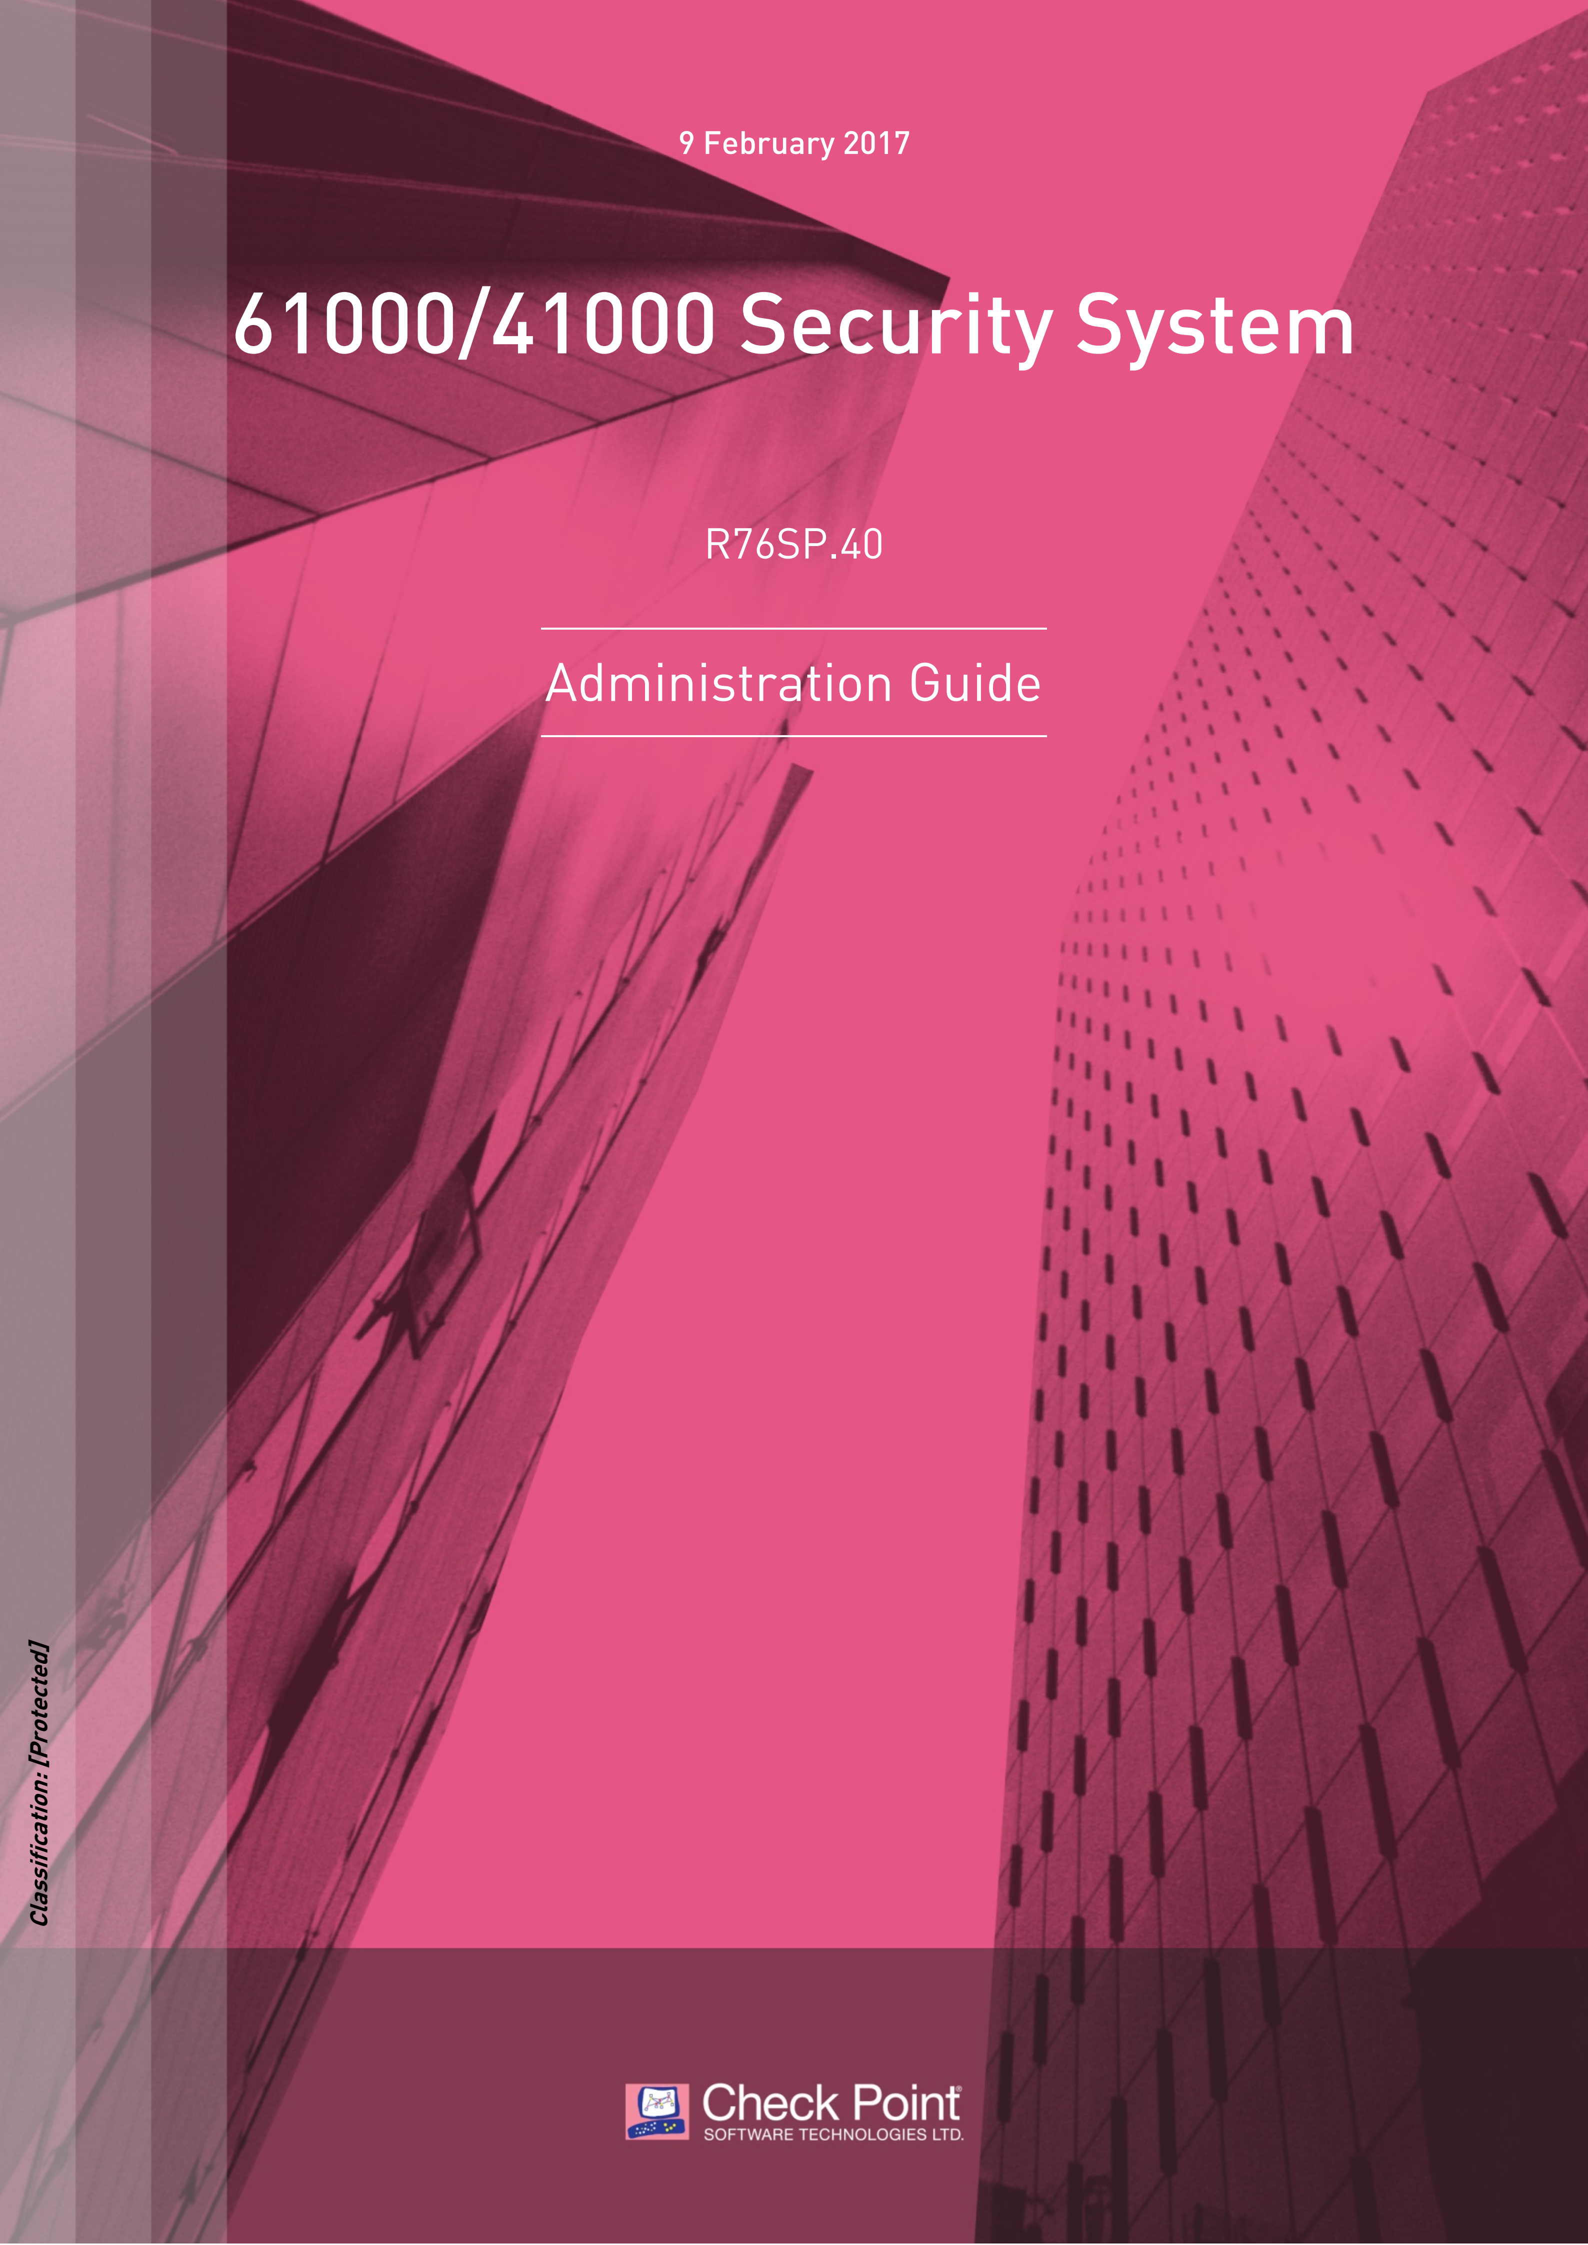

9 February 2017

61000/41000 Security System

R76SP.40

Administration Guide

Classification: [Protected]

Check Point
SOFTWARE TECHNOLOGIES LTD.

© 2017 Check Point Software Technologies Ltd.

All rights reserved. This product and related documentation are protected by copyright and distributed under licensing restricting their use, copying, distribution, and decompilation. No part of this product or related documentation may be reproduced in any form or by any means without prior written authorization of Check Point. While every precaution has been taken in the preparation of this book, Check Point assumes no responsibility for errors or omissions. This publication and features described herein are subject to change without notice.

RESTRICTED RIGHTS LEGEND:

Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 and FAR 52.227-19.

TRADEMARKS:

Refer to the Copyright page http://www.checkpoint.com/copyright.html for a list of our trademarks.

Refer to the Third Party copyright notices http://www.checkpoint.com/3rd_party_copyright.html for a list of relevant copyrights and third-party licenses.

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.

Latest Version of this Document

Download the latest version of this document http://supportcontent.checkpoint.com/documentation_download?ID=47984.

To learn more, visit the Check Point Support Center http://supportcenter.checkpoint.com.

Feedback

Check Point is engaged in a continuous effort to improve its documentation.

Please help us by sending your comments mailto:cp_techpub_feedback@checkpoint.com?subject=Feedback on 61000/41000 Security System R76SP.40 Administration Guide.

Revision History

Date	Description
09 February 2017	Updated the details on: Configuring IPv6 Static Routes - CLI (set ipv6 static-route) (on page 17), Working with Jumbo Frames (on page 264), Unique MAC Identifier Utility Options (on page 29), Destination-Based Routing (on page 57), Setting Blade-Range (on page 78), Known Limitations of asg diag Verification Tests (on page 138), Working with Management Aggregation (on page 210)
12 September 2016	Improved the explanation of the Fast Accelerator ("Using the Fast Accelerator (sim fastaccel)" on page 247).
7 June 2016	First release of this document.

Contents

Terms

Active/Standby

A High Availability cluster where only one member handles connections.

Administrator

A SmartDashboard or SmartDomain Manager user with permissions to manage Check Point security products and the network environment.

Affinity

The assignment of a specified process, Firewall instance, VSX Virtual System, interface or IRQ with one or more CPU cores.

Bond

A virtual interface that contains ("enslaves") two or more physical interfaces for redundancy and load sharing. The physical interfaces share one IP address and one MAC address.

BPDU

Bridge Protocol Data Unit. Data messages that are sent between switches in an extended LAN that uses a Spanning Tree Protocol (STP) topology.

Bridge Mode

A Security Gateway or Virtual System that works as layer-2 bridge device for easy deployment in an existing topology.

CCP

Cluster Control Protocol. Proprietary Check Point protocol that manages synchronization between High Availability between cluster members.

Chassis

The container that contains the all the components of a 61000/41000 Security System.

Cluster

1) Two or more Security Gateways or servers synchronized for High Availability or Load Sharing. 2) In a virtualized environment: a set of ESX/i hosts used for High Availability or Load Sharing.

Cluster Member

A Security Gateway that is part of a cluster.

ClusterXL

Check Point software-based cluster solution for Security Gateway redundancy and Load Sharing.

CMM

Chassis Monitoring Module. Hardware component that controls and monitors Chassis operation: fan speed, Chassis and module temperature, and component hot-swapping.

CoreXL

A performance-enhancing technology for Security Gateways on multi-core processing platforms.

Failover

A redundancy operation, where one cluster member automatically takes over for a failed member.

Firewall

The software and hardware that protects a computer network by analyzing the incoming and outgoing network traffic (packets).

Firewall Instance

On a Security Gateway with CoreXL enabled, the Firewall kernel is replicated multiple times. Each replicated copy, or firewall instance, runs on one processing core. These instances handle traffic concurrently, and each instance is a complete and independent inspection kernel.

GARP

Gratuitous Address Resolution Protocol. An ARP request or reply that is not normally required by the ARP specification (RFC 826).

Hybrid System

A 61000/41000 Security System that includes SGMs that have different quantities of CPU cores and configured CoreXL instances.

Link Aggregation

A technology that joins multiple physical interfaces together into one virtual interface, known as a bond interface. Also known as **interface bonding**.

Management Server

A Security Management Server or Multi-Domain Server that manages one or more Security Gateways and security policies.

Multi Domain Log Server

Physical server that contains the log database for all Domains.

Multi-Domain Security Management

A centralized management solution for large-scale, distributed environments with many different Domain networks.

Multi-Domain Server

A physical server that contains system information and Policy databases for all Domains in an enterprise environment.

Packet

A formatted unit of data that moves on computer networks.

PEM

Power Entry Module. Hardware component that supplies DC power to the Chassis with EMC filtering and over-current protection.

Permission Profile

A predefined group of SmartConsole access permissions assigned to Domains and administrators. This feature lets you configure complex permissions for many administrators with one definition.

Policy

A collection of rules that control network traffic and enforce organization guidelines for data protection and access to resources with packet inspection.

Primary Multi-Domain Server

The first Multi-Domain Server that you define and log into in a High Availability deployment.

PSU

Power Supply Unit. Hardware component that supplies AC power to the chassis with filtering and over-current protection.

Secondary Multi-Domain Server

All Multi-Domain Servers in a High Availability deployment created after the Primary Multi-Domain Server.

Security Gateway

A computer or appliance that inspects traffic and enforces Security Policies for connected network resources.

Security Management Server

The server that manages, creates, stores, and distributes the security policy to Security Gateways.

SGM

Security Gateway Module. 61000/41000 Security System hardware component 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.

SIC

Secure Internal Communication. The process by which networking components authenticate over SSL between themselves and the Security Management Server, as the Internal Certificate Authority (ICA), for secure communication. The Security Management Server issues a certificate, which components use to validate the identity of others.

SmartDashboard

A Check Point client used to create and manage the security policy.

SmartUpdate

SmartConsole client used to centrally upgrade and manage Check Point software and licenses.

SMO

Single Management Object. A Check Point technology that manages the 61000/41000 Security System as one large Security Gateway with one management IP address. All management tasks, are handled by one SGM (the SMO Master), which updates all other SGMs. All management tasks, such as Security Gateway configuration, policy installation, remote connections and logging are handled by the SMO master.

SMO Master

The physical SGM that handles management tasks for all SGMs in a 61000/41000 Security System environment. By default, the SGM with the lowest ID number assigned this role.

SNMP

Simple Network Management Protocol. A protocol used to monitor the activity of hardware and software in a network.

SNMP Counter

An SNMP object with an integer value that increases by one when a specified event occurs. Counters are typically used as performance metrics, such as network throughput, dropped packets, or error events.

SNMP Trap

A notification of an event generated by an SNMP-enabled device and sent to the SNMP server.

SSM

Security Switch Module. Hardware component that manages the flow of network traffic to and from the Security Gateway Modules.

Standby Domain Server

All Domain Management Servers for a Domain that are not designated as the Active Domain Management Server.

Standby Multi-Domain Server

All Multi-Domain Servers in a High Availability deployment that cannot manage global policies and objects. Standby Multi-Domain Servers are synchronized with the active Multi-Domain Server.

Traffic

The flow of data between network resources.

Virtual Device

A logical object that emulates the functionality of a type of physical network object.

Virtual Switch

Also vSwitch. A software abstraction of a physical Ethernet switch that can connect to physical switches through physical network adapters, to join virtual networks with physical networks. Can also be a Distributed Virtual Switch (dvSwitch), for definition and use on multiple ESXi hosts.

Virtual System

A virtual device that implements the functionality of a Security Gateway.

Virtual System Context

An independent VSX routing domain.

VLAN

Virtual Local Area Network. Open servers or appliances connected to a virtual network, which are not physically connected to the same network.

VLAN Trunk

A connection between two switches that contains multiple VLANs.

VPN

Virtual Private Network. A secure, encrypted connection between networks and remote clients on a public infrastructure, to give authenticated remote users and sites secured access to an organization's network and resources.

VSX

Virtual System Extension. Check Point virtual networking solution, hosted on a single computer or cluster containing 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.

Warp Link

An interface between a Virtual System and a Virtual Switch or Virtual Router that is created automatically in a VSX topology.

Introduction

In This Section:

Syntax Notation

This table shows the syntax characters used in this document.

Character	Name	Description
|	Pipe	OR
{ }	Curly brackets	Set of OR or AND operators
[ ]	Square brackets	Optional parameter
<*variable*>	Angle brackets	Variable
>	Right angle bracket	Prompt: Run command in clish or gclish (Use in procedures or examples only)
#	Hashtag	Prompt: Run command in the Expert mode (Use in procedures or examples only)
	none	Required parameter or option

Text in `monospace`: Enter exactly as shown.

Text in *italics*: Variable name. Enter a valid value.

Licensing

For information on how to monitor and administer licenses, see licenses in the *R76 Gaia Administrator's Guide* http://supportcontent.checkpoint.com/documentation_download?ID=22928.

Run all licensing commands in Global clish.

Managing the Network

In This Section:

Working with IPv6

IPv6 support is disabled by default. You must enable IPv6 support on the 61000/41000 Security System before you can configure IPv6 addresses and static routes.

To prepare your 61000/41000 Security System to work with IPv6:

1. Enable IPv6 support.
2. Install and activate an IPv6 license on the Security Management Server.
3. Create IPv6 objects in SmartDashboard.
4. Create IPv6 rules for Firewall and other Check Point Software Blades.
5. Reboot all SGMs.

Enabling/Disabling IPv6 Support (ipv6-state)

Use `ipv6-state` to:

- Enable IPv6 support for the all SGMs in the 61000/41000 Security System.
- Disable IPv6 support for the all SGMs in the 61000/41000 Security System.
- Show the IPv6 support status for all SGMs in the 61000/41000 Security System.

To complete the configuration you must reboot all SGMs at the same time. If you have a Chassis High Availability environment, you can enable IPv6 and reboot the SGMs one Chassis at a time. This feature makes it possible for network traffic to continue during configuration procedure.

Syntax

```
> set ipv6-state on|off
> show ipv6-state
```

Parameter	Description
`on|off`	`on` = Enable IPv6 support `off` = Disable IPv6 support

To Enable IPv6 Support on a Single Chassis system:

1. Log into the 61000/41000 Security System.
2. Run:
   ```
   > set ipv6-state on
   ```
3. Reboot all SGMs:
   ```
   > reboot -b all
   ```
4. Do the instructions on the screen.
5. Run:
   ```
   > show ipv6-state
   ```
   Make sure that IPv6 is enabled for all SGMs.

To Enable IPv6 on a Dual Chassis System:

This procedure lets you reboot one Chassis at a time to prevent unnecessary downtime.

1. Log into the 61000/41000 Security System.
2. Run:
   ```
   > set ipv6-state on
   ```
3. Reboot all SGMs on the Standby Chassis:
   ```
   > reboot -b <standby_chassis_name>
   ```
4. When the reboot completes, failover to the Standby Chassis:
   ```
   > asg chassis_admin -c <active_chassis_id> down
   ```
   The failover closes all active connections, which must be re-established.
5. Reboot all SGMs on the newly designated Standby Chassis:
   ```
   > reboot -b <new_standby_chassis_name>
   ```

Configuring IPv6 Static Routes - CLI (set ipv6 static-route)

Use `set ipv6 static-route` to add, change, or delete IPv6 static routes.

Syntax

```
> set ipv6 static-route <source_ip> nexthop gateway <gw_ip> [priority
<p_val>] on|off [interface <gw_if> [priority <p_val>]] on
```

```
> set ipv6 static-route <source_ip> nexthop [<gw_ip>] blackhole|reject|off
```

Parameter	Description
`gateway`	Defines the next hop path.

