecureXL to accelerate traffic on all interfaces (resets the non-accelerated configuration). **Notes:** ■ This command does not support Falcon Acceleration Cards. ■ To see the required information about the interfaces, run these commands in the specified order: `fw getifs` `fw ctl iflist` ■ To see if the "`fwaccel cfg -a ...`" command failed, run this command: `tail -n 10 /var/log/messages`
`-b {on | off}`	Controls the SecureXL Drop Templates match (sk66402): ■ `on` - Enables the SecureXL Drop Templates match ■ `off` - Disables the SecureXL Drop Templates match **Note** - In R82, SecureXL does not support this parameter yet..
`-c <Number>`	Configures the maximum number of connections, when SecureXL disables the templates.
`-d <Number>`	Configures the maximum number of delete retries.
`-e <Number>`	Configures the maximum number of general errors.
`-i {on | off}`	Configures SecureXL to ignore API version mismatch: ■ `on` - Ignore API version mismatch. ■ `off` - Do not ignore API version mismatch (this is the default).

Parameter	Description
`-l <Number>`	Configures the maximum number of entries in the SecureXL templates database. Valid values are: ■ 0 - To disable the limit (this is the default). ■ Between 10 and 524288 - To configure the limit. **Important** - If you configure a limit, you must stop and start the acceleration for this change to take effect. Run the *"fwaccel off" on page 128* command and then the *"fwaccel on" on page 133* command.
`-m <Seconds>`	Configures the timeout for entries in the SecureXL templates database. Valid values are: ■ 0 - To disable the timeout (this is the default). ■ Between 10 and 524288 - To configure the timeout.
`-p {on | off}`	Configures the offload of Connection Templates (if possible): ■ `on` - Enables the offload of new templates (this is the default). ■ `off` - Disables the offload of new templates.
`-r <Number>`	Configures the maximum number of retries for SecureXL API calls.
`-v <Seconds>`	Configures the interval between SecureXL statistics request. Valid values are: ■ 0 - To disable the interval. ■ 1 and greater - To configure the interval.
`-w {on | off}`	Configures the support for warnings about the IPS protection **Sequence Verifier**: ■ `on` - Enable the support for these warnings. ■ `off` - Disables the support for these warnings.

fwaccel conns

Description

The "`fwaccel conns`" and "`fwaccel6 conns`" commands show the list of the SecureXL connections on the local Security Gateway, or Cluster Member.

Warning - If the number of concurrent connections is large, when you run these commands, they can consume memory and CPU at very high level (see sk118716).

Important:

- In a Cluster, you must configure all the Cluster Members in the same way.
- On Scalable Platforms, you must connect to the applicable Security Group.
- In the VSNext mode / Traditional VSX mode, you must go to the context of an applicable Virtual Gateway / Virtual System.
  - In Gaia Clish, run: `set virtual-system <VSID>`
  - In the Expert mode, run: `vsenv <VSID>`

Syntax for IPv4

```
fwaccel conns
      -h
      -f <filter>
      -m <Number of Entries>
      -s
      -z
```

Syntax for IPv6

```
fwaccel6 conns
      -h
      -f <Filter>
      -m <Number of Entries>
      -s
      -z
```

Parameters

Parameter	Description
`-h`	Shows the applicable built-in help.

Parameter	Description
`-f <Filter>`	Shows the SecureXL Connections Table entries based on the specified filter flags. **Notes:** ■ To see the available filter flags, run: `fwaccel conns -h` ■ Each filter flag is one letter - capital, or small. ■ You can specify more than one flag. For example: `fwaccel conns -f AaQq`

Parameter	Description
	Available filter flags are: ■ `A` - Shows accounted connections (for which SecureXL counted the number of packets and bytes). ■ `a` - Shows not accounted connections. ■ `C` - Shows encrypted (VPN) connections. ■ `c` - Shows clear-text (not encrypted) connections. ■ `F` - Shows connections that SecureXL forwarded to Firewall. **Note** - In R82, SecureXL does not support this parameter. ■ `f` - Shows cut-through connections (which SecureXL accelerated). Note - In R82, SecureXL does not support this parameter. ■ `H` - Shows connections offloaded to the acceleration card. **Note** - R82, does not support the acceleration card (Known Limitation PMTR-18774). ■ `h` - Shows connections created in the SAM card. **Note** - R82, does not support the SAM card (Known Limitation PMTR-18774). ■ `H` - Shows connections created in the NVIDIA ConnectX 100G QSFP28 2-port Card. ■ `L` - Shows connections, for which SecureXL created links in its connections table (Server-to-Client entries for the original Client-to-Server entry). ■ `l` - Shows connections, for which SecureXL did not create links in its connections table. ■ `N` - Shows connections that undergo NAT. **Note** - In R82, SecureXL does not support this parameter. ■ `n` - Shows connections that do not undergo NAT. **Note** - R82, SecureXL does not support this parameter. ■ `Q` - Shows connections that undergo QoS. ■ `q` - Shows connections that do not undergo QoS. ■ `S` - Shows connections that undergo PXL. ■ `s` - Shows connections that do not undergo PXL. ■ `U` - Shows unidirectional connections. ■ `u` - Shows bidirectional connections.
`-m <`*`Number of Entries`*`>`	Specifies the maximum number of connections to show. **Note** - In R82, SecureXL does not support this parameter.
`-s`	Shows the summary of SecureXL Connections Table (number of connections). **Warning** - Depending on the number of current connections, might consume memory at very high level.

Parameter	Description
`-z`	Shows the summary of connections that are not accelerated in hardware. **Notes:** ■ This applies only to supported Check Point appliances that support LightSpeed hardware acceleration in 10/25/40/100G QSFP28 Ports. See the *LightSpeed 10/25/40/100G QSFP28 Ports Administration Guide*. ■ To see additional information, run the CPView tool (*"cpview" on page 362*) and go to **Advanced** > **SecureXL** > **LightSpeed2Host-Reasons**.

Example - Default output from a non-VSX Gateway

```
[Expert@MyGW:0]# fwaccel conns
Source          SPort Destination     DPort PR Flags     C2S i/f S2C i/f   Inst Identity
--------------- ----- --------------- ----- -- ----------- ------- ------- ---- -------
      1.1.1.200 50586       1.1.1.100 18191  6 F.............. 2/2     2/-     3        0
  192.168.0.244 35925   192.168.0.242 18192  6 F.............. 1/1     -/-     1        0
   192.168.0.93   257   192.168.0.242 53932  6 F.............. 1/1     1/-     0        0
  192.168.0.242    22   172.30.168.15 57914  6 F.............. 1/1     -/-     2        0
  192.168.0.244 34773   192.168.0.242 18192  6 F.............. 1/1     -/-     2        0
   192.168.0.88   138   192.168.0.255   138 17 F.............. 1/1     -/-     0        0
      1.1.1.100 18191       1.1.1.200 55336  6 F.............. 2/2     2/-     4        0
  192.168.0.242 18192   192.168.0.244 38567  6 F.............. 1/1     -/-     4        0
  192.168.0.242 53932    192.168.0.93   257  6 F.............. 1/1     1/-     0        0
  192.168.0.242 18192   192.168.0.244 62714  6 F.............. 1/1     -/-     1        0
  192.168.0.244 33558   192.168.0.242 18192  6 F.............. 1/1     -/-     5        0
      1.1.1.200 36359       1.1.1.100 18191  6 F.............. 2/2     2/-     5        0
      1.1.1.200 55336       1.1.1.100 18191  6 F.............. 2/2     2/-     4        0
  192.168.0.242 60756    192.168.0.93   257  6 F.............. 1/1     1/-     4        0
      1.1.1.100 18191       1.1.1.200 36359  6 F.............. 2/2     2/-     5        0
      1.1.1.100 18191       1.1.1.200 50586  6 F.............. 2/2     2/-     3        0
  192.168.0.244 38567   192.168.0.242 18192  6 F.............. 1/1     -/-     4        0
  192.168.0.242 18192   192.168.0.244 32877  6 F.............. 1/1     -/-     5        0
  192.168.0.242 53806   192.168.47.45    53 17 F.............. 1/1     1/-     3        0
  192.168.0.242 18192   192.168.0.244 33558  6 F.............. 1/1     -/-     5        0
  172.30.168.15 57914   192.168.0.242    22  6 F.............. 1/1     -/-     2        0
  192.168.0.255   138    192.168.0.88   138 17 F.............. 1/1     -/-     0        0
   192.168.0.93   257   192.168.0.242 60756  6 F.............. 1/1     1/-     4        0
      1.1.1.200 18192       1.1.1.100 37964  6 F.............. 2/2     -/-     1        0
      1.1.1.100 37964       1.1.1.200 18192  6 F.............. 2/2     -/-     1        0
  192.168.0.244 32877   192.168.0.242 18192  6 F.............. 1/1     -/-     5        0
  192.168.0.242 18192   192.168.0.244 34773  6 F.............. 1/1     -/-     2        0
  192.168.0.242 18192   192.168.0.244 35925  6 F.............. 1/1     -/-     1        0
  192.168.47.45    53   192.168.0.242 53806 17 F.............. 1/1     1/-     3        0
  192.168.0.244 62714   192.168.0.242 18192  6 F.............. 1/1     -/-     1        0

Idx Interface
--- ---------
  0 lo
  1 eth0
  2 eth1

Total number of connections: 30
[Expert@MyGW:0]#
```

fwaccel dbg

Description

The *fwaccel dbg* command controls the SecureXL debug. See *"SecureXL Debug Procedure" on page 297*.

Important:

- In a Cluster, you must configure all the Cluster Members in the same way.
- On Scalable Platforms, you must connect to the Gaia Portal of the applicable Security Group.
- If you do **not** use the complete debug procedure with the "`fw ctl kdebug`" command,
  and SecureXL works in Kernel Mode (KPPAK), then Security Gateway writes the debug outputs to these files:
  - `$FWDIR/log/fwk.elg` - Processing of traffic in the Firewall module
  - `/var/log/messages` - Additional information
- If you do **not** use the complete debug procedure with the "`fw ctl kdebug`" command,
  and SecureXL works in User Mode (UPPAK), then Security Gateway writes the debug outputs to these files:
  - `$FWDIR/log/fwk.elg` - Processing of traffic in the Firewall module
  - `/var/log/usim_x86.elg` - Processing of traffic in SecureXL
  - `/var/log/messages` - Processing of traffic in the ADP module (NVIDIA ConnectX 100G Cards)

Syntax in Gaia Clish or the Expert mode on a Security Gateway / ClusterXL:

```
fwaccel dbg
      -h
      -m <Name of SecureXL Debug Module>
      all
      + <Debug Flags>
      - <Debug Flags>
      reset
      -f {"<5-Tuple Debug Filter>" | reset}
      list
      resetall
```

Syntax in Gaia gClish on a Scalable Platform Security Group:

```
fwaccel dbg
      -h
      -m <Name of SecureXL Debug Module>
      all
      + <Debug Flags>
      - <Debug Flags>
      reset
      -f {"<5-Tuple Debug Filter>" | reset}
      list
      resetall
```

Syntax in the Expert mode on a Scalable Platform Security Group:

```
g_fwaccel dbg
      -h
      -m <Name of SecureXL Debug Module>
      all
      + <Debug Flags>
      - <Debug Flags>
      reset
      -f {"<5-Tuple Debug Filter>" | reset}
      list
      resetall
```

Parameters

Parameter	Description
`-h`	Shows the applicable built-in help.
`-m <Name of SecureXL Debug Module>`	Specifies the name of the SecureXL debug module. To see the list of available debug modules, run: `fwaccel dbg`
`all`	Enables all debug flags for the specified debug module.

Parameter	Description
`+ <Debug Flags>`	Enables the specified debug flags for the specified debug module: Syntax: `+ Flag1 [Flag2 Flag3 ... FlagN]` **Note** - You must press the space bar key after the plus (+) character.
`- <Debug Flags>`	Disables all debug flags for the specified debug module. Syntax: `- Flag1 [Flag2 Flag3 ... FlagN]` **Note** - You must press the space bar key after the minus (-) character.
`reset`	Resets all debug flags for the specified debug module to their default state.
`-f "<5-Tuple Debug Filter>"`	Configures the debug filter to show only debug messages that contain the specified connection. The filter is a string of five numbers separated with commas: `"<Source IP Address>,<Source Port>,<Destination IP Address>,<Destination Port>,<Protocol Number>"` **Notes:** ■ You can configure only one debug filter at one time. ■ You can use the asterisk "*" as a wildcard for an IP Address, Port number, or Protocol number. ■ For more information, see *IANA Service Name and Port Number Registry* and *IANA Protocol Numbers*.
`-f reset`	Resets the current debug filter.
`list`	Shows all enabled debug flags in all debug modules.
`resetall`	Reset all debug flags for all debug modules to their default state.

Enabling SecureXL debug flags during boot

From R81.20, you can configure SecureXL debug to start during boot.

Important - In a Cluster, you must configure all the Cluster Members in the same way.

Procedure

Step	Instructions
1	Connect to the command line on the Security Gateway / each Cluster Member. On Scalable Platforms, you must connect to the Gaia Portal of the applicable Security Group.
2	Log in to the Expert mode.
3	Create the require configuration files: ■ To collect the debug for IPv4 traffic: `touch $FWDIR/conf/fwaccel_dbg_flags.cfg` ■ To collect the debug for IPv6 traffic: `touch $FWDIR/conf/fwaccel6_dbg_flags.cfg`
4	Edit the applicable configuration file: ■ To collect the debug for IPv4 traffic: `vi $FWDIR/conf/fwaccel_dbg_flags.cfg` ■ To collect the debug for IPv6 traffic: `vi $FWDIR/conf/fwaccel6_dbg_flags.cfg`
5	Configure the applicable debug modules and the debug flags (see *"SecureXL Debug Modules and Debug Flags" on page 303*). Write each debug module and its flags on a separate line: `<Name of Debug Module #1> <Flag1> <Flag2> <Flag3> ... <FlagN>` `<Name of Debug Module #2> <Flag1> <Flag2> <Flag3> ... <FlagN>` Example: `default conn nat` `pkt drop nat`
6	Save the changes in the file and exit the editor.

Step	Instructions
7	On the Scalable Platform Security Group, you must copy the updated file to all Security Group Members: `asg_cp2blades $FWDIR/conf/fwaccel_dbg_flags.cfg` `asg_cp2blades $FWDIR/conf/fwaccel6_dbg_flags.cfg`
8	Reboot the Security Gateway / each Cluster Member / all Security Group Members.
9	Wait for the issue to occur.
10	Connect to the command line on the Security Gateway / each Cluster Member. On Scalable Platforms, you must connect to the Gaia Portal of the applicable Security Group.
11	Log in to the Expert mode.
12	Reset all the SecureXL debug flags in all SecureXL debug modules: ■ On the Security Gateway / each Cluster Member: `fwaccel dbg resetall` ■ On the Scalable Platform Security Group: `g_fwaccel dbg resetall`
13	Remove all entries from the configuration files: ■ On the Security Gateway / each Cluster Member: `echo '' > $FWDIR/conf/fwaccel_dbg_flags.cfg` `echo '' > $FWDIR/conf/fwaccel6_dbg_flags.cfg` ■ On the Scalable Platform Security Group: `echo '' > $FWDIR/conf/fwaccel_dbg_flags.cfg` `asg_cp2blades $FWDIR/conf/fwaccel_dbg_flags.cfg` `echo '' > $FWDIR/conf/fwaccel6_dbg_flags.cfg` `asg_cp2blades $FWDIR/conf/fwaccel6_dbg_flags.cfg`

Step	Instructions
14	Collect the output files. If SecureXL works in Kernel Mode (KPPAK): ■ `$FWDIR/log/fwk.elg` ■ `/var/log/messages` If SecureXL works in User Mode (UPPAK): ■ `$FWDIR/log/fwk.elg` ■ `/var/log/usim_x86.elg` ■ `/var/log/messages`

Examples

Example 1 - Default output

```
[Expert@MyGW:0]# fwaccel dbg
Usage: fwaccel dbg [-m <...>] [resetall | reset | list | all | +/- <flags>]
   -m <module>             - module of debugging
   -h                      - this help message
   resetall                - reset all debug flags for all modules
   reset                   - reset all debug flags for module
   all                     - set all debug flags for module
   list                    - list all debug flags for all modules
   -f reset | "<5-tuple>"  - filter debug messages
   + <flags>               - set the given debug flags
   - <flags>               - unset the given debug flags

Debug flags can be enabled at boot by adding them to
$FWDIR/conf/fwaccel_dbg_flags.cfg and $FWDIR/conf/fwaccel6_dbg_flags.cfg.
The file format is one line per module.

For example, to enable "tcp_state" and "routing" for the "pkt" module:
    echo "pkt tcp_state routing" >> $FWDIR/conf/fwaccel_dbg_flags.cfg

List of available modules and flags:

Module: default (default)
init drv tag lock cpdrv routing kdrv tcp_sv svm iter conn htab del update acct conf stat
queue ioctl corr util rngs relations ant conn_app rngs_print infra_ids offload nat

Module: db
get save del tmpl tmo init ant profile nmr nmt warning

Module: api
init add update del acct conf stat vpn notif tmpl sv pxl qos gtp infra tmpl_info upd_conf
upd_if_inf add_sa del_sa del_all_sas misc get_features get_tab get_stat reset_stat tag long_
ver del_all_tmpl get_state upd_link_sel

Module: pkt
f2f frag spoof acct notif tcp_state tcp_state_pkt sv cpls routing drop pxl qos user deliver
vlan pkt nat wrp corr caf bhm geneve sctp

Module: infras
reorder pm

Module: tmpl
dtmpl_get dtmpl_notif tmpl

Module: vpn
vpnpkt linksel routing vpn ls

Module: nac
db db_get pkt pkt_ex signature offload idnt ioctl nac

Module: cpaq
init client server exp cbuf opreg transport transport_utils broadcast

Module: synatk
init conf conn log pkt proxy state msg

Module: adp
rt nh eth heth wrp inf mbs bpl bplinf mbeinf if drop bond xmode ipsctl ac_print cpfifo qconf
qcomm filter packet mcast hw_offload hw_expn rte_api

Module: dos
fw1-cfg fw1-pkt sim-cfg sim-pkt detailed-pkt detailed-cfg drop cache

Module: gtp
pkt policy tables api drop notif general
```

```
Module: usdisp
error conn packet api msg state packet_err counter event quota ioctl lock clb uid queue
fwstats cachetab vpn temp_conns prio route dumbo

Module: dpdk_lib


Module: dpdk_pmd


Module: dpdk_other


[Expert@MyGW:0]#
```

Example 2 - Enabling and disabling of debug flags

```
[Expert@MyGW:0]# fwaccel dbg -m default + conn
Debug flags updated.
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dbg list

[Expert@MyGW:0]# fwaccel dbg list
 Module: default (400)
conn

Module: db (0)


Module: api (0)


Module: pkt (0)


Module: infras (0)


Module: tmpl (0)


Module: vpn (0)


Module: nac (0)


Module: cpaq (0)


Module: synatk (0)


Module: adp (0)


Module: dos (0)


Module: gtp (0)


Module: usdisp (0)


Module: dpdk_lib (0)


Module: dpdk_pmd (0)


Module: dpdk_other (0)


Debug filter not set.
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dbg -m default - conn
Debug flags updated.
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel dbg list

Module: default (0)


Module: db (0)


Module: api (0)


Module: pkt (0)


Module: infras (0)


Module: tmpl (0)


Module: vpn (0)


Module: nac (0)


Module: cpaq (0)


Module: synatk (0)


Module: adp (0)


Module: dos (0)


Module: gtp (0)


Module: usdisp (0)


Module: dpdk_lib (0)


Module: dpdk_pmd (0)


Module: dpdk_other (0)


Debug filter not set.
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dbg -m default reset
Debug flags updated.
[Expert@MyGW:0]#
```

Example 3 - Resetting all debug flags in all debug modules

```
[Expert@MyGW:0]# fwaccel dbg resetall
Debug state was reset to default.
[Expert@MyGW:0]#
```

Example 4 - Configuring debug filter for an SSH connection from 192.168.20.30 to 172.16.40.50

```
[Expert@MyGW:0]# fwaccel dbg -f 192.168.20.30,*,172.16.40.50,22,6
Debug filter was set.
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dbg list

... ...

Debug filter: "<*,*,*,*,*>"
[Expert@MyGW:0]#
```

fwaccel dos

Description

The "`fwaccel dos`" (for IPv4) and "`fwaccel6 dos`" (for IPv6) commands control the Rate Limiting for DoS mitigation techniques in SecureXL on the local Security Gateway, or Cluster Member.

See *"Rate Limiting for DoS Mitigation" on page 36*.

.

Important:

- In a Cluster, you must configure all the Cluster Members in the same way.
- In the VSNext mode / Traditional VSX mode, you must go to the context of an applicable Virtual Gateway / Virtual System.
  - In Gaia Clish, run: `set virtual-system <VSID>`
  - In the Expert mode, run: `vsenv <VSID>`
- On Scalable Platforms (ElasticXL, Maestro, Scalable Chassis), you must run the required commands only in this way:
  - On the Security Group command line, only on the SMO Security Group Member.
  - In the Global Gaia Clish (`gclish`), must run these commands:
    - `fwaccel dos <Options>`
    - `fwaccel6 dos <Options>`
  - In the Expert mode, must run these commands (start with the "`g_`" prefix):
    - `g_fwaccel dos <Options>`
    - `g_fwaccel6 dos <Options>`

Command Line

Platform	Command to Run	Description of Command
Security Gateway, ClusterXL, Traditional VSX Gateway, Traditional VSX Cluster	In Gaia Clish and the Expert mode: `fwaccel dos <Options>`	Manages all **IPv4** DoS / Rate Limiting features.
	In Gaia Clish and the Expert mode: `fwaccel6 dos <Options>`	Manages all **IPv6** DoS / Rate Limiting features.
ElasticXL Cluster, Security Group in Maestro, Security Group on Scalable Chassis, VSNext Virtual Gateway	In Gaia gClish: `fwaccel dos <Options>` In the Expert mode: `g_fwaccel dos <Options>`	Manages all **IPv4** DoS / Rate Limiting features.
	In Gaia gClish: `fwaccel6 dos <Options>` In the Expert mode: `g_fwaccel6 dos <Options>`	Manages all **IPv6** DoS / Rate Limiting features.

Syntax for IPv4

```
fwaccel dos
      config <options>
      deny <options>
      drop_frags <options>
      drop_opts <options>
      ioc_deny <options>
      ioc_deny_ext <options>
      ioc_monitor <options>
      ioc_monitor_ext <options>
      pbox <options>
      rate <options>
      stats <options>
```

Syntax for IPv6

```
fwaccel6 dos
      config <options>
      deny <options>
      drop_frags <options>
      drop_opts <options>
      ioc_deny <options>
      ioc_deny_ext <options>
      ioc_monitor <options>
      ioc_monitor_ext <options>
      pbox <options>
      rate <options>
      stats <options>
```

Parameters

Parameter	Description
`config <options>`	Shows the DoS mitigation configuration in SecureXL. See *"fwaccel dos config" on page 79*.
`deny <options>`	Controls the IP deny-list in SecureXL. See *"fwaccel dos deny" on page 84*.
`drop_frags <options>`	Controls the drop of IP Fragments in SecureXL. See *fwaccel dos drop_frags*.
`drop_opts <options>`	Controls the drop of IP Options in SecureXL. See *fwaccel dos drop_opts*.
`ioc_deny <options>`	Controls the IP addresses that are blocked by Threat Prevention IoC Feeds based on feed files.
`ioc_deny_ext <options>`	Controls the IP addresses that are blocked by Threat Prevention external IoC Feeds.
`ioc_monitor <options>`	Controls the IP addresses that are monitored by Threat Prevention IoC Feeds based on feed files.
`ioc_monitor_ext <options>`	Controls the IP addresses that are monitored by Threat Prevention external IoC Feeds.
`pbox <options>`	Controls the Penalty Box in SecureXL. See *"fwaccel dos pbox" on page 92*.

Parameter	Description
`rate <options>`	Controls the Rate Limiting policy in SecureXL. See *"fwaccel dos rate" on page 100*.
`stats <options>`	Shows and clears the DoS real-time statistics in SecureXL. See *"fwaccel dos stats" on page 122*.

fwaccel dos config

Description

The "`fwaccel dos config`" (for IPv4) and "`fwaccel6 dos config`" (for IPv6) commands show the global configuration parameters of the Rate Limiting for DoS mitigation in SecureXL.

These global parameters apply to all configured Rate Limiting rules.

Important:

- In a Cluster, you must configure all the Cluster Members in the same way.
- On Scalable Platforms, you must connect to the applicable Security Group.
  On Scalable Platforms (ElasticXL, Maestro, Scalable Chassis), you must run the required commands only in this way:
  - On the Security Group command line, only on the SMO Security Group Member.
  - In the Global Gaia Clish (`gclish`), must run these commands:
    - `fwaccel dos <Options>`
    - `fwaccel6 dos <Options>`
  - In the Expert mode, must run these commands (start with the "`g_`" prefix):
    - `g_fwaccel dos <Options>`
    - `g_fwaccel6 dos <Options>`
- In the VSNext mode / Traditional VSX mode, you must go to the context of an applicable Virtual Gateway / Virtual System.
  - In Gaia Clish, run: `set virtual-system <VSID>`
  - In the Expert mode, run: `vsenv <VSID>`

Syntax

```
{fwaccel | fwaccel6} dos config
      get
      set <options>
      reset-to-default
```

Parameters and Options

Parameter or Option	Description
No Parameters	Shows the applicable built-in usage.
`get`	Shows the configuration parameters.

Parameter or Option	Description
`set <options>`	This parameter is deprecated starting in R82. Use these commands: ■ *fwaccel dos deny* ■ *fwaccel dos drop_frags* ■ *fwaccel dos drop_opts* ■ *fwaccel dos pbox* ■ *fwaccel dos rate*
`reset-to-default`	Resets the configuration parameters to their default values. **Note** - This command does not affect the rate limit rules, IP values in deny-lists or allow-lists.

Example

```
[Expert@MyGW>:0]# fwaccel dos config get
Rate Limit Rules:
    Status                            on (without policy)
    Internal Interfaces               off
    Monitor-Only                      off
    Log Drops                         on
    Max Notifications Per-Second      100 logs/second
    Rule Cache                        on

Penalty Box:
    Status                            off
    Internal Interfaces               off
    Monitor-Only                      off
    Log Drops                         on
    Max Notifications Per-Second      100 logs/second
    Send TCP Reset                    off
    Timeout for Blocked IPs           180 seconds
    Has Blocked IPs                   no
    Log when a new IP is blocked      on
    Drop rate to trigger on           500 packets/second

Deny List:
    Status                            on (without policy)
    Internal Interfaces               off
    Monitor-Only                      off
    Log Drops                         on
    Max Notifications Per-Second      100 logs/second
    Send TCP Reset                    off
    Name                              Deny List

Disallow IPv4 Fragments:
    Status                            off
    Internal Interfaces               off
    Monitor-Only                      off
    Log Drops                         on
    Max Notifications Per-Second      100 logs/second

Disallow IP Options:
    Status                            off
    Internal Interfaces               off
    Monitor-Only                      off
    Log Drops                         on
    Max Notifications Per-Second      100 logs/second

IOC deny list (from files):
    Status                            on (without policy)
    Internal Interfaces               on
    Monitor-Only                      off
    Log Drops                         on
    Send TCP Reset                    off

IOC monitor-only list (from files):
    Status                            on (without policy)
    Internal Interfaces               on
    Monitor-Only                      on
    Log Drops                         on
    Send TCP Reset                    off

IOC deny list (from external feeds):
    Status                            on (without policy)
    Internal Interfaces               on
    Monitor-Only                      off
    Log Drops                         on
    Send TCP Reset                    off

IOC monitor-only list (from external feeds):
    Status                            on (without policy)
```

```
    Internal Interfaces               on
    Monitor-Only                      on
    Log Drops                         on
    Send TCP Reset                    off

[Expert@MyGW>:0]#
```

fwaccel dos deny

Description

The "`fwaccel dos deny`" (for IPv4) and "`fwaccel6 dos deny`" (for IPv6) commands control the IP deny-list in SecureXL.

The deny-list blocks all traffic to and from the specified IP addresses.

The deny-list drops occur in SecureXL, which is more efficient than an Access Control Policy to drop the packets.

Important:

- By default, the IP deny-list feature is enabled, without a Rate Limiting policy.
- By design, if you change the IP addresses in the Deny List with command line options and not through the corresponding files in the `$FWDIR/conf/deny_lists/` directory, then these changes do not survive a reboot.
- The Deny List scales up to millions of IP addresses.
- To enforce the IP deny-list in SecureXL, you must first enable the IP deny-lists. See these commands:
  - *"fwaccel dos config" on page 79*
  - *"fw sam_policy" on page 231* (configures more granular rules)
- In a Cluster, you must configure all the Cluster Members in the same way.
- On Scalable Platforms, you must connect to the applicable Security Group.
  On Scalable Platforms (ElasticXL, Maestro, Scalable Chassis), you must run the required commands only in this way:
  - On the Security Group command line, only on the SMO Security Group Member.
  - In the Global Gaia Clish (`gclish`), must run these commands:
    - `fwaccel dos <Options>`
    - `fwaccel6 dos <Options>`
  - In the Expert mode, must run these commands (start with the "`g_`" prefix):
    - `g_fwaccel dos <Options>`
    - `g_fwaccel6 dos <Options>`
- In the VSNext mode / Traditional VSX mode, you must go to the context of an applicable Virtual Gateway / Virtual System.
  - In Gaia Clish, run: `set virtual-system <VSID>`
  - In the Expert mode, run: `vsenv <VSID>`

Syntax

```
{fwaccel | fwaccel6} dos deny
      {-h | --help}
      allow
            {-h | --help}
            {-a | --add} <IP Address>[/<Subnet Mask Length>]
            {-d | --delete} <IP Address>[/<Subnet Mask Length>]
            {-F | --flush}
            {-l | --load} /<Path>/<Name of File>
            {-s | --show}
      {-a | --add} <IP Address>
      {-c | --show-config}
      {-d | --delete} <IP Address>
      {-E | --set-enabled} {on | off}
      {-F | --flush}
      {-G | --set-log-drops} {on | off}
      {-I | --set-enforce-internal} {on | off}
      {-l | --load} /<Path>/<Name of File>
      {-L | --load-default}
      {-M | --set-monitor-only} {on | off}
      {-N | --set-name} "<Name of IP Deny-list>"
      {-O | --set-notif-rate} <Number>
      {-R | --set-tcp-rst} {on | off}
      {-s | --show}
```

Parameters

Parameter	Description
No Parameters	Shows the applicable built-in usage.
`-h` `--help}`	Shows the applicable built-in usage.

Parameter	Description
`allow <options>`	Adds an IP address of a host or a network to a persistent "Allow List", so this IP address is not affected by the DoS / Rate Limiting protection: ■ `-h` `--help` Shows the applicable built-in usage. ■ `-a <IP Address>[/<Subnet Mask Length>]` `--add <IP Address>[/<Subnet Mask Length>]` Add an IP address in the CIDR notation to the override allow-list. • *`<IP Address>`* The IP address of a network or a host. • *`/<Subnet Mask Length>`* Must specify the length of the subnet mask from `/1` to `/32`. Optional for a host IP address. Mandatory for a network IP address. **Important** - If you do not specify the subnet mask length explicitly, this command uses the subnet mask length `/32`. ■ `-d <IP Address>[/<Subnet Mask Length>]` `--delete <IP Address>[/<Subnet Mask Length>]` Deletes an IP address in the CIDR notation from the override allow-list. ■ `-F` `--flush` Removes (flushes) all IP addresses from the override allow-list. ■ `-l /<Path>/<Name of File>` `--load /<Path>/<Name of File>` Loads the IP addresses into the override allow-list from the specified file. This file must contain IP addresses of hosts or networks in the CIDR notation, each IP address on a new line. ■ `-s` `--show` Shows the configured allow-list.
`-a <IP Address>` `--add <IP Address>`	Adds the specified IP address to the deny-list. **Note** - To add more than one IP address, run this command for each applicable IP address.

Parameter	Description
`-c` `--show-config`	Shows the current configuration.
`-d <IP Address>` `--delete <IP Address>`	Removes the specified IP addresses from the deny-list. **Note** - To remove more than one IP address, run this command for each applicable IP address.
`-E {on | off}` `--set-enabled {on | off}`	Enables (`on`) or disables (`off`) the feature. **Notes:** ■ By default, the IP deny-list feature is enabled without a Rate Limiting policy. ■ This change survives a reboot.
`-F` `--flush`	Removes (flushes) all IP addresses from the IP deny-list. **Notes:** ■ You can use this parameter "`{-F | --flush}`" with the parameter "`{-a | --add}`". ■ You can use this parameter "`{-F | --flush}`" with the parameter "`{-d | --delete}`". ■ You can use this parameter "`{-F | --flush}`" with the parameter "`{-l | --load}`".
`-G {on | off}` `--set-log-drops {on | off}`	Enables (`on`) or disables (`off`) the logging of packet drops. **Notes:** ■ By default, the Security Gateway generates the "Drop" logs for traffic that the DoS / Rate Limiting feature blocked. ■ By default, logging of packet drops is enabled.
`-I {on | off}` `--set-enforce-internal {on | off}`	Enables (`on`) or disables (`off`) the enforcement on interfaces, whose topology is configured as "Internal" in the Security Gateway object. **Notes:** ■ By default, DoS / Rate Limiting enforcement is disabled on interfaces, for which you configured the "Internal" topology in the Security Gateway / Cluster object. This is because the internal interfaces are assumed to be connected to trusted networks. ■ This change survives a reboot.

Parameter	Description
`-l /<Path>/<Name of File> --load /<Path>/<Name of File>`	Loads the IP addresses from the specified file. When dealing with large deny lists, the "`add`" command is cumbersome. Running a large number of "`add`" commands simultaneously (for example, with a shell script) can cause additional load on the Security Gateway's CPU. To configure large deny lists, it is better to add the list of IP addresses in a file, and then load the file in a single operation. **Notes:** ■ This file must contain IP addresses of hosts or networks in the CIDR notation, each IP address on a new line. ■ To add a comment line, it must start with the pound character "`#`". ■ The "`fwaccel`" command silently ignores all IPv6 addresses in the file. ■ The "`fwaccel6`" command silently ignores all IPv4 addresses in the file. ■ You may load multiple files at the same time.
`-L --load-default}`	Load all files from the `$FWDIR/conf/deny_lists/` directory into the IP deny-list. **Note** - The Security Gateway runs this command automatically during each boot.
`-M {on | off} --set-monitor-only {on | off}`	Enables (`on`) or disables (`off`) the monitor-only mode for the IP deny-list. In the monitor-only mode you can test the IP deny-list without blocking the traffic. The Security Gateway does not block traffic, but still generates a log. **Notes:** ■ By default, the monitor-only mode is disabled. ■ This change survives a reboot. ■ This command affects only the IP deny-list (does not affect the `fw samp` rules, etc.). ■ In addition to the Monitor-only mode, DoS / Rate Limiting has a more granular option to monitor packets on a rule-by-rule basis by specifying the action to be "`notify`" instead of the default action "`drop`".

Parameter	Description
`-N "<Name of IP Deny-list>" --set-name "<Name of IP Deny-list>"`	Configures the name for the IP deny-list. This name appears in the Security Gateway logs. **Notes:** ■ The default name is "`Deny List`". ■ This change survives a reboot. ■ Maximum name length is 79 characters. ■ You must use only ASCII characters.
`-O <Number> --set-notif-rate <Number>`	Configures the maximum number of logs per second for packet drops. When DoS / Rate Limiting blocks many packets, it can be important to limit the maximum number of the drop logs that the Security Gateway generates per second. **Notes:** ■ The default logging rate is 100 logs/second. ■ This change survives a reboot.
`-R {on | off} --set-tcp-rst {on | off}`	Enables (`on`) or disables (`off`) the response with the TCP [RST] packet for TCP connections that the IP deny-list blocked. **Notes:** ■ By default, SecureXL does not send the TCP [RST] packet for blocked TCP connections. ■ This change survives a reboot.
`-s --show`	Shows the IP addresses in the IP deny-list.

Example from a non-VSX Gateway

```
[Expert@MyGW:0]# fwaccel dos deny -c
Deny List:
    Status                          on (without policy)
    Internal Interfaces             off
    Monitor-Only                    off
    Log Drops                       on
    Max Notifications Per-Second    100 logs/second
    Send TCP Reset                  off
    Name                            Deny List
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dos deny -s
The deny list is empty
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dos deny -a 1.1.1.1
Adding 1.1.1.1
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dos deny -s
1.1.1.1
[Expert@MyGW:0]# fwaccel dos deny -a 2.2.2.2
Adding 2.2.2.2
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dos deny -s
2.2.2.2
1.1.1.1
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dos deny -d 2.2.2.2
Deleting 2.2.2.2
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dos deny -s
1.1.1.1
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dos deny -F
All deny list entries deleted
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dos deny -s
The deny list is empty
[Expert@MyGW:0]#
```

fwaccel dos pbox

Description

The "`fwaccel dos pbox`" (for IPv4) and "`fwaccel6 dos pbox`" (for IPv6) commands control the Penalty Box deny-list in SecureXL.

The SecureXL Penalty Box is a mechanism that performs an early drop of packets that arrive from suspected sources. The purpose of this feature is to allow the Security Gateway to cope better under high traffic load, possibly caused by a DoS/DDoS attack.

The SecureXL Penalty Box detects clients that send packets, which the Access Control Policy drops, and clients that violate the IPS protections. If the SecureXL Penalty Box detects a specific client frequently, it puts that client in a penalty box. From that point, SecureXL drops all packets that arrive from the blocked source IP address.

The Penalty Box allow-list in SecureXL configures the source IP addresses, which the SecureXL Penalty Box never blocks.

Important:

- By default, the Penalty Box is disabled.
- To enforce the Penalty Box in SecureXL, you must first enable the Penalty Box. See:
  - *"fwaccel dos config" on page 79*
- In a Cluster, you must configure all the Cluster Members in the same way.
- On Scalable Platforms, you must connect to the applicable Security Group.
  On Scalable Platforms (ElasticXL, Maestro, Scalable Chassis), you must run the required commands only in this way:
  - On the Security Group command line, only on the SMO Security Group Member.
  - In the Global Gaia Clish (`gclish`), must run these commands:
    - `fwaccel dos <Options>`
    - `fwaccel6 dos <Options>`
  - In the Expert mode, must run these commands (start with the "`g_`" prefix):
    - `g_fwaccel dos <Options>`
    - `g_fwaccel6 dos <Options>`
- In the VSNext mode / Traditional VSX mode, you must go to the context of an applicable Virtual Gateway / Virtual System.
  - In Gaia Clish, run: `set virtual-system <VSID>`
  - In the Expert mode, run: `vsenv <VSID>`
- In a Scalable Platform, when you add a new Security Group Member to a Security Group, the new Security Group Member pulls these configuration files:
  - `$FWDIR/conf/pbox-whitelist-v4.conf`
  - `$FWDIR/conf/pbox-whitelist-v6.conf`

Syntax

```
{fwaccel | fwaccel6} dos pbox
      {-h | --help}
      allow
            {-h | --help}
            {-a | --add} <IP Address>[/<Subnet Mask>]
            {-d | --delete} <IP Address>[/<Subnet Mask>]
            {-F | --flush}
            {-l | --load} /<Path>/<Name of File>
            {-s | --show}
      {-c | --show-config}
      {-E | --set-enabled} {on | off}
      {-F | --flush}
      {-G | --set-log-drops} {on | off}
      {-I | --set-enforce-internal} {on | off}
      {-L | --set-log-reported} {on | off}
      {-M | --set-monitor-only} {on | off}
      {-O | --set-notif-rate} <Number>}
      {-P | --set-drops-threshold} <Number>}
      {-R | --set-tcp-rst} {on | off}
      {-s | --show}
      {-T | --set-timeout} <Number>}
```

Parameters

Parameter	Description
No Parameters	Shows the applicable built-in usage.
`-h` `--help}`	Shows the applicable built-in usage.

Parameter	Description
`allow <options>`	**Important** - Before you use a 3rd-party or automatic blacklists, add trusted networks and hosts to this allow-list to avoid outages. Adds an IP address of a host or a network to a persistent "Allow List", so this IP address is not affected by the DoS / Rate Limiting protection: ■ `-h` `--help` Shows the applicable built-in usage. ■ `-a <IP Address>[/<Subnet Mask Length>]` `--add <IP Address>[/<Subnet Mask Length>]` Add an IP address in the CIDR notation to the override allow-list. • `<IP Address>` The IP address of a network or a host. • `/<Subnet Mask Length>` Must specify the length of the subnet mask from `/1` to `/32`. Optional for a host IP address. Mandatory for a network IP address. **Important** - If you do not specify the subnet mask length explicitly, this command uses the subnet mask length `/32`. ■ `-d <IP Address>[/<Subnet Mask Length>]` `--delete <IP Address>[/<Subnet Mask Length>]` Deletes an IP address in the CIDR notation from the override allow-list. ■ `-F` `--flush` Removes (flushes) all IP addresses from the override allow-list. ■ `-l /<Path>/<Name of File>` `--load /<Path>/<Name of File>` Loads the IP addresses into the override allow-list from the specified file. This file must contain IP addresses of hosts or networks in the CIDR notation, each IP address on a new line. ■ `-s` `--show` Shows the configured allow-list.
`-a <IP Address>` `--add <IP Address>`	Adds the specified IP address to the deny-list. **Note** - To add more than one IP address, run this command for each applicable IP address.

Parameter	Description
`-c` `--show-config`	Shows the current configuration.
`-d <IP Address>` `--delete <IP Address>`	Removes the specified IP addresses from the deny-list. **Note** - To remove more than one IP address, run this command for each applicable IP address.
`-E {on | off}` `--set-enabled {on | off}`	Enables (`on`) or disables (`off`) the feature. **Notes:** ■ By default, the Penalty Box is disabled. ■ This change survives a reboot.
`-F` `--flush`	Removes (flushes) all IP addresses from the Penalty Box. **Notes:** ■ You can use this parameter "`{-F | --flush}`" with the parameter "`{-a | --add}`". ■ You can use this parameter "`{-F | --flush}`" with the parameter "`{-d | --delete}`". ■ You can use this parameter "`{-F | --flush}`" with the parameter "`{-l | --load}`".
`-G {on | off}` `--set-log-drops {on | off}`	Enables (`on`) or disables (`off`) the logging of packet drops. **Notes:** ■ By default, the Security Gateway generates the "Drop" logs for traffic that the DoS / Rate Limiting feature blocked. ■ By default, logging of packet drops is enabled.
`-I {on | off}` `--set-enforce-internal {on | off}`	Enables (`on`) or disables (`off`) the enforcement on interfaces, whose topology is configured as "Internal" in the Security Gateway object. **Notes:** ■ By default, DoS / Rate Limiting enforcement is disabled on interfaces, for which you configured the "Internal" topology in the Security Gateway / Cluster object. This is because the internal interfaces are assumed to be connected to trusted networks. ■ This change survives a reboot.

Parameter	Description
`-L {on | off} --set-log-reported {on | off}`	Enables (`on`) or disables (`off`) the logging of IP addresses that were added to the Penalty Box. **Notes:** ■ By default, logging of IP addresses that were added to the Penalty Box is enabled. ■ This change survives a reboot.
`-M {on | off} --set-monitor-only {on | off}`	Enables (`on`) or disables (`off`) the monitor-only mode for the IP deny-list. In the monitor-only mode you can test the IP deny-list without blocking the traffic. The Security Gateway does not block traffic, but still generates a log. **Notes:** ■ By default, the monitor-only mode is disabled. ■ This change survives a reboot. ■ This command affects only the IP deny-list (does not affect the `fw samp` rules, etc.). ■ In addition to the Monitor-only mode, DoS / Rate Limiting has a more granular option to monitor packets on a rule-by-rule basis by specifying the action to be "`notify`" instead of the default action "`drop`".
`-O <Number> --set-notif-rate <Number>`	Configures the maximum number of logs per second for packet drops. When DoS / Rate Limiting blocks many packets, it can be important to limit the maximum number of the drop logs that the Security Gateway generates per second. **Notes:** ■ The default logging rate is 100 logs/second. ■ This change survives a reboot.
`-P <Number> --set-drops-threshold <Number>`	Configures the minimum number of dropped packets per second from a source to trigger the Penalty Box for that source. **Notes:** ■ The default drop rate is 500 packets/second. ■ This change survives a reboot.

Parameter	Description
`-R {on | off}` `--set-tcp-rst {on | off}`	Enables (`on`) or disables (`off`) the response with the TCP [RST] packet for TCP connections that the IP deny-list blocked. **Notes:** ■ By default, SecureXL does not send the TCP [RST] packet for blocked TCP connections. ■ This change survives a reboot.
`-s` `--show`	Shows the IP addresses in the IP deny-list.
`-T <Number>` `--set-timeout <Number>`	Configures the timeout (in seconds) for blocked IP addresses in the Penalty Box. After this timeout reaches 0, SecureXL removes the blocked IP address from the Penalty Box. **Notes:** ■ The default timeout is 180 seconds. ■ This change survives a reboot.

Example 1 - Default Configuration

```
[Expert@MyGW:0]# fwaccel dos pbox -c
Penalty Box:
    Status                          off
    Internal Interfaces             off
    Monitor-Only                    off
    Log Drops                       on
    Max Notifications Per-Second    100 logs/second
    Send TCP Reset                  off
    Timeout for Blocked IPs         180 seconds
    Has Blocked IPs                 no
    Log when a new IP is blocked    on
    Drop rate to trigger on         500 packets/second
[Expert@MyGW:0]#
```

Example 2 - Adding a host IP address without the optional subnet mask length

```
[Expert@MyGW:0]# fwaccel dos pbox allow -a 192.168.20.40
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dos pbox allow -s
192.168.20.40/32
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dos pbox allow -F
[Expert@MyGW:0]# fwaccel dos pbox allow -s
[Expert@MyGW:0]#
```

Example 3 - Adding a host IP address with the optional subnet mask length

```
[Expert@MyGW:0]# fwaccel dos pbox allow -a 192.168.20.40/32
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dos pbox allow -s
192.168.20.40/32
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dos pbox allow -F
[Expert@MyGW:0]# fwaccel dos pbox allow -s
[Expert@MyGW:0]#
```

Example 4 - Adding a network IP address with the mandatory subnet mask length

```
[Expert@MyGW:0]# fwaccel dos pbox allow -a 192.168.20.0/24
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dos pbox allow -s
192.168.20.0/24
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dos pbox allow -F
[Expert@MyGW:0]# fwaccel dos pbox allow -s
[Expert@MyGW:0]#
```

Example 5 - Deleting a host entry

```
[Expert@MyGW:0]# fwaccel dos pbox allow -a 192.168.20.40/32
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dos pbox allow -a 192.168.20.70/32
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dos pbox allow -s
192.168.20.40/32
192.168.20.70/32
[Expert@MyGW:0]# fwaccel dos pbox allow -d 192.168.20.70/32
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dos pbox allow -s
192.168.20.40/32
[Expert@MyGW:0]#
```

fwaccel dos rate

Description

The "`fwaccel dos rate`" (for IPv4) and "`fwaccel6 dos rate`" (for IPv6) commands show and install the Rate Limiting policy in SecureXL.

Important:

- By default, the Rate Limiting policy feature is enabled without any rules.
- In a Cluster, you must configure all the Cluster Members in the same way.
- On Scalable Platforms, you must connect to the applicable Security Group.
  On Scalable Platforms (ElasticXL, Maestro, Scalable Chassis), you must run the required commands only in this way:
  - On the Security Group command line, only on the SMO Security Group Member.
  - In the Global Gaia Clish (`gclish`), must run these commands:
    - `fwaccel dos <Options>`
    - `fwaccel6 dos <Options>`
  - In the Expert mode, must run these commands (start with the "`g_`" prefix):
    - `g_fwaccel dos <Options>`
    - `g_fwaccel6 dos <Options>`
- In the VSNext mode / Traditional VSX mode, you must go to the context of an applicable Virtual Gateway / Virtual System.
  - In Gaia Clish, run: `set virtual-system <VSID>`
  - In the Expert mode, run: `vsenv <VSID>`

Notes

- If you install a new rate limiting policy with more than one rule, it automatically enables the rate limiting feature.

  To disable the rate limiting feature manually, run this command (see *"fwaccel dos config" on page 79*):

  ```
  {fwaccel | fwaccel6} dos config set --disable-rate-limit
  ```

- To delete the current rate limiting policy, install a new policy with zero rules.

Syntax

```
{fwaccel | fwaccel6} dos rate {-h | --help}
```

```
{fwaccel | fwaccel6} dos rate
      add --help
      add [<SIC Connection>] <Options> <Match Conditions> <Limit>
<Tracking>
      add batch /<Path>/<Name of File>
```

```
{fwaccel | fwaccel6} dos rate
      counters --help
      counters
      counters '<Rule Index>'
      counters '<Rule UID>'
```

```
{fwaccel | fwaccel6} dos rate
      del --help
      del all
      del '<Rule UID>'
      del batch /<Path>/<Name of File>
```

```
{fwaccel | fwaccel6} dos rate
      get --help
      get {-r | -l | -o /<Path>/<Name of File>}
      get [-n] -u '<Rule UID>'
      get [-n] -k <Search Key>
      get [-n] -v <Search Value>
      get --show-tab
      get --show-counters
```

```
{fwaccel | fwaccel6} dos rate
      {-c | --show-config}
      {-E | --set-enabled} {on | off}
      {-G | --set-log-drops} {on | off}
      {-I | --set-enforce-internal} {on | off}
      {-M | --set-monitor-only} {on | off}
      {-O | --set-notif-rate} <Number>}
      {-R | --set-rule-cache} {on | off}
```

Parameters

Parameter	Description
No Parameters	Shows the applicable built-in usage.
`-h` `--help}`	Shows the applicable built-in usage.
`add [<SIC Connection>] <Options> <Match Conditions> <Limit> <Tracking>`	Adds one IPv4 rule to the Rate Limiting policy. See *"Adding an IPv4 rule to the Rate Limiting policy" on page 104*.
`add batch /<Path>/<Name of File>`	Adds IPv4 rules in a batch mode to the Rate Limiting policy. See *"Adding IPv4 rules in a batch mode to the Rate Limiting policy" on page 116*.
`counters <options>`	Shows the counters for all rules in the Rate Limiting policy. See *"Viewing DoS / Rate Limiting Counters" on page 119*. To reset counters, see *fwaccel dos stats*.
`del all`	Deletes all rules from the Rate Limiting policy.
`del '<Rule UID>'`	Deletes the specified rule from the Rate Limiting policy. **Important** - The quote marks (single or double) and angle brackets (`'<...>'`) are mandatory. The '*<Rule UID>*' value is generated automatically when you add a rule. To see the UID for each rule, run: `{fwaccel | fwaccel6} dos rate get` Example: `fwaccel dos rate del '<5779378c,00000000,64291eac,00005584>'`
`del batch /<Path>/<Name of File>`	Deletes the specified rules in a batch mode from the Rate Limiting policy. 1. Get the UID for each rule: `{fwaccel | fwaccel6} dos rate get` 2. Save the required UID value in a plain-text file, one UID on each line. 3. Delete the specified rules: `{fwaccel | fwaccel6} del batch /<Path>/<Name of File>`

Parameter	Description
`get <options>`	Shows information about the rules the Rate Limiting policy. See *"Viewing information about the rules the Rate Limiting policy" on page 116*.
`-c` `--show-config`	Shows the current configuration. Example: `[Expert@MyGW:0]# fwaccel dos rate -c` `Rate Limit Rules:` `  Status  on (without policy)` `  Internal Interfaces  off` `  Monitor-Only  off` `  Log Drops  on` `  Max Notifications Per-Second  100 logs/second` `  Rule Cache  on` `[Expert@MyGW:0]#`
`-E {on | off}` `--set-enabled {on | off}`	Enables (`on`) or disables (`off`) the feature. **Notes:** ■ By default, the Rate Limiting policy feature is enabled without any rules. ■ This change survives a reboot.
`-G {on | off}` `--set-log-drops {on | off}`	Enables (`on`) or disables (`off`) the logging of packet drops. **Notes:** ■ By default, the Security Gateway generates the "Drop" logs for traffic that the DoS / Rate Limiting feature blocked. ■ By default, logging of packet drops is enabled.
`-I {on | off}` `--set-enforce-internal {on | off}`	Enables (`on`) or disables (`off`) the enforcement on interfaces, whose topology is configured as "Internal" in the Security Gateway object. **Notes:** ■ By default, DoS / Rate Limiting enforcement is disabled on interfaces, for which you configured the "Internal" topology in the Security Gateway / Cluster object. This is because the internal interfaces are assumed to be connected to trusted networks. ■ This change survives a reboot.

Parameter	Description
`-M {on | off}` `--set-monitor-only {on | off}`	Enables (`on`) or disables (`off`) the monitor-only mode for the IP deny-list. In the monitor-only mode you can test the drops of IP Fragments without blocking the traffic. The Security Gateway does not block traffic, but still generates a log. **Notes:** ■ By default, the monitor-only mode is disabled. ■ This change survives a reboot. ■ This command affects only the IP deny-list (does not affect the `fw samp` rules, etc.). ■ In addition to the Monitor-only mode, DoS / Rate Limiting has a more granular option to monitor packets on a rule-by-rule basis by specifying the action to be "`notify`" instead of the default action "`drop`".
`-O <Number>` `--set-notif-rate <Number>`	Configures the maximum number of logs per second for packet drops. When DoS / Rate Limiting blocks many packets, it can be important to limit the maximum number of the drop logs that the Security Gateway generates per second. **Notes:** ■ The default logging rate is 100 logs/second. ■ This change survives a reboot.
`-R {on | off}` `--set-rule-cache {on | off}`	Enables (`on`) or disables (`off`) the rule cache. **Notes:** ■ By default, the rule cache is enabled for maximum performance. ■ This change survives a reboot.

Adding an IPv4 rule to the Rate Limiting policy

Exaplanation

Syntax

```
{fwaccel | fwaccel6} dos rate add [<SIC Connection>] <Options>
<Match Conditions> <Limit> <Tracking>
```

Parameters

<SIC Connection>

You can use this optional parameter only when running the command from the Management Server's command-line to configure DoS / Rate Limiting policy rules on a Security Gateway / each Cluster Member.

Use on of these:

- `-S <IP Address of Management Interface>`

  Specifies the IP address of the management interface (default=localhost) on a Security Gateway / Cluster Member.

  Example:

  ```
  fwaccel dos rate add -S 192.168.1.101 source any pkt-rate
  100000
  ```

- `-s <SIC Name>`

  Specifies the SIC Name of a VSNext Virtual Gateway / Traditional VSX Gateway / Traditional VSX Cluster Member.

  Example:

  ```
  fwaccel dos rate add -S 192.168.1.101 -s my_virtual_system_1
  source any pkt-rate 100000
  ```

<Options>

These parameters affect the general behavior of a given policy rule.

Important - If you would like to use these options, you must specify them in the CLI syntax in front of the "<*Match Conditions*>" or the "<*Limit*>".

Parameter	Description
`-t <Timeout>`	Automatically deletes the rule after the specified number of seconds. Default value: None Example: `fwaccel dos rate add -t 2 source any pkt-rate 100000`

Parameter	Description
`-a {d | n | b}`	Specifies the action. One of these: ■ `-a d` `-action drop` Drop all packets that violate this rule (this is the default). ■ `-a n` `-action notify` Notify, but do not drop packets that violate this rule (meaning, Monitor-only mode). ■ `-a b` `-action bypass` Bypass DoS / Rate Limiting policy for all packets that match this rule (meaning, Allow List). Limits and the "track" option are not relevant for bypass rules and should not be specified. Example: `fwaccel dos rate add -a b source cidr:192.168.0.0/16`
`-l {r | a}`	Specifies the Security Gateway log type: ■ `-l r` `-log regular` Regular log (this is the default). ■ `-l a` `-log alert` Alert. Example: `fwaccel dos rate add -l a source any pkt-rate 100000`
`-n '<`*`Rule Name`*`>'`	Specifies the rule name. **Important** - The quote marks (single or double) are mandatory. Example: `fwaccel dos rate add -n "This is a Rule Name" source any pkt-rate 100000`

Parameter	Description
`-c '<Comment>'`	Specifies the tule comment. **Important** - The quote marks (single or double) are mandatory. Example: `fwaccel dos rate add -c "This is a Comment" source any pkt-rate 100000`
`-o '<Originator>'`	Specifies the rule originator. **Important** - The quote marks (single or double) are mandatory. Example: `fwaccel dos rate add -o "John Doe" source any pkt-rate 100000`
`-f <Security Gateway>`	Specifies the Security Gateway on which to enforce this rule. Typically this option is only used when administrating rules remotely. Valid values are: ■ `"all"` Enforce this rule on all Security Gateways (this is the default). ■ *<Name of Security Gateway Object>* Enforce this rule on the specified Security Gateway. ■ *<Name of Network Group Object>* Enforce this rule on all Security Gateways listed in the specified Network Group object. Example: `fwaccel dos rate add -S 192.168.1.101 -f test_gateway_1 source any pkt-rate 100000`

<Match Conditions>

Rate limiting policy rules can specify 3 different types of match conditions:

- Source IP address
- Destination IP address
- IP protocol, or IP Protocol + destination port

If a type of match condition is not specified, it defaults to "any", but at least one source or destination match condition must be included in the rule.

If a source match condition or a destination match condition does not have a format specified, it defaults to "`cidr:`" or "`range:`".

Multiple match conditions can be combined to match specific network traffic.

The 3 different types of match conditions are described below:

- Source Match Condition
- Destination Match Condition
- Service Match Condition

Source Match Condition

The source match conditions are used to specify which source IP addresses match the rule.

The default value is "`any`".

Source	Description
`any`	Any/all source IP addresses match this DoS / Rate Limiting rule. This is the default if no source match condition is specified. Example: `fwaccel dos rate add source any pkt-rate 100000`
`range:<`*`MIN`*`>[-<`*`MAX`*`>]`	Specifies a range of IP addresses. If <*MAX*> is not specified, then the range represents a single IP address. Example: `fwaccel dos rate add source range:10.0.0.1-10.0.0.254 pkt-rate 100000`
`cidr:<`*`ADDRESS`*`>[/<`*`MASK`*`>]`	Specifies a network using the CIDR notation. If the mask parameter is not provided, then the expression represents a specific host address. Example: `fwaccel dos rate add source cidr:10.0.0.0/24 service 6/22 pkt-rate 100000`

Source	Description
`cc:<COUNTRY_CODE>`	**Note** - This parameter is not supported (Known Limitation PMTR-87460). Two-letter code defined in ISO 3166-1 alpha-2. The rule matches the country code to the IP addresses assigned to this country, based on the Geo IP database. Example: `fwaccel dos rate add source cc:US pkt-rate 100000`
`asn:<AUTONOMOUS_SYSTEM_NUMBER>`	**Note** - This parameter is not supported (Known Limitation PMTR-87460). Valid value syntax is AS`nnnn`, where `nnnn` is a number unique to the specific organization. The rule matches the AS number of the organization to the IP addresses that are assigned to this organization, based on the Geo IP database. Example: `fwaccel dos rate add source asn:AS1234 pkt-rate 100000`
`source-negated {true | false}`	If not specified, the default is "`false`". When set to "`true`", it inverts the match condition such that the rule matches all source IP addresses except for the given value. Example: `fwaccel dos rate add source-negated true source cidr:10.0.0.0/8 pkt-rate 100000`

Destination Match Condition

The destination match conditions are used to specify which destination IP addresses match the rule.

The default value is "`any`".

Destination	Description
`any`	Any/all source IP addresses match this DoS / Rate Limiting rule. This is the default if no source match condition is specified. Example: `fwaccel dos rate add destination any pkt-rate 100000`
`range:<`*`MIN`*`>[-<`*`MAX`*`>]`	Specifies a range of IP addresses. If *<MAX>* is not specified, then the range represents a single IP address. Example: `fwaccel dos rate add destination range:10.0.0.1-10.0.0.254 pkt-rate 100000`
`cidr:<`*`ADDRESS`*`>[/<`*`MASK`*`>]`	Specifies a network using the CIDR notation. If the mask parameter is not provided, then the expression represents a specific host address. Example: `fwaccel dos rate add destination cidr:10.0.0.0/24 service 6/22 pkt-rate 100000`
`cc:<`*`COUNTRY_CODE`*`>`	**Note** - This parameter is not supported (Known Limitation PMTR-87460). Two-letter code defined in ISO 3166-1 alpha-2. The rule matches the country code to the IP addresses assigned to this country, based on the Geo IP database. Example: `fwaccel dos rate add destination cc:US pkt-rate 100000`
`asn:<`*`AUTONOMOUS_SYSTEM_NUMBER`*`>`	**Note** - This parameter is not supported (Known Limitation PMTR-87460). Valid value syntax is AS`nnnn`, where `nnnn` is a number unique to the specific organization. The rule matches the AS number of the organization to the IP addresses that are assigned to this organization, based on the Geo IP database. Example: `fwaccel dos rate add destination asn:AS1234 pkt-rate 100000`

Destination	Description
`source-negated {true | false}`	If not specified, the default is "`false`". When set to "`true`", it inverts the match condition such that the rule matches all source IP addresses except for the given value. Example: `fwaccel dos rate add destination-negated true destination cidr:10.0.0.0/8 pkt-rate 100000`

Service Match Condition

The service match conditions are used to specify which network services (meaning, `<ipproto>:<port>`) match the rule.

The default value is any protocol, any port.

Service	Description
`any`	All IP protocols and ports match this DoS / Rate Limiting rule. This is the default if no service match condition is specified. Example: `fwaccel dos rate add source cidr:10.0.0.0/8 service any pkt-rate 100000`
<PROTO>	IANA IP Protocol number. See *IANA Protocol Numbers*. For example, "6" for TCP. Causes this rule to apply only to the given protocol . Valid values are 1 through 255. Example: `fwaccel dos rate add source cidr:10.0.0.0/8 service 1 pkt-rate 100000`
<MIN_PROTO>-*<MAX_PROTO>*	For example, "1-6" would specify all IP protocols with numbers between 1 and 6. Causes this rule to match all protocol numbers within the given range. Example: `fwaccel dos rate add source cidr:10.0.0.0/8 service 1-6 pkt-rate 100000`

Service	Description
`<PROTO>/<PORT>`	Only valid for the TCP and UDP protocols. Causes the rule to match the specific IP protocol and destination port number. Example: `fwaccel dos rate add source cidr:10.0.0.0/8 service 6/22 pkt-rate 100000`
`<PROTO>/<MIN_PORT>-<MAX_PORT>`	Only valid for the TCP and UDP protocols. Causes the rule to match the specific IP protocol, and range of destination port numbers. Example: `fwaccel dos rate add source cidr:10.0.0.0/8 service 6/0-1024 pkt-rate 100000`
`service-negated {true | false}`	If not specified, the default is "`false`". When set to "`true`", it inverts the match condition such that the rule matches all services/ports except for the given value. Example: `fwaccel dos rate add source cidr:10.0.0.0/8 service-negated true service 6/0-1024 pkt-rate 100000`

<Limit>

The limit defines the trigger threshold for a policy rule (for example, maximum allowable packets-per-second).

Only a single limit is allowed for each rule.