Parameter	Description
<*source_ip*>	Defines the source IPv6 IP and subnet.
<gw_ip>	Identifies the next hop gateway by its IP address.
<*gw_if*>	Identifies the next hop gateway by the interface that connects to it. Use this option only if the next hop gateway has an unnumbered interface.
`priority`	Assigns a path priority when there are many different paths. The available path with the lowest priority value is selected. The gateway with the lowest priority value is selected.
`interface`	Identifies the next hop gateway by the interface that connects to it. Use this option only if the next hop gateway has an unnumbered interface.
<*p_val*>	Priority for a route or interface. Valid values: Integers between 1 and 8 Default - `1`
`on`	Enables the specified route or next hop.
`off`	Deletes the specified route or next hop. If you specify a next hop, only the specified path is deleted. If no next hop is specified, the route and all related paths are deleted.
`blackhole`	Drops packets, but does not send an error message.
`reject`	Drops packets and sends an error message to the traffic source.

Note - There are no `add` or `show` commands for the static route feature.

Troubleshooting

Symptoms:

- You cannot configure the VPN Software Blade.
- This message shows: "VPN blade demands gateway's IP address corresponding to the interface's IP addresses."

Cause:

IPv6 is active, but the main IPv6 address is not configured.

Solution:

Configure the main IPv6 address in **General Properties**.

CLI Procedures - IPv6 Static Routes

This section includes some basic procedures for managing static routes using the CLI.

To show IPv6 static routes:

Run:

```
> show ipv6 route static
```

Output

```
Codes: C - Connected, S - Static, B - BGP, Rg - RIPng, A - Aggregate,
       O - OSPFv3 IntraArea (IA - InterArea, E - External),
       K - Kernel Remnant, H - Hidden, P - Suppressed

S     3100:55::1/64       is directly connected
S     3200::/64           is a blackhole route
S     3300:123::/64       is a blackhole route
S     3600:20:20:11::/64  is directly connected, eth3
```

To add an IPv6 static route:

Run:

> set ipv6 static-route <*dest_ip*> nexthop gateway <*gw_ip*> on
> set ipv6 static-route <*dest_ip*> nexthop gateway <*gw_ip*> interface <*gw_if*> on

Parameter	Description
<dest_ip>	Destination IPv6 address
<gw_ip>	Next hop *gateway* IPv6 address
<gw_if>	Next hop *gateway* interface name

Examples:

```
> set ipv6 static-route 3100:192::0/64 nexthop gateway 3900:172::1 on
> set ipv6 static-route 3100:192::0/64 nexthop gateway 3900:172::1 interface eth3
on
```

To add an IPv6 static route with paths and priorities:

Run:

> set static-route <*dest_ip*> nexthop gateway <*gw_ip*> priority <*p_val*>

Parameter	Description
<dest_ip>	Destination IP address
<gw_ip>	Next hop *gateway* IP address
<p_val>	Integer between 1 and 8 Default - 1

Run this command for each path. Assign a priority value to each. You can define two or more paths with the same priority. That specifies a backup path with the same priority.

Example:

```
> set ipv6 static-route 3100:192::0/64 nexthop gateway 3900:172::1 priority 3 on
```

To add an IPv6 static route where packets are dropped:

Run:

```
> set ipv6 static-route <dest_ip> nexthop <mess_option>
```

Parameter	Description
<dest_ip>	Destination IP address
<mess_option>	Sets whether to send an error message Allowed values: • `Reject` - Drops packets and sends an error message to the traffic source. • `Blackhole` - Drops `packets`, but does not send an error message

Examples:

```
> set ipv6 static-route 3100:192::0/64 nexthop reject
> set ipv6 static-route 3100:192::0/64 nexthop blackhole
```

To delete an IPv6 route and all related paths:

Run:

```
> set ipv6 static-route <dest_ip> off
```

Example:

```
> set ipv6 static-route 3100:192::0/64 off
```

To delete a path only:

Run:

```
> set static-route <dest_ip> nexthop gateway <gw_ip> off
```

Parameter	Description
`<dest_ip>`	Destination IP address
`<gw_ip>`	Next hop *gateway* IP address or interface name

Example:

```
> set ipv6 static-route 3100:192::0/64 nexthop gateway 3900:172::1 off
```

Configuring the 6in4 Internet Transition Mechanism

Use these commands to move IPv6 traffic over a network that does not support IPv6. The commands use the 6in4 Internet transition protocol to encapsulateIPv6 traffic onto IPv4 links.

To create 6in4 virtual interfaces:

Run:

```
> add interface <physical-if> 6in4 <6in4-id> remote <remote_ipv4_ip> [ttl <ttl>]
> set interface <sit_if_name> ipv6-address <ipv6_address> mask-length 64
```

To add the interface:

Syntax

`> add interface <physical_if> 6in4 <6in4_id> remote <remote_ipv4> [ttl <ttl>]`

Parameter	Description
<physical_if>	The physical interface traffic leaves the system from. For example: `eth1-01`
<6in4_id>	A numerical identifier for the 6in4 Virtual Interface.
<remote_ipv4_ip>	IPv4 address of the remote peer.
<ttl>	Time-to-live: the number of router hops before packets are discarded.

Example

```
> add interface eth1-01 6in4 999 remote 50.50.50.10
1_01:
Success
```

The virtual (`sit_6in4_`) interface is created for `eth1-01` on all SGMs, even though you specified a single physical interface (`eth1-01`) in the command line. To see the virtual interfaces for each SGM, run: `show interface eth1-01 6in4s`

To set the interface:

Syntax

`> set interface <sit_if_name> ipv6-address <ipv6_address> mask-length 64`

Parameter	Description
<sit_if_name>	The name of the virtual interface, which begins: `sit_6in4_`*<ID_num_from_previous_command>*
<ipv6_address>	IPv6 address

Example

```
> set interface sit_6in4_999 ipv6-address 30:30:30::1 mask-length 64
```

Output

```
1_01:
Success
```

To delete the 6in4 Virtual Interface:

Run:

`> delete interface <physical_if> 6in4 <6in4_id>`

Example

```
> delete interface eth1-01 6in4 999
```

Output

```
1_01:
success
```

asg search and 6in4

- `asg search` on IPv4 confirms if the SGM connection is active or backup and which Chassis has more than 1 SGM.
- `asg search` on IPv6 addresses shows 1 SGM on the Active Chassis and 1 SGM on the Standby Chassis

Working with Bridge Mode

C*heck* Point security devices support bridge interfaces that implement native, Layer-2 bridging. Bridge interfaces let network administrators deploy security devices in an existing topology without reconfiguring the IP routing scheme. This is an important advantage for large-scale, complex environments.

Configure Ethernet interfaces (including aggregated interfaces) on your C*heck* Point security device to work like ports on a physical bridge. The interfaces then send traffic with Layer-2 addressing. You can configure some interfaces as bridge interfaces, while other interfaces on the same device work as Layer-3 devices. Traffic between bridge interfaces is inspected at Layer-2.

- Bridge Mode is only supported with 2 interfaces.
- Bridge setup supports only the manual-general distribution mode.
- BPDU forwarding is not supported with VLAN tagging. For more information, see Disabling BPDU Forwarding (on page 26).
- The 6*1000/41000* Security System does not generate BPDU (STP) frames. It forwards BPDUs between bridge slave interfaces.
- For UserCheck to work properly, the Bridge Group must use an IP on the same subnet as clients or routers, that connect to the 6*1000/41000* Security System.

Working with Chassis High Availability in Bridge Mode

When a Dual Chassis 6*1000/41000* Security System deployment is in the Active/Standby mode, only the Active Chassis handles traffic. The 6*1000/41000* Security System maintains a MAC shadow table that caches MAC addresses handled by the system. When a Chassis failover occurs, the new Active Chassis generates advertisement packets with the cached MAC addresses. This causes the remote switches to forward traffic through a different interface, due to the updated MAC address table. The chassis in Standby mode stops forwarding BPDU frames of the spanning tree and only the new Active Chassis forwards these frames.

MAC tables

These are the MAC tables:

- OS - Not synchronized across SGMs
- Firewall - Synchronized across SGMs

To show the OS MAC table:

In Expert Mode, run:

```
# brctl showmacs <bridge_name>
```

To show the Firewall MAC table:

In clish, run:

```
> fw tab –t fdb_shadow
```

Special Advertisement Packets

When a Chassis fails over, Special Advertisement Packets are sent. They have this structure:

- Source IP - 8.7.6.5
- Destination IP - 4.3.2.1
- Destination port - 8116

Using the SSM60 in Bridge Mode

To use the SSM60 with Bridge mode:

1. Run:
   ```
   # g_update_conf_file simkern.conf bridge_mode_on_ssm60=1
   ```
2. Reboot the system.

Active/Active Bridge Mode

By default, Active/Active Bridge Mode does not support asymmetric traffic between Chassis. When asymmetric traffic is enabled, one Chassis handles client-to-server traffic and the other handles server-to-client traffic.

To enable asymmetric traffic:

In Expert Mode, run:

```
# g_update_conf_file fwkern.conf fwha_both_chassis_pass_traffic=1
# g_fw ctl set int fwha_both_chassis_pass_traffic 1
```

Important - `fwha_both_chassis_pass_traffic` can decrease performance.

Active/Active Bridge Mode Topologies

Active/Active Bridge mode supports these topologies:

- Layer-2 connectivity between Chassis. This topology requires Spanning Tree Protocol on the switches. The Spanning Tree Protocol (STP) is a network protocol that ensures a loop-free topology for Ethernet networks by sending special data frames called Bridge Protocol Data Units (BPDUs). The BPDUs help the switches decide which port to block in case of a loop detection. The BPDUs reach the switch from a different interface when they pass through the bridge interface of the gateway, which results in a successful blockage.

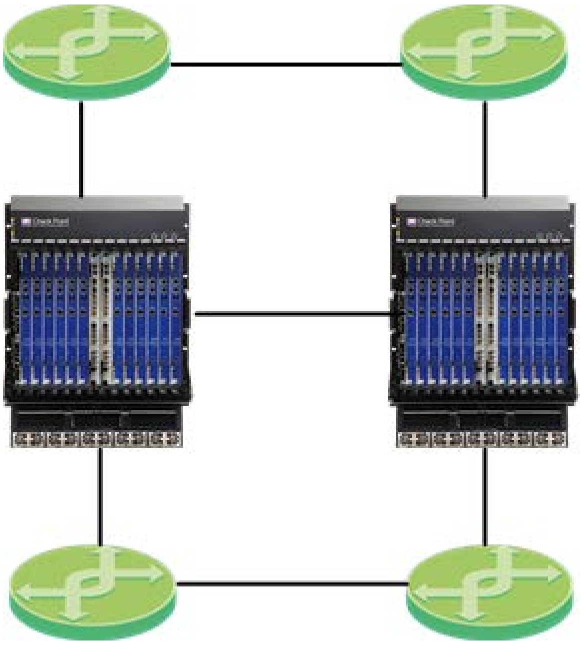

- No Layer-2 connectivity between Chassis. This topology does not require Spanning Tree Protocol on the switches. It is usually a router-based topology where a dynamic routing protocol decides through which segment to route the traffic.

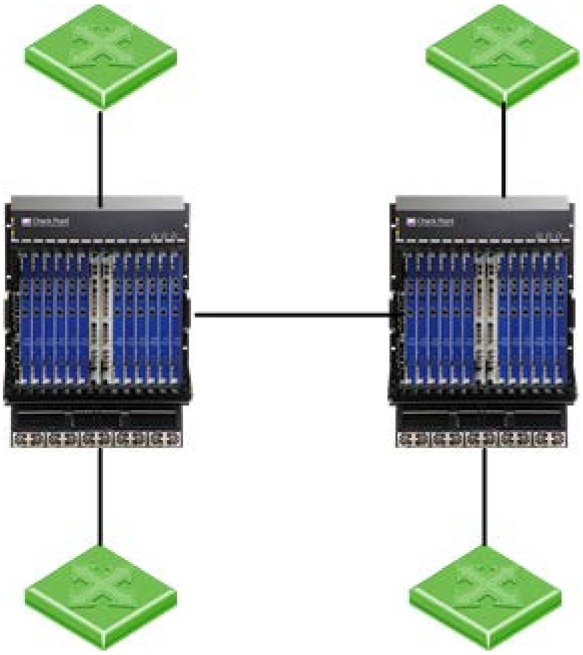

BPDU

The BDPU maximum age timer, controls the maximum time before a bridge port saves its BPDU information. It is set to 20 seconds by default, the time it takes to reach failover. You can change the BDPU maximum age timer to 6 seconds.

For example: On Cisco switches, use `spanning-tree vlan` on each VLAN to configure the BDPU maximum age timer.

Syntax

```
> spanning-tree vlan <vlan_id> max-age <age>
```

Parameter	Description
<*vlan_id*>	VLAN ID
<*age*>	BDPU maximum age in seconds Allowed values: 6-40

For more information, see the Cisco documentation.

Active/Standby Bridge Mode

To enable Active/Standby Bridge Mode:

In gclish, run:

```
# add bridging group <bridge_if_name>
# add bridging group <bridge_if_name> interface <if_1_ID>
# add bridging group <bridge_if_name> interface <if_2_ID>
```

Configuring Bridge Interfaces in SGW Mode

Use these commands to work with Bridge interfaces:

Syntax

```
> add bridging group <group_id> interface <if_name>
> delete bridging group <group_id> interface <if_name>
> show bridging group <group_id>
```

Parameter	Description
<*group_id*>	Integer that identifies the bridging group
<*if_name*>	Interface name as configured on the system

Example

```
> add bridging group 2 interface eth3
> show bridging group 2
```

Output

```
Bridge Configuration
    Bridge Interfaces
        eth3
```

To use VLAN interfaces in a bridging group:

1. Add the VLAN to the physical interface:

   ```
   > add interface <if_name> vlan <vlan_id>
   ```

2. Add the interface VLAN to the bridging group:

   ```
   > add bridging group <group_id> interface <if_name>.<vlan_id>
   ```

Configuring Bridge Interfaces in VSX Mode

Configure bridge mode in a virtual system to define it when you first create the object.

To configure Active/Standby Bridge Mode in a Virtual System:

1. In **Virtual System General Properties**, select **Bridge Mode**.
2. Click **Next**.

   The Virtual System Network Configuration window opens.
3. Configure the external and internal interfaces for the Virtual System.
4. Click **Next**.
5. Click **Finish**.

Disabling BPDU Forwarding

When VLAN translation is configured, BPDU frames can send the incorrect VLAN number to switch ports through the bridge. This mismatch can cause the switch port to block traffic.

To resolve this, disable BPDU forwarding in a way that survives reboot. This solution also works well for Layer-2 Virtual Systems.

To permanently disable BPDU forwarding:

1. Open `/etc/rc.d/init.d/network` in a text editor.
2. Search for:

   ```
   /etc/init.d/functions
   ```

3. Add this new line.

   ```
   /sbin/sysctl -w net.bridge.bpdu_forwarding=0
   ```

4. Exit the editor and save the file.
5. Copy the file to all SGMs.

   ```
   > asg_cp2blades /etc/rc.d/init.d/network
   ```

6. Reboot the system.

   If you are using a dual Chassis 61000/41000 Security System, reboot the Standby Chassis first and then reboot the Active Chassis.

To learn more, see sk98927 http://supportcontent.checkpoint.com/solutions?id=sk98927.

IPv6 Neighbor Discovery

Neighbor discovery works over the ICMPv6 Neighbor Discovery protocol, which is the functional equivalent of the IPv4 ARP protocol. ICMPv6 Neighbor Discovery Protocol must be explicitly allowed for all bridged networks in your Firewall rules. This is different from ARP, for which traffic is always allowed regardless of the Rule Base.

This is an example of a rule that allows ICMPv6 Neighbor Discovery protocol:

- **Source** - *Bridged_Network*
- **Destination** - *Bridged_Network*
- **Services & Applications** - **neighbor-advertisement**, **neighbor-solicitation**, **router-advertisement**, **router-solicitation**, **redirect6**
- **Action** - **Accept**

Working with Link State Propagation

Link State Propagation (LSP) lets you combine physical SSM interfaces into groups. An external switch or router connected to a gateway, fails over quickly when you use dynamic routing. LSP is disabled by default.

This release adds support for LSP Port Groups. An LSP Port Group is a set of one or more interfaces in Interface Groups. If all interfaces in an Interface Group are DOWN, the other Interface Groups in the LSP Port Group automatically go DOWN. If at least one interface in each Interface Group is UP, all Interface Groups in the LSP Port Group stay in the UP state.

To enable or disable Link State Propagation, run: `asg_lsp_util` <*enable* | *disable*>

Defining LSP Port Groups

Define LSP Port Groups in `/etc/lsp_groups.conf`. Each line in this file defines one LSP Port Group, with one or more Interface Groups, delimited by a comma. An Interface Group has one or more interfaces, delimited by a plus sign (+).

Configuration file syntax:

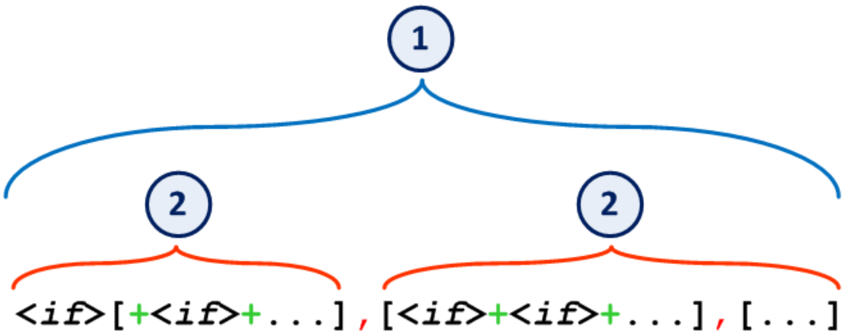

Item	Description
1	LSP Port Group (full syntax)
2	Interface Group
<if>	Physical Interface

Examples (4 SSM Installation):

```
eth1-01+eth2-01,eth3-01+eth4-01
eth1-02+eth1-03+eth1-04+eth1-05,eth3-02+eth4-02,eth3-03+eth4-03
```

In the first example, the LSP Port Group has two Interface Groups with two interfaces:

- Interface Group 1 contains *eth1-01* and *eth2-01*
- Interface Group 2 contains *eth3-01* and *eth4-01*

In the second example, the LSP port Group has three Interface Groups, one with four interfaces and others with two interfaces each.

To add an LSP Port Group:

1. Open `/etc/lsp_groups.conf` in a text editor.

   If `lsp_groups.conf` is not in the directory, create it now.
2. Add one line for each LSP Port Group.
3. Run `asg_lsp_util disable` and then `run asg_lsp_util enable`

   This step in necessary for the system to detect the change.

To delete an LSP Port Group:

1. Open `/etc/lsp_groups.conf` in a text editor.
2. Delete the applicable LSP Port Group line in the file.
3. Run `asg_lsp_util disable` and then `run asg_lsp_util enable`

   This step in necessary for the system to detect the change.

Note - Do not delete the configuration file or the only LSP port group line in the file. If you do not use LSP, disable it.

Configuring a Unique MAC Identifier

When there are more than one 61000/41000 Security System or ClusterXL systems on a Layer-2 segment, the Unique MAC Identifier must be different for each system. The Unique MAC Identifier is assigned by default during the initial setup. The last octet of the management interface MAC address is the Unique MAC Identifier.

The last octet of the management interface MAC address is set for these data interface types:

- Interfaces with names in the `ethX-YZ` format
- Bond interfaces
- VSX wrp interfaces
- VLAN interfaces

If there is no configured management interface, the Unique MAC Identifier is assigned the default value 254.

You can use `asg_unique_mac_utility` to set:

- Data interface Unique MAC Identifier
- Host name

To manually set the Unique MAC Identifier:

1. Run:
   `> asg_unique_mac_utility`
2. Select an option from the menu and do the instructions on the screen.

```
---------------------------------------------
|             Unique MAC Utility            |
---------------------------------------------
| HOSTNAME     [61000_GW]                   |
| Unique MAC   [254]                        |
---------------------------------------------


Choose one of the following options:
------------------------------------
1) Set Hostname with Unique MAC wizard
2) Apply Unique MAC from current HOSTNAME
3) Manual set Unique MAC
4) Revert to Unique MAC Factory Default
5) Exit
```

Note - You must reboot the system to apply the new Unique MAC Identifier.

Unique MAC Identifier Utility Options

1. **Set Host name with Unique MAC wizard.**

   The `_asg` suffix and the setup number, between 1 and 254, are added to the setup name.

   For example:

Setup Name	Suffix	Setup number
61000_GW	_asg	22

   This creates a new Host name with a Unique MAC Identifier of 22.

New Host Name	Unique MAC Identifier
61000_GW_asg22	22

   The setup number replaces the Unique MAC Identifier value of 254.

   After reboot, all data interface MAC addresses now have the new Unique MAC Identifier value 16.

   For example: `eth1-01 00:1C:7F:XY:ZW:`**16**

   Note - The last octet for `eth1-01` (shown in bold) is 16 hex (22 decimal).

2. **Apply Unique MAC from current Host name.**

   Assign a new Unique MAC Identifier to the interfaces. The new Unique MAC Identifier is created from the setup number in the host name.

   The current host name must first comply with the setup `name/asg suffix/setup` number convention.

3. **Manual Set Unique MAC.**

   Change the unique MAC with your own data, without changing the host name value.

   The existing host name does not have to comply with the setup `name/asg suffix/setup` number convention.

4. **Revert to Unique MAC Factory Default.**

   Set the Unique MAC identifier to the default value of 254.

Verifying the New MAC Address

Use these commands to make sure that the Unique MAC Identifier was changed:

- `# mac_verifier -v`
- `# ifconfig` *<if_name>*

Example

```
# ifconfig eth1-01
```

```
eth1-01  Link encap:Ethernet  HWaddr 00:1C:7F:81:01:16
         inet6 addr: fe80::21c:7fff:fe81:116/64 Scope:Link
         UP BROADCAST RUNNING SLAVE MULTICAST  MTU:1500 Metric:1
         RX packets:154820 errors:0 dropped:0 overruns:0 frame:0
         TX packets:23134 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:0 RX bytes:15965660 (15.2 MiB)
         TX bytes:2003398 (1.9 MiB)
```

Note - The last octet for eth1-01, shown in bold, is 16 hex (22 decimal).

Configuring VLANs

Use these commands to configure VLANs. These commands do not work in a VSX environment.

Syntax

```
> add interface <if_name> vlan <vlan_id>
> set interface <if_name>.<vlan_id> ip-address <ip_addr> mask-length <mask-len>
> delete interface <if_name> vlan <vlan_id>
```

Parameter	Description
<if_name>	Name of the physical interface
<vlan_id>	VLAN ID number
<ip_addr>	VLAN IPv4 or IPv6 Address
<mask-len>	Network mask length

Example -

Adding a VLAN interface

```
> add interface eth2-03 vlan 444
```

Output

```
1_01:
success
```

Example

Configuring a VLAN interface

```
set interface eth2-03.444 ipv4-address 203.0.113.1 mask-length 24
```

Output

```
1_01:
success
```

Example

Deleting a VLAN interface

```
delete interface eth2-03 vlan 444
```

Output

```
1_01:
success
```

To show the VLAN interfaces on a physical interface:

Run:

```
> show interface <interface> vlans
```

Example

```
> show interface eth2-03 vlans
```

Output

```
1_01:
eth2-03.444
```

Changing the Management Interface

Use this command to change the management interface for the SGMs.

Note - This procedure is applicable for Security Gateway environments only. Management interface changes are not supported for VSX.

To change the Management Interface:

1. Make sure that the management interface cable is connected to the network.
2. Run these commands in order:
   - `> set management interface <new_mng_if>`
   - `> delete interface <old_mng_if> ipv4-address`
   - `> set interface <new_mng_if> ipv4-address <ip> mask-length <length>`
   - `> set interface <new_mng_if> state on|off`

   Note - Do these commands through a console connection to ensure connectivity when you change the interfaces.
3. In SmartDashboard, get the new topology for the 61000/41000 Security System object.

4. Install policy.

Parameter	Description
<*new_mng_if*>	Interface name of the new management interface. For example: `eth1-Mgmt3`
<*old_mng_if*>	Interface name of the existing management interface that is to be changed or deleted. For example: `eth1-Mgmt2`.
<*ip*>	Interface IPv4 address
<*length*>	Interface net mask
`state`	Interface state (`on`/`off`)

Working with ECMP

To manually define a static route to a number of next-hop gateways, use Equal-cost multi-path routing (ECMP). By load-balancing traffic over multiple paths, to get to the destination network defined in the static route, it may offer substantial increases in bandwidth.

Syntax

```
> set static-route <network> nexthop gateway address <gw_ip> on
```

Parameter	Description
<*network*>	The IP address of the destination network
<*gw_ip*>	The IP address of the next-hop gateway

Example

```
> set static-route 50.50.50.0/24 nexthop gateway address 20.20.20.101 on
> set static-route 50.50.50.0/24 nexthop gateway address 20.20.20.102 on
> set static-route 50.50.50.0/24 nexthop gateway address 20.20.20.103 on
```

Notes

To get to addresses on the 50.50.50.0/24 network, packets must first be forwarded to one of these gateways:

- `20.20.20.101`
- `20.20.20.102`
- `20.20.20.103`

To make sure static routes to the next-hop gateways are being enforced:

Run:

```
> show route static
1_01:
Codes: C - Connected, S - Static, R - RIP, B - BGP,O - OSPF IntraArea (IA - InterArea,
E - External, N - NSSA)A - Aggregate, K - Kernel Remnant, H - Hidden, P - Suppressed

S 0.0.0.0/0  via 192.168.33.1, eth2-01, cost 0, age 2092
5.5.5.0/24  via 20.20.20.101, eth1-01, cost 0, age 322
            via 20.20.20.102, eth1-01
            via 20.20.20.103, eth1-01
```

The output shows that the static route to 50.50.50.0/24 is through three next-hop gateways.

Enhanced Failover of ECMP Static Routes

The SGM of every gateway of a static route, is pinged to detect its availability. On detection of an unreachable Next Hop gateway, the enhanced routing feature automatically starts failover and it is deleted from the routing table. When the gateway is again reachable, it is re-added to the routing table.

Syntax

```
> set static-route <network>/<subnet_len> ping on
```

Note - You can configure enhanced ECMP failover after you configure an ECMP static route.

Parameter	Description
<*network*>	The IP address of the destination network
<*subnet_len*>	The subnet length of the destination network

To adjust ping behavior:

```
> set ping count <val>
> set ping interval <val>
```

Parameter	Description
`count` <*val*>	Number of packets to be sent before next hop is declared dead
`interval` <*val*>	Time in seconds to wait between two consecutive pings

Example

1. Set ECMP for destination 5.5.5.0/24.

   ```
   > set static-route 5.5.5.0/24 nexthop gateway address 10.33.85.2 on
   > set static-route 5.5.5.0/24 nexthop gateway address 10.33.85.4 on
   > set static-route 5.5.5.0/24 nexthop gateway address 10.33.85.100 on
   > show route
   1_01:
   Codes: C - Connected, S - Static, R - RIP, B - BGP,
          O - OSPF IntraArea (IA - InterArea, E - External, N - NSSA)
          A - Aggregate, K - Kernel Remnant, H - Hidden, P - Suppressed

   S     0.0.0.0/0           via 192.168.33.1, eth2-01, cost 0, age 2092
         5.5.5.0/24          via 10.33.85.2, eth1-01, cost 0, age 322
                             via 10.33.85.4, eth1-01
                             via 10.33.85.100, eth1-01
   ```

2. Enable failover ECMP on all static routes configured for destination 5.5.5.0/24.

   ```
   > set static-route 5.5.5.0/24 ping on
   ```

   Make sure the configuration is correct. When next-hop 10.33.85.2 is unreachable (no ICMP replies), after 3 pings (by default) it is removed from the routing table.

   ```
   [Expert@CH_Lena-ch02-01]# tcpdump -nepi eth1-01 host 10.33.85.2
   tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
   listening on eth1-01, link-type EN10MB (Ethernet), capture size 96 bytes
   14:40:48.388032 00:1c:7f:a1:01:55 > 00:50:56:a7:7f:f5, ethertype IPv4
   (0x0800), length 62: 10.33.85.1 > 10.33.85.2: ICMP echo request, id 53007,
   seq 43981, length 28
   14:40:58.388425 00:1c:7f:a1:01:55 > 00:50:56:a7:7f:f5, ethertype IPv4
   (0x0800), length 62: 10.33.85.1 > 10.33.85.2: ICMP echo request, id 53007,
   seq 43981, length 28
   14:41:08.387895 00:1c:7f:a1:01:55 > 00:50:56:a7:7f:f5, ethertype IPv4
   (0x0800), length 62: 10.33.85.1 > 10.33.85.2: ICMP echo request, id 53007,
   seq 43981, length 28
   ```

   The route is deleted from the routing table.

   ```
   01 > show route
   1_01:
   Codes: C - Connected, S - Static, R - RIP, B - BGP,
          O - OSPF IntraArea (IA - InterArea, E - External, N - NSSA)
          A - Aggregate, K - Kernel Remnant, H - Hidden, P - Suppressed

        0.0.0.0/0           via 192.168.33.1, eth2-01, cost 0, age 2511
   S     5.5.5.0/24          via 10.33.85.4, eth1-01, cost 0, age 52
                             via 10.33.85.100, eth1-01
   ```

   When 10.33.85.2 is reached, `tcpdump` shows that it replies to the ping requests, and is re-added to the routing table.

   ```
   [Expert@CH_Lena-ch02-01]# tcpdump -nepi eth1-01 host 10.33.85.2
   tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
   listening on eth1-01, link-type EN10MB (Ethernet), capture size 96 bytes
   ```

```
14:38:08.388224 00:1c:7f:a1:01:55 > 00:50:56:a7:7f:f5, ethertype IPv4
(0x0800), length 62: 10.33.85.1 > 10.33.85.2: ICMP echo request, id 53007,
seq 43981, length 28
14:38:08.388462 00:50:fc:58:80:0a > 00:1c:7f:0f:00:fe, ethertype IPv4
(0x0800), length 62: 10.33.85.2 > 10.33.85.1: ICMP echo reply, id 53007,
seq 43981, length 28
14:38:18.387762 00:1c:7f:a1:01:55 > 00:50:56:a7:7f:f5, ethertype IPv4
(0x0800), length 62: 10.33.85.1 > 10.33.85.2: ICMP echo request, id 53007,
seq 43981, length 28
14:38:18.387980 00:50:fc:58:80:0a > 00:1c:7f:0f:00:fe, ethertype IPv4
(0x0800), length 62: 10.33.85.2 > 10.33.85.1: ICMP echo reply, id 53007,
seq 43981, length 28
14:38:28.388161 00:1c:7f:a1:01:55 > 00:50:56:a7:7f:f5, ethertype IPv4
(0x0800), length 62: 10.33.85.1 > 10.33.85.2: ICMP echo request, id 53007,
seq 43981, length 28
14:38:28.388382 00:50:fc:58:80:0a > 00:1c:7f:0f:00:fe, ethertype IPv4
(0x0800), length 62: 10.33.85.2 > 10.33.85.1: ICMP echo reply, id 53007,
seq 43981, length 28

> show route
1_01:
Codes: C - Connected, S - Static, R - RIP, B - BGP,
       O - OSPF IntraArea (IA - InterArea, E - External, N - NSSA)
       A - Aggregate, K - Kernel Remnant, H - Hidden, P - Suppressed

S     0.0.0.0/0           via 192.168.33.1, eth2-01, cost 0, age 2092
     5.5.5.0/24          via 10.33.85.2, eth1-01, cost 0, age 322
                          via 10.33.85.4, eth1-01
                          via 10.33.85.100, eth1-01
```

Validation

1. Run from gclish:

   `> show route`

2. Make sure that only ECMP static routes with reachable Next Hops show.

   Run: `tcpdump`

3. Make sure that every few seconds there is a ping request on the interface with static route, and that the ping is on.

Working with the ARP Table (asg_arp)

`asg_arp` shows the ARP cache for the whole 61000/41000 Security System or for the specified SGMs, Interface, MAC address, and Host name. You can show summary or detailed (verbose) information. You can also run MAC address verification on both Chassis.

Syntax

```
# asg_arp -h
# asg_arp [-b <sgm_ids>] [-v]  [--verify] [-i <if>] [-m <mac>] [<hostname>]
# asg_arp --legacy
```

Parameter	Description
`-h`	Shows command syntax and help information
`-v`	Verbose - Shows detailed SGM cache information
`-b` <*sgm_ids*>	Works with SGMs and/or Chassis as specified by <*sgm_ids*>. <*sgm_ids*> can be: • No <*sgm_ids*> specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`)
`-i` <*if*>	Shows the ARP cache for the specified interface
`-m` <*mac*>	Shows the ARP cache for the specified MAC address
<*hostname*>	Shows the ARP cache for the specified host name
`--verify`	Run MAC address verification on both Chassis and show the results
`--legacy`	Shows the ARP cache for each SGM in the legacy format

Verbose Mode Output

This example shows the ARP cash in the detailed (verbose) mode for the active Chassis.

```
# asg_arp -v
Address        HWtype HWaddress         Flags   Iface       SGMs
172.23.9.198   ether  00:0C:29:87:AF:15    C    eth1-Mgmt1  1_1, 1_3, 1_4, 1_5
192.0.2.5      ether  00:1C:7F:05:04:FE    C    Sync        1_1, 1_3, 1_4
172.23.9.4     ether  00:17:65:3C:30:43    C    eth1-Mgmt1  1_1
192.0.2.3      ether  00:1C:7F:03:04:FE    C    Sync        1_1, 1_5
192.0.2.4      ether  00:1C:7F:04:04:FE    C    Sync        1_1, 1_3, 1_5
192.0.2.1      ether  00:1C:7F:01:04:FE    C    Sync        1_3, 1_4, 1_5
24.24.24.1     ether  00:04:23:C0:0E:98    C    eth2-01     1_3, 1_5
14.14.14.3     ether  00:04:23:C0:0F:5B    C    eth1-01     1_3, 1_5
198.51.100.32  ether  00:A0:12:99:E6:22    C    eth1-CIN    1_5
198.51.100.232 ether  00:A0:12:99:65:E2    C    eth2-CIN    1_5
198.51.100.33  ether  00:18:49:01:B3:82    C    eth1-CIN    1_5
```

Verifying MAC Addresses

This example shows the output of the MAC address verification on the active Chassis.

```
# asg_arp --verify
Address      HWtype   HWaddress            Flags Mask  Iface        SGMs
172.23.9.4   ether    00:17:65:3C:30:43    C           eth1-Mgmt4   2_02
192.0.2.16   ether    00:1C:7F:10:04:FE    C           Sync         2_03,2_04
192.0.2.17   ether    00:1C:7F:11:04:FE    C           Sync         2_02,2_04
192.0.2.18   ether    00:1C:7F:12:04:FE    C           Sync         2_02,2_03
cmm          ether    00:18:49:01:6D:89    C           eth1-CIN     2_02
ssm1         ether    00:A0:12:A4:63:41    C           eth1-CIN     2_02
ssm2         .        (incomplete)         .           eth2-CIN     2_02

Starting mac address verification on local chassis... (Chassis 2)
```

```
No inconsistency found on local chassis
```

Legacy Mode Output

This example shows the legacy mode output, for each SGM.

```
# asg_arp --legacy
1_01:
Address        HWtype  HWaddress           Flags Mask        Iface
172.23.9.198   ether   00:0C:29:87:AF:15   C                 eth1-Mgmt1
192.0.2.5      ether   00:1C:7F:05:04:FE   C                 Sync
172.23.9.4     ether   00:17:65:3C:30:43   C                 eth1-Mgmt1
192.0.2.3      ether   00:1C:7F:03:04:FE   C                 Sync
192.0.2.4      ether   00:1C:7F:04:04:FE   C                 Sync
1_03:
Address        HWtype  HWaddress           Flags Mask        Iface
192.0.2.5      ether   00:1C:7F:05:04:FE   C                 Sync
24.24.24.1     ether   00:04:23:C0:0E:98   C                 eth2-01
192.0.2.4      ether   00:1C:7F:04:04:FE   C                 Sync
192.0.2.1      ether   00:1C:7F:01:04:FE   C                 Sync
172.23.9.198   ether   00:0C:29:87:AF:15   C                 eth1-Mgmt1
14.14.14.3     ether   00:04:23:C0:0F:5B   C                 eth1-01
1_04:
Address        HWtype  HWaddress           Flags Mask        Iface
192.0.2.1      ether   00:1C:7F:01:04:FE   C                 Sync
172.23.9.198   ether   00:0C:29:87:AF:15   C                 eth1-Mgmt1
192.0.2.5      ether   00:1C:7F:05:04:FE   C                 Sync
1_05:
Address        HWtype  HWaddress           Flags Mask        Iface
ssm1           ether   00:A0:12:99:E6:22   C                 eth1-CIN
192.0.2.3      ether   00:1C:7F:03:04:FE   C                 Sync
172.23.9.198   ether   00:0C:29:87:AF:15   C                 eth1-Mgmt1
14.14.14.3     ether   00:04:23:C0:0F:5B   C                 eth1-01
192.0.2.4      ether   00:1C:7F:04:04:FE   C                 Sync
ssm2           ether   00:A0:12:99:65:E2   C                 eth2-CIN
192.0.2.1      ether   00:1C:7F:01:04:FE   C                 Sync
cmm            ether   00:18:49:01:B3:82   C                 eth1-CIN
24.24.24.1     ether   00:04:23:C0:0E:98   C                 eth2-01
```

Working with Proxy ARP for Manual NAT

A gateway can respond to ARP requests on behalf of other hosts, with Proxy ARP. For more information about Proxy ARP configuration, see sk30197 https://supportcenter.checkpoint.com/supportcenter/portal?eventSubmit_doGoviewsolutiondetails=&solutionid=sk30197.

To configure the proxy ARP mechanism on the 61000/41000 Security System:

1. Add these to `$FWDIR/conf/local.arp` on the local SGM
   a) The IPs that the 61000/41000 Security System should answer for ARP requests
   b) The respective MAC addresses to be advertised

For example, to reply to ARP requests for IP 192.168.10.100 on interface eth2-01 with MAC address 00:1C:7F:82:01:FE, add the entry below to `local.arp`:

```
192.168.10.100 00:1C:7F:82:01:FE
```

Note - The interface VMAC is different between Chassis when working on a Dual Chassis setup. When editing `local.arp`, MAC values must be taken from the local SGM.

1. Distribute the updated `local.arp` to all SGMs

   ```
   # local_arp_update
   ```

   This command distributes `local.arp` to all SGMs in the system, and automatically changes the MAC values for SGMs on another Chassis.
2. Enable the **Merge manual proxy ARP configuration** option in **SmartDashboard > Global Properties > NAT**.
3. Install policy to apply the updated proxy ARP entries.

Notes:

- When you add an SGM to a system with the Proxy ARP configured, the `local.arp` file is automatically copied to the new SGM from the SMO.
- When you change `local.arp` on a Virtual System, the changes apply to that Virtual System only.
- Proxy ARP is also required when configuring Connect Control on the 61000/41000 Security System.

Verification

To make sure that all the entries in `local.arp` are applied correctly on the system, run:

```
# asg_local_arp_verifier
```

To compare the entries manually, run:

```
# g_fw ctl arp
```

Configuring Port Speed

Configuring SSM Port Speed

Use the `asg_port_speed` command to configure and verify SSM data ports. This command saves the settings in `/etc/ssm_port_speed.conf` and automatically copies it to all SGMs. Run this command in the Expert mode.

Syntax

To configure the 40g ports mode on all SSMs, run:

```
# asg_port_speed 40g_mode <ssm_id> on | off
```

Parameter	Description
`<ssm_id>`	SSM Identification number
`on`	All SSMs work in the 40G mode
`off`	All SSMs work in the 4x10G mode

To configure the interface speed, run:

```
# asg_port_speed set <ifn> <speed> 10G|1G|0
```

Parameter	Description
`<ifn>`	Interface name (for example: eth1-01)
`<speed>`	Interface speed `10G` - 10 Gb/second `1G` - 1 Gb/second `0` - Automatically selected based on detected hardware

To apply the specified configuration file on all SSMs, run:

```
# asg_port_speed config <file_name>
```

To create a configuration file based on the current system status, run:

```
# asg_port_speed create_conf
```

To run SSM verifications and show the configuration settings, run:

```
# asg_port_speed verify
```

This command confirms that:

- The configuration file is the same on all SGMs
- The SSM port speed on all Chassis is the same as defined in the configuration file.

Command Examples

Verification

```
# asg_port_speed verify
+------------------------------------------------+
|Port speed verifier                             |
+---------------+----------+----------+----------+
|Interface      |Conf.     |Chassis1  |Result    |
+---------------+----------+----------+----------+
|eth1-01        |10G       |10G       |OK        |
+---------------+----------+----------+----------+
|eth1-02        |10G       |10G       |OK        |
+---------------+----------+----------+----------+
|eth1-03        |10G       |10G       |OK        |
+---------------+----------+----------+----------+
|eth1-04        |10G       |10G       |OK        |
+---------------+----------+----------+----------+
|eth1-05        |10G       |10G       |OK        |
+---------------+----------+----------+----------+
|eth1-06        |10G       |10G       |OK        |
+---------------+----------+----------+----------+
|eth1-07        |10G       |10G       |OK        |
+---------------+----------+----------+----------+
|eth1-09        |40G       |40G       |OK        |
+---------------+----------+----------+----------+
|eth1-10        |auto      |auto      |OK        |
+---------------+----------+----------+----------+
|eth1-11        |auto      |auto      |OK        |
+---------------+----------+----------+----------+
|eth1-12        |auto      |auto      |OK        |
+---------------+----------+----------+----------+
|eth1-13        |40G       |40G       |OK        |
+---------------+----------+----------+----------+
|eth1-14        |auto      |auto      |OK        |
+---------------+----------+----------+----------+
|eth1-15        |auto      |auto      |OK        |
+---------------+----------+----------+----------+
|eth1-16        |auto      |auto      |OK        |
+---------------+----------+----------+----------+
|eth2-01        |10G       |10G       |OK        |
+---------------+----------+----------+----------+
|eth2-02        |10G       |10G       |OK        |
+---------------+----------+----------+----------+
|eth2-03        |10G       |10G       |OK        |
+---------------+----------+----------+----------+
|eth2-04        |10G       |10G       |OK        |
+---------------+----------+----------+----------+
|eth2-05        |10G       |10G       |OK        |
+---------------+----------+----------+----------+
|eth2-06        |10G       |10G       |OK        |
+---------------+----------+----------+----------+
|eth2-07        |10G       |10G       |OK        |
+---------------+----------+----------+----------+
|eth2-09        |40G       |40G       |OK        |
+---------------+----------+----------+----------+
|eth2-10        |auto      |auto      |OK        |
+---------------+----------+----------+----------+
|eth2-11        |auto      |auto      |OK        |
+---------------+----------+----------+----------+
|eth2-12        |auto      |auto      |OK        |
+---------------+----------+----------+----------+
|eth2-13        |40G       |40G       |OK        |
+---------------+----------+----------+----------+
|eth2-14        |auto      |auto      |OK        |
+---------------+----------+----------+----------+
|eth2-15        |auto      |auto      |OK        |
+---------------+----------+----------+----------+
|eth2-16        |auto      |auto      |OK        |
+---------------+----------+----------+----------+
|SSM1 40G mode  |on        |on        |OK        |
+---------------+----------+----------+----------+
|SSM2 40G mode  |on        |on        |OK        |
+---------------+----------+----------+----------+
Comparing SSMs configuration with conf file...              [   OK   ]
Comparing /etc/ssm_port_speed.conf on all SGMS...           [   OK   ]
```

To Set Interface Speed:

```
#asg_port_speed set eth1-05 1G
Executing "ccutil set_port_speed 1 05 1000" on SGM 1...     [   OK   ]
Updating /etc/ssm_port_speed.conf...                        [   OK   ]
Copy /etc/ssm_port_speed.conf to all SGMs...                [   OK   ]
```

To Set Port Speed:

```
#asg_port_speed 40g_mode 1 on

Changing 40G mode on SSM1 will revert SSM to manufactory defaults. All current
configuration of SSM1 will be deleted.
Run "asg_port_speed config /etc/ssm_port_speed.conf" to apply current configuration
Proceed with configuration?(y/n)
>y

Updating /etc/ssm_port_speed.conf...                        [   OK   ]
Copy /etc/ssm_port_speed.conf to all SGMs...                [   OK   ]
Executing "ccutil set_qsfp_ports_mode 1 40G" on SGM 1...    [   OK   ]
```

Management Port Speed Configuration

To set the speed of a management port on a dual Chassis configuration:

Run this procedure on the SSM of both Chassis.

1. Connect to the SSM.

   To learn how to connect to the SSM, see SSM160 CLI (on page 299).

2. Run these commands:

   ```
   # config
   # port <port>
   # speed <speed>
   # commit
   # end
   ```

3. Make sure that the port speed is correct:

   ```
   # show port <port> <speed>
   ```

Parameter	Description
<port>	In SSM160 use: • 1/5/3 for ethx-mgmt03 • 1/5/4 for ethX-mgmt04 In SSM60 use: • 1/5/1 for ethx-mgmt01 • 1/5/2 for ethX-mgmt02
<speed>	Speed in Mbps Valid values: • `1000` • `100`

Example

```
> T-HUB4#config
Entering configuration mode terminal

--- WARNING ------------------------------------------------------
Running db may be inconsistent. Enter private configuration mode and
install a saved configuration.
------------------------------------------------------------------
T-HUB4(config)#port 1/5/4

--- WARNING ------------------------------------------------------
Running db may be inconsistent. Enter private configuration mode and
install a saved configuration.
------------------------------------------------------------------
T-HUB4(config-port-1/5/4)#speed 100

--- WARNING ------------------------------------------------------
Running db may be inconsistent. Enter private configuration mode and
install a saved configuration.
------------------------------------------------------------------
T-HUB4(config-port-1/5/4)#commit
% No modifications to commit.

--- WARNING ------------------------------------------------------
```

```
Running db may be inconsistent. Enter private configuration mode and
install a saved configuration.
------------------------------------------------------------------
T-HUB4(config-port-1/5/4)#end

T-HUB4#show port 1/5/4

============================================================================
Ethernet Interface
============================================================================
Interface          : 1/5/4
Description        :
Admin State        : up              Port State              : up
Config Duplex      : auto            Operational Duplex      : full
Config Speed       : 100             Operational Speed(Mbps) : 100
-------------------------------------------------------------------------------
Flow Control       : disabled
Dual Port          : No              Active Link        : RJ45
-------------------------------------------------------------------------------
Default VLAN       : 1               MTU[Bytes]         : 1544
MAC Learning       :
LAG ID             : N/A
============================================================================
```

Configuring Multicast Routing

Multicast is a method of sending IP datagrams   in one transmission. The Multicast group address sends and receives multicast messages. Sources use the group address as the IP destination address in their data packets. Receivers use the group address to show that they want to receive packets sent to that group.