Limit	Description
`concurrent-conns`	Maximum number of simultaneously active connections allowed by this rule. Example: `fwaccel dos rate add source cidr:10.0.0.0/8 service 6/22 concurrent-conns 5000`

Limit	Description
`concurrent-conns-ratio`	Limit the maximum number of connections to a percentage of the total number of active connections passing through the Security Gateway. The value is expressed as a ratio in parts-per-65536. For example, a value of 6553, would allow a ratio of about 10%. **Note** - To prevent false positives, a global minimum of 1000 connections must be simultaneously active (within the specific Security Gateway) before this limit will activate. Example: `fwaccel dos rate add source cidr:10.0.0.0/8 service 6/22 concurrent-conns-ratio 6553`
`new-conn-rate`	Maximum number of connections-per-second allowed by this rule. The minimum value is 1. Example: `fwaccel dos rate add source cidr:10.0.0.0/8 service 6/22 new-conn-rate 100`
`new-conn-rate-ratio`	Limit the maximum connections-per-second to a percentage of the total connections-per-second passing through the Security Gateway. The value is expressed as a ratio in parts-per-65536. For example, a value of 6553, would allow a ratio of about 10%. **Note** - To prevent false positives, a global minimum of 100 new connections-per-second must flow through the Security Gateway before this limit will activate. Example: `fwaccel dos rate add source cidr:10.0.0.0/8 service 6/22 new-conn-rate-ratio 6553`
`pkt-rate`	Maximum number of packets-per-second allowed by this rule. A value of zero will block all traffic that matches this rule. Example: `fwaccel dos rate add source cidr:10.0.0.0/8 service 6/22 pkt-rate 5000`

Limit	Description
`pkt-rate-ratio`	Limit the maximum packets-per-second to a percentage of the total packets-per-second passing through the Security Gateway. The value is expressed as a ratio in parts-per-65536. For example, a value of 6553, would allow a ratio of about 10%. **Note** - To prevent false positives, a global minimum of 1000 packets-per-second must flow through the Security Gateway before this limit will activate. Example: `fwaccel dos rate add source cidr:10.0.0.0/8 service 6/22 pkt-rate-ratio 6553`
`byte-rate`	Maximum number of bytes-per-second allowed by this rule. A value of zero will block all traffic that matches this rule. Example: `fwaccel dos rate add source cidr:10.0.0.0/8 service 6/22 byte-rate 500000`
`byte-rate-ratio`	Limit the maximum bytes-per-second to a percentage of the total bytes-per-second passing through the Security Gateway. The value is expressed as a ratio in parts-per-65536. For example, a value of 6553, would allow a ratio of about 10%. **Note** - To prevent false positives, a global minimum of 100000 bytes-per-second must flow through the Security Gateway before this limit will activate. Example: `fwaccel dos rate add source cidr:10.0.0.0/8 service 6/22 byte-rate-ratio 6553`

<Tracking>

The tracking mode controls how the limits are applied to traffic that matches the rule.

Tracking Mode	Description
Default (no tracking)	The all traffic matching the rule is blocked if the rule is violated Example: `fwaccel dos rate add source cidr:10.0.0.0/8 service 6/22 new-conn-rate 100`

Tracking Mode	Description
`source`	Each source IP address is tracked separately. If traffic from a given source IP address violates this rule, only that source IP address is blocked. Also, one source IP will not affect the rate limit counters of another source IP. Example: `fwaccel dos rate add source cidr:10.0.0.0/8 service 6/22 new-conn-rate 100 track source`
`source-service`	Each source IP address plus service is tracked separately. If traffic from a given source IP and service violates this rule, only that specific service and source IP are blocked. Other services from the same source IP are unaffected by the rule (unless they violate the rule also). Example: `fwaccel dos rate add source cidr:10.0.0.0/8 service 6/22 new-conn-rate 100 track source-service`
`destination`	Each destination IP address is tracked separately. If traffic to a given destination IP address violates this rule, only that destination IP address is blocked. Also, one destination IP will not affect the rate limit counters of another destination IP. Example: `fwaccel dos rate add source cidr:10.0.0.0/8 service 6/22 new-conn-rate 100 track destination`
`destination-service`	Each destination IP address plus service is tracked separately. If traffic to a given destination IP and service violates this rule, only that specific service and destination IP are blocked. Other services for the same destination IP are unaffected by the rule (unless they violate the rule also). Example: `fwaccel dos rate add source cidr:10.0.0.0/8 service 6/22 new-conn-rate 100 track destination-service`

Adding IPv4 rules in a batch mode to the Rate Limiting policy

Procedure

1. Add the list of rules to a plain-text file (one rule on each line).

   Remove the "`fwaccel dos rate add`" command prefix from each rule.

   Example:

   ```
   service any source cidr:192.168.2.0/24 pkt-rate 1000
   service any source cidr:192.168.2.17 concurrent-conns 150
   ```

2. Use the batch command to load the file:

   ```
   {fwaccel | fwaccel6} dos rate add batch /<Path>/<Name of
   File>
   ```

   Example:

   ```
   fwaccel dos rate add batch /var/log/rule_config.txt
   ```

Viewing information about the rules the Rate Limiting policy

Explanation

Syntax:

```
{fwaccel | fwaccel6} dos rate
      get --help
      get {-r | -l | -o /<Path>/<Name of File>}
      get [-n] -u '<Rule UID>'
      get [-n] -k <Search Key>
      get [-n] -v <Search Value>
      get --show-tab
      get --show-counters
```

Parameters:

Parameter	Description
`--help}`	Shows the applicable built-in usage.

Parameter	Description
`-r`	Shows parameters of each rule in the legacy form - each entire rule on a separate line. This is the default output format. Example: `operation=add uid=<66b101e7,00000000,cd18e2cd,000019bf> target=all timeout=none action=drop log=regular source=cidr:192.168.17.2 service=6/22 concurrent-conns=5`
`-l`	Shows parameters of each rule on separate lines. Example: `uid` `<66b101e7,00000000,cd18e2cd,000019bf>` `target` `all` `timeout` `2147483647` `action` `drop` `log` `regular` `source` `cidr:192.168.17.2` `service` `6/22` `concurrent-conns` `5`
`-o /<`*`Path`*`>/<`*`Name of File`*`>`	Saves the output in the specified file. Format of the rules will be similar to what is expected to add rules in the Batch Mode.
`-n`	Negates the search expression.
`-u '<`*`Rule UID`*`>'`	Specifies the Rule UID to show only a specific rule. **Important** - The quote marks (single or double) and angle brackets (`'<...>'`) are mandatory. To see the UID for each rule, run: `{fwaccel | fwaccel6} dos rate get`

Parameter	Description
`-k <Search Key>`	The search key. One of these: ■ `source` ■ `destination` ■ `service` ■ `concurrent-conns` ■ `concurrent-conns-ratio` ■ `pkt-rate` ■ `pkt-rate-ratio` ■ `byte-rate` ■ `byte-rate-ratio` ■ `new-conn-rate` ■ `new-conn-rate-ratio`
`-v <Search Value>`	Specifies the search value.
`--show-tab`	Also shows the output of the "`fwaccel tab`" command. See *fwaccel tab*.
`--show-counters`	Also shows the output of the "`fwaccel dos rate counters`" command.

Examples:

- Show all DoS / Rate Limiting rules:

```
[Expert@MyGW:0]# fwaccel dos rate get
fwaccel dos rate add -i
'<66b101e7,00000000,cd18e2cd,000019bf>' -action drop -log
regular source cidr:192.168.17.2 service 6/22 concurrent-
conns 5
fwaccel dos rate add -i
'<66ad0499,00000000,cd18e2cd,0000677f>' -action notify -log
regular source cidr:10.10.10.10 pkt-rate 20 service any
fwaccel dos rate add -i
'<66ad0470,00000000,cd18e2cd,00006683>' -action notify -log
regular source cidr:1.1.1.10 pkt-rate 20 service any
(3 rules found)
[Expert@MyGW:0]#
```

- Show all DoS / Rate Limiting rules other than ("not") rules with the source value of

"cidr:1.1.1.10":

```
[Expert@MyGW:0]# fwaccel dos rate get -n -k source -v
cidr:1.1.1.10
fwaccel dos rate add -i
'<66b101e7,00000000,cd18e2cd,000019bf>' -action drop -log
regular source cidr:192.168.17.2 service 6/22 concurrent-
conns 5
fwaccel dos rate add -i
'<66ad0499,00000000,cd18e2cd,0000677f>' -action notify -log
regular source cidr:10.10.10.10 pkt-rate 20 service any
(2 rules found)
[Expert@MyGW:0]#
```

- Show all DoS / Rate Limiting rules with the service value of "6/22":

```
[Expert@MyGW:0]# fwaccel dos rate get -k service -v 6/22
fwaccel dos rate add -i
'<66b101e7,00000000,cd18e2cd,000019bf>' -action drop -log
regular source cidr:192.168.17.2 service 6/22 concurrent-
conns 5
(1 rules found)
[Expert@MyGW:0]#
```

Viewing DoS / Rate Limiting Counters

Explanation

Syntax:

```
{fwaccel | fwaccel6} dos rate
      counters --help
      counters
      counters '<Rule Index>'
      counters '<Rule UID>'
```

Parameters:

Parameter	Description
`--help}`	Shows the applicable built-in usage.

Parameter	Description
`counters '<Rule Index>'`	Shows counters for the specified rule. The "`<Rule Index>`" value is assigned automatically when you add a rule - this is the rule's ordinal number. To see the Index for each rule, run: `{fwaccel | fwaccel6} dos rate get` **Notes:** ■ The count starts from the top from the number one. ■ Use the special Rule Index value "0" to see the global counters. These counters are used for ratio calculations. For example, when a rule includes a "`pkt-rate-ratio`" limit.
`counters '<Rule UID>'`	Shows counters for the specified rule. The "`<Rule Index>`" value is assigned automatically when you add a rule. To see the UID for each rule, run: `{fwaccel | fwaccel6} dos rate get`

Notes:

- To reset rule counters, see *fwaccel dos stats*.

Examples:

- Show the global counters (the Rule Index "0"):

```
[Expert@MyGW:0]# fwaccel dos rate counters 0
(Global Counters)
Concurrent Connections              0
Connection Rate                     0
Packets                             0
Bytes                               0
[Expert@MyGW:0]#
```

- Show counters for the rule with the Rule Index "1":

```
[Expert@MyGW:0]# fwaccel dos rate counters 1

Rule UID
<66b101e7,00000000,cd18e2cd,000019bf>
Policy                              1
FW Index                            1
SecureXL Index                      1
Timeout                             unlimited
Max Concurrent Connections          5
New Connection Rate                 unlimited
Packet Rate                         unlimited
Byte Rate                           unlimited
Max Concurrent Connections Ratio    unlimited
New Connection Rate Ratio           unlimited
Packet Rate Ratio                   unlimited
Byte Rate Ratio                     unlimited
Action                              drop
Log Type                            regular
Concurrent Connections              0
Connection Rate                     0
Packets                             0
Bytes                               0
Packets Dropped                     0 (since ratelimiting
policy installed)
Packets Matched                     0 (since ratelimiting
policy installed)
Violated Limits                     (none)
[Expert@MyGW:0]#
```

fwaccel dos stats

Description

The "`fwaccel dos stats`" (for IPv4) and "`fwaccel6 dos stats`" (for IPv6) commands show and clear the DoS real-time statistics in SecureXL.

Important:

- On Scalable Platforms, you must connect to the applicable Security Group.
  On Scalable Platforms (ElasticXL, Maestro, Scalable Chassis), you must run the required commands only in this way:
  - On the Security Group command line, only on the SMO Security Group Member.
  - In the Global Gaia Clish (`gclish`), must run these commands:
    - `fwaccel dos <Options>`
    - `fwaccel6 dos <Options>`
  - In the Expert mode, must run these commands (start with the "`g_`" prefix):
    - `g_fwaccel dos <Options>`
    - `g_fwaccel6 dos <Options>`
- In the VSNext mode / Traditional VSX mode, you must go to the context of an applicable Virtual Gateway / Virtual System.
  - In Gaia Clish, run: `set virtual-system <VSID>`
  - In the Expert mode, run: `vsenv <VSID>`

Syntax

```
{fwaccel | fwaccel6} dos stats
      {-h | --help}
      clear
      get
```

Parameters

Parameter	Description
No Parameters	Shows the applicable built-in usage.
`-h` `--help}`	Shows the applicable built-in usage.
`clear`	Clears the real-time statistics counters.
`get`	Shows the real-time statistics counters. **Note** - To see the counters for DoS / Rate Limiting rules, run "`fwaccel dos rate counters`".

Example - Get the current DoS statistics

```
[Expert@MyGW:0]# fwaccel dos stats get

Firewall Instances in Aggregate:
    Memory Usage:                       35376
    Total Active Connections:  (FW connection limiting inactive)
    New Connections/Second:    (FW connection limiting inactive)
    Number of Elements in Tables:
        Penalty Box Violating IPs:                      0
        Rate Limit Source Only Tracks:                  0
        Rate Limit Source and Service Tracks:           0
        Rate Limit Dest Only Tracks:                    0
        Rate Limit Dest and Service Tracks:             0

SecureXL:
    Memory Usage:                       12852
    Packets/Second:                     (rate limiting inactive)
    Bytes/Second:                       (rate limiting inactive)
    Reasons Packets Dropped:                    Monitored Only:
        IP Fragment:                0                               0
        IP Option:                  0                               0
        Penalty Box:                0                               0
        Deny List:                  0                               0
        IOC Deny List:              0                               0
        Rate Limit:                 0                               0
    Number of Elements in Tables:
        Penalty Box IPs:                                0
        Deny List IPs:                                  0
        IOC Deny List IPs:                              0
        IOC Monitor-Only IPs:                           0
        IOC External Deny List IPs:                     0
        IOC External Monitor-Only IPs:                  0
        Rate Limit Matches:                             0
        Rate Limit Source Only Tracks:                  0
        Rate Limit Source and Service Tracks:           0
        Rate Limit Dest Only Tracks:                    0
        Rate Limit Dest and Service Tracks:             0
[Expert@MyGW:0]#
```

fwaccel feature

Description

The "`fwaccel feature`" and "`fwaccel6 feature`" commands enable and disable the specified SecureXL features.

Important:

- If you disable a SecureXL feature, SecureXL does not accelerate the applicable traffic anymore.
- This change does **not** survive reboot.
- In VSX Gateway, this change is global and applies to all Virtual Systems.
- In a Cluster, you must configure all the Cluster Members in the same way.
- On Scalable Platforms, you must connect to the applicable Security Group.

Syntax for IPv4

```
fwaccel feature <Name of Feature>
      get
      off
      on
```

Syntax for IPv6

```
fwaccel6 feature <Name of Feature>
      get
      off
      on
```

Parameters

Parameter	Description
No Parameters	Shows the applicable built-in usage.
<Name of Feature>	Specifies the SecureXL feature. R82 SecureXL supports only this feature: ■ Name: `sctp` ■ Description: Stream Control Transmission Protocol (SCTP) - see sk35113
`get`	Shows the current state of the specified SecureXL feature.
`off`	Disables the specified SecureXL feature. This means that SecureXL does not accelerate the applicable traffic anymore.
`on`	Enables the specified SecureXL feature. This means that SecureXL accelerates the applicable traffic again.

Disabling the '`sctp`' feature permanently

See *"Working with Kernel Parameters" on page 479*.

1. Add this line to the `$FWDIR/boot/modules/fwkern.conf` file:

   ```
   sim_sctp_disable_by_default=1
   ```

2. Reboot.

Example 1 - Default output

```
[Expert@MyGW:0]# fwaccel feature
Usage: fwaccel feature <name> {on|off|get}

Available features: sctp
[Expert@MyGW:0]#
```

Example 2 - Disabling and enabling a feature

```
[Expert@MyGW:0]# fwaccel feature sctp get
sim_sctp_disable_by_default = 0
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel feature sctp off
Set operation succeeded
[Expert@MyGW:0]# fwaccel feature sctp get
sim_sctp_disable_by_default = 1
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel feature sctp on
Set operation succeeded
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel feature sctp get
sim_sctp_disable_by_default = 0
[Expert@MyGW:0]#
```

fwaccel ip_mr_cache

Description

The "`fwaccel ip_mr_cache`" command shows the IPv4 multicast routing cache when SecureXL works in the User Mode (UPPAK) (see *"Configuring SecureXL" on page 20*).

Important:

- On Scalable Platforms, you must connect to the applicable Security Group.
- In the VSNext mode / Traditional VSX mode, you must go to the context of an applicable Virtual Gateway / Virtual System.
  - In Gaia Clish, run: `set virtual-system <VSID>`
  - In the Expert mode, run: `vsenv <VSID>`

Syntax for IPv4

```
fwaccel ip_mr_cache
```

Example

```
[Expert@MyGW:0]# fwaccel ip_mr_cache
Group           Origin          Iif     Wrong    Last Used  Flags      Refcount  Oifs
[Expert@MyGW:0]#
```

fwaccel off

Description

The *fwaccel off* and *fwaccel6 off* commands stop the SecureXL on-the-fly.

Starting from R80.20, you can stop the SecureXL only *temporarily*. The SecureXL starts automatically when you start Check Point services (with the `cpstart` command), or reboot the Security Gateway.

Important:

- Disable the SecureXL only for debug purposes, if Check Point Support explicitly instructs you to do so.
- If you disable the SecureXL, this change does **not** survive reboot.
  SecureXL remains disabled until you enable it again on-the-fly, or reboot the Security Gateway.
- If you disable the SecureXL, this change applies only to new connections that arrive after you disable the acceleration.
  SecureXL continues to accelerate the connections that are already accelerated.
  Other non-connection oriented processing continues to function (for example, virtual defragmentation, VPN decrypt).
- On a VSX Gateway:
  - If you wish to stop the acceleration only for a specific Virtual System, go to the context of that Virtual System.
    In Gaia Clish, run: `set virtual-system <VSID>`
    In Expert mode, run: `vsenv <VSID>`
  - If you wish to stop the acceleration for all Virtual Systems, you must use the "`-a`" parameter.
    In this case, it does not matter from which Virtual System context you run this command.
- In a Cluster, you must configure all the Cluster Members in the same way.

Syntax for IPv4

```
fwaccel off [-a] [-q]
```

Syntax for IPv6

```
fwaccel6 off [-a] [-q]
```

Parameters

Parameter	Description
`-a`	On a VSX Gateway, stops acceleration on all Virtual Systems.
`-q`	Suppresses the output (does not show a returned output).

Possible returned output

- `SecureXL device disabled`
- `SecureXL device is not active`
- `Failed to disable SecureXL device`
- `fwaccel_off: failed to set process context <VSID>`

Example 1 - Output from a non-VSX Gateway

```
[Expert@MyGW:0]# fwaccel off
SecureXL device disabled.
[Expert@MyGW:0]#
```

```
fwaccel off
```

Example 2 - Output from a VSX Gateway for a specific Virtual System

```
[Expert@MyVSXGW:1]# vsx stat -v
VSX Gateway Status
==================
Name:            VSX2_192.168.3.242
Access Control Policy: VSX_GW_VSX
Installed at:    17Sep2018 13:17:14
Threat Prevention Policy: <No Policy>
SIC Status:      Trust

Number of Virtual Systems allowed by license:          25
Virtual Systems [active / configured]:                  2 / 2
Virtual Routers and Switches [active / configured]:     0 / 0
Total connections [current / limit]:                    4 / 44700

Virtual Devices Status
======================

 ID   | Type &amp;amp; Name     | Access Control Policy | Installed at    | Threat Prevention Policy
| SIC Stat
-----+--------------------+-----------------------+-----------------+--------------------------
+---------
   1 | S VS1              | VS1_Policy            | 17Sep2018 12:47 | <No Policy>
| Trust
   2 | S VS2              | VS2_Policy            | 17Sep2018 12:47 | <No Policy>
| Trust

Type: S - Virtual System, B - Virtual System in Bridge mode,
      R - Virtual Router, W - Virtual Switch.

[Expert@MyVSXGW:1]#
[Expert@MyVSXGW:1]# vsenv 1
Context is set to Virtual Device VS1 (ID 1).
[Expert@MyVSXGW:1]#
[Expert@MyVSXGW:1]# fwaccel stat
+-------------------------------------------------------------------------------+
|Id|Name     |Status     |Interfaces               |Features                     |
+-------------------------------------------------------------------------------+
|0 |KPPAK    |enabled    |eth1,eth2,eth3           |Acceleration,Cryptography    |
|  |         |           |                         |                             |
|  |         |           |                         |Crypto: Tunnel,UDPEncap,MD5, |
|  |         |           |                         |SHA1,3DES,DES,AES-128,AES-256,|
|  |         |           |                         |ESP,LinkSelection,DynamicVPN,|
|  |         |           |                         |NatTraversal,AES-XCBC,SHA256,|
|  |         |           |                         |SHA384,SHA512                |
+-------------------------------------------------------------------------------+

Accept Templates : enabled
Drop Templates   : disabled
NAT Templates    : enabled
LightSpeed Accel : disabled
[Expert@MyVSXGW:1]#

[Expert@MyVSXGW:1]# fwaccel off
SecureXL device disabled. (Virtual ID 1)
[Expert@MyVSXGW:1]#
[Expert@MyVSXGW:1]# fwaccel stat
+-------------------------------------------------------------------------------+
|Id|Name     |Status     |Interfaces               |Features                     |
+-------------------------------------------------------------------------------+
|0 |KPPAK    |disabled   |eth1,eth2,eth3           |Acceleration,Cryptography    |
|  |         |           |                         |                             |
|  |         |           |                         |Crypto: Tunnel,UDPEncap,MD5, |
|  |         |           |                         |SHA1,3DES,DES,AES-128,AES-256,|
|  |         |           |                         |ESP,LinkSelection,DynamicVPN,|
|  |         |           |                         |NatTraversal,AES-XCBC,SHA256,|
|  |         |           |                         |SHA384,SHA512                |
+-------------------------------------------------------------------------------+
```

```
Accept Templates : enabled
Drop Templates   : disabled
NAT Templates    : enabled
LightSpeed Accel : disabled
[Expert@MyVSXGW:1]#
```

Example 3 - Output from a VSX Gateway for all Virtual Systems

```
[Expert@MyVSXGW:1]# vsx stat -v
VSX Gateway Status
==================
Name:            VSX2_192.168.3.242
Access Control Policy: VSX_GW_VSX
Installed at:    17Sep2018 13:17:14
Threat Prevention Policy: <No Policy>
SIC Status:      Trust

Number of Virtual Systems allowed by license:          25
Virtual Systems [active / configured]:                  2 / 2
Virtual Routers and Switches [active / configured]:     0 / 0
Total connections [current / limit]:                    4 / 44700

Virtual Devices Status
======================

 ID  | Type &amp;amp; Name     | Access Control Policy | Installed at    | Threat Prevention Policy
| SIC Stat
-----+---------------------+-----------------------+-----------------+--------------------------
+---------
   1 | S VS1               | VS1_Policy            | 17Sep2018 12:47 | <No Policy>
| Trust
   2 | S VS2               | VS2_Policy            | 17Sep2018 12:47 | <No Policy>
| Trust

Type: S - Virtual System, B - Virtual System in Bridge mode,
      R - Virtual Router, W - Virtual Switch.

[Expert@MyVSXGW:1]#
[Expert@MyVSXGW:1]# vsenv 1
Context is set to Virtual Device VS1 (ID 1).
[Expert@MyVSXGW:1]#
[Expert@MyVSXGW:1]# fwaccel off -a
SecureXL device disabled. (Virtual ID 0)
SecureXL device disabled. (Virtual ID 1)
SecureXL device disabled. (Virtual ID 2)
[Expert@MyVSXGW:1]#
```

fwaccel on

Description

The *fwaccel on* and *fwaccel6 on* commands start the acceleration on-the-fly, if it was previously stopped with the *fwaccel off* or *fwaccel6 off* command (see *"fwaccel off" on page 128*).

Important:

- On a VSX Gateway:
  - If you wish to start the acceleration only for a specific Virtual System, go to the context of that Virtual System.
    In Gaia Clish, run: `set virtual-system <VSID>`
    In Expert mode, run: `vsenv <VSID>`
  - If you wish to start the acceleration for all Virtual Systems, you must use the "`-a`" parameter.
    In this case, it does not matter from which Virtual System context you run this command.
- In a Cluster, you must configure all the Cluster Members in the same way.

Syntax for IPv4

```
fwaccel on [-a] [-q]
```

Syntax for IPv6

```
fwaccel6 on [-a] [-q]
```

Parameters

Parameter	Description
`-a`	On a VSX Gateway, starts the acceleration on all Virtual Systems.
`-q`	Suppresses the output (does not show a returned output).

Possible returned output

- `SecureXL device is enabled.`
- `Failed to start SecureXL.`
- `No license for SecureXL.`
- `SecureXL is disabled by the firewall. Please try again later.`

- The installed SecureXL device is not compatible with the installed firewall (version mismatch).
- The SecureXL device is in the process of being stopped. Please try again later.
- SecureXL cannot be started while "flows" are active.
- SecureXL is already started.
- SecureXL will be started after a policy is loaded.
- fwaccel: Failed to check FloodGate-1 status. Acceleration will not be started.
- FW-1: SecureXL acceleration cannot be started while QoS is running in express mode.

  Please disable FloodGate-1 express mode or SecureXL.
- FW-1: SecureXL acceleration cannot be started while QoS is running with citrix printing rule.

  Please remove the citrix printing rule to enable SecureXL.
- FW-1: SecureXL acceleration cannot be started while QoS is running with UAS rule.

  Please remove the UAS rule to enable SecureXL.
- FW-1: SecureXL acceleration cannot be started while QoS is running.

  Please remove the QoS blade to enable SecureXL.
- Failed to enable SecureXL device
- fwaccel_on: failed to set process context <*VSID*>

Example 1 - Output from a non-VSX Gateway

```
[Expert@MyGW:0]# fwaccel on
SecureXL device is enabled.
[Expert@MyGW:0]#
```

```
fwaccel on
```

Example 2 - Output from a VSX Gateway for a specific Virtual System

```
[Expert@MyVSXGW:1]# vsx stat -v
VSX Gateway Status
==================
Name:            VSX2_192.168.3.242
Access Control Policy: VSX_GW_VSX
Installed at:    17Sep2018 13:17:14
Threat Prevention Policy: <No Policy>
SIC Status:      Trust

Number of Virtual Systems allowed by license:          25
Virtual Systems [active / configured]:                  2 / 2
Virtual Routers and Switches [active / configured]:     0 / 0
Total connections [current / limit]:                    4 / 44700

Virtual Devices Status
======================

 ID   | Type &amp;amp; Name     | Access Control Policy | Installed at    | Threat Prevention Policy
| SIC Stat
-----+--------------------+-----------------------+-----------------+--------------------------
+---------
   1 | S VS1              | VS1_Policy            | 17Sep2018 12:47 | <No Policy>
| Trust
   2 | S VS2              | VS2_Policy            | 17Sep2018 12:47 | <No Policy>
| Trust

Type: S - Virtual System, B - Virtual System in Bridge mode,
      R - Virtual Router, W - Virtual Switch.

[Expert@MyVSXGW:1]#
[Expert@MyVSXGW:1]# vsenv 1
Context is set to Virtual Device VS1 (ID 1).
[Expert@MyVSXGW:1]#
[Expert@MyVSXGW:1]# fwaccel stat
+-------------------------------------------------------------------------------+
|Id|Name     |Status     |Interfaces               |Features                     |
+-------------------------------------------------------------------------------+
|0 |KPPAK    |disabled   |eth1,eth2,eth3           |Acceleration,Cryptography    |
|  |         |           |                         |                             |
|  |         |           |                         |Crypto: Tunnel,UDPEncap,MD5, |
|  |         |           |                         |SHA1,3DES,DES,AES-128,AES-256,|
|  |         |           |                         |ESP,LinkSelection,DynamicVPN,|
|  |         |           |                         |NatTraversal,AES-XCBC,SHA256,|
|  |         |           |                         |SHA384,SHA512                |
+-------------------------------------------------------------------------------+

Accept Templates : enabled
Drop Templates   : disabled
NAT Templates    : enabled
LightSpeed Accel : disabled
[Expert@MyVSXGW:1]#

[Expert@MyVSXGW:1]# fwaccel on
SecureXL device enabled. (Virtual ID 1)
[Expert@MyVSXGW:1]#
[Expert@MyVSXGW:1]# fwaccel stat
+-------------------------------------------------------------------------------+
|Id|Name     |Status     |Interfaces               |Features                     |
+-------------------------------------------------------------------------------+
|0 |KPPAK    |enabled    |eth1,eth2,eth3           |Acceleration,Cryptography    |
|  |         |           |                         |                             |
|  |         |           |                         |Crypto: Tunnel,UDPEncap,MD5, |
|  |         |           |                         |SHA1,3DES,DES,AES-128,AES-256,|
|  |         |           |                         |ESP,LinkSelection,DynamicVPN,|
|  |         |           |                         |NatTraversal,AES-XCBC,SHA256,|
|  |         |           |                         |SHA384,SHA512                |
+-------------------------------------------------------------------------------+
```

```
Accept Templates : enabled
Drop Templates   : disabled
NAT Templates    : enabled
LightSpeed Accel : disabled
[Expert@MyVSXGW:1]#
```

Example 3 - Output from a VSX Gateway for all Virtual Systems

```
[Expert@MyVSXGW:1]# vsx stat -v
VSX Gateway Status
==================
Name:            VSX2_192.168.3.242
Access Control Policy: VSX_GW_VSX
Installed at:    17Sep2018 13:17:14
Threat Prevention Policy: <No Policy>
SIC Status:      Trust

Number of Virtual Systems allowed by license:          25
Virtual Systems [active / configured]:                  2 / 2
Virtual Routers and Switches [active / configured]:     0 / 0
Total connections [current / limit]:                    4 / 44700

Virtual Devices Status
======================

 ID  | Type &amp;amp; Name     | Access Control Policy | Installed at    | Threat Prevention Policy
| SIC Stat
-----+---------------------+-----------------------+-----------------+--------------------------
+---------
   1 | S VS1               | VS1_Policy            | 17Sep2018 12:47 | <No Policy>
| Trust
   2 | S VS2               | VS2_Policy            | 17Sep2018 12:47 | <No Policy>
| Trust

Type: S - Virtual System, B - Virtual System in Bridge mode,
      R - Virtual Router, W - Virtual Switch.

[Expert@MyVSXGW:1]#
[Expert@MyVSXGW:1]# vsenv 1
Context is set to Virtual Device VS1 (ID 1).
[Expert@MyVSXGW:1]#
[Expert@MyVSXGW:1]# fwaccel on -a
[Expert@MyVSXGW:1]#
```

fwaccel ranges

Description

The *fwaccel ranges* and *fwaccel6 ranges* commands show the SecureXL loaded ranges:

- Ranges of Rule Base source IP addresses
- Ranges of Rule Base destination IP addresses
- Ranges of Rule Base destination ports and protocols

The Security Gateway creates these ranges during the policy installation. The Firewall creates and offloads ranges to SecureXL when any of these features is enabled:

- Rulebase ranges for Drop Templates
- Anti-Spoofing enforcement ranges on per-interface basis
- NAT64 ranges
- NAT46 ranges

These ranges are related to matching of connections to SecureXL Drop Templates. These ranges represent the **Source**, **Destination** and **Service** columns of the Rule Base.

These ranges are not exactly the same as the Rule Base, because as there are objects that cannot be represented as real (deterministic) IP addresses. For example, Domain objects and Dynamic objects. The Security Gateway converts such non-deterministic objects to "Any" IP address.

In addition, implied rules are represented in these ranges, except for some specific implied rules.

You can use these commands for troubleshooting.

Syntax for IPv4

```
fwaccel ranges
      -h
      -a
      -l
      -p <Range ID>
      -s <Range ID>
```

Syntax for IPv6

```
fwaccel6 ranges
      -h
      -a
      -l
      -p <Range ID>
      -s <Range ID>
```

Parameters

Parameter	Description
`-h`	Shows the applicable built-in usage.
`-a` *or* No Parameters	Shows the full information for all loaded ranges. **Note** - In the list of SecureXL Drop Templates (output of the *"fwaccel templates" on page 190* command), each Drop Template is assembled from ranges indexes. To see mapping between range index and the range itself, run this command "`fwaccel ranges -a`". This way you understand better the practical ranges for Drop Templates and when it is appropriate to use them.
`-l`	Shows the list of loaded ranges: ■ 0 - Ranges of Rule Base source IP addresses ■ 1 - Ranges of Rule Base destination IP addresses ■ 2 - Ranges of Rule Base destination ports and protocols
`-p <Range ID>`	Shows the full information for the specified range.
`-s <Range ID>`	Shows the summary information for the specified range.

Examples

Example 1 - Show the list of ranges from a non-VSX Gateway

```
[Expert@MyGW:0]# fwaccel ranges -l
SecureXL device 0:
        0 Rule base source ranges (ip):
        1 Rule base destination ranges (ip):
        2 Rule base dport ranges (port, proto):
[Expert@MyGW:0]#
```

Example 2 - Show the full information for all loaded ranges from a non-VSX Gateway

```
[Expert@MyGW:0]# fwaccel ranges
SecureXL device 0:
    Rule base source ranges (ip):
        (0) 0.0.0.0 - 192.168.204.0
        (1) 192.168.204.1 - 192.168.204.1
        (2) 192.168.204.2 - 192.168.204.39
        (3) 192.168.204.40 - 192.168.204.40
        (4) 192.168.204.41 - 192.168.254.39
        (5) 192.168.254.40 - 192.168.254.40
        (6) 192.168.254.41 - 255.255.255.255
    Rule base destination ranges (ip):
        (0) 0.0.0.0 - 192.168.204.0
        (1) 192.168.204.1 - 192.168.204.1
        (2) 192.168.204.2 - 192.168.204.39
        (3) 192.168.204.40 - 192.168.204.40
        (4) 192.168.204.41 - 192.168.254.39
        (5) 192.168.254.40 - 192.168.254.40
        (6) 192.168.254.41 - 255.255.255.255
    Rule base dport ranges (port, proto):
        (0) 0, 0 - 138, 6
        (1) 139, 6 - 139, 6
        (2) 140, 6 - 18189, 6
        (3) 18190, 6 - 18190, 6
        (4) 18191, 6 - 18191, 6
        (5) 18192, 6 - 18192, 6
        (6) 18193, 6 - 19008, 6
        (7) 19009, 6 - 19009, 6
        (8) 19010, 6 - 136, 17
        (9) 137, 17 - 138, 17
        (10) 139, 17 - 65535, 65535
[Expert@MyGW:0]#
```

Example 3 - Show the full information for the specified range from a non-VSX Gateway

```
[Expert@MyGW:0]# fwaccel ranges -p 0
SecureXL device 0:
    Rule base source ranges (ip):
        (0) 0.0.0.0 - 192.168.204.0
        (1) 192.168.204.1 - 192.168.204.1
        (2) 192.168.204.2 - 192.168.204.39
        (3) 192.168.204.40 - 192.168.204.40
        (4) 192.168.204.41 - 192.168.254.39
        (5) 192.168.254.40 - 192.168.254.40
        (6) 192.168.254.41 - 255.255.255.255
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel ranges -p 1
SecureXL device 0:
    Rule base destination ranges (ip):
        (0) 0.0.0.0 - 192.168.204.0
        (1) 192.168.204.1 - 192.168.204.1
        (2) 192.168.204.2 - 192.168.204.39
        (3) 192.168.204.40 - 192.168.204.40
        (4) 192.168.204.41 - 192.168.254.39
        (5) 192.168.254.40 - 192.168.254.40
        (6) 192.168.254.41 - 255.255.255.255
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel ranges -p 2
SecureXL device 0:
    Rule base dport ranges (port, proto):
        (0) 0, 0 - 138, 6
        (1) 139, 6 - 139, 6
        (2) 140, 6 - 18189, 6
        (3) 18190, 6 - 18190, 6
        (4) 18191, 6 - 18191, 6
        (5) 18192, 6 - 18192, 6
        (6) 18193, 6 - 19008, 6
        (7) 19009, 6 - 19009, 6
        (8) 19010, 6 - 136, 17
        (9) 137, 17 - 138, 17
        (10) 139, 17 - 65535, 65535
[Expert@MyGW:0]#
```

Example 4 - Show the summary information for the specified range from a non-VSX Gateway

```
[Expert@MyGW:0]# fwaccel ranges -s 0
SecureXL device 0:
    List name "Rule base source ranges (ip):", ID 0, Number of ranges 7
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel ranges -s 1
SecureXL device 0:
    List name "Rule base destination ranges (ip):", ID 1, Number of ranges 7
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel ranges -s 2
SecureXL device 0:
    List name "Rule base dport ranges (port, proto):", ID 2, Number of ranges 11
[Expert@MyGW:0]#
```

Example 5 - Show the list of ranges from a VSX Gateway

```
[Expert@MyVSXGW:2]# vsenv 0
Context is set to Virtual Device VSX2_192.168.3.242 (ID 0).
[Expert@MyVSXGW:0]# fwaccel ranges -l
SecureXL device 0:
        0 Anti spoofing ranges eth0:
        1 Anti spoofing ranges eth1:
[Expert@MyVSXGW:0]# vsenv 1
Context is set to Virtual Device VS1 (ID 1).
[Expert@MyVSXGW:1]# fwaccel ranges -l
SecureXL device 0:
        0 Anti spoofing ranges eth3:
        1 Anti spoofing ranges eth2.52:
[Expert@MyVSXGW:1]# vsenv 2
Context is set to Virtual Device VS2 (ID 2).
[Expert@MyVSXGW:2]# fwaccel ranges -l
SecureXL device 0:
        0 Anti spoofing ranges eth4:
        1 Anti spoofing ranges eth2.53:
[Expert@MyVSXGW:2]#
```

Example 6 - Show the full information for all loaded ranges from a VSX Gateway

```
[Expert@MyVSXGW:2]# vsenv 0
Context is set to Virtual Device VSX2_192.168.3.242 (ID 0).
[Expert@MyVSXGW:0]# fwaccel ranges
SecureXL device 0:
    Anti spoofing ranges eth0:
        (0) 0.0.0.0 - 10.20.29.255
        (1) 10.20.31.0 - 126.255.255.255
        (2) 128.0.0.0 - 192.168.2.255
        (3) 192.168.3.1 - 192.168.3.241
        (4) 192.168.3.243 - 192.168.3.254
        (5) 192.168.4.0 - 223.255.255.255
        (6) 240.0.0.0 - 255.255.255.254
    Anti spoofing ranges eth1:
        (0) 10.20.30.1 - 10.20.30.241
        (1) 10.20.30.243 - 10.20.30.254
[Expert@MyVSXGW:0]#
[Expert@MyVSXGW:1]# vsenv 1
Context is set to Virtual Device VS1 (ID 1).
[Expert@MyVSXGW:1]# fwaccel ranges
SecureXL device 0:
    Anti spoofing ranges eth3:
        (0) 40.50.60.0 - 40.50.60.255
        (1) 192.168.196.17 - 192.168.196.17
        (2) 192.168.196.19 - 192.168.196.30
    Anti spoofing ranges eth2.52:
        (0) 70.80.90.0 - 70.80.90.255
        (1) 192.168.196.1 - 192.168.196.1
        (2) 192.168.196.3 - 192.168.196.14
[Expert@MyVSXGW:1]#
[Expert@MyVSXGW:1]# vsenv 2
Context is set to Virtual Device VS2 (ID 2).
[Expert@MyVSXGW:2]# fwaccel ranges
SecureXL device 0:
    Anti spoofing ranges eth4:
        (0) 100.100.100.0 - 100.100.100.255
        (1) 192.168.196.17 - 192.168.196.17
        (2) 192.168.196.19 - 192.168.196.30
    Anti spoofing ranges eth2.53:
        (0) 192.168.196.1 - 192.168.196.1
        (1) 192.168.196.3 - 192.168.196.14
        (2) 200.200.200.0 - 200.200.200.255
[Expert@MyVSXGW:2]#
```

Example 7 - Show the summary information for the specified range from a VSX Gateway

```
[Expert@MyVSXGW:2]# vsenv 1
Context is set to Virtual Device VS1 (ID 1).
[Expert@MyVSXGW:1]#
[Expert@MyVSXGW:1]# fwaccel ranges -s 0
SecureXL device 0:
    List name "Anti spoofing ranges eth3:", ID 0, Number of ranges 3
[Expert@MyVSXGW:1]#
[Expert@MyVSXGW:1]# fwaccel ranges -s 1
SecureXL device 0:
    List name "Anti spoofing ranges eth2.52:", ID 1, Number of ranges 3
[Expert@MyVSXGW:1]#
[Expert@MyVSXGW:1]# fwaccel ranges -s 2
SecureXL device 0:
        The requested range table is empty
[Expert@MyVSXGW:1]#
[Expert@MyVSXGW:1]# vsenv 2
Context is set to Virtual Device VS2 (ID 2).
[Expert@MyVSXGW:1]#
[Expert@MyVSXGW:2]# fwaccel ranges -s 0
SecureXL device 0:
    List name "Anti spoofing ranges eth4:", ID 0, Number of ranges 3
[Expert@MyVSXGW:1]#
[Expert@MyVSXGW:2]# fwaccel ranges -s 1
SecureXL device 0:
    List name "Anti spoofing ranges eth2.53:", ID 1, Number of ranges 3
[Expert@MyVSXGW:1]#
[Expert@MyVSXGW:2]# fwaccel ranges -s 2
SecureXL device 0:
        The requested range table is empty
[Expert@MyVSXGW:2]#
```

fwaccel stat

Description

The *fwaccel stat* and *fwaccel6 stat* commands show this information on the Security Gateway, or Cluster Member:

Shows this information:

Column	Information
Id	Always 0.
Name	SecureXL mode (KPPAK or UPPAK).
Status	Enabled or Disabled.
Interfaces	Names of accelerated interfaces.
Features	Names of accelerated features.

In addition, these command show the status of these features (enabled or disabled):

- SecureXL Accept Templates.
- SecureXL Drop Templates.
- SecureXL NAT Templates.
- SecureXL status of LightSpeed Acceleration.

Syntax for IPv4

```
fwaccel stat
```

Syntax for IPv6

```
fwaccel6 stat
```

Example

- This output is from a Security Gateway (non-VSX).
- SecureXL works in the Kernel Mode (KPPAK) mode.
- SecureXL is enabled.
- SecureXL accelerates the interfaces eth0 and eth1.

```
[Expert@MyGW:0]# fwaccel stat
+-------------------------------------------------------------------------------+
|Id|Name     |Status     |Interfaces               |Features                     |
+-------------------------------------------------------------------------------+
|0 |KPPAK    |enabled    |eth0,eth1                |Acceleration,Cryptography    |
|  |         |           |                         |                             |
|  |         |           |                         |Crypto: Tunnel,UDPEncap,MD5, |
|  |         |           |                         |SHA1,3DES,DES,AES-128,AES-256,|
|  |         |           |                         |ESP,LinkSelection,DynamicVPN, |
|  |         |           |                         |NatTraversal,AES-XCBC,SHA256, |
|  |         |           |                         |SHA384,SHA512                |
+-------------------------------------------------------------------------------+

Accept Templates : enabled
Drop Templates   : disabled
NAT Templates    : enabled
LightSpeed Accel : disabled
[Expert@MyGW:0]#
```

fwaccel stats

Description

The *fwaccel stats* and *fwaccel6 stats* commands show acceleration statistics for IPv4 on the local Security Gateway, or Cluster Member.

Syntax for IPv4

```
fwaccel stats
      [-c]
      [-d]
      [-l]
      [-m]
      [-n]
      [-o]
      [-p]
      [-q]
      [-r]
      [-s]
      [-x]
```

Syntax for IPv6

```
fwaccel6 stats
      [-c]
      [-d]
      [-l]
      [-m]
      [-n]
      [-o]
      [-p]
      [-q]
      [-r]
      [-s]
      [-x]
```

Parameters

Parameter	Description
`-c`	Shows the statistics for Cluster Correction.
`-d`	Shows the statistics for drops from device.
`-l`	Shows the statistics in legacy mode - as one table.
`-m`	Shows the statistics for multicast traffic.
`-n`	Shows the statistics for Identity Awareness (NAC).
`-o`	Shows the statistics for Reorder Infrastructure.  **Note** - Scalable Platforms do not support this parameter.
`-p`	Shows the statistics for SecureXL violations (F2F packets).
`-q`	Shows the statistics notifications the SecureXL sent to the Firewall.
`-r`	Resets all the counters.
`-s`	Shows the statistics summary only.
`-x`	Shows the statistics for PXL. **Note** - PXL is the technology name for combination of SecureXL and PSL (Passive Streaming Library).

In addition, see:

- *"Description of the Statistics Counters in the "fwaccel stats" Output" on page 149*
- *"Example Outputs of the "fwaccel stats" Commands" on page 158*

Description of the Statistics Counters in the "fwaccel stats" Output

The "Accelerated Path" section

Counter	Description
`accel packets`	Number of accelerated packets.
`accel bytes`	Number of accelerated bytes.
`outbound packets`	Number of outbound packets.
`outbound bytes`	Number of outbound bytes.
`conns created`	Number of connections the SecureXL created.
`conns deleted`	Number of connections the SecureXL deleted.
`C total conns`	Total number of connections the SecureXL currently handles.
`C templates`	*Not in use* Total number of SecureXL templates the SecureXL currently handles.
`C TCP conns`	Number of TCP connections the SecureXL currently handles.
`C non TCP conns`	Number of non-TCP connections the SecureXL currently handles.
`conns from templates`	*Not in use* Number of connections the SecureXL created from SecureXL templates.
`nat conns`	Number of NAT connections.
`dropped packets`	Number of packets the SecureXL dropped.
`dropped bytes`	Number of bytes the SecureXL dropped.
`nat templates`	*Not in use*
`port alloc templates`	*Not in use*
`conns from nat tmpl`	*Not in use*
`port alloc conns`	*Not in use*

Counter	Description
`fragments received`	Number of received fragments.
`fragments transmit`	Number of transmitted fragments.
`fragments dropped`	Number of dropped fragments.
`fragments expired`	Number of expired fragments.
`IP options stripped`	Number of packets, from SecureXL stripped IP options.
`IP options restored`	Number of packets, in which SecureXL restored IP options.
`IP options dropped`	Number of packets with IP options that SecureXL dropped.
`corrs created`	Number of corrections the SecureXL made.
`corrs deleted`	Number of corrections the SecureXL deleted.
`C corrections`	Number of corrections the SecureXL currently handles.
`corrected packets`	Number of corrected packets.
`corrected bytes`	Number of corrected bytes.

The "Accelerated VPN Path" section

Counter	Description
`C crypt conns`	Number of encrypted connections the SecureXL currently handles.
`enc bytes`	Number of encrypted traffic bytes.
`dec bytes`	Number of decrypted traffic bytes.
`ESP enc pkts`	Number of ESP encrypted packets.
`ESP enc err`	Number of ESP encryption errors.
`ESP dec pkts`	Number of ESP decrypted packets.
`ESP dec err`	Number of ESP decryption errors.
`ESP other err`	Number of ESP general errors.
`espudp enc pkts`	*Not in use*
`espudp enc err`	*Not in use*
`espudp dec pkts`	*Not in use*
`espudp dec err`	*Not in use*
espudp other err	*Not in use*

The "Medium Streaming Path" section

Counter	Description
`PXL packets`	Number of PXL packets. PXL is combination of SecureXL and Passive Streaming Library (PSL), which is an IPS infrastructure that transparently listens to TCP traffic as network packets, and rebuilds the TCP stream out of these packets. Passive Streaming can listen to all TCP traffic, but process only the data packets, which belong to a previously registered connection.
`PXL async packets`	Number of PXL packets the SecureXL handled asynchronously.
`PXL bytes`	Number of PXL bytes.
`C PXL conns`	Number of PXL connections the SecureXL currently handles.
`C PXL templates`	*Not in use* Number of PXL templates.
`PXL FF conns`	Number of PXL Fast Forward connections.
`PXL FF packets`	Number of PXL Fast Forward packets.
`PXL FF bytes`	Number of PXL Fast Forward bytes.
`PXL FF acks`	Number of PXL Fast Forward acknowledgments.

The "Inline Streaming Path" section

Counter	Description
`PSL Inline packets`	Number of accelerated PSL packets.
`PSL Inline bytes`	Number of accelerated PSL bytes.
`CPAS Inline packets`	Number of accelerated CPAS packets.
`CPAS Inline bytes`	Number of accelerated CPAS bytes.

The "QoS General Information" section

Counter	Description
`Total QoS Conns`	Total number of QoS connections.
`QoS Classify Conns`	Number of classified QoS connections.
`QoS Classify flow`	Number of classified QoS flows.
`Reclassify QoS polic`	Number of reclassify QoS requests.

The "Firewall QoS Path" section

Counter	Description
`Enqueued IN packets`	Number of waiting packets in Firewall QoS inbound queue.
`Enqueued OUT packets`	Number of waiting packets in Firewall QoS outbound queue.
`Dequeued IN packets`	Number of processed packets in Firewall QoS inbound queue.
`Dequeued OUT packets`	Number of processed packets in Firewall QoS outbound queue.
`Enqueued IN bytes`	Number of waiting bytes in Firewall QoS inbound queue.
`Enqueued OUT bytes`	Number of waiting bytes in Firewall QoS outbound queue.
`Dequeued IN bytes`	Number of processed bytes in Firewall QoS inbound queue.
`Dequeued OUT bytes`	Number of processed bytes in Firewall QoS outbound queue.

The "Firewall QoS Path" section

Counter	Description
`Enqueued IN packets`	Number of waiting packets in SecureXL QoS inbound queue.
`Enqueued OUT packets`	Number of waiting packets in SecureXL QoS outbound queue.
`Dequeued IN packets`	Number of processed packets in SecureXL QoS inbound queue.
`Dequeued OUT packets`	Number of processed packets in SecureXL QoS outbound queue.
`Enqueued IN bytes`	Number of waiting bytes in SecureXL QoS inbound queue.
`Enqueued OUT bytes`	Number of waiting bytes in SecureXL QoS outbound queue.
`Dequeued IN bytes`	Number of processed bytes in SecureXL QoS inbound queue.
`Dequeued OUT bytes`	Number of processed bytes in SecureXL QoS outbound queue.

The "Firewall Path" section

Counter	Description
`F2F packets`	Number of packets that SecureXL forwarded to the Firewall kernel in Slow Path.
`F2F bytes`	Number of bytes that SecureXL forwarded to the Firewall kernel in Slow Path.
`TCP violations`	Number of packets, which are in violation of the TCP state.
`C anticipated conns`	Number of anticipated connections SecureXL currently handles.
`port alloc f2f`	*Not in use*
`F2V conn match pkts`	Number of packets that matched a SecureXL connection and SecureXL forwarded to the Firewall kernel.
`F2V packets`	Number of packets that SecureXL forwarded to the Firewall kernel and the Firewall re-injected back to SecureXL.
`F2V bytes`	Number of bytes that SecureXL forwarded to the Firewall kernel and the Firewall re-injected back to the SecureXL.

The "GTP" section

Counter	Description
`gtp tunnels created`	Number of created GTP tunnels.
`gtp tunnels`	Number of GTP tunnels the SecureXL currently handles.
`gtp accel pkts`	Number of accelerated GTP packets.
`gtp f2f pkts`	Number of GTP packets the SecureXL forwarded to the Firewall kernel.
`gtp spoofed pkts`	Number of spoofed GTP packets.
`gtp in gtp pkts`	Number of GTP-in-GTP packets.
`gtp signaling pkts`	Number of signaling GTP packets.
`gtp tcpopt pkts`	Number of GTP packets with TCP Options.
`gtp apn err pkts`	Number of GTP packets with APN errors.

The "General" section

Counter	Description
`memory used`	*Not in use*
`free memory`	*Not in use*
`C used templates`	*Not in use*
`pxl tmpl conns`	*Not in use*
`C conns from tmpl`	*Not in use* Number of current connections that SecureXL created from SecureXL Templates.
`C tcp handshake conn`	Number of current TCP connections that are not yet established.
`C tcp established co`	Number of established TCP connections the SecureXL currently handles.
`C tcp closed conns`	Number of closed TCP connections the SecureXL currently handles.
`C tcp pxl handshake`	Number of not yet established PXL TCP connections the SecureXL currently handles.
`C tcp pxl establishe`	Number of established PXL TCP connections the SecureXL currently handles.
`C tcp pxl closed con`	Number of closed PXL TCP connections the SecureXL currently handles.
`outbound pxl packets`	*Not in use*

Example Outputs of the "fwaccel stats" Commands

Example: fwaccel stats -s

Example of statistics summary:

```
Accelerated conns/Total conns : 0/0 (0%)
Accelerated pkts/Total pkts   : 0/8 (0%)
F2Fed pkts/Total pkts         : 8/8 (100%)
F2V pkts/Total pkts           : 0/8 (0%)
CPASXL pkts/Total pkts        : 0/8 (0%)
PSLXL pkts/Total pkts         : 0/8 (0%)
QOS inbound pkts/Total pkts   : 0/8 (0%)
QOS outbound pkts/Total pkts  : 0/8 (0%)
Corrected pkts/Total pkts     : 0/8 (0%)
```

Example: fwaccel stats

Example of the default output:

```
Name                          Value         Name                          Value
----------------------------  ------------  ----------------------------  ------------

Accelerated Path
-------------------------------------------------------------------------------
accel packets                            0    accel bytes                              0
outbound packets                         0    outbound bytes                           0
conns created                            0    conns deleted                            0
C total conns                            0    C TCP conns                              0
C non TCP conns                          0    nat conns                                0
dropped packets                          0    dropped bytes                            0
fragments received                       0    fragments transmit                       0
fragments dropped                        0    fragments expired                        0
IP options stripped                      0    IP options restored                      0
IP options dropped                       0    corrs created                            0
corrs deleted                            0    C corrections                            0
corrected packets                        0    corrected bytes                          0

Accelerated VPN Path
-------------------------------------------------------------------------------
C crypt conns                            0    enc bytes                                0
dec bytes                                0    ESP enc pkts                             0
ESP enc err                              0    ESP dec pkts                             0
ESP dec err                              0    ESP other err                            0
espudp enc pkts                          0    espudp enc err                           0
espudp dec pkts                          0    espudp dec err                           0
espudp other err                         0

Medium Streaming Path
-------------------------------------------------------------------------------
CPASXL packets                           0    PSLXL packets                            0
CPASXL async packets                     0    PSLXL async packets                      0
CPASXL bytes                             0    PSLXL bytes                              0
C CPASXL conns                           0    C PSLXL conns                            0
CPASXL conns created                     0    PSLXL conns created                      0
PXL FF conns                             0    PXL FF packets                           0
PXL FF bytes                             0    PXL FF acks                              0
PXL no conn drops                        0

Inline Streaming Path
-------------------------------------------------------------------------------
PSL Inline packets                       0    PSL Inline bytes                         0
CPAS Inline packets                      0    CPAS Inline bytes                        0

QoS Paths
-------------------------------------------------------------------------------
QoS General Information:
------------------------
Total QoS Conns                          0    QoS Classify Conns                       0
QoS Classify flow                        0    Reclassify QoS policy                    0

FireWall QoS Path:
------------------
Enqueued IN packets                      0    Enqueued OUT packets                     0
Dequeued IN packets                      0    Dequeued OUT packets                     0
Enqueued IN bytes                        0    Enqueued OUT bytes                       0
Dequeued IN bytes                        0    Dequeued OUT bytes                       0

Accelerated QoS Path:
---------------------
Enqueued IN packets                      0    Enqueued OUT packets                     0
Dequeued IN packets                      0    Dequeued OUT packets                     0
Enqueued IN bytes                        0    Enqueued OUT bytes                       0
Dequeued IN bytes                        0    Dequeued OUT bytes                       0

Firewall Path
-------------------------------------------------------------------------------
F2F packets                          35324    F2F bytes                          1797781
TCP violations                           0    F2V conn match pkts                      0
F2V packets                              0    F2V bytes                                0
```

```
GTP
-------------------------------------------------------------------------------
gtp tunnels created                     0    gtp tunnels                        0
gtp accel pkts                          0    gtp f2f pkts                       0
gtp spoofed pkts                        0    gtp in gtp pkts                    0
gtp signaling pkts                      0    gtp tcpopt pkts                    0
gtp apn err pkts                        0

General
-------------------------------------------------------------------------------
memory used                      38798784    C tcp handshake conns              0
C tcp established conns                 0    C tcp closed conns                 0
C tcp pxl handshake conns               0    C tcp pxl established conns        0
C tcp pxl closed conns                  0    outbound cpasxl packets            0
outbound pslxl packets                  0    outbound cpasxl bytes              0
outbound pslxl bytes                    0    DNS DoR stats                      0

(*) Statistics marked with C refer to current value, others refer to total value
```

Example: fwaccel stats -c

Example of statistics for Cluster Correction:

```
Cluster Correction stats:

 Name                    Value           Name                     Value
----------------------  ------------    ----------------------  ------------
Sent pkts (total)                0    Sent with metadata                 0
Received pkts (total)            0    Received with metadata             0
Sent bytes                       0    Received bytes                     0
Send errors                      0    Receive errors                     0
```

Example: fwaccel stats -d

Example of statistics for drops from device:

```
Reason                Value              Reason                Value
--------------------  ---------------    --------------------  ---------------
general reason                   0    CPASXL decision                    0
PSLXL decision                   0    clr pkt on vpn                     0
encrypt failed                   0    drop template                      0
decrypt failed                   0    interface down                     0
cluster error                    0    XMT error                          0
anti spoofing                    0    local spoofing                     0
sanity error                     0    monitored spoofed                  0
QOS decision                     0    C2S violation                      0
S2C violation                    0    Loop prevention                    0
DOS Fragments                    0    DOS IP Options                     0
DOS Blacklists                   0    DOS Penalty Box                    0
DOS Rate Limiting                0    Syn Attack                         0
Reorder                          0    Expired Fragments                  0
```

Example: fwaccel stats -l

Example of the output in legacy mode (as one table):

```
Name                          Value         Name                          Value
---------------------------  ------------  ---------------------------  ------------
-                                    0     accel packets                        0
accel bytes                          0     outbound packets                     0
outbound bytes                       0     conns created                        0
conns deleted                        0     C total conns                        0
C TCP conns                          0     C non TCP conns                      0
nat conns                            0     dropped packets                      0
dropped bytes                        0     fragments received                   0
fragments transmit                   0     fragments dropped                    0
fragments expired                    0     IP options stripped                  0
IP options restored                  0     IP options dropped                   0
corrs created                        0     corrs deleted                        0
C corrections                        0     corrected packets                    0
corrected bytes                      0     C crypt conns                        0
enc bytes                            0     dec bytes                            0
ESP enc pkts                         0     ESP enc err                          0
ESP dec pkts                         0     ESP dec err                          0
ESP other err                        0     espudp enc pkts                      0
espudp enc err                       0     espudp dec pkts                      0
espudp dec err                       0     espudp other err                     0
acct update interval              3600     CPASXL packets                       0
PSLXL packets                        0     CPASXL async packets                 0
PSLXL async packets                  0     CPASXL bytes                         0
PSLXL bytes                          0     C CPASXL conns                       0
C PSLXL conns                        0     CPASXL conns created                 0
PSLXL conns created                  0     PXL FF conns                         0
PXL FF packets                       0     PXL FF bytes                         0
PXL FF acks                          0     PXL no conn drops                    0
PSL Inline packets                   0     PSL Inline bytes                     0
CPAS Inline packets                  0     CPAS Inline bytes                    0
Total QoS Conns                      0     QoS Classify Conns                   0
QoS Classify flow                    0     Reclassify QoS policy                0
Enqueued IN packets                  0     Enqueued OUT packets                 0
Dequeued IN packets                  0     Dequeued OUT packets                 0
Enqueued IN bytes                    0     Enqueued OUT bytes                   0
Dequeued IN bytes                    0     Dequeued OUT bytes                   0
Enqueued IN packets                  0     Enqueued OUT packets                 0
Dequeued IN packets                  0     Dequeued OUT packets                 0
Enqueued IN bytes                    0     Enqueued OUT bytes                   0
Dequeued IN bytes                    0     Dequeued OUT bytes                   0
F2F packets                      35383     F2F bytes                      1801493
TCP violations                       0     F2V conn match pkts                  0
F2V packets                          0     F2V bytes                            0
gtp tunnels created                  0     gtp tunnels                          0
gtp accel pkts                       0     gtp f2f pkts                         0
gtp spoofed pkts                     0     gtp in gtp pkts                      0
gtp signaling pkts                   0     gtp tcpopt pkts                      0
gtp apn err pkts                     0     memory used                   38798784
C tcp handshake conns                0     C tcp established conns              0
C tcp closed conns                   0     C tcp pxl handshake conns            0
C tcp pxl established conns          0     C tcp pxl closed conns               0
outbound cpasxl packets              0     outbound pslxl packets               0
outbound cpasxl bytes                0     outbound pslxl bytes                 0
DNS DoR stats                        0
(*) Statistics marked with C refer to current value, others refer to total value
```

Example: fwaccel stats -m

Example of statistics for multicast traffic:

```
Name                 Value              Name                 Value
-------------------  ---------------    -------------------  ---------------
in packets                         0    out packets                        0
if restricted                      0    conns with down if                 0
f2f packets                        0    f2f bytes                          0
dropped packets                    0    dropped bytes                      0
accel packets                      0    accel bytes                        0
mcast conns                        0
```

Example: fwaccel stats -n

Example of statistics for Identity Awareness (NAC):

```
Name                 Value              Name                 Value
-------------------  ---------------    -------------------  ---------------
NAC packets                        0    NAC bytes                          0
NAC connections                    0    compliance failure                 0
```

Example: fwaccel stats -o

Example of statistics for Reorder Infrastructure:

```
Appliaction: F2V
        Statistic                       Value
----------------------------------     --------------------
                        Queued pkts                     0
                     Max queued pkts                    0
                     Timer triggered                    0
          Callback hahndling unhold                     0
 Callback hahndling unhold and drop                     0
           Callback hahndling reset                     0
              Dequeued pkts resumed                     0
                Queue ent allocated                     0
                    Queue ent freed                     0
                   Queues allocated                     0
                       Queues freed                     0
                     Ack notif sent                     0
              Ack respones handling                     0
              Dequeued pkts dropped                     0
       Reached max queued pkt limit                     0
                  Set timer failed                      0
                Error already held                      0
             Queue ent alloc failed                     0
                 Queue alloc failed                     0
                   Ack notif failed                     0
       Ack respones handling failed                     0
--------------------------------------------------

Appliaction: Route
        Statistic                       Value
----------------------------------     --------------------
                        Queued pkts                     0
                     Max queued pkts                    0
                     Timer triggered                    0
          Callback hahndling unhold                     0
 Callback hahndling unhold and drop                     0
           Callback hahndling reset                     0
              Dequeued pkts resumed                     0
                Queue ent allocated                     0
                    Queue ent freed                     0
                   Queues allocated                     0
                       Queues freed                     0
                     Ack notif sent                     0
              Ack respones handling                     0
              Dequeued pkts dropped                     0
       Reached max queued pkt limit                     0
                  Set timer failed                      0
                Error already held                      0
             Queue ent alloc failed                     0
                 Queue alloc failed                     0
                   Ack notif failed                     0
       Ack respones handling failed                     0
--------------------------------------------------

Appliaction: New connection
        Statistic                       Value
----------------------------------     --------------------
                        Queued pkts                     0
                     Max queued pkts                    0
                     Timer triggered                    0
          Callback hahndling unhold                     0
 Callback hahndling unhold and drop                     0
           Callback hahndling reset                     0
              Dequeued pkts resumed                     0
                Queue ent allocated                     0
                    Queue ent freed                     0
                   Queues allocated                     0
                       Queues freed                     0
                     Ack notif sent                     0
              Ack respones handling                     0
              Dequeued pkts dropped                     0
       Reached max queued pkt limit                     0
                  Set timer failed                      0
```

```
                Error already held                    0
             Queue ent alloc failed                   0
                 Queue alloc failed                   0
                   Ack notif failed                   0
       Ack respones handling failed                   0
-------------------------------------------------

Appliaction: F2P
        Statistic                       Value
----------------------------------     --------------------
                        Queued pkts                     0
                     Max queued pkts                    0
                     Timer triggered                    0
          Callback hahndling unhold                     0
 Callback hahndling unhold and drop                     0
           Callback hahndling reset                     0
              Dequeued pkts resumed                     0
                Queue ent allocated                     0
                    Queue ent freed                     0
                  Queues allocated                      0
                      Queues freed                      0
                    Ack notif sent                      0
              Ack respones handling                     0
              Dequeued pkts dropped                     0
       Reached max queued pkt limit                     0
                   Set timer failed                     0
                Error already held                      0
             Queue ent alloc failed                     0
                 Queue alloc failed                     0
                   Ack notif failed                     0
       Ack respones handling failed                     0
-------------------------------------------------
```

Note - Scalable Platforms do not support this parameter.

Example: fwaccel stats -p

Example of statistics for SecureXL violations (F2F packets):

```
F2F packets:
--------------
Violation             Packets            Violation             Packets
--------------------  ---------------    --------------------  ---------------
pkt has IP options                0    ICMP miss conn                    3036
TCP-SYN miss conn                 8    TCP-other miss conn              32224
UDP miss conn                  3772    other miss conn                      0
VPN returned F2F                  0    uni-directional viol                 0
possible spoof viol               0    TCP state viol                       0
out if not def/accl               0    bridge, src=dst                      0
routing decision err              0    sanity checks failed                 0
fwd to non-pivot                  0    broadcast/multicast                  0
cluster message                   0    cluster forward                      0
chain forwarding                  0    F2V conn match pkts                  0
general reason                    0    route changes                        0
```

Example: fwaccel stats -q

Example of statistics for notifications the SecureXL sent to the Firewall:

```
Notification           Packets           Notification           Packets
--------------------  --------------    ---------------------  --------------
ntSAAboutToExpire                0    ntSAExpired                           0
ntMSPIError                      0    ntNoInboundSA                         0
ntNoOutboundSA                   0    ntDataIntegrityFailed                 0
ntPossibleReplay                 0    ntReplay                              0
ntNextProtocolError              0    ntCPIError                            0
ntClearTextPacket                0    ntFragmentation                       0
ntUpdateUdpEncTable              0    ntSASync                              0
ntReplayOutOfWindow              0    ntVPNTrafficReport                    0
ntConnDeleted                    0    ntConnUpdate                          0
ntPacketDropped                  0    ntSendLog                             0
ntRefreshGTPTunnel               0    ntMcastDrop                           0
ntAccounting                     0    ntAsyncIndex                          0
ntACkReordering                  0    ntAccelAckInfo                        0
ntMonitorPacket                  0    ntPacketCapture                       0
ntCpasPacketCapture              0    ntPSLGlueUpdateReject                 0
ntSeqVerifyDrop                  0    ntPacketForwardBefore                 0
ntICMPMessage                    0    ntQoSReclassifyPacket                 0
ntQoSResumePacket                0    ntVPNEncHaLinkFailure                 0
ntVPNEncLsLinkFailure            0    ntVPNEncRouteChange                   0
ntVPNDecVerRouteChang            0    ntVPNDecRouteChange                   0
ntMuxSimToFw                     0    ntPSLEventLog                         0
ntSendCPHWDStats             14871    ntPacketTaggingViolat                 0
ntDosNotify                     28    ntSynatkNotify                        0
ntSynatkStats                    0    ntQoSEventLog                         0
ntPrintGetParam                  0
```

Example: fwaccel stats -x

Example of statistics for PXL:

```
PXL Release Context statistics:

 Name                    Value           Name                     Value
----------------------  ------------    -----------------------  ------------
End Handler                       0    Post Sync                             0
Stop Stream                       0    kbuf fail                             0
Set field failure                 0    Notif set field fail                  0
Non SYN seq fail                  0    Tmpl kbuf fail                        0
Tmpl set field fail               0    Segment Injection                     0
Init app fail                     0    Expiration                            0
Newconn set field fail            0    Newconn fail                          0
CPHWD dec                         0    No PSL policy                         0

PXL Exception statistics:

 Name                    Value           Name                     Value
----------------------  ------------    -----------------------  ------------
urgent packets                    0    invalid SYN retrans                   0
SYN seq not init                  0    old pkts out win                      0
old pkts out win trunc            0    old pkts out win strip                0
new pkts out win                  0    incorrect retrans                     0
TCP pkts with bad csum            0    ACK unprocessed data                  0
old ACK out win                   0    Max segments reached                  0
No resources                      0    Hold timeout                          0
```

fwaccel synatk

Description

The *fwaccel synatk* and *fwaccel6 synatk* commands control the Accelerated SYN Defender on the local Security Gateway, or Cluster Member.

Important:

- See sk120476 for information about the 'SYN Attack' protection in SmartConsole.
- In a Scalable Platform, the same SecureXL command must run on all Security Group Members.
  Therefore, you must run the SecureXL commands in either Gaia gClish, or Expert mode.
  - In Gaia gClish, run the "`fwaccel ...`" and "`fwaccel6 ...`" commands.
  - In the Expert mode, run the "`g_fwaccel ...`" and "`g_fwaccel6 ...`" commands.
- In a Scalable Platform, when you add a new Security Group Member to a Security Group, the new Security Group Member pulls the "`fwaccel synatk`" configuration that you saved it in a configuration file - in the default file `$FWDIR/conf/synatk.conf`, or in the file specified with the "`fwaccel synatk -c`" command.

Syntax for IPv4

```
fwaccel synatk
      -a
      -c <options>
      -d
      -e
      -g
      -m
      -t <options>
      config
      monitor <options>
      state <options>
      whitelist <options>
```

Syntax for IPv6

```
fwaccel6 synatk
      -a
      -c <options>
      -d
      -e
      -g
      -m
      -t <options>
      config
      monitor <options>
      state <options>
      whitelist <options>
```

Parameters

Parameter	Description
No Parameters	Shows the applicable built-in usage.
`-a`	Applies the configuration from the default file. See *"fwaccel synatk -a" on page 169*.
`-c <options>`	Applies the configuration from the specified file. See *"fwaccel synatk -c <Configuration File>" on page 170*.
`-d`	Disables the Accelerated SYN Defender on all interfaces. See *"fwaccel synatk -d" on page 171*.
`-e`	Enables the Accelerated SYN Defender on interfaces with topology "External". Enables the Accelerated SYN Defender in Monitor (Detect only) mode on interfaces with topology "Internal". See *"fwaccel synatk -e" on page 172*.
`-g`	Enables the Accelerated SYN Defender on all interfaces. See *"fwaccel synatk -g" on page 173*.
`-m`	Enables the Accelerated SYN Defender in Monitor (Detect only) mode on all interfaces. In this state, the Accelerated SYN Defender only sends a log when it recognizes a TCP SYN Flood attack. See *"fwaccel synatk -m" on page 174*.

Parameter	Description
`-t <options>`	Configures the threshold numbers of half-opened TCP connections that trigger the Accelerated SYN Defender. See *"fwaccel synatk -t <Threshold>" on page 175*.
`config`	Shows the current Accelerated SYN Defender configuration. See *"fwaccel synatk config" on page 177*.
`monitor <options>`	Shows the Accelerated SYN Defender status. See *"fwaccel synatk monitor" on page 180*.
`state <options>`	Controls the Accelerated SYN Defender states. See *"fwaccel synatk state" on page 185*.

fwaccel synatk -a

Description

The "*fwaccel synatk -a*" and "*fwaccel6 synatk -a*" commands apply the Accelerated SYN Defender configuration from the default `$FWDIR/conf/synatk.conf` file.

Notes:

- Both IPv4 and IPv6 use the same configuration file.
- Interface specific state settings that you define in the configuration file, override the settings that you define with these commands:
  - *"fwaccel synatk -d" on page 171*
  - *"fwaccel synatk -e" on page 172*
  - *"fwaccel synatk -g" on page 173*
  - *"fwaccel synatk -m" on page 174*

Syntax for IPv4

```
fwaccel synatk -a
```

Syntax for IPv6

```
fwaccel6 synatk -a
```

fwaccel synatk -c <Configuration File>

Description

The "*fwaccel synatk -c <Configuration File>*" and "*fwaccel6 synatk -c <Configuration File>*" commands apply the Accelerated SYN Defender configuration from the specified file.

Important - If you use this parameter, then it must be the first parameter in the syntax.

Notes:

- Both IPv4 and IPv6 use the same configuration file.
- The state settings of a specific interface that you define in the configuration file, override the settings that you define with these commands:
  - *"fwaccel synatk -d" on page 171*
  - *"fwaccel synatk -e" on page 172*
  - *"fwaccel synatk -g" on page 173*
  - *"fwaccel synatk -m" on page 174*

Syntax for IPv4

```
fwaccel synatk -c <Configuration File>
```

Syntax for IPv6

```
fwaccel6 synatk -c <Configuration File>
```

Parameters

Parameter	Description
<Configuration File>	Specifies the full path and the name of the file. For reference, see the default file: `$FWDIR/conf/synatk.conf`

fwaccel synatk -d

Description

The "*fwaccel synatk -d*" and "*fwaccel6 synatk -d*" commands disable the Accelerated SYN Defender on all interfaces.

Notes:

- This command:
  1. Modifies the default configuration file `$FWDIR/conf/synatk.conf`, or the configuration file specified with the "`-c`" parameter.
  2. Loads the modified file.
  3. Does not show any output.
- Output of the *"fwaccel synatk monitor" on page 180* command shows:
  - In the row "`Configuration`": `Disabled`
  - In the column "`Enforce`": `Disable`
  - In the column "`State (sec)`": `Disable`
- Output of the *"fwaccel synatk config" on page 177* command shows:
  - In the row "`enabled`": `0`
  - In the row "`enforce`": `0`

Syntax for IPv4

```
fwaccel synatk -d
```

Syntax for IPv6

```
fwaccel6 synatk -d
```

fwaccel synatk -e

Description

The "*fwaccel synatk -e*" and "*fwaccel6 synatk -e*" commands:

- Enable the Accelerated SYN Defender on interfaces with topology "External".
- Enable the Accelerated SYN Defender in Monitor (Detect only) mode on interfaces with topology "Internal".

Notes:

- This command:
  1. Modifies the default configuration file `$FWDIR/conf/synatk.conf`, or the configuration file specified with the "`-c`" parameter.
  2. Loads the modified file.
- Output of the *"fwaccel synatk monitor" on page 180* command shows for "External" interfaces:
  - Configuration: `Enforcing`
  - Enforce: `Prevent`
  - State: `Ready` (may change later depending on what the SYN Defender detects)
- Output of the *"fwaccel synatk monitor" on page 180* command shows for "Internal" interfaces:
  - Configuration: `Enforcing`
  - Enforce: `Detect`
  - State: `Monitor`
- Output of the *"fwaccel synatk config" on page 177* command shows:
  - `enabled 1`
  - `enforce 1`

Syntax for IPv4

```
fwaccel synatk -e
```

Syntax for IPv6

```
fwaccel6 synatk -e
```

fwaccel synatk -g

Description

The "*fwaccel synatk -g*" and "*fwaccel6 synatk -g*" commands enable the Accelerated SYN Defender on all interfaces.

Notes:

- This command:
  1. Modifies the default configuration file `$FWDIR/conf/synatk.conf`, or the configuration file specified with the "`-c`" parameter.
  2. Loads the modified file.
- Output of the *"fwaccel synatk monitor" on page 180* command shows for "External" interfaces:
  - Configuration: `Enforcing`
  - Enforce: `Prevent`
  - State: `Ready` (may change later depending on what the SYN Defender detects)
- Output of the *"fwaccel synatk monitor" on page 180* command shows for "Internal" interfaces:
  - Configuration: `Enforcing`
  - Enforce: `Detect`
  - State: `Monitor`
- Output of the *"fwaccel synatk config" on page 177* command shows:
  - `enabled 1`
  - `enforce 2`

Syntax for IPv4

```
fwaccel synatk -g
```

Syntax for IPv6

```
fwaccel6 synatk -g
```

fwaccel synatk -m

Description

The "*fwaccel synatk -m*" and "*fwaccel6 synatk -m*" commands enable the Accelerated SYN Defender in Monitor (Detect only) mode on all interfaces.

In this state, the Accelerated SYN Defender only sends a log when it recognizes a TCP SYN Flood attack.

Notes:

- This command:
  1. Modifies the default configuration file `$FWDIR/conf/synatk.conf`, or the configuration file specified with the "`-c`" parameter.
  2. Loads the modified file.
- Output of the *"fwaccel synatk monitor" on page 180* command shows:
  - Configuration: `Monitoring`
  - Enforce: `Detect`
  - State: `Monitor`
- Output of the *"fwaccel synatk config" on page 177* command shows:
  - `enabled 1`
  - `enforce 0`

Syntax for IPv4

```
fwaccel synatk -m
```

Syntax for IPv6

```
fwaccel6 synatk -m
```

fwaccel synatk -t <Threshold>

Description

The "*fwaccel synatk -t <Threshold>*" and "*fwaccel6 synatk -t <Threshold>*" commands configure the threshold numbers of half-opened TCP connections that trigger the Accelerated SYN Defender.

Notes:

- This command:
  1. Modifies the default configuration file `$FWDIR/conf/synatk.conf`, or the configuration file specified with the "`-c`" parameter.
  2. Loads the modified file.
- Threshold values are independent for IPv4 and IPv6.

Syntax for IPv4

```
fwaccel synatk -t <Threshold>
```

Syntax for IPv6

```
fwaccel6 synatk -t <Threshold>
```

Thresholds

- The **Global high attack threshold** number is configured to the specified value *<Threshold>*.

  This is the number of half-open TCP connections on all interfaces required for the Accelerated SYN Defender to engage.

  - Valid values: 100 and greater
  - Default: 10000

- The **High attack threshold** number is configured to 1/2 of the specified value *<Threshold>*.

  This is the high number of half-open TCP connections on an interface required for the Accelerated SYN Defender to engage.

  - Valid values: (Low attack threshold) < (High attack threshold) <= (Global high attack threshold)
  - Default: 5000

- The **Low attack threshold** number is configured to 1/10 of the specified value *<Threshold>*.

  This is the low number of half-open TCP connections on an interface required for the Accelerated SYN Defender to engage.

  - Valid values: 10 and greater
  - Default: 1000

fwaccel synatk config

Description

The "*fwaccel synatk config*" and "*fwaccel6 synatk config*" commands show the current Accelerated SYN Defender configuration.

Syntax for IPv4

```
fwaccel synatk config
```

Syntax for IPv6

```
fwaccel6 synatk config
```

Example

```
[Expert@MyGW:0]# fwaccel synatk config
enabled 0
enforce 1
global_high_threshold 10000
periodic_updates 1
cookie_resolution_shift 6
min_frag_sz 80
high_threshold 5000
low_threshold 1000
score_alpha 100
monitor_log_interval (msec) 60000
grace_timeout (msec) 30000
min_time_in_active (msec) 60000
[Expert@MyGW:0]#
```

Description of Configuration Parameters

Parameter	Description
`enabled`	Shows if the Accelerated SYN Defender is enabled or disabled. ■ Valid values: 0 (disabled), 1 (enabled) ■ Default: 0
`enforce`	When the Accelerated SYN Defender is enabled, shows it enforces the protection. Valid values: ■ 0 - The Accelerated SYN Defender is in Monitor (Detect only) mode on all interfaces. ■ 1 - The Accelerated SYN Defender is engaged only on external interfaces when the number of half-open TCP connections exceeds the threshold. ■ 2 - The Accelerated SYN Defender is engaged on both external and internal interfaces when the number of half-open TCP connections exceeds the threshold.
`global_high_threshold`	Global high attack threshold number. See the *"fwaccel synatk -t <Threshold>" on page 175* command.
`periodic_updates`	For internal Check Point use only. ■ Valid values: 0 (disabled), 1 (enabled) ■ Default: 1
`cookie_resolution_shift`	For internal Check Point use only. ■ Valid values: 1-7 ■ Default: 6
`min_frag_sz`	During the TCP SYN Flood attack, the Accelerated SYN Defender prevents TCP fragments smaller than this minimum size value. ■ Valid values: 80 and greater ■ Default: 80
`high_threshold`	High attack threshold number. See the *"fwaccel synatk -t <Threshold>" on page 175* command.
`low_threshold`	Low attack threshold number. See the *"fwaccel synatk -t <Threshold>" on page 175* command.

Parameter	Description
`score_alpha`	For internal Check Point use only. ■ Valid values: 1-127 ■ Default: 100
`monitor_log_interval (msec)`	Interval, in milliseconds, between successive warning logs in the Monitor (Detect only) mode. ■ Valid values: 1000 and greater ■ Default: 60000
`grace_timeout (msec)`	Maximum time, in milliseconds, to stay in the Grace state (which is a transitional state between Ready and Active). In the Grace state, the Accelerated SYN Defender stops challenging Clients for TCP SYN Cookie, but continues to validate TCP SYN Cookies it receives from Clients. ■ Valid values: 10000 and greater ■ Default: 30000
`min_time_in_active (msec)`	Minimum time, in milliseconds, to stay in the Active mode. In the Active mode, the Accelerated SYN Defender is actively challenging TPC SYN packets with SYN Cookies. ■ Valid values: 10000 and greater ■ Default: 60000

fwaccel synatk monitor

Description

The "*fwaccel synatk monitor*" and "*fwaccel6 synatk monitor*" commands show the Accelerated SYN Defender status.

Important - To enable the Accelerated SYN Defender in Monitor (Detect only) mode on all interfaces, you must run the *"fwaccel synatk -m" on page 174* command.

Syntax for IPv4

```
fwaccel synatk monitor
      [-p]
      [-p] -a
      [-p] -s
      [-p] -v
```

Syntax for IPv6

```
fwaccel6 synatk monitor
      [-p]
      [-p] -a
      [-p] -s
      [-p] -v
```

Parameters

Important - You can specify only one of these parameters: `-a`, `-s`, or `-v`.

Parameter	Description
`-p`	Shows the Accelerated SYN Defender status for each SecureXL instance ("`PPAK ID: 0`" is the Host Security Appliance).
`[-p] -a`	Shows the Accelerated SYN Defender statistics for all interfaces (for each SecureXL instance).
`[-p] -s`	Shows the attack state in short form (for each SecureXL instance).
`[-p] -v`	Shows the attack state in verbose form (for each SecureXL instance).

Examples

Example 1 - Default output before and after enabling the Accelerated SYN Defender

```
[Expert@MyGW:0]# fwaccel synatk monitor
+-------------------------------------------------------------------------------+
| SYN Defender status                                                           |
+-------------------------------------------------------------------------------+
| Configuration                                                        Disabled |
| Status                                                                 Normal |
| Non established connections                                                 0 |
| Global Threshold                                                        10000 |
| Interface Threshold                                                      5000 |
+-------------------------------------------------------------------------------+
| IF              | Topology | Enforce | State (sec)  | Non-established conns |
|                 |          |         |              | Peak      | Current   |
+-------------------------------------------------------------------------------+
| eth0            | External | Disable | Disable      | N/A       | N/A       |
| eth1            | Internal | Disable | Disable      | N/A       | N/A       |
+-------------------------------------------------------------------------------+
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel synatk -m
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel synatk monitor
+-------------------------------------------------------------------------------+
| SYN Defender status                                                           |
+-------------------------------------------------------------------------------+
| Configuration                                                     Monitoring |
| Status                                                                 Normal |
| Non established connections                                                 0 |
| Global Threshold                                                        10000 |
| Interface Threshold                                                      5000 |
+-------------------------------------------------------------------------------+
| IF              | Topology | Enforce | State (sec)  | Non-established conns |
|                 |          |         |              | Peak      | Current   |
+-------------------------------------------------------------------------------+
| eth0            | External | Detect  | Monitor      | 0         | 0         |
| eth1            | Internal | Detect  | Monitor      | 0         | 0         |
+-------------------------------------------------------------------------------+
[Expert@MyGW:0]#
```

Example 2 - Showing the Accelerated SYN Defender status for each SecureXL instance

```
[Expert@MyGW:0]# fwaccel synatk monitor -p
+-------------------------------------------------------------------------------+
| SYN Defender status                                                           |
+-------------------------------------------------------------------------------+
| Configuration                                                      Monitoring |
| Status                                                                 Normal |
| Non established connections                                                 0 |
| Global Threshold                                                        10000 |
| Interface Threshold                                                      5000 |
+-------------------------------------------------------------------------------+
| IF              | Topology | Enforce | State (sec)  | Non-established conns |
|                 |          |         |              | Peak      | Current   |
+-------------------------------------------------------------------------------+
| eth0            | External | Detect  | Monitor      | 0         | 0         |
| eth1            | Internal | Detect  | Monitor      | 0         | 0         |
+-------------------------------------------------------------------------------+

PPAK ID: 0
----------
+-------------------------------------------------------------------------------+
| SYN Defender status                                                           |
+-------------------------------------------------------------------------------+
| Configuration                                                      Monitoring |
| Status                                                                 Normal |
| Non established connections                                                 0 |
| Global Threshold                                                        10000 |
| Interface Threshold                                                      5000 |
+-------------------------------------------------------------------------------+
| IF              | Topology | Enforce | State (sec)  | Non-established conns |
|                 |          |         |              | Peak      | Current   |
+-------------------------------------------------------------------------------+
| eth0            | External | Detect  | Monitor      | 0         | 0         |
| eth1            | Internal | Detect  | Monitor      | 0         | 0         |
+-------------------------------------------------------------------------------+
[Expert@MyGW:0]#
```

Example 3 - Showing the Accelerated SYN Defender statistics for all interfaces and for each SecureXL instance.

```
[Expert@MyGW:0]# fwaccel synatk monitor -p -a
Global:
  status            attached
  nr_active                0

Firewall
----------
Per-interface:
                          eth0        eth1
                      ----------  ----------
topology                External    Internal
state                    Monitor     Monitor
syn ready                      0           0
syn active prev                0           0
syn active curr                0           0
active_score                   0           0
msec grace                     0           0
msec active                    0           0
sent cookies                   0           0
fail validations               0           0
succ validations               0           0
early packets                  0           0
no conn data                   0           0
bogus syn                      0           0
peak non-estab                 0           0
int sent cookies               0           0
int succ validations           0           0
msec interval                  0           0

PPAK ID: 0
----------
Per-interface:
                          eth0        eth1
                      ----------  ----------
topology                External    Internal
state                    Monitor     Monitor
syn ready                      0           0
syn active prev                0           0
syn active curr                0           0
active_score                   0           0
msec grace                     0           0
msec active                    0           0
sent cookies                   0           0
fail validations               0           0
succ validations               0           0
early packets                  0           0
no conn data                   0           0
bogus syn                      0           0
peak non-estab                 0           0
int sent cookies               0           0
int succ validations           0           0
msec interval                  0           0
[Expert@MyGW:0]#
```

Example 4 - Showing the attack state in short form (for each SecureXL instance)

```
[Expert@MyGW:0]# fwaccel synatk monitor -p -s
M,N,0,0

PPAK ID: 0
----------
M,N,0,0
[Expert@MyGW:0]#
```

Example 5 - Showing the attack state in verbose form (for each SecureXL instance)

```
[Expert@MyGW:0]# fwaccel synatk monitor -p -v
+-------------------------------------------------------------------------------+
| SYN Defender statistics                                                       |
+-------------------------------------------------------------------------------+
| Status                                                                Normal  |
| Spoofed SYN/sec                                                            0  |
+-------------------------------------------------------------------------------+
 PPAK ID: 0
----------
+-------------------------------------------------------------------------------+
| SYN Defender statistics                                                       |
+-------------------------------------------------------------------------------+
| Status                                                                Normal  |
| Spoofed SYN/sec                                                            0  |
+-------------------------------------------------------------------------------+
[Expert@MyGW:0]#
```

fwaccel synatk state

Description

The "*fwaccel synatk state*" and "*fwaccel6 synatk state*" commands control the Accelerated SYN Defender states.

The states are independent for IPv4 and IPv6.

Important - This command is **not** intended for end-user usage. Transitions between states (Ready, Grace, and Active) occur automatically. This command provides a way to force temporarily a state transition on an interface or group of interfaces.

Syntax for IPv4

```
fwaccel synatk state
      -h
      -a
      -d
      -g
      -i {all | external | internal | <Name of Interface>}
      -m
      -r
```

Syntax for IPv6

```
fwaccel6 synatk state
      -h
      -a
      -d
      -g
      -i {all | external | internal | <Name of Interface>}
      -m
      -r
```

Parameters

Important - You can specify only one of these parameters: `-a`, `-d`, `-g`, `-m`, or `-r`.

Parameter	Description
`-h`	Shows the applicable built-in usage.
`-a`	Sets the state to Active.
`-d`	Sets the state to Disabled.
`-g`	Sets the state to Grace.
`-i all`	Applies the change to all interfaces (this is the default).
`-i external`	Applies the change only to external interfaces.
`-i internal`	Applies the change only to internal interfaces.
`-i <Name of Interface>`	Applies the change to the specified interface.
`-m`	Sets the state to Monitor (Detect only) mode.
`-r`	Sets the state to Ready.

fwaccel tab

Description

The *fwaccel tab* and *fwaccel6 tab* commands show the contents of the specified SecureXL kernel table.

Notes:

- Dynamic tables, such as the `connections` table can change while this command prints their contents.
  This may cause some values to be missed or reported twice.
- For some tables, the command prints their contents on the screen.
- For some tables, the command prints their contents to the `/var/log/messages` file.
- Also, see the `fw tab` command.

Syntax for IPv4

```
fwaccel tab [-f] [-m <Number of Rows>] -t <Name of Kernel Table>
fwaccel tab -s -t <Name of Kernel Table>
```

Syntax for IPv6

```
fwaccel6 tab [-f] [-m <Number of Rows>] -t <Name of Kernel Table>
fwaccel6 tab -s -t <Name of Kernel Table>
```

Parameters

Parameter	Description
No Parameters	Shows the applicable built-in usage.
`-f`	Formats the output. We recommend to always use this parameter.
`-m <Number of Rows>`	Specifies how many rows to show from the kernel table. Note - The command counts from the top of the table. Default : 1000
`-s`	Shows summary information only.

Parameter	Description
`-t <Name of Kernel Table>`	Specifies the kernel table. This command supports only these kernel tables: ■ connections ■ dos_ip_blacklists ■ dos_pbox ■ dos_pbox_violating_ips ■ dos_rate_matches ■ dos_rate_track_src ■ dos_rate_track_src_svc ■ drop_templates ■ frag_table ■ gtp_apns ■ gtp_tunnels ■ if_by_name ■ inbound_SAs ■ invalid_replay_counter ■ ipsec_mtu_icmp ■ mcast_drop_conns ■ outbound_SAs ■ PMTU_table ■ <Profile> ■ reset_table ■ vpn_link_selection ■ vpn_trusted_ifs

Examples

```
[Expert@MyGW:0]# fwaccel tab -f -m 200 -t connections
Table connections is empty
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel tab -t inbound_SAs
Table contents written to /var/log/messages.
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel tab -t outbound_SAs
Table contents written to /var/log/messages.
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel tab -t vpn_link_selection
Table contents written to /var/log/messages.
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel tab -t drop_templates
Table drop_templates is empty
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel tab -t vpn_trusted_ifs
Table contents written to /var/log/messages.
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel tab -t profile
Table contents written to /var/log/messages.
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel tab -t mcast_drop_conns
Table contents written to /var/log/messages.
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel tab -t invalid_replay_counter
Table contents written to /var/log/messages.
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel tab -t ipsec_mtu_icmp
Table contents written to /var/log/messages.
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel tab -t gtp_tunnels
Table contents written to /var/log/messages.
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel tab -t gtp_apns
Table contents written to /var/log/messages.
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel tab -t if_by_name
Table contents written to /var/log/messages.
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel tab -t PMTU_table
Table PMTU_table is empty
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel tab -t frag_table
Table frag_table is empty
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel tab -t reset_table
Table reset_table is empty
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel tab -t dos_ip_blacklists
Table dos_ip_blacklists is not active for SecureXL device 0.
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel tab -t dos_pbox
Table dos_pbox is not active for SecureXL device 0.
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel tab -t dos_rate_matches
Table dos_rate_matches is not active for SecureXL device 0.
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel tab -t dos_rate_track_src
Table dos_rate_track_src is not active for SecureXL device 0.
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel tab -t dos_rate_track_src_svc
Table dos_rate_track_src_svc is not active for SecureXL device 0.
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel tab -t dos_pbox_violating_ips
Table dos_pbox_violating_ips is not active for SecureXL device 0.
[Expert@MyGW:0]#
```

fwaccel templates

Description

The *fwaccel templates* and *fwaccel6 templates* commands show the contents of the SecureXL templates tables:

- Accept Templates
- SecureXL Drop Templates
- Firewall Drop Templates

Important:

- By default, all Drop Templates are disabled.
  To enable the SecureXL Drop Templates and Firewall Drop Templates:
  1. In SmartConsole, open the Security Gateway / Cluster object.
  2. In the left tree, click the **Optimizations** pane.
  3. Select **Enable drop optimization**.
  4. Click **OK**.
  5. Install the Access Control policy.
- To disable only the Firewall Drop Templates:
  1. Connect to the command line on the Security Gateway / each Cluster Member.
  2. Set the value of the kernel parameter `fw_drop_templates_enabled` to `0` (zero):
     ```
     fw ctl set -f int fw_drop_templates_enabled 0
     ```

  To enable the Firewall Drop Templates again, run:
  ```
  fw ctl set -f int fw_drop_templates_enabled 1
  ```

Important - Based on the number of current templates, these commands can consume memory at very high level.

Syntax for IPv4

```
fwaccel [-i <SecureXL ID>] templates
      [-h]
      [-c]
      [-d [{-R | -S}]
      [-m <Number of Rows>]
      [-r]
      [-R]
      [-s]
      [-S]
```

Syntax for IPv6

```
fwaccel6 templates
      [-h]
      [-c]
      [-d [{-R | -S}]
      [-m <Number of Rows>]
      [-r]
      [-R]
      [-s]
      [-S]
```

Parameters

Parameter	Description
`-i <SecureXL ID>`	Specifies the SecureXL instance ID (for IPv4 only).
No Parameters	Shows the contents of the SecureXL Accept Templates table (Table Name - `cphwd_tmpl`, Table ID - 8111).
`-h`	Shows the applicable built-in usage.
`-c`	Shows the contents of the Firewall Drop Templates table (connections).
`-d`	Shows the contents of the SecureXL Drop Templates table.
`-d -R`	Shows statistics for the reasons for failures to create SecureXL Drop Templates.
`-d -S`	Shows statistics for the SecureXL Drop Templates.
`-m <Number of Rows>`	Specifies how many rows to show from the templates table. Note - The command counts from the top of the table. Default : 1000
`-r`	Shows the contents of the Firewall Drop Templates table (ranges).
`-R`	Shows the reasons for failures to create SecureXL Connections Templates.
`-s`	Shows the total number of offloaded SecureXL Connections Templates.
`-S`	Shows statistics for the offloaded SecureXL Connections Templates.

Accept Templates flags

One or more of these flags appears in the output:

Flag	Description
A	Connection is accounted (SecureXL counts the number of packets and bytes).
B	Connection is created for a rule that contains an Identity Awareness object, or for a rule below that rule.
E	Connection is created for a NAT rule that contains an Identity Awareness object.
I	Identity Awareness (NAC) is enabled for this connection.
M	Connection is created for a rule that contains a Domain object, or for a rule below that rule.
N	Connection undergoes NAT.
O	Connection is created for a rule that contains a Dynamic object, or for a rule below that rule.
P	Connection is created for a rule that may match a Service with an explicitly configured Source port.
Q	QoS is enabled for this connection.
R	Connection is created for a rule that contains a Traceroute object, or for a rule below that rule.
S	PXL (combination of SecureXL and PSL (Passive Streaming Library)) is enabled for this connection.
T	Connection is created for a rule that contains a Time object, or for a rule below that rule.
U	Connection is unidirectional.
X	Connection is created for a NAT rule that contains a translated Dynamic object.
Z	Connection is created for a rule that contains a Security Zone object, or for a rule below that rule.

Drop Templates flags

One or more of these flags appears in the output:

Flag	Description
D	SecureXL Drop template exists for this connection.
L	Log and Drop action for this connection.

Firewall Drop Templates flags

One or more of these flags appears in the output:

Flag	Description
M	Connection is created for a rule that contains a Domain object, or for a rule below that rule.
O	Connection is created for a rule that contains a Dynamic object, or for a rule below that rule.
U	Connection is created for a rule that contains a Updatable object, or for a rule below that rule.

Examples

Example 1 - Default output

```
[Expert@MyGW:0]# fwaccel templates
Source          SPort Destination     DPort PR Flags         LCT  DLY C2S i/f S2C i/f
--------------- ----- --------------- ----- -- ------------  ---- --- ------- -------
192.168.10.20       * 192.168.10.50      80  6            0     0   0 eth5/eth1 eth1/eth5
[Expert@MyGW:0]#
```

Example 2 - Drop Templates

```
[Expert@MyGW:0]# fwaccel templates -d
The SecureXL drop templates table is empty
[Expert@MyGW:0]#
```

Example 3 - Summary of SecureXL Connections Templates

```
[Expert@MyGW:0]# fwaccel templates -s
Total number of templates: 1
[Expert@MyGW:0]#
```

Example 4 - Templates statistics

```
[Expert@MyGW:0]# fwaccel templates -S

Templates stats:

 Name                     Value           Name                     Value
--------------------     ------------    --------------------     ------------
C templates                          0    conns from templates                 0
nat templates                        0    conns from nat tmpl                  0
C CPASXL templates                   0    C PSLXL templates                    0
C used templates                     0    cpasxl tmpl conns                    0
pslxl tmpl conns                     0    C conns from tmpl                    0

[Expert@MyGW:0]#
```

fwaccel ver

Description

Shows this information:

- Firewall Version and Build
- Accelerator Version
- Firewall API version
- Accelerator API version

Syntax

```
fwaccel ver
```

Example

```
Expert@MyGW:0]# fwaccel ver
Firewall version: R82 - Build 123
Acceleration Device: Performance Pack
Accelerator Version 2.1
Firewall API version: 3.0NG (19/11/2015)
Accelerator API version: 3.0NG (19/11/2015)
[Expert@MyGW:0]#
```

fw monitor

Description

Firewall Monitor is the Check Point traffic capture tool.

In a Security Gateway, traffic passes through different inspection points - Chain Modules in the Inbound direction and then in the Outbound direction (see the "`fw ctl chain`" command.

The FW Monitor tool captures the traffic at each Chain Module in both directions.

You can later analyze the captured traffic with the same FW Monitor tool, or with special tools like Wireshark.

Notes:

- Only one instance of "`fw monitor`" can run at a time.
- You can stop the "`fw monitor`" instance in one of these ways:
  - In the shell, in which the "`fw monitor`" instance runs, press **CTRL + C** keys
  - In another shell, run this command: `fw monitor -U`
- Each time you run the FW Monitor, it compiles its temporary policy files (`$FWDIR/tmp/monitorfilter.*`).
- The FW Monitor also show the traffic accelerated with SecureXL.
- For more information, see sk30583.

Important:

- You can run this command in the Expert mode or in Gaia Clish (Gaia gClish on Scalable Platforms).
- On Scalable Platforms, you must connect to the applicable Security Group.

Syntax for IPv4

```
fw monitor {-h | -help}
```

```
fw monitor [-d] [-D] [-ci <Number of Inbound Packets>] [-co
<Number of Outbound Packets>] [-e <INSPECT Expression> | -f
{<INSPECT Filter File> | -}] [-F "<Source IP>,<Source Port>,<Dest
IP>,<Dest Port>,<Protocol Number>"] [-H "<IP Address>"] [-i] [-l
<Length>] [-m {i,I,o,O,e,E}] [-o <Output File> [-w]] [[-pi
<Position>] [-pI <Position>] [-po <Position>] [-pO <Position>] | -
p all [-a]] [-T] [-u | -s] [-U] [-v <VSID>] [-x <Offset>
[,<Length>] [-w]] [-b <Size of Buffer>]
```

Syntax for IPv6

```
fw6 monitor {-h | -help}
```

```
fw6 monitor [-d] [-D] [-ci <Number of Inbound Packets>] [-co
<Number of Outbound Packets>] [-e <INSPECT Expression> | -f
{<INSPECT Filter File> | -}] [-F "<Source IP>,<Source Port>,<Dest
IP>,<Dest Port>,<Protocol Number>"] [-H "<IP Address>"] [-i] [-l
<Length>] [-m {i,I,o,O,e,E}] [-o <Output File> [-w]] [[-pi
<Position>] [-pI <Position>] [-po <Position>] [-pO <Position>] | -
p all [-a]] [-T] [-u | -s] [-U] [-v <VSID>] [-x <Offset>
[,<Length>] [-w]] [-b <Size of Buffer>]
```

Parameters

Parameter	Description
`{-h | -help}`	Shows the built-in usage.
`-b <Size of Buffer>`	Specifies the size of the memory buffer for FW Monitor. An approximate formula for total required RAM memory: `(4) x (Number of active CoreXL Firewall instances) x (Size of Buffer)` The default FW Monitor algorithm for allocating the memory buffer required to capture traffic: 1. Try to allocate the memory buffer with the default size. 2. If the available RAM is not enough, try allocate the memory buffer with a smaller size. If the available RAM is not enough on a Security Gateway, FW Monitor can fail to allocate the required memory buffer. With this parameter, you can specify the size of the memory buffer to "help" FW Monitor in its caculations. **Note** - Each CoreXL Firewall instance has two internal buffers, and there is one general buffer that merges the data from all CoreXL Firewall instances. **Default:** 8192 (in kilobytes) **Range:** 128 - 8192 (in kilobytes)

Parameter	Description
`-d` `-D`	Runs the command in debug mode and shows some information about how the FW Monitor starts and compiles the specified INSPECT filter: ■ `-d` Simple debug output. ■ `-D` Verbose output. **Note** - You can specify both parameters to show more information.
`-ci <Number of Inbound Packets>` `-co <Number of Outbound Packets>`	Specifies how many packets to capture. The FW Monitor stops the traffic capture if it counted the specified number of packets. ■ `-ci` Specifies the number of inbound packets to count. ■ `-co` Specifies the number of outbound packets to count **Best Practice** - You can use the "`-ci`" and the "`-co`" parameters together. This is especially useful during large volumes of traffic. In such scenarios, FW Monitor may bind so many resources (for writing to the console, or to a file) that recognizing the break sequence (**CTRL+C**) might take a very long time.

Parameter	Description
`-e <INSPECT Expression>` *or* `-f {<INSPECT Filter File> | -}`	Captures only specific packets of non-accelerated traffic: ■ "`-e <INSPECT Expression>`" Defines the INSPECT filter expression on the command line. ■ "`-f <INSPECT Filter File>`" Reads the INSPECT filter expression from the specified file. You must enter the full path and name of the plain-text file that contains the INSPECT filter expression. ■ "`-f -`" Reads the INSPECT filter expression from the standard input. After you enter the INSPECT filter expression, you must enter the `^D` (**CTRL+D**) as the EOF (End Of File) character. **Warning** - These INSPECT filters do **not** apply to the accelerated traffic. **Important** - Make sure to enclose the INSPECT filter expression correctly in single quotes (ASCII value 39) or double quotes (ASCII value 34). **Notes:** ■ Refer to the `$FWDIR/lib/fwmonitor.def` file for useful macro definitions. ■ See syntax examples below (*"Examples for the "-e" parameter" on page 214*).
`-F "<Source IP>,<Source Port>,<Dest IP>,<Dest Port>,<Protocol Number>"`	Specifies the capture filter (for both accelerated and non-accelerated traffic): ■ `<Source IP>` - Specifies the source IP address ■ `<Source Port>` - Specifies the source Port Number (see *IANA Service Name and Port Number Registry*) ■ `<Dest IP>` - Specifies the destination IP address ■ `<Dest Port>` - Specifies the destination Port Number (see *IANA Service Name and Port Number Registry*) ■ `<Protocol Number>` - Specifies the Protocol Number (see *IANA Protocol Numbers*)

Parameter	Description
	Notes: ■ See syntax examples below (*"Examples for the "-F" parameter" on page 228*). ■ The "`-F`" parameter uses the Kernel Debug Filters. For more information, see Kernel Debug Filters. • For the Source IP address: `simple_debug_filter_saddr_<N> "<IP Address>"` • For the Source Ports: `simple_debug_filter_sport_<N> <1-65535>` • For the Destination IP address: `simple_debug_filter_daddr_<N> "<IP Address>"` • For the Destination Ports: `simple_debug_filter_dport_<N> <1-65535>` • For the Protocol Number: `simple_debug_filter_proto_<N> <0-254>` ■ Value 0 means "`any`". ■ This parameter supports up to 5 capture filters (up to 5 instances of the "`-F`" parameter in the syntax). The FW Monitor performs the logical "OR" between all specified simple capture filters.
`-H "<IP Address>"`	Creates an IP address filter. For more information, see Kernel Debug Filters. This parameter supports up to 3 capture filters (up to 3 instances of the "`-H`" parameter in the syntax). Example - Capture traffic only to and from the Host 1.1.1.1: `fw monitor -H "1.1.1.1"`

Parameter	Description
`-i`	Flushes the standard output. **Note** - This parameter is valid only with the "`-v <VSID>`" parameter. **Best Practice** - Use this parameter to make sure FW Monitor immediately writes the captured data for each packet to the standard output. This is especially useful if you want to kill a running FW Monitor process, and want to be sure that FW Monitor writes all the data to the specified file.
`-l <Length>`	Specifies the maximum length of the captured packets. FW Monitor reads only the specified number of bytes from each packet. **Notes:** ■ This parameter is optional. ■ With this parameter you can capture only the headers from each packet (for example, IP and TCP) and omit the payload. This decreases the size of the output file. This also helps the internal FW Monitor buffer not to fill too fast. ■ Make sure to capture the minimum required number of bytes, to capture the Layer 3 IP header and Layer 4 Transport header.

Parameter	Description
`-m {i, I, o, O, e, E}`	Specifies the capture mask (inspection point) in relation to Chain Modules, in which the FW Monitor captures the traffic. These are the inspection points, through which each packet passes on a Security Gateway. ■ `-m i` Pre-Inbound only (before the packet enters a Chain Module in the inbound direction) ■ `-m I` Post-Inbound only (after the packet passes a Chain Module in the inbound direction) ■ `-m o` Pre-Outbound only (before the packet enters a Chain Module in the outbound direction) ■ `-m O` Post-Outbound only (after the packet passes through a Chain Module in the outbound direction) ■ `-m e` Pre-Outbound VPN only (before the packet enters a VPN Chain Module in the outbound direction) ■ `-m E` Post-Outbound VPN only (after the packet passes through a VPN Chain Module in the outbound direction)

Parameter	Description
	Notes: ■ You can specify several capture masks (for example, to see NAT on the egress packets, enter "`... -m o O ...`"). ■ You can use this capture mask parameter "`-m {i, I, o, O, e, E}`" together with the chain module position parameter "`-p{i | I | o | O}`". ■ In the inbound direction: • All chain positions *before* the FireWall Virtual Machine module are Pre-Inbound (the "`fw ctl chain`" command shows this module as "`fw VM inbound`"). • All chain modules *after* the FireWall Virtual Machine module are Post-Inbound. ■ In the outbound direction: • All chain position *before* the FireWall Virtual Machine module are Pre-Outbound. • All chain modules *after* the FireWall Virtual Machine module are Post-Outbound. ■ By default, the FW Monitor captures the traffic only in the FireWall Virtual Machine module. ■ The packet direction relates to each specific packet, and not to the connection's direction. ■ The letters "`q`" and "`Q`" after the inspection point mean that the QoS policy is applied to the interface.
	Example packet flows: ■ From a Client to a Server through the FireWall Virtual Machine module: `[Client] --> ("i") {FW VM attached to eth1} ("I") [Security Gateway] ("o") {FW VM attached to eth2} ("O") --> [Server]` ■ From a Server to a Client through the FireWall Virtual Machine module: `[Client] <-- ("O") {FW VM attached to eth1} ("o") [Security Gateway] ("I") {FW VM attached to eth2} ("i") <-- [Server]`

Parameter	Description
`-o <Output File>`	Specifies the output file, to which FW Monitor writes the captured raw data. **Important** - If you do not specify the path explicitly, FW Monitor creates this output file in the current working directory. Because this output file can grow very fast to very large size, we always recommend to specify the full path to the largest partition `/var/log/`. The format of this output file is the same format used by tools like `snoop` (refer to RFC 1761). You can later analyze the captured traffic with the same FW Monitor tool, or with special tools like Wireshark.
`-pi <Position>` `-pI <Position>` `-po <Position>` `-pO <Position>` *or* `-p all [-a]`	Inserts the FW Monitor Chain Module at the specified position between the kernel Chain Modules (see the "`fw ctl chain`" command). If the FW Monitor writes the captured data to the specified output file (with the parameter "`-o <Output File>`"), it also writes the position of the FW Monitor chain module as one of the fields. You can insert the FW Monitor Chain Module in these positions only: ■ `-pi <Position>` Inserts the FW Monitor Chain Module in the specified Pre-Inbound position. ■ `-pI <Position>` Inserts the FW Monitor Chain Module in the specified Post-Inbound position. ■ `-po <Position>` Inserts the FW Monitor Chain Module in the specified Pre-Outbound position. ■ `-pO <Position>` Inserts the FW Monitor Chain Module in the specified Post-Outbound position ■ `-p all [-a]` Inserts the FW Monitor Chain Module at all positions (both Inbound and Outbound). **Warning** - This parameter causes very high load on the CPU, but provides the most complete traffic capture. The "`-a`" parameter specifies to use absolute chain positions. This parameter changes the chain ID from a relative value (which only makes sense with the matching output from the "`fw ctl chain`" command) to an absolute value.

Parameter	Description
	Notes: ■ *<Position>* can be one of these: • A relative position number In the output of the "`fw ctl chain`" command, refer to the numbers in the leftmost column (for example, 0, 5, 14). • A relative position alias In the output of the "`fw ctl chain`" command, refer to the internal chain module names in the rightmost column in the parentheses (for example, `sxl_in`, `fw`, `cpas`). • An absolute position In the output of the "`fw ctl chain`" command, refer to the numbers in the second column from the left (for example, -7fffffff, -1fffff8, 7f730000). In the syntax, you must write these numbers in the hexadecimal format (for example, -0x7fffffff, -0x1fffff8, 0x7f730000). ■ You can use this chain module position parameter "`-p{i | I| o | O} ...`" together with the capture mask parameter "`-m {i, I, o, O, e, E}`". ■ In the inbound direction: • All chain positions *before* the FireWall Virtual Machine module are Pre-Inbound (the "`fw ctl chain`" command shows this module as "`fw VM inbound`"). • All chain modules *after* the FireWall Virtual Machine module are Post-Inbound. ■ In the outbound direction: • All chain position *before* the FireWall Virtual Machine module are Pre-Outbound. • All chain modules *after* the FireWall Virtual Machine module are Post-Outbound. ■ By default, the FW Monitor captures the traffic only in the FireWall Virtual Machine module. ■ The chain module position parameters "`-p{i | I| o | O} ...`" parameters do **not** apply to the accelerated traffic, which is still monitored at the default inbound and outbound positions. ■ For more information about the inspection points, see the applicable table below.

Parameter	Description
`-T`	Shows the timestamp for each packet: `DDMMMYYYY HH:MM:SS.mmmmmm` **Best Practice** - Use this parameter if you do not save the output to a file, but print it on the screen.
`-u` *or* `-s`	Shows UUID for each packet (it is only possible to print either the UUID, or the SUUID - not both): ■ `-u` Prints connection's Universal-Unique-ID (UUID) for each packet ■ `-s` Prints connection's Session UUID (SUUID) for each packet
`-U`	Removes the simple capture filters specified with this parameter: `-F "<Source IP>,<Source Port>,<Dest IP>,<Dest Port>,<Protocol Number>"`
`-v <VSID>`	On a VSX Gateway or VSX Cluster Member, captures the packets on the specified Virtual System or Virtual Router. By default, FW Monitor captures the packets on all Virtual Systems and Virtual Routers. Example: `fw monitor -v 4 -e "accept;" -o /var/log/fw_mon.cap`
`-w`	Captures the entire packet, instead of only the header. Must be used together with one of these parameters: ■ `-o <Output File>` ■ `-x <Offset>[,<Length>]`

Parameter	Description
`-x <Offset> [,<Length>]`	Specifies the position in each packet, where the FW Monitor starts to capture the data from each packet. Optionally, it is also possible to limit the amount of data the FW Monitor captures. ■ *<Offset>* Specifies how many bytes to skip from the beginning of each packet. FW Monitor starts to capture the data from each packet only after the specified number of bytes. ■ *<Length>* Specifies the maximum length of the captured packets. FW Monitor reads only the specified number of bytes from each packet. For example, to skip over the IP header and TCP header, enter "`-x 52,96`"

Inspection points in Security Gateway and in the FW Monitor output

Note - The Inbound and Outbound traffic direction relates to each specific packet, and not to the connection.

- *Inbound*

Name of inspection point	Relation to the FireWall Virtual Machine	Notion of inspection point in the FW Monitor output
Pre-Inbound	Before the inbound FireWall VM	i (for example, `eth4:i`)
Post-Inbound	After the inbound FireWall VM	I (for example, `eth4:I`)
Pre-Inbound VPN	Inbound before decrypt	id (for example, `eth4:id`)
Post-Inbound VPN	Inbound after decrypt	ID (for example, `eth4:ID`)
Pre-Inbound QoS	Inbound before QoS	iq (for example, `eth4:iq`)
Post-Inbound QoS	Inbound after QoS	IQ (for example, `eth4:IQ`)

- *Outbound*

Name of inspection point	Relation to the FireWall Virtual Machine	Notion of inspection point in the FW Monitor output
Pre-Outbound	Before the outbound FireWall VM	o (for example, `eth4:o`)
Post-Outbound	After the outbound FireWall VM	O (for example, `eth4:O`)
Pre-Outbound VPN	Outbound before encrypt	e (for example, `eth4:e`)
Post-Outbound VPN	Outbound after encrypt	E (for example, `eth4:E`)
Pre-Outbound QoS	Outbound before QoS	oq (for example, `eth4:oq`)
Post-Outbound QoS	Outbound after QoS	OQ (for example, `eth4:OQ`)

Generic Examples

Example 1 - Default syntax

```
[Expert@MyGW:0]# fw monitor
 monitor: getting filter (from command line)
 monitor: compiling
monitorfilter:
Compiled OK.
 monitor: loading
 monitor: monitoring (control-C to stop)
[vs_0][fw_1] eth0:i[40]: 192.168.204.1 -> 192.168.204.40 (TCP) len=40 id=31789
TCP: 53901 -> 22 ....A. seq=761113cd ack=f92e2a13
[vs_0][fw_1] eth0:I[40]: 192.168.204.1 -> 192.168.204.40 (TCP) len=40 id=31789
TCP: 53901 -> 22 ....A. seq=761113cd ack=f92e2a13
[vs_0][fw_1] eth0:i[40]: 192.168.204.1 -> 192.168.204.40 (TCP) len=40 id=31790
TCP: 53901 -> 22 ....A. seq=761113cd ack=f92e2a47
... ... ...
 monitor: caught sig 2
 monitor: unloading
[Expert@MyGW:0]#
```

Example 2 - Showing timestamps in the output for each packet

```
[Expert@MyGW:0]# fw monitor -T
 monitor: getting filter (from command line)
 monitor: compiling
monitorfilter:
Compiled OK.
 monitor: loading
 monitor: monitoring (control-C to stop)
[vs_0][fw_1] 12Sep2018 19:08:05.453947 eth0:oq[124]: 192.168.3.53 -> 172.20.168.16 (TCP)
len=124 id=38414
TCP: 22 -> 64424 ...PA. seq=1c23924a ack=3c951092
[vs_0][fw_1] 12Sep2018 19:08:05.453960 eth0:OQ[124]: 192.168.3.53 -> 172.20.168.16 (TCP)
len=124 id=38414
TCP: 22 -> 64424 ...PA. seq=1c23924a ack=3c951092
[vs_0][fw_1] 12Sep2018 19:08:05.454059 eth0:oq[252]: 192.168.3.53 -> 172.20.168.16 (TCP)
len=252 id=38415
TCP: 22 -> 64424 ...PA. seq=1c23929e ack=3c951092
[vs_0][fw_1] 12Sep2018 19:08:05.454064 eth0:OQ[252]: 192.168.3.53 -> 172.20.168.16 (TCP)
len=252 id=38415
TCP: 22 -> 64424 ...PA. seq=1c23929e ack=3c951092
[vs_0][fw_1] 12Sep2018 19:08:05.454072 eth0:oq[252]: 192.168.3.53 -> 172.20.168.16 (TCP)
len=252 id=38416
TCP: 22 -> 64424 ...PA. seq=1c239372 ack=3c951092
[vs_0][fw_1] 12Sep2018 19:08:05.454074 eth0:OQ[252]: 192.168.3.53 -> 172.20.168.16 (TCP)
len=252 id=38416
TCP: 22 -> 64424 ...PA. seq=1c239372 ack=3c951092
[vs_0][fw_1] 12Sep2018 19:08:05.463165 eth0:iq[40]: 172.20.168.16 -> 192.168.3.53 (TCP)
len=40 id=17398
TCP: 64424 -> 22 ....A. seq=3c951092 ack=1c239446
[vs_0][fw_1] 12Sep2018 19:08:05.463177 eth0:IQ[40]: 172.20.168.16 -> 192.168.3.53 (TCP)
len=40 id=17398
TCP: 64424 -> 22 ....A. seq=3c951092 ack=1c239446
 monitor: unloading
[Expert@MyGW:0]#
```

Example 3 - Capturing only three Pre-Inbound packets at the FireWall Virtual Machine module

```
[Expert@MyGW:0]# fw monitor -m i -ci 3
 monitor: getting filter (from command line)
 monitor: compiling
monitorfilter:
Compiled OK.
 monitor: loading
 monitor: monitoring (control-C to stop)
[vs_0][fw_1] eth0:i[40]: 192.168.204.1 -> 192.168.204.40 (TCP) len=40 id=31905
TCP: 53901 -> 22 ....A. seq=76111bb5 ack=f92e683b
[vs_0][fw_1] eth0:i[40]: 192.168.204.1 -> 192.168.204.40 (TCP) len=40 id=31906
TCP: 53901 -> 22 ....A. seq=76111bb5 ack=f92e68ef
[vs_0][fw_1] eth0:i[40]: 192.168.204.1 -> 192.168.204.40 (TCP) len=40 id=31907
TCP: 53901 -> 22 ....A. seq=76111bb5 ack=f92e69a3
 monitor: unloading
Read 3 inbound packets and 0 outbound packets
[Expert@MyGW:0]#
```

Example 4 - Inserting the FW Monitor chain is before the chain #2 and capture only three Pre-Inbound packets

```
[Expert@MyGW:0]# fw ctl chain
in chain (15):
        0: -7fffffff (0000000000000000) (00000000) SecureXL inbound (sxl_in)
        1: -7ffffffe (0000000000000000) (00000000) SecureXL inbound CT (sxl_ct)
        2: -7f800000 (ffffffff8b6718c0) (ffffffff) IP Options Strip (in) (ipopt_strip)
        3: - 1fffff8 (ffffffff8b66f6f0) (00000001) Stateless verifications (in) (asm)
        4: - 1fffff7 (ffffffff8b66f210) (00000001) fw multik misc proto forwarding
        5:         0 (ffffffff8b8506a0) (00000001) fw VM inbound  (fw)
        6:         2 (ffffffff8b671d10) (00000001) fw SCV inbound (scv)
        7:         4 (ffffffff8b061ed0) (00000003) QoS inbound offload chain module
        8:         5 (ffffffff8b564d30) (00000003) fw offload inbound (offload_in)
        9:        10 (ffffffff8b842710) (00000001) fw post VM inbound  (post_vm)
        10:    100000 (ffffffff8b7fd6c0) (00000001) fw accounting inbound (acct)
        11:  22000000 (ffffffff8b0638d0) (00000003) QoS slowpath inbound chain mod (fg_sched)
        12:  7f730000 (ffffffff8b3c40b0) (00000001) passive streaming (in) (pass_str)
        13:  7f750000 (ffffffff8b0e5b40) (00000001) TCP streaming (in) (cpas)
        14:  7f800000 (ffffffff8b671870) (ffffffff) IP Options Restore (in) (ipopt_res)
out chain (14):
        0: -7f800000 (ffffffff8b6718c0) (ffffffff) IP Options Strip (out) (ipopt_strip)
        1: - 1fffff0 (ffffffff8b0d0190) (00000001) TCP streaming (out) (cpas)
        2: - 1ffff50 (ffffffff8b3c40b0) (00000001) passive streaming (out) (pass_str)
        3: - 1f00000 (ffffffff8b66f6f0) (00000001) Stateless verifications (out) (asm)
        4: -     1ff (ffffffff8aeec0a0) (00000001) NAC Packet Outbound (nac_tag)
        5:         0 (ffffffff8b8506a0) (00000001) fw VM outbound (fw)
        6:        10 (ffffffff8b842710) (00000001) fw post VM outbound  (post_vm)
        7:  15000000 (ffffffff8b062540) (00000003) QoS outbound offload chain modul (fg_pol)
        8:  21000000 (ffffffff8b0638d0) (00000003) QoS slowpath outbound chain mod (fg_sched)
        9:  7f000000 (ffffffff8b7fd6c0) (00000001) fw accounting outbound (acct)
        10:  7f700000 (ffffffff8b0e4660) (00000001) TCP streaming post VM (cpas)
        11:  7f800000 (ffffffff8b671870) (ffffffff) IP Options Restore (out) (ipopt_res)
        12:  7f900000 (0000000000000000) (00000000) SecureXL outbound (sxl_out)
        13:  7fa00000 (0000000000000000) (00000000) SecureXL deliver (sxl_deliver)
[Expert@MyGW:0]#
[Expert@MyGW:0]# fw monitor -pi 2 -ci 3
 monitor: getting filter (from command line)
 monitor: compiling
monitorfilter:
Compiled OK.
 monitor: loading
in chain (17):
        0: -7fffffff (0000000000000000) (00000000) SecureXL inbound (sxl_in)
        1: -7ffffffe (0000000000000000) (00000000) SecureXL inbound CT (sxl_ct)
        2: -7f800001 (ffffffff8b6774d0) (ffffffff) fwmonitor (i/f side)
        3: -7f800000 (ffffffff8b6718c0) (ffffffff) IP Options Strip (in) (ipopt_strip)
        4: - 1fffff8 (ffffffff8b66f6f0) (00000001) Stateless verifications (in) (asm)
        5: - 1fffff7 (ffffffff8b66f210) (00000001) fw multik misc proto forwarding
        6:         0 (ffffffff8b8506a0) (00000001) fw VM inbound  (fw)
        7:         2 (ffffffff8b671d10) (00000001) fw SCV inbound (scv)
        8:         4 (ffffffff8b061ed0) (00000003) QoS inbound offload chain module
        9:         5 (ffffffff8b564d30) (00000003) fw offload inbound (offload_in)
        10:        10 (ffffffff8b842710) (00000001) fw post VM inbound  (post_vm)
        11:    100000 (ffffffff8b7fd6c0) (00000001) fw accounting inbound (acct)
        12:  22000000 (ffffffff8b0638d0) (00000003) QoS slowpath inbound chain mod (fg_sched)
        13:  70000000 (ffffffff8b6774d0) (ffffffff) fwmonitor (IP  side)
        14:  7f730000 (ffffffff8b3c40b0) (00000001) passive streaming (in) (pass_str)
        15:  7f750000 (ffffffff8b0e5b40) (00000001) TCP streaming (in) (cpas)
        16:  7f800000 (ffffffff8b671870) (ffffffff) IP Options Restore (in) (ipopt_res)
out chain (16):
        0: -7f800000 (ffffffff8b6718c0) (ffffffff) IP Options Strip (out) (ipopt_strip)
        1: -70000000 (ffffffff8b6774d0) (ffffffff) fwmonitor (i/f side)
        2: - 1fffff0 (ffffffff8b0d0190) (00000001) TCP streaming (out) (cpas)
        3: - 1ffff50 (ffffffff8b3c40b0) (00000001) passive streaming (out) (pass_str)
        4: - 1f00000 (ffffffff8b66f6f0) (00000001) Stateless verifications (out) (asm)
        5: -     1ff (ffffffff8aeec0a0) (00000001) NAC Packet Outbound (nac_tag)
        6:         0 (ffffffff8b8506a0) (00000001) fw VM outbound (fw)
        7:        10 (ffffffff8b842710) (00000001) fw post VM outbound  (post_vm)
        8:  15000000 (ffffffff8b062540) (00000003) QoS outbound offload chain modul (fg_pol)
        9:  21000000 (ffffffff8b0638d0) (00000003) QoS slowpath outbound chain mod (fg_sched)
        10:  70000000 (ffffffff8b6774d0) (ffffffff) fwmonitor (IP side)
        11:  7f000000 (ffffffff8b7fd6c0) (00000001) fw accounting outbound (acct)
        12:  7f700000 (ffffffff8b0e4660) (00000001) TCP streaming post VM (cpas)
```

```
        13:  7f800000 (ffffffff8b671870) (ffffffff) IP Options Restore (out) (ipopt_res)
        14:  7f900000 (0000000000000000) (00000000) SecureXL outbound (sxl_out)
        15:  7fa00000 (0000000000000000) (00000000) SecureXL deliver (sxl_deliver)
 monitor: monitoring (control-C to stop)
[vs_0][fw_1] eth0:oq1 (TCP streaming (out))[1228]: 192.168.204.40 -> 192.168.204.1 (TCP)
len=1228 id=37575
TCP: 22 -> 51702 ...PA. seq=34e2af31 ack=e6c995ce
[vs_0][fw_1] eth0:OQ10 (TCP streaming post VM)[1228]: 192.168.204.40 -> 192.168.204.1 (TCP)
len=1228 id=37575
TCP: 22 -> 51702 ...PA. seq=34e2af31 ack=e6c995ce
[vs_0][fw_1] eth0:iq2 (IP Options Strip (in))[40]: 192.168.204.1 -> 192.168.204.40 (TCP)
len=40 id=32022
TCP: 51702 -> 22 ....A. seq=e6c995ce ack=34e2af31
[vs_0][fw_1] eth0:IQ13 (TCP streaming (in))[40]: 192.168.204.1 -> 192.168.204.40 (TCP) len=40
id=32022
TCP: 51702 -> 22 ....A. seq=e6c995ce ack=34e2af31
[vs_0][fw_1] eth0:oq1 (TCP streaming (out))[1356]: 192.168.204.40 -> 192.168.204.1 (TCP)
len=1356 id=37576
TCP: 22 -> 51702 ...PA. seq=34e2b3d5 ack=e6c995ce
[vs_0][fw_1] eth0:OQ10 (TCP streaming post VM)[1356]: 192.168.204.40 -> 192.168.204.1 (TCP)
len=1356 id=37576
TCP: 22 -> 51702 ...PA. seq=34e2b3d5 ack=e6c995ce
[vs_0][fw_1] eth0:iq2 (IP Options Strip (in))[40]: 192.168.204.1 -> 192.168.204.40 (TCP)
len=40 id=32023
TCP: 51702 -> 22 ....A. seq=e6c995ce ack=34e2b8f9
[vs_0][fw_1] eth0:IQ13 (TCP streaming (in))[40]: 192.168.204.1 -> 192.168.204.40 (TCP) len=40
id=32023
TCP: 51702 -> 22 ....A. seq=e6c995ce ack=34e2b8f9
[vs_0][fw_1] eth0:oq1 (TCP streaming (out))[1356]: 192.168.204.40 -> 192.168.204.1 (TCP)
len=1356 id=37577
TCP: 22 -> 51702 ...PA. seq=34e2b8f9 ack=e6c995ce
[vs_0][fw_1] eth0:OQ10 (TCP streaming post VM)[1356]: 192.168.204.40 -> 192.168.204.1 (TCP)
len=1356 id=37577
TCP: 22 -> 51702 ...PA. seq=34e2b8f9 ack=e6c995ce
[vs_0][fw_1] eth0:oq1 (TCP streaming (out))[412]: 192.168.204.40 -> 192.168.204.1 (TCP)
len=412 id=37578
TCP: 22 -> 51702 ...PA. seq=34e2be1d ack=e6c995ce
[vs_0][fw_1] eth0:OQ10 (TCP streaming post VM)[412]: 192.168.204.40 -> 192.168.204.1 (TCP)
len=412 id=37578
TCP: 22 -> 51702 ...PA. seq=34e2be1d ack=e6c995ce
[vs_0][fw_1] eth0:iq2 (IP Options Strip (in))[40]: 192.168.204.1 -> 192.168.204.40 (TCP)
len=40 id=32024
TCP: 51702 -> 22 ....A. seq=e6c995ce ack=34e2bf91
[vs_0][fw_1] eth0:IQ13 (TCP streaming (in))[40]: 192.168.204.1 -> 192.168.204.40 (TCP) len=40
id=32024
TCP: 51702 -> 22 ....A. seq=e6c995ce ack=34e2bf91
[vs_0][fw_1] eth0:oq1 (TCP streaming (out))[716]: 192.168.204.40 -> 192.168.204.1 (TCP)
len=716 id=37579
TCP: 22 -> 51702 ...PA. seq=34e2bf91 ack=e6c995ce
[vs_0][fw_1] eth0:OQ10 (TCP streaming post VM)[716]: 192.168.204.40 -> 192.168.204.1 (TCP)
len=716 id=37579
TCP: 22 -> 51702 ...PA. seq=34e2bf91 ack=e6c995ce
 monitor: unloading
Read 3 inbound packets and 5 outbound packets
[Expert@MyGW:0]#
```

Example 5 - Showing list of Chain Modules with the FW Monitor, when you do not change the default capture positions

```
[Expert@MyGW:0]# fw ctl chain
in chain (17):
        0: -7fffffff (0000000000000000) (00000000) SecureXL inbound (sxl_in)
        1: -7ffffffe (0000000000000000) (00000000) SecureXL inbound CT (sxl_ct)
        2: -7f800000 (ffffffff8b6718c0) (ffffffff) IP Options Strip (in) (ipopt_strip)
        3: -70000000 (ffffffff8b6774d0) (ffffffff) fwmonitor (i/f side)
        4: - 1fffff8 (ffffffff8b66f6f0) (00000001) Stateless verifications (in) (asm)
        5: - 1fffff7 (ffffffff8b66f210) (00000001) fw multik misc proto forwarding
        6:         0 (ffffffff8b8506a0) (00000001) fw VM inbound  (fw)
        7:         2 (ffffffff8b671d10) (00000001) fw SCV inbound (scv)
        8:         4 (ffffffff8b061ed0) (00000003) QoS inbound offload chain module
        9:         5 (ffffffff8b564d30) (00000003) fw offload inbound (offload_in)
        10:        10 (ffffffff8b842710) (00000001) fw post VM inbound  (post_vm)
        11:    100000 (ffffffff8b7fd6c0) (00000001) fw accounting inbound (acct)
        12:  22000000 (ffffffff8b0638d0) (00000003) QoS slowpath inbound chain mod (fg_sched)
        13:  70000000 (ffffffff8b6774d0) (ffffffff) fwmonitor (IP  side)
        14:  7f730000 (ffffffff8b3c40b0) (00000001) passive streaming (in) (pass_str)
        15:  7f750000 (ffffffff8b0e5b40) (00000001) TCP streaming (in) (cpas)
        16:  7f800000 (ffffffff8b671870) (ffffffff) IP Options Restore (in) (ipopt_res)
out chain (16):
        0: -7f800000 (ffffffff8b6718c0) (ffffffff) IP Options Strip (out) (ipopt_strip)
        1: -70000000 (ffffffff8b6774d0) (ffffffff) fwmonitor (i/f side)
        2: - 1fffff0 (ffffffff8b0d0190) (00000001) TCP streaming (out) (cpas)
        3: - 1ffff50 (ffffffff8b3c40b0) (00000001) passive streaming (out) (pass_str)
        4: - 1f00000 (ffffffff8b66f6f0) (00000001) Stateless verifications (out) (asm)
        5: -     1ff (ffffffff8aeec0a0) (00000001) NAC Packet Outbound (nac_tag)
        6:         0 (ffffffff8b8506a0) (00000001) fw VM outbound (fw)
        7:        10 (ffffffff8b842710) (00000001) fw post VM outbound  (post_vm)
        8:  15000000 (ffffffff8b062540) (00000003) QoS outbound offload chain modul (fg_pol)
        9:  21000000 (ffffffff8b0638d0) (00000003) QoS slowpath outbound chain mod (fg_sched)
        10:  70000000 (ffffffff8b6774d0) (ffffffff) fwmonitor (IP side)
        11:  7f000000 (ffffffff8b7fd6c0) (00000001) fw accounting outbound (acct)
        12:  7f700000 (ffffffff8b0e4660) (00000001) TCP streaming post VM (cpas)
        13:  7f800000 (ffffffff8b671870) (ffffffff) IP Options Restore (out) (ipopt_res)
        14:  7f900000 (0000000000000000) (00000000) SecureXL outbound (sxl_out)
        15:  7fa00000 (0000000000000000) (00000000) SecureXL deliver (sxl_deliver)
[Expert@MyGW:0]#
```

Examples for the "-e" parameter

Example 1 - Capture everything

```
[Expert@HostName]# fw monitor -e "accept;" -o /var/log/fw_
mon.cap
```

Example 2 - Capture traffic to / from specific hosts

To specify a host, you can use one of these expressions:

- Use "`host(<IP_Address_in_Doted_Decimal_format>)`", which applies to both Source IP address and Destination IP address
- Use a specific Source IP address "`src=<IP_Address_in_Doted_Decimal_format>`" and a specific Destination IP address "`dst=<IP_Address_in_Doted_Decimal_format>`"

Example filters:

- Capture everything between host X and host Y:

```
[Expert@HostName]# fw monitor -e "host(x.x.x.x) and host
(y.y.y.y), accept;" -o /var/log/fw_mon.cap
```

```
[Expert@HostName]# fw monitor -e "((src=x.x.x.x ,
dst=y.y.y.y) or (src=y.y.y.y , dst=x.x.x.x)), accept;" -o
/var/log/fw_mon.cap
```

- Capture everything between hosts X,Z and hosts Y,Z in *all* Firewall kernel chains:

```
[Expert@HostName]# fw monitor -p all -e "((src=x.x.x.x or
dst=z.z.z.z) and (src=y.y.y.y or dst=z.z.z.z)), accept ;" -o
/var/log/fw_mon.cap
```

- Capture everything to/from host X or to/from host Y or to/from host Z:

```
[Expert@HostName]# fw monitor -e "host(x.x.x.x) or host
(y.y.y.y) or host(z.z.z.z), accept;" -o /var/log/fw_mon.cap
```

```
[Expert@HostName]# fw monitor -e "((src=x.x.x.x or
dst=x.x.x.x) or (src=y.y.y.y or dst=y.y.y.y) or (src=z.z.z.z
or dst=z.z.z.z)), accept;" -o /var/log/fw_mon.cap
```

Example 3 - Capture traffic to / from specific ports

Note - You must specify port numbers in Decimal format. Refer to the `/etc/services` file on the Security Gateway, or to *IANA Service Name and Port Number Registry*.