For example, if some content is related to group 239.1.1.1, the source sends data packets destined to 239.1.1.1. Receivers for that content show that they are interested in receiving data packets sent to group 239.1.1.1. The receiver joins 239.1.1.1.

Dynamic Multicast Routing (PIM Dense Mode) Configuration

1. For each interface that uses PIM Dense mode, run:

   `> set pim interface <if_name> on`

2. Set PIM mode to Dense:

   `> set pim mode dense`

To change the PIM Multicast Routing mode between dense and sparse:

1. For each applicable interface, run:

   `> set pim interface <if_name> off`

2. For each applicable interface, run:

   For dense mode:

   `> set pim mode dense`

   For spare mode:

   `> set pim mode sparse`

3. For each applicable interface, run:

   `> set pim interface <if_name> on`

Important - You must use this procedure to change the mode. Failure to do so can cause unexpected behavior.

Validation

Run:

```
> show pim interfaces
```

Example

```
> set pim interface eth1-01 on
1_01:
success
> set pim interface eth1-02 on
1_01:
success
> set pim interface eth2-01 on
1_01:
success
> set pim mode dense
1_01:
success
> show pim interfaces
1_01:
Status flag: V - virtual  address option enabled
Mode flag: SR - state  refresh enabled
Interface      Status   State    Mode      DR Address       DR Pri     NumNbrs
eth2-01        Up       DR       dense     2.2.2.10         1          0
eth1-01        Up       DR       dense     12.12.12.10      1          0
eth1-02        Up       DR       dense     22.22.22.10      1          0
```

Multicast Restrictions

Multicast groups are addresses or address ranges. Multicast access restrictions can be defined on each interface, to allow or block multicast groups.

Configuration

In SmartDashboard edit the **Gateway Properties** > **Topology** > **Add or Edit interface** > **Multicast Restrictions** tab.

Parameter	Description
Drop multicast packets whose destination is in the list	Specifies that outgoing packets from this interface to the listed multicast destinations are dropped.
Drop all multicast packets except those whose destination is in the list	Specifies that outgoing packets from this interface to all multicast destinations except those listed, are dropped.
Add	Add a Multicast address or address range to the list.
Remove	Remove a selected Multicast address or address range from the list.

Parameter	Description
Tracking	Choose whether and how to track when Multicast packets are dropped.

Limitations

Multicast restrictions are not supported on bridge interfaces.

Multicast Acceleration

Multicast Acceleration allows SecureXL to accelerate multicast flow, in Fan-out scenarios as well.

Configuration

Multicast Acceleration is enabled by default. Use these commands to enable or disable it:

```
> sim feature mcast_route_v2 {on|off}
> fwaccel off
> fwaccel on
```

Limitations

Multicast acceleration supports IPv4 only.

Validation and Debugging

```
> fwaccel stat
-*- 4 blades: 1_01 1_02 2_01 2_02 -*-
Accelerator Status : on
Accept Templates   : enabled
Drop Templates     : disabled
NAT Templates      : enabled
Accelerator Features : Accounting, NAT, Cryptography, Routing,
                       HasClock, Templates, Synchronous, IdleDetection,
                       Sequencing, TcpStateDetect, AutoExpire,
                       DelayedNotif, TcpStateDetectV2, CPLS,McastRouting,
                       WireMode, DropTemplates, NatTemplates,
                       Streaming, MultiFW, AntiSpoofing, DoS Defender,
                       ViolationStats, Nac, AsychronicNotif, McastRoutingV2,
                       ConnectionsLimit
Cryptography Features : Tunnel, UDPEncapsulation, MD5, SHA1, NULL,
                        3DES, DES, CAST, CAST-40, AES-128, AES-256,
                        ESP, LinkSelection, DynamicVPN, NatTraversal,
                        EncRouting, AES-XCBC, SHA256
```

To show the accelerator's connections table:

Run:

```
> fwaccel conns
```

To show multicast statistics:

Run:

```
> fwaccel stats –m
```

To enable SIM debug:

Run:

```
> sim dbg -m drv + routing
```

Example

This example disables the feature.

```
> sim feature mcast_route_v2 off
-*- 4 blades: 1_01 1_02 1_03 1_04 -*-
Feature will be disabled the next time acceleration is started/restarted

> fwaccel off
-*- 4 blades: 1_01 1_02 1_03 1_04 -*-
SecureXL device disabled.

> fwaccel on
-*- 4 blades: 1_01 1_02 1_03 1_04 -*-
SecureXL device is enabled.

> fwaccel stat
-*- 4 blades: 1_01 1_02 1_03 1_04 -*-
Accelerator Status : on
Accept Templates   : enabled
Drop Templates     : disabled
NAT Templates      : enabled
Accelerator Features : Accounting, NAT, Cryptography, Routing,
                       HasClock, Templates, Synchronous, IdleDetection,
                       Sequencing, TcpStateDetect, AutoExpire,
                       DelayedNotif, TcpStateDetectV2, CPLS, McastRouting,
                       WireMode, DropTemplates, NatTemplates,
                       Streaming, MultiFW, AntiSpoofing, DoS Defender,
                       ViolationStats, Nac, AsychronicNotif
Cryptography Features : Tunnel, UDPEncapsulation, MD5, SHA1, NULL,
                        3DES, DES, CAST, CAST-40, AES-128, AES-256,
                        ESP, LinkSelection, DynamicVPN, NatTraversal,
                        EncRouting, AES-XCBC, SHA256
```

Working with Routing Tables (asg_route)

`asg_route` is an advanced utility that collects and shows routing information on all SGMs. It also makes sure that route information in the 61000/41000 Security System database is the same as the operating system routing table. This can cause routing errors if not corrected. The command also makes sure that routing information is the same on all SGMs.

This command lets you filter and customize the collected information based on different criteria, such as:

- Specified SGMs or Chassis
- Virtual Systems
- IPv4 and IPv6 addresses
- Dynamic routing protocols
- Static routes
- Source-based routes
- Inactive routes

You can run a summary report that shows the number of routes in different categories and protocols. The summary report also makes sure that the routing information is the same on all SGMs.

Basic Syntax

```
> asg_route -h
> asg_route -v
> asg_route [-a] [-b <sgm_ids>] [-6] [-vs <vs_ids>] --inactive [<filter>]
> asg_route [-a] [-b <sgm_ids>] [-6] [-vs <vs_ids>] --comp_os_db
```

Parameter	Description
`-h`	Show command syntax, help information and examples.
`-v`	Collect route information from all SGMs and save to a file at: `/var/log/asg_route/all_routes`
`-b` *<sgm_ids>*	Works with SGMs and/or Chassis as specified by *<sgm_ids>*. *<sgm_ids>* can be: • No *<sgm_ids>* specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`)
`-6`	Show IPv6 routes only (default shows IPv4 routes only).
`-a`	Show all SGMs, including those that are in the admin down state.
`--vs` *<vs_ids>*	Show the routing table only for the specified Virtual System. This option is available only for VSX environments. *<vs_ids>* can be: • No *<vs_ids>* (default) - Shows the current Virtual System context. • One Virtual System. • A comma-separated list of Virtual Systems (1, 2, 4, 5). • A range of Virtual Systems (VS 3-5). • `all` - Shows all Virtual Systems. **Note:** This parameter is only relevant in a VSX environment.
-inactive	Optional inactive route filter parameters ("Using the Advanced Filters" on page 50)
<filter>	Optional advanced routing parameters ("Using the Advanced Filters" on page 50)
`--compare-os-db`	Compares the routing data in the database with the operating system and shows: • All routes in the database that are in the operating system routing table • All routes in the operating system routing table that are not in the database

Note - You can put many basic options together on one line, but you can use only one `advanced_filter` option.

Using an SGM Filter

Example 1

This example shows a simple filter for one SGM. The route type is a one letter code in the left column and the route type codes are at the end of the list.

```
> asg_route -b 1_01
Collecting routing information, may take few seconds...
==============================================================================

Fetching Routes info from SGMs:
1_01

Routes:
C         127.0.0.0/8         is directly connected, lo
C         130.0.0.0/24        is directly connected, eth1-CIN
C         172.23.9.0/24       is directly connected, eth1-Mgmt4
C         192.0.2.0/24        is directly connected, Sync
S         0.0.0.0/0           via 172.23.9.4, eth1-Mgmt4, cost 0

Types: C - Connected, S - Static, R - RIP, B - BGP,
       O - OSPF IntraArea (IA - InterArea, E - External, N - NSSA)
       A - Aggregate, K - Kernel Remnant, H - Hidden, P - Suppressed
       SBR - Source-Based Routes
```

Example 2

This example shows a complex SGM filter that includes 4 SGMs. The results show route inconsistencies between the 61000/41000 Security System database and the operating system.

```
> asg_route -b 1_1,2_1-2_3
Collecting routing information, may take few seconds...
==============================================================================

Fetching Routes info from SGMs:
1_01,2_01,2_02,2_03

-------------------------------------------------------
Status:  DB Routes info is NOT identical on all SGMs
         OS Routes info is NOT identical on all SGMs
-------------------------------------------------------

Identical DB Routes: (21 records)
C         10.33.86.0/24       is directly connected, bond2.160
C         10.33.87.0/24       is directly connected, bond2.163
C         10.33.89.0/24       is directly connected, bond2.165
C         127.0.0.0/8         is directly connected, lo
C         192.0.2.0/24        is directly connected, Sync
C         192.168.15.128/25   is directly connected, eth1-Mgmt4
C         192.168.33.0/24     is directly connected, bond1.33
C         192.168.34.0/24     is directly connected, bond1.34
C         198.51.100.0/25     is directly connected, eth1-CIN
C         198.51.100.128/25   is directly connected, eth2-CIN
C         2.2.2.0/24          is directly connected, bond2.166
S         0.0.0.0/0           via 192.168.33.1, bond1.33, cost 0
S         16.0.0.0/24         via 10.33.86.16, bond2.160, cost 0
S         16.0.1.0/24         via 10.33.86.16, bond2.160, cost 0
S         16.0.2.0/24         via 10.33.86.16, bond2.160, cost 0
S         16.0.3.0/24         via 10.33.86.16, bond2.160, cost 0
```

```
S         16.0.4.0/24         via 10.33.86.16, bond2.160, cost 0
S         16.0.5.0/24         via 10.33.86.16, bond2.160, cost 0
S         16.0.6.0/24         via 10.33.86.16, bond2.160, cost 0
S         16.0.8.0/24         via 10.33.86.16, bond2.160, cost 0
S         194.29.40.138/32    via 192.168.15.254, eth1-Mgmt4, cost 0

Inconsistent DB Routes:
1_01:
-

2_01:
R         10.33.96.0/24       via 192.168.33.96, bond1.33, cost 2, tag 13142
R         15.0.2.0/24         via 192.168.33.96, bond1.33, cost 2, tag 13142

2_02:
-

2_03:
R         10.33.96.0/24       via 192.168.33.96, bond1.33, cost 2, tag 13142
R         15.0.2.0/24         via 192.168.33.96, bond1.33, cost 2, tag 13142

Types: C - Connected, S - Static, R - RIP, B - BGP,
       O - OSPF IntraArea (IA - InterArea, E - External, N - NSSA)
       A - Aggregate, K - Kernel Remnant, H - Hidden, P - Suppressed
       SBR - Source-Based Routes
```

Using the Summary Option (--summary)

The `--summary` parameter shows this summary information:

- Total number of routes by route type
- Summary of routes that are the same on the database and the operating system routing table
- Summary of routes where the database and the operating system are different
- OSPF interfaces and neighbors
- BGP peers

Example

```
> asg_route --summary
Collecting routing information, may take few seconds...
OSPF interfaces -
-*- 6 blades: 1_02 1_03 1_04 2_01 2_02 2_03 -*-
Name       IP Address      Area ID         State   DR Interface   BDR Interface
bond1.34   192.168.34.86   0.0.0.86        DR      192.168.34.86  0.0.0.0
bond2.163  10.33.87.1      0.0.0.91        BDR     10.33.87.88    10.33.87.1

Status: OK
==============================================================================
OSPF neighbors -
-*- 6 blades: 1_02 1_03 1_04 2_01 2_02 2_03 -*-
Neighbor        Pri     State           Address         Interface
10.33.87.88     1       FULL/DR         10.33.87.88     10.33.87.1

Status: OK
==============================================================================
BGP peers -
-*- 1 blade: 1_02 (DR Manager) -*-
PeerID            AS      State   ActRts  Routes  InUpds  OutUpds Uptime
192.168.33.96     86      Active  0       0       0       0       00:00:00

-*- 5 blades: 1_03 1_04 2_01 2_02 2_03 -*-
```

```
PeerID                  AS      State
92.168.33.96           86      Idle
192.168.34.33           161     Idle
192.168.33.94           162     Idle
192.168.34.94           162     Idle

Status: OK
==============================================================================

Fetching Summary info from SGMs:
1_02,1_03,1_04,2_01,2_02,2_03

-------------------------------------------------------
Status:  DB Summary info is NOT identical on all SGMs
        OS Summary info is identical on all SGMs
-------------------------------------------------------

Identical DB Summary: (7 records)
Total        628
aggregate    0
connected    11
igrp         0
ospf         602
rip          2
static       10

-------------------------------------------------------
Identical OS Summary: (649 records)
```

Comparing the OS Routing Table with the Database (--compare-os-db)

Use the `--compare-os-db` option to compare the routing data in the database with the operating system routing table. The output shows:

- All routes in the database that are in the operating system routing table
- All routes in the operating system routing table that are not in the database

Example

```
> asg_route --compare-os-db
Collecting routing information, may take few seconds...
==============================================================================

Fetching Routes info from SGMs:
1_01

>> Found inconsistency between routes in DB & OS

DB Routes that does not exists in OS: (7 records)
O E       10.33.92.0/24       via 10.33.87.88, bond2.163, cost 2:0
O E       12.1.145.0/24       via 10.33.87.88, bond2.163, cost 2:0
O E       12.1.146.0/24       via 10.33.87.88, bond2.163, cost 2:0
O E       12.1.147.0/24       via 10.33.87.88, bond2.163, cost 2:0
O E       12.1.148.0/24       via 10.33.87.88, bond2.163, cost 2:0
O E       12.1.149.0/24       via 10.33.87.88, bond2.163, cost 2:0
O E       12.1.150.0/24       via 10.33.87.88, bond2.163, cost 2:0
```

```
OS Routes that does not exist in DB: (6 records)
9.9.9.9 via 10.33.87.88 dev bond2.163  proto gated
12.3.0.0/24 via 10.33.87.88 dev bond2.163  proto gated
12.3.1.0/24 via 10.33.87.88 dev bond2.163  proto gated
12.3.2.0/24 via 10.33.87.88 dev bond2.163  proto gated
12.3.3.0/24 via 10.33.87.88 dev bond2.163  proto gated
12.3.4.0/24 via 10.33.87.88 dev bond2.163  proto gated
```

Using the Advanced Filters

Advanced filters let you customize the routing table display to show only the routes that you want to see. This release includes these advanced filter criteria:

Advanced Filter Criterion	Description
`--route`	Shows active routes filtered by a specified parameter
`--inactive`	Shows inactive routes filtered by a specified parameter
`--dyn-route`	Shows specified OSPF and BGP route information and makes sure that there are no inconsistencies between SGMs

Each advanced filter type has many different parameters that you can use to show a precisely filtered route list.

Advanced Filter Syntax and Parameters

You can combine many basic options on one line, but you can only use one advanced filter option at a time.

`> asg_route [<basic_options>] -n |--dyn-route <dyn_route_par>`

<dyn_route_par>	Description
`ospf`	Shows OSPF interfaces and neighbors
`rip`	Shows RIP interfaces and neighbors
`bgp`	Shows BGP peers

`> asg_route [<basic_options>] -r | --route <adv_par>`

<adv_par>	Description
`aggregate`	Shows active aggregate routes
`bgp`	Shows BGP peers
`destination` *<ip_addr>*	Shows routes to the specified destination
`direct`	Shows directly connected routes
`exact` *<ip_addr/ mask>*	Shows a route from the specified IP address
`subnets` *<ip_addr/ mask>*	Shows routes to the specified network and subnets

<*adv_par*>	Description
`ospf`	Shows OSPF interfaces and neighbors
`static`	Shows static routes
`rip`	Shows RIP interfaces and neighbors
`all`	Shows all routes (including inactive routes)

```
> asg_route [<basic_options>] -i | --inactive <inact_route_par>
```

<*inact_route_par*>	Description
`aggregate`	Shows active aggregate routes
`bgp`	Shows BGP routes
`direct`	Shows directly connected routes
`ospf`	Shows routes received from OSPF
`static`	Shows static routes
`rip`	Shows RIP Routes
`all`	Shows all routes (including inactive routes)

Advanced Filter Examples

Example 1 - BGP routes for all SGMs

```
> asg_route -b all --route bgp
Collecting routing information, may take few seconds...

===============================================================================
Fetching Routes info from SGMs:
1_01

Routes:
B         10.33.88.0/24       via 192.168.34.33, bond1.34, cost -1
B         10.33.94.0/24       via 192.168.33.94, bond1.33, cost -1
B         10.34.94.0/24       via 192.168.34.94, bond1.34, cost -1

Types: C - Connected, S - Static, R - RIP, B - BGP,
       O - OSPF IntraArea (IA - InterArea, E - External, N - NSSA)
       A - Aggregate, K - Kernel Remnant, H - Hidden, P - Suppressed
       SBR - Source-Based Routes

-------------------------------------------------------
===============================================================================
```

Example 2 - Dynamic Routing filter for OSPF neighbors

```
> asg_route --dyn-route ospf
Collecting routing information, may take few seconds...
OSPF neighbors -
-*- 1 blade: 1_01 -*-
Neighbor        Pri     State           Address         Interface
```

```
10.33.94.1      1       FULL/BDR        192.168.33.94   192.168.33.86
10.33.87.88     1       FULL/BDR        10.33.87.88     10.33.87.1

Status: OK
```

Example 3 - Inactive OSPF Routes

```
> asg_route --inactive ospf
Collecting routing information, may take few seconds...
===============================================================================

Fetching Routes info from SGMs:
1_01

Routes:
O    H i  10.33.87.0/24       is an unusable route
O    H i  192.168.33.0/24     is an unusable route
O    H i  192.168.34.0/24     is an unusable route
O E    i  194.29.40.138/32    via 10.33.87.88, bond2.163, cost 2:0

Types: C - Connected, S - Static, R - RIP, B - BGP,
       O - OSPF IntraArea (IA - InterArea, E - External, N - NSSA)
       A - Aggregate, K - Kernel Remnant, H - Hidden, P - Suppressed
       SBR - Source-Based Routes

-------------------------------------------------------
===============================================================================
```

Note - Do not use `-v` with an advanced filter or the command ignores the advanced filter and shows all the routes.

Dynamic Routing

When Dynamic Routing is enabled, one SGM is designated as the Dynamic Routing peer. This SGM is called the Dynamic Routing Manager (DR Manager). The DR Manager communicates with DR peers and updates the SGMs DR information. Before an SGM goes to the UP state, it updates its dynamic routing information from the DR Manager.

The SSM sends Dynamic Routing packets to an arbitrary SGM based on the SSM's distribution decision, not necessarily the DR Manager. If the SGM is not the DR Manager, the packets are forwarded to the DR Manager.

Unicast Routing

When an SGM that is not the DR Manager receives unicast IP routing packets, the SGM forwards them to the DR Manager. The DR Manager then communicates with its DR peers and updates the other SGMs DR information.

Use `asg_route` to administer unicast routing.

Multicast Routing

When an SGM receives multicast IP routing packets, the SGM forwards them to all other SGMs. Each SGM handles these packets on its own.

Use `asg_pim`, `asg_pim_neighbors`, and `asg_igmp` to administer multicast routing.

To see which SGM is the Dynamic Routing Manager:

Run:

```
> asg stat -i tasks
Chassis ID: 1
-------------
Task (Task ID)          SGM ID

SMO        (0)          1(local)
DR Manager (4)          1(local)
UIPC       (5)          1(local)
General    (1)          2
LACP       (2)          3
CH Monitor (3)          4

Chassis ID: 2
-------------
Task (Task ID)          SGM ID

UIPC       (5)          1
General    (1)          2
LACP       (2)          3
CH Monitor (3)          4
```

Limitation

- Only IPv4 routing protocols are supported.

Multiple OSPFv2 Instances

This release supports multiple OSPF instances. You can configure as many multiple OSPF instances as necessary. Each instance contains a fully independent OSPF database. The routes from one instance are not disclosed to other instances.

Note - In some cases, show commands refer to the vrf instance and not the OSPF instance ID.

Prerequisites

- If you use a Virtual System only for connectivity between Virtual Systems with *Per Virtual System High Availability* or *VSLS*, you **must** connect an interface to the Virtual System. See sk36980 http://supportcontent.checkpoint.com/solutions?id=sk36980 for details.
- Make sure the router ID is the same for all SGMs, but unique for each VS in the network.
- Make sure the OSPF configuration is the same on all SGMs.

Enabling OSPFv2 Multiple Instances

Make sure the default OSPF instance is configured on at least one interface:

```
> set ospf ospf-instance default on
> set ospf ospf-instance default area backbone on
> set ospf ospf-instance default interface <interface_name> area backbone on
```

You can use routemaps for route redistribution between instances, the same way you use them for redistribution between protocols.