To specify a port, you can use one of these expressions:

- Use "`port(<IANA_Port_Number>)`", which applies to both Source Port and Destination Port
- Use a specific Source Port "`sport=<IANA_Port_Number>`" and a specific Destination Port "`dport=<IANA_Port_Number>`"
- In addition:
  - For specific TCP port, you can use "`tcpport(<IANA_Port_Number>)`", which applies to both Source TCP Port and Destination TCP Port
  - For specific UDP port, you can use "`udpport(<IANA_Port_Number>)`", which applies to both Source UDP Port and Destination UDP Port

Example filters:

- Capture everything to/from port X:

```
[Expert@HostName]# fw monitor -e "port(x), accept;" -o
/var/log/fw_mon.cap
```

```
[Expert@HostName]# fw monitor -e "(sport=x or dport=x),
accept;" -o /var/log/fw_mon.cap
```

- Capture everything except port X:

```
[Expert@HostName]# fw monitor -e "((sport=!x) or
(dport=!x)), accept;" -o /var/log/fw_mon.cap
```

```
[Expert@HostName]# fw monitor -e "not (sport=x or dport=x),
accept;" -o /var/log/fw_mon.cap
```

- Capture everything except SSH:

```
[Expert@HostName]# fw monitor -e "((sport!=22) or
(dport!=22)), accept;" -o /var/log/fw_mon.cap
```

```
[Expert@HostName]# fw monitor -e "not (sport=22 or
dport=22), accept;" -o /var/log/fw_mon.cap
```

```
[Expert@HostName]# fw monitor -e "not tcpport(22), accept;"
-o /var/log/fw_mon.cap
```

- Capture everything to/from host X except SSH:

```
[Expert@HostName]# fw monitor -e "(host(x.x.x.x) and
(sport!=22 or dport!=22)), accept;" -o /var/log/fw_mon.cap
```

```
[Expert@HostName]# fw monitor -e "((src=x.x.x.x or
dst=x.x.x.x) and (not (sport=22 or dport=22))), accept;" -o
/var/log/fw_mon.cap
```

```
[Expert@HostName]# fw monitor -e "(host(x.x.x.x) and not
tcpport(22)), accept;" -o /var/log/fw_mon.cap
```

- Capture everything except NTP:

```
[Expert@HostName]# fw monitor -e "not udpport(123), accept;"
-o /var/log/fw_mon.cap
```

Example 4 - Capture traffic over specific protocol

Note - You must specify protocol numbers in Decimal format. Refer to the `/etc/protocols` file on the Security Gateway, or to *IANA Protocol Numbers*.

To specify a protocol, you can use one of these expressions:

- Use "`ip_p=<IANA_Protocol_Number>`"

  Examples:

  - To specify TCP protocol with byte offset, use "`ip_p=6`"
  - To specify UDP protocol with byte offset, use "`ip_p=11`"
  - To specify ICMP protocol with byte offset, use "`ip_p=1`"

- Use "`accept [9:1]=<IANA_Protocol_Number>`"

  Examples:

  - To specify TCP protocol with byte offset, use "`accept [9:1]=6`"
  - To specify UDP protocol with byte offset, use "`accept [9:1]=11`"
  - To specify ICMP protocol with byte offset, use "`accept [9:1]=1`"

- In addition, you can explicitly use these expressions to specify protocols:

  Summary Table

Which protocol to specify	On which port(s) traffic is captured	Expression
TCP	N / A	`"tcp, accept;"`
UDP	N / A	`"udp, accept;"`
ICMPv4	N / A	`"icmp, accept;"` or `"icmp4, accept;"`
ICMPv6	N / A	`"icmp6, accept;"`
HTTP	TCP 80	`"http, accept;"`
HTTPS	TCP 443	`"https, accept;"`
PROXY	TCP 8080	`"proxy, accept;"`
DNS	UDP 53	`"dns, accept;"`
IKE	UDP 500	`"ike, accept;"`
NAT-T	UDP 4500	`"natt, accept;"`
ESP and IKE	IP proto 50 and UDP 500	`"vpn, accept;"`

Which protocol to specify	On which port(s) traffic is captured	Expression
All VPN-related data: a. ESP b. IPsec over UDP c. IKE d. NAT-T e. CRL f. RDP g. Tunnel Test h. Topology i. L2TP j. SCV k. Multi-Portal l. and so on	a. IP proto 50 b. UDP 2746 c. UDP 500 d. UDP 4500 e. TCP 18264 f. UDP 259 g. UDP 18234 h. TCP 264 i. TCP 1701 j. UDP 18233 k. TCP 443 + TCP 444 l. and so on	`"vpnall, accept;"`
Multi-Portal connections	TCP 443 and TCP 444	`"multi, accept;"`
SSH	TCP 22	`"ssh, accept;"`
FTP	TCP 20 and TCP 21	`"ftp, accept;"`
Telnet	TCP 23	`"telnet, accept;"`
SMTP	TCP 25	`"smtp, accept;"`
POP3	TCP 110	`"pop3, accept;"`

Example filters:

- Filter to capture everything on protocol X:

```
[Expert@HostName]# fw monitor -e "ip_p=X, accept;" -o
/var/log/fw_mon.cap
```

- Filter to capture rverything on protocol X and port Z on protocol Y:

```
[Expert@HostName]# fw monitor -e "(ip_p=X) or (ip_p=Y, port
(Z)), accept;" -o /var/log/fw_mon.cap
```

- Filter to capture capture everything TCP between host X and host Y:

```
[Expert@HostName]# fw monitor -e "ip_p=6, host(x.x.x.x) or
host(y.y.y.y), accept;" -o /var/log/fw_mon.cap
```

```
[Expert@HostName]# fw monitor -e "tcp, host(x.x.x.x) or host
(y.y.y.y), accept;" -o /var/log/fw_mon.cap
```

```
[Expert@HostName]# fw monitor -e "accept [9:1]=6 ,
((src=x.x.x.x , dst=y.y.y.y) or (src=y.y.y.y ,
dst=x.x.x.x));"
```

```
[Expert@HostName]# fw monitor -e "ip_p=6, ((src=x.x.x.x ,
dst=y.y.y.y) or (src=y.y.y.y , dst=x.x.x.x)), accept;" -o
/var/log/fw_mon.cap
```

Example 5 - Capture traffic with specific protocol options

Note - Refer to the `$FWDIR/lib/tcpip.def` file on Security Gateway.

Summary Table for IPv4

Option Description	Expression	Example
Source IPv4 address of the IPv4 packet	`ip_src = <IPv4_Address>`	`fw monitor -e "ip_src = 192.168.22.33, accept;"`
Destination IPv4 address of the IPv4 packet	`ip_dst = <IPv4_Address>`	`fw monitor -e "ip_dst = 192.168.22.33, accept;"`
Time To Live of the IPv4 packet	`ip_ttl = <Number>`	`fw monitor -e "ip_ttl = 255, accept;"`
Total Length of the IPv4 packet in bytes	`ip_len = <Length_in_Bytes>`	`fw monitor -e "ip_len = 64, accept;"`

Option Description	Expression	Example
TOS field of the IPv4 packet	`ip_tos = <Number>`	`fw monitor -e "ip_ tos = 0, accept;"`
IANA Protocol Number (either in Dec or in Hex) encapsulated in the IPv4 packet	`ip_p = <IANA_ Protocol_ Number>`	Example for TCP: `fw monitor -e "ip_p = 6, accept;"` Examples for UDP: `fw monitor -e "ip_p = 17, accept;"` `fw monitor -e "ip_p = 0x11, accept;"` Example for ICMPv4: `fw monitor -e "ip_p = 1, accept;"`

Summary Table for IPv6

Option Description	Expression	Example
Source IPv6 address of the IPv6 packet	ip_src6p = <IPv6_ Address>	`fw monitor -e "ip_src6p = 0:0:0:0:0:ffff:c0a8:1621, accept;"`
Destination IPv6 address of the IPv6 packet	ip_dst6p = <IPv6_ Address>	`fw monitor -e "ip_dst6p = 0:0:0:0:0:ffff:c0a8:1621, accept;"`
Payload Length of the IPv6 packet in bytes	ip_len6 = <Length_in_ Bytes>	`fw monitor -e "ip_len6 = 1000, accept;"`
Hop Limit ("Time To Live") of the IPv6 packet	ip_ttl6 = <Number>	`fw monitor -e "ip_ttl6 = 255, accept;"`
Next Header of the IPv6 packet - encapsulated IANA Protocol Number	ip_p6 = <IANA_ Protocol_ Number>	`fw monitor -e "ip_p6 = 6, accept;"`

Summary Table for TCP

Option Description	Expression	Example
SYN flag is set in TCP packet	`syn`	`fw monitor -e "ip_p = 6, syn, accept;"`
ACK flag is set in TCP packet	`ack`	`fw monitor -e "ip_p = 6, ack, accept;"`
RST flag is set in TCP packet	`rst`	`fw monitor -e "ip_p = 6, rst, accept;"`
FIN flag is set in TCP packet	`fin`	`fw monitor -e "ip_p = 6, fin, accept;"`
First packet of TCP connection (SYN flag is set, but ACK flag is not set in TCP packet)	`first`	`fw monitor -e "ip_p = 6, first, accept;"`
Not the first packet of TCP connection (SYN flag is not set in TCP packet)	`not_first`	`fw monitor -e "ip_p = 6, not_first, accept;"`
Established TCP connection (either ACK flag is set, or SYN flag is not set in TCP packet)	`established`	`fw monitor -e "ip_p = 6, established, accept;"`
Last packet of TCP connection (both ACK flag and FIN flag are set in TCP packet)	`last`	`fw monitor -e "ip_p = 6, last, accept;"`
End of TCP connection (either RST flag is set, or FIN flag is set in TCP packet)	`tcpdone`	`fw monitor -e "ip_p = 6, tcpdone, accept;"`

<table>
<tr><th>Option Description</th><th>Expression</th><th>Example</th></tr>
<tr><td>General way to match the flags inside in TCP packets</td><td><code>th_flags = <<i>Sum_of_Flags_Hex_Values</i>></code></td><td>
<table>
<tr><th>TCP Flag</th><th>Example</th></tr>
<tr><td>SYN (0x2)</td><td><code>fw monitor -e "th_ flags = 0x2, accept;"</code></td></tr>
<tr><td>ACK (0x10)</td><td><code>fw monitor -e "th_ flags = 0x10, accept;"</code></td></tr>
<tr><td>PSH (0x8)</td><td><code>fw monitor -e "th_ flags = 0x8, accept;"</code></td></tr>
<tr><td>FIN (0x1)</td><td><code>fw monitor -e "th_ flags = 0x1, accept;"</code></td></tr>
<tr><td>RST (0x4)</td><td><code>fw monitor -e "th_ flags = 0x4, accept;"</code></td></tr>
<tr><td>URG (0x20)</td><td><code>fw monitor -e "th_ flags = 0x20, accept;"</code></td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Option Description</th><th>Expression</th><th>Example</th></tr>
<tr><td></td><td></td><td>
<table>
<tr><th>TCP Flag</th><th>Example</th></tr>
<tr><td>SYN + ACK</td><td><code>fw monitor -e "th_ flags = 0x12, accept;"</code></td></tr>
<tr><td>PSH + ACK</td><td><code>fw monitor -e "th_ flags = 0x18, accept;"</code></td></tr>
<tr><td>FIN + ACK</td><td><code>fw monitor -e "th_ flags = 0x11, accept;"</code></td></tr>
<tr><td>RST + ACK</td><td><code>fw monitor -e "th_ flags = 0x14, accept;"</code></td></tr>
</table>
</td></tr>
<tr><td>TCP source port</td><td><code>th_sport = <<i>Port_Number</i>></code></td><td><code>fw monitor -e "th_ sport = 59259, accept;"</code></td></tr>
<tr><td>TCP destination port</td><td><code>th_dport = <<i>Port_Number</i>></code></td><td><code>fw monitor -e "th_ dport = 22, accept;"</code></td></tr>
<tr><td>TCP sequence number (either in Dec or in Hex)</td><td><code>th_seq = <<i>Number</i>></code></td><td>Example for Dec format:
<code>fw monitor -e "th_seq = 3937833514, accept;"</code>
Example for Hex format:
<code>fw monitor -e "th_seq = 0xeab6922a, accept;"</code></td></tr>
</table>

Option Description	Expression	Example
TCP acknowledged number (either in Dec or in Hex)	`th_ack = <Number>`	Example for Dec format: `fw monitor -e "th_ack = 509054325, accept;"` Example for Hex format: `fw monitor -e "th_ack = 0x1e578d75, accept;"`

Summary Table for UDP

Option Description	Expression	Example
UDP source port	`uh_sport = <Port_Number>`	`fw monitor -e "uh_sport = 53, accept;"`
UDP destination port	`uh_dport = <Port_Number>`	`fw monitor -e "uh_dport = 53, accept;"`

Summary Table for ICMPv4

Option Description	Expression	Example
ICMPv4 packets with specified Type	`icmp_type = <Number>`	`fw monitor -e "icmp_type = 0, accept;"`
ICMPv4 packets with specified Code	`icmp_code = <Number>`	`fw monitor -e "icmp_code = 0, accept;"`
ICMPv4 packets with specified Identifier	`icmp_id = <Number>`	`fw monitor -e "icmp_id = 20583, accept;"`
ICMPv4 packets with specified Sequence number	`icmp_seq = <Number>`	`fw monitor -e "icmp_seq = 1, accept;"`
ICMPv4 Echo Request packets (Type 8, Code 0)	`echo_req`	`fw monitor -e "echo_req, accept;"`
ICMPv4 Echo Reply packets (Type 0, Code 0)	`echo_reply`	`fw monitor -e "echo_reply, accept;"`
ICMPv4 Echo Request and ICMPv4 Echo Reply packets	`ping`	`fw monitor -e "ping, accept;"`

Option Description	Expression	Example
Traceroute packets as implemented in Unix OS (UDP packets on ports above 30000 and with TTL<30; or ICMP Time exceeded packets)	`traceroute`	`fw monitor -e "traceroute, accept;"`
Traceroute packets as implemented in Windows OS (ICMP Request packets with TTL<30; or ICMP Time exceeded packets)	`tracert`	`fw monitor -e "tracert, accept;"`
Length of ICMPv4 packets	`icmp_ip_len = <length>`	`fw monitor -e "icmp_ip_len = 84, accept;"`

Summary Table for ICMPv6

Option Description	Expression	Example
ICMPv6 packets with specified Type	`icmp6_type = <Number>`	`fw monitor -e "icmp6_type = 1, accept;"`
ICMPv6 packets with specified Code	`icmp6_code = <Number>`	`fw monitor -e "icmp6_code = 3, accept;"`

Example 6 - Capture specific bytes in packets

Syntax:

```
fw monitor -e "accept [ <Offset> : <Length> , <Byte Order> ]
<Relational-Operator> <Value>;"
```

Parameters:

Parameter	Explanation
<Offset>	Specifies the offset relative to the beginning of the IP packet from where the value should be read.

Parameter	Explanation
<Length>	Specifies the number of bytes: ■ `1` = byte ■ `2` = word ■ `4` = dword If length is not specified, FW Monitor assumes 4 (dword).
<Byte Order>	Specifies the byte order: ■ `b` = big endian, or network order ■ `l` = little endian, or host order If order is not specified, FW Monitor assumes little endian byte order.
<Relational-Operator	Relational operator to express the relation between the packet data and the value: ■ `<` - less than ■ `>` - greater than ■ `<=` - less than or equal to ■ `>=` - greater than ■ `=` or `is` - equal to ■ `!=` or `is not` - not equal to
<Value>	One of the data types known to INSPECT (for example, an IP address, or an integer).

Explanations:

- The IP-based protocols are stored in the IP packet as a byte at offset 9.
  - To filter based on a Protocol encapsulated into IP, use this syntax:

    ```
    [Expert@HostName]# fw monitor -e "accept [9:1]=<IANA_
    Protocol_Number>;"
    ```

- The Layer 3 IP Addresses are stored in the IP packet as double words at offset 12 (Source address) and at offset 16 (Destination address).

- To filter based on a Source IP address, use this syntax:

```
[Expert@HostName]# fw monitor -e "accept [12:4,b]=<IP_
Address_in_Doted_Decimal_format>;"
```

- To filter based on a Destination IP address, use this syntax:

```
[Expert@HostName]# fw monitor -e "accept [16:4,b]=<IP_
Address_in_Doted_Decimal_format>;"
```

- The Layer 4 Ports are stored in the IP packet as a word at offset 20 (Source port) and at offset 22 (Destination port).
  - To filter based on a Source port, use this syntax:

```
[Expert@HostName]# fw monitor -e "accept
[20:2,b]=<Port_Number_in_Decimal_format>;"
```

  - To filter based on a Destination port, use this syntax:

```
[Expert@HostName]# fw monitor -e "accept
[22:2,b]=<Port_Number_in_Decimal_format>;"
```

Example filters:

- Capture everything between host X and host Y:

```
[Expert@HostName]# fw monitor -e "accept (([12:4,b]=x.x.x.x
, [16:4,b]=y.y.y.y) or ([12:4,b]=y.y.y.y ,
[16:4,b]=x.x.x.x));"
```

- Capture everything on port X:

```
[Expert@HostName]# fw monitor -e "accept [20:2,b]=x or
[22:2,b]=x;" -o /var/log/fw_mon.cap
```

Example 7 - Capture traffic to/from specific network

You must specify the *network address* and *length of network mask* (number of bits).

There are 3 options:

Traffic direction	Expression
To or From a network	"net(<*Network_IP_Address*>, <*Mask_Length*>), accept;"

Traffic direction	Expression
To a network	"to_net(<*Network_IP_Address*>, <*Mask_Length*>), accept;"
From a network	"from_net(<*Network_IP_Address*>, <*Mask_Length*>), accept;"

Example filters:

- Capture everything to/from network 192.168.33.0 / 24:

```
[Expert@HostName]# fw monitor -e "net(192.168.33.0, 24), accept;"
```

- Capture everything sent to network 192.168.33.0 / 24:

```
[Expert@HostName]# fw monitor -e "to_net(192.168.33.0, 24), accept;"
```

- Capture everything sent from network 192.168.33.0 / 24:

```
[Expert@HostName]# fw monitor -e "from_net(192.168.33.0, 24), accept;"
```

Example 8 - Filter out irrelevant "noise"

Filter in only TCP protocol, and HTTP and HTTPS ports

Filter out the SSH and FW Logs

```
[Expert@HostName]# fw monitor -e "accept (ip_p=6) and (not (sport=22 or dport=22)) and (not (sport=257 or dport=257)) and ((dport=80 or dport=443) or (sport=80 or sport=443);" -o /var/log/fw_mon.cap
```

Examples for the "-F" parameter

You can specify up to 5 capture filters with this parameter (up to 5 instances of the "`-F`" parameter in the syntax).

The FW Monitor performs the logical "OR" between all specified simple capture filters.

Value 0 is used as "`any`".

Example 1 - Capture everything

```
[Expert@HostName]# fw monitor -F "0,0,0,0,0" -o /var/log/fw_
mon.cap
```

Example 2 - Capture traffic to / from specific hosts

- Capture all traffic from Source IP x.x.x.x (any port) to Destination IP y.y.y.y (any port), over all protocols:

  ```
  [Expert@HostName]# fw monitor -F "x.x.x.x,0,y.y.y.y,0,0" -o
  /var/log/fw_mon.cap
  ```

- Capture all traffic between Host x.x.x.x (any port) and Host y.y.y.y (any port), over all protocols:

  ```
  [Expert@HostName]# fw monitor -F "x.x.x.x,0,y.y.y.y,0,0" -F
  "y.y.y.y,0, x.x.x.x ,0,0" -o /var/log/fw_mon.cap
  ```

Example 3 - Capture traffic to / from specific ports

- Capture traffic from any Source IP from Source Port X to any Destination IP to Destination Port Y, over all protocols:

  ```
  [Expert@HostName]# fw monitor -F "0,x,0,y,0" -o /var/log/fw_
  mon.cap
  ```

- Capture traffic between all hosts, between Port X and Port Y, over all protocols:

  ```
  [Expert@HostName]# fw monitor -F "0,x,0,y,0" -F "0,y,0,x,0"
  -o /var/log/fw_mon.cap
  ```

Example 4 - Capture traffic over specific protocol

- Capture traffic between all hosts, between all ports, over a Protocol with assigned number X:

  ```
  [Expert@HostName]# fw monitor -F "0,0,0,0,x" -o /var/log/fw_
  mon.cap
  ```

Example 5 - Capture traffic between specific hosts between specific ports over specific protocol

```
[Expert@HostName]# fw monitor -F "a.a.a.a,b,c.c.c.c,d,e" -F
"c.c.c.c,d,a.a.a.a,b,e" -o /var/log/fw_mon.cap
```

To capture only HTTP traffic between the Client 1.1.1.1 and the Server 2.2.2.2:

```
fw montior -F "1.1.1.1,0,2.2.2.2,80,6" -F
"2.2.2.2,80,1.1.1.1,0,6" -o /var/log/fw_mon.cap
```

fw sam_policy

Description

Manages the Suspicious Activity Policy editor that works with these types of rules:

- Suspicious Activity Monitoring (SAM) rules.

  See sk112061: How to create and view Suspicious Activity Monitoring (SAM) Rules.

- Rate Limiting rules.

  See sk182350 - How to configure Rate Limiting rules for DoS Mitigation.

Also, see these commands:

- `sam_alert` (see the *R82 CLI Reference Guide*)

Notes:

- These commands are interchangeable:
  - For IPv4: "`fw sam_policy`" and "`fw samp`".
  - For IPv6: "`fw6 sam_policy`" and "`fw6 samp`".
- You can run these commands in Gaia Clish, or Expert mode.
- Security Gateway stores the SAM Policy rules in the `$FWDIR/database/sam_policy.db` file.
- Security Gateway stores the SAM Policy management settings in the `$FWDIR/database/sam_policy.mng` file.

Important:

- Configuration you make with these commands, survives reboot.
- You can run this command in the Expert mode or in Gaia Clish (Gaia gClish on Scalable Platforms).
- In a Cluster, you must configure all the Cluster Members in the same way.
- On Scalable Platforms, you must connect to the applicable Security Group.
- VSNext mode and Traditional VSX mode do **not** support Suspicious Activity Monitoring (SAM) Rules. See sk79700.

Best Practice - The SAM Policy rules consume some CPU resources on Security Gateway. Set an expiration for rules that gives you time to investigate, but does not affect performance. Keep only the required SAM Policy rules. If you confirm that an activity is risky, edit the Security Policy, educate users, or otherwise handle the risk.

Syntax for IPv4

```
fw [-d] sam_policy
      add <options>
      batch
      del <options>
      get <options>
```

```
fw [-d] samp
      add <options>
      batch
      del <options>
      get <options>
```

Syntax for IPv6

```
fw6 [-d] sam_policy
      add <options>
      batch
      del <options>
      get <options>
```

```
fw6 [-d] samp
      add <options>
      batch
      del <options>
      get <options>
```

Parameters

Parameter	Description
`-d`	Runs the command in debug mode. Use only if you troubleshoot the command itself. **Best Practice** - If you use this parameter, then redirect the output to a file, or use the `script` command to save the entire CLI session.
`add <options>`	Adds one Rate Limiting rule one at a time. See *"fw sam_policy add" on page 234*.
`batch`	Adds or deletes many Rate Limiting rules at a time. See *"fw sam_policy batch" on page 247*.
`del <options>`	Deletes one configured Rate Limiting rule one at a time. See *"fw sam_policy del" on page 249*.
`get <options>`	Shows all the configured Rate Limiting rules. See *"fw sam_policy get" on page 253*.

fw sam_policy add

Description

The "`fw sam_policy add`" and "`fw6 sam_policy add`" commands:

- Add one Suspicious Activity Monitoring (SAM) rule at a time.
- Add one Rate Limiting rule at a time.

Notes:

- These commands are interchangeable:
  - For IPv4: "`fw sam_policy`" and "`fw samp`".
  - For IPv6: "`fw6 sam_policy`" and "`fw6 samp`".
- You can run these commands in Gaia Clish, or Expert mode.
- Security Gateway stores the SAM Policy rules in the `$FWDIR/database/sam_policy.db` file.
- Security Gateway stores the SAM Policy management settings in the `$FWDIR/database/sam_policy.mng` file.

Important:

- Configuration you make with these commands, survives reboot.
- You can run this command in the Expert mode or in Gaia Clish (Gaia gClish on Scalable Platforms).
- In a Cluster, you must configure all the Cluster Members in the same way.
- On Scalable Platforms, you must connect to the applicable Security Group.
- VSNext mode and Traditional VSX mode do **not** support Suspicious Activity Monitoring (SAM) Rules. See sk79700.

Best Practice - The SAM Policy rules consume some CPU resources on Security Gateway. Set an expiration for rules that gives you time to investigate, but does not affect performance. Keep only the required SAM Policy rules. If you confirm that an activity is risky, edit the Security Policy, educate users, or otherwise handle the risk.

Syntax to configure a Suspicious Activity Monitoring (SAM) rule for IPv4

```
fw [-d] sam_policy add [-u] -a {d|n|b} [-l {r|a}] [-t <Timeout>]
[-f <Target>] [-n <"Rule Name">] [-c <"Rule Comment">] [-o <"Rule
Originator">] [-z "<Zone>"] ip <IP Filter Arguments>
```

Syntax to configure a Suspicious Activity Monitoring (SAM) rule for IPv6

```
fw6 [-d] sam_policy add [-u] -a {d|n|b} [-l {r|a}] [-t <Timeout>]
[-f <Target>] [-n <"Rule Name">] [-c <"Rule Comment">] [-o <"Rule
Originator">] [-z "<Zone>"] ip <IP Filter Arguments>
```

Syntax to configure a Rate Limiting rule for IPv4

```
fw [-d] sam_policy add [-u] -a {d|n|b} [-l {r|a}] [-t <Timeout>]
[-f <Target>] [-n <"Rule Name">] [-c <"Rule Comment">] [-o <"Rule
Originator">] [-z "<Zone>"] quota <Quota Filter Arguments>
```

Syntax to configure a Rate Limiting rule for IPv6

```
fw6 [-d] sam_policy add [-u] -a {d|n|b} [-l {r|a}] [-t <Timeout>]
[-f <Target>] [-n <"Rule Name">] [-c <"Rule Comment">] [-o <"Rule
Originator">] [-z "<Zone>"] quota <Quota Filter Arguments
```

Parameters

Parameter	Description
`-d`	Runs the command in debug mode. Use only if you troubleshoot the command itself. **Best Practice** - If you use this parameter, then redirect the output to a file, or use the `script` command to save the entire CLI session.
`-u`	Optional. Specifies that the rule category is `User-defined`. Default rule category is `Auto`.
`-a {d | n | b}`	Mandatory. Specifies the rule action if the traffic matches the rule conditions: ■ `d` - Drop the connection. ■ `n` - Notify (generate a log) about the connection and let it through. ■ `b` - Bypass the connection - let it through without checking it against the policy rules. **Note** - Rules with action set to *Bypass* cannot have a log or limit specification. Bypassed packets and connections do not count towards overall number of packets and connection for limit enforcement of type ratio.
`-l {r | a}`	Optional. Specifies which type of log to generate for this rule for all traffic that matches: ■ `-r` - Generate a regular log ■ `-a` - Generate an alert log

Parameter	Description
`-t <Timeout>`	Optional. Specifies the time period (in seconds), during which the rule will be enforced. Default timeout is indefinite.
`-f <Target>`	Optional. Specifies the target Security Gateways, on which to enforce the Rate Limiting rule. `<Target>` can be one of these: ■ `all` - This is the default option. Specifies that the rule should be enforced on all managed Security Gateways. ■ Name of the Security Gateway or Cluster object - Specifies that the rule should be enforced only on this Security Gateway or Cluster object (the object name must be as defined in the SmartConsole). ■ Name of the Group object - Specifies that the rule should be enforced on all Security Gateways that are members of this Group object (the object name must be as defined in the SmartConsole).
`-n "<Rule Name>"`	Optional. Specifies the name (label) for this rule. **Notes:** ■ You must enclose this string in double quotes. ■ The length of this string is limited to 128 characters. ■ Before each space or a backslash character in this string, you must write a backslash (\) character. Example: `"This\ is\ a\ rule\ name\ with\ a\ backslash\ \\"`
`-c "<Rule Comment>"`	Optional. Specifies the comment for this rule. **Notes:** ■ You must enclose this string in double quotes. ■ The length of this string is limited to 128 characters. ■ Before each space or a backslash character in this string, you must write a backslash (\) character. Example: `"This\ is\ a\ comment\ with\ a\ backslash\ \\"`

Parameter	Description
`-o "<Rule Originator>"`	Optional. Specifies the name of the originator for this rule. **Notes:** ■ You must enclose this string in double quotes. ■ The length of this string is limited to 128 characters. ■ Before each space or a backslash character in this string, you must write a backslash (\) character. Example: `"Created\ by\ John\ Doe"`
`-z "<Zone>"`	Optional. Specifies the name of the Security Zone for this rule. **Notes:** ■ You must enclose this string in double quotes. ■ The length of this string is limited to 128 characters.
`ip <IP Filter Arguments>`	Mandatory (use this `ip` parameter, or the `quota` parameter). Configures the *Suspicious Activity Monitoring (SAM)* rule. Specifies the IP Filter Arguments for the SAM rule (you must use at least one of these options): `[-C] [-s <Source IP>] [-m <Source Mask>] [-d <Destination IP>] [-M <Destination Mask>] [-p <Port>] [-r <Protocol>]` See the explanations below.

Parameter	Description
`quota <Quota Filter Arguments>`	Mandatory (use this `quota` parameter, or the `ip` parameter). Configures the *Rate Limiting* rule. Specifies the Quota Filter Arguments for the Rate Limiting rule (see the explanations below): ■ `[flush true]` ■ `[source-negated {true | false}] source <Source>` ■ `[destination-negated {true | false}] destination <Destination>` ■ `[service-negated {true | false}] service <Protocol and Port numbers>` ■ `[<Limit1 Name> <Limit1 Value>] [<Limit2 Name> <Limit2 Value>] ...[<LimitN Name> <LimitN Value>]` ■ `[track <Track>]` **Important:** ■ The Quota rules are not applied immediately to the Security Gateway. They are only registered in the Suspicious Activity Monitoring (SAM) policy database. To apply all the rules from the SAM policy database immediately, add "`flush true`" in the `fw samp add` command syntax. ■ Explanation: For new connections rate (and for any rate limiting in general), when a rule's limit is violated, the Security Gateway also drops all packets that match the rule. The Security Gateway computes new connection rates on a per-second basis. At the start of the 1-second timer, the Security Gateway allows all packets, including packets for existing connections. If, at some point, during that 1 second period, there are too many new connections, then the Security Gateway blocks all remaining packets for the remainder of that 1-second interval. At the start of the next 1-second interval, the counters are reset, and the process starts over - the Security Gateway allows packets to pass again up to the point, where the rule's limit is violated.

Explanation for the *IP Filter Arguments* syntax for Suspicious Activity Monitoring (SAM) rules

Argument	Description
`-C`	Specifies that open connections should be closed.
`-s <Source IP>`	Specifies the Source IP address.
`-m <Source Mask>`	Specifies the Source subnet mask (in dotted decimal format - x.y.z.w).
`-d <Destination IP>`	Specifies the Destination IP address.
`-M <Destination Mask>`	Specifies the Destination subnet mask (in dotted decimal format - x.y.z.w).
`-p <Port>`	Specifies the port number (see *IANA Service Name and Port Number Registry*).
`-r <Protocol>`	Specifies the protocol number (see *IANA Protocol Numbers*).

Explanation for the *Quota Filter Arguments* syntax for Rate Limiting rules

Argument	Description
`flush true`	Specifies to compile and load the quota rule to the SecureXL immediately.
`[source-negated {true | false}] source <Source>`	Specifies the source type and its value: ■ `any` The rule is applied to packets sent from all sources. ■ `range:<IP Address>` or `range:<IP Address Start>-<IP Address End>` The rule is applied to packets sent from: • Specified IPv4 addresses (x.y.z.w) • Specified IPv6 addresses (xxxx:yyyy:....:zzzz) ■ `cidr:<IP Address>/<Prefix>` The rule is applied to packets sent from: • IPv4 address with Prefix from 0 to 32 • IPv6 address with Prefix from 0 to 128 ■ `cc:<Country Code>` The rule matches the country code to the source IP addresses assigned to this country, based on the Geo IP database. The two-letter codes are defined in ISO 3166-1 alpha-2. ■ `asn:<Autonomous System Number>` The rule matches the AS number of the organization to the source IP addresses that are assigned to this organization, based on the Geo IP database. The valid syntax is *ASnnnn*, where *nnnn* is a number unique to the specific organization. **Notes:** ■ Default is: `source-negated false` ■ The `source-negated true` processes all source types, *except* the specified type.

Argument	Description
`[destination-negated {true | false}] destination <Destination>`	Specifies the destination type and its value: ■ `any` The rule is applied to packets sent to all destinations. ■ `range:<IP Address>` or `range:<IP Address Start>-<IP Address End>` The rule is applied to packets sent to: • Specified IPv4 addresses (x.y.z.w) • Specified IPv6 addresses (xxxx:yyyy:...:zzzz) ■ `cidr:<IP Address>/<Prefix>` The rule is applied to packets sent to: • IPv4 address with Prefix from 0 to 32 • IPv6 address with Prefix from 0 to 128 ■ `cc:<Country Code>` The rule matches the country code to the destination IP addresses assigned to this country, based on the Geo IP database. The two-letter codes are defined in ISO 3166-1 alpha-2. ■ `asn:<Autonomous System Number>` The rule matches the AS number of the organization to the destination IP addresses that are assigned to this organization, based on the Geo IP database. The valid syntax is *ASnnnn*, where *nnnn* is a number unique to the specific organization. **Notes:** ■ Default is: `destination-negated false` ■ The `destination-negated true` will process all destination types except the specified type

Argument	Description
`[service-negated {true | false}] service <Protocol and Port numbers>`	Specifies the Protocol number (see *IANA Protocol Numbers*) and Port number (see *IANA Service Name and Port Number Registry*): ■ `<Protocol>` IP protocol number in the range 1-255 ■ `<Protocol Start>-<Protocol End>` Range of IP protocol numbers ■ `<Protocol>/<Port>` IP protocol number in the range 1-255 and TCP/UDP port number in the range 1-65535 ■ `<Protocol>/<Port Start>-<Port End>` IP protocol number and range of TCP/UDP port numbers from 1 to 65535 **Notes:** ■ Default is: `service-negated false` ■ The `service-negated true` will process all traffic except the traffic with the specified protocols and ports

Argument	Description
`[<Limit 1 Name> <Limit 1 Value>] [<Limit 2 Name> <Limit 2 Value>] ... [<Limit N Name> <Limit N Value>]`	Specifies quota limits and their values. **Note** - Separate multiple quota limits with spaces. ■ `concurrent-conns <Value>` Specifies the maximum number of concurrent active connections that match this rule. ■ `concurrent-conns-ratio <Value>` Specifies the maximum ratio of the *concurrent-conns* value to the total number of active connections through the Security Gateway, expressed in parts per 65536 (formula: `N / 65536`). ■ `pkt-rate <Value>` Specifies the maximum number of packets per second that match this rule. ■ `pkt-rate-ratio <Value>` Specifies the maximum ratio of the *pkt-rate* value to the rate of all connections through the Security Gateway, expressed in parts per 65536 (formula: `N / 65536`). ■ `byte-rate <Value>` Specifies the maximum total number of bytes per second in packets that match this rule. ■ `byte-rate-ratio <Value>` Specifies the maximum ratio of the *byte-rate* value to the bytes per second rate of all connections through the Security Gateway, expressed in parts per 65536 (formula: `N / 65536`). ■ `new-conn-rate <Value>` Specifies the maximum number of connections per second that match the rule. ■ `new-conn-rate-ratio <Value>` Specifies the maximum ratio of the *new-conn-rate* value to the rate of all connections per second through the Security Gateway, expressed in parts per 65536 (formula: `N / 65536`).

Argument	Description
`[track <Track>]`	Specifies the tracking option: ■ `source` Counts connections, packets, and bytes for specific source IP address, and not cumulatively for this rule. ■ `source-service` Counts connections, packets, and bytes for specific source IP address, and for specific IP protocol and destination port, and not cumulatively for this rule.

Examples

Example 1 - Rate Limiting rule with a range

```
fw sam_policy add -a d -l r -t 3600 quota service any source range:172.16.7.11-172.16.7.13 new-conn-rate 5 flush true
```

Explanations:

- This rule drops packets for all connections (`-a d`) that exceed the quota set by this rule, including packets for existing connections.
- This rule logs packets (`-l r`) that exceed the quota set by this rule.
- This rule will expire in 3600 seconds (`-t 3600`).
- This rule limits the rate of creation of new connections to 5 connections per second (`new-conn-rate 5`) for any traffic (`service any`) from the source IP addresses in the range 172.16.7.11 - 172.16.7.13 (`source range:172.16.7.11-172.16.7.13`).

  Note - The limit of the total number of log entries per second is configured with the *fwaccel dos config set -n <rate>* command.
- This rule will be compiled and loaded on the SecureXL, together with other rules in the Suspicious Activity Monitoring (SAM) policy database immediately, because this rule includes the "`flush true`" parameter.

Example 2 - Rate Limiting rule with a service specification

```
fw sam_policy add -a n -l r quota service 1,50-51,6/443,17/53 service-negated true source cc:QQ byte-rate 0
```

Explanations:

- This rule logs and lets through all packets (`-a n`) that exceed the quota set by this rule.
- This rule does not expire (the `timeout` parameter is not specified). To cancel it, you must delete it explicitly.
- This rule applies to all packets except (`service-negated true`) the packets with IP protocol number 1, 50-51, 6 port 443 and 17 port 53 (`service 1,50-51,6/443,17/53`).
- This rule applies to all packets from source IP addresses that are assigned to the country with specified country code (`cc:QQ`).
- This rule does not let any traffic through (`byte-rate 0`) except the packets with IP protocol number 1, 50-51, 6 port 443 and 17 port 53.
- This rule will not be compiled and installed on the SecureXL immediately, because it does not include the "`flush true`" parameter.

Example 3 - Rate Limiting rule with ASN

```
fw sam_policy -a d quota source asn:AS64500,cidr:[::FFFF:C0A8:1100]/120 service any pkt-rate 0
```

Explanations:

- This rule drops (`-a d`) all packets that match this rule.
- This rule does not expire (the `timeout` parameter is not specified). To cancel it, you must delete it explicitly.
- This rule applies to packets from the Autonomous System number 64500 (`asn:AS64500`).
- This rule applies to packets from source IPv6 addresses FFFF:C0A8:1100/120 (`cidr:[::FFFF:C0A8:1100]/120`).
- This rule applies to all traffic (`service any`).
- This rule does not let any traffic through (`pkt-rate 0`).
- This rule will not be compiled and installed on the SecureXL immediately, because it does not include the "`flush true`" parameter.

Example 4 - Rate Limiting rule with an Allow List

```
fw sam_policy add -a b quota source range:172.16.8.17-172.16.9.121 service 6/80
```

Explanations:

- This rule bypasses (`-a b`) all packets that match this rule.

  Note - The Access Control Policy and other types of security policy rules still apply.
- This rule does not expire (the `timeout` parameter is not specified). To cancel it, you must delete it explicitly.
- This rule applies to packets from the source IP addresses in the range 172.16.8.17 - 172.16.9.121 (`range:172.16.8.17-172.16.9.121`).
- This rule applies to packets sent to TCP port 80 (`service 6/80`).
- This rule will not be compiled and installed on the SecureXL immediately, because it does not include the "`flush true`" parameter.

Example 5 - Rate Limiting rule with tracking

```
fw sam_policy add -a d quota service any source-negated true source cc:QQ concurrent-conns-ratio 655 track source
```

Explanations:

- This rule drops (`-a d`) all packets that match this rule.
- This rule does not log any packets (the `-l r` parameter is not specified).
- This rule does not expire (the `timeout` parameter is not specified). To cancel it, you must delete it explicitly.
- This rule applies to all traffic (`service any`).
- This rule applies to all sources except (`source-negated true`) the source IP addresses that are assigned to the country with specified country code (`cc:QQ`).
- This rule limits the maximum number of concurrent active connections to 655/65536=~1% (`concurrent-conns-ratio 655`) for any traffic (`service any`) except (`service-negated true`) the connections from the source IP addresses that are assigned to the country with specified country code (`cc:QQ`).
- This rule counts connections, packets, and bytes for traffic only from sources that match this rule, and not cumulatively for this rule.
- This rule will not be compiled and installed on the SecureXL immediately, because it does not include the "`flush true`" parameter.

fw sam_policy batch

Description

The "`fw sam_policy batch`" and "`fw6 sam_policy batch`" commands:

- Add and delete many Suspicious Activity Monitoring (SAM) rules at a time.
- Add and delete many Rate Limiting rules at a time.

Notes:

- These commands are interchangeable:
  - For IPv4: "`fw sam_policy`" and "`fw samp`".
  - For IPv6: "`fw6 sam_policy`" and "`fw6 samp`".
- You can run these commands in Gaia Clish, or Expert mode.
- Security Gateway stores the SAM Policy rules in the `$FWDIR/database/sam_policy.db` file.
- Security Gateway stores the SAM Policy management settings in the `$FWDIR/database/sam_policy.mng` file.

Important:

- Configuration you make with these commands, survives reboot.
- You can run this command in the Expert mode or in Gaia Clish (Gaia gClish on Scalable Platforms).
- In a Cluster, you must configure all the Cluster Members in the same way.
- On Scalable Platforms, you must connect to the applicable Security Group.
- VSNext mode and Traditional VSX mode do **not** support Suspicious Activity Monitoring (SAM) Rules. See sk79700.

Best Practice - The SAM Policy rules consume some CPU resources on Security Gateway. Set an expiration for rules that gives you time to investigate, but does not affect performance. Keep only the required SAM Policy rules. If you confirm that an activity is risky, edit the Security Policy, educate users, or otherwise handle the risk.

Procedure

1. **Start the batch mode**

   - For IPv4, run:

     ```
     fw sam_policy batch << EOF
     ```

   - For IPv6, run:

     ```
     fw6 sam_policy batch << EOF
     ```

2. **Enter the applicable commands**

   - Enter one "`add`" or "`del`" command on each line, on as many lines as necessary.

     Start each line with only "`add`" or "`del`" parameter (not with "`fw samp`").
   - Use the same set of parameters and values as described in these commands:
     - *"fw sam_policy add" on page 234*
     - *"fw sam_policy del" on page 249*
   - Terminate each line with a Return (ASCII 10 - Line Feed) character (press Enter).

3. **End the batch mode**

   Type `EOF` and press Enter.

Example of a Rate Limiting rule for IPv4

```
[Expert@HostName]# fw samp batch <<EOF

add -a d -l r -t 3600 -c "Limit\ conn\ rate\ to\ 5\ conn/sec from\ these\ sources" quota service
any source range:172.16.7.13-172.16.7.13 new-conn-rate 5

del <501f6ef0,00000000,cb38a8c0,0a0afffe>

add -a b quota source range:172.16.8.17-172.16.9.121 service 6/80

EOF
[Expert@HostName]#
```

fw sam_policy del

Description

The "`fw sam_policy del`" and "`fw6 sam_policy del`" commands:

- Delete one configured Suspicious Activity Monitoring (SAM) rule at a time.
- Delete one configured Rate Limiting rule at a time.

Notes:

- These commands are interchangeable:
  - For IPv4: "`fw sam_policy`" and "`fw samp`".
  - For IPv6: "`fw6 sam_policy`" and "`fw6 samp`".
- You can run these commands in Gaia Clish, or Expert mode.
- Security Gateway stores the SAM Policy rules in the `$FWDIR/database/sam_policy.db` file.
- Security Gateway stores the SAM Policy management settings in the `$FWDIR/database/sam_policy.mng` file.

Important:

- Configuration you make with these commands, survives reboot.
- You can run this command in the Expert mode or in Gaia Clish (Gaia gClish on Scalable Platforms).
- In a Cluster, you must configure all the Cluster Members in the same way.
- On Scalable Platforms, you must connect to the applicable Security Group.
- VSNext mode and Traditional VSX mode do **not** support Suspicious Activity Monitoring (SAM) Rules. See sk79700.

Best Practice - The SAM Policy rules consume some CPU resources on Security Gateway. Set an expiration for rules that gives you time to investigate, but does not affect performance. Keep only the required SAM Policy rules. If you confirm that an activity is risky, edit the Security Policy, educate users, or otherwise handle the risk.

Syntax for IPv4

```
fw [-d] sam_policy del '<Rule UID>'
```

Syntax for IPv6

```
fw6 [-d] sam_policy del '<Rule UID>'
```

Parameters

Parameter	Description
`-d`	Runs the command in debug mode. Use only if you troubleshoot the command itself. **Best Practice** - If you use this parameter, then redirect the output to a file, or use the `script` command to save the entire CLI session.
`'<Rule UID>'`	Specifies the UID of the rule you wish to delete. **Important:** ■ The quote marks and angle brackets (`'<...>'`) are mandatory. ■ To see the Rule UID, run the *"fw sam_policy get" on page 253* command.

Procedure

1. **List all the existing rules in the Suspicious Activity Monitoring policy database**

   List all the existing rules in the Suspicious Activity Monitoring policy database.

   - For IPv4, run:

     ```
     fw sam_policy get
     ```

   - For IPv6, run:

     ```
     fw6 sam_policy get
     ```

   The rules show in this format:

   ```
   operation=add uid=<Value1,Value2,Value3,Value4> target=...
   timeout=... action=... log= ... name= ... comment=...
   originator= ... src_ip_addr=... req_tpe=...
   ```

   Example for IPv4:

   ```
   operation=add uid=<5ac3965f,00000000,3403a8c0,0000264a>
   target=all timeout=300 action=notify log=log name=Test\ Rule
   comment=Notify\ about\ traffic\ from\ 1.1.1.1
   originator=John\ Doe src_ip_addr=1.1.1.1 req_tpe=ip
   ```

2. **Delete a rule from the list by its UID**

   - For IPv4, run:

     ```
     fw [-d] sam_policy del '<Rule UID>'
     ```

   - For IPv6, run:

     ```
     fw6 [-d] sam_policy del '<Rule UID>'
     ```

   Example for IPv4:

   ```
   fw samp del '<5ac3965f,00000000,3403a8c0,0000264a>'
   ```

3. **Add the flush-only rule**

   - For IPv4, run:

     ```
     fw samp add -t 2 quota flush true
     ```

   - For IPv6, run:

     ```
     fw6 samp add -t 2 quota flush true
     ```

   Explanation:

   The "`fw samp del`" and "`fw6 samp del`" commands only remove a rule from the persistent database. The Security Gateway continues to enforce the deleted rule until the next time you compiled and load a policy. To force the rule deletion immediately, you must enter a flush-only rule right after the "`fw samp del`" and "`fw6 samp del`" command. This flush-only rule immediately deletes the rule you specified in the previous step, and times out in 2 seconds.

   Best Practice - Specify a short timeout period for the flush-only rules. This prevents accumulation of rules that are obsolete in the database.

fw sam_policy get

Description

The "`fw sam_policy get`" and "`fw6 sam_policy get`" commands:

- Show all the configured Suspicious Activity Monitoring (SAM) rules.
- Show all the configured Rate Limiting rules.

Notes:

- These commands are interchangeable:
  - For IPv4: "`fw sam_policy`" and "`fw samp`".
  - For IPv6: "`fw6 sam_policy`" and "`fw6 samp`".
- You can run these commands in Gaia Clish, or Expert mode.
- Security Gateway stores the SAM Policy rules in the `$FWDIR/database/sam_policy.db` file.
- Security Gateway stores the SAM Policy management settings in the `$FWDIR/database/sam_policy.mng` file.

Important:

- Configuration you make with these commands, survives reboot.
- You can run this command in the Expert mode or in Gaia Clish (Gaia gClish on Scalable Platforms).
- In a Cluster, you must configure all the Cluster Members in the same way.
- On Scalable Platforms, you must connect to the applicable Security Group.
- VSNext mode and Traditional VSX mode do **not** support Suspicious Activity Monitoring (SAM) Rules. See sk79700.

Best Practice - The SAM Policy rules consume some CPU resources on Security Gateway. Set an expiration for rules that gives you time to investigate, but does not affect performance. Keep only the required SAM Policy rules. If you confirm that an activity is risky, edit the Security Policy, educate users, or otherwise handle the risk.

Syntax for IPv4

```
fw [-d] sam_policy get [-l] [-u '<Rule UID>'] [-k '<Key>' -t
<Type> [+{-v '<Value>'}] [-n]]
```

Syntax for IPv6

```
fw6 [-d] sam_policy get [-l] [-u '<Rule UID>'] [-k '<Key>' -t
<Type> [+{-v '<Value>'}] [-n]]
```

Parameters

Note - All these parameters are optional.

Parameter	Description
`-d`	Runs the command in debug mode. Use only if you troubleshoot the command itself. **Best Practice** - If you use this parameter, then redirect the output to a file, or use the `script` command to save the entire CLI session.
`-l`	Controls how to print the rules: ■ In the default format (without "`-l`"), the output shows each rule on a separate line. ■ In the list format (with "`-l`"), the output shows each parameter of a rule on a separate line. ■ See *"fw sam_policy add" on page 234*.
`-u '<Rule UID>'`	Prints the rule specified by its Rule UID or its zero-based rule index. The quote marks and angle brackets ('<...>') are mandatory.
`-k '<Key>'`	Prints the rules with the specified predicate key. The quote marks are mandatory.
`-t <Type>`	Prints the rules with the specified predicate type. For Rate Limiting rules, you must always use "`-t in`".
`+{-v '<Value>'}`	Prints the rules with the specified predicate values. The quote marks are mandatory.
`-n`	Negates the condition specified by these predicate parameters: ■ `-k` ■ `-t` ■ `+-v`

Examples

Example 1 - Output in the default format

```
[Expert@HostName:0]# fw samp get

operation=add uid=<5ac3965f,00000000,3403a8c0,0000264a> target=all timeout=300 action=notify
log=log name=Test\ Rule comment=Notify\ about\ traffic\ from\ 1.1.1.1 originator=John\ Doe
src_ip_addr=1.1.1.1 req_tpe=ip
```

Example 2 - Output in the list format

```
[Expert@HostName:0]# fw samp get -l

uid
<5ac3965f,00000000,3403a8c0,0000264a>
target
all
timeout
2147483647
action
notify
log
log
name
Test\ Rule
comment
Notify\ about\ traffic\ from\ 1.1.1.1
originator
John\ Doe
src_ip_addr
1.1.1.1
req_type
ip
```

Example 3 - Printing a rule by its Rule UID

```
[Expert@HostName:0]# fw samp get -u '<5ac3965f,00000000,3403a8c0,0000264a>'
0
operation=add uid=<5ac3965f,00000000,3403a8c0,0000264a> target=all timeout=300 action=notify
log=log name=Test\ Rule comment=Notify\ about\ traffic\ from\ 1.1.1.1 originator=John\ Doe
src_ip_addr=1.1.1.1 req_tpe=ip
```

Example 4 - Printing rules that match the specified filters

```
[Expert@HostName:0]# fw samp get
no corresponding SAM policy requests
[Expert@HostName:0]#
[Expert@HostName:0]# fw samp add -a d -l r -t 3600 quota service any source
range:172.16.7.11-172.16.7.13 new-conn-rate 5 flush true
[Expert@HostName:0]#
[Expert@HostName:0]# fw samp add -a n -l r quota service 1,50-51,6/443,17/53 service-negated
true source cc:QQ byte-rate 0
[Expert@HostName:0]#
[Expert@HostName:0]# fw samp add -a b quota source range:172.16.8.17-172.16.9.121 service
6/80
[Expert@HostName:0]#
[Expert@HostName:0]# fw samp add -a d quota service any source-negated true source cc:QQ
concurrent-conns-ratio 655 track source
[Expert@HostName:0]#
[Expert@HostName:0]# fw samp get
operation=add uid=<5bab3acf,00000000,3503a8c0,00003ddc> target=all timeout=indefinite
action=drop service=any source-negated=true source=cc:QQ concurrent-conns-ratio=655
track=source req_type=quota
operation=add uid=<5bab3ac6,00000000,3503a8c0,00003dbf> target=all timeout=3586 action=drop
log=log service=any source=range:172.16.7.11-172.16.7.13 new-conn-rate=5 flush=true req_
type=quota
operation=add uid=<5bab3acc,00000000,3503a8c0,00003dd7> target=all timeout=indefinite
action=bypass source=range:172.16.8.17-172.16.9.121 service=6/80 req_type=quota
operation=add uid=<5bab3ac9,00000000,3503a8c0,00003dd5> target=all timeout=indefinite
action=notify log=log service=1,50-51,6/443,17/53 service-negated=true source=cc:QQ byte-
rate=0 req_type=quota
[Expert@HostName:0]#
[Expert@HostName:0]# fw samp get -k 'service' -t in -v '6/80'
operation=add uid=<5bab3acc,00000000,3503a8c0,00003dd7> target=all timeout=indefinite
action=bypass source=range:172.16.8.17-172.16.9.121 service=6/80 req_type=quota
[Expert@HostName:0]#
[Expert@HostName:0]# fw samp get -k 'service-negated' -t in -v 'true'
operation=add uid=<5bab3ac9,00000000,3503a8c0,00003dd5> target=all timeout=indefinite
action=notify log=log service=1,50-51,6/443,17/53 service-negated=true source=cc:QQ byte-
rate=0 req_type=quota
[Expert@HostName:0]#
[Expert@HostName:0]# fw samp get -k 'source' -t in -v 'cc:QQ'
operation=add uid=<5bab3acf,00000000,3503a8c0,00003ddc> target=all timeout=indefinite
action=drop service=any source-negated=true source=cc:QQ concurrent-conns-ratio=655
track=source req_type=quota
operation=add uid=<5bab3ac9,00000000,3503a8c0,00003dd5> target=all timeout=indefinite
action=notify log=log service=1,50-51,6/443,17/53 service-negated=true source=cc:QQ byte-
rate=0 req_type=quota
[Expert@HostName:0]#
[Expert@HostName:0]# fw samp get -k source -t in -v 'cc:QQ' -n
operation=add uid=<5bab3ac6,00000000,3503a8c0,00003dbf> target=all timeout=3291 action=drop
log=log service=any source=range:172.16.7.11-172.16.7.13 new-conn-rate=5 flush=true req_
type=quota
operation=add uid=<5bab3acc,00000000,3503a8c0,00003dd7> target=all timeout=indefinite
action=bypass source=range:172.16.8.17-172.16.9.121 service=6/80 req_type=quota
[Expert@HostName:0]#
[Expert@HostName:0]# fw samp get -k 'source-negated' -t in -v 'true'
operation=add uid=<5baa94e0,00000000,860318ac,00003016> target=all timeout=indefinite
action=drop service=any source-negated=true source=cc:QQ concurrent-conns-ratio=655
track=source req_type=quota
[Expert@HostName:0]#
[Expert@HostName:0]# fw samp get -k 'byte-rate' -t in -v '0'
operation=add uid=<5baa9431,00000000,860318ac,00002efd> target=all timeout=indefinite
action=notify log=log service=1,50-51,6/443,17/53 service-negated=true source=cc:QQ byte-
rate=0 req_type=quota
[Expert@HostName:0]#
[Expert@HostName:0]# fw samp get -k 'flush' -t in -v 'true'
operation=add uid=<5baa9422,00000000,860318ac,00002eea> target=all timeout=2841 action=drop
log=log service=any source=range:172.16.7.11-172.16.7.13 new-conn-rate=5 flush=true req_
type=quota
[Expert@HostName:0]#
[Expert@HostName:0]# fw samp get -k 'concurrent-conns-ratio' -t in -v '655'
operation=add uid=<5baa94e0,00000000,860318ac,00003016> target=all timeout=indefinite
```

```
action=drop service=any source-negated=true source=cc:QQ concurrent-conns-ratio=655
track=source req_type=quota
[Expert@HostName:0]#
```

The /proc/ppk/ and /proc/ppk6/ entries

Description

SecureXL supports Linux **/proc** entries. The read-only entries in the **/proc/ppk/** and **/proc/ppk6/** contain various data about SecureXL.

Syntax for IPv4

```
[Expert@MyGW:0]# ls -lR /proc/ppk/
[Expert@MyGW:0]# cat /proc/ppk/<Name of File>
[Expert@MyGW:0]# cat /proc/ppk/<SecureXL Instance ID>/<Name of File>
```

Syntax for IPv6

```
[Expert@MyGW:0]# ls -lR /proc/ppk6/
[Expert@MyGW:0]# cat /proc/ppk6/<Name of File>
[Expert@MyGW:0]# cat /proc/ppk6/<SecureXL Instance ID>/<Name of File>
```

Files

File	Description
`affinity`	Contains status and the thresholds for SecureXL New Affinity mechanism. See *"/proc/ppk/affinity" on page 261*.
`conf`	Contains the SecureXL configuration and basic statistics. See *"/proc/ppk/conf" on page 262*.
`conns`	Contains the list of the SecureXL connections. See *"/proc/ppk/conns" on page 263*.
`cpls`	Contains SecureXL configuration for ClusterXL Load Sharing (CPLS). See *"/proc/ppk/cpls" on page 264*.
`cqstats`	Contains statistics for SecureXL connections queue. See *"/proc/ppk/cqstats" on page 265*.
`drop_ statistics`	Contains SecureXL statistics for dropped packets. See *"/proc/ppk/drop_statistics" on page 266*.

File	Description
`ifs`	Contains the list of interfaces that SecureXL uses. See *"/proc/ppk/ifs" on page 267*.
`mcast_ statistics`	Contains SecureXL statistics for multicast traffic. See *"/proc/ppk/mcast_statistics" on page 272*.
`nac`	Contains SecureXL statistics for Identity Awareness Network Access Control (NAC) traffic. See *"/proc/ppk/nac" on page 273*.
`notify_ statistics`	Contains SecureXL statistics for notifications SecureXL sent to Firewall about accelerated connections. See *"/proc/ppk/notify_statistics" on page 274*.
`profile_ cpu_stat`	Contains IDs of the CPU cores and status of Traffic Profiling See *"/proc/ppk/profile_cpu_stat" on page 276*.
`rlc`	Contains SecureXL statistics for drops due to Rate Limiting for DoS Mitigation. See *"/proc/ppk/rlc" on page 277*.
`statistics`	Contains SecureXL overall statistics. See *"/proc/ppk/statistics" on page 278*.
`stats`	Contains the IRQ numbers and names of interfaces the SecureXL uses. See *"/proc/ppk/stats" on page 280*.
`viol_ statistics`	Contains SecureXL statistics for violations - packets SecureXL forwarded (F2F) to the Firewall. See *"/proc/ppk/viol_statistics" on page 281*.

/proc/ppk/affinity

Description

Contains the number of accelerated packets per second and rate of encrypted bytes.

Syntax for IPv4

```
[Expert@MyGW:0]# ls -lR /proc/ppk/
[Expert@MyGW:0]# cat /proc/ppk/affinity
```

Syntax for IPv6

```
[Expert@MyGW:0]# ls -lR /proc/ppk6/
[Expert@MyGW:0]# cat /proc/ppk6/affinity
```

Example for IPv4

```
[Expert@MyGW:0]# cat /proc/ppk/affinity
Current accelerated PPS        : 0
Current enc. bytes rate        : 0
[Expert@MyGW:0]#
```

/proc/ppk/conf

Description

Contains the SecureXL configuration and basic statistics.

Syntax for IPv4

```
[Expert@MyGW:0]# ls -lR /proc/ppk/
[Expert@MyGW:0]# cat /proc/ppk/conf
[Expert@MyGW:0]# cat /proc/ppk/<SecureXL Instance ID>/conf
```

Syntax for IPv6

```
[Expert@MyGW:0]# ls -lR /proc/ppk6/
[Expert@MyGW:0]# cat /proc/ppk6/conf
[Expert@MyGW:0]# cat /proc/ppk6/<SecureXL Instance ID>/conf
```

Example for IPv4

```
[Expert@MyGW:0]# cat /proc/ppk/conf
Flags                          : 0x00000592
Accounting Update Interval     : 3600
Conn Refresh Interval          : 512
SA Sync Notification Interval  : 200000
UDP Encapsulation Port         : 2746
Min TCP MSS                    : 0
TCP End Timeout                : 5
Connection Limit               : 18446744073709551615

Total Number of conns          : 0
Number of Crypt conns          : 0
Number of TCP conns            : 0
Number of Non-TCP conns        : 0
Total Number of corrs          : 0

Debug flags :
 0      : 0x1
 1      : 0x1
 2      : 0x1
 3      : 0x1
 4      : 0x1
 5      : 0x1
 6      : 0x1
 7      : 0x1
 8      : 0x100
 9      : 0x8
 10      : 0x1
 11      : 0x10
[Expert@MyGW:0]#
```

/proc/ppk/conns

Description

Contains the list of the SecureXL connections.

Important - This file is for future use. Refer to the *"fwaccel conns" on page 57* command.

Syntax for IPv4

```
[Expert@MyGW:0]# ls -lR /proc/ppk/
[Expert@MyGW:0]# cat /proc/ppk/conns
[Expert@MyGW:0]# cat /proc/ppk/<SecureXL Instance ID>/conns
```

Syntax for IPv6

```
[Expert@MyGW:0]# ls -lR /proc/ppk6/
[Expert@MyGW:0]# cat /proc/ppk6/conns
[Expert@MyGW:0]# cat /proc/ppk6/<SecureXL Instance ID>/conns
```

/proc/ppk/cpls

Description

Contains SecureXL configuration for ClusterXL Load Sharing (CPLS).

Important - This file is for future use. Refer to the "`fwaccel cfg -h`" command (see *"fwaccel cfg" on page 53*).

Syntax for IPv4

```
[Expert@MyGW:0]# ls -lR /proc/ppk/
[Expert@MyGW:0]# cat /proc/ppk/cpls
```

Syntax for IPv6

```
[Expert@MyGW:0]# ls -lR /proc/ppk6/
[Expert@MyGW:0]# cat /proc/ppk6/cpls
```

Example for IPv4

```
[Expert@MyGW:0]# cat /proc/ppk/cpls
fwha_conf_flags: 638
fwha_df_type: 0
fwha_member_id: 0
fwha_port: 8116
FWHAP MAC magic: 0
Forwarding MAC magic: 0
My state: ACTIVE
udp_enc_port: 0
selection table size: 0
[Expert@MyGW:0]#
```

/proc/ppk/cqstats

Description

Contains statistics for SecureXL connections queue.

Syntax for IPv4

```
[Expert@MyGW:0]# ls -lR /proc/ppk/
[Expert@MyGW:0]# cat /proc/ppk/cqstats
[Expert@MyGW:0]# cat /proc/ppk/<SecureXL Instance ID>/cqstats
```

Syntax for IPv6

```
[Expert@MyGW:0]# ls -lR /proc/ppk6/
[Expert@MyGW:0]# cat /proc/ppk6/cqstats
[Expert@MyGW:0]# cat /proc/ppk6/<SecureXL Instance ID>/cqstats
```

Example for IPv4

```
[Expert@MyGW:0]# cat /proc/ppk/cqstats
Name                  Value              Name                  Value
--------------------  ---------------    --------------------  ---------------
Queued pkts                          0    Queue fail                           0
Dequeue & f2f                        0    Dequeue & drop                       0
Dequeue & resume                     0    Async index req                      0
Err Async index req                  0    Async index cb                       0
Err Async index cb                   0    Queue alloc fail                     0
Queue empty err                      0
[Expert@MyGW:0]#
```

/proc/ppk/drop_statistics

Description

Contains SecureXL statistics for dropped packets.

Note - This is the same information that the "`fwaccel stats -d`" command shows (see *"fwaccel stats" on page 147*).

Syntax for IPv4

```
[Expert@MyGW:0]# ls -lR /proc/ppk/
```

```
[Expert@MyGW:0]# cat /proc/ppk/drop_statistics
```

```
[Expert@MyGW:0]# cat /proc/ppk/<SecureXL Instance ID>/drop_
statistics
```

Syntax for IPv6

```
[Expert@MyGW:0]# ls -lR /proc/ppk6/
```

```
[Expert@MyGW:0]# cat /proc/ppk6/drop_statistics
```

```
[Expert@MyGW:0]# cat /proc/ppk6/<SecureXL Instance ID>/drop_
statistics
```

Example for IPv4

```
[Expert@MyGW:0]# cat /proc/ppk/drop_statistics
Reason                Packets            Reason                Packets
--------------------  ---------------    --------------------  ---------------
general reason                      0    CPASXL decision                     0
PSLXL decision                      0    clr pkt on vpn                      0
encrypt failed                      0    drop template                       0
decrypt failed                      0    interface down                      0
cluster error                       0    XMT error                           0
anti spoofing                       0    local spoofing                      0
sanity error                        0    monitored spoofed                   0
QOS decision                        0    C2S violation                       0
S2C violation                       0    Loop prevention                     0
DOS Fragments                       0    DOS IP Options                      0
DOS Blacklists                      0    DOS Penalty Box                     0
DOS Rate Limiting                   0    Syn Attack                          0
Reorder                             0    Defrag timeout                      0
[Expert@MyGW:0]#
```

/proc/ppk/ifs

Description

Contains the list of interfaces that SecureXL uses.

Syntax for IPv4

```
[Expert@MyGW:0]# ls -lR /proc/ppk/
```

```
[Expert@MyGW:0]# cat /proc/ppk/ifs
```

Syntax for IPv6

```
[Expert@MyGW:0]# ls -lR /proc/ppk6/
```

```
[Expert@MyGW:0]# cat /proc/ppk6/ifs
```

Example for IPv4

```
[Expert@MyGW:0]# cat /proc/ppk/ifs
 No | Interface | Address         | IRQ | F   | SIM F | Dev                | Output Func
| Features
-------------------------------------------------------------------------------------------------
-------------
  2 | eth0      |    192.168.3.52 |  67 |   1 |   480 | 0xffff81023e5df000 | 0x000013a0
  3 | eth1      |     10.20.30.52 |  83 |   1 |   488 | 0xffff81023dd0c000 | 0x000013a0
  4 | eth2      |     40.50.60.52 |  59 |   1 |   480 | 0xffff810237f88000 | 0x000013a0
  5 | eth3      |         0.0.0.0 |  67 |   1 |    80 | 0xffff810239b3d000 | 0x000013a0
  6 | eth4      |         0.0.0.0 |  91 |   1 |    80 | 0xffff81023841f000 | 0x000013a0
  7 | eth5      |         0.0.0.0 |  83 |   1 |   480 | 0xffff8102396fe000 | 0x000013a0
  8 | eth6      |         0.0.0.0 |  59 |   1 |   480 | 0xffff810239a4d000 | 0x000013a0
 10 | bond0     |     70.80.90.52 |   0 |   1 |   280 | 0xffff8101f1a0e000 | 0x000013a0
[Expert@MyGW:0]#
```

Example for IPv6

```
[Expert@MyGW:0]# cat /proc/ppk6/ifs
  No | Interface | Address         | IRQ | F   | SIM F | Dev                | Output Func
| Features
-------------------------------------------------------------------------------------------------------
-------------
  2 | eth0      |           fe80:0:0:0:250:56ff:fea3:1807 |  67 |   1 |   480 |
0xffff81023e5df000 | 0x000013a0
  3 | eth1      |           fe80:0:0:0:250:56ff:fea3:15a4 |  83 |   1 |   480 |
0xffff81023dd0c000 | 0x000013a0
  4 | eth2      |           fe80:0:0:0:250:56ff:fea3:2f50 |  59 |   1 |   480 |
0xffff810237f88000 | 0x000013a0
  5 | eth3      |                         0:0:0:0:0:0:0:0 |  67 |   1 |    80 |
0xffff810239b3d000 | 0x000013a0
  6 | eth4      |                         0:0:0:0:0:0:0:0 |  91 |   1 |    80 |
0xffff81023841f000 | 0x000013a0
  7 | eth5      |           fe80:0:0:0:250:56ff:fea3:75a9 |  83 |   1 |   480 |
0xffff8102396fe000 | 0x000013a0
  8 | eth6      |           fe80:0:0:0:250:56ff:fea3:5d4c |  59 |   1 |   480 |
0xffff810239a4d000 | 0x000013a0
 10 | bond0     |           fe80:0:0:0:250:56ff:fea3:287b |   0 |   1 |   280 |
0xffff8101f1a0e000 | 0x000013a0
[Expert@MyGW:0]#
```

Explanation about the configuration flags in the "F" and "SIM F" columns

The "**F**" column shows the internal configuration flags that Firewall set on these interfaces.

The "**SIM F**" column shows the internal configuration flags that SecureXL set on these interfaces.

Flag	Description
0x001	If this flag is set, the SecureXL drops the packet at the end of the inbound inspection, if the packet is a "cut-through" packet. In outbound, SecureXL forwards all the packets to the network.
0x002	If this flag is set, the SecureXL sends an applicable notification when a TCP state change occurs (connection is established or torn down).
0x004	If this flag is set, the SecureXL it sets the UDP header's checksum field correctly when the SecureXL encapsulates an encrypted packet (UDP encapsulation). If this flag is not set, SecureXL sets the UDP header's checksum field to zero. It is safe to ignore this flag, if it is set to 0 (SecureXL continues to calculate the UDP packet's checksum).
0x008	If this flag is set, the SecureXL does not create new connections that match a template, and SecureXL drops the packet that matches the template, when the number of entries in the Connections Table reaches the specified limit. If this flag is not set, the SecureXL forwards the packet to the Firewall.

Flag	Description
0x010	If this flag is set, the SecureXL forwards fragments to the Firewall.
0x020	If this flag is set, the SecureXL does not create connections from TCP templates anymore. The Firewall offloads connections to SecureXL when necessary. This flag only disables the creation of TCP templates.
0x040	If this flag is set, the SecureXL notifies the Firewall at intervals, so it refreshes the accelerated connections in the Firewall kernel tables.
0x080	If this flag is set, the SecureXL does not create connections from non-TCP templates anymore. The Firewall offloads connections to SecureXL when necessary. This flag only disables the creation of non-TCP templates.
0x100	If this flag is set, the SecureXL allows sequence verification violations for connections that did not complete the TCP 3-way handshake process. If this flag is not set, SecureXL must forward the violating packets to the Firewall.
0x200	If this flag is set, the SecureXL allows sequence verification violations for connections that completed the TCP 3-way handshake process. If this flag is not set, SecureXL must forward the violating packets to the Firewall.
0x400	If this flag is set, the SecureXL forwards TCP [RST] packets to the Firewall.
0x0001	If this flag is set, the SecureXL notifies the Firewall about HitCount data.
0x0002	If this flag is set, the VSX Virtual System works as a junction, rather than a regular Virtual System (only the local Virtual System flag is applicable).
0x0004	If this flag is set, the SecureXL disables the reply counter of inbound encrypted traffic. At a result, SecureXL kernel module works in the same way as the VPN kernel module.
0x0008	If this flag is set, the SecureXL enables the MSS Clamping. Refer to the kernel parameters "`fw_clamp_tcp_mss`" and "`fw_clamp_vpn_mss`" in sk101219.
0x0010	If this flag is set, the SecureXL disables the "No Match Ranges" (NMR) Templates (see sk117755).

Flag	Description
0x0020	If this flag is set, the SecureXL disables the "No Match Time" (NMT) Templates (see sk117755).
0x0040	If this flag is set, the SecureXL does not send Drop Templates notifications about dropped packets to the Firewall (to update the drop counters). For example, if you set the value of the kernel parameter "`activate_optimize_drops_support_now`" to 1, it disables the Drop Templates notifications.
0x0080	If this flag is set, the SecureXL enables the MultiCore support for IPsec VPN (see sk118097).
0x0100	If this flag is set, the SecureXL enables the support for CoreXL Dynamic Dispatcher (see sk105261).
0x0800	If this flag is set, the SecureXL does not enforce the Path MTU Discovery for IP multicast packets.
0x1000	If this flag is set, the SecureXL disables the SIM "drop_templates" feature.
0x2000	If this flag is set, it indicates that an administrator enabled the Link Selection Load Sharing feature.
0x4000	If this flag is set, the SecureXL disables the asynchronous notification feature.
0x8000	If this flag is set, it indicates that the capacity of the Firewall Connections Table is unlimited.

Examples:

Value	Description
0x039	Means the sum of these flags: ■ 0x001 ■ 0x008 ■ 0x010 ■ 0x020
0x00008a16	Means the sum of these flags: ■ 0x0002 ■ 0x0004 ■ 0x0010 ■ 0x0200 ■ 0x0800 ■ 0x8000
0x00009a16	Means the sum of these flags: ■ 0x0002 ■ 0x0004 ■ 0x0010 ■ 0x0200 ■ 0x0800 ■ 0x1000 ■ 0x8000

/proc/ppk/mcast_statistics

Description

Contains SecureXL statistics for multicast traffic.

Note - This is the same information that the "`fwaccel stats -m`" command shows (see *"fwaccel stats" on page 147*).

Syntax for IPv4

```
[Expert@MyGW:0]# ls -lR /proc/ppk/
```

```
[Expert@MyGW:0]# cat /proc/ppk/mcast_statistics
```

```
[Expert@MyGW:0]# cat /proc/ppk/<SecureXL Instance ID>/mcast_
statistics
```

Syntax for IPv6

```
[Expert@MyGW:0]# ls -lR /proc/ppk6/
```

```
[Expert@MyGW:0]# cat /proc/ppk6/mcast_statistics
```

```
[Expert@MyGW:0]# cat /proc/ppk6/<SecureXL Instance ID>/mcast_
statistics
```

Example for IPv4

```
[Expert@MyGW:0]# cat /proc/ppk/mcast_statistics
Name                  Value              Name                  Value
--------------------  ----------------   --------------------  ---------------
in packets                       10100   out packets                       0
if restricted                        0   conns with down if                0
f2f packets                          0   f2f bytes                         0
dropped packets                      0   dropped bytes                     0
accel packets                        0   accel bytes                       0
mcast conns                          0
[Expert@MyGW:0]#
```

/proc/ppk/nac

Description

Contains SecureXL statistics for Identity Awareness Network Access Control (NAC) traffic.

Note - This is the same information that the "`fwaccel stats -n`" command shows (see *"fwaccel stats" on page 147*).

Syntax for IPv4

```
[Expert@MyGW:0]# ls -lR /proc/ppk/
[Expert@MyGW:0]# cat /proc/ppk/nac
[Expert@MyGW:0]# cat /proc/ppk/<SecureXL Instance ID>/nac
```

Syntax for IPv6

```
[Expert@MyGW:0]# ls -lR /proc/ppk6/
[Expert@MyGW:0]# cat /proc/ppk6/nac
[Expert@MyGW:0]# cat /proc/ppk6/<SecureXL Instance ID>/nac
```

Example for IPv4

```
[Expert@MyGW:0]# cat /proc/ppk/nac
Name                  Value              Name                  Value
--------------------  ---------------    --------------------  ---------------
NAC packets                          0    NAC bytes                            0
NAC connections                      0    complience failure                   0
[Expert@MyGW:0]#
```

/proc/ppk/notify_statistics

Description

Contains SecureXL statistics for notifications SecureXL sent to Firewall about accelerated connections.

Syntax for IPv4

```
[Expert@MyGW:0]# ls -lR /proc/ppk/
```

```
[Expert@MyGW:0]# cat /proc/ppk/notify_statistics
```

```
[Expert@MyGW:0]# cat /proc/ppk/<SecureXL Instance ID>/notify_
statistics
```

Syntax for IPv6

```
[Expert@MyGW:0]# ls -lR /proc/ppk6/
```

```
[Expert@MyGW:0]# cat /proc/ppk6/notify_statistics
```

```
[Expert@MyGW:0]# cat /proc/ppk6/<SecureXL Instance ID>/notify_
statistics
```

Example for IPv4

```
[Expert@MyGW:0]# cat /proc/ppk/notify_statistics
Notification           Packets           Notification           Packets
--------------------  --------------    --------------------  --------------
ntSAAboutToExpire                 0    ntSAExpired                         0
ntMSPIError                       0    ntNoInboundSA                       0
ntNoOutboundSA                    0    ntDataIntegrityFailed               0
ntPossibleReplay                  0    ntReplay                            0
ntNextProtocolError               0    ntCPIError                          0
ntClearTextPacket                 0    ntFragmentation                     0
ntUpdateUdpEncTable               0    ntSASync                            0
ntReplayOutOfWindow               0    ntVPNTrafficReport                  0
ntConnDeleted                     0    ntConnUpdate                        0
ntPacketDropped                   0    ntSendLog                           0
ntRefreshGTPTunnel                0    ntMcastDrop                         0
ntAccounting                      0    ntAsyncIndex                        0
ntACkReordering                   0    ntAccelAckInfo                      0
ntMonitorPacket                   0    ntPacketCapture                     0
ntCpasPacketCapture               0    ntPSLGlueUpdateReject               0
ntSeqVerifyDrop                   0    ntPacketForwardBefore               0
ntICMPMessage                     0    ntQoSReclassifyPacket               0
ntQoSResumePacket                 0    ntVPNEncHaLinkFailure               0
ntVPNEncLsLinkFailure             0    ntVPNEncRouteChange                 0
ntVPNDecVerRouteChang             0    ntVPNDecRouteChange                 0
ntMuxSimToFw                      0    ntPSLEventLog                       0
ntSendCPHWDStats              39375    ntPacketTaggingViolat               0
ntDosNotify                       0    ntSynatkNotify                      0
ntSynatkStats                     0    ntQoSEventLog                       0
ntPrintGetParam                   0
[Expert@MyGW:0]#
```

/proc/ppk/profile_cpu_stat

Description

This file is for Check Point use only.

Contains IDs of the CPU cores and status of Traffic Profiling:

- The first column shows the IDs of the CPU cores.
- The second column shows the status of Traffic Profiling for the applicable CPU core.

Syntax for IPv4

```
[Expert@MyGW:0]# ls -lR /proc/ppk/
```

```
[Expert@MyGW:0]# cat /proc/ppk/profile_cpu_stat
```

```
[Expert@MyGW:0]# cat /proc/ppk/<SecureXL Instance ID>/profile_cpu_
stat
```

Syntax for IPv6

```
[Expert@MyGW:0]# ls -lR /proc/ppk6/
```

```
[Expert@MyGW:0]# cat /proc/ppk6/profile_cpu_stat
```

```
[Expert@MyGW:0]# cat /proc/ppk6/<SecureXL Instance ID>/profile_
cpu_stat
```

Example for IPv4 from a Security Gateway with 4 CPU cores

```
[Expert@MyGW:0]# cat /proc/ppk/profile_cpu_stat
0 0
1 0
2 0
3 0
[Expert@MyGW:0]#
```

/proc/ppk/rlc

Description

Contains SecureXL statistics for drops due to Rate Limiting for DoS Mitigation.

Syntax for IPv4

```
[Expert@MyGW:0]# ls -lR /proc/ppk/
[Expert@MyGW:0]# cat /proc/ppk/rlc
```

Syntax for IPv6

```
[Expert@MyGW:0]# ls -lR /proc/ppk6/
[Expert@MyGW:0]# cat /proc/ppk6/rlc
```

Example for IPv4

```
[Expert@MyGW:0]# cat /proc/ppk/rlc
Total drop packets : 0
Total drop bytes : 0
[Expert@MyGW:0]#
```

/proc/ppk/statistics

Description

Contains SecureXL overall statistics.

To see these statistics in a better way, run the *"fwaccel stats" on page 147* command.

Syntax for IPv4

```
[Expert@MyGW:0]# ls -lR /proc/ppk/
[Expert@MyGW:0]# cat /proc/ppk/statistics
[Expert@MyGW:0]# cat /proc/ppk/<SecureXL Instance ID>/statistics
```

Syntax for IPv6

```
[Expert@MyGW:0]# ls -lR /proc/ppk6/
[Expert@MyGW:0]# cat /proc/ppk6/statistics
[Expert@MyGW:0]# cat /proc/ppk6/<SecureXL Instance ID>/statistics
```

Example for IPv4

```
[Expert@MyGW:0]# cat /proc/ppk/statistics
Name                  Value             Name                  Value
--------------------  ---------------   --------------------  ---------------
accel packets                     0   accel bytes                       0
outbound packets                  0   outbound bytes                    0
conns created                     0   conns deleted                     0
current total conns               0   TCP conns                         0
non TCP conns                     0   nat conns                         0
dropped packets                 728   dropped bytes                107978
fragments received                0   fragments transmit                0
fragments dropped                 0   fragments expired                 0
IP options stripped               0   IP options restored               0
IP options dropped                0   corrs created                     0
corrs deleted                     0   C corrections                     0
corrected packets                 0   corrected bytes                   0
crypt conns                       0   enc bytes                         0
dec bytes                         0   ESP enc pkts                      0
ESP enc err                       0   ESP dec pkts                      0
ESP dec err                       0   ESP other err                     0
espudp enc pkts                   0   espudp enc err                    0
espudp dec pkts                   0   espudp dec err                    0
espudp other err                  0   acct update interval           3600
CPASXL packets                    0   PSLXL packets                     0
CPASXL async packets              0   PSLXL async packets               0
CPASXL bytes                      0   PSLXL bytes                       0
CPASXL conns                      0   PSLXL conns                       0
CPASXL conns created              0   PSLXL conns created               0
PXL FF conns                      0   PXL FF packets                    0
PXL FF bytes                      0   PXL FF acks                       0
PXL no conn drops                 0   PSL Inline packets                0
PSL Inline bytes                  0   CPAS Inline packets               0
CPAS Inline bytes                 0   Total QoS conns                   0
CLASSIFY                          0   CLASSIFY_FLOW                     0
RECLASSIFY_POLICY                 0   Enq-IN FW pkts                    0
Enq-OUT FW pkts                   0   Deq-IN FW pkts                    0
Deq-OUT FW pkts                   0   Enq-IN FW bytes                   0
Enq-OUT FW bytes                  0   Deq-IN FW bytes                   0
Deq-OUT FW bytes                  0   Enq-IN AXL pkts                   0
Enq-OUT AXL pkts                  0   Deq-IN AXL pkts                   0
Deq-OUT AXL pkts                  0   Enq-IN AXL bytes                  0
Enq-OUT AXL bytes                 0   Deq-IN AXL bytes                  0
Deq-OUT AXL bytes                 0   F2F packets                       0
F2F bytes                         0   TCP violations                    0
F2V conn match pkts               0   F2V packets                       0
F2V bytes                         0   gtp tunnels created               0
gtp tunnels                       0   gtp accel pkts                    0
gtp f2f pkts                      0   gtp spoofed pkts                  0
gtp in gtp pkts                   0   gtp signaling pkts                0
gtp tcpopt pkts                   0   gtp apn err pkts                  0
memory used                38799384   C tcp handshake conn              0
C tcp estab. conns                0   C tcp closed conns                0
C tcp pxl hnshk conn              0   C tcp pxl est. conn               0
C tcp pxl closed                  0   ob cpasxl packets                 0
ob pslxl packets                  0   ob cpasxl bytes                   0
ob pslxl bytes                    0   DNS DoR stats                     0
trimmed pkts
[Expert@MyGW:0]#
```

/proc/ppk/stats

Description

Contains the IRQ numbers and names of interfaces the SecureXL uses.

Syntax for IPv4

```
[Expert@MyGW:0]# ls -lR /proc/ppk/
[Expert@MyGW:0]# cat /proc/ppk/stats
```

Syntax for IPv6

```
[Expert@MyGW:0]# ls -lR /proc/ppk6/
[Expert@MyGW:0]# cat /proc/ppk6/stats
```

Example for IPv4

```
[Expert@MyGW:0]# cat /proc/ppk/stats
IRQ | Interface
---------------------------
 18   eth0
 16   eth1
 17   eth2
 18   eth3
 19   eth4
[Expert@MyGW:0]#
```

/proc/ppk/viol_statistics

Description

Contains SecureXL statistics for violations - packets SecureXL forwarded (F2F) to the Firewall.

Note - This is the same information that the "`fwaccel stats -p`" command shows (see *"fwaccel stats" on page 147*).

Syntax for IPv4

```
[Expert@MyGW:0]# ls -lR /proc/ppk/
```

```
[Expert@MyGW:0]# cat /proc/ppk/viol_statistics
```

Syntax for IPv6

```
[Expert@MyGW:0]# ls -lR /proc/ppk6/
```

```
[Expert@MyGW:0]# cat /proc/ppk6/viol_statistics
```

Example for IPv4

```
[Expert@MyGW:0]# cat /proc/ppk/viol_statistics
Violation             Packets            Violation             Packets
--------------------  ---------------    --------------------  ---------------
pkt has IP options                0    ICMP miss conn                   4
TCP-SYN miss conn               356    TCP-other miss conn        1386954
UDP miss conn                943355    other miss conn                  0
VPN returned F2F                  0    uni-directional viol             0
possible spoof viol               0    TCP state viol                   0
out if not def/accl               0    bridge, src=dst                  0
routing decision err              0    sanity checks failed             0
fwd to non-pivot                  0    broadcast/multicast              0
cluster message              250859051    cluster forward                 0
chain forwarding                  0    F2V conn match pkts              0
general reason                    0    route changes                    0

[Expert@MyGW:0]#
```

SecureXL Debug

To understand how SecureXL processes the traffic, enable the SecureXL debug while the traffic passes through the Security Gateway / ClusterXL / Scalable Platform Security Group.

Warning - Debug increases the load on the CPU on the Security Gateway / Cluster Members / Security Group Members. We recommend you schedule a maintenance window to debug the SecureXL.

For the complete debug procedure, see the *R82 Quantum Security Gateway Guide* > Chapter "**Kernel Debug**".

fwaccel dbg

Description

The *fwaccel dbg* command controls the SecureXL debug. See *"SecureXL Debug Procedure" on page 297*.

Important:

- In a Cluster, you must configure all the Cluster Members in the same way.
- On Scalable Platforms, you must connect to the Gaia Portal of the applicable Security Group.
- If you do **not** use the complete debug procedure with the "`fw ctl kdebug`" command,
  and SecureXL works in Kernel Mode (KPPAK), then Security Gateway writes the debug outputs to these files:
  - `$FWDIR/log/fwk.elg` - Processing of traffic in the Firewall module
  - `/var/log/messages` - Additional information
- If you do **not** use the complete debug procedure with the "`fw ctl kdebug`" command,
  and SecureXL works in User Mode (UPPAK), then Security Gateway writes the debug outputs to these files:
  - `$FWDIR/log/fwk.elg` - Processing of traffic in the Firewall module
  - `/var/log/usim_x86.elg` - Processing of traffic in SecureXL
  - `/var/log/messages` - Processing of traffic in the ADP module (NVIDIA ConnectX 100G Cards)

Syntax in Gaia Clish or the Expert mode on a Security Gateway / ClusterXL:

```
fwaccel dbg
      -h
      -m <Name of SecureXL Debug Module>
      all
      + <Debug Flags>
      - <Debug Flags>
      reset
      -f {"<5-Tuple Debug Filter>" | reset}
      list
      resetall
```

Syntax in Gaia gClish on a Scalable Platform Security Group:

```
fwaccel dbg
      -h
      -m <Name of SecureXL Debug Module>
      all
      + <Debug Flags>
      - <Debug Flags>
      reset
      -f {"<5-Tuple Debug Filter>" | reset}
      list
      resetall
```

Syntax in the Expert mode on a Scalable Platform Security Group:

```
g_fwaccel dbg
      -h
      -m <Name of SecureXL Debug Module>
      all
      + <Debug Flags>
      - <Debug Flags>
      reset
      -f {"<5-Tuple Debug Filter>" | reset}
      list
      resetall
```

Parameters

Parameter	Description
`-h`	Shows the applicable built-in help.
`-m <Name of SecureXL Debug Module>`	Specifies the name of the SecureXL debug module. To see the list of available debug modules, run: `fwaccel dbg`
`all`	Enables all debug flags for the specified debug module.

Parameter	Description
`+ <Debug Flags>`	Enables the specified debug flags for the specified debug module: Syntax: `+ Flag1 [Flag2 Flag3 ... FlagN]` **Note** - You must press the space bar key after the plus (+) character.
`- <Debug Flags>`	Disables all debug flags for the specified debug module. Syntax: `- Flag1 [Flag2 Flag3 ... FlagN]` **Note** - You must press the space bar key after the minus (-) character.
`reset`	Resets all debug flags for the specified debug module to their default state.
`-f "<5-Tuple Debug Filter>"`	Configures the debug filter to show only debug messages that contain the specified connection. The filter is a string of five numbers separated with commas: `"<Source IP Address>,<Source Port>,<Destination IP Address>,<Destination Port>,<Protocol Number>"` **Notes:** ■ You can configure only one debug filter at one time. ■ You can use the asterisk "*" as a wildcard for an IP Address, Port number, or Protocol number. ■ For more information, see *IANA Service Name and Port Number Registry* and *IANA Protocol Numbers*.
`-f reset`	Resets the current debug filter.
`list`	Shows all enabled debug flags in all debug modules.
`resetall`	Reset all debug flags for all debug modules to their default state.

Enabling SecureXL debug flags during boot

From R81.20, you can configure SecureXL debug to start during boot.

Important - In a Cluster, you must configure all the Cluster Members in the same way.

Procedure

Step	Instructions
1	Connect to the command line on the Security Gateway / each Cluster Member. On Scalable Platforms, you must connect to the Gaia Portal of the applicable Security Group.
2	Log in to the Expert mode.
3	Create the require configuration files: ■ To collect the debug for IPv4 traffic: `touch $FWDIR/conf/fwaccel_dbg_flags.cfg` ■ To collect the debug for IPv6 traffic: `touch $FWDIR/conf/fwaccel6_dbg_flags.cfg`
4	Edit the applicable configuration file: ■ To collect the debug for IPv4 traffic: `vi $FWDIR/conf/fwaccel_dbg_flags.cfg` ■ To collect the debug for IPv6 traffic: `vi $FWDIR/conf/fwaccel6_dbg_flags.cfg`
5	Configure the applicable debug modules and the debug flags (see *"SecureXL Debug Modules and Debug Flags" on page 303*). Write each debug module and its flags on a separate line: `<Name of Debug Module #1> <Flag1> <Flag2> <Flag3> ... <FlagN>` `<Name of Debug Module #2> <Flag1> <Flag2> <Flag3> ... <FlagN>` Example: `default conn nat` `pkt drop nat`
6	Save the changes in the file and exit the editor.

Step	Instructions
7	On the Scalable Platform Security Group, you must copy the updated file to all Security Group Members: `asg_cp2blades $FWDIR/conf/fwaccel_dbg_flags.cfg` `asg_cp2blades $FWDIR/conf/fwaccel6_dbg_flags.cfg`
8	Reboot the Security Gateway / each Cluster Member / all Security Group Members.
9	Wait for the issue to occur.
10	Connect to the command line on the Security Gateway / each Cluster Member. On Scalable Platforms, you must connect to the Gaia Portal of the applicable Security Group.
11	Log in to the Expert mode.
12	Reset all the SecureXL debug flags in all SecureXL debug modules: ■ On the Security Gateway / each Cluster Member: `fwaccel dbg resetall` ■ On the Scalable Platform Security Group: `g_fwaccel dbg resetall`
13	Remove all entries from the configuration files: ■ On the Security Gateway / each Cluster Member: `echo '' > $FWDIR/conf/fwaccel_dbg_flags.cfg` `echo '' > $FWDIR/conf/fwaccel6_dbg_flags.cfg` ■ On the Scalable Platform Security Group: `echo '' > $FWDIR/conf/fwaccel_dbg_flags.cfg` `asg_cp2blades $FWDIR/conf/fwaccel_dbg_flags.cfg` `echo '' > $FWDIR/conf/fwaccel6_dbg_flags.cfg` `asg_cp2blades $FWDIR/conf/fwaccel6_dbg_flags.cfg`

Step	Instructions
14	Collect the output files. If SecureXL works in Kernel Mode (KPPAK): ■ `$FWDIR/log/fwk.elg` ■ `/var/log/messages` If SecureXL works in User Mode (UPPAK): ■ `$FWDIR/log/fwk.elg` ■ `/var/log/usim_x86.elg` ■ `/var/log/messages`

Examples

Example 1 - Default output

```
[Expert@MyGW:0]# fwaccel dbg
Usage: fwaccel dbg [-m <...>] [resetall | reset | list | all | +/- <flags>]
   -m <module>             - module of debugging
   -h                      - this help message
   resetall                - reset all debug flags for all modules
   reset                   - reset all debug flags for module
   all                     - set all debug flags for module
   list                    - list all debug flags for all modules
   -f reset | "<5-tuple>"  - filter debug messages
   + <flags>               - set the given debug flags
   - <flags>               - unset the given debug flags

Debug flags can be enabled at boot by adding them to
$FWDIR/conf/fwaccel_dbg_flags.cfg and $FWDIR/conf/fwaccel6_dbg_flags.cfg.
The file format is one line per module.

For example, to enable "tcp_state" and "routing" for the "pkt" module:
    echo "pkt tcp_state routing" >> $FWDIR/conf/fwaccel_dbg_flags.cfg

List of available modules and flags:

Module: default (default)
init drv tag lock cpdrv routing kdrv tcp_sv svm iter conn htab del update acct conf stat
queue ioctl corr util rngs relations ant conn_app rngs_print infra_ids offload nat

Module: db
get save del tmpl tmo init ant profile nmr nmt warning

Module: api
init add update del acct conf stat vpn notif tmpl sv pxl qos gtp infra tmpl_info upd_conf
upd_if_inf add_sa del_sa del_all_sas misc get_features get_tab get_stat reset_stat tag long_
ver del_all_tmpl get_state upd_link_sel

Module: pkt
f2f frag spoof acct notif tcp_state tcp_state_pkt sv cpls routing drop pxl qos user deliver
vlan pkt nat wrp corr caf bhm geneve sctp

Module: infras
reorder pm

Module: tmpl
dtmpl_get dtmpl_notif tmpl

Module: vpn
vpnpkt linksel routing vpn ls

Module: nac
db db_get pkt pkt_ex signature offload idnt ioctl nac

Module: cpaq
init client server exp cbuf opreg transport transport_utils broadcast

Module: synatk
init conf conn log pkt proxy state msg

Module: adp
rt nh eth heth wrp inf mbs bpl bplinf mbeinf if drop bond xmode ipsctl ac_print cpfifo qconf
qcomm filter packet mcast hw_offload hw_expn rte_api

Module: dos
fw1-cfg fw1-pkt sim-cfg sim-pkt detailed-pkt detailed-cfg drop cache

Module: gtp
pkt policy tables api drop notif general
```

```
Module: usdisp
error conn packet api msg state packet_err counter event quota ioctl lock clb uid queue
fwstats cachetab vpn temp_conns prio route dumbo

Module: dpdk_lib


Module: dpdk_pmd


Module: dpdk_other


[Expert@MyGW:0]#
```

Example 2 - Enabling and disabling of debug flags

```
[Expert@MyGW:0]# fwaccel dbg -m default + conn
Debug flags updated.
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dbg list

[Expert@MyGW:0]# fwaccel dbg list
 Module: default (400)
conn

Module: db (0)


Module: api (0)


Module: pkt (0)


Module: infras (0)


Module: tmpl (0)


Module: vpn (0)


Module: nac (0)


Module: cpaq (0)


Module: synatk (0)


Module: adp (0)


Module: dos (0)


Module: gtp (0)


Module: usdisp (0)


Module: dpdk_lib (0)


Module: dpdk_pmd (0)


Module: dpdk_other (0)


Debug filter not set.
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dbg -m default - conn
Debug flags updated.
[Expert@MyGW:0]#
```

```
[Expert@MyGW:0]# fwaccel dbg list

Module: default (0)


Module: db (0)


Module: api (0)


Module: pkt (0)


Module: infras (0)


Module: tmpl (0)


Module: vpn (0)


Module: nac (0)


Module: cpaq (0)


Module: synatk (0)


Module: adp (0)


Module: dos (0)


Module: gtp (0)


Module: usdisp (0)


Module: dpdk_lib (0)


Module: dpdk_pmd (0)


Module: dpdk_other (0)


Debug filter not set.
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dbg -m default reset
Debug flags updated.
[Expert@MyGW:0]#
```

Example 3 - Resetting all debug flags in all debug modules

```
[Expert@MyGW:0]# fwaccel dbg resetall
Debug state was reset to default.
[Expert@MyGW:0]#
```

Example 4 - Configuring debug filter for an SSH connection from 192.168.20.30 to 172.16.40.50

```
[Expert@MyGW:0]# fwaccel dbg -f 192.168.20.30,*,172.16.40.50,22,6
Debug filter was set.
[Expert@MyGW:0]#
[Expert@MyGW:0]# fwaccel dbg list

... ...

Debug filter: "<*,*,*,*,*>"
[Expert@MyGW:0]#
```

SecureXL Debug Procedure

By default, SecureXL writes the output debug information to the `/var/log/messages` file.

To collect the applicable SecureXL debug and to make its analysis easier, follow the steps below.

Note - For the complete debug procedure, see the *R82 Quantum Security Gateway Guide* > Chapter "**Kernel Debug**".

Important:

- We strongly recommend to schedule a full maintenance window to minimize the impact on your production traffic.
- We strongly recommend to connect over serial console to your Security Gateway / each Cluster Member / Scalable Platform Security Group Members. This is to prevent a possible issue when you cannot work with the CLI because of a high load on the CPU.
- In cluster, you must collect this debug from all Cluster Members in the same way.
- Debug the specific SecureXL instance only when you are sure that only that SecureXL instance processes the traffic.

Procedure

1. **Connect to the command line on your Security Gateway / each Cluster Member / Scalable Platform Security Group**

   Use an SSH or a console connection.

   Best Practice - Use a console connection.

   Note - On Scalable Platforms, you must connect to the Gaia Portal of the applicable Security Group.

2. **Log in to the Expert mode**

   If the default shell is Gaia Clish, run:

   ```
   expert
   ```

3. **Reset all kernel debug flags in all kernel debug modules**

- On the Security Gateway / each Cluster Member, run:

```
fw ctl debug 0
```

- On the Scalable Platform Security Group, run:

```
g_fw ctl debug 0
```

4. **Reset all the SecureXL debug flags in all SecureXL debug modules**

   - On the Security Gateway / each Cluster Member:

```
fwaccel dbg resetall
```

   - On the Scalable Platform Security Group:

```
g_fwaccel dbg resetall
```

5. **Allocate the kernel debug buffer**

   - On the Security Gateway / each Cluster Member, run:

```
fw ctl debug -buf 8200 [-v {"<List of VSIDs>" | all}]
```

   - On the Scalable Platform Security Group, run:

```
g_fw ctl debug -buf 8200 [-v {"<List of VSIDs>" | all}]
```

   Note - The optional part "`-v {"<List of VSIDs>" | all}`" is to specify the applicable Virtual Systems on a VSX Gateway or VSX Cluster Member.

6. **Make sure the Security Gateway allocated the kernel debug buffer**

   - On the Security Gateway / each Cluster Member, run:

```
fw ctl debug | grep buffer
```

   - On the Scalable Platform Security Group, run:

```
g_fw ctl debug | grep buffer
```

7. **Configure the applicable kernel debug modules and kernel debug flags**

- On the Security Gateway / each Cluster Member, run:

```
fw ctl debug -m <Name of Kernel Debug Module> {all | +
<Kernel Debug Flags>}
```

- On the Scalable Platform Security Group, run:

```
g_fw ctl debug -m <Name of Kernel Debug Module> {all | +
<Kernel Debug Flags>}
```

8. **Configure the applicable SecureXL debug modules and SecureXL debug flags**

   - On the Security Gateway / each Cluster Member:

```
fwaccel dbg -m <Name of SecureXL Debug Module> {all | +
<SecureXL Debug Flags>}
```

   - On the Scalable Platform Security Group:

```
g_fwaccel dbg -m <Name of SecureXL Debug Module> {all |
+ <SecureXL Debug Flags>}
```

   See *"SecureXL Debug Modules and Debug Flags" on page 303*.

9. **Examine the kernel debug configuration for kernel debug modules**

   - On the Security Gateway / each Cluster Member, run:

```
fw ctl debug
```

   - On the Scalable Platform Security Group, run:

```
g_fw ctl debug
```

10. **Examine the SecureXL debug configuration for SecureXL debug modules**

    - On the Security Gateway / each Cluster Member:

```
fwaccel dbg list
```

    - On the Scalable Platform Security Group:

```
g_fwaccel dbg list
```

11. **Remove all entries from both the Firewall Connections table and SecureXL Connections table**

Important:

- This step makes sure that you collect the debug of the real issue that is not affected by the existing connections.
- **This command deletes all existing connections. This interrupts all connections, including the SSH.**
  Run this command only if you are connected over a serial console to your Security Gateway / each Cluster Member / Scalable Platform Security Group Members.

- On the Security Gateway / each Cluster Member, run:

```
fw tab -t connections -x -y
```

- On the Scalable Platform Security Group, run:

```
g_fw tab -t connections -x -y
```

12. **Remove all entries from the Firewall Templates table**

Note - This command does **not** interrupt the existing connections. This step makes sure that you collect the debug of the real issue that is not affected by the existing connection templates.

- On the Security Gateway / each Cluster Member, run:

```
fw tab -t cphwd_tmpl -x -y
```

- On the Scalable Platform Security Group, run:

```
g_fw tab -t cphwd_tmpl -x -y
```

13. **Start the kernel debug**

In Gateway mode:

- On the Security Gateway / each Cluster Member, run:

```
fw ctl kdebug -T -f -o /var/log/kernel_debug.txt
```

- On the Scalable Platform Security Group, run:

```
g_fw ctl kdebug -T -f -o /var/log/kernel_debug.txt
```

In VSX mode - for specific Virtual Systems:

- On the VSX Gateway / each VSX Cluster Member, run:

```
fw ctl kdebug -v {"<List of VSIDs>" | all} -k -T -f -o
/var/log/kernel_debug.txt
```

- On the Scalable Platform Security Group in VSX mode, run:

```
g_fw ctl kdebug -v {"<List of VSIDs>" | all} -k -T -f -o
/var/log/kernel_debug.txt
```

14. **Replicate the issue, or wait for the issue to occur**

    Perform the steps that cause the issue to occur, or wait for it to occur.

15. **Stop the kernel debug**

    Press **CTRL+C**.

16. **Reset all kernel debug flags in all kernel debug modules**

    - On the Security Gateway / each Cluster Member, run:

```
fw ctl debug 0
```

    - On the Scalable Platform Security Group, run:

```
g_fw ctl debug 0
```

17. **Reset all the SecureXL debug flags in all SecureXL debug modules**

    - On the Security Gateway / each Cluster Member:

```
fwaccel dbg resetall
```

    - On the Scalable Platform Security Group:

```
g_fwaccel dbg resetall
```

18. **Examine the kernel debug configuration to make sure it returned to the default**

- On the Security Gateway / each Cluster Member, run:

```
fw ctl debug
```

- On the Scalable Platform Security Group, run:

```
g_fw ctl debug
```

19. **Examine the SecureXL debug configuration to make sure it returned to the default**

    - On the Security Gateway / each Cluster Member:

```
fwaccel dbg list
```

    - On the Scalable Platform Security Group:

```
g_fwaccel dbg list
```

20. **Analyze the debug output files**

    Transfer these files from the Security Gateway / each Cluster Member / each Security Group Member to your computer:

`/var/log/kernel_debug.txt`
`/var/log/messages*`
`$FWDIR/log/fwk.elg*`
`/var/log/usim_x86.elg*`

    Best Practice - Compress these files with the "`tar -zxvf`" command and transfer the archive from the Security Gateway / each Cluster Member / each Security Group Member to your computer. If you transfer to an FTP server, do so in the binary mode.

SecureXL Debug Modules and Debug Flags

To see the available SecureXL debug modules and their debug flags, run the *"fwaccel dbg" on page 283* command.

Module "default"

Flag	Description
`acct`	Connection accounting information
`ant`	Anticipated connections
`conf`	Configuration of the SecureXL (for example, interfaces)
`conn`	Processing of connections
`conn_app`	Processing of connections
`corr`	Correction layer
`cpdrv`	*Currently not in use*
`del`	Deletion of connections
`drv`	Driver information
`htab`	Hash table
`infra_ids`	Allocating IDs for a given range in Identity Awareness
`init`	Initialization
`ioctl`	Changes in the configuration, which were initiated from the user space
`iter`	Connection table iterator
`kdrv`	Driver information
`lock`	Lock initializing and finalizing
`nat`	Processing of NAT connections
`offload`	Offloading of connections from the Firewall to the SecureXL
`queue`	Connections queue
`relations`	Related connections (such as FTP data connections)
`rngs`	Handling of SecureXL ranges

Flag	Description
`rngs_print`	Printing of SecureXL ranges
`routing`	Handling of SecureXL routing
`stat`	Handling of SecureXL statistics
`svm`	Registering templates or connections for System Counters in Security Gateway object in SmartConsole
`tag`	Tags that were added to the packets by the SecureXL before forwarding them to the Firewall
`tcp_sv`	Verification of sequence in TCP packets
`update`	Updates of connections
`util`	Utilization

Module "pkt" (Packet)

Flag	Description
`acct`	Connection accounting information
`bhm`	BHM Statistics for each CPU
`caf`	Mirror and Decrypt feature - Mirror only of all traffic
`corr`	Correction layer
`cpls`	ClusterXL Load Sharing
`deliver`	Packet delivery
`drop`	Packets dropped by SecureXL
`f2f`	Reason for forwarding a packet to the Firewall
`frag`	Processing of fragments
`geneve`	Processing of Generic Network Virtualization Encapsulation (Geneve) packets in GRE and VxLAN interfaces
`nat`	Processing of NAT connections
`notif`	Notifications sent to the Firewall

Flag	Description
pkt	Processing of packets
pxl	PXL (PacketXL) handling - API between the SecureXL and PSL (Packet Streaming Layer), which is a TCP Streaming engine that parses TCP streams
qos	QoS acceleration
routing	Handling of SecureXL routing
sctp	Handling of Stream Control Transmission Protocol (SCTP) packets
spoof	Handling of SecureXL Anti-Spoofing
sv	Validation of sequence in TCP packets
tcp_ state	Validation of TCP state in TCP packets
tcp_ state_ pkt	Validation of TCP packets
user	Details of a packet
vlan	Handling of VLAN tags
wrp	Handling of WRP interfaces in VSX

Module "db" (Database)

Flag	Description
ant	Anticipated connections
del	Deleting of data from the SecureXL database
get	Retrieving of data from the SecureXL database
init	Initializing and finalizing of SecureXL database
nmr	"No Match Ranges" templates, which allow SecureXL Accept Templates for rules that contain Dynamic objects or Domain objects (or for rules located below such rules)

Flag	Description
`nmt`	"No Match Time" templates, which allow SecureXL Accept Templates for rules that contain Time objects (or for rules located below such rules)
`profile`	Operations on profile table
`save`	Saving of data to the SecureXL database
`tmo`	Handling of timeouts for SecureXL database entries
`tmpl`	Handling of SecureXL templates database
`warning`	General failures to process a packet

Module "api" (Application Programmable Interface)

Flag	Description
`acct`	Connection accounting information
`add`	Adding of connections
`add_sa`	Offloading of VPN SA to SecureXL
`conf`	Configuration of the SecureXL (for example, interfaces)
`del`	Deletion of connections
`del_all_sas`	Deletion of all VPN SAs from SecureXL
`del_all_tmpl`	Deletion of the SecureXL Templates
`del_sa`	Deletion of VPN SA from SecureXL
`get_features`	Getting features buffer (in SecureXL initialization)
`get_stat`	Retrieving of SecureXL statistics
`get_state`	Getting the connection state from SecureXL
`get_tab`	Some extra printouts when processing SecureXL tables
`gtp`	Processing of GPRS Tunnelling Protocol (GTP) tunnel connections

Flag	Description
`infra`	SecureXL infrastructure
`init`	Enabling and disabling of SecureXL
`long_ver`	Prints additional verbose information about connections
`misc`	Prints additional information about SecureXL internals
`notif`	Notifications sent to the Firewall
`pxl`	PXL (PacketXL) handling - API between the SecureXL and PSL (Packet Streaming Layer), which is a TCP Streaming engine that parses TCP streams
`qos`	QoS acceleration
`reset_stat`	Prints statistics IDs that are reset
`stat`	Handling of SecureXL statistics
`sv`	Validation of sequence in TCP packets
`tag`	Tags that were added to the packets by the SecureXL before forwarding them to the Firewall
`tmpl`	Handling of SecureXL Templates
`tmpl_info`	Information about SecureXL Templates
`upd_conf`	Update of SecureXL in ClusterXL Load Sharing
`upd_if_inf`	Prints some text that shows if SecureXL updated information about interfaces
`upd_link_sel`	Updates of VPN Link Selection
`update`	Updates of connections
`vpn`	Processing of VPN connection

Module "adp" (NVIDIA ConnectX 100G Cards)

Flag	Description
`ac_print`	Prints additional information
`bond`	Information about Bond interfaces
`bpl`	Information about packet processing in the backplane
`bplinf`	Information about packet processing in the backplane
`cpfifo`	*Currently is not used*
`drop`	Information about packet drops in the backplane
`eth`	Information about ports from the NVIDIA ConnectX 100G Card's point of view
`filter`	Information about packet processing in the network interface
`heth`	Information about ports from the Host Security Appliance's point of view
`hw_expn`	Information about packet processing in the NVIDIA ConnectX 100G Card
`hw_offload`	Offloading of connections from the Host Security Appliance to the SecureXL
`if`	Information about interfaces
`inf`	Information about slots and ports
`ipsctl`	Information about slots in the Host Security Appliance
`mbeinf`	Information about packet processing in the backplane
`mbs`	Information about packet processing in the backplane
`mcast`	Handling of multicast traffic
`nh`	Handling of next hop routing
`packet`	Information about packet processing in the NVIDIA ConnectX 100G Card
`qcomm`	Information about communication queues
`qconf`	Information about packets in the communication queues
`rt`	Handling of general routing

Flag	Description
`rte_api`	Detailed information about packet processing in the NVIDIA ConnectX 100G Card
`wrp`	Handling of WRP interfaces in VSX
`xmode`	Events in the known neighbors database

Module "infras" (Identity Awareness - Identities Infrastructure)

Flag	Description
`pm`	Pattern Matcher
`reorder`	Reordering of packets in queue

Module "nac" (Identity Awareness - Network Access Control)

Flag	Description
`db`	Updating, adding, deleting of identities
`db_get`	Updating, fetching, searching of identities
`idnt`	Identity Tags
`ioctl`	Changes in the configuration, which were initiated from the user space
`nac`	Network Access Control
`offload`	Offloading of connections from the Firewall to the SecureXL
`pkt`	Forwarding of connections to Firewall (when identity is not found or revoked, or NAC packet tagging verification failed)
`pkt_ex`	NAC packet-tagging verification
`signature`	Signing of packets

Module "vpn" (VPN)

Flag	Description
`linksel`	VPN Link Selection

Flag	Description
`ls`	Forwarding of packets between Cluster Members in a ClusterXL Load Sharing mode
`routing`	VPN Encryption routing information
`vpn`	Processing of VPN connections
`vpnpkt`	Processing of VPN packets

Module "cpaq" (Internal Asynchronous Queue)

Flag	Description
`broadcast`	Processing of broadcast packets
`cbuf`	Information about queue buffers
`client`	Information about queue clients
`error`	General errors
`exp`	Information about expiration of queue items
`init`	Initializing of queue
`opreg`	*Currently not in use*
`server`	Information about queue servers
`transport`	Information about sending messages in queue
`transport_utils`	Additional information about sending messages in queue

Module "dos" (Denial of Service Defender)

Flag	Description
`cache`	*Currently not in use*
`detailed-cfg`	Detailed information about DoS Rate Limiting configuration. **Important** - This debug flag is not suitable for large traffic volumes because it prints a large number of messages. This causes high load on the CPU.

Flag	Description
`detailed-pkt`	Detailed information about DoS Rate Limiting packet flow. **Important** - This debug flag is not suitable for large traffic volumes because it prints a large number of messages. This causes high load on the CPU.
`drop`	Dropped packets
`fw1-cfg`	Information about DoS Rate Limiting configuration in the Firewall kernel module
`fw1-pkt`	Information about DoS Rate Limiting packet flow in the Firewall kernel module
`sim-cfg`	Information about DoS Rate Limiting configuration in the SecureXL kernel module
`sim-pkt`	Information about DoS Rate Limiting packet flow in the SecureXL kernel module

Module "synatk" (Accelerated SYN Defender)

Flag	Description
`conf`	Receiving and updating of Accelerated SYN Defender module's configuration
`conn`	Handling of TCP connections
`init`	Initializing of the Accelerated SYN Defender module
`log`	Prints time of the last sent monitor log and interval between the monitor logs
`msg`	Information about internal messages in the Accelerated SYN Defender module
`pkt`	Handling of TCP packets
`proxy`	*Currently not in use*
`state`	Information about states of the Accelerated SYN Defender module

Module "tmpl" (Drop Templates)

Flag	Description
`dtmpl_get`	Getting of Drop Templates
`dtmpl_notif`	Notifications about Drop Templates
`tmpl`	Information about Drop Templates

Module "gtp" (GPRS Tunnelling Protocol (GTP))

Flag	Description
`api`	Information about GTP tunnels in kernel tables
`drop`	Drops of GTP packets
`general`	General information about APN and kernel tables
`notif`	Information about notification messages
`pkt`	Processing of GTP packets
`policy`	APN information
`tables`	Operations in kernel tables

Module "usdisp" (User Space Dispatcher)

Flag	Description
`api`	*Currently not in use*
`cachetab`	*Currently not in use*
`clb`	*Currently not in use*
`conn`	Processing of connections
`counter`	*Currently not in use*
`dumbo`	*Currently not in use*
`error`	General errors
`event`	*Currently not in use*

Flag	Description
`fwstats`	*Currently not in use*
`ioctl`	IOCTL control messages (communication between kernel and user space processes)
`lock`	*Currently not in use*
`msg`	*Currently not in use*
`packet`	Processing of packets
`packet_err`	Processing of packets
`prio`	*Currently not in use*
`queue`	*Currently not in use*
`quota`	*Currently not in use*
`route`	*Currently not in use*
`state`	*Currently not in use*
`temp_conns`	Shows the selected CoreXL Firewall instance
`uid`	*Currently not in use*
`vpn`	*Currently not in use*

Module "dpdk_lib" (Data Plane Development Kit (DPDK) - Library)

Reserved for future use.

Module "dpdk_pmd" (Data Plane Development Kit (DPDK) - Poll Mode Drivers)

Reserved for future use.

Module "dpdk_other" (Data Plane Development Kit (DPDK) - Other)

Reserved for future use.

CoreXL

CoreXL is a performance-enhancing technology for Security Gateways on multi-core platforms.

CoreXL makes it possible for the CPU cores to perform multiple tasks concurrently. This enhances the Security Gateway performance.

CoreXL provides almost linear scalability of performance, according to the number of processing cores on a single machine. The increase in performance does not require changes to management or to network topology.

On a Security Gateway / Cluster Members / Scalable Platform Security Group with CoreXL enabled, the Firewall kernel is replicated multiple times.

Each replicated copy of the Firewall kernel, or CoreXL Firewall instance, runs on one CPU core.

These CoreXL Firewall instances handle traffic concurrently, and each CoreXL Firewall instance is a complete and independent Firewall inspection kernel. When CoreXL is enabled, all the Firewall kernel instances in the Security Gateway process traffic through the same interfaces and apply the same security policy.

CoreXL Firewall instances work with SecureXL instances.

Enabling and Disabling CoreXL

Important Notes for Cluster:

- You must configure the CoreXL in the same way on all the cluster members.
- If you enable CoreXL, disable CoreXL, or change the number of CoreXL Firewall instances, you should treat this change as a version upgrade. Schedule a full maintenance window and follow the instructions in the *R82 Installation and Upgrade Guide* - Chapter *Upgrading ClusterXL Deployments*. Perform either a *Minimum Effort Upgrade* procedure (requires downtime), or a *Zero Downtime Upgrade* procedure (no downtime, but active connections are lost). Instead of the version upgrade, configure the CoreXL on each cluster member.

To change the CoreXL configuration

Step	Instructions
1	Connect to the command line on the Security Gateway. **Note** - On Scalable Platforms, you must connect to the Gaia Portal of the applicable Security Group.
2	Log in to Gaia Clish or Expert mode. **Note** - On Scalable Platforms, you must run the applicable commands in Gaia gClish or the Expert mode of the applicable Security Group.
3	Run: `cpconfig`
4	Enter the number of the **Check Point CoreXL** option.
5	Enter the number of the applicable option: `(1) Change the number of firewall instances` `(2) Change the number of IPv6 firewall instances` `(3) Disable Check Point CoreXL`
6	Follow the instructions on the screen.
7	Exit from the `cpconfig` menu.

Step	Instructions
8	Reboot. ■ On the Security Gateway (each Cluster Member), run: `reboot` ■ On the Scalable Platform Security Group, run in Gaia gClish: `reboot` ■ On the Scalable Platform Security Group, run in the Expert mode: `g_reboot -a`

Default Configuration of CoreXL

Important - This default configuration applies **only** to Security Gateways that do not support Dynamic Balancing of CoreXL Instances. See *"Dynamic Balancing of CoreXL Instances" on page 341*.

When you enable CoreXL, the default number of CoreXL Firewall instances is based on the total number of CPU cores.

The default affinity setting for all interfaces is automatic when SecureXL is enabled. See *"Allocation of Processing CPU Cores" on page 331*.

Traffic from all interfaces is directed to the CPU cores that run the CoreXL Secure Network Distributor (SND).

Default number of IPv4 CoreXL Firewall instances:

Number of CPU cores	Default number of CoreXL IPv4 FW instances	Default number of Secure Network Distributors (SNDs)
1	1 (CoreXL is disabled)	1 (CoreXL is disabled)
2	2	2
4	3	1
6-20	Number of CPU cores, minus 2	2
More than 20	Number of CPU cores, minus 4. However, no more than 40. **Note** - This limit applies only to the Kernel Mode Firewall (KMFW).	4

The numbers of CoreXL Firewall instances start from zero.

The numbers of CPU cores start from the highest CPU ID allowed by the current Check Point license on your Security Gateway / Cluster Member / Scalable Platform Security Group.