Example

Enable Multiple OSPFv2 instances:

```
set ospf ospf-instance default on
set ospf ospf-instance default area backbone on
set ospf ospf-instance default interface bond0.173 area backbone on
set ospf ospf-instance default interface eth1-05 area backbone on
set ospf ospf-instance default export-routemap <ROUTEMAP_NAME> preference
1 on
set ospf ospf-instance 174 on
set ospf ospf-instance 174 area backbone on
set ospf ospf-instance 174 area 255.1.1.1 on
set ospf ospf-instance 174 interface bond0.174 area 255.1.1.1 on
```

Redistribute network between instances and restrict specific network:

```
set routemap <ROUTEMAP_NAME> id 1 on
set routemap <ROUTEMAP_NAME> id 1 restrict
set routemap <ROUTEMAP_NAME> id 1 match protocol ospf2
set routemap <ROUTEMAP_NAME> id 1 match ospf-instance 174 on
set routemap <ROUTEMAP_NAME> id 1 match network 75.75.12.0/24 all
set routemap <ROUTEMAP_NAME> id 3 on
set routemap <ROUTEMAP_NAME> id 3 allow
set routemap <ROUTEMAP_NAME> id 3 match ospf-instance 174 on
set routemap <ROUTEMAP_NAME> id 3 match network 75.75.0.0/16 all
set routemap <ROUTEMAP_NAME> id 3 match protocol ospf2
```

Disabling OSPFv2 Multiple Instances

To disable OSPF multiple interfaces, run:

```
set ospf ospf-instance <id> interface <interface_name> area backbone off
set ospf ospf-instance <id> area backbone off
set ospf ospf-instance <id> off
```

Monitoring

To see the OSPF interfaces, run: `# show ospf interfaces`

Sample output

```
OSPF instance default:
Name IP Address Area ID State NC DR Interface BDR Interface Errors
eth3.19 10.99.12.100 0.0.0.0 DR 0 10.99.12.100 0.0.0.0 30649
OSPF instance 1:
Name IP Address Area ID State NC DR Interface BDR Interface Errors
eth5.2011 10.95.0.161 0.0.0.0 BDR 1 10.95.0.164 10.95.0.161 8434
eth1 10.99.12.70 0.0.0.0 BDR 1 10.99.12.67 10.99.12.70 5
```

```
eth3.25 10.99.26.1 0.0.0.0 DR 1 10.99.26.1 10.99.26.2 471
eth5.29 10.99.26.130 0.0.0.0 BDR 1 10.99.26.133 10.99.26.130 0
```

To see the neighbors of an instance, run:

```
> show ospf ospf-instance <instance_id> neighbors
```

Sample output

```
> show ospf ospf-instance 1 neighbors
OSPF instance 1
Neighbor state flag: G - graceful restart
Instance Neighbor ID Pri State Dead Address Interface Errors
1 1.1.201.201 252 FULL/DR 33 10.95.0.164 10.95.0.161 0
1 1.1.11.11 1 FULL/DR 33 10.99.12.67 10.99.12.70 5
1 10.99.26.2 1 FULL/BDR 35 10.99.26.2 10.99.26.1 0
1 165.10.10.57 1 FULL/DR 39 10.99.26.133 10.99.26.130 0
```

To enable logging:

Run in clish: `> set trace ospf all on`

To see logs:

`/var/log/routed.log` (cyclic log)

`/var/log/messages`

Known Limitations

- OSPF Multiple Instances are not supported with IPv6.
- There is only one Router-ID for the entire Security Gateway/Virtual System.
- Up to 12 OSPF instances are supported.
- When doing OSPF calculations, the routing daemon can be busy and not respond to the CLI commands. This can result in unexpected output. Repeat the command if there is no response after two or more seconds.
- If you create static route SmartDashboard, it must have a valid and available IP address. Otherwise, redistribution results can be inconsistent and the router-ID process can be unstable.
- If the OSPF database on a Virtual System has two or more of the same route prefixes with equal cost, it adds the route with the lowest next hop IP address to the routing table.

Configuring DHCP Relay (set bootp)

Use BOOTP/DHCP to configure DHCP relay for a specified interface.

BOOTP/DHCP Relay extends BOOTP and DHCP operations across multiple hops in a routed network. With standard BOOTP, all LAN interfaces are loaded from one configuration server on the LAN. BOOTP Relay sends configuration requests to and from configuration servers located outside the LAN.

BOOTP/DHCP Relay has these advantages over standard BOOTP/DHCP:

- Relay client configuration requests

  Configure an interface on the Check Point system to relay requests to multiple servers. This provides redundancy.

  Configuration requests are sent to all configured relay servers simultaneously.

- Load balancing

  Configure interfaces to relay client configuration requests to different relay servers.

- Centrally manage client configuration over multiple LANs

  This is very useful in large enterprise environments.

Syntax

```
> set bootp interface <if_name> [primary default|<ip>] [wait-time <seconds>]
[relay-to <ip1>,<ip2>...] on|off
```

Parameter	Description
`interface` *<if_name>*	The interface name as defined by the system. Press **Tab** after you enter this parameter to see a list of valid interface names.
`primary` *<ip>* or `primary default`	The IP address of the Security Gateway interface that always gets requests from the DHCP client. If you do not define a Primary IP address, the system automatically uses the IP address of the interface that the DHCP request comes from. You can use the `default` value instead of an IP address. This forces the system to use the IP address of the interface that the DHCP request comes from. This is useful when you want to change the `wait-time` parameter, but not define a Primary IP.
`wait-time` *<seconds>*	The minimum wait time, in seconds, before a BOOTP request can be sent. This includes the elapsed time after the client starts to boot. This delay lets a local configuration server reply, before it sends the relay to a remote server. The `wait-time` keyword is optional. The system assumes that an integer after the `primary` value is the `wait-time` value. Valid values: 0 - 65535 Default - 60
`relay-to` *<ip>*	The IP address of the relay server to which BOOTP requests are sent. You can specify more than one server.
`on|off`	Enables or disables BOOTP on the specified interface.

Example 1

This example enables DHCP Relay on `eth0-4` with default values and no Primary IP. The IP address is automatically assigned by the DHCP server.

```
> set bootp interface eth0-04 on
```

Example 2

This example enables DHCP Relay on `eth0-04` and defines the Primary IP address as `30.30.30.1`. The wait time is the default value (60 seconds).

```
> set bootp interface eth0-04 primary 30.30.30.1 wait-time default on
```

Example 3

This example enables DHCP Relay on `eth1-o4` and sends BOOTP requests to the relay server at `20.20.20.200`.

```
> set bootp interface eth1-04 relay-to 20.20.20.200 on
```

Verification

Use this command to monitor and troubleshoot the BOOTP implementation:

```
> show bootp interface|interfaces|stats
```

Parameter	Description
`interface`	BOOTP/DHCP Relay Interface
`interfaces`	All BOOTP/DHCP Relay Interfaces
`stats`	BOOTP/DHCP Relay Statistics

Destination-Based Routing

Destination-Based Routing has advantages over Source-Based Routing.

Advantages

- Protects against cache pollution by an attacker.
- Better performance than Source-Based Routing.

Limitations

- ECMP requires Source-Based Routing.
- If you use Policy-Based Routing in a FROM rule, Destination-Based Routing is disabled.

Configuring Destination-Based Routing

You can enable Destination-Based Routing permanently or temporarily.

Use `asg_dst_route` to manage Destination-Based Routing.

Syntax

```
> asg_dst_route [-e|-d|-a|-v][-b <id>] [--g {increase |decrease}]
```

Parameter	Description
`-b` *<id>*	Valid values: No SGM or Chassis specified `all` - shows all SGMs and Chassis One ID of an SGM or name of a Chassis Comma-separated list (1_1,1_4) Range (1_1 - 1_4) The Active Chassis (chassis_active)
`-g {increase |decrease}`	Increase or decrease route cache garbage aggressiveness
`-e`	Enable Destination-Based Routing
`-d`	Disable Destination-Based Routing
`-a`	Restore Destination-Based Routing to the default
`-v`	Show the current and persistent status of Destination-Based Routing

Example

```
> asg_dst_route -v
```

Output

```
+------------------------------------------------------------------------+
| SGM        | Current Status            | Persistency                   |
+------------------------------------------------------------------------+
| 1_01       | Source+Dest based route   | Source+Dest based (due to pbr rule) |
| 1_02       | Source+Dest based route   | Source+Dest based (due to pbr rule) |
+------------------------------------------------------------------------+
```

You can use `echo` *<value>* `> /proc/sys/net/ipv4/route/src_mask` to temporarily enable or disable Destination-Based Routing.

Syntax

```
> echo <value> > /proc/sys/net/ipv4/route/src_mask
```

Parameter	Description
<value>	Enable or disable Destination-Based Routing Allowed values: • `0` - Enable Destination-Based Routing • `-1` - Disable Destination-Based Routing

Destination-Based Routing Statistics

asg_dst_route -s

Use `asg_dst_route -s` to show summary Destination-Based Routing statistics.

Syntax

```
asg_dst_route -s
```

Output

```
+------------------------------------------------+
| SGM  | Cache load | Hit rate   | Effectiveness |
+------------------------------------------------+
| 2_01 | 4 %        | 58 %       | 100 %         |
| 2_02 | 4 %        | 50 %       | 100 %         |
+------------------------------------------------+
```

Column	Description
SGM	SGM ID
Cache load	Calculated percentage of how hard the route cache works The route cache starts to clean when this exceeds 100%.
Hit rate	Percent of total lookups found in the route cache in the last 3 seconds
Effectiveness	Estimate of the effectiveness of the route cache This values drops if the cache load increases or hit rate decreases.

Managing the 61000/41000 Security System

In This Section:

Administration

Working with Global Commands

The 61000/41000 Security System operating system includes a set of global commands that apply to all or specified SGMs in a system.

- gclish commands apply globally to all SGMs, by default.
- Some gclish commands are applicable to the 61000/41000 Security System and its components.
- gclish commands do not apply to SGMs that are DOWN. If you run a `set` command while an SGM is DOWN, the command does not update that SGM. The SGM synchronizes its database during startup and applies the changes after reboot.
- clish commands are documented in the R76 Gaia Administration Guide http://supportcontent.checkpoint.com/documentation_download?ID=22928. Most of these commands are also available in the 61000/41000 Security System.

Note - Documentation for the Chassis feature is found in the Hardware Monitoring and Chassis High Availability ("Working with Active/Standby High Availability" on page 191) sections.

Global Commands

`auditlog`

- Enabled by default.
- All commands are recorded in the audit log.
- To learn more about the audit log, see Looking at the Audit Log ("Looking at the Audit Log File (asg_auditlog)" on page 181).

`config-lock`

- Protects the gclish database by locking it. Each SGM has a single lock.

- To set gclish operations for an SGM, the SGM must hold the `config-lock`.

  To set `config-lock`, run:

  ```
  # set config-lock on override
  ```

- gclish traffic runs on the Sync interface, port 1129/TCP.

`blade-range`

- Runs commands on specified SGMs.
- Runs gclish embedded commands only on this subset of SGMs.
- We do not recommend that you use the `blade-range` command, because all SGMs must have identical configurations.

Check Point Global Commands

Global commands are scripts that run commands on more than one SGM. This section includes Check Point product-related commands, such as `fw`, `sim`, `fwaccel,` and `cpconfig`.

fwaccel, fwaccel6

`fwaccel` and `fwaccel6` dynamically enable or disable acceleration for IPv4 traffic while the 61000/41000 Security System is in operation. `fwaccel6` is used for IPv6 traffic and resets to the default value of `fwaccel` after reboot.

Run `fwaccel` and `fwaccel6` from gclish, to comparatively show combined information from all SGMs.

Syntax

```
> fwaccel {on|off|stat|stats [-s} [-d] |conns [-s] -m <max_entries> [-b <sgm_ids>]
> fwaccel templates[-s] [-m <max_entries>] [-b <sgm_ids>]

> fwaccel6 {on|off|stat|stats [-s} [-d] |conns [-s] -m <max_entries> [-b <sgm_ids>]
> fwaccel6 templates[-s] [-m <max_entries>] [-b <sgm_ids>]
```

Parameter	Description
`-b`	Works with SGMs and/or Chassis as specified by <*sgm_ids*>. <*sgm_ids*> can be: • No <*sgm_ids*> specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`) Note: You can only select SGMs from one Chassis with this option.
`on`	Starts acceleration
`off`	Stops acceleration

Parameter	Description
`stat`	Shows the acceleration device status and the status of the Connection Templates on the local Security Gateway.
`stats`	Shows acceleration statistics.
`stats -s`	Shows more summarized statistics.
`stats -d`	Shows dropped packet statistics.
`conns`	Shows all connections.
`conns -m` <*max_entries*>	Limits the number of connections displayed by the conns command to <*max_entries*>.
`templates`	Shows all connection templates.
`templates -m` <*max_entries*>	Limits the number of templates displayed by the templates command to <*max_entries*>.
`templates -s`	Shows the number of templates currently defined in the accelerator.

Example

```
> fwaccel stats
Displaying aggregated data from blades: all

Name                  Value              Name                  Value
--------------------  ---------------    --------------------  --------------
Accelerated Path
-----------------------------------------------------------------------------
accel packets             6518    accel bytes                  870476
conns created            38848    conns deleted                 38043
C total conns              801    C templates                       0
C TCP conns                493    C delayed TCP conns               0
C non TCP conns            308    C delayed nonTCP con              0
conns from templates         0    temporary conns                   0
nat conns                    0    C nat conns                       0
dropped packets              0    dropped bytes                     0
nat templates                0    port alloc templates              0
conns from nat tmpl          0    port alloc conns                  0
Policy deleted tmpl          0    C Policy deleted tmp              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
AH enc pkts                  0    AH enc err                        0
AH dec pkts                  0    AH dec err                        0
AH other err                 0    espudp enc pkts                   0
espudp enc err               0    espudp dec pkts                   0
espudp dec err               0    espudp other err                  0

Medium Path
-----------------------------------------------------------------------------
```

```
PXL packets                  0    PXL async packets                 0
PXL bytes                    0    PXL conns                         0
C PXL conns                  0    C PXL templates                   0

Firewall Path
-----------------------------------------------------------------------------
F2F packets           10077862    F2F bytes                1185051123
F2F conns                38839    C F2F conns                    800
TCP violations               0    C partial conns                  0
C anticipated conns          0

General
-----------------------------------------------------------------------------
memory used                  0    free memory                      0

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

Monitor Mode

`fwaccel_m` continuously monitors `fwaccel` output in real time. In Monitor Mode, the screen shows changes in parameters as highlighted text. You cannot run commands or other operations while in Monitor mode.

To close Monitor mode, press **Ctl-c**.

Example

```
> fwaccel_m stats -p
```

Output

```
Displaying aggregated data from blades: all

F2F packets:
--------------
Violation             Packets            Violation             Packets
--------------------  ---------------    --------------------  ---------------
pkt is a fragment                   0    pkt has IP options              19286
ICMP miss conn                     33    TCP-SYN miss conn               28713
TCP-other miss conn            125290    UDP miss conn                95373635
other miss conn                268865    VPN returned F2F                    0
ICMP conn is F2Fed               5390    TCP conn is F2Fed               73812
UDP conn is F2Fed                9131    other conn is F2Fed              4827
unidirectional viol                0    possible spoof viol                 0
TCP state viol                      0    out if not def/accl                 0
bridge, src=dst                     0    routing decision err               82
sanity checks failed                0    temp conn expired                   0
fwd to non-pivot                    0    broadcast/multicast                 0
cluster message                     0    partial conn                        1
PXL returned F2F                    0    cluster forward                     0
chain forwarding                    0    general reason                      0
port alloc f2f                      0    sticky SA F2F                       0
```

fw, fw6

The `fw` and `fw6` commands are global scripts that run the `fw` and `fw6` commands on each SGM.

Example 1

```
> fw ctl
```

Output

```
-*- 6 blades: 1_01 1_02 1_03 2_01 2_02 2_03 -*-
Usage: fw ctl command args...
Commands: install, uninstall, pstat, iflist, arp, debug, kdebug, bench
chain, conn
```

Example 2

```
> fw ctl iflist
```

Output

```
-*- 6 blades: 1_01 1_02 1_03 2_01 2_02 2_03 -*-
0 : BPEth0
1 : BPEth1
2 : eth1-Mgmt4
3 : eth2-Mgmt4
4 : eth1-01
5 : eth1-CIN
6 : eth2-CIN
8 : eth2-01
16 : Sync
17 : eth1-Mgmt1
18 : eth2-Mgmt1
```

fw dbgfile

Use these commands to debug the system:

- `fw dbgfile collect` - Collects firewall debugging information (`fw ctl debug`).
- `fw dbgfile view -` Shows the collected debugging information

Syntax

```
> fw dbgfile collect -f <debug_file_path> [-buf <buf_size>] [<fw_flags>]
> fw dbgfile view [<debug_file_path>] [-o <agg_file_path>]
```

Parameter	Description
<debug_file_path>	Full path of the debug file
`-buf` *<buf_size>*	Buffer size
<fw_flags>	Firewall flags
-o <agg_file_path>	Use an aggregate debug file *<agg_file_path>* - Full path of the aggregate debug file

Example - Collect Debug Information

```
> fw dbgfile collect -f /home/admin/temp.dbg -buf 2300 -m kiss + pmdump -m fw + xlate
```

Example - See Debug Information

```
> fw dbgfile view /home/admin/temp.dbg
```

Global Commands Generated by CMM

The CMM monitors and controls Chassis components and activates and shuts down SGMs and SSMs.

SGMs are shut down in serious situations, for example, when the Sync Interface cannot access the SGM. In that case, the `reboot` command does not work.

Commands that control SGM power from the CMM:

- `> asg_reboot` <*global_command_flags*> – Restart SGMs
- `> asg_hard_shutdown` <*global_command_flags*> – Turn off SGMs
- `> asg_hard_start` <*global_command_flags*> – Turn on SGMs

To learn more about <*global_command_flags*>, see Global Operating System Commands. You can run global commands from `gclish` and the Expert mode.

Example

```
> asg_reboot -b 1_03,2_05
You are about to perform hard reboot on SGMs: 1_03,2_05
It might cause performance hit for a period of time

Are you sure? (Y - yes, any other key - no) Y

Hard reboot requires auditing
Enter your full name: User1
Enter reason for hard reboot [Maintenance]:
WARNING: Hard reboot on SGMs: 1_03,2_05, User: User1, Reason: Maintenance

Rebooting SGMs: 1_03,2_05
```

Notes

- At least one SGM must be UP and running on the remote Chassis to run these commands.
- To learn how to restart an SSM from the CMM, see asg_chassis_ctrl (Check Point - http://www.checkpoint.com).

General Global Commands

Global commands run commands on more than one SGM. The global commands syntax is shown in Global Operating System Commands.

These commands are available in the gclish and clish:

gclish name	bash name
`update_conf_file`	`g_update_conf_file`
`global`	`global_help`
`asg_cp2blades`	`asg_cp2blades`
`asg_clear_table`	`asg_clear_table`
`asg_clear_messages`	`asg_clear_messages`
`asg_blade_stats`	`asg_blade_stats`

Update Configuration Files (update_conf_file)

Use this command to add, update, and remove variables from configuration files. If the file does not exist, this command creates it.

Syntax

> `update_conf_file` *<file_name>* *<var>*=*<value>*

Parameter	Description
<file_name>	Full path and name of the configuration file to update You do not need to specify the path for these files: • `fwkern.conf` • `simkern.conf`
<var>	Variable to update
<value>	New value for the variable

Examples

```
> update_conf_file /home/admin/MyConfFile.txt var1=hello
> cat /home/admin/MyConfFile.txt
-*- 3 blades: 2_01 2_02 2_03 -*-
var1=hello

> update_conf_file /home/admin/MyConfFile.txt var2=24h
> cat /home/admin/MyConfFile.txt
-*- 3 blades: 2_01 2_02 2_03 -*-
var2=24h
var1=hello

> update_conf_file /home/admin/MyConfFile.txt var1=goodbye
> cat /home/admin/MyConfFile.txt
-*- 3 blades: 2_01 2_02 2_03 -*-
var2=24h
var1=goodbye

> update_conf_file /home/admin/MyConfFile.txt var2=
> cat /home/admin/MyConfFile.txt
-*- 3 blades: 2_01 2_02 2_03 -*-
var1=goodbye
```

Configuration file required format:

This command works with configuration files composed of lines where each line defines one variable.

<variable>=*<value>*

Variable name must not include: =

Note - `fwkern.conf` and `simkern.conf` use this format.

Setting Sim Kernel Parameters

Use the `sim_param` command to change or show sim parameter values. Run these commands in the Expert mode.

Syntax

```
sim_param show [<filter>]
sim_param get <parameter>
sim_param set <parameter> <value>
sim_param save <file_name>
```

Parameter	Description
`show`	Shows a detailed list of all sim parameters
<*filter*>	Shows only those sim parameters that contain the specified text string
`get`	Shows the value of the specified sim parameter
`set`	Set the specified sim parameter to the specified value
<*parameter*>	Sim parameter to set or show
<*value*>	Sim parameter value
`save`	Save the sim parameters to the specified file name
<*file_name*>	Sim parameter file name

Note - To keep parameters from changing, manually edit the applicable parameters in `$PPKDIR/boot/modules/simkern.conf`. Use the `g_update_conf_file` command.

sim_param Examples

Example 1

This example shows the usage of the `sim_param set` command.

```
# sim_param set sim_mcast_silent_spoof 0
sim_mcast_silent_spoof successfully changed to 0.
```

Example 2

This example shows a filtered list of sim parameters. You can see that the `sim_mcast_silent_spoof` value changed to 0 from the above example.

```
# sim_param show mcast
+------------------------------------------+-----------------+----------+-----------+
|Name                                      |Value            |Default   |Permission |
+------------------------------------------+-----------------+----------+-----------+
|sim_build_mcast_entry_disabled            |1                |Identical |R/W        |
|sim_drop_standby_mcast                    |1                |Identical |R/W        |
|sim_db_mcast_drop_tmo                     |15               |Identical |R/W        |
|sim_drop_mcast_exceptions                 |<NULL>           |Identical |R/W        |
|sim_mcast_packets_to_f2f                  |10               |Identical |R/W        |
|sim_mcast_drop_refresh_route              |1                |Identical |R/W        |
|sim_drop_mcast_on_standby                 |1                |Identical |R/W        |
|reserved_mcast_check                      |1                |Identical |R/W        |
|sim_mcast_drop_refresh_f2f                |1                |Identical |R/W        |
|sim_mcast_silent_spoof                    |0                |1         |R/W        |
+------------------------------------------+-----------------+----------+-----------+
```

Example 3

This example shows how to show a list of all the sim parameters and their current values. This command is also useful to show all of the correct parameter names for use with the `sim_param set` command.