Refer to the **ID** and **CPU** columns in this example:

```
> fw ctl multik stat

ID | Active   | CPU    | Connections | Peak
----------------------------------------------
 0 | Yes      | 7      |           5 |       21
 1 | Yes      | 6      |           3 |       23
 2 | Yes      | 5      |           5 |       25
 3 | Yes      | 4      |           4 |       21
 4 | Yes      | 3      |           5 |       21
 5 | Yes      | 2      |           5 |       20
```

```
>
> fw6 ctl multik stat

ID | Active   | CPU    | Connections | Peak
----------------------------------------------
 0 | Yes      | 7      |           0 |        4
 1 | Yes      | 6      |           0 |        4
>
```

Maximum number of IPv4 CoreXL Firewall instances

Gaia kernel edition	Check Point Appliance	Open Server
64-bit	40	40

Notes:

- Starting in R80.20 and R80.20SP, the Gaia kernel edition is 64-bit only.
- The total number of IPv4 CoreXL Firewall instances and IPv6 CoreXL Firewall instances cannot exceed the numbers in the table above. This limit applies only to the Kernel Mode Firewall (KMFW).

Configuring IPv4 and IPv6 CoreXL Firewall instances

In This Section:

Important Notes for Cluster:

- You must configure the CoreXL in the same way on all the cluster members.
- If you enable CoreXL, disable CoreXL, or change the number of CoreXL Firewall instances, you should treat this change as a version upgrade. Schedule a full maintenance window and follow the instructions in the *R82 Installation and Upgrade Guide* - Chapter *Upgrading ClusterXL Deployments*. Perform either a *Minimum Effort Upgrade* procedure (requires downtime), or a *Zero Downtime Upgrade* procedure (no downtime, but active connections are lost). Instead of the version upgrade, configure the CoreXL on each cluster member.

IPv4 and IPv6 CoreXL Firewall Instances

After you enable Gaia IPv6 support on the Security Gateway / Scalable Platform Security Group (see *R82 Gaia Administration Guide*), configure the CPU cores to run different combinations of IPv4 and IPv6 CoreXL Firewall instances:

- The number of IPv4 CoreXL Firewall instances you can configure is from a minimum of two to a maximum equal to the total number of CPU cores on the Security Gateway / Scalable Platform Security Group:

  ```
  2 <= (Number of IPv4 CoreXL Firewall instances) <= (Total
  Number of CPU cores)
  ```

- By default, the number of IPv6 CoreXL Firewall instances is set to two.

  When the SMT (Hyper-Threading) is enabled, the default number of IPv6 CoreXL Firewall instances is four.

- The number of IPv6 CoreXL Firewall instances you can configure is from a minimum of two to a maximum equal to the total number of IPv4 CoreXL Firewall instances.

The number of IPv6 CoreXL Firewall instances cannot be greater than the number of IPv4 CoreXL Firewall instances:

```
2 <= (Number of IPv6 CoreXL Firewall instances) <= (Total
Number of IPv4 CoreXL Firewall instances)
```

- The total number of IPv4 *and* IPv6 CoreXL Firewall instances cannot be greater than forty:

  Note - This limit applies only to the Kernel Mode Firewall (KMFW).

  ```
  (Number of IPv4 CoreXL Firewall instances) + (Number of IPv6
  CoreXL Firewall instances) <= 40
  ```

Configuring the Number of IPv4 CoreXL Firewall Instances

Step	Instructions
1	Connect to the command line on the Security Gateway / each Cluster Member. **Note** - On Scalable Platforms, you must connect to the Gaia Portal of the applicable Security Group.
2	Log in to Gaia Clish or Expert mode. **Note** - On Scalable Platforms, you must run the applicable commands in Gaia gClish or the Expert mode of the applicable Security Group.
3	Run: `cpconfig`
4	Enter the number of the **Check Point CoreXL** option.
5	Enter **1** to select **Change the number of firewall instances**.
6	Enter the total number of IPv4 CoreXL Firewall instances you wish the Security Gateway to run. **Note** - You can only select a number from the range shown. Follow the instructions on the screen.
7	Exit from the `cpconfig` menu.
8	Reboot. ■ On the Security Gateway (each Cluster Member), run: `reboot` ■ On the Scalable Platform Security Group, run in Gaia gClish: `reboot` ■ On the Scalable Platform Security Group, run in the Expert mode: `g_reboot -a`

Configuring the Number of IPv6 CoreXL Firewall Instances

Step	Instructions
1	Connect to the command line on the Security Gateway / each Cluster Member. **Note** - On Scalable Platforms, you must connect to the Gaia Portal of the applicable Security Group.
2	Log in to Gaia Clish or Expert mode. **Note** - On Scalable Platforms, you must run the applicable commands in Gaia gClish or the Expert mode of the applicable Security Group.
3	Run: `cpconfig`
4	Enter the number of the **Check Point CoreXL** option.
5	Enter **2** to select **Change the number of IPv6 firewall instances**.
6	Enter the total number of IPv6 CoreXL Firewall instances you wish the Security Gateway to run. **Note** - You can only select a number from the range shown. Follow the instructions on the screen.
7	Exit from the `cpconfig` menu.
8	Reboot. ■ On the Security Gateway (each Cluster Member), run: `reboot` ■ On the Scalable Platform Security Group, run in Gaia gClish: `reboot` ■ On the Scalable Platform Security Group, run in the Expert mode: `g_reboot -a`

Example CoreXL Configuration

Security Gateway has four CPU cores.

By default, there are three IPv4 CoreXL Firewall instances and two IPv6 CoreXL Firewall instances:

CPU Core	IPv4 CoreXL Firewall instances	IPv6 CoreXL Firewall instances
CPU 0	N / A	N / A
CPU 1	`fw4_2`	N / A
CPU 2	`fw4_1`	`fw6_1`
CPU 3	`fw4_0`	`fw6_0`

- IPv4 CoreXL Firewall instances: The minimum allowed number is two and the maximum is four
- IPv6 CoreXL Firewall instances: The minimum allowed number is two and the maximum is three

To increase the number of IPv6 CoreXL Firewall instances to four, first you must increase the number of IPv4 CoreXL Firewall instances to the maximum of four and reboot:

```
CoreXL is currently enabled with 3 IPv4 firewall instances and 2 IPv6 firewall instances.

(1) Change the number of firewall instances
(2) Change the number of IPv6 firewall instances
(3) Disable Check Point CoreXL

(4) Exit
Enter your choice (1-4) : 1

How many IPv4 firewall instances would you like to enable (2 to 4) [3] ? 4

CoreXL was enabled successfully with 4 firewall instances.
Important: This change will take effect after reboot.
```

After the reboot, the CoreXL configuration on the Security Gateway looks like this:

CPU Core	IPv4 CoreXL Firewall instances	IPv6 CoreXL Firewall instances
CPU 0	`fw4_3`	N / A
CPU 1	`fw4_2`	N / A
CPU 2	`fw4_1`	`fw6_1`
CPU 3	`fw4_0`	`fw6_0`

Increase the number of IPv6 CoreXL Firewall instances to four and reboot:

```
CoreXL is currently enabled with 4 IPv4 firewall instances and 2 IPv6 firewall instances.

(1) Change the number of firewall instances
(2) Change the number of IPv6 firewall instances
(3) Disable Check Point CoreXL

(4) Exit
Enter your choice (1-4) : 2

How many IPv6 firewall instances would you like to enable (2 to 4)[2] ? 4

CoreXL was enabled successfully with 3 IPv6 firewall instances.
Important: This change will take effect after reboot.
```

After the reboot, the CoreXL configuration on the Security Gateway looks like this:

CPU Core	IPv4 CoreXL Firewall instances	IPv6 CoreXL Firewall instances
CPU 0	`fw4_3`	`fw6_3`
CPU 1	`fw4_2`	`fw6_2`
CPU 2	`fw4_1`	`fw6_1`
CPU 3	`fw4_0`	`fw6_0`

CoreXL Limitations

- R82 CoreXL does not support:
  - Overlapping NAT
  - VPN Traditional Mode
- The global CoreXL Firewall instance #0 (`fw_worker_0`) always processes all the 6in4 traffic.

Configuring Affinity Settings

In This Section:

Introduction

The script `$FWDIR/scripts/fwaffinity_apply` on Security Gateway (Scalable Platform Security Group Members) executes automatically during boot and controls the affinity settings. When you make a change in the affinity settings, the changes do not take effect until you either reboot the Security Gateway (Scalable Platform Security Group), or manually execute the `$FWDIR/scripts/fwaffinity_apply` script.

The `$FWDIR/scripts/fwaffinity_apply` script configures the affinity of interfaces based on the settings in the `$FWDIR/conf/fwaffinity.conf` configuration file. To change these affinity settings, edit that configuration file.

In addition, see *"taskset_us_all" on page 431*.

The $FWDIR/conf/fwaffinity.conf Configuration File

The configuration file `$FWDIR/conf/fwaffinity.conf` controls CoreXL affinity settings.

Each line in this plain-text file uses the same format:

```
<Type> <ID> <CPU_ID>
```

Where:

Field	Allowed Value	Description
<Type>	**i**	Configures the affinity of an interface.
	n	Configures the affinity of a Check Point daemon.
	k	Configures the affinity of a CoreXL Firewall instance.
<ID>	Name of Interface	If **<type> = i**.

Field	Allowed Value	Description
	Name of Daemon	If **<type> = n**.
	ID of CoreXL Firewall instance	If **<type> = k**.
	default	Configures affinities for interfaces that are not specified other lines.
<CPU_ID>	Number (ID) of CPU core	Specifies the ID numbers of processing CPU cores, to which you affine an interface, a Check Point daemon, or a CoreXL Firewall instance.
	all	Specifies all processing CPU cores as available to configure the affinity of an interface, a Check Point daemon, or a CoreXL Firewall instance.
	auto	Configures Automatic mode. See *"Allocation of Processing CPU Cores" on page 331*.
	ignore	No specified affinity. This is useful to exclude an interface from the "**default**" configuration.

Notes:

- The default configuration in this file is:

```
i default auto
```

- Possible combinations:

- To configure the affinity for an interface:

```
i <Name of Interface> {<CPU ID Number> | all | ignore |
auto}
i default {<CPU ID Number> | all | ignore | auto}
```

- To configure the affinity of a Check Point daemon:

```
n <Name of Daemon> {<CPU ID Number> | all | ignore |
auto}
```

- To configure the affinity of a CoreXL Firewall instance:

```
k <ID of CoreXL Firewall instance> {<CPU ID Number> | all
| ignore | auto}
```

- To view the IRQs of all interfaces, run:
  - On a Security Gateway (each Cluster Member), run in Gaia Clish or the Expert mode:

```
fw ctl affinity -l -v -a
```

  - On a Scalable Platform Security Group, run in Gaia gClish:

```
fw ctl affinity -l -v -a
```

  - On a Scalable Platform Security Group, run in the Expert mode:

```
g_fw ctl affinity -l -v -a
```

  See *"fw ctl affinity" on page 396*.

- Interfaces that share an IRQ cannot have different CPU cores as their affinities.

  This also applies when one interface is included in the **default** affinity setting.

  You must either configure the same affinity of all interfaces, or use **ignore** for one of these interfaces.

- On a Scalable Platform Security Group, after you edit the `$FWDIR/conf/fwaffinity.conf` file, you must copy it to all Security Group Members:

```
asg_cp2blades $FWDIR/conf/fwaffinity.conf
```

The $FWDIR/scripts/fwaffinity_apply Script

Syntax

- To execute this shell script on a Security Gateway (each Cluster Member), run in the Expert mode:

```
$FWDIR/scripts/fwaffinity_apply <Parameter>
```

- To execute this shell script on a Scalable Platform Security Group, run in the Expert mode:

```
g_all $FWDIR/scripts/fwaffinity_apply <Parameter>
```

Parameters

Parameter	Description
`-q`	Quiet mode - prints only error messages (standard output goes to `/dev/null`).
`-t i` `-t n` `-t k`	Applies affinity only for the specified type: ■ `-t i` - For interfaces ■ `-t n` - For Check Point daemons ■ `-t k` - For CoreXL Firewall instances

Performance Tuning

This section describes how to fine tune the CoreXL performance.

Allocation of Processing CPU Cores

The CoreXL software architecture includes the Secure Network Distributor (SND).

The SND is responsible for these:

- Processing the incoming traffic from the network interfaces.
- Accelerating authorized packets (when SecureXL is enabled).
- Distributing non-accelerated packets between the CoreXL Firewall instances.

The association of a specific interface with a specific processing CPU core is called the interface's *affinity* with that CPU core. This affinity causes the interface's traffic to be directed to that CPU core and the CoreXL SND to run on that CPU core.

The association of a specific CoreXL Firewall instance with a specific CPU core is called the CoreXL Firewall instance's *affinity* with that CPU core.

The association of a specific user space process with a specific CPU core is called the process's *affinity* with that CPU core.

The default affinity setting for all interfaces is Automatic. Automatic affinity means that if SecureXL is enabled, the affinity of each interface is changed at specific intervals and balanced between the available CPU cores. If SecureXL is disabled, the default affinities of all interfaces are with one available CPU core. In both cases, all processing CPU cores that run a CoreXL Firewall instance, or configured as the affinity of a different user space process, is considered unavailable, and the affinity of interfaces is not set to those CPU cores.

In some cases, which we discuss in the sections below, it can be necessary to change the distribution of CoreXL Firewall instances, the CoreXL SND, and other user space processes, between the processing CPU cores. To do so, you change the affinities of different NICs (interfaces) or user space processes. To make sure CoreXL operates at an efficient level, traffic from all interfaces must be directed to CPU cores that do not run CoreXL Firewall instances. Therefore, if you change affinities for interfaces or other user space processes, you must configure the corresponding number of CoreXL Firewall instances. In addition, you must make sure that the CoreXL Firewall instances run on other processing CPU cores.

Usually, we do not recommend for a CoreXL SND and a CoreXL Firewall instance to use the same CPU core. It is necessary for the CoreXL SND and a CoreXL Firewall instance to use a CPU core when Security Gateway (Scalable Platform Security Group) runs on a platform with only two CPU cores.

Adding Processing CPU Cores to the Hardware

If you increase the number of processing CPU cores on the computer, it does **not** automatically increase the number of CoreXL Firewall instances.

You must manually configure the applicable number of CoreXL Firewall instances in the `cpconfig` menu (see *"Configuring IPv4 and IPv6 CoreXL Firewall instances" on page 319*).

Important Notes for Cluster:

- You must configure the CoreXL in the same way on all the cluster members.
- If you enable CoreXL, disable CoreXL, or change the number of CoreXL Firewall instances, you should treat this change as a version upgrade. Schedule a full maintenance window and follow the instructions in the *R82 Installation and Upgrade Guide* - Chapter *Upgrading ClusterXL Deployments*. Perform either a *Minimum Effort Upgrade* procedure (requires downtime), or a *Zero Downtime Upgrade* procedure (no downtime, but active connections are lost). Instead of the version upgrade, configure the CoreXL on each cluster member.

Allocating Additional CPU Cores to the CoreXL SND

The default configuration of CoreXL Firewall instances and the CoreXL SND instances might not be optimal for your needs.

If the default number of CoreXL SND instances is not enough to process the incoming traffic, and your Security Gateway has enough CPU cores, you can decrease the number of CoreXL Firewall instances. This automatically allocates additional CPU cores to run the CoreXL SND instances.

This scenario is likely to occur if much of the traffic is accelerated by SecureXL. In this case, the task load of the CoreXL SND instances may be disproportionate to that of the CoreXL Firewall instances.

To check if the SND is slowing down the traffic:

Step	Instructions
1	Identify the processing CPU core, to which the interfaces direct their traffic. ■ On a Security Appliance, run in Gaia Clish or the Expert mode: `fw ctl affinity -l -r` ■ On a Scalable Platform Security Group, run in Gaia gClish: `fw ctl affinity -l -r` ■ On a Scalable Platform Security Group, run in the Expert mode: `g_fw ctl affinity -l -r`
2	Under heavy traffic conditions, run the `top` command. Examine the values for the different CPU cores in the `idle` column. ■ On a Security Appliance, run in the Expert mode: `top` ■ On a Scalable Platform Security Group, run in the Expert mode: `g_top`

Best Practice - We recommend to allocate an additional CPU core to the CoreXL SND only if **all** these conditions are met:

- There are at least 8 processing CPU cores.
- In the output of the `top` command, the `idle` values for the CPU cores that run the CoreXL SND instances are in the 0%-5% range.
- In the output of the `top` command, the sum of the `idle` values for the CPU cores that run the CoreXL Firewall instances is significantly higher than 100%.

If at least one of the above conditions is not met, the default CoreXL configuration is sufficient.

To allocate an additional processing CPU core to the CoreXL SND:

Item	Description
1	Decrease the number of CoreXL Firewall instances in the `cpconfig` menu. See *"Configuring IPv4 and IPv6 CoreXL Firewall instances" on page 319*.
2	Configure interface affinities to the remaining CPU cores. See *"Configuring Affinities for Interfaces" on page 338*.
3	Reboot to apply the new configuration.

Allocating a CPU Core for Heavy Logging

If the Security Gateway generates very large number of logs, it may be advisable to allocate a processing CPU core to the **fwd** daemon, which generates the logs.

Note - This change decreases the number of CPU cores available for CoreXL Firewall instances.

Important Notes for Cluster:

- You must configure the CoreXL in the same way on all the cluster members.
- If you enable CoreXL, disable CoreXL, or change the number of CoreXL Firewall instances, you should treat this change as a version upgrade. Schedule a full maintenance window and follow the instructions in the *R82 Installation and Upgrade Guide* - Chapter *Upgrading ClusterXL Deployments*. Perform either a *Minimum Effort Upgrade* procedure (requires downtime), or a *Zero Downtime Upgrade* procedure (no downtime, but active connections are lost). Instead of the version upgrade, configure the CoreXL on each cluster member.

To allocate a processing CPU core to the *fwd* daemon:

See *"Configuring Affinity Settings" on page 326*.

Step	Instructions
1	Connect to the command line on the Security Gateway (each Cluster Member). **Note** - On Scalable Platforms, you must connect to the Gaia Portal of the applicable Security Group.
2	Log in to the Expert mode.
3	Run: `cpconfig`
4	Enter the number of the **Check Point CoreXL** option.
5	Decrease the number of CoreXL Firewall instances. See *"Configuring IPv4 and IPv6 CoreXL Firewall instances" on page 319*.
6	Exit from the `cpconfig` menu.

Step	Instructions
7	Examine which processing CPU cores run the CoreXL Firewall instances and which CPU cores handle the traffic from interfaces. ■ On the Security Gateway (each Cluster Member), run: `fw ctl affinity -l -r` ■ On the Scalable Platform Security Group, run: `g_fw ctl affinity -l -r` See *"fw ctl affinity" on page 396*.
8	Back up the `$FWDIR/conf/fwaffinity.conf` file. ■ On the Security Gateway (each Cluster Member), run: `cp -v $FWDIR/conf/fwaffinity.conf{,_BKP}` ■ On the Scalable Platform Security Group, run: `g_cp -v $FWDIR/conf/fwaffinity.conf{,_BKP}`
9	Edit the `$FWDIR/conf/fwaffinity.conf` file. The same syntax applies to the Security Gateway (each Cluster Member) and the Scalable Platform Security Group: `vi $FWDIR/conf/fwaffinity.conf`
10	Allocate one of the remaining CPU cores to the **fwd** daemon. To do so, configure the affinity of the **fwd** daemon to that CPU core. `n fwd <CPU ID>` For example, to affine the **fwd** daemon to CPU core **#2**, add this line: `n fwd 2` **Note** - It is important to avoid the CPU cores that run the CoreXL SND instances only if these CPU cores are explicitly defined for the affinities for interfaces. If affinity of interfaces is configured in the Automatic mode, the **fwd** daemon can use all CPU cores that do not run CoreXL Firewall instances. Traffic from interfaces is automatically diverted to other CPU cores.
11	Save the changes in the file and exit the editor.
12	On the Scalable Platform Security Group, copy the `$FWDIR/conf/fwaffinity.conf` configuration file to all other Security Group Members: `asg_cp2blades $FWDIR/conf/fwaffinity.conf`

Step	Instructions
13	Load the new configuration. ■ To load it immediately: • On the Security Gateway (each Cluster Member), run: `$FWDIR/scripts/fwaffinity_apply` • On the Scalable Platform Security Group, run: `g_all $FWDIR/scripts/fwaffinity_apply` ■ To load it later, reboot. • On the Security Gateway (each Cluster Member), run: `reboot` • On the Scalable Platform Security Group, run: `g_reboot -a`

Configuring Affinities for Interfaces

The association of a specific interface with a specific processing CPU core is called the interface's *affinity* with that CPU core. This affinity causes the interface's traffic to be directed to that CPU core and the CoreXL SND to run on that CPU core.

Security Gateway loads (Scalable Platform Security Group Members load) affinities for interfaces during the boot from the CoreXL configuration file `$FWDIR/conf/fwaffinity.conf`. In this configuration file, lines that begin with the letter "`i`", define the affinities for interfaces.

Workflow:

Step	Instructions
1	Check which processing CPU cores run the CoreXL Firewall instances and which CPU cores handle the traffic from interfaces: ■ On a Security Gateway (each Cluster Member), run in Gaia Clish or the Expert mode: `fw ctl affinity -l -r` ■ On a Scalable Platform Security Group, run in Gaia gClish: `fw ctl affinity -l -r` ■ On a Scalable Platform Security Group, run in the Expert mode: `g_fw ctl affinity -l -r` See *"fw ctl affinity" on page 396*.
2	Allocate the remaining CPU cores to run the CoreXL SND instances. To do so, configure the affinity of interfaces to the applicable CPU cores. For more information, see *"Allocation of Processing CPU Cores" on page 331*. **Notes:** ■ To set the affinity of VLAN interfaces, use their physical interfaces. ■ If you allocate more than one processing CPU core to the CoreXL SND, it is necessary to configure affinities for interfaces explicitly to the remaining CPU cores. If you have multiple interfaces, decide which interfaces to affine to which CPU cores. Try to achieve a balance of expected traffic between the CPU cores. Examine the traffic balance with the `top` command.

Configuring affinities for interfaces explicitly:

Step	Instructions
1	Connect to the command line on the Security Gateway (each Cluster Member / Scalable Platform Security Group).
2	Log in to the Expert mode.
3	Configure the affinity of each interface in the `$FWDIR/conf/fwaffinity.conf` file. See *"Configuring Affinity Settings" on page 326*. For each interface, there must be a separate line that begins with the letter "`i`". Each of these lines must have this syntax: `i <Name of Interface> <CPU ID>` For example, if it is necessary that the traffic from `eth0` and `eth1` (`eth1-05` and `eth1-07`) goes to CPU core `#0`, and the traffic from `eth2` (`eth1-09`) goes to CPU core `#1`, add these lines: ■ On the Security Gateway (each Cluster Member): `i eth0 0` `i eth1 0` `i eth2 1` ■ On the Scalable Platform Security Group: `i eth1-05 0` `i eth1-07 0` `i eth1-09 1`

Step	Instructions
	Alternatively, you can choose to configure affinities for interface explicitly for only one processing CPU core, and define other CPU cores as the default affinity of the remaining interfaces. `i default <CPU ID>` For example, if it is necessary that the traffic from `eth2` (`eth1-05`) goes to CPU core `#1`, and the traffic from all other interfaces goes to CPU core `#0`, add these lines: ■ On the Security Gateway (each Cluster Member): `i eth2 1` `i default 0` ■ On the Scalable Platform Security Group: `i eth1-05 1` `i default 0`
4	Load the new configuration. ■ To load it immediately: • On the Security Gateway (each Cluster Member), run: `$FWDIR/scripts/fwaffinity_apply` • On the Scalable Platform Security Group, run: `g_all $FWDIR/scripts/fwaffinity_apply` ■ To load it later, reboot. • On the Security Gateway (each Cluster Member), run: `reboot` • On the Scalable Platform Security Group, run: `g_reboot -a`

Best Practice - If you allocate only one CPU core to the CoreXL SND, it is best to have that CPU core selected automatically. To do so, leave the default automatic interface affinity and do not configure explicit affinities for interfaces to CPU cores. Make sure the `$FWDIR/conf/fwaffinity.conf` file contains this line:

```
i default auto
```

Make sure that the `$FWDIR/conf/fwaffinity.conf` file does not contain other lines that begin with "`i`", so that there are no explicit affinities for interfaces configured. This makes sure that Security Gateway directs (Scalable Platform Security Group Members direct) all traffic to the remaining CPU cores.

Best Practice - In addition, see *"Multi-Queue" on page 436*.

Dynamic Balancing of CoreXL Instances

Introduction

On Check Point Appliances, R80.40 added the ability to change the number of CoreXL Firewall and SND instances without reboot (Dynamic Balancing).

Important:

- By default, this feature is *enabled*.
- We do **not** recommend manual configuration of CoreXL Firewall and SND instances, because such configuration *disables* the CoreXL Dynamic Balancing.
  To enable the CoreXL Dynamic Balancing again, you must disable it and enable it.
- For CoreXL Dynamic Balancing requirements, see sk164155.

When CoreXL Dynamic Balancing is enabled, Security Gateways / Cluster Members / Scalable Platform Security Group Members monitor the average CPU utilization of CoreXL Firewall and SND instances and automatically increases or decreases the number of CoreXL Firewall instances.

The Dynamic Balancing Daemon (*dsd*) has three stages in each iteration:

1. Examine the current CPU utilization.
2. Decide if and what changes to make based on the current CPU utilization.
3. If needed, change the current CoreXL configuration in one of these ways:
   - Add a CoreXL Firewall instance.

     This change is possible only under these conditions:

     a. Average difference in CPU utilization between CoreXL Firewall and SND instances is greater than 10%.

     b. The current number of CoreXL Firewall instances is less than it was during the boot.

- Add a CoreXL SND instance.

  This change stops a CoreXL Firewall instance and moves it to another CPU core.

  This change is possible only under these conditions:

  a. Average difference in CPU utilization between CoreXL Firewall and SND instances is greater than 10%.

  b. CoreXL Firewall instances consume the CPU cores at less than 40%.

  c. There is an available CPU core.

Syntax

Important:

- There are commands for Gaia Clish (Gaia gClish) and for the Expert mode.
- In a Cluster, you must configure all the Cluster Members in the same way.

Enabling the feature

This command enables the CoreXL Dynamic Balancing.

Shell	Security Gateway (each Cluster Member)	Security Group on a Scalable Platform
Gaia Clish	`set dynamic-balancing state enable` `reboot`	N / A
Gaia gClish	N / A	`set dynamic-balancing state enable` `reboot`
Expert mode	`dynamic_balancing -o enable` `reboot`	`g_dynamic_balancing -o enable` `g_reboot -a` **Best Practice** - To keep the Security Group active, we recommended to reboot the Security Group Members gradually.

Important:

- After you enable this feature for the first time, the Security Gateway (Scalable Platform Security Group) may require a reboot in these cases:
  - The current CoreXL configuration is not the default
  - More CoreXL SND instances are required for the current CPU load
- After the boot, you can stop, start, and restart this feature without a reboot.

Stopping the feature

This command stops the CoreXL Dynamic Balancing ("freezes" it).

Shell	Security Gateway (each Cluster Member)	Security Group on a Scalable Platform
Gaia Clish	`set dynamic-balancing state stop`	N / A
Gaia gClish	N / A	`set dynamic-balancing state stop`
Expert mode	`dynamic_balancing -o stop`	`g_dynamic_balancing -o stop`

Important:

- When you stop this feature, the Security Gateway (Scalable Platform Security Group) uses the last CoreXL Balancing configuration.
- This change does **not** require a reboot.
- This change survives the reboot.
- The status of the CoreXL Dynamic Balancing appears as "off".

Starting the feature

This command starts the CoreXL Dynamic Balancing after it was stopped.

Shell	Security Gateway (each Cluster Member)	Security Group on a Scalable Platform
Gaia Clish	`set dynamic-balancing state start`	N / A
Gaia gClish	N / A	`set dynamic-balancing state start`
Expert mode	`dynamic_balancing -o start`	`g_dynamic_balancing -o start`

Important:

- When you start this feature, the Security Gateway (Scalable Platform Security Group) continues to change the CoreXL Balancing configuration automatically based on the CPU utilization.
- This change does **not** require a reboot.
- This change survives the reboot.

Resetting the feature

This command resets the CoreXL configuration to the default and keep the CoreXL Dynamic Balancing enabled.

This command is equivalent to the "`disable`" command followed by the "`enable`" command.

Shell	Security Gateway (each Cluster Member)	Security Group on a Scalable Platform
Gaia Clish	`set dynamic-balancing state reset`	N / A
Gaia gClish	N / A	`set dynamic-balancing state reset`
Expert mode	`dynamic_balancing -r`	`g_dynamic_balancing -r`

Important:

- After this feature restarts, the CoreXL configuration returns to the default (see *"Default Configuration of CoreXL" on page 317*).
- This change does **not** require a reboot.

Disabling the feature

This command disables the CoreXL Dynamic Balancing.

Shell	Security Gateway (each Cluster Member)	Security Group on a Scalable Platform
Gaia Clish	`set dynamic-balancing state disable` `reboot`	N / A
Gaia gClish	N / A	`set dynamic-balancing state disable` `reboot`
Expert mode	`dynamic_balancing -o disable` `reboot`	`g_dynamic_balancing -o disable` `g_reboot -a` **Best Practice** - To keep the Security Group active, we recommended to reboot the Security Group Members gradually.

Important:

- When you disable this feature, the CoreXL configuration returns to the default (see *"Default Configuration of CoreXL" on page 317*).
- After you disable this feature, the Security Gateway (Scalable Platform Security Group) requires a reboot.
  The command shows the applicable message.

Monitoring

- You can monitor the *status* of the CoreXL Dynamic Balancing with CLI commands:

Shell	Security Gateway (each Cluster Member)	Security Group on a Scalable Platform
Gaia Clish	`show dynamic-balancing state`	N / A
Gaia gClish	N / A	`show dynamic-balancing state`
Expert mode	`dynamic_balancing -p`	`g_dynamic_balancing -p`

- You can monitor the *status* of the CoreXL Dynamic Balancing in the CPView tool:

  Procedure

  1. Connect to the command line on the Security Gateway (Scalable Platform Security Group).
  2. On the Scalable Platform Security Group, go to Gaia gClish or log in to the Expert mode.
  3. Run:

     ```
     cpview
     ```

  4. From the top, click:

     SysInfo

  5. Examine this field:

     DS Status

     - **On** - Means the CoreXL Dynamic Balancing is enabled
     - **Off** - Means the CoreXL Dynamic Balancing is disabled

- You can monitor the *performance* of the CoreXL Dynamic Balancing in the CPView tool:

  Procedure

  1. Connect to the command line on the Security Gateway (Scalable Platform Security Group).
  2. On the Scalable Platform Security Group, go to Gaia gClish or log in to the Expert mode.
  3. Run:

     ```
     cpview
     ```

  4. From the top, click:

     CPU > **Overview** > **Host**

  5. Examine these sections:

     - **Overview** - Shows the current number of CoreXL instances and the average CPU utilization
     - **CPU** - Shows the CPU cores, the CoreXL instance types they run, and the CPU utilization in different categories

- You can monitor the CoreXL Firewall instances with this command:

Shell	Security Gateway (each Cluster Member)	Security Group on a Scalable Platform
Gaia Clish	`fw ctl multik stat`	N / A
Gaia gClish	N / A	`fw ctl multik stat`
Expert mode	`fw ctl multik stat`	`g_fw ctl multik stat`

- You can monitor the CoreXL Affinity with this command:

Shell	Security Gateway (each Cluster Member)	Security Group on a Scalable Platform
Gaia Clish	`fw ctl affinity -l -r -a`	N / A
Gaia gClish	N / A	`fw ctl affinity -l -r -a`
Expert mode	`fw ctl affinity -l -r -a`	`g_fw ctl affinity -l -r -a`

- You can examine these log files:
  - When the CoreXL Dynamic Balancing changes the CoreXL configuration, it writes the applicable entries in the `$FWDIR/log/dsd.elg` file.
  - When the CoreXL Dynamic Balancing starts, it writes the applicable entries in the `$FWDIR/log/dynamic_split.elg` file.

CoreXL Firewall Mode - User Space or Kernel Space

Kernel Space Firewall (KSFW) is the infrastructure in which CoreXL Firewall instances run in the kernel.

User Space Firewall (USFW) is the infrastructure in which CoreXL Firewall instances run in the user space. This mode is available from R80.30 with Gaia kernel 3.10.

Important - For the complete information about the User Space Firewall (USFW) mode, see sk167052.

To change the Firewall Mode:

Step	Instructions
1	Connect to the command line on the Security Gateway / each Cluster Member. **Note** - On Scalable Platforms, you must connect to the Gaia Portal of the applicable Security Group.
2	Log in to Gaia Clish or Expert mode. **Note** - On Scalable Platforms, you must run the applicable commands in Gaia gClish or the Expert mode of the applicable Security Group.
3	Run: `cpconfig`
4	Enter the number of the **Check Point CoreXL** option.
5	Enter **3** to select **Change firewall mode**.
6	Follow the instructions on the screen.
7	Exit from the `cpconfig` menu.
8	Reboot. ■ On the Security Gateway (each Cluster Member), run: `reboot` ■ On the Scalable Platform Security Group, run in Gaia gClish: `reboot` ■ On the Scalable Platform Security Group, run in the Expert mode: `g_reboot -a`

CoreXL Commands

This section describes different CLI commands CoreXL.

Syntax Legend for CLI Commands

Whenever possible, this guide lists commands, parameters and options in the alphabetical order.

This guide uses this convention in the Command Line Interface (CLI) syntax:

Character	Description
TAB	Shows the available nested subcommands: `main command` `→ nested subcommand 1` `→ → nested subsubcommand 1-1` `→ → nested subsubcommand 1-2` `→ nested subcommand 2` Example: `cpwd_admin` `config` `-a <options>` `-d <options>` `-p` `-r` `del <options>` Meaning, you can run only **one** of these commands: ■ This command: `cpwd_admin config -a <options>` ■ Or this command: `cpwd_admin config -d <options>` ■ Or this command: `cpwd_admin config -p` ■ Or this command: `cpwd_admin config -r` ■ Or this command: `cpwd_admin del <options>`
Curly brackets or braces **{ }**	Enclose a list of available commands or parameters, separated by the vertical bar **|**. User can enter only one of the available commands or parameters.

Character	Description
Angle brackets **< >**	Enclose a variable. User must explicitly specify a supported value.
Square brackets or brackets **[ ]**	Enclose an optional command or parameter, which user can also enter.

cp_conf corexl

Description

Enables or disables CoreXL.

Important:

- This command is for Check Point use only.
  To configure CoreXL, use the **Check Point CoreXL** option in the *"cpconfig" on page 358* menu.
- After all changes in CoreXL configuration on the Security Gateway / Cluster Member / Security Group, you must reboot it.
- In a Cluster, you must configure all the Cluster Members in the same way.
- On Scalable Platforms, you must connect to the applicable Security Group.

Syntax on a Security Gateway / Cluster Member in Gaia Clish or the Expert mode

- To enable CoreXL with 'n' IPv4 Firewall instances and optionally 'k' IPv6 Firewall instances:

```
cp_conf corexl [-v] enable [n] [-6 k]
```

- To disable CoreXL:

```
cp_conf corexl [-v] disable
```

Syntax on a Scalable Platform Security Group in Gaia gClish

- To enable CoreXL with 'n' IPv4 Firewall instances and optionally 'k' IPv6 Firewall instances:

```
cp_conf corexl [-v] enable [n] [-6 k]
```

- To disable CoreXL:

```
cp_conf corexl [-v] disable
```

Syntax on a Scalable Platform Security Group in the Expert mode

- To enable CoreXL with 'n' IPv4 Firewall instances and optionally 'k' IPv6 Firewall instances:

```
g_all cp_conf corexl [-v] enable [n] [-6 k]
```

- To disable CoreXL:

```
g_all cp_conf corexl [-v] disable
```

The related command is: *"fwboot corexl" on page 418*.

Parameters

Parameter	Description
-v	Leaves the high memory (vmalloc) unchanged.
n	Denotes the number of IPv4 CoreXL Firewall instances.
k	Denotes the number of IPv6 CoreXL Firewall instances.

Example

Currently, the Security Gateway runs two IP4v CoreXL Firewall instances (`KERN_INSTANCE_NUM = 2`).

We change the number of IP4v CoreXL Firewall instances to three.

```
[Expert@MyGW:0]# fw ctl multik stat
ID | Active  | CPU    | Connections | Peak
----------------------------------------------
 0 | Yes     | 2      |           7 |       28
 1 | Yes     | 1      |           0 |       11
[Expert@MyGW:0]#
[Expert@MyGW:0]# cat /etc/fw.boot/boot.conf
CTL_IPFORWARDING 1
DEFAULT_FILTER_PATH 0
KERN_INSTANCE_NUM 2
COREXL_INSTALLED 1
KERN6_INSTANCE_NUM 2
IPV6_INSTALLED 0
CORE_OVERRIDE 4
[Expert@MyGW:0]#
[Expert@MyGW:0]# cp_conf corexl -v enable 3
[Expert@MyGW:0]#
[Expert@MyGW:0]# cat /etc/fw.boot/boot.conf
CTL_IPFORWARDING 1
DEFAULT_FILTER_PATH 0
KERN_INSTANCE_NUM 3
COREXL_INSTALLED 1
KERN6_INSTANCE_NUM 2
IPV6_INSTALLED 0
CORE_OVERRIDE 4
[Expert@MyGW:0]#
[Expert@MyGW:0]# reboot
.. ... ...
[Expert@MyGW:0]# fw ctl multik stat
ID | Active  | CPU    | Connections | Peak
----------------------------------------------
 0 | Yes     | 3      |           7 |       28
 1 | Yes     | 2      |           0 |       11
 2 | Yes     | 1      |           4 |       10
[Expert@MyGW:0]#
```

cpconfig

Description

This command starts the Check Point Configuration Tool.

This tool configures specific settings for the installed Check Point products.

Important:

- In a Cluster, you must configure all the Cluster Members in the same way.
- On Scalable Platforms, you must connect to the applicable Security Group.

Syntax on a Security Gateway / Cluster Member in Gaia Clish or the Expert mode

```
cpconfig
```

Syntax on a Scalable Platform Security Group in Gaia gClish or the Expert mode

```
cpconfig
```

Menu Options

Note - The options shown depend on the configuration and installed products.

Menu Option	Description
Licenses and contracts	Manages Check Point licenses and contracts on this Security Gateway or Cluster Member.
SNMP Extension	Obsolete. Do **not** use this option anymore. To configure SNMP, see the *R82 Gaia Administration Guide* - Chapter *System Management* - Section *SNMP*.
PKCS#11 Token	Register a cryptographic token, for use by Gaia Operating System. See details of the token, and test its functionality.
Random Pool	Configures the RSA keys, to be used by Gaia Operating System.

Menu Option	Description
Secure Internal Communication	Manages SIC on the Security Gateway or Cluster Member. This change requires a restart of Check Point services on the Security Gateway or Cluster Member. For more information, see: ■ The *R82 Security Management Administration Guide*. ■ sk65764: How to reset SIC.
Enable cluster membership for this gateway	Enables the cluster membership on the Security Gateway. This change requires a reboot of the Security Gateway. For more information, see the: ■ *R82 Installation and Upgrade Guide*. ■ *R82 ClusterXL Administration Guide*. **Note** - This section does **not** apply to Scalable Platforms (ElasticXL, Maestro, and Chassis).
Disable cluster membership for this gateway	Disables the cluster membership on the Security Gateway. This change requires a reboot of the Security Gateway. For more information, see the: ■ *R82 Installation and Upgrade Guide*. ■ *R82 ClusterXL Administration Guide*. **Note** - This section does **not** apply to Scalable Platforms (ElasticXL, Maestro, and Chassis).
Enable Check Point Per Virtual System State	Enables Virtual System Load Sharing on the VSX Cluster Member. For more information, see the *R82 VSX Administration Guide*. **Note** - This section does **not** apply to Scalable Platforms (ElasticXL, Maestro, and Chassis).
Disable Check Point Per Virtual System State	Disables Virtual System Load Sharing on the VSX Cluster Member. For more information, see the *R82 VSX Administration Guide*. **Note** - This section does **not** apply to Scalable Platforms (ElasticXL, Maestro, and Chassis).

Menu Option	Description
Enable Check Point ClusterXL for Bridge Active/Standby	Enables Check Point ClusterXL for Bridge mode. This change requires a reboot of the Cluster Member. For more information, see the: ■ *R82 Installation and Upgrade Guide*. ■ *R82 ClusterXL Administration Guide*. **Note** - This section does **not** apply to Scalable Platforms (ElasticXL, Maestro, and Chassis).
Disable Check Point ClusterXL for Bridge Active/Standby	Disables Check Point ClusterXL for Bridge mode. This change requires a reboot of the Cluster Member. For more information, see the: ■ *R82 Installation and Upgrade Guide*. ■ *R82 ClusterXL Administration Guide*. **Note** - This section does **not** apply to Scalable Platforms (ElasticXL, Maestro, and Chassis).
Check Point CoreXL	Manages CoreXL and Firewall mode on the Security Gateway / Cluster Member / Scalable Platform Security Group. After all changes in CoreXL configuration, you must reboot the Security Gateway / Cluster Member / Security Group. For more information, see *"CoreXL" on page 314*.
Automatic start of Check Point Products	Shows and controls which of the installed Check Point products start automatically during boot.
Exit	Exits from the Check Point Configuration Tool.

Example 1 - Menu on a single Security Gateway

```
[Expert@MySingleGW:0]# cpconfig
This program will let you re-configure
your Check Point products configuration.


Configuration Options:
----------------------
(1) Licenses and contracts
(2) SNMP Extension
(3) PKCS#11 Token
(4) Random Pool
(5) Secure Internal Communication
(6) Enable cluster membership for this gateway
(7) Check Point CoreXL
(8) Automatic start of Check Point Products

(9) Exit

Enter your choice (1-9) :
```

Example 2 - Menu on a Cluster Member

```
[Expert@MyClusterMember:0]# cpconfig
This program will let you re-configure
your Check Point products configuration.


Configuration Options:
----------------------
(1) Licenses and contracts
(2) SNMP Extension
(3) PKCS#11 Token
(4) Random Pool
(5) Secure Internal Communication
(6) Disable cluster membership for this gateway
(7) Enable Check Point Per Virtual System State
(8) Enable Check Point ClusterXL for Bridge Active/Standby
(9) Check Point CoreXL
(10) Automatic start of Check Point Products

(11) Exit

Enter your choice (1-11) :
```

cpview

Overview of CPView

Description

CPView is a text based built-in utility on a Check Point computer.

CPView Utility shows statistical data that contain both general system information (CPU, Memory, Disk space) and information for different Software Blades (only on a Security Gateway / ClusterXL / Scalable Platform Security Group).

The CPView continuously updates the data in easy to access views.

On a Security Gateway / ClusterXL / Scalable Platform Security Group, you can use this statistical data to monitor the performance.

For more information, see sk101878.

Syntax

```
cpview --help
```

CPView User Interface

The CPView user interface has three sections:

Section	Description
Header	This view shows the time the statistics in the third view are collected. It updates when you refresh the statistics.
Navigation	This menu bar is interactive. Move between menus with the arrow keys and mouse. A menu can have sub-menus and they show under the menu bar.
View	This view shows the statistics collected in that view. These statistics update at the refresh rate.

Using CPView

Use these keys to navigate the CPView:

Key	Description
Arrow keys	Moves between menus and views. Scrolls in a view.
Home	Returns to the **Overview** view.
Enter	Changes to the **View Mode**. On a menu with sub-menus, the **Enter** key moves you to the lowest level sub-menu.
Esc	Returns to the **Menu Mode**.
Q	Quits CPView.

Use these keys to change CPView interface options:

Key	Description
R	Opens a window where you can change the refresh rate. The default refresh rate is 2 seconds.
W	Changes between wide and normal display modes. In wide mode, CPView fits the screen horizontally.
S	Manually sets the number of rows or columns.
M	Switches on/off the mouse.
P	Pauses and resumes the collection of statistics.

Use these keys to save statistics, show help, and refresh statistics:

Key	Description
C	Saves the current page to a file. The file name format is: `cpview_<ID of the cpview process>.cap<Number of the capture>`
H	Shows a tooltip with CPView options.
Space bar	Immediately refreshes the statistics.

dynamic_balancing

Description

On Check Point Appliances, R80.40 added the ability to change the number of CoreXL Firewall and SND instances without reboot (Dynamic Balancing).

Important:

- By default, this feature is *enabled*.
- We do **not** recommend manual configuration of CoreXL Firewall and SND instances, because such configuration *disables* the CoreXL Dynamic Balancing.
  To enable the CoreXL Dynamic Balancing again, you must disable it and enable it.
- For CoreXL Dynamic Balancing requirements, see sk164155.

The "`dynamic_balancing`" command in the Expert mode (and the command "`set dynamic-balancing state`" in Gaia Clish) controls the Dynamic Balancing of CoreXL Firewall and SND instances on the local Security Gateway, or Cluster Member.

For more information, see *"Dynamic Balancing of CoreXL Instances" on page 341*.

Important:

- In a Cluster, you must configure all the Cluster Members in the same way.
- On Scalable Platforms, you must connect to the applicable Security Group.

Syntax

Shell	Security Gateway (each Cluster Member)	Security Group on a Scalable Platform
Gaia Clish	`set dynamic-balancing state` `disable` `enable` `reset` `start` `stop` `show dynamic-balancing state`	N / A
Gaia gClish	N / A	`set dynamic-balancing state` `disable` `enable` `reset` `start` `stop` `show dynamic-balancing state`
Expert mode	`dynamic_balancing` `-o disable` `-o enable` `-o start` `-o stop` `-p` `-r`	`g_dynamic_balancing` `-o disable` `-o enable` `-o start` `-o stop` `-p` `-r`

Parameters

Parameter	Description
No Parameters	Shows the applicable built-in help.

Parameter	Description
`disable`	Disables the CoreXL Dynamic Balancing. **Important:** ■ When you disable this feature, the CoreXL configuration returns to the default. ■ After you disable this feature, the Security Gateway (Scalable Platform Security Group) **requires** a reboot. The command shows the applicable message. **Best Practice** - To keep the Scalable Platform Security Group active, we recommended to reboot the Security Group Members gradually.
`enable`	Enables the CoreXL Dynamic Balancing. **Important:** ■ After you enable this feature, the Security Gateway (Scalable Platform Security Group) **requires** a reboot. The command shows the applicable message. ■ After you enable this feature for the first time, the Security Gateway (Scalable Platform Security Group) may require a reboot in these cases: • The current CoreXL configuration is not the default • More CoreXL SND instances are required for the current CPU load ■ After the boot, you can stop, start, and this feature without a reboot. **Best Practice** - To keep the Scalable Platform Security Group active, we recommended to reboot the Security Group Members gradually.
`reset` or `-r`	Resets the CoreXL configuration to the default and keeps the CoreXL Dynamic Balancing enabled. This command is equivalent to the "`disable`" command followed by the "`enable`" command. **Important:** ■ After this feature resets, the CoreXL configuration returns to the default. ■ This change does **not** require a reboot.

Parameter	Description
`start`	Starts the CoreXL Dynamic Balancing after it was stopped. **Important:** ■ When you start this feature, the Security Gateway (Scalable Platform Security Group) continues to change the CoreXL Balancing configuration automatically based on the CPU utilization. ■ This change does **not** require a reboot. ■ This change survives the reboot. ■ The status of the CoreXL Dynamic Balancing appears as "off".
`stop`	Stops the CoreXL Dynamic Balancing. **Important:** ■ When you stop this feature, the Security Gateway (Scalable Platform Security Group) uses the last CoreXL Balancing configuration. ■ This change does **not** require a reboot. ■ This change survives the reboot.
`show dynamic-balancing state` or `-p`	Shows the current state of the CoreXL Dynamic Balancing (enabled, disabled, started, or stopped).

Example

```
[Expert@MyGW:0]# dynamic_balancing -p
Dynamic Balancing is currently On
[Expert@MyGW:0]#
```

fw ctl multik

Description

The "`fw ctl multik`" and "`fw6 ctl multik`" commands control CoreXL for IPv4 and IPv6, respectively.

Syntax for IPv4

```
fw ctl multik
      add_bypass_port <options>
      del_bypass_port <options>
      dynamic_dispatching <options>
      gconn <options>
      get_instance <options>
      heavy_conn_analyzer
      print_heavy_conn
      prioq <options>
      queues
      show_bypass_ports
      snd_dist <options>
      stat
      start
      stop
      utilize
```

Syntax for IPv6

```
fw6 ctl multik
      add_bypass_port <options>
      del_bypass_port <options>
      dynamic_dispatching <options>
      gconn <options>
      get_instance <options>
      heavy_conn_analyzer
      print_heavy_conn
      prioq <options>
      queues
      show_bypass_ports
      snd_dist <options>
      stat
      start
      stop
      utilize
```

Parameters

Parameter	Description
`add_bypass_port <options>`	Adds the specified TCP and UDP ports to the CoreXL Dynamic Dispatcher bypass list. See *"fw ctl multik add_bypass_port" on page 371*.
`del_bypass_port <options>`	Removes the specified TCP and UDP ports from the CoreXL Dynamic Dispatcher bypass list. See *"fw ctl multik del_bypass_port" on page 373*.
`dynamic_ dispatching <options>`	Shows and controls CoreXL Dynamic Dispatcher (see sk105261). See *"fw ctl multik dynamic_dispatching" on page 375*.
`gconn <options>`	Shows statistics about CoreXL Global Connections. See *"fw ctl multik gconn" on page 376*.
`get_instance <options>`	Shows CoreXL Firewall instance that processes the specified IPv4 connection. See *"fw ctl multik get_instance" on page 381*.
`heavy_conn_ analyzer print_heavy_conn`	Shows the table with Heavy Connections (that consume the most CPU resources) in the CoreXL Dynamic Dispatcher. See *"fw ctl multik print_heavy_conn" on page 383*.
`prioq <options>`	Configures the CoreXL Firewall Priority Queues (see sk105762). See *"fw ctl multik prioq" on page 385*.
`queues`	Shows the CoreXL Queues and their utilization. This command is supported only when the SecureXL works in the User Mode (UPPAK). See *"fw ctl multik queues" on page 386*.
`show_bypass_ ports`	Shows the TCP and UDP ports configured in the bypass port list of the CoreXL Dynamic Dispatcher. See *"fw ctl multik show_bypass_ports" on page 387*.
`snd_dist <options>`	Shows advanced information about distribution of connections by CoreXL SND instances. See *"fw ctl multik snd_dist" on page 388*.
`stat`	Shows the CoreXL status. See *"fw ctl multik stat" on page 391*.

Parameter	Description
`start`	Starts all CoreXL Firewall instances on-the-fly. See *"fw ctl multik start" on page 393*.
`stop`	Stops all CoreXL Firewall instances temporarily. See *"fw ctl multik stop" on page 394*.
`utilize`	Shows the CoreXL queue utilization for each CoreXL Firewall instance. See *"fw ctl multik utilize" on page 395*.

fw ctl multik add_bypass_port

Description

Adds the specified TCP and UDP ports to the bypass port list of the CoreXL Dynamic Dispatcher.

For more information about the CoreXL Dynamic Dispatcher, see sk105261.

Important - This command saves the configuration in the `$FWDIR/conf/dispatcher_bypass.conf` file. You must **not** edit this file manually.

Syntax

```
fw ctl multik add_bypass_port <Port Number 1>,<Port Number 2>,...,<Port Number N>
```

Parameters

Parameter	Description
<Port Number>	Specifies the numbers of TCP and UDP ports to add to the list. **Important** - You can add 10 ports maximum.

Example

```
[Expert@MyGW:0]# fw ctl multik show_bypass_ports
dynamic dispatcher bypass port list:
[Expert@MyGW:0]#
[Expert@MyGW:0]# cat $FWDIR/conf/dispatcher_bypass.conf
dynamic_dispatcher_bypass_ports_number = 0
[Expert@MyGW:0]#
[Expert@MyGW:0]# fw ctl multik add_bypass_port 8888
[Expert@MyGW:0]#
[Expert@MyGW:0]# fw ctl multik show_bypass_ports
dynamic dispatcher bypass port list:
(8888)
[Expert@MyGW:0]
[Expert@MyGW:0]# cat $FWDIR/conf/dispatcher_bypass.conf
dynamic_dispatcher_bypass_ports_number = 1
dynamic_dispatcher_bypass_port_table=8888
[Expert@MyGW:0]
[Expert@MyGW:0]# fw ctl multik add_bypass_port 9999
[Expert@MyGW:0]#
[Expert@MyGW:0]# fw ctl multik show_bypass_ports
dynamic dispatcher bypass port list:
(8888,9999)
[Expert@MyGW:0]
[Expert@MyGW:0]# cat $FWDIR/conf/dispatcher_bypass.conf
dynamic_dispatcher_bypass_ports_number = 2
dynamic_dispatcher_bypass_port_table=8888,9999
[Expert@MyGW:0]
```

fw ctl multik del_bypass_port

Description

Removes the specified TCP and UDP ports from the bypass port list of the CoreXL Dynamic Dispatcher.

For more information about the CoreXL Dynamic Dispatcher, see sk105261.

Important - This command saves the configuration in the `$FWDIR/conf/dispatcher_bypass.conf` file. You must **not** edit this file manually.

Syntax

```
fw ctl multik del_bypass_port <Port Number 1>,<Port Number 2>,...,<Port Number N>
```

Parameters

Parameter	Description
<Port Number>	Specifies the numbers of TCP and UDP ports to remove from the list.

Example

```
[Expert@MyGW:0]# fw ctl multik show_bypass_ports
dynamic dispatcher bypass port list:
[Expert@MyGW:0]#
[Expert@MyGW:0]# cat $FWDIR/conf/dispatcher_bypass.conf
dynamic_dispatcher_bypass_ports_number = 0
[Expert@MyGW:0]#
[Expert@MyGW:0]# fw ctl multik add_bypass_port 8888
[Expert@MyGW:0]#
[Expert@MyGW:0]# fw ctl multik show_bypass_ports
dynamic dispatcher bypass port list:
(8888)
[Expert@MyGW:0]
[Expert@MyGW:0]# cat $FWDIR/conf/dispatcher_bypass.conf
dynamic_dispatcher_bypass_ports_number = 1
dynamic_dispatcher_bypass_port_table=8888
[Expert@MyGW:0]
[Expert@MyGW:0]# fw ctl multik add_bypass_port 9999
[Expert@MyGW:0]#
[Expert@MyGW:0]# fw ctl multik show_bypass_ports
dynamic dispatcher bypass port list:
(8888,9999)
[Expert@MyGW:0]
[Expert@MyGW:0]# cat $FWDIR/conf/dispatcher_bypass.conf
dynamic_dispatcher_bypass_ports_number = 2
dynamic_dispatcher_bypass_port_table=8888,9999
[Expert@MyGW:0]
[Expert@MyGW:0]# fw ctl multik add_bypass_port 9999
[Expert@MyGW:0]#
[Expert@MyGW:0]# fw ctl multik show_bypass_ports
dynamic dispatcher bypass port list:
(8888)
[Expert@MyGW:0]
[Expert@MyGW:0]# cat $FWDIR/conf/dispatcher_bypass.conf
dynamic_dispatcher_bypass_ports_number = 1
dynamic_dispatcher_bypass_port_table=8888
[Expert@MyGW:0]
```

fw ctl multik dynamic_dispatching

Description

Shows and controls the CoreXL Dynamic Dispatcher that dynamically assigns new connections to a CoreXL Firewall instances based on the utilization of CPU cores.

For more information, see sk105261.

Syntax for IPv4

```
fw ctl multik dynamic_dispatching
      get_mode
      off
      on
```

Syntax for IPv6

```
fw6 ctl multik dynamic_dispatching
      get_mode
      off
      on
```

Parameters

Parameter	Description
`get_mode`	Shows the current state of the CoreXL Dynamic Dispatcher.
`off`	Disables the CoreXL Dynamic Dispatcher.
`on`	Enables the CoreXL Dynamic Dispatcher.

Example

```
[Expert@MyGW:0]# fw ctl multik dynamic_dispatching get_mode
Current mode is Off
[Expert@MyGW:0]#
[Expert@MyGW:0]# fw ctl multik dynamic_dispatching on
New mode is: On
Please reboot the system
[Expert@MyGW:0]#
```

fw ctl multik gconn

Description

Shows statistics about CoreXL Global Connections that Security Gateway stores in the kernel table `fw_multik_ld_gconn_table`.

The CoreXL Global Connections table contains information about which CoreXL Firewall instance owns which connections.

Notes:

- This command does not support VSX.
- This command does not support IPv6.

Syntax

```
fw [-d] ctl multik gconn
      -h
      -p
      -sec
      -seg <Number>
```

Parameters

Parameter	Description
`-d`	Runs the command in debug mode. Use only if you troubleshoot the command itself. **Best Practice** - If you use this parameter, then redirect the output to a file, or use the `script` command to save the entire CLI session.
none	Shows the interactive menu for the CoreXL Firewall Priority Queues.
`-h`	Shows the built-in help.

Parameter	Description
`-p`	Shows the additional information about each CoreXL Firewall instance, including the information about Firewall Priority Queues: ■ `I/O` (In or Out) ■ `Inst. ID` (CoreXL Firewall instance ID) ■ `Flags` ■ `Seq` (Sequence) ■ `Hold_ref` (Hold reference) ■ `Prio` (Firewall Priority Queues mode) ■ `last_enq_jiff` (Jiffies since last enqueue) ■ `queue_indx` (Queue index number) ■ `conn_tokens` (Connection Tokens)
`-s`	Shows the total number of global connections.
`-sec`	Shows the additional information about each CoreXL Firewall instance: ■ `I/O` (In or Out) ■ `Inst. ID` (CoreXL Firewall instance ID) ■ `Flags` ■ `Seq` (Sequence) ■ `Hold_ref` (Hold reference)
`-seg` *<Number>*	Shows the default information about the specified Global Connections Segment.