```
# sim_param show
+-------------------------------------------+-----------------+----------+-----------+
|Name                                       |Value            |Default   |Permission |
+-------------------------------------------+-----------------+----------+-----------+
|sim_reuse_tcp_conn                         |1                |Identical |R/W        |
|sim_gtp_inner_frag_accel                   |1                |Identical |R/W        |
|sim_drop_percentage_to_check_overall_drops|35               |Identical |R/W        |
|sim_bond_refresh_interval_ha               |1                |Identical |R/W        |
|sim_aff_min_accel_pkts_rate                |250000           |Identical |R/W        |
|bridge_mode_on_ssm60                       |0                |Identical |R/W        |
|           •                               |                 |          |           |
|           •                               |                 |          |           |
|           •                               |                 |          |           |
|sim_ntquota_pxl_only                       |0                |Identical |R/W        |
|sim_mcast_silent_spoof                     |1                |Identical |R/W        |
|sim_hlqos_log_interval                     |2                |Identical |R/W        |
+-------------------------------------------+-----------------+----------+-----------+
```

Setting Firewall Kernel Parameters (g_fw ctl set)

Use these commands to set or show specified Firewall kernel parameters. Run these commands in the Expert mode.

Syntax

```
g_fw ctl get <type> <parameter_name>
g_fw ctl set <type> <parameter_name> <value>
```

Parameter	Description
`get`	Shows the specified parameter and its value
`set`	Change the parameter value to the specified value
<*type*>	Type of parameter value: `int` - Integer value `string` - String value **Note**: You must enter the correct parameter type or the command returns an error message. Run this command to see a list of valid parameters: `# modinfo $FWDIR/modules/fwmod.2.6.18.cp.x86_64.o`
<*parameter_name*>	Parameter name
<*value*>	Parameter value

Note - To make changes persistent, you must manually edit the applicable parameters in `$FWDIR/boot/modules/fwkern.conf`. Use the `g_update_conf_file` command to do this.

Copy Files to Blades (asg_cp2blades)

Use this command to copy files from the current SGM to other SGMs.

Syntax

> asg_cp2blades [-b <*sgm_ids*>][-s] <*source_path*> [<*dest_path*>]

Parameter	Description
-b <*sgm_ids*>	Works with SGMs and/or Chassis as specified by <*sgm_ids*>. <*sgm_ids*> can be: • No <*sgm_ids*> specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`)
-s	Save a local copy of the old file on each SGM The copy is saved in the same directory as the new file. The old file has the same name with this at the end: .bak.<*date*>.<*time*>
<*source_path*>	Full path and name of the file to copy
<*dest_path*>	Full path of the destination If not specified, the command copies the file to the relative source file location.

Example

```
> asg_cp2blades /home/admin/note.txt
```

Output

```
Operation completed successfully
```

Verification

```
> cat /home/admin/note.txt
-*- 3 blades: 2_01 2_02 2_03 -*-
hello world
```

global help

This command shows the list of global commands you can use in gclish, and how they are generally used.

Syntax

```
> global help
```

Output

```
> global help
Usage: <command_name> [-b SGMs] [-a -l -r --] <native command arguments>
Executes the specified command on specified blades.

Optional Arguments:
  -b   blades: in one of the following formats
              1_1,1_4 or 1_1-1_4 or 1_01,1_03-1_08,1_10
              all (default)
              chassis1
              chassis2
              chassis_active
  -a        : Force execution on all SGMs (incl. down SGMs).
  -l        : Execute only on local blade.
  -r        : Execute only on remote SGMs.

Command list:
arp cat cp cpconfig cplic cpstart cpstop dmesg ethtool fw fw6 fwaccel fwaccel6
fwaccel6_m fwaccel_m ls md5sum mv netstat reboot sim sim6 snapshot_recover
snapshot_show_current tail tcpdump top unlock update_conf_file vpn asg
```

asg_clear_table

Use this command to delete connections from the firewall connection table. The command runs up to 15 times, or until there are less than 50 connections.

Syntax

```
> asg_clear_table [-b <sgm_ids>]
```

Parameter	Description
-b <*sgm_ids*>	Works with SGMs and/or Chassis as specified by <*sgm_ids*>. <*sgm_ids*> can be: • No <*sgm_ids*> specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`) Note: You can only select SGMs from one Chassis with this option.

Note - If you connected to the machine with SSH, your connection is disconnected

asg_clear_messages

Use `asg_clear_messages` to clear all messages in `/var/log/messages` files.

Syntax

```
> asg_clear_messages
```

Output

```
This action will erase the messages in /var/log/messages
and will be executed on blades: all
Are you sure? (Y - yes, any other key - no) y
Command completed successfully
```

show interface

Example

```
> show interface eth1-01 ipv4-address
1_01:
ipv4-address 4.4.4.10/24

1_02:
ipv4-address 4.4.4.10/24

1_03:
ipv4-address 4.4.4.10/24

1_04:
ipv4-address 4.4.4.10/24

1_05:
Blade 1_05 is down. See "/var/log/messages".

2_01:
ipv4-address 4.4.4.10/24

2_02:
ipv4-address 4.4.4.10/24

2_03:
ipv4-address 4.4.4.10/24

2_04:
ipv4-address 4.4.4.10/24

2_05:
ipv4-address 4.4.4.10/24
```

Configuring Chassis state (asg chassis_admin -c)

Use this command to put a Chassis in the administrative UP or DOWN state. You must have administrator permission to do this.

When a Chassis is in the Administrative DOWN state:

- Backup connections for SGMs are lost.
- New connections are not synchronized with the DOWN Chassis.

Syntax

```
> asg chassis_admin -c <chassis_id> down|up
```

Parameter	Description
<chassis_id>	Chassis identification number (1 or 2)
down|up	Chassis state

Example

```
> asg chassis_admin
You are about to perform Chassis_admin down on Chassis: 2
Are you sure? (Y - yes, any other key - no) y
Chassis_admin down requires auditing
Enter your full name: John
Enter reason for chassis_admin down [Maintenance]: test
WARNING: Chassis_admin down on Chassis: 2, User: John, Reason: test
Chassis 2 is going DOWN...
Chassis 2 state is DOWN
```

Notes

- This command is audited in the `asg log audit`.
- Run this command to see the Chassis state:
  `> asg stat /monitor`
- In a Dual Chassis environment, a Chassis in the administrative DOWN state causes degradation of the system performance

Synchronizing SGM Time (asg_ntp_sync_config)

`asg_ntp_sync_config` uses an NTP server to synchronize the local time for all SGMs and the CMM.

- If the refresh time is less than the default of 300 seconds, refresh occurs every 300 seconds.
- Disable the **replies_from_any_port** property for the **NTP over UDP** service, to allow time synchronization for all SGMs.
  - In GuiDBEdit, search for the **NTP/UDP** service.
  - Go to the **replies_from_any_port** property.
  - Change the property to **false**.
  - Install the policy.

Syntax

```
> asg_ntp_sync_config set primary|secondary <ip>|<hostname> [-v <version>] [-r <timeout>]
> asg_ntp_sync_config disable|enable|delete|show
> asg_ntp_sync_config -h
```

Parameter	Description
set	Configure an NTP server
primary	The system uses this NTP server by default

Parameter	Description
`secondary`	The system uses this if the primary NTP server is not available
`NTP Server <`*ip*`>|<`*hostname*`>`	NTP server IP address or host name
`-v <`*version*`>`	Server version of the NTP Service (default = NTPv4)
`timeout`	Timeout in seconds between refreshes (default = 300 seconds)
`show`	Show NTP Server configuration
`disable`	Disable NTP service
`enable`	Enable NTP service
`delete`	Delete primary or secondary NTP Service
`-h`	Show syntax and help information

Validation

To confirm that the time is the same on all SGMs enter:

```
> show time
```

To confirm that all SGMs start NTP connections, run:

`tcpdump` on port 123/UDP for the applicable interfaces

Configuring SGMs (asg_blade_config)

Manage SGMs with `asg_blade_config.`

- Copy the SGM configuration from a different SGM
- Change the synchronization start IP address
- Reset the system uptime value
- Get a policy from the Security Management server

Syntax

```
# asg_blade_config pull_config [policy|all] [-force] <ip>
# asg_blade_config full_sync <ip>
# asg_blade_config set_sync_start_ip <ip>
# asg_blade_config reset_uptime
# asg_blade_config reset_uptime_user
# asg_blade_config get_smo_ip
# asg_blade_config is_in_security_group
# asg_blade_config is_in_pull_conf_group|config fetch_smc
# asg_blade_config upgrade_start <new_version> [cu]
# asg_blade_config upgrade_stop|upgrade_stat|upgrade_cu
```

Parameter	Description
`pull_config`	Copy the configuration from another SGM.
`full_sync` <*ip*>	Run a full synchronization from another SGM. <*ip*> - Synchronization interface on remote SGM
`set_sync_start_ip` <*ip*>	Changes the Synchronization start IP address from the local SGM to the specified IP address.
`reset_uptime`	Resets the system uptime value on all SGMs to the current time.
`reset_uptime_user`	An interactive command that resets the uptime for all SGMs to a user configured time.
`get_smo_ip`	Return the Synchronization IP address of the Single Management Object, as defined in SmartDashboard. This address is not shown in SmartDashboard.
`is_in_security_group`	Make sure that the local SGM is in the Security Group.
`is_in_pull_conf_group`	Make sure that the local SGM is in the Pulling Configuration Group. If not, the SGM cannot copy the configuration and policy.
`config fetch_smc`	Get the policy from the Security Management Server, and send it to all SGMs.
`upgrade_start` <*new_version*> `[cu]`	Start the upgrade procedure. <*new_version*> - New version name. `[cu]` - Specifies the Connectivity upgrade.
`upgrade_stop`	Stop the upgrade procedure.
`upgrade_stat`	Shows the upgrade procedure and policy status.
`upgrade_cu`	Change from Zero Downtime upgrade to Connectivity upgrade.

Troubleshooting `asg_blade_config`

To troubleshoot problems associated with the `asg_blade_config` command, examine the logs stored at: `/var/log/blade_config`

For example, if the SGM unexpectedly reboots, you can search the log file for the word `reboot` to learn why.

Backup and Restore

It is a best practice to back up your 61000/41000 Security System operating system configuration periodically and before you upgrade or make major changes to the system. You can always restore a saved configuration as necessary. The backup is saved to a .tgz file.

Restoring a Configuration

- To restore a backup from a locally held file, run:

  `> set backup restore local <file>`

- To restore a backup from a remote server using ftp, run:

  `> set backup restore ftp <ip_address> username <name> password <password> <file>`

- To restore a backup from a remote server using tftp, run:

  `> set backup restore tftp <ip_address> file <file>`

- To restore a backup from a remote server using scp, run:

  `> set backup restore scp ip <ip_address> username <name> password <password> <file>`

Parameter	Description
<*file*>	Name of the backup file
<*ip_address*>	The IP address of the ftp. tftp, or scp remote server
<*name*>	User name to log in to the remote server
<*password*>	Remote server password

Example:

```
> set  backup restore  ftp ip 192.0.2.24 username user1 password pass1 file
backup_gw-24_17_4_2012_11_07.tgz

Restoring from backup package. Use the command 'show backups' to monitor restoring
progress.
Please reboot the machine when it's finished.
```

Configuring SGM state (asg sgm_admin)

Use this command to manually change the state, UP or DOWN, for one or more SGMs.

Syntax

```
> asg sgm_admin -b <sgm_ids> up|{down [-a]} [-p]
> asg sgm_admin -h
```

Parameter	Description
`-b` <*sgm_ids*>	Works with SGMs and/or Chassis as specified by <*sgm_ids*>. <*sgm_ids*> can be: • No <*sgm_ids*> specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`)

-p	Persistent. The setting is kept after reboot
-a	Synchronize accelerated connections to other SGMs
-h	Show command syntax and help information

Example

```
> asg sgm_admin -b 2_03 -p
You are about to perform blade_admin up on blades: 2_03

Are you sure? (Y - yes, any other key - no) y

Blade_admin up requires auditing
Enter your full name: Fred
Enter reason for blade_admin up [Maintenance]: test
WARNING: Blade_admin up on blades: 2_03, User: Fred, Reason: test

Performing blade_admin up on blades: 2_03
[2_03]Setting blade to normal operation ...
[2_03]pulling configuration from: 192.0.2.16 (may take few seconds)
[2_03]Blade current state is ACTIVE
```

Notes

- When an SGM is in the **Administrative DOWN** state:
  - gclish commands do not run on this SGM.
  - Traffic is not sent to this SGM.
  - `asg stat` shows the SGM as **DOWN (admin)**.
- When an SGM is changed to Administrative UP, it automatically synchronizes the configuration from a different SGM that is in the UP state.
- This command generates log entries. To show the logs, run:
  ```
  > asg log audit
  ```
- This command is useful for debugging. We do not recommend that you use it in production environments because it causes performance degradation.

Image Management

You can:

- **Revert** to a saved image. This restores the system, including the configuration of the installed products.
- **Delete** an image from the local system.
- **Export** an existing image. This creates a compressed version of the image. You can download the exported image to a different computer and delete the exported image from the Gaia computer. This saves disk space. You must not rename the exported image. If you rename a snapshot image, it is not possible to revert to it.
- **Import** an exported image.
- See a list of saved images.

Global Image Management - (snapshot)

Use this command to create, import, export, and show snapshots for all SGMs in the 61000/41000 Security System.

To create a new image:

```
> add snapshot <snapshot_name> desc <description>
```

To monitor the snapshot creation process or view a list of existing snapshots:

```
> show snapshots
```

To delete an image:

```
> delete snapshot <snapshot_name>
```

To export or import an image, or to revert to an image:

```
> set snapshot import|export <snapshot_name> path <path>
> set snapshot revert <snapshot_name>
```

To show image information:

```
> show snapshot <snapshot_name> all|date|desc|size
```

Parameter	Description
`snapshot` *<snapshot_name>*	Name of the image
`desc` *<desc>*	Description of the image
`snapshot export` *<snapshot_name>*	Name of the image to export
`snapshot import` *<snapshot_name>*	Name of the image to import
`path` *<path>*	Location for the exported image For example: `/var/log`
`all`	All image details
`date`	Date the image was made
`desc`	Description of the image
`size`	Size of the image

Notes

- You must have sufficient available space on the backup partition to create snapshot images for all SGMs. The required available disk space is the actual size of the root partition, multiplied by 1.15.
- The available space required in the export file storage location is the size of the snapshot multiplied by two.
- The minimum size of a snapshot is 2.5G. Therefore, the minimum available space necessary in the export file storage location is 5G.

Image Management for Specified SGMs (g_snapshot)

Show and revert snapshots for specified SGMs or Chassis. This is different from `snapshot`, which works for all SGMs together. You must run this command from Expert mode.

Syntax

```
# g_snapshot [-b <sgm_ids>] show|[revert <snapshot_name>]
```

Parameter	Description
`show`	Shows saved snapshots for the specified SGMs or Chassis.
`revert`	Restore specified SGMs or Chassis to the specified snapshot.
<snapshot_name>	Snapshot file name
<sgm_ids>	Works with SGMs and/or Chassis as specified by *<sgm_ids>*. *<sgm_ids>* can be: • No *<sgm_ids>* specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`)

Examples

- `# g_snapshot -b 1_1,1_4 revert My_Snapshot`

  This example restores SGMs 1_1 and 1_4 to `My_Snapshot`.
- `# g_snapshot –b chassis2 revert My_Snapshot`

  This example restores Chassis2 to `My_Snapshot`.
- `# g_snapshot -b chassis1 show`

  This example shows the saved snapshots for all SGMs on Chassis1.

Setting Blade-Range

Use the blade-range command to activate Software Blades.

Syntax

```
> set blade-range <Chassis-ID>_<Blade-ID> - <Chassis-ID>_<Blade-ID>
```

Parameter	Description
Chassis-ID	valid values: `1` or `2`
Blade-ID	valid values: • `1` to `12` • `all` (does not work on VSX)

Port Mirroring (SPAN Port)

Port Mirroring lets a gateway listen to traffic on a mirror port or SPAN port on a switch. The mirror port on a Check Point gateway is typically configured to monitor and analyze network traffic with no effect on the physical network. The mirror port duplicates the network traffic and records the activity in logs.

You can use mirror ports to:

- Monitor the use of applications in your organization, as a permanent part of your deployment
- Evaluate the capabilities of the Application Control and IPS Software Blades before you purchase them

The mirror port does not enforce a policy. You can only use it to see the monitoring and detection capabilities of the blades.

Benefits of a mirror port include:

- There is no risk to your production environment.
- It requires minimal set-up configuration.
- It does not require expensive TAP equipment.

Configuring Port Mirroring on a Security Gateway

To configure a port mirroring log:

1. Create a new bridge group:
   ```
   > add bridging group 0
   ```
2. Add the interface to bridging group `br0:`
   ```
   > add bridging group 0 interface <if_name>
   ```
   <*if_name*> = Interface name
3. In SmartDashboard, manually add the bridge interface to the 61000/41000 Security System gateway object.
4. Change the bridge interface name to **br0**.
5. Select **Global Properties** from the **Policy** menu.
6. Select **Stateful Inspection** and clear these options:
   - **Drop out of state TCP packets**
   - **Drop out of state ICMP packets**
7. Install the policy.
8. From the 61000/41000 Security System command line, define the interface as a SPAN port:
   ```
   > asg_span_port set <br_if_name>
   ```
9. Reboot all SGMs.
10. In **Global Properties > Stateful Inspection > Exceptions**, add an exception for the 61000/41000 Security System.

We recommend that you run `asg if` to make sure that the bridge and its related interface are up and running.

Configuring Port Mirroring for a VSX Gateway

To configure port mirroring for a VSX Gateway:

1. In SmartDashboard, create a new Virtual System in the Bridge mode.
2. Add an interface for the SPAN port that is connected to the physical port of the SSM.
3. Select **Global Properties** from the **Policy** menu.
4. Select **Stateful Inspection** and clear these options:
   - **Drop out of state TCP packets**
   - **Drop out of state ICMP packets**
5. Install the policy on the Virtual System.
6. Open an SSH connection to the VSX Gateway.
7. From the new Virtual System context, run:

   ```
   > asg_span_port set
   ```
8. Reboot all SGMs.

Disabling Port Mirroring on a VSX Gateway

To disable port mirroring on a VSX Gateway:

1. Go to the Bridge Mode Virtual System context.
2. Run:

   ```
   > asg_span_port unset <br_if_name>
   ```

Recommended

In SmartDashboard:

1. Go to **Policy** > **Global Properties** > **Stateful Inspection**.
2. Select both **Drop out of state packets** options.

   We recommend that you clear these options:
   - **Drop out of state TCP packets**
   - **Drop out of state ICMP packets**
3. Install policy on the Virtual Systems.
4. Reboot all SGMs.

Additional Port Mirroring Configuration Steps

We recommend doing these additional steps for the specified scenarios:

1. In Application Control and URL Filtering policies, change the destination default settings from **Internet** to **Any.**
2. For IPS, turn off the **Sequence Verifier**.
3. Set the **Distribution Mode** to **General**
4. From gclish run: > `set distribution configuration manual-general`

Security

Resetting the Administrator Password

If you forget your administrator password, you can use the Emergendisk utility to restore the initial system administrator username and password (`admin/admin`). Run Emergendisk on the Single Management Object (SMO).

To reset the administrator password:

1. Make sure that the SMO is in the Admin UP state, and then set all other SGMS to Admin DOWN.

   Pull these Admin DOWN SGMs out from the Chassis.
2. Insert the Emergendisk device into a USB port on the SMO.
3. From the SMO CLI, go to the Expert mode.
4. At the prompt, run: `reboot`

   When the "`Automatic boot in 6 seconds`" message shows, press any key to continue.

   If this message does not show, change the boot sequence in the BIOS so that **USB device** is the first device, and reboot again.
5. From the Emergendisk menu, select: `Reset Admin Password`

   When his message shows:

   ```
   Admin password successfully reset
   Please remove disk or any other media and press enter to restart
   ```

   Remove the USB device.
6. Press **Enter** to reboot

   The administrator username/password is now set to `admin/admin`.
7. Change the administrator password.
8. Replace the SGMs into the Chassis and put them in the Admin UP state.

The system automatically copies the new password from the SMO to the other SGMs.

For more information about the Emergendisk utility, see the *Emergendisk* section in the *Gaia Administration Guide*.

Generic Routing Encapsulation – GRE (asg_gre)

Generic Routing Encapsulation (GRE) is a tunneling protocol that can encapsulate several network layer protocols inside virtual point-to-point links over an IP network.

Syntax

```
# asg_gre load | stat | verify
```

To configure GRE:

Edit this configuration file:

```
$FWDIR/conf/gre_loader.conf
```

Tunnel configuration:

```
tunnel=<tunnel_ifname> local_tun_addr=<local_tunnel_ip>
remote_tun_addr=<remote_tunnel_ip> phy_ifname=<physical_ifname>
local_addr=<local_physical_addr> remote_addr=<remote_physical_addr> ttl=<ttl>
```

Route configuration:

```
tunnel_route=<tunnel_ifname> remote_tun_addr=<remote_tunnel_ip>
network=<network>
```

Parameter	Description
<*tunnel_ifname*>	Tunnel interface name
<*local_tunnel_ip*>	Local tunnel IP address
<*physical_ifname*>	Physical interface name
<*local_physical_addr*>	Local physical address
<*remote_physical_addr*>	Remove physical address
<*ttl*>	Time To Live
<*remote_tunnel_ip*>	Remote tunnel IP
<*network*>	IP and subnet mask that define the network for the route

Configuration Example:

Configure tunnel interface with these parameters:

- Tunnel interface name: "GREtun"
- Local tunnel address: 10.0.0.3
- Remote tunnel address: 10.0.0.4
- Physical interface: eth2-01
- Local address: 40.40.40.1
- Remote address: 40.40.40.2
- ttl: 64

Add these lines to:

```
tunnel=GREtun local_tun_addr=10.0.0.3 remote_tun_addr=10.0.0.4 phy_ifname=eth2-01
local_addr=40.40.40.1 remote_addr=40.40.40.2 ttl=64
tunnel_route=GREtun remote_tun_addr=10.0.0.4 network=50.50.50.0/24
```

Note - All parameters are required.

To load the new configuration:

Run:

```
# asg_gre
```

Output:

```
# asg_gre load
Copying configuration file to all blades... done
```

```
1_01:
Clearing existing GRE tunnels...
Loading GRE module... Done
Loading tunnel interface: GREtun
Loading route: 50.50.50.11/32 via 10.0.0.4 (GREtun)
Loading tunnel interface: GREtuA
Loading tunnel interface: GREtuB
Loading tunnel interface: GREtuC
Configuration loaded
1_02:
Clearing existing GRE tunnels...
Loading GRE module... Done
Loading tunnel interface: GREtun
Loading route: 50.50.50.11/32 via 10.0.0.4 (GREtun)
Loading tunnel interface: GREtuA
Loading tunnel interface: GREtuB
Loading tunnel interface: GREtuC
Configuration loaded
1_03:
Clearing existing GRE tunnels...
Loading GRE module... Done
Loading tunnel interface: GREtun
Loading route: 50.50.50.11/32 via 10.0.0.4 (GREtun)
Loading tunnel interface: GREtuA
Loading tunnel interface: GREtuB
Loading tunnel interface: GREtuC
Configuration loaded
1_04:
Clearing existing GRE tunnels...
Loading GRE module... Done
Loading tunnel interface: GREtun
Loading route: 50.50.50.11/32 via 10.0.0.4 (GREtun)
Loading tunnel interface: GREtuA
Loading tunnel interface: GREtuB
Loading tunnel interface: GREtuC
Configuration loaded
```

Role Based Administration (RBA)

The access to gclish features is controlled by Role Based Administration (RBA). Each user is assigned a role. Each role has a set of read-only features and read-write features. The user is not exposed to any features, other than the ones assigned to his role.

RBA configuration and properties for the 61000/41000 Security System are the same as for Gaia. See the *Gaia Administration Guide* http://supportcontent.checkpoint.com/documentation_download?ID=22928 for more details.

Notes:

- Extended commands have no read/write notion. But, when you add an extended command to a role, add it as a write. The users assigned to this role can execute it, regardless of its implications.
- Each extended command should be separately added to role. Because `asg` is the "entrance" to the 61000/41000 Security System, it usually needs to be added to all roles.
- A user's uid must be zero to run to run extended commands. This property is enforced when adding new users.
- Do not edit the `/etc/passwd` file. Only do RBA configuration with gclish.

Example:

```
> add rba role myRole domain-type System readonly-features chassis,interface
readwrite-features route
> add user myUser uid 0 homedir /home/myUser
> set user myUser password
> add rba user myUser roles myRole
> show rba role myRole
```

RADIUS Authentication

RADIUS (Remote Authentication Dial-In User Service) is a client/server authentication system that supports remote-access applications. User profiles are kept in a central database on a RADIUS authentication server. Client computers or applications connect to the RADIUS server to authenticate users.

You can configure the 61000/41000 Security System to work as a RADIUS client. The 61000/41000 Security System does not include RADIUS server functionality. You can configure the 61000/41000 Security System to authenticate users even when they are not defined locally. See Configuring Non-local RADIUS Users.

You can configure your 61000/41000 Security System computer to connect to multiple RADIUS servers. If the first server in the list is unavailable, the next RADIUS server in the priority list connects.

You can delete a server at any time.

To set the 61000/41000 Security System as a Radius client

Use the `aaa radius-servers` commands to add, configure, and delete Radius authentication servers.

To configure RADIUS for use in a single authentication profile:

```
> add aaa radius-servers priority <priority> host <host> [port <port>] prompt-secret
timeout <timeout>
> add aaa radius-servers priority <priority> host <host> [port <port>] secret
<secret> timeout <timeout>
```

Example: Adding a new radius server 1.1.1.1 which listens on port 1812

```
> add aaa radius-servers priority 1 host 1.1.1.1 port 1812 prompt-secret
timeout 3
```

To delete a RADIUS configuration:

```
> delete aaa radius-servers priority <priority>
```

To change the configuration of a RADIUS entry:

```
> set aaa radius-servers priority <priority> host <host>
> set aaa radius-servers priority <priority> new-priority <priority>
> set aaa radius-servers priority <priority> port <port>
> set aaa radius-servers priority <priority> prompt-secret
> set aaa radius-servers priority <priority> secret <secret>
> set aaa radius-servers priority <priority> timeout <timeout>
```

Note - The configuration is done based on priority and not the server ID or name.

To see a list of all servers associated with an authentication profile:

```
> show aaa radius-servers list
```

To see the RADIUS server configuration:

```
> show aaa radius-servers priority <priority > host
> show aaa radius-servers priority <priority> port
> show aaa radius-servers priority <priority> timeout
```

Parameter	Description
`priority` *<priority>*	RADIUS server priority as an integer between 0 and 999 (default=0). When there two or more RADIUS servers, Gaia connects to the server with the highest priority. Low numbers have the higher priority.
`new-priority` *<priority>*	New RADIUS server priority as an integer between 0 and 999 (default=0). When there two or more RADIUS servers, Gaia connects to the server with the highest priority. Low numbers have the higher priority.
`host` *<host>*	RADIUS server IP address in dot-delimited format.
`port` *<port>*	UDP port on the RADIUS server. This value must match the port as configured on the RADIUS server. Typically this 1812 (default) or 1645 (non-standard but a commonly used alternative).
`prompt secret`	Shared secret (password) text string. The system prompts you to enter the value.
`timeout` *<timeout>*	The number of seconds to wait for the server to respond. The default value 3 seconds.
`secret` *<secret>*	The shared secret used to authenticate the RADIUS server and the local client. You must define this value on your RADIUS server.

Note - After RADIUS client configuration, every authentication request is forwarded to the RADIUS server. Therefore, every account that is configured locally must be configured on the RADIUS server as well.

Configuring Non-local RADIUS Users

To allow login with non-local user to the 61000/41000 Security System, you must define a default role for all non-local users that are configured in the Radius server.

The default role can include a combination of:

- Administrative (read/write) access to some features
- Monitoring (read-only) access to other features
- No access to other features.

Syntax

```
> add rba role radius-group-any domain-type System
readonly-features|readwrite-features <list>
```

Parameter	Description
`readonly-features` *<list>*	Comma separated list of Gaia features that have read only permissions in the specified role

Parameter	Description
`readwrite-features` *<list>*	Comma separated list of Gaia features that have read/write permissions in the specified role

Example

```
> add rba role radius-group-any domain-type System readonly-features arp
```

Verification

Connect to the 61000/41000 Security System with a non-local user:

```
MyLaptop > ssh my_radius_user@my_61k_server
```

After successful authentication, the user `my_radius_user` is assigned the role `radius-group-any` granted all the privileges defined in the `radius-group-any` role.

Configuring Local Radius users (with specific role)

You can configure users to have different roles by creating new users on the 61000/41000 Security System and assigning them the required role.

We recommended that you keep the local user's password blank.

Adding a New Radius User (add user)

You can add new Radius users.

Syntax

```
> add user <username> uid 0 homedir <path>
```

Parameter	Description
<username>	Login name of the user
<path>	Full path for the user home directory

Example

```
> add user local uid 0 homedir /home/local
```

Assigning a User Roll (add rba user)

You can choose a role from preexisting roles, or create a new role and give it custom permissions.

Syntax

```
> add rba user <username> roles <rolename>
```

Parameter	Description
<username>	User name
<rolename>	Role to assign to the user

Adding a role

You can add new roles and give them custom permissions.

Syntax

```
> add rba role <rolename> domain-type System readonly-features <readonly_list>
readwrite-features <readwrite_list>
```

Parameter	Description
<rolename>	Role name
<readonly_list>	Comma separated list of features to grant read only permissions for
<readwrite_list>	Comma separated list of features to grant read/write permissions for

Example

```
> add rba role radius domain-type System readonly-features Chassis,configuration
readwrite-features aaa-servers
```

Configuring TACACS + Servers - CLI (aaa)

Description Use the `aaa tacacs-servers` commands to configure one or more TACACS+ authentication servers.

Syntax To add a TACACS+ server:

```
add aaa tacacs-servers priority VALUE server VALUE key VALUE timeout
VALUE
```

To change the configuration of a TACACS+ server entry:

```
set aaa tacacs-servers priority VALUE
   key VALUE
   new-priority VALUE
   server VALUE
   timeout VALUE
set aaa tacacs-servers state VALUE
```

To delete TACACS+ server from the list of servers:

```
delete aaa tacacs-servers priority VALUE
```

To see the configuration of the TACACS+ servers

```
show aaa tacacs-servers
   list
   priority VALUE server
   priority VALUE timeout
   state
```

Parameters

Parameter	Description
`priority VALUE`	The priority of the TACACS+ server. Must be unique for this operating system. The priority is used to: • Determine the order in which Gaia makes contact with the servers. The server with the lowest priority number is first. For example, if three TACACS+ servers have a priority of 1, 5, and 10 respectively. Gaia makes contact with the servers in that order, and uses the first server that responds. Identify the server in commands. A command with `priority 1` applies to the server with priority 1. • **Range:** Integers 1 - 20 • **Default:** No default.
`server VALUE`	The TACACS+ server IPv4 address. • **Default:** No default.
`key VALUE`	The shared secret used for authentication between the authentication server and the Gaia client. Enter the shared secret text string without a backslash. Make sure that the shared string defined on the Gaia client matches that which is defined on the authentication server. • **Range:** Text strings, up to 256 characters, without any whitespace characters. • **Default:** No default.
`timeout VALUE`	The maximum number of seconds to wait for the server to respond. • **Range:** 1-45. • **Default:** 5
`new-priority VALUE`	The new priority.
`state VALUE`	**Range:** **On** - Enable TACACS+ authentication for all servers. **Off** - Disable TACACS+ authentication for all servers.
`list`	The list of TACACS+ servers that this system is configured to use.

Example

```
set aaa tacacs-servers priority 2 server 10.10.10.99 key MySharedSecretKey timeout 10
```

Logging and Monitoring

In This Section:

Monitoring the Network

Monitoring Service Traffic (asg profile)

Use `asg profile` to monitor traffic for each service that passes through the 61000/41000 Security System. This information is equivalent to SmartView Monitor traffic monitoring. This command has a minimal performance hit.

Syntax

```
> asg profile [ --delay <timeout>] [ -b <sgm_ids> ] [-v | -p | -g] [--rel] [--tcp
| --udp] [--ipv6 | --ipv4]
> asg profile -m
> asg profile --enable
> asg profile --disable
> asg profile --help
```

Parameter	Description
`--delay` *<timeout>*	Information refresh interval (seconds).
`-b` *<sgm_ids>*	Works with SGMs and/or Chassis as specified by *<sgm_ids>*. *<sgm_ids>* can be: • No *<sgm_ids>* specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`)

Parameter	Description
`-v | -p | -g`	The default view (with none of these options) shows values for each service, for throughput, packet rate, connection rate and the number of concurrent connections. Alternatively, you can choose one of these options: `-v` - Show verbose service statistics. `-p` - Show service statistics for these paths: • Acceleration (Accelerated by a SecureXL device) • Medium • Firewall `-g` - Show graph view of BPS per service
`--rel`	Show the results as a percentage. For the `-v`, `–p`, and default views.
`--tcp | --udp`	Choose one of these options: `--tcp` - Show TCP statistics only `--udp` - Show UDP statistics only
`--ipv6 | --ipv4`	Choose one of these options: `--ipv4` - Show ipv4 statistics only. `--ipv6` - Show ipv6 statistics only.

Parameter	Description
`-m`	Run in a convenient interactive menu mode.
`--enable`	Enable statistics collection.
`--disable`	Disable statistics collection.
`-help`	Show command syntax and help information.

Example

```
> asg profile -m
Aggregated statistics of SGMs: 1_1 Virtual Systems: 0
+---------------------------------------------------------------------------------------+
|Service distribution summary                                                           |
+-------------------------+----------+-----------+---------------+----------------------+
|Service                  |Throughput|Packet rate|Connection rate|Concurrent connections|
+-------------------------+----------+-----------+---------------+----------------------+
|8116/udp cp-cluster      |116.2 K   |112        |0              |0                     |
+-------------------------+----------+-----------+---------------+----------------------+
|22/tcp ssh               |4.5 K     |5          |0              |0                     |
+-------------------------+----------+-----------+---------------+----------------------+
|33628/tcp                |2.0 K     |1          |0              |0                     |
+-------------------------+----------+-----------+---------------+----------------------+
|33635/tcp                |1.2 K     |0          |0              |0                     |
+-------------------------+----------+-----------+---------------+-----------------------+
|33624/tcp                |1.2 K     |0          |0              |0                      |
+-------------------------+----------+-----------+---------------+-----------------------+
|33630/tcp                |400       |0          |0              |0                      |
+-------------------------+----------+-----------+---------------+-----------------------+
|33626/tcp                |400       |0          |0              |0                      |
+-------------------------+----------+-----------+---------------+-----------------------+
|33632/tcp                |336       |0          |0              |0                      |
+-------------------------+----------+-----------+---------------+-----------------------+
|67/udp bootps            |288       |0          |0              |0                      |
+-------------------------+----------+-----------+---------------+-----------------------+
|257/tcp set              |48        |0          |0              |2                      |
+-------------------------+----------+-----------+---------------+-----------------------+

+---------------------------------------------------------------------------------------+
|Totals                                                                                 |
+-------------------------+----------+-----------+---------------+-----------------------+
|Total tcp                |10.2 K    |9          |0              |8                      |
|Total udp                |116.5 K   |112        |0              |0                      |
|Total other              |0         |0          |0              |2                      |
+-------------------------+----------+-----------+---------------+-----------------------+
|System                   |126.7 K   |121        |0              |10                     |
+-------------------------+----------+-----------+---------------+-----------------------+

Time: Sun Jul 07 14:34:30 IDT 2013
SGMs: 1_1 1_2
VSs: 0 1
Choose one of the following option:(Bold options are current view)
n) Normal View
    a) Absolute Values
    r) Relative Values
v) Verbose View
V) Move to a different Virtual System
p) Path View
g) Graph View
O) Online
H) History
S) Move to next sgm
b) Back one menu
e) Exit
```

Note - This example shows the normal (not verbose) view with absolute values. The highest throughput and packet rate is from the service `8116/udp cp-cluster`. To show this view,

type: `a`

Monitoring the 61000/41000 Security System (asg_archive)

The `asg_archive` utility collects 61000/41000 Security System status and activity information in real-time, which is periodically saved to a history file. The system refreshes the data and saves history files automatically based on predefined time intervals for each status information type. You can change the refresh time intervals based on your requirements.

The `asg_archive` utility shows newest and historical statistics for each SGM or VSX Virtual System. You can easily change the SGM and/or Virtual System that shows. You can enable or disable data collection globally for all status types or for specified status types. You can also assign the data collection process to a specified CPU to help prevent negative performance impact.

Syntax

```
> asg_archive
> asg_archive --height <max_lines>
> asg_archive {--enable|--disable}
> asg_archive --status
> asg_archive --config [<collector> {enable|disable} [<seconds>]]
> asg_archive --refresh <timeout>
> asg_archive --cpu [<cpu_id>]
> asg_archive --remote <path>
```

Parameter	Description
No Parameter	Shows the System Status and the Options menu.
`--height`	Set the maximum number of lines in the output.
`--enable`	Start all data collectors, except those that were manually disabled with: `asg_archive –config`
`--disable`	Disable all information collectors.
`--status`	Show if `asg_archive` is enabled or disabled.
`--config`	Show or set the configuration of information collectors. <*collector*> - Name of the information collector, as shown in the `asg_archive --config` output. Enclose the name in double quotes. <*seconds*> - Enter a refresh period, in seconds, for the specified collector. If you do not enter a refresh, the default value is applied automatically.
`--refresh` <*timeout*>	Show or set the default refresh time, in seconds, which applies when no value is specified with the `--config` parameter.
`--cpu` <*cpu_id*>	Show or select the default CPU assigned to the data collection process. This can help prevent unnecessary performance impact caused by this command.
`--remote` <*path*>	Read archive files from a specified remote Security Gateway. Specify the path to this Security Gateway.
`--help`	Show the command syntax and help text. This option automatically closes the interactive mode and goes back to the command line.

Working with Interactive Mode

When you run `asg_archive`, the system enters Interactive Mode and shows a menu. You select an option and the applicable status information shows on the upper portion of the screen. Some menu items have sub-menus with more choices. Use the arrow keys to scroll through the status information. The menu is always available on the lower portion of the screen. This example shows the memory status (option 3-m).

```
+----------------------------------------------------------------------------+
|Resource Table                                                              |
+------------+----------------+------------+------------+------------------+
|SGM ID      |Resource Name   |Usage       |Threshold   |Total             |
+------------+----------------+------------+------------+------------------+
|1_01        |Memory          |20%         |50%         |31.3G             |
|            |HD: /           |22%         |80%         |19.4G             |
|            |HD: /var/log    |1%          |80%         |58.1G             |
|            |HD: /boot       |19%         |80%         |288.6M            |
+------------+----------------+------------+------------+------------------+

Time: Tue Jan 14 12:13:30 IST 2014
SGMs: 1_1 1_2 1_3 1_4 1_5 2_1 2_2 2_3 2_4 2_5
VSs: 0 1 2
```

```
Choose one of the following option:(Bold options are current view)
1) System Status
2) Performance
3) Hardware & Resources
    m) Memory
    f) FW Memory Allocation
    c) CPU Usage
    t) Top Process
    h) Hardware
4) SXL Statistics
5) Diagnostic
6) Logs
7) SYN Attack
8) Network
O) Online
H) History
S) Move to next SGM
V) Move to next VS
b) Back one menu
e) Exit
```

To select a menu item, enter the number or letter to the left of the item. The letters are **case sensitive**. If there is a sub-menu, the first option automatically shows in the upper section of the screen. To select a different option, enter the applicable letter. Some options open another sub-menu.

The numbered options show status and system information. The letter options, at the bottom of the menu, are operations that control the information display.

Menu Option	Description
O	**Online** - Shows the current status for the selected item
H	**History** - Shows status historical status information saved in the history files. Select the sub-menu item to show the specified history file.
S	**Move to next SGM** - Use this option to show the SGMs in sequential order.
V	**Move to next Virtual System** - Use this option to show the different Virtual Systems in sequential order.
b	**Back one menu** - Go back to the main menu or a higher sub-menu.
e	**Exit** - Close the interactive mode and go back to the command line.

Working with Interface Status (asg if)

Use this command to show information for interfaces for the 61000/41000 Security System. The command output shows:

- IPv4, IPv6, and MAC address
- Interface type
- State
- Currently defined interface speed
- MTU
- Duplex status

You can also use this command to do these interface management tasks:

- Set the interface speed
- Enable or disable the interface

Syntax

```
> asg if -h
> asg if [-i <interface> [-v] [enable|disable] [set_speed {0|1000|10000}] [-ip ]
```

Parameter	Description
`-h`	Show command syntax.
`-i` <*interface*>	Interface status for the specified interface or a comma-separated list of interfaces. If this parameter is not specified, the status for all interfaces shows.
`-v`	Verbose - Shows detailed output.
`enable | disable`	Enable or disable the specified interface.
`set_speed`	Set interface port speed. Valid values: • `0` • `1000` • `10000`
`-ip`	Interface IPv4 or IPv6 address.

Global view of all interfaces (asg if)

Use `asg if` to show the current status of all defined interfaces on the system.

```
> asg if
+---------------------------------------------------------------------------------------+
|Interfaces Data                                                                        |
+---------------------------------------------------------------------------------------+
|Interface  |IPv4 Address      |Info        |State        |Speed    |MTU       |Duplex  |
|           |MAC  Address      |            |(ch1)        |         |          |        |
+-----------+------------------+------------+-------------+---------+----------+--------+
|bond1      |17.17.17.10       |Bond Master |(down)       |NA       |NA        |NA      |
|           |00:1c:7f:81:05:fe |            |slaves:      |         |          |        |
|           |                  |            |eth1-05(down)|         |          |        |
|           |                  |            |eth2-05(down)|         |          |        |
+-----------+------------------+------------+-------------+---------+----------+--------+
|  eth1-05  |-                 |Bond slave  |(down)       |10G      |1500      |Full    |
|           |00:1c:7f:81:05:fe |            |master:      |         |          |        |
|           |                  |            |bond1(down)  |         |          |        |
+-----------+------------------+------------+-------------+---------+----------+--------+
|  eth2-05  |-                 |Bond slave  |(down)       |10G      |1500      |Full    |
|           |00:1c:7f:81:05:fe |            |master:      |         |          |        |
|           |                  |            |bond1(down)  |         |          |        |
+-----------+------------------+------------+-------------+---------+----------+--------+
|bond1.201  |18.18.18.10       |Vlan        |(down)       |NA       |NA        |NA      |
|           |00:1c:7f:81:05:fe |            |             |         |          |        |
+-----------+------------------+------------+-------------+---------+----------+--------+
|br0        |-                 |Bridge Mast |(up)         |NA       |NA        |NA      |
|           |00:1c:7f:81:07:fe |            |ports:       |         |          |        |
|           |                  |            |eth2-07(down)|         |          |        |
|           |                  |            |eth1-07(down)|         |          |        |
+-----------+------------------+------------+-------------+---------+----------+--------+
|  eth1-07  |-                 |Bridge port |(down)       |10G      |1500      |Full    |
|           |00:1c:7f:81:07:fe |            |master:      |         |          |        |
|           |                  |            |br0(up)      |         |          |        |
+-----------+------------------+------------+-------------+---------+----------+--------+
|  eth2-07  |-                 |Bridge port |(down)       |10G      |1500      |Full    |
|           |00:1c:7f:82:07:fe |            |master:      |         |          |        |
|           |                  |            |br0(up)      |         |          |        |
+-----------+------------------+------------+-------------+---------+----------+--------+
|eth1-01    |15.15.15.10       |Ethernet    |(up)         |10G      |1500      |Full    |
|           |00:1c:7f:81:01:fe |            |             |         |          |        |
+-----------+------------------+------------+-------------+---------+----------+--------+
|eth1-Mgmt4 |172.23.9.67       |Ethernet    |(up)         |10G      |1500      |Full    |
|           |00:d0:c9:ca:c7:fa |            |             |         |          |        |
+-----------+------------------+------------+-------------+---------+----------+--------+
|eth2-01    |25.25.25.10       |Ethernet    |(up)         |10G      |1500      |Full    |
|           |00:1c:7f:82:01:fe |            |             |         |          |        |
+-----------+------------------+------------+-------------+---------+----------+--------+
|Sync       |192.0.2.1         |Bond Mas    |(up)         |NA       |NA        |NA      |
|           |00:1c:7f:01:04:fe |            |slaves:      |         |          |        |
|           |                  |            |eth1-Sync(up)|         |          |        |
|           |                  |            |eth2-Sync(up)|         |          |        |
+-----------+------------------+------------+-------------+---------+----------+--------+
|  eth1-Sync|-                 |Bond slave  |(up)         |10G      |1500      |Full    |
|           |00:1c:7f:01:04:fe |            |master:      |         |          |        |
|           |                  |            |Sync(up)     |         |          |        |
+-----------+------------------+------------+-------------+---------+----------+--------+
|  eth2-Sync|-                 |Bond slave  |(up)         |10G      |1500      |Full    |
|           |00:1c:7f:01:04:fe |            |master:      |         |          |        |
|           |                  |            |Sync(up)     |         |          |        |
+-----------+------------------+------------+-------------+---------+----------+--------+
```

Notes

- This sample output shows:
  - This sync interface is a bond-Master
  - Interfaces are UP or DOWN
- To add a comment to an interface, run:

  > `set interface` <*if_name*> `comment` <*comment_text*>

Verbose mode

The verbose mode shows extended information, including information retrieved from the switch. You can use the verbose mode for one interface or a comma-separated list of interfaces. This operation can take a few seconds for each interface.