Example 1 - Default information

```
[Expert@MyGW:0]# fw ctl multik gconn
Default:

================================================================================================================
============================
| Segm | Src IP | S.port | Dst IP | D.port | Proto | Flags | PP |Ref Cnt(I/O)|Inst|PPAK ID|clstr
mem ID|Rec. ref|Rec. Type|

================================================================================================================
============================
|  0   | 192.168.3.52    | 18192 | 192.168.3.240   | 46082 | 6  |FP .. ..| No | 0/0 |  1   | 32 |
 0     |    0   | UNDEF |
|  0   | 192.168.3.52    | 54216 | 192.168.3.240   | 257   | 6  |FP .. ..| No | 0/0 |  1   | 32 |
 0     |    0   | UNDEF |
|  0   | 192.168.3.240   | 53925 | 192.168.3.53    | 18192 | 6  |FP .. ..| No | 0/0 |  0   | 32 |
 1     |    0   | UNDEF |
|  0   | 192.168.3.240   | 257   | 192.168.3.52    | 54216 | 6  |FP .. ..| No | 0/0 |  1   | 32 |
 0     |    0   | UNDEF |
|  0   | 192.168.3.53    | 18192 | 192.168.3.240   | 64216 | 6  |FP .. ..| No | 0/0 |  1   | 32 |
 15    |    0   | UNDEF |
|  0   | 0.0.0.0         | 8116  | 192.168.3.53    | 8116  | 17 |FP .. ..| No | 0/0 |  1   | 32 |
 1     |    0   | UNDEF |
|  0   | 0.0.0.0         | 8116  | 192.168.3.52    | 8116  | 17 |FP .. ..| No | 0/0 |  1   | 32 |
 0     |    0   | UNDEF |
|  0   | 192.168.3.240   | 64216 | 192.168.3.53    | 18192 | 6  |FP .. ..| No | 0/0 |  1   | 32 |
 15    |    0   | UNDEF |
|  0   | 192.168.3.52    | 8116  | 0.0.0.0         | 8116  | 17 |FP .. ..| No | 0/0 |  1   | 32 |
 0     |    0   | UNDEF |
|  0   | 172.20.168.16   | 63800 | 192.168.3.53    | 22    | 6  |FP .. ..| No | 0/0 |  0   | 32 |
 1     |    0   | UNDEF |
|  0   | 192.168.3.240   | 46082 | 192.168.3.52    | 18192 | 6  |FP .. ..| No | 0/0 |  1   | 32 |
 0     |    0   | UNDEF |
|  0   | 192.168.3.53    | 8116  | 0.0.0.0         | 8116  | 17 |FP .. ..| No | 0/0 |  1   | 32 |
 1     |    0   | UNDEF |
|  0   | 192.168.3.53    | 22    | 172.20.168.16   | 63800 | 6  |FP .. ..| No | 0/0 |  0   | 32 |
 1     |    0   | UNDEF |
|  0   | 192.168.3.53    | 18192 | 192.168.3.240   | 53925 | 6  |FP .. ..| No | 0/0 |  0   | 32 |
 1     |    0   | UNDEF |

================================================================================================================
============================
FP - from pool.    T - temporary connection.    PP - pending pernament.
[Expert@MyGW:0]#
```

Example 2 - Summary information only

```
[Expert@MyGW:0]# fw ctl multik gconn -s
Summary:
        Total number of global connections: 12
[Expert@MyGW:0]#
```

Example 3 - Additional information about each CoreXL Firewall instance, including the information about Firewall Priority Queues

```
[Expert@MyGW:0]# fw ctl multik gconn -p
Instance section prio info:

================================================================================================
================================================================================================
=======
| Segm | Src IP | S.port | Dst IP | D.port | Proto | Flags | PP |Ref Cnt(I/O)|Inst|PPAK ID|clstr
mem ID|Rec. ref|Rec. Type|Inst. Section: I/O|Inst. ID|Flags| Seq | Hold_ref |Prio:|last_enq_
jiff|queue_indx|conn_tokens

================================================================================================
================================================================================================
=======
|  0  | 192.168.3.52    | 18192 | 192.168.3.240   | 46082 | 6  |FP .. ..| No | 0/0 |  1  | 32 |
 0    |   0   | UNDEF |Inst. Section: Out |  1  | Perm |  0    |  0  |Prio:|  0  |  -1  |  0  |
|  0  | 192.168.3.240   | 53925 | 192.168.3.53    | 18192 | 6  |FP .. ..| No | 0/0 |  0  | 32 |
 1    |   0   | UNDEF |Inst. Section: In  |  0  | Perm |  0    |  0  |Prio:|  0  |  -1  |  0  |
|  0  | 192.168.3.240   | 257   | 192.168.3.52    | 35883 | 6  |FP .. ..| No | 0/0 |  1  | 32 |
 0    |   0   | UNDEF |Inst. Section: In  |  1  | Perm |  0    |  0  |Prio:|  0  |  -1  |  0  |
|  0  | 192.168.3.53    | 18192 | 192.168.3.240   | 64216 | 6  |FP .. ..| No | 0/0 |  1  | 32 |
 15   |   0   | UNDEF |Inst. Section: Out |  1  | Perm |  0    |  0  |Prio:|  0  |  -1  |  0  |
|  0  | 0.0.0.0         | 8116  | 192.168.3.53    | 8116  | 17 |FP .. ..| No | 0/0 |  1  | 32 |
 1    |   0   | UNDEF |Inst. Section: In  |  1  | Perm |  0    |  0  |Prio:|  0  |  -1  |  0  |
|  0  | 0.0.0.0         | 8116  | 192.168.3.52    | 8116  | 17 |FP .. ..| No | 0/0 |  1  | 32 |
 0    |   0   | UNDEF |Inst. Section: In  |  1  | Perm |  0    |  0  |Prio:|  0  |  -1  |  0  |
|  0  | 192.168.3.240   | 64216 | 192.168.3.53    | 18192 | 6  |FP .. ..| No | 0/0 |  1  | 32 |
 15   |   0   | UNDEF |Inst. Section: In  |  1  | Perm |  0    |  0  |Prio:|  0  |  -1  |  0  |
|  0  | 192.168.3.52    | 8116  | 0.0.0.0         | 8116  | 17 |FP .. ..| No | 0/0 |  1  | 32 |
 0    |   0   | UNDEF |Inst. Section: Out |  1  | Perm |  0    |  0  |Prio:|  0  |  -1  |  0  |
|  0  | 172.20.168.16   | 63800 | 192.168.3.53    | 22    | 6  |FP .. ..| No | 0/0 |  0  | 32 |
 1    |   0   | UNDEF |Inst. Section: In  |  0  | Perm |  494  |  0  |Prio:|  0  |  -1  |  0  |
|  0  | 192.168.3.240   | 46082 | 192.168.3.52    | 18192 | 6  |FP .. ..| No | 0/0 |  1  | 32 |
 0    |   0   | UNDEF |Inst. Section: In  |  1  | Perm |  0    |  0  |Prio:|  0  |  -1  |  0  |
|  0  | 192.168.3.52    | 35883 | 192.168.3.240   | 257   | 6  |FP .. ..| No | 0/0 |  1  | 32 |
 0    |   0   | UNDEF |Inst. Section: Out |  1  | Perm |  0    |  0  |Prio:|  0  |  -1  |  0  |
|  0  | 192.168.3.53    | 8116  | 0.0.0.0         | 8116  | 17 |FP .. ..| No | 0/0 |  1  | 32 |
 1    |   0   | UNDEF |Inst. Section: Out |  1  | Perm |  0    |  0  |Prio:|  0  |  -1  |  0  |
|  0  | 192.168.3.53    | 22    | 172.20.168.16   | 63800 | 6  |FP .. ..| No | 0/0 |  0  | 32 |
 1    |   0   | UNDEF |Inst. Section: Out |  0  | Perm |  280  |  0  |Prio:|  0  |  -1  |  0  |
|  0  | 192.168.3.53    | 18192 | 192.168.3.240   | 53925 | 6  |FP .. ..| No | 0/0 |  0  | 32 |
 1    |   0   | UNDEF |Inst. Section: Out |  0  | Perm |  219  |  0  |Prio:|  0  |  -1  |  0  |

================================================================================================
================================================================================================
=======
FP - from pool.    T - temporary connection.    PP - pending pernament.    In - inbound.    Out
- outbound.
[Expert@MyGW:0]#
```

Example 4 - Additional information about each CoreXL Firewall instance

```
[Expert@MyGW:0]# fw ctl multik gconn -sec
Instance section:

================================================================================================================
================================================================================
| Segm | Src IP | S.port | Dst IP | D.port | Proto | Flags | PP |Ref Cnt(I/O)|Inst|PPAK ID|clstr
mem ID|Rec. ref|Rec. Type|Inst. Section: I/O|Inst. ID|Flags| Seq | Hold_ref |

================================================================================================================
================================================================================
|  0  | 192.168.3.52    | 18192 | 192.168.3.240   | 46082 | 6  |FP .. ..| No | 0/0 |  1  | 32 |
 0    |   0   | UNDEF |Inst. Section: Out |  1  | Perm |  0    |  0  |
|  0  | 192.168.3.52    | 52864 | 192.168.3.240   | 257   | 6  |FP .. ..| No | 0/0 |  2  | 32 |
 0    |   0   | UNDEF |Inst. Section: Out |  2  | Perm |  0    |  0  |
|  0  | 192.168.3.240   | 53925 | 192.168.3.53    | 18192 | 6  |FP .. ..| No | 0/0 |  0  | 32 |
 1    |   0   | UNDEF |Inst. Section: In  |  0  | Perm |  0    |  0  |
|  0  | 192.168.3.53    | 18192 | 192.168.3.240   | 64216 | 6  |FP .. ..| No | 0/0 |  1  | 32 |
 15   |   0   | UNDEF |Inst. Section: Out |  1  | Perm |  0    |  0  |
|  0  | 192.168.3.53    | 60186 | 192.168.3.240   | 257   | 6  |FP .. ..| No | 0/0 |  1  | 32 |
 1    |   0   | UNDEF |Inst. Section: Out |  1  | Perm |  76   |  0  |
|  0  | 0.0.0.0         | 8116  | 192.168.3.53    | 8116  | 17 |FP .. ..| No | 0/0 |  1  | 32 |
 1    |   0   | UNDEF |Inst. Section: In  |  1  | Perm |  0    |  0  |
|  0  | 0.0.0.0         | 8116  | 192.168.3.52    | 8116  | 17 |FP .. ..| No | 0/0 |  1  | 32 |
 0    |   0   | UNDEF |Inst. Section: In  |  1  | Perm |  0    |  0  |
|  0  | 192.168.3.240   | 64216 | 192.168.3.53    | 18192 | 6  |FP .. ..| No | 0/0 |  1  | 32 |
 15   |   0   | UNDEF |Inst. Section: In  |  1  | Perm |  0    |  0  |
|  0  | 192.168.3.52    | 8116  | 0.0.0.0         | 8116  | 17 |FP .. ..| No | 0/0 |  1  | 32 |
 0    |   0   | UNDEF |Inst. Section: Out |  1  | Perm |  0    |  0  |
|  0  | 172.20.168.16   | 63800 | 192.168.3.53    | 22    | 6  |FP .. ..| No | 0/0 |  0  | 32 |
 1    |   0   | UNDEF |Inst. Section: In  |  0  | Perm |  479  |  0  |
|  0  | 192.168.3.240   | 46082 | 192.168.3.52    | 18192 | 6  |FP .. ..| No | 0/0 |  1  | 32 |
 0    |   0   | UNDEF |Inst. Section: In  |  1  | Perm |  0    |  0  |
|  0  | 192.168.3.53    | 8116  | 0.0.0.0         | 8116  | 17 |FP .. ..| No | 0/0 |  1  | 32 |
 1    |   0   | UNDEF |Inst. Section: Out |  1  | Perm |  0    |  0  |
|  0  | 192.168.3.240   | 257   | 192.168.3.52    | 52864 | 6  |FP .. ..| No | 0/0 |  2  | 32 |
 0    |   0   | UNDEF |Inst. Section: In  |  2  | Perm |  0    |  0  |
|  0  | 192.168.3.53    | 22    | 172.20.168.16   | 63800 | 6  |FP .. ..| No | 0/0 |  0  | 32 |
 1    |   0   | UNDEF |Inst. Section: Out |  0  | Perm |  257  |  0  |
|  0  | 192.168.3.53    | 18192 | 192.168.3.240   | 53925 | 6  |FP .. ..| No | 0/0 |  0  | 32 |
 1    |   0   | UNDEF |Inst. Section: Out |  0  | Perm |  219  |  0  |
|  0  | 192.168.3.240   | 257   | 192.168.3.53    | 60186 | 6  |FP .. ..| No | 0/0 |  1  | 32 |
 1    |   0   | UNDEF |Inst. Section: In  |  1  | Perm |  0    |  0  |

================================================================================================================
================================================================================
FP - from pool.    T - temporary connection.    PP - pending pernament.    In - inbound.    Out
- outbound.
[Expert@MyGW:0]#
```

fw ctl multik get_instance

Description

Shows CoreXL Firewall instance that processes the specified IPv4 connection.

Important - This command works only if the CoreXL Dynamic Dispatcher is disabled (see sk105261).

Syntax

- **To show the CoreXL Firewall instance that processes the specified IPv4 connection:**

```
fw ctl multik get_instance sip=<Source IPv4 Address>
dip=<Destination IPv4 Address> proto=<Protocol Number>
```

- **To show the CoreXL Firewall instance that processes the specified range of IPv4 connections:**

```
fw ctl multik get_instance sip=<Source IPv4 Address Start> -
<Source IPv4 Address End> dip=<Destination IPv4 Address Start>
- <Destination IPv4 Address End> proto=<Protocol Number>
```

Parameters

Parameter	Description
<Source IPv4 Address>	Source IPv4 address of the specified connection
<Source IPv4 Address Start>	First source IPv4 address of the specified range of IPv4 addresses
<Source IPv4 Address End>	Last source IPv4 address of the specified range of IPv4 addresses
<Destination IPv4 Address>	Destination IPv4 address of the specified connection
<Destination IPv4 Address Start>	First destination IPv4 address of the specified range of IPv4 addresses
<Destination IPv4 Address End>	Last destination IPv4 address of the specified range of IPv4 addresses
<Protocol Number>	See *IANA Protocol Numbers*. For example: ■ 1 = ICMP ■ 6 = TCP ■ 17 = UDP

Example for a specified IPv4 connection

```
[Expert@MyGW:0]# fw ctl multik get_instance sip=192.168.2.3 dip=172.30.241.66 proto=6
protocol: 6
192.168.2.3 -> 172.30.241.66 => 3
[Expert@MyGW:0]#
```

Example for a specified range of IPv4 connections

```
[Expert@MyGW:0]# fw ctl multik get_instance sip=192.168.2.3-192.168.2.8 dip=172.30.241.66
proto=6
protocol: 6
192.168.2.3 -> 172.30.241.66 => 3
192.168.2.4 -> 172.30.241.66 => 0
192.168.2.5 -> 172.30.241.66 => 3
192.168.2.6 -> 172.30.241.66 => 5
192.168.2.7 -> 172.30.241.66 => 4
192.168.2.8 -> 172.30.241.66 => 5
[Expert@MyGW:0]#
```

fw ctl multik print_heavy_conn

Description

These commands show the table with Heavy Connections (that consume the most CPU resources) in the CoreXL Dynamic Dispatcher.

For more information about the CoreXL Dynamic Dispatcher, see sk105261.

CoreXL suspects that a connection is "heavy" if it meets these conditions:

- Security Gateway detected the suspected connection during the last 24 hours
- The suspected connection lasts more than 10 seconds
- CoreXL Firewall instance that processes this connection causes a CPU load of over 60%
- The suspected connection utilizes more than 50% of the total work the applicable CoreXL Firewall instance does

The output table shows this information about the Heavy Connections:

- Source IP address
- Source Port
- Destination IP address
- Destination Port
- Protocol Number
- CoreXL Firewall instance ID that processes this connection
- CoreXL Firewall instance load on the CPU
- Connection's relative load on the CoreXL Firewall instance

Notes:

- These commands shows the suspected heavy connections even if they are already closed.
- In the *"CPView" on page 459* utility, go to **CPU > Top-Connections > InstancesX-Y > InstanceZ**. Refer to the **Top Connections** section.

Syntax

```
fw [-d] ctl multik print_heavy_conn
```

```
fw [-d] ctl multik heavy_conn_analyzer
```

Parameters

Parameter	Description
`-d`	Runs the command in debug mode. Use only if you troubleshoot the command itself. **Best Practice** - If you use this parameter, then redirect the output to a file, or use the `script` command to save the entire CLI session.

Example

```
[Expert@MyGW:0]# fw ctl multik print_heavy_conn
Source: 192.168.20.31; SPort: 51006; Dest: 172.30.40.55; DPort: 80; IPP: 6; Instance 1; Instance
Load 61%; Connection instance load 100%
Source: 192.168.20.31; SPort: 50994; Dest: 172.30.40.55; DPort: 80; IPP: 6; Instance 1; Instance
Load 61%; Connection instance load 100%
Source: 192.168.20.31; SPort: 50992; Dest: 172.30.40.55; DPort: 80; IPP: 6; Instance 1; Instance
Load 61%; Connection instance load 100%
[Expert@MyGW:0]#
```

fw ctl multik prioq

Description

Configures the CoreXL Firewall Priority Queues. For more information, see sk105762.

Important - This command saves the configuration in the `$FWDIR/conf/prioq_mode.conf` file. You must **not** edit this file manually.

Syntax for IPv4

```
fw ctl multik prioq [{0 | 1 | 2}]
```

Syntax for IPv6

```
fw6 ctl multik prioq [{0 | 1 | 2}]
```

Parameters

Parameter	Description
No Parameters	Shows the interactive menu for configuration of the CoreXL Firewall Priority Queues.
0	Disables the CoreXL Firewall Priority Queues.
1	Enables the CoreXL Firewall Priority Queues in the Evaluator-only mode.
2	Enables the CoreXL Firewall Priority Queues.

Example

```
[Expert@MyGW:0]# fw ctl multik prioq
Current mode is Off

Available modes:
0.      Off
1.      Evaluator-only
2.      On

Choose the desired mode number: (or 3 to Quit)
[Expert@MyGW:0]#
```

fw ctl multik queues

Description

Shows the CoreXL Queues and their utilization.

This command is supported only when the SecureXL works in the User Mode (UPPAK).

Syntax for IPv4

```
fw [-d] ctl multik queues
```

Syntax for IPv6

```
fw6 [-d] ctl multik queues
```

Parameters

Parameter	Description
`-d`	Runs the command in debug mode. Use only if you troubleshoot the command itself. **Best Practice** - If you use this parameter, then redirect the output to a file, or use the `script` command to save the entire CLI session.

fw ctl multik show_bypass_ports

Description

Shows the TCP and UDP ports configured in the bypass port list of the CoreXL Dynamic Dispatcher with the *"fw ctl multik add_bypass_port" on page 371* command.

For more information about the CoreXL Dynamic Dispatcher, see sk105261.

Important - This command reads the configuration from the `$FWDIR/conf/dispatcher_bypass.conf` file. You must **not** edit this file manually.

Syntax

```
fw ctl multik show_bypass_ports
```

Example

```
[Expert@MyGW:0]# fw ctl multik show_bypass_ports
dynamic dispatcher bypass port list:
(9999,8888)
[Expert@MyGW:0]#
```

fw ctl multik snd_dist

Description

Shows advanced information about distribution of connections by CoreXL SND instances.

Syntax for IPv4

```
fw [-d] ctl multik snd_dist {-h | -d | -e | -i <Number> | -p | -r
| -sn | -st | -t}
```

Syntax for IPv6

```
fw6 [-d] ctl multik snd_dist {-h | -d | -e | -i <Number> | -p | -r
| -sn | -st | -t}
```

Parameters

Parameter	Description
`fw -d`	Runs the command in debug mode. Use only if you troubleshoot the command itself. **Best Practice** - If you use this parameter, then redirect the output to a file, or use the `script` command to save the entire CLI session.
`-h`	Shows the built-in help.
`-e`	Enables the collection of the distribution information.
`-p`	Prints the distribution information.
`-r`	Restarts the collection of the distribution information (resets all counters).
`-d`	Disables the collection of the distribution information and deletes the collected information.
`-i <Number>`	Prints the specified number of 5-tuple iterations to the `dmesg` ring buffer.
`-t`	Toggles the debug - printing of 5-tuple iterations to the `dmesg` ring buffer. By default, this debug is disabled. Each time you run the command with this parameter, it changes the status of this debug.
`-st`	Shows the status of the debug ("`-t`") in the text format: ■ `Snd_dist port debugs is not set` ■ `Snd_dist port debugs is finished` ■ `Snd_dist port debugs is running`
`-sn`	Shows the status of the debug ("`-t`") in the numerical format. ■ `0` - The debug is finished ■ `1` - The debug is not set ■ `2` - The debug is running

Example

```
[Expert@MyGW:0]# fw ctl multik snd_dist -p
Statistics is off
[Expert@MyGW:0]#

[Expert@MyGW:0]# fw ctl multik snd_dist -e
[Expert@MyGW:0]#

[Expert@MyGW:0]# fw ctl multik snd_dist -p
Requests via hash              5046
Requests via user hint         0
Total requests                 5046
Number of active queues        1
Number of none instance cpus   1

Distribution by OPCODE:
=================================================================================================
|OPCODE                                    |requested via hash     |requested via user hint|
=================================================================================================
|CPHWD_API_OPCODE_GET_STATISTICS           |367                    |0                      |
=================================================================================================
|CPHWD_API_OPCODE_CPVIEW                   |163                    |0                      |
=================================================================================================
|CPHWD_API_OPCODE_GET_HWACCEL_STATUS       |24                     |0                      |
=================================================================================================
[Expert@MyGW:0]#

[Expert@MyGW:0]# fw ctl multik snd_dist -t
Snd_dist port debugs is on
[Expert@MyGW:0]#

[Expert@MyGW:0]# fw ctl multik snd_dist -st
Snd_dist port debugs is running
[Expert@MyGW:0]#

[Expert@MyGW:0]# fw ctl multik snd_dist -t
Snd_dist port debugs is off
[Expert@MyGW:0]#
```

fw ctl multik stat

Description

Shows information for each CoreXL Firewall instance.

Syntax for IPv4

```
fw [-d] ctl multik stat
```

Syntax for IPv6

```
fw6 [-d] ctl multik stat
```

Information in the output

- The ID number of each CoreXL Firewall instance (numbers starts from zero).
- The state of each CoreXL Firewall instance.
- The ID number of CPU core, on which the CoreXL Firewall instance runs (numbers starts from the highest available CPU ID).
- The number of concurrent connections the CoreXL Firewall instance currently handles.
- The peak number of concurrent connections the CoreXL Firewall instance handled from the time it started.

Parameters

Parameter	Description
`-d`	Runs the command in debug mode. Use only if you troubleshoot the command itself. **Best Practice** - If you use this parameter, then redirect the output to a file, or use the `script` command to save the entire CLI session.

Example

```
[Expert@MyGW:0]# fw ctl multik stat
ID | Active  | CPU    | Connections | Peak
----------------------------------------------
 0 | Yes     | 7      |           5 |       21
 1 | Yes     | 6      |           3 |       23
 2 | Yes     | 5      |           5 |       25
 3 | Yes     | 4      |           4 |       21
 4 | Yes     | 3      |           5 |       21
 5 | Yes     | 2      |           5 |       20
[Expert@MyGW:0]#

[Expert@MyGW:0]# fw6 ctl multik stat
ID | Active  | CPU    | Connections | Peak
----------------------------------------------
 0 | Yes     | 7      |           0 |        4
 1 | Yes     | 6      |           0 |        4
[Expert@MyGW:0]#
```

fw ctl multik start

Description

Starts all CoreXL Firewall instances on-the-fly, if they were stopped with the *"fw ctl multik stop" on page 394* command.

Syntax for IPv4

```
fw ctl multik start
```

Syntax for IPv6

```
fw6 ctl multik start
```

Example

```
[Expert@MyGW:0]# fw ctl multik stat
ID | Active  | CPU    | Connections | Peak
----------------------------------------------
 0 | No      | -      |           6 |       13
 1 | No      | -      |           3 |       11
 2 | No      | -      |           4 |       13
[Expert@MyGW:0]#
[Expert@MyGW:0]# fw ctl multik start
Instance 1 started (2 of 3 are active)
[Expert@MyGW:0]#
[Expert@MyGW:0]# fw ctl multik start
Instance 2 started (3 of 3 are active)
[Expert@MyGW:0]#
[Expert@MyGW:0]# fw ctl multik stat
ID | Active  | CPU    | Connections | Peak
----------------------------------------------
 0 | Yes     | 3      |           5 |       13
 1 | Yes     | 2      |           4 |       11
 2 | Yes     | 1      |           4 |       13
[Expert@MyGW:0]#
[Expert@MyGW:0]# fw ctl multik start
All instances are already active
[Expert@MyGW:0]#
```

fw ctl multik stop

Description

Stops all CoreXL Firewall instances on-the-fly.

Important - To start all CoreXL Firewall instances on-the-fly, run the *"fw ctl multik start" on page 393* command.

Syntax for IPv4

```
fw ctl multik stop
```

Syntax for IPv6

```
fw6 ctl multik stop
```

Example

```
[Expert@MyGW:0]# fw ctl multik stat
ID | Active  | CPU    | Connections | Peak
----------------------------------------------
 0 | Yes     | 3      |           5 |       13
 1 | Yes     | 2      |           4 |       11
 2 | Yes     | 1      |           4 |       13
[Expert@MyGW:0]#
[Expert@MyGW:0]# fw ctl multik stop
Instance 2 stopped (2 of 3 are active)
[Expert@MyGW:0]#
[Expert@MyGW:0]# fw ctl multik stop
Instance 1 stopped (1 of 3 are active)
[Expert@MyGW:0]#
[Expert@MyGW:0]# fw ctl multik stat
ID | Active  | CPU    | Connections | Peak
----------------------------------------------
 0 | Yes     | 3      |           4 |       13
 1 | No      | -      |           3 |       11
 2 | No      | -      |           7 |       13
[Expert@MyGW:0]#
[Expert@MyGW:0]# fw ctl multik stop
All instances are already inactive
[Expert@MyGW:0]#
[Expert@MyGW:0]# fw ctl multik stat
ID | Active  | CPU    | Connections | Peak
----------------------------------------------
 0 | No      | -      |           6 |       13
 1 | No      | -      |           3 |       11
 2 | No      | -      |           4 |       13
[Expert@MyGW:0]#
```

fw ctl multik utilize

Description

Shows the CoreXL queue utilization for each CoreXL Firewall instance.

Note - This command does not support VSX.

Syntax for IPv4

```
fw ctl multik utilize
```

Syntax for IPv6

```
fw6 ctl multik utilize
```

Example

```
[Expert@MyGW:0]# fw ctl multik utilize
ID | Utilize(%)   | Queue Elements
-----------------------------------
 0 |            1 |             30
 1 |            0 |             10
 2 |            0 |             17
[Expert@MyGW:0]#
[Expert@MyGW:0]# fw6 ctl multik utilize
ID | Utilize(%)   | Queue Elements
-----------------------------------
 0 |            0 |              0
 1 |            0 |              0
[Expert@MyGW:0]#
```

fw ctl affinity

The "`fw ctl affinity`" and "`fw6 ctl affinity`" commands show and configure the CoreXL affinity settings for:

- Interfaces
- User-space processes
- CoreXL Firewall instances

In addition, see *"taskset_us_all" on page 431*.

Running the 'fw ctl affinity -l' command in Gateway Mode

Description

The "`fw ctl affinity -l`" (for IPv4) and "`fw6 ctl affinity -l`" (for IPv6) commands show the current CoreXL affinity settings on a Security Gateway for:

- Interfaces
- User-space processes
- CoreXL Firewall instances

Syntax

- **To see the built-in help:**

  ```
  fw ctl affinity
  ```

- **To show all the existing affinities:**

  ```
  {fw | fw6} ctl affinity -l [-a] [-v] [-r] [-q]
  ```

- **To show the affinity for a specified interface:**

  ```
  {fw | fw6} ctl affinity -l -i <Interface Name>
  ```

- **To show the affinity for a specified CoreXL Firewall instance:**

  ```
  {fw | fw6} ctl affinity -l -k <CoreXL Firewall instance ID>
  ```

- **To show the affinity for a specified user-space process by its PID:**

  ```
  fw ctl affinity -l -p <Process ID>
  ```

- **To show the affinity for a specified user-space process by its name:**

  ```
  fw ctl affinity -l -n <Process Name>
  ```

- **To show the number of system CPU cores allowed by the installed CoreXL license:**

  ```
  fw -d ctl affinity -corelicnum
  ```

Parameters

Parameter	Description
`-i <Interface Name>`	Shows the affinity for the specified interface.
`-k <CoreXL Firewall instance ID>`	Shows the affinity for the specified CoreXL Firewall instance.
`-p <Process ID>`	Shows the affinity for the Check Point user-space process (for example: *fwd*, *vpnd)* specified by its PID.
`-n <Process Name>`	Shows the affinity for the Check Point user-space process (for example: *fwd*, *vpnd)* specified by its name.
`all`	Shows the affinity for all CPU cores (numbers start from zero).
`<CPU ID0> ... <CPU IDn>`	Shows the affinity for the specified CPU cores (numbers start from zero).
`-a`	Shows all current CoreXL affinities.
`-v`	Shows verbose output with IRQ numbers of interfaces.
`-r`	Shows the CoreXL affinities in reverse order.
`-q`	Suppresses the errors in the output.

Example 1

```
[Expert@MyGW:0]# fw ctl affinity -l
eth0: CPU 0
eth1: CPU 0
eth2: CPU 0
eth3: CPU 0
fw_0: CPU 7
fw_1: CPU 6
fw_2: CPU 5
fw_3: CPU 4
fw_4: CPU 3
fw_5: CPU 2
fwd: CPU 2 3 4 5 6 7
fgd50: CPU 2 3 4 5 6 7
status_proxy: CPU 2 3 4 5 6 7
rad: CPU 2 3 4 5 6 7
cpstat_monitor: CPU 2 3 4 5 6 7
mpdaemon: CPU 2 3 4 5 6 7
cpsead: CPU 2 3 4 5 6 7
cserver: CPU 2 3 4 5 6 7
rtmd: CPU 2 3 4 5 6 7
fwm: CPU 2 3 4 5 6 7
cpsemd: CPU 2 3 4 5 6 7
cpca: CPU 2 3 4 5 6 7
cprid: CPU 2 3 4 5 6 7
cpd: CPU 2 3 4 5 6 7
[Expert@MyGW:0]#
```

Example 2

```
[Expert@MyGW:0]# fw ctl affinity -l -a -v
Interface eth0 (irq 67): CPU 0
Interface eth1 (irq 75): CPU 0
Interface eth2 (irq 83): CPU 0
Interface eth3 (irq 59): CPU 0
fw_0: CPU 7
fw_1: CPU 6
fw_2: CPU 5
fw_3: CPU 4
fw_4: CPU 3
fw_5: CPU 2
fwd: CPU 2 3 4 5 6 7
fgd50: CPU 2 3 4 5 6 7
status_proxy: CPU 2 3 4 5 6 7
rad: CPU 2 3 4 5 6 7
cpstat_monitor: CPU 2 3 4 5 6 7
mpdaemon: CPU 2 3 4 5 6 7
cpsead: CPU 2 3 4 5 6 7
cserver: CPU 2 3 4 5 6 7
rtmd: CPU 2 3 4 5 6 7
fwm: CPU 2 3 4 5 6 7
cpsemd: CPU 2 3 4 5 6 7
cpca: CPU 2 3 4 5 6 7
cprid: CPU 2 3 4 5 6 7
cpd: CPU 2 3 4 5 6 7
[Expert@MyGW:0]#
```

Example 3

```
[Expert@MyGW:0]# fw ctl affinity -l -a -v -r
CPU 0:  eth0 (irq 67) eth1 (irq 75) eth2 (irq 83) eth3 (irq 59)
CPU 1:
CPU 2:  fw_5
        fwd fgd50 status_proxy rad cpstat_monitor mpdaemon cpsead cserver rtmd fwm cpsemd cpca
cprid cpd
CPU 3:  fw_4
        fwd fgd50 status_proxy rad cpstat_monitor mpdaemon cpsead cserver rtmd fwm cpsemd cpca
cprid cpd
CPU 4:  fw_3
        fwd fgd50 status_proxy rad cpstat_monitor mpdaemon cpsead cserver rtmd fwm cpsemd cpca
cprid cpd
CPU 5:  fw_2
        fwd fgd50 status_proxy rad cpstat_monitor mpdaemon cpsead cserver rtmd fwm cpsemd cpca
cprid cpd
CPU 6:  fw_1
        fwd fgd50 status_proxy rad cpstat_monitor mpdaemon cpsead cserver rtmd fwm cpsemd cpca
cprid cpd
CPU 7:  fw_0
        fwd fgd50 status_proxy rad cpstat_monitor mpdaemon cpsead cserver rtmd fwm cpsemd cpca
cprid cpd
All:
[Expert@MyGW:0]#
```

Example 4

```
[Expert@MyGW:0]# fw ctl affinity -l -i eth0
eth0: CPU 0
[Expert@MyGW:0]#
```

Example 5

```
[Expert@MyGW:0]# ps -ef | grep -v grep | egrep "PID|fwd"
UID        PID  PPID  C STIME TTY          TIME CMD
admin    26641 26452  0 Mar27 ?        00:06:56 fwd
[Expert@MyGW:0]#
[Expert@MyGW:0]# fw ctl affinity -l -p 26641
Process 26641: CPU 2 3 4 5 6 7
[Expert@MyGW:0]#
[Expert@MyGW:0]# fw ctl affinity -l -n fwd
fwd: CPU 2 3 4 5 6 7
[Expert@MyGW:0]#
```

Example 6

```
[Expert@MyGW:0]# fw ctl affinity -l -k 1
fw_1: CPU 6
[Expert@MyGW:0]#
```

Example 7

```
[Expert@MyGW:0]# fw -d ctl affinity -corelicnum
[5363 4134733584]@MyGW[4 Apr 18:11:03] Number of system CPUs 8
[5363 4134733584]@MyGW[4 Apr 18:11:03] cplic_get_navailable_cpus: fw_get_allowed_cpus_num
returned invalid value (100000) - all cpus considered as allowed!!!
4
[5363 4134733584]@MyGW[4 Apr 18:11:03] cpKeyTaskManager::~cpKeyTaskManager: called.
[Expert@MyGW:0]#
```

Running the 'fw ctl affinity -l' command in VSX Mode

Description

The "`fw ctl affinity -l`" (for IPv4) and `fw6 ctl affinity -l` (for IPv6) commands show the CoreXL affinity settings on a VSNext Security Group / Traditional VSX Gateway for:

- Interfaces
- User-space processes
- CoreXL Firewall instances

Note - Before running the "`fw ctl affinity -l -x`" commands, you must go to the context of the applicable Virtual Gateway / Traditional Virtual System / Traditional Virtual Router with the Gaia Clish command "`set virtual-system <VSID>`".

Syntax

- **To show the affinities in VSX mode (you can combine the optional parameters):**

```
{fw | fw6} ctl affinity -l -x
      [-vsid <VSID ranges>]
      [-cpu <CPU ID ranges>]
      [-flags {e | k | t | n | h | o}]
```

- **To show the number of system CPU cores allowed by the installed CoreXL license:**

```
fw -d ctl affinity -corelicnum
```

Parameters

Parameter	Description
`-vsid <VSID ranges>`	Shows the affinity for: ■ The specified single Virtual System (for example, `-vsid 7`) ■ The specified several Virtual Systems (for example, `-vsid 0-2 4`) **Important** - If you omit the `-vsid` parameter, the command runs in the current virtual context.
`<CPU ID ranges>`	Shows the affinity for: ■ The specified single CPU (for example, `-cpu 7`) ■ The specified several CPU cores (for example, `-cpu 0-2 4`)
`-flags {e | k | t | n | h | o}`	The `-flags` parameter requires at least one of these arguments: ■ `e` - Do not print the exception processes ■ `k` - Do not print the kernel threads ■ `t` - Print all process threads ■ `n` - Print the process name instead of the `/proc/<PID>/cmdline` ■ `h` - Print the CPU mask in Hex format ■ `o` - Print the output into the file called `/tmp/affinity_list_output` **Important** - You must specify multiple arguments together. For example: `-flags tn`

Example 1

```
[Expert@VSX_GW:0]# fw ctl affinity -l -x -cpu 0
-------------------------------------------------------------------
|PID     |VSID |               CPU                    |SRC|V|KT |EXC| NAME
-------------------------------------------------------------------
|      2 |   0 |                                    0 |   | | K |   |
|      3 |   0 |                                    0 |   | | K |   |
|      4 |   0 |                                    0 |   | | K |   |
|     14 |   0 |                                    0 |   | | K |   |
|     99 |   0 |                                    0 |   | | K |   |
|    278 |   0 |                                    0 |   | | K |   |
|    382 |   0 |                                    0 |   | | K |   |
|    674 |   0 |                                    0 |   | | K |   |
|   2195 |   0 |                                    0 |   | | K |   |
|   6348 |   0 |                                    0 |   | | K |   |
|   6378 |   0 |                                    0 |   | | K |   |
-------------------------------------------------------------------
PID   - represents the pid of the process
VSID  - represents the virtual device id
CPU   - represents the CPUs assigned to the specific process
SRC   - represents the source configuration file of the process - (V)SID / (I)nstance /
(P)rocess
V     - represents validity,star means that the actual affinity is different than the configured
affinity
KT    - represents whether the process is a kernel thread
EXC   - represents whether the process belongs to the process exception list (vsaffinity_
exception.conf)
[Expert@VSX_GW:0]#
```

Example 2

```
[Expert@VSX_GW:0]# fw ctl affinity -l -x -vsid 1
-------------------------------------------------------------------
|PID     |VSID |               CPU                    |SRC|V|KT |EXC| NAME
-------------------------------------------------------------------
|   3593 |   1 |                                1 2 3 |   | |   |   | httpd
|  10997 |   1 |                                1 2 3 |   | |   |   | cvpn_rotatelogs
|  11005 |   1 |                                1 2 3 |   | |   |   | httpd
|  22294 |   1 |                                1 2 3 |   | |   |   | routed
|  22328 |   1 |                                1 2 3 |   | |   |   | fwk_wd
|  22333 |   1 |                                1 2 3 | P | |   |   | fwk
|  22488 |   1 |                                1 2 3 |   | |   |   | cpd
|  22492 |   1 |                                1 2 3 |   | |   |   | fwd
|  22504 |   1 |                                1 2 3 |   | |   |   | cpviewd
|  22525 |   1 |                                1 2 3 |   | |   |   | mpdaemon
|  22527 |   1 |                                1 2 3 |   | |   |   | ci_http_server
|  30629 |   1 |                                1 2 3 |   | |   |   | vpnd
|  30631 |   1 |                                1 2 3 |   | |   |   | pdpd
|  30632 |   1 |                                1 2 3 |   | |   |   | pepd
|  30635 |   1 |                                1 2 3 |   | |   |   | fwpushd
|  30743 |   1 |                                1 2 3 |   | |   |   | dbwriter
|  30748 |   1 |                                1 2 3 |   | |   |   | cvpnproc
|  30752 |   1 |                                1 2 3 |   | |   |   | MoveFileServer
|  30756 |   1 |                                1 2 3 |   | |   |   | CvpnUMD
|  30760 |   1 |                                1 2 3 |   | |   |   | Pinger
|  30764 |   1 |                                1 2 3 |   | |   |   | IdlePinger
|  30770 |   1 |                                1 2 3 |   | |   |   | cvpnd
-------------------------------------------------------------------
[Expert@VSX_GW:0]#
```

Running the 'fw ctl affinity -s' command in Gateway Mode

Description

The "`fw ctl affinity -s`" command configures the CoreXL affinity settings on a Security Gateway for:

- Interfaces
- User-space processes
- CoreXL Firewall instances

Note - The command saves these configuration changes in the `$FWDIR/conf/fwaffinity.conf` configuration file.

Syntax

- **To see the built-in help:**

```
fw ctl affinity
```

- **To configure the affinity for a specified interface by its name:**

```
fw ctl affinity -s -i <Interface Name>
      all
      <CPU ID0> [ <CPU ID1> ... <CPU IDn> ]
```

- **To configure the affinity for a specified CoreXL Firewall instance:**

```
fw ctl affinity -s -k <CoreXL Firewall instance ID>
      all
      <CPU ID0> [ <CPU ID1> ... <CPU IDn> ]
```

- **To configure the affinity for a specified user-space process by its PID:**

```
fw ctl affinity -s -p <Process ID>
      all
      <CPU ID0> [ <CPU ID1> ... <CPU IDn> ]
```

- **To configure the affinity for a specified user-space process by its name:**

```
fw ctl affinity -s -n <Process Name>
      all
      <CPU ID0> [ <CPU ID1> ... <CPU IDn> ]
```

Parameters

Parameter	Description
`-i <Interface Name>`	Configures the affinity for the specified interface.
`-k <CoreXL Firewall instance ID>`	Configures the affinity for the specified CoreXL Firewall instance.
`-p <Process ID>`	Configures the affinity for the Check Point user-space process (for example: *fwd*, *vpnd)* specified by its PID.
`-n <Process Name>`	Configures the affinity for the Check Point user-space process (for example: *fwd*, *vpnd)* specified by its name. **Important** - The process name is case-sensitive.
`all`	Configures the affinity for all CPU cores (numbers start from zero).
`<CPU ID0> ... <CPU IDn>`	Configures the affinity for the specified CPU cores (numbers start from zero).

Example 1 - Affine the interface eth1 to the CPU core #1

```
[Expert@MyGW:0]# fw ctl affinity -s -i eth1 1
eth1: CPU 1 - set successfully
Multi-queue affinity was not changed.  For More info, see sk113834.
[Expert@MyGW:0]#
```

Example 2 - Affine the CoreXL Firewall instance #1 to the CPU core #2

```
[Expert@MyGW:0]# fw ctl affinity -s -k 1 2
fw_1: CPU 2 - set successfully
Multi-queue affinity was not changed.  For More info, see sk113834.
[Expert@MyGW:0]#
```

Example 3 - Affine the process CPD by its PID to the CPU core #2

```
[Expert@MyGW:0]# cpwd_admin list | egrep "PID|cpd"
APP         PID    STAT  #START  START_TIME             MON  COMMAND
CPD         6080   E     1       [13:46:27] 17/9/2018   Y    cpd
[Expert@MyGW:0]#
[Expert@MyGW:0]# fw ctl affinity -s -p 6080 2
Process 6080: CPU 2 - set successfully
Multi-queue affinity was not changed.  For More info, see sk113834.
[Expert@MyGW:0]#
```

Example 4 - Affine the process CPD by its name to the CPU core #2

```
[Expert@MyGW:0]# fw ctl affinity -s -n cpd 2
cpd: CPU 2 - set successfully
Multi-queue affinity was not changed.  For More info, see sk113834.
[Expert@MyGW:0]#
```

Running the 'fw ctl affinity -s' command in VSX Mode

Description

The "`fw ctl affinity -s`" command configures the CoreXL affinity settings on a VSNext Security Group / Traditional VSX Gateway for:

- Interfaces
- User-space processes
- CoreXL Firewall instances

Syntax

- **To see the built-in help:**

```
fw ctl affinity
```

- **To configure the affinities of VSNext Virtual Gateways / Traditional VSX Virtual Systems:**

```
fw ctl affinity -s -d [-vsid <VSID ranges> ] -cpu <CPU ID
ranges>
```

- **To configure the affinities of a specified user-space process:**

```
fw ctl affinity -s -d -pname <Process Name> [-vsid <VSID
ranges>]
      -cpu all
      -cpu <CPU ID ranges>
```

- **To configure the affinities of specified FWK daemon instances (user-space Firewall):**

```
fw ctl affinity -s -d -inst <Instances Ranges> -cpu <CPU ID
ranges>
```

- **To configure the affinities of all FWK instances (user-space Firewalls):**

```
fw ctl affinity -s -d -fwkall <Number of CPUs>
```

- **To reset the affinities to defaults:**

```
fw ctl affinity
      -vsx_factory_defaults
      -vsx_factory_defaults_no_prompt
```

Important

- The command saves these configuration changes in the `$FWDIR/conf/fwaffinity.conf` configuration file.
- When you configure affinity of an interface, it automatically configures the affinities of all other interfaces that share the same IRQ to the same CPU core.

Parameters

Parameter	Description
`-vsid <VSID ranges>`	Configures the affinity for: ■ One specified Virtual System. For example: `-vsid 7` ■ Several specified Virtual Systems. For example: `-vsid 0-2 4` **Note** - If you omit the "`-vsid`" parameter, the command uses the current virtual context.
`<CPU ID ranges>`	Configures the affinity to: ■ One specified CPU core. For example: `-cpu 7` ■ Several specified CPU cores. For example: `-cpu 0-2 4` **Important** - Numbers of CPU cores start from zero.
`-pname <Process Name>`	Configures the affinity for the Check Point daemon specified by its name (for example: `fwd`, `vpnd`). **Important** - The process name is case-sensitive.
`-inst <Instances Ranges>`	Configures the affinity for: ■ One specified FWK daemon instance. For example: `-inst 7` ■ Several specified FWK daemon instances. For example: `-inst 0 2 4`

Parameter	Description
`-fwkall <Number of CPUs>`	Configures the affinity for all running FWK daemon instances to the specified number of CPU cores. If it is necessary to affine all running FWK daemon instances to all CPU cores, enter the number of all available CPU cores.
`-vsx_factory_defaults`	Deletes all existing affinity settings and creates the default affinity settings during the next reboot. **Important** - Before this operation, the command prompts the user whether to proceed. You must reboot to complete the operation.
`-vsx_factory_defaults_no_prompt`	Deletes all current affinity settings and creates the default affinity settings during the next reboot. **Important** - Before this operation, the command does **not** prompt the user whether to proceed. You must reboot to complete the operation.

Example 1 - Affine the Virtual Devices #0,1,2,4,7,8 to the CPU cores #0,1,2,4

```
[Expert@VSX_GW:0]# fw ctl affinity -s -d -vsid 0-2 4 6-8 -cpu 0-2 4
VDevice 0-2 4 6-8 : CPU 0 1 2 4 - set successfully
Multi-queue affinity was not changed.  For More info, see sk113834.
[Expert@VSX_GW:0]#
```

Example 2 - Affine the process CPD by its name for Virtual Devices #0-12 to the CPU core #7

```
[Expert@VSX_GW:0]# fw ctl affinity -s -d -pname cpd -vsid 0-12 -cpu 7
VDevice 0-12 : CPU 7 - set successfully
Multi-queue affinity was not changed.  For More info, see sk113834.
Warning: some of the VSIDs did not exist
[Expert@VSX_GW:0]#
```

Example 3 - Affine the FWK daemon instances #0,2,4 to the CPU core #5

```
[Expert@VSX_GW:0]# fw ctl affinity -s -d -inst 0 2 4 -cpu 5
VDevice 0 2 4: CPU 5 - set successfully
Multi-queue affinity was not changed.  For More info, see sk113834.
[Expert@VSX_GW:0]#
```

Example 4 - Affine all FWK daemon instances to the last two CPU cores

```
[Expert@VSX_GW:0]# fw ctl affinity -s -d -fwkall 2
VDevice 0-2 : CPU 2 3 - set successfully
Multi-queue affinity was not changed.  For More info, see sk113834.
[Expert@VSX_GW:0]#
```

Example 5 - Affine all FWK daemon instances to all CPU cores

```
[Expert@VSX_GW:0]# fw ctl affinity -s -d -fwkall 4
There are configured processes/FWK instances
(y) will override all currently configured affinity and erase the configuration files
(n) will set affinity only for unconfigured processes/threads
Do you want to override existing configurations (y/n) ? y
VDevice 0-2 : CPU all - set successfully
Multi-queue affinity was not changed.  For More info, see sk113834.
[Expert@VSX_GW:0]#
```

fw -i

Description

By default, the "`fw`" commands apply to the entire Security Gateway.

The `fw` commands show aggregated information for all CoreXL Firewall instances.

The `fw -i` commands apply to the specified CoreXL Firewall instance.

Important:

- You can run this command in the Expert mode or in Gaia Clish (Gaia gClish on Scalable Platforms).
- In a Cluster, you must configure all the Cluster Members in the same way.
- On Scalable Platforms, you must connect to the applicable Security Group.

Syntax

```
fw -i <ID of CoreXL Firewall instance> <Command>
```

Parameters

Parameter	Description
<ID of CoreXL Firewall instance>	Specifies the ID of the CoreXL Firewall instance. To see the available IDs, run the "`fw ctl multik stat`"*"fw ctl multik stat" on page 391* command.
<Command>	Only these commands support the `fw -i` syntax: ■ `fw -i <ID> conntab ...` ■ `fw -i <ID> ctl get ...` ■ `fw -i <ID> ctl leak ...` ■ `fw -i <ID> ctl pstat ...` ■ `fw -i <ID> ctl set ...` ■ `fw -i <ID> monitor ...` ■ `fw -i <ID> tab ...` For details and additional parameters for any of these commands, refer to the corresponding entry for each command.

Example 1 - Show the Connections table for CoreXL Firewall instance #1

```
fw -i 1 tab -t connections
```

Example 2 - Show various internal statistics for CoreXL Firewall instance #1

```
fw -i 1 ctl pstat
```

fwboot bootconf

Description

Configures boot security options.

Notes:

- You must run this command from the Expert mode.
- The settings are saved in the `$FWDIR/boot/boot.conf` file.

  Warning - To avoid issues, do not edit the `$FWDIR/boot/boot.conf` file manually. Edit the file only with this command.
- Refer to these related commands:
  - *"fwboot corexl" on page 418*

Syntax to show the current boot security options

```
[Expert@HostName:0]# $FWDIR/boot/fwboot bootconf
      get_corexl
      get_core_override
      get_def
      get_ipf
      get_ipv6
      get_kernnum
      get_kern6num
```

Syntax to configure the boot security options

```
[Expert@HostName:0]# $FWDIR/boot/fwboot bootconf
      set_corexl {0 | 1}
      set_core_override <number>
      set_def [</path/filename>]
      set_ipf {0 | 1}
      set_ipv6 {0 | 1}
      set_kernnum <number>
      set_kern6num <number>
```

Parameters

Parameter	Description
No Parameters	Shows the built-in help with available parameters.

Parameter	Description
`get_corexl`	Shows if the CoreXL is enabled or disabled: ■ 0 - disabled ■ 1 - enabled **Note** - In the `$FWDIR/boot/boot.conf` file, refer to the value of the `COREXL_INSTALLED`.
`get_core_override`	Shows the number of overriding CPU cores. The SMT (HyperThreading) feature (sk93000) uses this configuration to set the number of CPU cores after reboot. **Note** - In the `$FWDIR/boot/boot.conf` file, refer to the value of the `CORE_OVERRIDE`.
`get_def`	Shows the configured path and the name of the Default Filter policy file (default is `$FWDIR/boot/default.bin`). **Note** - In the `$FWDIR/boot/boot.conf` file, refer to the value of the `DEFAULT_FILTER_PATH`.
`get_ipf`	Shows if the IP Forwarding during boot is enabled or disabled: ■ 0 - disabled (Security Gateway does not forward traffic between its interfaces during boot) ■ 1 - enabled **Note** - In the `$FWDIR/boot/boot.conf` file, refer to the value of the `CTL_IPFORWARDING`.
`get_ipv6`	Shows if the IPv6 support is enabled or disabled: ■ 0 - disabled ■ 1 - enabled **Note** - In the `$FWDIR/boot/boot.conf` file, refer to the value of the `IPV6_INSTALLED`.
`get_kernnum`	Shows the configured number of IPv4 CoreXL Firewall instances. **Note** - In the `$FWDIR/boot/boot.conf` file, refer to the value of the `KERN_INSTANCE_NUM`.
`get_kern6num`	Shows the configured number of IPv6 CoreXL Firewall instances. **Note** - In the `$FWDIR/boot/boot.conf` file, refer to the value of the `KERN6_INSTANCE_NUM`.

Parameter	Description
`set_corexl {0 | 1}`	Enables or disables CoreXL: ■ 0 - disables ■ 1 - enables **Notes:** ■ In the `$FWDIR/boot/boot.conf` file, refer to the value of the `COREXL_INSTALLED`. ■ To configure CoreXL, use the *"cpconfig" on page 358* menu.
`set_core_ override <number>`	Configures the number of overriding CPU cores. The SMT (HyperThreading) feature (sk93000) uses this configuration to set the number of CPU cores after reboot. **Note** - In the `$FWDIR/boot/boot.conf` file, refer to the value of the `CORE_OVERRIDE`.
`set_def [< /path/filename >]`	Configures the path and the name of the Default Filter policy file (default is `$FWDIR/boot/default.bin`). **Notes:** ■ In the `$FWDIR/boot/boot.conf` file, refer to the value of the `DEFAULT_FILTER_PATH`. ■ If you do not specify the path and the name explicitly, then the value of the `DEFAULT_FILTER_PATH` is set to 0. As a result, Security Gateway does not load a Default Filter during boot. **Best Practice** - The best location for this file is the `$FWDIR/boot/` directory.
`set_ipf {0 | 1}`	Configures the IP forwarding during boot: ■ 0 - disables (forbids the Security Gateway to forward traffic between its interfaces during boot) ■ 1 - enables **Note** - In the `$FWDIR/boot/boot.conf` file, refer to the value of the `CTL_IPFORWARDING`.

Parameter	Description
`set_ipv6 {0 | 1}`	Enables or disables the IPv6 Support: ■ 0 - disables ■ 1 - enables **Notes:** ■ In the `$FWDIR/boot/boot.conf` file, refer to the value of the `IPV6_INSTALLED`. ■ Configure the IPv6 Support in Gaia Portal, or Gaia Clish. See the *R82 Gaia Administration Guide*.
`set_kernnum <number>`	Configures the number of IPv4 CoreXL Firewall instances. **Notes:** ■ In the `$FWDIR/boot/boot.conf` file, refer to the value of the `KERN_INSTANCE_NUM`. ■ To configure CoreXL, use the *"cpconfig" on page 358* menu.
`set_kern6num <number>`	Configures the number of IPv6 CoreXL Firewall instances. **Notes:** ■ In the `$FWDIR/boot/boot.conf` file, refer to the value of the `KERN6_INSTANCE_NUM`. ■ To configure CoreXL, use the *"cpconfig" on page 358* menu.

fwboot corexl

Description

Configures and monitors the CoreXL.

Note - The settings are saved in the `$FWDIR/boot/boot.conf` file.

Warning - To avoid issues, do not edit the `$FWDIR/boot/boot.conf` file manually. Edit the file only with this command.

Syntax to show CoreXL configuration

```
[Expert@HostName:0]# $FWDIR/boot/fwboot corexl
      core_count
      curr_instance4_count
      curr_instance6_count
      def_instance4_count
      def_instance6_count
      eligible
      installed
      max_instance4_count
      max_instances4_32bit
      max_instances4_64bit
      max_instance6_count
      max_instances_count
      max_instances_32bit
      max_instances_64bit
      min_instance_count
      unsupported_features
```

Syntax to configure CoreXL

Important:

- The configuration commands are for Check Point use only. To configure CoreXL, use the **Check Point CoreXL** option in the *"cpconfig" on page 358* menu.
- After all changes in CoreXL configuration on the Security Gateway, you must reboot it.
- In a Cluster, you must configure all the Cluster Members in the same way.

```
[Expert@HostName:0]# $FWDIR/boot/fwboot corexl
      def_by_allowed [n]
      default
      [-v] disable
      [-v] enable [n] [-6 k]
      vmalloc_recalculate
```

Parameters

Parameter	Description
No Parameters	Shows the built-in help with available parameters.
`core_count`	Returns the number of CPU cores on this computer. **Example** `[Expert@MyGW:0]# $FWDIR/boot/fwboot corexl core_count` `[Expert@MyGW:0]# echo $?` `4` `[Expert@MyGW:0]#` `[Expert@MyGW:0]# cat /proc/cpuinfo | grep processor` `processor : 0` `processor : 1` `processor : 2` `processor : 3` `[Expert@MyGW:0]#`

Parameter	Description
`curr_ instance4_ count`	Returns the current configured number of IPv4 CoreXL Firewall instances. **Example** `[Expert@MyGW:0]# $FWDIR/boot/fwboot corexl curr_ instance4_count` `[Expert@MyGW:0]# echo $?` `3` `[Expert@MyGW:0]#` `[Expert@MyGW:0]# fw ctl multik stat` `ID | Active | CPU | Connections | Peak` `----------` ` 0 | Yes | 3 | 11 | 18` ` 1 | Yes | 2 | 12 | 18` ` 2 | Yes | 1 | 13 | 18` `[Expert@MyGW:0]#`
`curr_ instance6_ count`	Returns the current configured number of IPv6 CoreXL Firewall instances. **Example** `[Expert@MyGW:0]# $FWDIR/boot/fwboot corexl curr_ instance6_count` `[Expert@MyGW:0]# echo $?` `2` `[Expert@MyGW:0]#` `[Expert@MyGW:0]# fw6 ctl multik stat` `ID | Active | CPU | Connections | Peak` `----------` ` 0 | Yes | 3 | 11 | 18` ` 1 | Yes | 2 | 12 | 18` `[Expert@MyGW:0]#`
`def_by_ allowed [n]`	Sets the default configuration for CoreXL according to the specified allowed number of CPU cores.
`default`	Sets the default configuration for CoreXL.

Parameter	Description
`def_ instance4_ count`	Returns the default number of IPv4 CoreXL Firewall instances for this Security Gateway. **Example** `[Expert@MyGW:0]# $FWDIR/boot/fwboot corexl def_ instance4_count` `[Expert@MyGW:0]# echo $?` `3` `[Expert@MyGW:0]#`
`def_ instance6_ count`	Returns the default number of IPv6 CoreXL Firewall instances for this Security Gateway. **Example** `[Expert@MyGW:0]# $FWDIR/boot/fwboot corexl def_ instance6_count` `[Expert@MyGW:0]# echo $?` `2` `[Expert@MyGW:0]#`
`[-v] disable`	Disables CoreXL. ■ `-v` - Leaves the high memory (`vmalloc`) unchanged. See the *"cp_conf corexl" on page 355* command.
`eligible`	Returns whether CoreXL can be enabled on this Security Gateway. ■ 0 - CoreXL cannot be enabled ■ 1 - CoreXL can be enabled **Example** `[Expert@MyGW:0]# $FWDIR/boot/fwboot corexl eligible` `[Expert@MyGW:0]# echo $?` `1` `[Expert@MyGW:0]#`

Parameter	Description
`[-v] enable [n] [-6 k]`	Enables CoreXL with 'n' IPv4 Firewall instances and optionally 'k' IPv6 Firewall instances. ■ `-v` - Leaves the high memory (`vmalloc`) unchanged. ■ `n` - Denotes the number of IPv4 CoreXL Firewall instances. ■ `k` - Denotes the number of IPv6 CoreXL Firewall instances. See the *"cp_conf corexl" on page 355* command.
`installed`	Returns whether CoreXL is installed (enabled) on this Security Gateway. ■ 0 - CoreXL is not enabled ■ 1 - CoreXL is enabled **Example** `[Expert@MyGW:0]# $FWDIR/boot/fwboot corexl installed` `[Expert@MyGW:0]# echo $?` `1` `[Expert@MyGW:0]#`
`max_ instance4_ count`	Returns the maximum allowed number of IPv4 CoreXL Firewall instances for this Security Gateway. **Example** `[Expert@MyGW:0]# $FWDIR/boot/fwboot corexl max_ instance4_count` `[Expert@MyGW:0]# echo $?` `4` `[Expert@MyGW:0]#`
`max_ instances 4_32bit`	Returns the maximum allowed number of IPv4 CoreXL Firewall instances for a Security Gateway that runs Gaia with 32-bit kernel. **Example** `[Expert@MyGW:0]# $FWDIR/boot/fwboot corexl max_ instances4_32bit` `[Expert@MyGW:0]# echo $?` `14` `[Expert@MyGW:0]#`

Parameter	Description
`max_instances4_64bit`	Returns the maximum allowed number of IPv4 CoreXL Firewall instances for a Security Gateway that runs Gaia with 64-bit kernel. **Example** `[Expert@MyGW:0]# $FWDIR/boot/fwboot corexl max_instances4_64bit` `[Expert@MyGW:0]# echo $?` `38` `[Expert@MyGW:0]#`
`max_instance6_count`	Returns the maximum allowed number of IPv6 CoreXL Firewall instances for this Security Gateway. **Example** `[Expert@MyGW:0]# $FWDIR/boot/fwboot corexl max_instance6_count` `[Expert@MyGW:0]# echo $?` `3` `[Expert@MyGW:0]#`
`max_instances_count`	Returns the total maximum allowed number of CoreXL Firewall instances (IPv4 and IPv6) for this Security Gateway. **Example** `[Expert@MyGW:0]# $FWDIR/boot/fwboot corexl max_instances_count` `[Expert@MyGW:0]# echo $?` `40` `[Expert@MyGW:0]#`
`max_instances_32bit`	Returns the total maximum allowed number of CoreXL Firewall instances for a Security Gateway that runs Gaia with 32-bit kernel. **Example** `[Expert@MyGW:0]# $FWDIR/boot/fwboot corexl max_instances_32bit` `[Expert@MyGW:0]# echo $?` `16` `[Expert@MyGW:0]#`

Parameter	Description
`max_instances_64bit`	Returns the total maximum allowed number of CoreXL Firewall instances for a Security Gateway that runs Gaia with 64-bit kernel. **Example** `[Expert@MyGW:0]# $FWDIR/boot/fwboot corexl max_instances_64bit` `[Expert@MyGW:0]# echo $?` `40` `[Expert@MyGW:0]#`
`min_instance_count`	Returns the minimum allowed number of IPv4 CoreXL Firewall instances. **Example** `[Expert@MyGW:0]# $FWDIR/boot/fwboot corexl min_instance_count` `[Expert@MyGW:0]# echo $?` `2` `[Expert@MyGW:0]#`
`vmalloc_recalculate`	Updates the value of the `vmalloc` parameter in the `/boot/grub/grub.conf` file.
`unsupported_features`	Returns 1 if at least one feature is configured, which CoreXL does not support. **Example** `[Expert@MyGW:0]# $FWDIR/boot/fwboot corexl unsupported_features` `corexl unsupported feature: QoS is configured.` `[Expert@MyGW:0]# echo $?` `1` `[Expert@MyGW:0]#`

fwboot cpuid

Description

Shows the number of available CPUs and CPU cores on this Security Gateway.

Syntax

```
[Expert@HostName:0]# $FWDIR/boot/fwboot cpuid
      {-h | -help | --help}
      -c
      --full
      ht_aware
      -n
      --possible
```

Parameters

Parameter	Description
No Parameters	Shows the IDs of the available CPU cores on this Security Gateway. **Example** `[Expert@MyGW:0]# $FWDIR/boot/fwboot cpuid` `3 2 1 0` `[Expert@MyGW:0]#`
`-c`	Counts the number of available CPU cores on this Security Gateway. The command stores the returned number as its exit code. **Example** `[Expert@MyGW:0]# $FWDIR/boot/fwboot cpuid -c` `[Expert@MyGW:0]# echo $?` `4` `[Expert@MyGW:0]#`

Parameter	Description
`--full`	Shows a full map of the available CPUs and CPU cores on this Security Gateway. **Example** `[Expert@MyGW:0]# $FWDIR/boot/fwboot cpuid --full` `cpuid phys_id core_id thread_id` ` 0 0 0 0` ` 1 2 0 0` ` 2 4 0 0` ` 3 6 0 0` `[Expert@MyGW:0]#`
`ht_aware`	Shows the CPU cores in the order of their awareness of Hyper-Threading. **Example** `[Expert@MyGW:0]# $FWDIR/boot/fwboot cpuid ht_aware` `3 2 1 0` `[Expert@MyGW:0]#`
`-n`	Counts the number of available CPUs on this Security Gateway. The command stores the returned number as its exit code. **Example** `[Expert@MyGW:0]# $FWDIR/boot/fwboot cpuid -n` `[Expert@MyGW:0]# echo $?` `4` `[Expert@MyGW:0]#`
`--possible`	Counts the number of possible CPU cores. The command stores the returned number as its exit code. **Example** `[Expert@MyGW:0]# $FWDIR/boot/fwboot cpuid --possible` `[Expert@MyGW:0]# echo $?` `4` `[Expert@MyGW:0]#`

fwboot ht

Important - This command is obsolete and is not supported. To configure SMT (HyperThreading) feature, follow sk93000.

fwboot multik_reg

Description

Shows the internal memory address of the registration function for the specified CoreXL Firewall instance.

Important - This command is for Check Point use only.

Note - You must run this command from the Expert mode.

Syntax

```
[Expert@HostName:0]# $FWDIR/boot/fwboot multik_reg <Number of CoreXL Firewall instance> {ipv4 | ipv6} [-d]
```

Parameters

Parameter	Description
No Parameters	Shows the built-in help with available parameters.
<Number of CoreXL Firewall instance>	Specifies the ID number of the CoreXL Firewall instance.
`ipv4`	Specifies to work with IPv4 CoreXL Firewall instances.
`ipv6`	Specifies to work with IPv6 CoreXL Firewall instances.
`-d`	Shows the decimal 64-bit address of the hook function.

Example

```
[Expert@MyGW:0]# fw ctl multik stat
ID | Active  | CPU    | Connections | Peak
----------------------------------------------
 0 | Yes     | 3      |          11 |       18
 1 | Yes     | 2      |          12 |       18
 2 | Yes     | 1      |          13 |       18
[Expert@MyGW:0]#


[Expert@MyGW:0]# $FWDIR/boot/fwboot multik_reg 0 ipv4
0
[Expert@MyGW:0]#

[Expert@MyGW:0]# $FWDIR/boot/fwboot multik_reg 1 ipv4
0xffffffff8a2a5690
[Expert@MyGW:0]#

[Expert@MyGW:0]# $FWDIR/boot/fwboot multik_reg 2 ipv4
0xffffffff8a2a5690
[Expert@MyGW:0]#
```

fwboot post_drv

Description

Loads the Firewall driver for CoreXL during boot.

Important:

- This command is for Check Point use only.
- If you run this command, Security Gateway can block all traffic. In such case, you must connect to the Security Gateway over a console and restart Check Point services with the "`cpstop`" and "`cpstart`" commands. Alternatively, you can reboot the Security Gateway.

Note - You must run this command from the Expert mode.

Syntax

```
[Expert@HostName:0]# $FWDIR/boot/fwboot post_drv {ipv4 | ipv6}
```

Parameters

Parameter	Description
No Parameters	Shows the built-in help with available parameters.
`ipv4`	Loads the IPv4 Firewall driver for CoreXL.
`ipv6`	Loads the IPv6 Firewall driver for CoreXL.

taskset_us_all

Description

The script `$FWDIR/bin/taskset_us_all` configures the affinity of all running User Space processes to CPU cores.

- You can enter the required CPU cores either as a list of CPU core IDs, or as a CPU bitmask.
- You can configure the affinity User Space processes to CPU cores that run as CoreXL Firewall instances, as CoreXL SND, or all CPU cores.
- You can configure a list of User Space processes that this script must ignore.
- The script creates this log file: `$FWDIR/log/taskset_us_all.elg`

Note - You must run this command from the Expert mode.

Related Information

- *"Configuring Affinity Settings" on page 326*
- *"fw ctl affinity" on page 396*

Syntax

```
[Expert@HostName:0]# taskset_us_all
      --help}
      {-i | --instance} [<IDs of Instances>]
      {-l | --list} <IDs of CPU Cores>
            [--ignore_default]
            [--ignore_exe
"<ProcessName1>|<ProcessName2>|...|<ProcessNameN>"
      {-m | --mask} <CPU BitMask>
      {-s | --snd} [--no_zero]
```

On a Scalable Platform Security Group, you must run: `g_all taskset_us_all <parameters>`

Parameters

Parameter	Type	Description
No Parameters	N / A	Shows the syntax to call the built-in help.
`--help`	Optional	Shows the built-in help with available parameters.
`{-i | --instance} [<IDs of Instances>]`	Optional	Cnfigures the CPU affinity only to CPU cores that run as CoreXL Firewall Instances. Example for the CoreXL Firewall Instances 0, 2, and 5: `--instance 0,2,5` **Notes:** ■ Mutually exclusive with the parameter '`-s | --snd`'. ■ Not supported in VSX mode.

Parameter	Type	Description
{-l | --list} <*IDs of CPU Cores*> [--ignore_default] [--ignore_exe "<*ProcessName1*>|<*ProcessName2*>|...|<*ProcessNameN*>"]	Mandatory	Configures the script to use the argument as a list of CPU core IDs. Examples for CPU cores from 0 to 3: ■ Comma-separated list: `--list 0,1,2,3` ■ Range of consecutive IDs: `--list 0-3` ■ Range and specific ID: `--list 0-2,3`

Parameter	Type	Description
		Notes: ■ The optional sub-parameter '`--ignore_default`' configures the script to ignore a predefined list of Check Point executable files. There are specific Check Point executable files, whose CPU affinity you must **not** change. If you change the CPU affinity of these specific processes, the performance of the Security Gateway **decreases**. The predefined list of Check Point executable files to ignore: `fwk`, `dmd_run`, `usim_x86`, `usim_wd_agent`, `usim_launcher` ■ The optional sub-parameter '`--ignore_exe "<`*`ProcessName1`*`>|<`*`ProcessName2`*`>|...|<`*`ProcessNameN`*`>"`' configures the script to ignore the specified executable files. Example: `--ignore_exe "bash|fwk|clishd"` To configure a constant list of ignored executable files, edit the `$FWDIR/bin/taskset_us_all` script and configure the value of the variable '`ignore_list=""`'.

Parameter	Type	Description
`{-m | --mask} <CPU BitMask>`	Mandatory	Configures the script to use the argument as a bitmask of CPU cores IDs. Example bitmask for CPU core IDs from 0 to 3 on a Security Gateway with 16 CPU cores: `0x0000000000000001 (for CPU 0) +` `0x0000000000000010 (for CPU 1) +` `0x0000000000000100 (for CPU 2) +` `0x0000000000001000 (for CPU 3) =` `----------` `0x0000000000001111` Example bitmask for CPU core IDs 5, 7, 9, and 15 on a Security Gateway with 16 CPU cores: `0x0000000000100000 (for CPU 5) +` `0x0000000010000000 (for CPU 7) +` `0x0000001000000000 (for CPU 9) +` `0x1000000000000000 (for CPU 15) =` `----------` `0x1000001010100000`
`{-s | --snd} [--no_zero]`	Optional	Configures the CPU affinity only to CPU cores that run as CoreXL SND. **Notes:** ■ Mutually exclusive with the parameter '`-i | --instance`'. ■ The optional sub-parameter '`--no_zero`' configures the script to ignore the CPU Core 0.

Multi-Queue

By default, each network interface has one traffic queue handled by one CPU.

You cannot use more CPU cores for acceleration than the number of interfaces handling traffic.

Multi-Queue configures more than one traffic queue for each network interface.

For each interface, more than one CPU core is used for acceleration.

Note - Multi-Queue is applicable only if SecureXL is enabled (this is the default).

Overview:

- Multi-Queue is enabled by default on all interfaces that use the supported drivers.
- The number of traffic queues on each supported interface is determined automatically, based on:
  - The number of available CPU cores that run CoreXL SND Instances.
  - The limitations of the interfaces and its driver.
- Traffic queues are automatically affined to the CPU cores that runs CoreXL SND Instances.
- Changes in Multi-Queue configuration do **not** require a reboot.
- You configure Multi-Queue on the command line - in Gaia Clish (Gaia gClish on Scalable Platforms), or in the Expert mode.

Multi-Queue Requirements and Limitations

- Multi-Queue only supports Security Gateways / Cluster Members with two or more CPU cores.
- Multi-Queue only supports interfaces that use these drivers:

Driver	Max Speed	Interface / Driver Description
igb	1 Gbps	Intel® Network Adapter Driver for PCIe 1 Gigabit Ethernet Network
ixgbe	10 Gbps	Intel® Network Adapter Driver for PCIe 10 Gigabit Ethernet Network
i40e	40 Gbps	Intel® Network Adapter Driver for PCIe 40 Gigabit Ethernet Network
i40evf	40 Gbps	Intel® i40e driver for Virtual Function Network Devices
ice	25 Gbps	Intel® Network Adapter Driver for E810 Series Devices
mlx5_core	40 Gbps	Mellanox® ConnectX® mlx5 core driver
ena	20 Gbps	Elastic Network Adapter in Amazon® EC2
virtio_net	10 Gbps	VirtIO paravirtualized device driver from KVM®
vmxnet3	10 Gbps	VMXNET Generation 3 driver from VMware®

- Multi-Queue does not use network interfaces that are currently in the down state.
- The number of traffic queues is limited by the number of CPU cores and the type of interface driver:

Interface Driver	Maximum Number of RX Queues
igb	2-16 (depends on the interface)
ixgbe	16
i40e	64
i40evf	4

Interface Driver	Maximum Number of RX Queues
ice	64
mlx5_core	60
ena	Configured automatically
virtio_net	Configured automatically
vmxnet3	Configured automatically

- In a Cluster, you must configure all the Cluster Members in the same way.

Note - Configuring Multi-Queue settings on an interface might temporarily interfere with traffic to that interface.
For example, configuring Multi-Queue settings on an interface while you are connected to that interface over SSH might close the SSH connection.

Deciding Whether to Enable the Multi-Queue

This section helps you decide if you can benefit from the Multi-Queue.

Best Practice - We recommend that you perform the steps below **before** you configure the Multi-Queue.

1. **Make sure that network interfaces support the Multi-Queue**

   Only network cards that use these drivers can support the Multi-Queue.

   See *"Multi-Queue Requirements and Limitations" on page 437*.

   Important - Before you upgrade these drivers, make sure that the latest version supports the Multi-Queue.

   Notes:
   - To view, which driver an interface uses, run this command in the Expert mode:
     - On the Security Gateway (each Cluster Member), run:
       ```
       ethtool -i <Name of Interface>
       ```
     - On the Scalable Platform Security Group, run:
       ```
       g_ethtool -i <Name of Interface>
       ```
   - When you install a new interface, you must run these two commands in the Expert mode:
     - On the Security Gateway (each Cluster Member), run:
       ```
       mq_mng --reconf
       reboot
       ```
     - On the Scalable Platform Security Group, run:
       ```
       g_all mq_mng --reconf
       g_reboot -a
       ```

2. **Make sure that SecureXL is enabled**

Step	Instructions
1	Connect to the command line on the Security Gateway / each Cluster Member / Scalable Platform Security Group.
2	Log in to the Gaia Clish, or the Expert mode. **Note** - On Scalable Platforms, you must run the applicable commands in Gaia gClish or the Expert mode of the applicable Security Group.

Step	Instructions
3	Get the SecureXL state (see *"fwaccel stat" on page 145*): ■ On a Security Gateway (each Cluster Member), run in Gaia Clish or the Expert mode: `fwaccel stat` ■ On a Scalable Platform Security Group, run in Gaia gClish: `fwaccel stat` ■ On a Scalable Platform Security Group, run in the Expert mode: `g_fwaccel stat`
4	Examine the **Status** column. **Example from a non-VSX Gateway** `[Expert@MyGW:0]# fwaccel stat` `+-----+---------+---------+----------+----------+` `|Id|Name |Status |Interfaces |Features |` `+-----+---------+---------+----------+----------+` `|0 |KPPAK |enabled |eth0,eth1 |Acceleration,Cryptography |` `| | | | | |` `| | | | |Crypto: Tunnel,UDPEncap,MD5, |` `| | | | |SHA1,3DES,DES,AES-128,AES-256,|` `| | | | |ESP,LinkSelection,DynamicVPN, |` `| | | | |NatTraversal,AES-XCBC,SHA256, |` `| | | | |SHA384,SHA512 |` `+-----+---------+---------+----------+----------+` `Accept Templates : enabled` `Drop Templates : disabled` `NAT Templates : enabled` `LightSpeed Accel : disabled` `[Expert@MyGW:0]#`

Step	Instructions
5	If the SecureXL is disabled, enable it (see *"fwaccel on" on page 133*): ■ On a Security Gateway (each Cluster Member), run in Gaia Clish or the Expert mode: `fwaccel on` ■ On a Scalable Platform Security Group, run in Gaia gClish: `fwaccel on` ■ On a Scalable Platform Security Group, run in the Expert mode: `g_fwaccel on`

3. **Examine the CPU roles allocation**

Step	Instructions
1	Connect to the command line on the Security Gateway / each Cluster Member / Scalable Platform Security Group.
2	Log in to the Gaia Clish, or the Expert mode. **Note** - On Scalable Platforms, you must run the applicable commands in Gaia gClish or the Expert mode of the applicable Security Group.

Step	Instructions
3	Get the list of CPU roles (see *"fw ctl affinity" on page 396*): ■ On a Security Gateway (each Cluster Member), run in Gaia Clish or the Expert mode: `fw ctl affinity -l [-a] [-v] [-r]` ■ On a Scalable Platform Security Group, run in Gaia gClish: `fw ctl affinity -l [-a] [-v] [-r]` ■ On a Scalable Platform Security Group, run in the Expert mode: `g_fw ctl affinity -l [-a] [-v] [-r]` **Example** CPU0 and CPU1 run the CoreXL SND instances: `[Expert@GW:0]# fw ctl affinity -l` `Mgmt: CPU 0` `eth1-04: CPU 1` `eth1-05: CPU 0` `eth1-06: CPU 1` `eth1-07: CPU 0` `fw_0: CPU 5` `fw_1: CPU 4` `fw_2: CPU 3` `fw_3: CPU 2` `[Expert@GW:0]#`

4. **Examine the CPU cores utilization**

Step	Instructions
1	Connect to the command line on the Security Gateway / each Cluster Member / Scalable Platform Security Group.
2	Log in to the Gaia Clish, or the Expert mode. **Note** - On Scalable Platforms, you must run the applicable commands in Gaia gClish or the Expert mode of the applicable Security Group.

Step	Instructions
3	Get the utilization of CPU cores: ■ On a Security Gateway (each Cluster Member), run in the Expert mode: `top` ■ On a Scalable Platform Security Group, run in the Gaia Clish: `top` ■ On a Scalable Platform Security Group, run in the Expert mode: `g_top`
4	Press 1 to show all the CPU cores. **Example** ■ CPU cores that run CoreXL SND instances (CPU0 and CPU1) are approximately 30% idle. ■ CPU cores that run CoreXL Firewall instances are approximately 70% idle. `top - 18:02:33 up 8 days, 1:18, 1 user, load average: 1.22, 1.38, 1.48` `Tasks: 137 total, 3 running, 134 sleeping, 0 stopped, 0 zombie` `Cpu0 : 2.0%us, 0.0%sy, 0.0%ni, 28.7%id, 5.9%wa, 0.0%hi, 63.4%si, 0.0%st` `Cpu1 : 0.0%us, 1.0%sy, 0.0%ni, 27.6%id, 0.0%wa, 0.0%hi, 71.4%si, 0.0%st` `Cpu2 : 2.0%us, 2.0%sy, 0.0%ni, 66.5%id, 0.0%wa, 4.0%hi, 25.5%si, 0.0%st` `Cpu3 : 1.0%us, 2.0%sy, 0.0%ni, 71.3%id, 0.0%wa, 0.0%hi, 25.7%si, 0.0%st` `Cpu4 : 5.0%us, 1.0%sy, 0.0%ni, 69.0%id, 0.0%wa, 0.0%hi, 25.0%si, 0.0%st` `Mem: 12224020k total, 70005820k used, 5218200k free, 273536k buffers` `Swap: 14707496k total, 0k used, 14707496k free, 484340k cached` `PID USER PR NI VIRT RES SHR S %CPU %MEM TIME+ COMMAND` `3301 root 15 0 0 0 0 R 31 0.0 747:04 [fw_worker_3]` `3326 root 15 0 0 0 0 R 29 0.0 593:35 [fw_worker_0]` `... ... ...`

5. **Decide if you can allocate more CPU cores to run the CoreXL SND instances**

To decide if you can allocate more CPU cores to run the CoreXL SND instances

If you have more active network interfaces than the CPU cores that run CoreXL SND instances, you can allocate more CPU cores to run more CoreXL SND instances.

We recommend to configure the Multi-Queue when:

a. CoreXL SND instances cause high CPU load (idle is less than 20%).

b. CoreXL Firewall instances cause low CPU load (idle is greater than 50%).

Note - You cannot assign more CPU cores to run CoreXL SND instances if you change interface IRQ affinity.

Multi-Queue Basic Configuration

In This Section:

You configure Multi-Queue on the command line in one of these shells:

- In the Expert mode
- In Gaia Clish
- In Gaia gClish on Scalable Platforms

Multi-Queue Configuration in the Expert mode

Description

The `mq_mng` utility shows and configures the Multi-Queue on supported interfaces.

Syntax

Important:

- In a Cluster, you must configure all the Cluster Members in the same way.
- You must run these commands in the Expert mode.
- On Scalable Platforms, you must connect to the applicable Security Group and use the "`g_all mq_mng`" commands.
- Change in the Multi-Queue mode can cause short packet loss.

- **To see the built-in help**

```
mq_mng {-h | --help}
```

- **To show the existing Multi-Queue configuration:**

```
mq_mng {-o | --show} [{-v | -vv}] [-a]
```

- **To configure the Multi-Queue for the specified driver:**

```
mq_mng {-s | --set-mode}
      auto
      manual
            {-i | --interface} <Names of Interfaces>
            {-c | --core} <IDs of CPU Cores>
      off
            [{-i | --interface} <Names of Interfaces>]
```

- **To apply the existing Multi-Queue policy:**

```
mq_mng {-r | --reconf}
```

Parameters

Parameter	Description
`-h | --help`	Shows built-in help.

Parameter	Description
`-o	--show`	Shows the existing Multi-Queue configuration. **Notes**: ■ `Available Queues` - Shows the maximum number of queues available for the given interface. ■ `Actual Queues` - Shows the current number of queues assigned to the given interface. ■ When CoreXL Dynamic-Balancing is disabled, the values of `Available Queues` and `Actual Queues` are the same for both automatic and manual Multi-Queue modes.
`-v	-vv`	Verbose output.
`-a`	Shows all interfaces in the output.
`-s	--set-mode`	Configures the Multi-Queue mode: ■ `auto` - Automatic mode (this is the default). Multi-Queue automatically configures the affinity of all supported interfaces to CPU cores that run CoreXL SND Instances. ■ `manual` - Manual mode. Administrator configures the affinity of interfaces to CPU cores that run CoreXL SND Instances. In this mode, you can specify interfaces, CPU cores, or both. ■ `off` - Disables the Multi-Queue on all or specified supported interfaces. **Important** - Change in the Multi-Queue mode can cause short packet loss.

Parameter	Description
	Notes: ■ To specify interfaces: • Use this syntax: `{-i | --interface} <Names of Interfaces>` • If you do not specify interfaces, then the configuration applies to all supported interfaces. • To specify a specific interface, enter its name (for example: `-i eth2`). • To specify several interfaces, enter their names separates with spaces (for example: `-i eth2 eth4`). ■ To specify CPU cores: • Use this syntax: `{-c | --core} <IDs of CPU Cores that run CoreXL SND Instances>` • To specify a specific CPU core, enter its ID number (for example: `-c 1`). • To specify several nonconsecutive CPU cores, enter their ID numbers separated with spaces (for example: `-c 1 3`) or commas (for example: `-c 1,3`). • To specify several consecutive CPU cores, enter their first and last ID numbers separated with a hyphen (for example: `-c 3-6`). ■ To see the current CoreXL affinity configuration, run the *"fw ctl affinity" on page 396* command (with applicable parameters). ■ To see the CoreXL Firewall Instances and which CPU cores they use, run the *"fw ctl multik stat" on page 391* command. ■ To see all available CPU cores, run: `cat /proc/cpuinfo | grep processor`
`-r | --reconf`	Applies the existing Multi-Queue policy.

Examples

Show the current Multi-Queue configuration on all interfaces

```
[Expert@Hostname:0]# mq_mng --show

Total 112 cores. Available for MQ 32 cores: Dynamic-Balancing enabled
i/f             driver          driver mode     state           mode (queues)   cores
                                                                actual/avail
-------------------------------------------------------------------------------------------
---
ethsBP1-01      mlx5_pci        DPDK            Up              Dynamic (10/32)
0,56,1,57,2,58,3,59,4,
                                                                                60
ethsBP1-02      mlx5_pci        DPDK            Up              Dynamic (10/32)
0,56,1,57,2,58,3,59,4,
                                                                                60
```

Show the current Multi-Queue verbose configuration on all interfaces

```
[Expert@Hostname:0]# mq_mng --show -v

Total 112 cores. Available for MQ 32 cores: Dynamic-Balancing enabled:
0,56,1,57,2,58,3,59,4,60,28,84,29,85,30,86,31,87,32,88,5,61,33,89,6,62,34,90,7,63,35,91
i/f             driver          driver mode     state           mode (queues)   cores
                                                                actual/avail
-------------------------------------------------------------------------------------------
---
ethsBP1-01      mlx5_pci        DPDK            Up              Dynamic (10/32)
0,56,1,57,2,58,3,59,4,
                                                                                60
ethsBP1-02      mlx5_pci        DPDK            Up              Dynamic (10/32)
0,56,1,57,2,58,3,59,4,
                                                                                60

core            interfaces      queue                   irq        rx packets      tx
packets
-------------------------------------------------------------------------------------------
---
0               ethsBP1-01      0                       None       18409592        3886402
                ethsBP1-02      0                       None       7495723         115534704
1               ethsBP1-01      2                       None       15926139        255839
                ethsBP1-02      2                       None       6633336         975
2               ethsBP1-01      4                       None       16105770        210370
                ethsBP1-02      4                       None       6003514         1096
3               ethsBP1-01      6                       None       19906359        39
                ethsBP1-02      6                       None       7304843         0
4               ethsBP1-01      8                       None       22296468        700
                ethsBP1-02      8                       None       7920505         0
56              ethsBP1-01      1                       None       18902160        166099
                ethsBP1-02      1                       None       6603803         47917
57              ethsBP1-01      3                       None       18508951        187803
                ethsBP1-02      3                       None       7513138         776
58              ethsBP1-01      5                       None       17152508        475
                ethsBP1-02      5                       None       6001529         0
59              ethsBP1-01      7                       None       18597442        55
                ethsBP1-02      7                       None       7320374         0
60              ethsBP1-01      9                       None       21868703        86
                ethsBP1-02      9                       None       7205076         0
```

Show the current Multi-Queue verbose configuration on the interface eth2

```
[Expert@Hostname:0]# mq_mng --show -v -i ethsBP1-01

Total 112 cores. Available for MQ 32 cores: Dynamic-Balancing enabled:
0,56,1,57,2,58,3,59,4,60,28,84,29,85,30,86,31,87,32,88,5,61,33,89,6,62,34,90,7,63,35,91
i/f             driver          driver mode     state           mode (queues)   cores
                                                                actual/avail
-------------------------------------------------------------------------------------------
---
BPEth0          mlx5_pci        Up              Dynamic (10/32) 0,56,1,57,2,58,3,59,4,
                                                                60

core            interfaces      queue                   irq        rx packets      tx
packets
-------------------------------------------------------------------------------------------
---
0               BPEth0          0                       None       18414645        3887732
1               BPEth0          2                       None       15929930        255999
2               BPEth0          4                       None       16110040        210483
3               BPEth0          6                       None       19911262        39
4               BPEth0          8                       None       22302258        700
56              BPEth0          1                       None       18906274        166165
57              BPEth0          3                       None       18513354        187850
58              BPEth0          5                       None       17156988        475
59              BPEth0          7                       None       18601956        55
60              BPEth0          9                       None       21874418        86
```

Set automatic Multi-Queue mode on all interfaces

```
mq_mng --set-mode auto
```

Set manual Multi-Queue mode on the interfaces eth1 and eth2 to CPU cores 0, 1, 2, 4, 5, and 6

```
mq_mng -s manual -i eth1 eth2 -c 0-2 4-6
```

Multi-Queue Configuration in Gaia Clish / Gaia gClish

Syntax

Important:

- In a Cluster, you must configure all the Cluster Members in the same way.
- You must run these commands in Gaia Clish / Gaia gClish.
- On Scalable Platforms, you must connect to the applicable Security Group.
- Change in the Multi-Queue mode can cause short packet loss.

- **To show the existing Multi-Queue configuration for the specified interface:**

```
show interface <Name of Interface> multi-queue [verbose]
```

- **To configure the Multi-Queue for the specified interface:**

```
set interface <Name of Interface> multi-queue
      auto
      manual core <IDs of CPU Cores that run CoreXL SND
Instances>
      off
```

Parameters

Parameter	Description
<Name of Interface>	Specifies the interface.
`verbose`	Verbose output that also includes: ■ IRQ numbers for traffic queues ■ Total number of RX and TX packets in traffic queues
`auto`	Configures the automatic Multi-Queue mode (this is the default). Multi-Queue automatically configures the affinity of the specified interface to CPU cores that run CoreXL SND Instances.

Parameter	Description
`manual core <IDs of CPU Cores>`	Configures the manual Multi-Queue mode. Administrator configures the affinity of the specified interface to CPU cores that run CoreXL SND Instances. **Notes:** ■ To specify a specific CPU core, enter its ID number (for example: `manual core 1`). ■ To specify several nonconsecutive CPU cores, enter their ID numbers separated with commas and without spaces (for example: `manual core 1,3`). ■ To specify several consecutive CPU cores, enter their first and last ID numbers separated with a hyphen (for example: `manual core 3-6`). ■ To see the current CoreXL affinity configuration, run the *"fw ctl affinity" on page 396* command (with applicable parameters). ■ To see the CoreXL Firewall Instances and which CPU cores they use, run the *"fw ctl multik stat" on page 391* command. ■ To see all available CPU cores, run: `cat /proc/cpuinfo | grep processor`
`off`	Disables the Multi-Queue on the specified interface.

Examples

Show Multi-Queue configuration on the interface eth2

```
MyGW> show interface eth2 multi-queue

Total 8 cores. Multiqueue 2 cores
i/f             type            state           config          cores
-------------------------------------------------------------------------
eth2            igb             Up              Auto            4,0

Note: The output does not include network interfaces that are currently in the down state.
MyGW>
```

Show Multi-Queue verbose configuration on the interface eth2

```
MyGW> show interface eth2 multi-queue verbose

Total 8 cores. Multiqueue 2 cores: 0,4
i/f             type            state           config          cores
-------------------------------------------------------------------------
eth2            igb             Up              Auto            4(62),0(79)

core            interfaces      queue               irq         rx packets      tx packets
-----------------------------------------------------------------------------------------
0               eth2            eth2-TxRx-1         79          212             80
4               eth2            eth2-TxRx-0         62          16232           18901
MyGW>
```

Set automatic Multi-Queue mode on the interface eth2

```
set interface eth2 multi-queue auto
```

Set manual Multi-Queue mode on the interface eth2 to CPU cores 0, 1, 2, 4, 5, and 6

```
set interface eth2 multi-queue manual core 0-2,4-6
```

Multi-Queue Special Scenarios and Configurations

This section provides instructions for configuring the Multi-Queue in special scenarios.

Default Number of Active RX Queues

Gateway Mode

Changing the number of CoreXL Firewall instances when the Multi-Queue is enabled on some, or all interfaces

For best performance, the Multi-Queue calculates the default number of active RX queues based on this formula:

```
Number of active RX queues = (Number of CPU cores) - (Number of
CoreXL Firewall instances)
```

This configuration is set automatically when you configure the Multi-Queue.

When you change the number of CoreXL Firewall instances, the number of active RX queues changes automatically, if it is not set manually.

VSX Mode

Changing the number of CPU cores, to which the FWK processes are assigned

For best performance, the Multi-Queue calculates the default number of active RX queues based on this formula:

```
Number of active RX queues = The lowest CPU ID, to which an FWK
process is assigned
```

Example

- The number of active RX queues is set to 2.
- This configuration is set automatically when you configure the Multi-Queue.
- It does not automatically update when you change the affinity of Virtual Systems.

```
[Expert@GW:0]# fw ctl affinity -l
Mgmt: CPU 0
eth1-05: CPU 0
eth1-06: CPU 1
VS_0 fwk: CPU 2 3 4 5
VS_1 fwk: CPU 2 3 4 5
[Expert@GW:0]#
```

Adding a Network Interface

When you add a network interface card to a Security Gateway / ClusterXL / Scalable Platform Security Group, the Multi-Queue configuration can change due to the way the operating system indexes the interfaces.

If you added a network interface card to a Security Gateway / ClusterXL / Scalable Platform Security Group, make sure to either configure the Multi-Queue again, or apply the existing Multi-Queue configuration:

- On a Security Gateway (each Cluster Member), run in the Expert mode:

  ```
  mq_mng --reconf
  ```

- On a Scalable Platform Security Group, run in the Expert mode:

  ```
  g_all mq_mng --reconf
  ```

Changing the Affinity of CoreXL Firewall instances

Best Practice - For best performance, we recommend that you do **not** assign both CoreXL SND instance and a CoreXL Firewall instance to the same CPU core.

Processing Packets that Arrive in the Wrong Order on an Interface that Works in Monitor Mode

Best Practice - If you enable Multi-Queue on an interface that works in Monitor Mode, then enable the Symmetric Hash for Receive-Side Scaling (RSS). This lets the corresponding interface drivers handle better packets that arrive in the wrong order (for example, TCP [SYN-ACK] that arrives before the TCP [SYN]). As a result, the same CPU core handles the applicable Client-to-Server and Server-to-Client packets.

Follow the instructions in sk101670 to download and run the special shell script `asym2sym.sh` on the Security Gateway / Cluster Members / Scalable Platform Security Group Members.

Multi-Queue Troubleshooting

Scenario	Explanation and next steps
After reboot, the wrong interfaces are configured for Multi-Queue.	This can happen after changing the physical interfaces on the Security Gateway / Scalable Platform Security Group. Use **one** of these options: ■ Apply the existing Multi-Queue configuration and reboot. • On the Security Gateway (each Cluster Member), run in Gaia Clish or the Expert mode: `mq_mng --reconf` `reboot` • On the Scalable Platform Security Group, run in Gaia gClish: `mq_mng --reconf` `reboot` • On the Scalable Platform Security Group, run in the Expert mode: `g_all mq_mng --reconf` `g_reboot -a` ■ Configure the Multi-Queue again.
When you change the status of interfaces, all the interface IRQs are assigned to CPU 0, or to all of the CPU cores.	This can happen when an interface status is changed to UP after the automatic affinity procedure runs (during each boot). Apply the existing Multi-Queue configuration. ■ On the Security Gateway (each Cluster Member), run in Gaia Clish or the Expert mode: `mq_mng --reconf` ■ On the Scalable Platform Security Group, run in Gaia gClish: `mq_mng --reconf` ■ On the Scalable Platform Security Group, run in the Expert mode: `g_all mq_mng --reconf`

Scenario	Explanation and next steps
In VSX mode, an **fwk** process runs on the same CPU core as some of the interface queues.	This can happen when the affinity of the Virtual System was manually changed but Multi-Queue was not reconfigured accordingly. Use **one** of these options: ■ Apply the existing Multi-Queue configuration and reboot. • On the Security Gateway (each Cluster Member), run in Gaia Clish or the Expert mode: `mq_mng --reconf` `reboot` • On the Scalable Platform Security Group, run in Gaia gClish: `mq_mng --reconf` `reboot` • On the Scalable Platform Security Group, run in the Expert mode: `g_all mq_mng --reconf` `g_reboot -a` ■ Configure the number of active RX queues manually
In Gateway mode, after you change the number of CoreXL Firewall instances, the Multi-Queue is disabled on all interfaces.	When you change the number of CoreXL Firewall instances, the number of active RX queues automatically changes based on this formula: `Active RX queues = (Number of CPU cores) - (Number of CoreXL Firewall instances)` If the difference between the number of CPU cores and the number of CoreXL Firewall instances is 1, Multi-Queue is disabled.

CPView

Overview of CPView

Description

CPView is a text based built-in utility on a Check Point computer.

CPView Utility shows statistical data that contain both general system information (CPU, Memory, Disk space) and information for different Software Blades (only on a Security Gateway / ClusterXL / Scalable Platform Security Group).

The CPView continuously updates the data in easy to access views.

On a Security Gateway / ClusterXL / Scalable Platform Security Group, you can use this statistical data to monitor the performance.

For more information, see sk101878.

Syntax

```
cpview --help
```

CPView User Interface

The CPView user interface has three sections:

Section	Description
Header	This view shows the time the statistics in the third view are collected. It updates when you refresh the statistics.
Navigation	This menu bar is interactive. Move between menus with the arrow keys and mouse. A menu can have sub-menus and they show under the menu bar.
View	This view shows the statistics collected in that view. These statistics update at the refresh rate.

Using CPView

Use these keys to navigate the CPView:

Key	Description
Arrow keys	Moves between menus and views. Scrolls in a view.
Home	Returns to the **Overview** view.
Enter	Changes to the **View Mode**. On a menu with sub-menus, the **Enter** key moves you to the lowest level sub-menu.
Esc	Returns to the **Menu Mode**.
Q	Quits CPView.

Use these keys to change CPView interface options:

Key	Description
R	Opens a window where you can change the refresh rate. The default refresh rate is 2 seconds.
W	Changes between wide and normal display modes. In wide mode, CPView fits the screen horizontally.
S	Manually sets the number of rows or columns.
M	Switches on/off the mouse.
P	Pauses and resumes the collection of statistics.

Use these keys to save statistics, show help, and refresh statistics:

Key	Description
C	Saves the current page to a file. The file name format is: `cpview_<ID of the cpview process>.cap<Number of the capture>`
H	Shows a tooltip with CPView options.
Space bar	Immediately refreshes the statistics.

CPU Spike Detective

The CPU Spike Detective is a tool that monitors the CPU utilization and saves information about the CPU utilization spikes it detects.

This tool does **not** impact the performance.

Use these commands in Gaia Clish:

Note - On Scalable Platforms, you must use Gaia gClish.

`show spike-detective[ESC][ESC]`
`set spike-detective[ESC][ESC]`
`delete spike-detective[ESC][ESC]`

For more information, see sk166454.

HyperFlow

Overview

Elephant flows are large (in total number of bytes) continuous connections that the TCP or UDP establishes.

For example, a download of a large file (such as a Linux ISO file) over the HTTP, HTTPS, FTP, or NFS protocol.

These large continuous connections consume the network capacity significantly in comparison to other types of data sessions.

Without the HyperFlow feature, a Security Gateway uses only one CPU core (one CoreXL Firewall instance) to inspect one elephant connection. In addition, traffic throughput decreases gradually as the CPU utilization increases on the Security Gateway.

The HyperFlow feature on Security Gateways R81.20 and higher handles such elephant connections on more than one CPU core in parallel.

The HyperFlow feature breaks the whole inspection task into smaller tasks and dispatches these smaller tasks to the available CPU cores:

The tasks without the HyperFlow	The tasks with the HyperFlow
1. Packet retrieval 2. Inbound Streaming 3. Protocol parsers 4. Context Management Interface / Infrastructure (CMI) 5. Pattern Match (PM) and Hash (MD5, SHA) 6. Software Blade logic 7. Outbound Streaming 8. Routing 9. Packet transmission	1. Inbound processing in CoreXL Firewall: a. Packet retrieval b. Inbound Streaming c. Protocol parsers d. Context Management Interface / Infrastructure (CMI) 2. Internal PPE processing (on many CPU cores): a. Pattern Match (PM) and Hash (MD5, SHA) b. Packet transmission 3. Outbound processing in CoreXL Firewall: a. Software Blade logic b. Outbound Streaming c. Routing

As a result, the HyperFlow feature:

- Increases throughput of elephant connections when Threat Prevention Software Blades are enabled (the Security Gateway takes less time to inspect elephant connections).

  This is possible only if the network infrastructure is not a "bottleneck".

- Automatically detects and dynamically allocates the CPU cores between main tasks on a Security Gateway.
- Improves response time from the CoreXL FWK processes while they inspects elephant connections (the idle time of the corresponding CPU cores increases).

Important:

- By default, the HyperFlow feature is enabled on Check Point Appliances that meet the requirements.
- By design, the HyperFlow feature works only in the User Space Firewall (USFW).
- By design, the HyperFlow feature engages only when needed, and when the total CPU load allows it.
  The total throughput has priority over elephant connections.

Notes:

- By design, a manual allocation of CPU cores is not necessary. Therefore, it is not possible.
  You can configure thresholds to control when HyperFlow is active or passive.
- By default, HyperFlow works in the standby mode.
  HyperFlow is triggered (becomes active) when a heavy connection is detected.
  HyperFlow becomes passive when the heavy connection is closed.

For additional information, see sk178070.

Requirements

1. Check Point Appliance models with at least 8 CPU logical cores.

   For the list of supported models, see sk178070.

2. Firewall in User Mode (USFW). See sk167052.

3. Enable the CoreXL Dynamic Balancing (see *"Dynamic Balancing of CoreXL Instances" on page 341*):

   ```
   dynamic_split -o enable
   ```

4. Configure SecureXL to work in Kernel Mode (KPPAK) (see *"Configuring SecureXL" on page 20*).

5. Enable the applicable Software Blades from one of these categories:

   - NGFW
   - NGTP
   - NGTP with HTTPS Inspection
   - NGTX

Glossary

Term	Description
CoreXL_FW	A CoreXL Firewall instance that handles the traffic concurrently. Each CoreXL Firewall instance is independent and replicated multiple times (see *"CoreXL" on page 314*). Each replicated CoreXL Firewall instance runs on one processing CPU core. **Note** - CPView shows this string on the **CPU** > **Overview** > **Host** page > in the column **Type**.
CoreXL_SND	A CoreXL Secure Network Distributor (SND) responsible for: ■ Processing incoming traffic from the network interfaces ■ Securely accelerating authorized packets (when SecureXL is running) ■ Distributing non-accelerated packets among FW instances **Note** - CPView shows this string on the **CPU** > **Overview** > **Host** page > in the column **Type**.
CoreXL_FW_RESERVED	A logical sibling of a CoreXL Firewall instance (FW worker) that handles a heavy connection. When a logical CPU core is utilized at a high level because it handles heavy connections, its logical sibling can be stopped to decrease its utilization of resources from the physical CPU core. Chain on events: 1. CoreXL assigns a CoreXL Firewall instance to inspect an elephant connection. 2. This CoreXL Firewall instance runs on a logical CPU core of a physical CPU. 3. To improve the internal performance of the physical CPU, the CoreXL Dynamic Balancing feature can stop the CoreXL Firewall instances on the sibling logical CPU cores of the original logical CPU core.
PPE	Parallel Processing Engine architecture. This is the HyperFlow dynamic infrastructure that allocates CPU cores as required to increase the throughput of Elephant connections. This is the thread that polls the interface queues and retrieves packets. Also known as Dual Mode Job Dispatcher (DMD).

Term	Description
PPE_MGR	Parallel Processing Engine Manager. Receives packet payload and dispatches jobs to **PPE**. Works on a complete physical CPU core (if there are 32 or more CPU cores). **Note** - CPView shows this string on the **CPU** > **Overview** > **Host** page > in the column **Type**. This string appears in CPView only when the HyperFlow is enabled, and an elephant connection passes through the Security Gateway.
PPE_MGR_ RESERVED	A sibling of **PPE_MGR**. Because **PPE_MGR** works on a complete physical CPU core (if there are 32 or more CPU cores), it is not possible to use other logical CPU cores on that physical CPU core. The logical siblings have the status "**PPE_MGR_RESERVED**" **Note** - CPView shows this string on the **CPU** > **Overview** > **Host** page > in the column **Type**. This string appears in CPView only when the HyperFlow is enabled, and an elephant connection passes through the Security Gateway.
PPE_WT	Parallel Processing Engine Worker Thread. Receives packet handling jobs from the Parallel Processing Engine Manager. Dispatches jobs to Worker Threads (WTs.) Works on a complete physical CPU core (if there are 32 or more CPU cores). **Note** - CPView shows this string on the **CPU** > **Overview** > **Host** page > in the column **Type**. This string appears in CPView only when the HyperFlow is enabled, and an elephant connection passes through the Security Gateway.
DPDK	Data Plane Development Kit.
PMD	Poll Mode Driver.
WT	Worker Thread. The thread that executes the packet handling logic. In plural: **WTs**

Syntax

Important:

- You must run this command in the Expert mode.
- In a Cluster, you must configure all the Cluster Members in the same way.
- On Scalable Platforms, you must connect to the Gaia Portal of the applicable Security Group.
- On Scalable Platforms, use the command "`g_connection_pipelining`" with the same parameters.

```
connection_pipelining
      advanced
            allow_accelerated_pipeline
            async
            default
            prevent_accelerated_pipeline
            sleep
            sync
            wake_up
      on
      off
      heaviest_conn
      pipelined
      status
```

Parameters

Parameter	Description
No Parameters	Shows the built-in help.
`connection_pipelining`	Must enter this command only on Security Gateways other than Scalable Platforms.
`g_connection_pipelining`	Must enter this command only on Scalable Platforms.
`advanced`	Shows the advanced options.

Parameter	Description
`allow_accelerated_pipeline`	Allows new connections to be opened as accelerated pipeline connections - the Security Gateway uses the new asynchronous parsers for connections. This is the default. **Notes:** ■ This command applies only to elephant connections that opened after you run this command. ■ When you run this command, the Security Gateway deletes all SecureXL Connection Templates with this command: `fw tab -t cphwd_tmpl -x -y` ■ This command does **not** require a reboot.
`async`	Configures the asynchronous flow mode (this is the default). In this mode, CoreXL Firewall instances send jobs to the PPE. **Notes:** ■ This command applies to existing elephant connections. ■ This command does **not** require a reboot.
`default`	Restores default settings
`heaviest_conn`	Shows the statistics for the heaviest connection with the maximum duration (number of packets and bytes).
`off`	Disables the feature. **Important** - This change requires a reboot.
`on`	Enables the feature. This is the default (on Check Point Appliances that meet the requirements). **Important** - This change requires a reboot.
`pipelined`	Shows the accelerated elephant connections in the pipeline.

Parameter	Description
`prevent_accelerated_pipeline`	Prevents new connections from being opened as accelerated pipeline connections. In this mode, the Security Gateway uses the legacy parsers for connections. **Notes:** ■ Use this command only to troubleshoot issues with elephant connections. ■ This command applies only to elephant connections that opened after you run this command. ■ This command does **not** require a reboot. ■ When you run this command, the Security Gateway deletes all SecureXL Connection Templates with this command: `fw tab -t cphwd_tmpl -x -y`
`sleep`	Configures the Job Dispatcher (PPE) and Working Threads (WTs) to sleep. **Notes:** ■ Use this command only to troubleshoot issues with elephant connections, if the issue persists after you configured the synchronous flow mode. ■ This command does **not** require a reboot.
`status`	Shows the status and configuration of the feature. The output shows these lines with the applicable values: ■ `Status of connection pipelining: <Status>` ■ `Flow mode: <Mode>` ■ `Status of PPE_MGR and PPE_WT: <Status>` ■ `Status of accelerated pipeline: <Status>`
`sync`	Configures the synchronous flow mode. In this mode, CoreXL Firewall instances do **not** send jobs to the PPE. **Notes:** ■ Use this mode only to troubleshoot issues with elephant connections. ■ This command applies to existing elephant connections. ■ This command does **not** require a reboot.
`wake_up`	Wakes up the Job Dispatcher (PPE) and Working Threads (WTs) from sleep. The PPE gets CPU cores available for allocation.

Monitoring in CPView

You can monitor how the HyperFlow performs on the Security Gateway.

Viewing if HyperFlow is enabled or disabled

1. Go to **CPU** > **Advanced**
2. If the tab **Hyperflow** appears, it means HyperFlow is enabled

Viewing if HyperFlow is active or sleeping

1. Go to **CPU** > **Advanced** > **Hyperflow** > **Overview**
2. Examine the field **PPE_MGR state**

Viewing the allocation of CPU cores

1. Go to **CPU** > **Overview** > **Host**
2. Examine the last section **CPU** (pay attention to the column **Idle**)

Viewing the CPU load from specific elephant connections

1. On the Security Gateway, examine the elephant connections in the past 24 hours and which CoreXL Firewall instance inspects them (see the "`[fw_<Number>]`" in the beginning of each line):

   ```
   fw ctl multik print_heavy_conn
   ```

   Example:

   ```
   ;[fw_5]: Conn: 192.168.10.20:60478 -> 172.30.40.50:80 IPP:
   6; Instance load: 63%; Connection instance load: 99%;
    ...<truncated>...
   ```

2. In CPView, go to **CPU** > **Top-Connections** > **Instances<X>-<Y>** (for example, **Instances0-5**) > **Instance<Z>** (for example, **Instance5**)

3. In the last section **Top Connections**, examine the columns "**% out of CPU**" and "**% out of WT CPU**"

   Example (this is only the applicable part of the output):

   ```
   Top Connections

   Connection                                    Protocol    % out
   of CPU % out of WT CPU
   192.168.10.20:60478 -> 172.30.40.50:80     TCP:http
    70.81%        48.97%
   ```

Viewing the status of the PPE

1. Go to **Advanced** > **HyperFlow** > **Overview** > **PPE_0**
2. Click the applicable tab
3. Examine these sections:

Section	Gauge	Description
PPE overview	**PPE state**	Shows the state of the feature - **Active** or **Asleep**
Pipeline status	**Free**	The number of free slots in the pipeline (maximum is 320)
	Active	Slots that PPE is currently using
	Job pending	Slots whose execution depends on another job
	Slot pending	Jobs whose execution depends on an available slot in the pipeline
Overload indicators	**No pipeline entry**	PPE_MGR had no free slots in the pipeline
	WT slot unavailable	The WTs have no available slots

Viewing the PPE messages from the Firewall

1. Go to **Advanced** > **HyperFlow** > **Firewall-messages**
2. Examine these sections:

Section	Gauge	Description
Sessions and data	**New session**	Number of new opened sessions
	Update session	Number of updated session (for example, because of policy installation)
	End session	Number of sessions that ended
	Current session	Number of currently opened sessions
	Data	Number of received data buffers
Errors		Various internal PPE errors
Messages received in a single read loop (Histogram)		In each loop, the PPE_MGR can read up to 32 items from the receiving queue. This section shows a histogram of the times the PPE_MGR read items from the queue and how many were "waiting" in the queue (up to 32)

Viewing the number of jobs and their execution time

1. Go to **Advanced** > **HyperFlow** > **Jobs** > **PPE**
2. Examine these sections:
   - **Sent to firewall (count)**
   - **Average execution time (Cycles)**

Viewing the queue utilization

1. Go to **Advanced** > **HyperFlow** > **Comm** > **PPE**
2. Examine these sections:
   - **Enqueue**
   - **Dequeue**

Viewing the number and state of Worker Threads

1. Go to **Advanced** > **HyperFlow** > **WT**
2. Examine these rows:
   - **Number of WTs**
   - **Worker ID State**

Limitations

- Firewall in Kernel Mode (KFW) is not supported.
- Open Servers and Virtual Machines are not supported.
- Cloud platforms are not supported.
- HyperFlow is not supported when the Security Gateway / Cluster is configured to work as an HTTP/HTTPS Proxy.
- When an elephant connection triggers the HyperFlow, the outputs of the "`top`" and "`ps`" commands may show that the HyperFlow user space processes consume some CPU cores at 99-100%.

  This happens because HyperFlow is constantly polling its queues to handle the incoming jobs. After the elephant connection closes, the output of these commands shows that the user space "`us`" consumption goes back to regular levels because HyperFlow stops processing the jobs.

  To see the actual load on the CPU, use *"CPView" on page 459* (`CPU` > `Overview` > `Host`), SNMP, or SmartConsole.

  This does **not** trigger inspection bypass because of a high CPU load.
- When the HyperFlow feature is enabled, the speed of an elephant flow may not increase in these cases:
  - This is a complex connection that requires special parsers – FTP, SCP, VoIP
  - This connection undergoes the Content Awareness or a Data Loss Prevention policy
  - This connection has dynamic content and goes through the Pattern Matcher 1st tier
  - Strict Hold is enabled (the `$FWDIR/conf/malware_config` file on the Security Gateway contains "`strict_hold_enable=1`")
  - This connection is HTTPS and the HTTPS Inspection is disabled
- When you enable only the Firewall Software Blade in the Security Gateway object, the HyperFlow feature does not improve the performance.

  This is because SecureXL accelerates connections that do **not** undergo the inspection, while the HyperFlow accelerates connections that undergo the inspection by various Software Blades.

Troubleshooting

- Log files (the main file is rotated every 10 megabytes):
  - `$FWDIR/log/connection_pipelining.elg`
  - `$FWDIR/log/dmd.elg`
  - `$FWDIR/log/dmd_controller.elg`
  - `$FWDIR/log/dsd.elg`
- Internal configuration files (you must **not** edit these files):
  - `$FWDIR/conf/connection_pipelining.conf`
  - `$FWDIR/conf/connection_pipelining_params.conf`

For additional information, see sk178070.

Command Line Reference

See the *R82 CLI Reference Guide*.

Working with Kernel Parameters

See the *R82 Quantum Security Gateway Guide* > Chapter "**Working with Kernel Parameters**".

Kernel Debug

See the *R82 Quantum Security Gateway Guide* > Chapter "**Kernel Debug**".

Glossary

A

Accelerated Path
Packet flow on the Host appliance, when the packet is completely handled by the SecureXL device. It is processed and forwarded to the network.

Affinity
The assignment of a specified CoreXL Firewall instance, VSX Virtual System, interface, user space process, or IRQ to one or more specified CPU cores.

Anti-Bot
Check Point Software Blade on a Security Gateway that blocks botnet behavior and communication to Command and Control (C&C) centers. Acronyms: AB, ABOT.

Anti-Spam
Check Point Software Blade on a Security Gateway that provides comprehensive protection for email inspection. Synonym: Anti-Spam & Email Security. Acronyms: AS, ASPAM.

Anti-Virus
Check Point Software Blade on a Security Gateway that uses real-time virus signatures and anomaly-based protections from ThreatCloud to detect and block malware at the Security Gateway before users are affected. Acronym: AV.

Application Control
Check Point Software Blade on a Security Gateway that allows granular control over specific web-enabled applications by using deep packet inspection. Acronym: APPI.

Audit Log
Log that contains administrator actions on a Management Server (login and logout, creation or modification of an object, installation of a policy, and so on).

B

Breakout Cable
An optical fiber cable that contains several jacketed simplex optical fibers that are packaged together inside an outer jacket. Synonyms: Fanout cable, Fan-Out cable, Splitter cable.

Bridge Mode
Security Gateway or Virtual System that works as a Layer 2 bridge device for easy deployment in an existing topology.

C

Chassis Management Module
On Scalable Chassis - a hardware component that controls and monitors 60000 / 40000 Appliance (Chassis) operation such as, fan speed, Chassis and module temperature, and component hot-swapping. Acronym: CMM.

Cluster
Two or more Security Gateways that work together in a redundant configuration - High Availability, or Load Sharing.

Cluster Member
Security Gateway that is part of a cluster.

Compliance
Check Point Software Blade on a Management Server to view and apply the Security Best Practices to the managed Security Gateways. This Software Blade includes a library of Check Point-defined Security Best Practices to use as a baseline for good Security Gateway and Policy configuration.

Content Awareness
Check Point Software Blade on a Security Gateway that provides data visibility and enforcement. Acronym: CTNT.

CoreXL
Performance-enhancing technology for Security Gateways on multi-core processing platforms. Multiple Check Point Firewall instances are running in parallel on multiple CPU cores.

CoreXL Dynamic Dispatcher
Improved CoreXL SND feature. Part of CoreXL that distributes packets between CoreXL Firewall instances. Traffic distribution between CoreXL Firewall instances is dynamically based on the utilization of CPU cores, on which the CoreXL Firewall instances are running. The dynamic decision is made for first packets of connections, by assigning each of the CoreXL Firewall instances a rank, and selecting the CoreXL Firewall instance with the lowest rank. The rank for each CoreXL Firewall instance is calculated according to its CPU utilization. The higher the CPU utilization, the higher the CoreXL Firewall instance's rank is, hence this CoreXL Firewall instance is less likely to be selected by the CoreXL SND.

CoreXL Firewall Instance
On a Security Gateway with CoreXL enabled, the Firewall kernel is copied multiple times. Each replicated copy, or firewall instance, runs on one processing CPU core. These firewall instances handle traffic at the same time, and each firewall instance is a complete and independent firewall inspection kernel. Synonym: CoreXL FW Instance.

CoreXL SND
Secure Network Distributer. Part of CoreXL that is responsible for: Processing incoming traffic from the network interfaces; Securely accelerating authorized packets (if SecureXL is enabled); Distributing non-accelerated packets between Firewall kernel instances (SND maintains global dispatching table, which maps connections that were assigned to CoreXL Firewall instances). Traffic distribution between CoreXL Firewall instances is statically based on Source IP addresses, Destination IP addresses, and the IP 'Protocol' type. The CoreXL SND does not really "touch" packets. The decision to stick to a particular FWK daemon is done at the first packet of connection on a very high level, before anything else. Depending on the SecureXL settings, and in most of the cases, the SecureXL can be offloading decryption calculations. However, in some other cases, such as with Route-Based VPN, it is done by FWK daemon.

CPUSE
Check Point Upgrade Service Engine for Gaia Operating System. With CPUSE, you can automatically update Check Point products for the Gaia OS, and the Gaia OS itself.

D

DAC Cable
Direct Attach Copper cable. A form of the high-speed shielded twinax copper cable with pluggable transceivers on both ends. Used to connect to network devices (switches, routers, or servers).

DAIP Gateway
Dynamically Assigned IP (DAIP) Security Gateway is a Security Gateway, on which the IP address of the external interface is assigned dynamically by the ISP.

Data Loss Prevention
Check Point Software Blade on a Security Gateway that detects and prevents the unauthorized transmission of confidential information outside the organization. Acronym: DLP.

Data Type
Classification of data in a Check Point Security Policy for the Content Awareness Software Blade.

Distributed Deployment
Configuration in which the Check Point Security Gateway and the Security Management Server products are installed on different computers.

Downlink Ports
Interfaces on the Quantum Maestro Orchestrator used to connect to Check Point Security Appliances. You use DAC cables, Fiber cables (with transceivers), or Breakout cables to connect between the Downlink ports and Security Appliances. The Check Point Management traffic (policy, logs, synchronization, and so on) co-exists with the data (user) traffic on the Downlink ports. Bandwidth is guaranteed for the Check Point Management traffic (portion of the downlink bandwidth). These ports form the system backplane (management, data plane, synchronization).

Dynamic Object
Special object type, whose IP address is not known in advance. The Security Gateway resolves the IP address of this object in real time.

E

Endpoint Policy Management
Check Point Software Blade on a Management Server to manage an on-premises Harmony Endpoint Security environment.

Expert Mode
The name of the elevated command line shell that gives full system root permissions in the Check Point Gaia operating system.

F

F2F
Denotes non-VPN connections that SecureXL forwarded to firewall. See "Firewall Path".

Firewall Path
Packet flow on the Host Security Appliance, when the SecureXL device is unable to process the packet. The packet is passed to the CoreXL layer and then to one of the CoreXL Firewall instances for full processing. This path also processes all packets when SecureXL is disabled. Synonym: Slow Path.

G

Gaia
Check Point security operating system that combines the strengths of both SecurePlatform and IPSO operating systems.

Gaia Clish
The name of the default command line shell in Check Point Gaia operating system. This is a restricted shell (role-based administration controls the number of commands available in the shell).

Gaia gClish
The name of the global command line shell in Check Point Gaia operating system for Security Appliances connected to Check Point Quantum Maestro Orchestrators and for Security Gateway Modules on Scalable Chassis. Commands you run in this shell apply to all Security Gateway Module / Security Appliances in the Security Group.

Gaia Portal
Web interface for the Check Point Gaia operating system.

H

Hotfix
Software package installed on top of the current software version to fix a wrong or undesired behavior, and to add a new behavior.

HTTPS Inspection
Feature on a Security Gateway that inspects traffic encrypted by the Secure Sockets Layer (SSL) protocol for malware or suspicious patterns. Synonym: SSL Inspection. Acronyms: HTTPSI, HTTPSi.

HyperSync
Check Point patented technology that makes sure that active connections are only synchronized to backup Security Appliances in the Security Group. HyperSync makes sure each connection flow has a backup within the Security Group.

I

ICA
Internal Certificate Authority. A component on Check Point Management Server that issues certificates for authentication.

Identity Awareness
Check Point Software Blade on a Security Gateway that enforces network access and audits data based on network location, the identity of the user, and the identity of the computer. Acronym: IDA.

Identity Logging
Check Point Software Blade on a Management Server to view Identity Logs from the managed Security Gateways with enabled Identity Awareness Software Blade.

Internal Network
Computers and resources protected by the Firewall and accessed by authenticated users.

IPS
Check Point Software Blade on a Security Gateway that inspects and analyzes packets and data for numerous types of risks (Intrusion Prevention System).

IPsec VPN
Check Point Software Blade on a Security Gateway that provides a Site to Site VPN and Remote Access VPN access.

IRQ Affinity
A state of binding an IRQ to one or more CPU cores.

J

Jumbo Hotfix Accumulator
Collection of hotfixes combined into a single package. Acronyms: JHA, JHF, JHFA.

K

Kerberos
An authentication server for Microsoft Windows Active Directory Federation Services (ADFS).

L

Log Server
Dedicated Check Point server that runs Check Point software to store and process logs.

Logging & Status
Check Point Software Blade on a Management Server to view Security Logs from the managed Security Gateways.

M

Maestro Orchestrator
A scalable Network Security System that connects multiple Check Point Security Appliances into a unified system. Synonyms: Orchestrator, Quantum Maestro Orchestrator, Maestro Hyperscale Orchestrator. Acronym: MHO.

Management Interface
(1) Interface on a Gaia Security Gateway or Cluster member, through which Management Server connects to the Security Gateway or Cluster member. (2) Interface on Gaia computer, through which users connect to Gaia Portal or CLI.

Management Server
Check Point Single-Domain Security Management Server or a Multi-Domain Security Management Server.

Manual NAT Rules
Manual configuration of NAT rules by the administrator of the Check Point Management Server.

Medium Path
Packet flow on the Host Security Appliance, when the packet is handled by the SecureXL device. The CoreXL layer passes the packet to one of the CoreXL Firewall instances to process it. Even when CoreXL is disabled, the SecureXL uses the CoreXL infrastructure to send the packet to the single CoreXL Firewall instance that still functions. When the Medium Path is available, the SecureXL fully accelerates the TCP handshake. Rule Base match is achieved for the first packet through an existing connection acceleration template. The SecureXL also fully accelerates the TCP [SYN-ACK] and TCP [ACK] packets. However, once data starts to flow, to stream it for Content Inspection, an FWK instance now handles the packets. The SecureXL sends all packets that contain data to FWK for data extraction in order to build the data stream. Only the SecureXL handles the TCP [RST], TCP [FIN] and TCP [FIN-ACK] packets, because they do not contain data that needs to be streamed. The Medium Path is available only when CoreXL is enabled. Exceptions are: IPS (some protections); VPN (in some configurations); Application Control; Content Awareness; Anti-Virus; Anti-Bot; HTTPS Inspection; Proxy mode; Mobile Access; VoIP; Web Portals. Synonym: PXL.

Mobile Access
Check Point Software Blade on a Security Gateway that provides a Remote Access VPN access for managed and unmanaged clients. Acronym: MAB.

Multi-Domain Log Server
Dedicated Check Point server that runs Check Point software to store and process logs in a Multi-Domain Security Management environment. The Multi-Domain Log Server consists of Domain Log Servers that store and process logs from Security Gateways that are managed by the corresponding Domain Management Servers. Acronym: MDLS.

Multi-Domain Server
Dedicated Check Point server that runs Check Point software to host virtual Security Management Servers called Domain Management Servers. Synonym: Multi-Domain Security Management Server. Acronym: MDS.

Multi-Queue
An acceleration feature on Security Gateway that configures more than one traffic queue for each network interface. Multi-Queue assigns more than one receive packet queue (RX Queue) and more than one transmit packet queue (TX Queue) to an interface. Multi-Queue is applicable only if SecureXL is enabled (this is the default). Acronym: MQ.

N

Network Object
Logical object that represents different parts of corporate topology - computers, IP addresses, traffic protocols, and so on. Administrators use these objects in Security Policies.

Network Policy Management
Check Point Software Blade on a Management Server to manage an on-premises environment with an Access Control and Threat Prevention policies.

O

Open Server
Physical computer manufactured and distributed by a company, other than Check Point.

Orchestrator
See "Maestro Orchestrator".

P

Power Entry Module
Hardware component that supplies DC power with EMC filtering and over-current protection on Scalable Chassis. Acronym: PEM.

Power Supply Unit
Hardware component that supplies AC power with filtering and over-current protection. Acronym: PSU.

Provisioning
Check Point Software Blade on a Management Server that manages large-scale deployments of Check Point Security Gateways using configuration profiles. Synonyms: SmartProvisioning, SmartLSM, Large-Scale Management, LSM.

PSL
Passive Streaming Library. Packets may arrive at Security Gateway out of order, or may be legitimate retransmissions of packets that have not yet received an acknowledgment. In some cases, a retransmission may also be a deliberate attempt to evade IPS detection by sending the malicious payload in the retransmission. Security Gateway ensures that only valid packets are allowed to proceed to destinations. It does this with the Passive Streaming Library (PSL) technology. (1) The PSL is an infrastructure layer, which provides stream reassembly for TCP connections. (2) The Security Gateway makes sure that TCP data seen by the destination system is the same as seen by code above PSL. (3) The PSL handles packet reordering, congestion, and is responsible for various security aspects of the TCP layer, such as handling payload overlaps, some DoS attacks, and others. (4) The PSL is capable of receiving packets from the Firewall chain and from the SecureXL. (5) The PSL serves as a middleman between the various security applications and the network packets. It provides the applications with a coherent stream of data to work with, free of various network problems or attacks. (6) The PSL infrastructure is wrapped with well-defined APIs called the Unified Streaming APIs, which are used by the applications to register and access streamed data.

PSLXL
Technology name for combination of SecureXL and PSL (Passive Streaming Library) in versions R80.20 and higher. In versions R80.10 and lower, this technology was called PXL (PacketXL).

Q

QoS
Check Point Software Blade on a Security Gateway that provides policy-based traffic bandwidth management to prioritize business-critical traffic and guarantee bandwidth and control latency.

R

Rule
Set of traffic parameters and other conditions in a Rule Base (Security Policy) that cause specified actions to be taken for a communication session.

Rule Base
All rules configured in a given Security Policy. Synonym: Rulebase.

S

Scalable Chassis
The container that contains the all the components of a 60000 / 40000 Appliance. Synonym: Chassis.

SecureXL
Check Point product on a Security Gateway that accelerates IPv4 and IPv6 traffic that passes through a Security Gateway.

Security Gateway
Dedicated Check Point server that runs Check Point software to inspect traffic and enforce Security Policies for connected network resources.

Security Gateway Module
On Scalable Chassis - a hardware component on a 60000 / 40000 Appliance (Chassis) that operates as a physical Security Gateway. A Chassis contains many Security Gateway Modules that work together as a single, high performance Security Gateway or VSX Gateway. In Maestro - a role of a Security Appliance. Part of the Security Group that contains the assigned Security Appliances. A Security Appliance in a Security Group has one IPv4 address and represents all assigned Security Appliances as one entity. Acronym: SGM.

Security Group
A logical group of Security Appliances (in Maestro) / Security Gateway Modules (on Scalable Chassis) that provides Active/Active cluster functionality. A Security Group can contain one or more Security Appliances / Security Gateway Modules. Security Groups work separately and independently from each other. To the production networks, a Security Group appears a single Security Gateway. In Maestro, each Security Group contains: (A) Applicable Uplink ports, to which your production networks are connected; (B) Security Appliances (the Quantum Maestro Orchestrator determines the applicable Downlink ports automatically); (C) Applicable management port, to which the Check Point Management Server is connected.

Security Group Member
Member of a Security Group in ElasticXL Cluster, Maestro, and Scalable Chassis. Acronym: SGM.

Security Management Server
Dedicated Check Point server that runs Check Point software to manage the objects and policies in a Check Point environment within a single management Domain. Synonym: Single-Domain Security Management Server.

Security Policy
Collection of rules that control network traffic and enforce organization guidelines for data protection and access to resources with packet inspection.

Security Switch Module
On Scalable Chassis - a hardware component on a 60000 / 40000 Appliance (Chassis) that manages the flow of network traffic to and from the Security Gateway Module in the Chassis. In Maestro - a role of the Quantum Maestro Orchestrator that manages the flow of network traffic to and from the Security Groups. Acronym: SSM.

SGM
In Maestro - a role of a Security Appliance. Part of the Security Group that contains the assigned Security Appliances. A Security Appliance in a Security Group has one IPv4 address and represents all assigned Security Appliances as one entity. For Scalable Chassis - see "Security Gateway Module".

Shared Management
Feature that makes it possible to assign the same Management Port (interface ethX-MgmtY) on a Quantum Maestro Orchestrator to different Security Groups. The assigned Management Port has a different IP address and a different MAC address in each Security Group, to which this port is assigned.

SIC
Secure Internal Communication. The Check Point proprietary mechanism with which Check Point computers that run Check Point software authenticate each other over SSL, for secure communication. This authentication is based on the certificates issued by the ICA on a Check Point Management Server.

Single Management Object
Single Security Gateway object in SmartConsole that represents a Security Group configured on a Quantum Maestro Orchestrator / Scalable Chassis. Acronym: SMO.

SmartConsole
Check Point GUI application used to manage a Check Point environment - configure Security Policies, configure devices, monitor products and events, install updates, and so on.

SmartDashboard
Legacy Check Point GUI client used to create and manage the security settings in versions R77.30 and lower. In versions R80.X and higher is still used to configure specific legacy settings.

SmartProvisioning
Check Point Software Blade on a Management Server (the actual name is "Provisioning") that manages large-scale deployments of Check Point Security Gateways using configuration profiles. Synonyms: Large-Scale Management, SmartLSM, LSM.

SmartUpdate
Legacy Check Point GUI client used to manage licenses and contracts in a Check Point environment.

SMO Master
The Security Appliance (in Maestro) / Security Gateway Module (on Scalable Chassis) in a Security Group that handles management tasks for all Security Appliances / Security Gateway Modules in the Security Group. By default, this role is assigned to the Security Appliance / Security Gateway Module with the lowest Member ID in the Security Group. See "SMO".

Software Blade
Specific security solution (module): (1) On a Security Gateway, each Software Blade inspects specific characteristics of the traffic (2) On a Management Server, each Software Blade enables different management capabilities.

SSM
In Maestro - a role of the Quantum Maestro Orchestrator that manages the flow of network traffic to and from the Security Groups. For Scalable Chassis - see "Security Switch Module".

Standalone
Configuration in which the Security Gateway and the Security Management Server products are installed and configured on the same server.

T

Threat Emulation
Check Point Software Blade on a Security Gateway that monitors the behavior of files in a sandbox to determine whether or not they are malicious. Acronym: TE.

Threat Extraction
Check Point Software Blade on a Security Gateway that removes malicious content from files. Acronym: TEX.

U

Updatable Object
Network object that represents an external service, such as Microsoft 365, AWS, Geo locations, and more.

Uplink Ports
Interfaces on the Quantum Maestro Orchestrator used to connect to external and internal networks. Gaia operating system shows these interfaces in Gaia Portal and in Gaia Clish. SmartConsole shows these interfaces in the corresponding SMO Security Gateway object.

URL Filtering
Check Point Software Blade on a Security Gateway that allows granular control over which web sites can be accessed by a given group of users, computers or networks. Acronym: URLF.

User Directory
Check Point Software Blade on a Management Server that integrates LDAP and other external user management servers with Check Point products and security solutions.

V

VSX
Virtual System Extension. Check Point virtual networking solution, hosted on a computer or cluster with virtual abstractions of Check Point Security Gateways and other network devices. These Virtual Devices provide the same functionality as their physical counterparts.

VSX Gateway
Physical server that hosts VSX virtual networks, including all Virtual Devices that provide the functionality of physical network devices. It holds at least one Virtual System, which is called VS0.

Z

Zero Phishing
Check Point Software Blade on a Security Gateway (R81.20 and higher) that provides real-time phishing prevention based on URLs. Acronym: ZPH.