```
# asg if -i eth1-01 -v
```

```
Collecting information, may take few seconds
+----------------------------------------------------------------------------------------+
|Interfaces Data                                                                         |
+----------------------------------------------------------------------------------------+
|Interface|IPv4 Address         |Info        |State         |Speed   |MTU       |Duplex  |
|         |MAC  Address         |            |(ch1)/(ch2)   |        |          |        |
|         |IPv6 Address (global)|            |              |        |          |        |
|         |IPv6 Address (local) |            |              |        |          |        |
+---------+---------------------+------------+--------------+--------+----------+--------+
|eth1-01  |-                    |Bond slave  |(up)/(up)     |10G     |1500      |Full    |
|         |00:1c:7f:a1:01:0     |            |master:       |        |          |        |
|         |-                    |            |bond1(up)/(up)|        |          |        |
|         |-                    |            |              |        |          |        |
+---------+---------------------+------------+--------------+--------+----------+--------+
|Comment                                                                                 |
+----------------------------------------------------------------------------------------+
|internal interface                                                                      |
+----------------------------------------------------------------------------------------+
|Traffic                                                                                 |
+----------------------------------------------------------------------------------------+
|media            |In traffic |In pkt(uni/mul/brd)|Out traffic    |Out pkt(uni/mul/brd)  |
+-----------------+-----------+-------------------+---------------+----------------------+
|FTLF8528P2BNV-EM |28.8Kbps   |0pps/38pps/5pps    |4.1Mbps        |0pps/355pps/0pps      |
+----------------------------------------------------------------------------------------+
|Errors (total/pps)                                                                      |
+----------------------------------------------------------------------------------------+
|OutDiscards                  |InDiscards         |InErrors       |OutErrors             |
+-----------------------------+-------------------+---------------+----------------------+
|0/0                          |0/0                |0/0            |0/0                   |
+-----------------------------+-------------------+---------------+----------------------+
```

Enabling/Disabling Interface Ports

Use the `asg if` command to enable or disable interface. You can only use `asg if` to enable or disable physical interfaces (for example: eth1-01). You cannot use this command for bonds, VLANs, or other virtual interfaces. This command works on the SSM level.

To disable an interface port:

Run:

```
# asg if -i eth1-01 disable
You are about to perform port state disable on eth1-01 on blades: all


Are you sure? (Y - yes, any other key - no) y

Port state disable on eth1-01 requires auditing
Enter your full name: y
Enter reason for port state disable on eth1-01 [Maintenance]: y
WARNING: Port state disable on eth1-01 on blades: all, User: y, Reason: y
interface eth1-01 is disabled
```

To enable an interface port:

Run:

```
# asg if -i eth1-01 enable
You are about to perform port state enable on eth1-01 on blades: all


Are you sure? (Y - yes, any other key - no) y

Port state disable on eth1-01 requires auditing
Enter your full name: y
Enter reason for port state disable on eth1-01 [Maintenance]: y
WARNING: Port state enable on eth1-01 on blades: all, User: y, Reason: y
interface eth1-01 is enabled
```

Connecting to a specific SGM (blade)

When you connect to the 61000/41000 Security System, you are actually connected to one of the SGMs. You can use `blade` to open a connection to a different Security Gateway Module. You must run `blade` in the Expert mode, which establishes a new SSH connection over the Sync interface.

Syntax

```
# blade [<chassis_id>_]<sgm_id>
```

Example

```
# blade 1_03
```

Output

```
Moving to blade 1_3
```

Notes

- When you only enter the SGM ID, the default Chassis is assumed.
- To go back to the last SGM, run: `exit`
- You can run more than one `blade` command to open many SSH sessions.

Setting the Port Speed

You can set the port speed for one interface port or a comma-separated list of ports.

```
# asg if -i eth1-01,eth2-01 set_speed 10000
You are about to perform port speed change to 10000 on eth1-01 eth2-01 on blades:
all


Are you sure? (Y - yes, any other key - no) y

Port speed change to 10000 on eth1-01 eth2-01 requires auditing
Enter your full name: y
Enter reason for port speed change to 10000 on eth1-01 eth2-01 [Maintenance]: y
WARNING: Port speed change to 10000 on eth1-01 eth2-01 on blades: all, User: y,
Reason: y
Interface eth1-01 speed was set to 10G
Interface eth2-01 speed was set to 10G
```

Showing Bond Interfaces (asg_bond)

The `asg_bond` command shows bond interfaces and runs LACP packet tests:

- MAC address consistency for each Chassis
- Slave state consistency for all SGMs
- Database consistency for all SGMs
- Make sure that the LACP aggregator ID between bond and slaves are compatible
- Verification of the LACP packet between neighbors and key comparison

You can run this command for specified bonds or for all bonds.

Syntax

```
# asg_bond [v] [ -i <filter>] [-help |-h]
```

Parameter	Description
`-h|--help`	Show command syntax.
`-i` *<filter>*	Enter a bond name or a string. The output shows all bonds that match the bond name or those names that contain the text string.
`-v`	Run LACP packet test for the specified interfaces.

Global List of all Bonds

Use this command without parameters to show all currently defined bonds.

```
# asg_bond
+------+-------------------------------------------+---------+--------+------------------+
|Name  |Address                 |Mode              |Slaves   |Result  |Comments          |
+------+------------------------+------------------+---------+--------+------------------+
|bond1 |(MAC)  00:1c:7f:81:02:fe|LACP 802.3ad      |eth1-02  |OK      |                 |
|      |(IPv4) 13.13.1.10       |Load Sharing      |eth1-03  |        |                 |
|      |                        |                  |eth2-03  |        |                 |
|      |                        |                  |eth2-02  |        |                 |
+------+------------------------+------------------+---------+--------+-----------------+
|bond3 |(MAC)  00:1c:7f:82:04:fe|XOR               |eth2-04  |OK      |                 |
|      |(IPv4) 23.23.1.10       |Load Sharing      |eth1-04  |        |                 |
+------+------------------------+------------------+---------+--------+-----------------+
|bond5 |(MAC)  00:1c:7f:81:07:fe|Round-Rubin       |eth1-07  |OK      |                 |
|      |(IPv4) 33.33.1.10       |Load Sharing      |eth2-07  |        |                 |
+------+------------------------+------------------+---------+--------+-----------------+
|bond7 |(MAC)  00:00:00:00:00:fe|Active-Backup     |         |OK      |- No slaves exist |
|      |                        |High Availability |         |        |                 |
+------+------------------------+------------------+---------+--------+-----------------+
```

Filtering a Bond Interface

This example shows the command output for the specified bond.

```
# asg_bond -i bond5
```

```
+--------+------------------------------+-------------+---------+--------+---------+
|Name    |Address                       |Mode         |Slaves   |Result  |Comments |
+--------+------------------------------+-------------+---------+--------+---------+
|bond5   |(MAC)  00:1c:7f:81:07:fe      |Round-Rubin  |eth1-07  |OK      |         |
|        |(IPv4) 33.33.1.10             |Load Sharing |eth2-07  |        |         |
+--------+------------------------------+-------------+---------+--------+---------+
```

Note - You can also specify a substring that is part of a bond name to show all bonds that contain the substring.

Verification Test

This example shows the verification test results for all bonds, including one with an error.

```
> asg_bond -v
```

```
Listening for LACP packets [...............................]   [ OK ]

+-----+-----------------------+-----------------+-------+------+------------------------+
|Name |Address                |Mode             |Slaves |Result|Comments                |
+-----+-----------------------+-----------------+-------+------+------------------------+
|bond1|(MAC) 00:1c:7f:81:02:fe |LACP 802.3ad     |eth1-02|Failed|eth1-02 missing LACP pkts|
|     |(IPv4)13.13.1.10       |Load Sharing     |eth1-03|      |eth1-03 missing LACP pkts|
|     |                       |                 |eth2-03|      |eth2-03 missing LACP pkts|
|     |                       |                 |eth2-02|      |eth2-02 missing LACP pkts|
+-----+-----------------------+-----------------+-------+------+------------------------+
|bond3|(MAC)  00:1c:7f:82:04:fe|XOR              |eth2-04|OK    |                        |
|     |(IPv4) 23.23.1.10      |Load Sharing     |eth1-04|      |                        |
+-----+-----------------------+-----------------+-------+------+------------------------+
|bond5|(MAC)  00:1c:7f:81:07:fe|Round-Rubin      |eth1-07|OK    |                        |
|     |(IPv4) 33.33.1.10      |Load Sharing     |eth2-07|      |                        |
+-----+-----------------------+-----------------+-------+------+------------------------+
|bond7|(MAC)  00:00:00:00:00:fe|Active-Backup    |       |OK    | - No slaves exist      |
|     |                       |High Availability|       |      |                        |
+-----+-----------------------+-----------------+-------+------+------------------------+
```

Notes

- The comments column shows a description of problems detected by the verification tests.
- Bond7 shows an incomplete definition with no slaves configured.

Setting the Minimum Number of Slaves in a Bond

You can monitor Bond interfaces with `asg stat`. A Bond interface is considered DOWN when the number of slaves in the bond that are `UP,` are less than the `min_slaves` value. You can change the `min_slaves` value in gclish.

Syntax

```
> set chassis high-availability bond <bond_port> min_slaves
```

Example

```
> set chassis high-availability bond bond1 min_slaves 2
```

Notes

- The default value for `min_slaves` is 1.
- The Bond is considered DOWN if the number of slaves in the UP state, is below the `min_slaves` value.

Showing Traffic Information (asg_ifconfig)

The `asg_ifconfig` command collects traffic statistics from all or a specified range of SGMs. The combined output shows the traffic distribution between SGMs and their interfaces (calculated during a certain period).

The `asg_ifconfig` command has these modes:

- **Native**

  Default setting. When the `analyze` or `banalyze` option is not specified the command behaves almost the same as the native Linux `ifconfig` command. However, the output shows statistics for all interfaces on all SGMs, and for interfaces on the local SGM.
- **Analyze**

  Shows accumulated traffic information and traffic distribution between SGMs.
- **Banalyze**

  Shows accumulated traffic information and traffic distribution between interfaces

Note:

- The `analyze` and `banalyze` parameters cannot be used together.
- If you run this command in a Virtual System context, you can only see the output that applies to that context.

Syntax

```
> asg_ifconfig [-b <sgm_ids>] [<interface>] [analyze|banalyze] [-d <delay>] [-v] [-a]
```

Parameter	Description
`Interface`	The name of the interface
`-b <sgm_ids>`	Works with SGMs and/or Chassis as specified by <*sgm_ids*>. <*sgm_ids*> can be: • No <*sgm_ids*> specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`)
`-d delay`	Delay, in seconds, between data samples (default = 5)

Parameter	Description
`-v`	Verbose mode: Shows traffic distribution between interfaces
`-a`	Shows total traffic volume By default (without `-a`), the average traffic volume per second shows.
`-h`	Shows help information and exit
`analyze`	Shows accumulated traffic information Use the `-v`, `-a,` and `-d <delay>` parameters to show traffic distribution between interfaces.

`banalyze`	Shows accumulated traffic information.
	Use the `-v`, `-a,` and `-d <delay>` parameters to show traffic distribution between interfaces.
	You can use these parameters to sort the traffic distribution table:
	`-rp` X packets
	`-rb` X bytes
	`-rd` X dropped packets
	`-tp` X packets
	`-tb` X bytes
	`-td` X dropped packet
	For example, if you sort with the `-rb` option, the higher values appear at the top of the RX bytes column in the traffic distribution table:
	`SGM ID  RX packets   RX bytes   RX dropped`
	`1_03                   70%`
	`1_02                   20%`
	`1_01                   10%`
	By default, the traffic distribution table is not sorted.

Native Usage

This example shows the total traffic sent and received by eth2-01 for all SGMs on Chassis 1 (Active Chassis). By default, the average traffic volume per second shows.

```
> asg_ifconfig -b chassis1 eth2-01

as1_02:
eth2-01     Link encap:Ethernet  HWaddr 00:1C:7F:81:01:EA
            UP BROADCAST RUNNING SLAVE MULTICAST  MTU:1500  Metric:1
            RX packets:94 errors:0 dropped:0 overruns:0 frame:0
            TX packets:63447 errors:0 dropped:0 overruns:0 carrier:0
            collisions:0 txqueuelen:0
            RX bytes:5305 (5.1 KiB)  TX bytes:5688078 (5.4 MiB)

1_03:
eth2-01     Link encap:Ethernet  HWaddr 00:1C:7F:81:01:EA
            UP BROADCAST RUNNING SLAVE MULTICAST  MTU:1500  Metric:1
            RX packets:137 errors:0 dropped:0 overruns:0 frame:0
            TX packets:26336 errors:0 dropped:0 overruns:0 carrier:0
            collisions:0 txqueuelen:0
            RX bytes:7591 (7.4 KiB)  TX bytes:2355386 (2.2 MiB)
```

```
1_04:
eth2-01     Link encap:Ethernet  HWaddr 00:1C:7F:81:01:EA
            UP BROADCAST RUNNING SLAVE MULTICAST  MTU:1500  Metric:1
            RX packets:124 errors:0 dropped:0 overruns:0 frame:0
            TX packets:3098 errors:0 dropped:0 overruns:0 carrier:0
            collisions:0 txqueuelen:0
            RX bytes:6897 (6.7 KiB)  TX bytes:378990 (370.1 KiB)

1_05:
eth2-01     Link encap:Ethernet  HWaddr 00:1C:7F:81:01:EA
            UP BROADCAST RUNNING SLAVE MULTICAST  MTU:1500  Metric:1
            RX packets:79 errors:0 dropped:0 overruns:0 frame:0
            TX packets:26370 errors:0 dropped:0 overruns:0 carrier:0
            collisions:0 txqueuelen:0
            RX bytes:4507 (4.4 KiB)  TX bytes:2216546 (2.1 MiB)
```

Using the Analyze Option

This example shows accumulated traffic volume statistics for eth2-Sync per SGM and the total for all SGMs. The traffic distribution for each SGM also shows. The `-a` option shows the total traffic volume instead of the average volume per second.

```
> asg_ifconfig eth2-Sync analyze -v -a
Command is executed on SGMs: chassis_active

1_01:
eth2-Sync   Link encap:Ethernet  HWaddr 00:1C:7F:01:04:FE
            UP BROADCAST RUNNING SLAVE MULTICAST  MTU:1500  Metric:1
            RX: packets:225018 bytes:36970520 (37.0 MiB)  dropped:0
            TX: packets:3522445 bytes:1381032583 (1.4 GiB)  dropped:0

1_02:
eth2-Sync   Link encap:Ethernet  HWaddr 00:1C:7F:02:04:FE
            UP BROADCAST RUNNING SLAVE MULTICAST  MTU:1500  Metric:1
            RX: packets:221395 bytes:35947248 (35.9 MiB)  dropped:0
            TX: packets:4674143 bytes:1850315554 (1.9 GiB)  dropped:0

1_03:
eth2-Sync   Link encap:Ethernet  HWaddr 00:1C:7F:03:04:FE
            UP BROADCAST RUNNING SLAVE MULTICAST  MTU:1500  Metric:1
            RX: packets:10 bytes:644 (644.0 b)  dropped:0
            TX: packets:67826313 bytes:7345458105 (7.3 GiB)  dropped:0

1_04:
eth2-Sync   Link encap:Ethernet  HWaddr 00:1C:7F:04:04:FE
            UP BROADCAST RUNNING SLAVE MULTICAST  MTU:1500  Metric:1
            RX: packets:13 bytes:860 (860.0 b)  dropped:0
            TX: packets:68489217 bytes:7487476060 (7.5 GiB)  dropped:0

1_05:
eth2-Sync   Link encap:Ethernet  HWaddr 00:1C:7F:05:04:FE
            UP BROADCAST RUNNING SLAVE MULTICAST  MTU:1500  Metric:1
            RX: packets:203386 bytes:19214238 (19.2 MiB)  dropped:0
            TX: packets:7164109 bytes:2740761091 (2.7 GiB)  dropped:0

=*= Accumulative =*=
eth2-Sync   Link encap:Ethernet
            UP BROADCAST RUNNING SLAVE MULTICAST  MTU:1500  Metric:1
            RX: packets:649822 bytes:92133510 (92.1 MiB)  dropped:0
            TX: packets:151676227 bytes:20805043393 (20.8 GiB)  dropped:0

=*= Traffic Distribution =*=
-----------------------------------------------------------------------------
     SGM ID RX packets   RX bytes RX dropped TX packets   TX bytes TX dropped
-----------------------------------------------------------------------------
       1_01      34.6%      40.1%       0.0%       2.3%       6.6%       0.0%
       1_02      34.1%      39.0%       0.0%       3.1%       8.9%       0.0%
       1_03       0.0%       0.0%       0.0%      44.7%      35.3%       0.0%
       1_04       0.0%       0.0%       0.0%      45.2%      36.0%       0.0%
       1_05      31.3%      20.9%       0.0%       4.7%      13.2%       0.0%
-----------------------------------------------------------------------------
```

Showing Multicast Information

Showing Multicast Routing - asg_mroute

The `asg_mroute` command shows this multicast routing information in a tabular format:

- **Source** - Source IP address
- **Dest** - Destination address
- **Iif** - Source interface
- **Oif** - Outbound interface

You can filter the output for specified interfaces and SGMs.

Syntax

```
> asg_mroute -h
> asg_mroute [-d <dest_route>] [-s <src_route>] [-i <src_if>][-b <sgm_ids>]
```

Parameter	Description
`-h`	Show command syntax.
`-d`	Destination multicast group IP address.
`-s`	Source IP address.
`-i`	Source interface name.
`-b` *<sgm_ids>*	Works with SGMs and/or Chassis as specified by *<sgm_ids>*. *<sgm_ids>* can be: • No *<sgm_ids>* specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`)

Example: Show all multicast routes

This example shows all multicast routes for all interfaces and SGMs.

```
> asg_mroute
+--------------------------------------------------------------------------------+
|Multicast Routing (All SGMs)                                                    |
+--------------------------------------------------------------------------------+
|Source                   |Dest                     |Iif            |Oif           |
+-------------------------+-------------------------+---------------+--------------+
|12.12.12.1               |225.0.90.90              |eth1-01        |eth1-02       |
+-------------------------+-------------------------+---------------+---------------+
|22.22.22.1               |225.0.90.90              |eth1-02        |eth1-01        |
+-------------------------+-------------------------+---------------+---------------+
|22.22.22.1               |225.0.90.91              |eth1-02        |eth1-01        |
+-------------------------+-------------------------+---------------+---------------+
```

When no optional parameters are specified, all routes, interfaces and SGMs are shown.

Example: Show only specified interfaces or SGMs

This example shows routes for the specified source IP address, Interface and destination IP address.

```
> asg_mroute  -s 22.22.22.1 -i eth1-02  -d 225.0.90.91
+-----------------------------------------------------------------------------------+
|Multicast Routing (All SGMs)                                                       |
+-----------------------------------------------------------------------------------+
|Source                   |Dest                     |Iif            |Oif            |
+-------------------------+-------------------------+---------------+---------------+
|22.22.22.1               |225.0.90.91              |eth1-02        |eth2-01        |
+-------------------------+-------------------------+---------------+---------------+
```

Showing PIM Information - (asg_pim)

The `asg_pim` command shows this PIM information in a tabular format:

- **Source** - Source IP address
- **Dest** - Destination IP address
- **Mode** - Both Dense Mode and Sparse Mode are supported
- **Flags** - Local source and MFC state indicators
- **In. intf** - Source interface
- **RPF** - Reverse Path Forwarding indicator
- **Out int** - Outbound interface
- **State** - Outbound interface state

You can filter the output for specified interfaces and SGMs.

Syntax

```
> asg_pim -h
> asg_pim [-b <sgm_ids>] [-i <if>] [-n <neighbor>]
> asg_pim neighbors
```

Parameter	Description
`-h`	Show command syntax.
`-b` *<sgm_ids>*	Works with SGMs and/or Chassis as specified by <*sgm_ids*>. *<sgm_ids>* can be: • No <*sgm_ids*> specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`)
`-i` *<if>*	Show only the specified source interface.

Parameter	Description
`-n` *<neighbor>*	Show only the specified PIM neighbor. This parameter is relevant only with the `neighbors` option.
`neighbors`	Runs verification test to make sure that PIM neighbors are the same on all SGMs and shows this information: - **Verification** - Results of verification test. - **Neighbor** - PIM neighbor. - **Interface** - Interface name. - **Holdtime** - Time in seconds to hold a connection open during peer negotiation. - **Expires** - Minimum and Maximum expiration values for all SGMs.

Example: Show PIM information for all interfaces and SGMs

This example shows PIM information and multicast routes for all interfaces and SGMs.

```
> asg_pim
+--------------------------------------------------------------------------------------+
|PIM (All SGMs)                                                                        |
+--------------------------------------------------------------------------------------+
|source     |dest        |Mode      |Flags|In. intf |RPF       |Out. intf   |State     |
+-----------+------------+----------+-----+---------+----------+------------+----------+
|12.12.12.1 |225.0.90.90 |Dense-Mode|L|M  |eth1-01  |none      |            |          |
+-----------+------------+----------+-----+---------+----------+------------+----------+
|22.22.22.1 |225.0.90.90 |Dense-Mode|L|M  |eth1-02  |none      |eth1-01     |Forwarding|
+-----------+------------+----------+-----+---------+----------+------------+----------+
|22.22.22.1 |225.0.90.91 |Dense-Mode|L|M  |eth1-02  |none      |eth1-01     |Forwarding|
|           |            |          |     |         |          |eth2-01     |Forwarding|
+-----------+------------+----------+-----+---------+----------+------------+----------+
Flags: L - Local source, M - MFC State
```

- When no optional parameters are specified, all routes, interfaces and SGMs are shown.
- In this version, both the Dense Mode and the Sparse Mode are supported.

Example: Show PIM Information for the specified interface on all SGMs.

```
> asg_pim -i eth1-02 -b all
+--------------------------------------------------------------------------------------+
|PIM (All SGMs)                                                                        |
+--------------------------------------------------------------------------------------+
|SGM 1_01                                                                              |
+--------------------------------------------------------------------------------------+
|source     |dest        |Mode      |Flags|In. intf |RPF       |Out. intf   |State     |
+-----------+------------+----------+-----+---------+----------+------------+----------+
|22.22.22.1 |225.0.90.90 |Dense-Mode|L|M  |eth1-02  |none      |eth1-01     |Forwarding|
+-----------+------------+----------+-----+---------+----------+------------+----------+
|22.22.22.1 |225.0.90.91 |Dense-Mode|L    |eth1-02  |none      |eth1-01     |Forwarding|
|           |            |          |     |         |          |eth2-01     |Forwarding|
+-----------+------------+----------+-----+---------+----------+------------+----------+
|SGM 1_02                                                                              |
+--------------------------------------------------------------------------------------+
|source     |dest        |Mode      |Flags|In. intf |RPF       |Out. intf   |State     |
+-----------+------------+----------+-----+---------+----------+------------+----------+
|22.22.22.1 |225.0.90.90 |Dense-Mode|L|M  |eth1-02  |none      |eth1-01     |Forwarding|
+-----------+------------+----------+-----+---------+----------+------------+----------+
|22.22.22.1 |225.0.90.91 |Dense-Mode|L|M  |eth1-02  |none      |eth1-01     |Forwarding|
|           |            |          |     |         |          |eth2-01     |Forwarding|
+-----------+------------+----------+-----+---------+----------+------------+----------+
```

Example: Neighbors option

```
> asg_pim neighbors
+--------------------------------------------------------------------------------------+
|PIM Neighbors (All SGMs)                                                              |
+--------------------------------------------------------------------------------------+
|Verification:                                                                         |
|Neighbors Verification: Passed - Neighbors are identical on all blades                |
+--------------------+--------------------+--------------------+-----------------------+
|Neighbor            |Interface           |Holdtime            |Expires(min-max)       |
+--------------------+--------------------+--------------------+-----------------------+
|11.1.1.1            |bond1               |105                 |11:36:45-11:37:59      |
+--------------------+--------------------+--------------------+-----------------------+
```

Showing IGMP Information (asg_igmp)

Use this command to show IGMP information in a tabular format. You can filter the output for specified interfaces and SGMs. If no blade is specified, the command runs a verification to make sure that IGMP data is the same on all SGMs:

- **Group verification** - Makes sure that the groups exist on all SGMs. If a group is missing on some SGMs, a message shows which group is missing on which blade.
- **Global properties** - Makes sure that the flags, address and other information are the same on all SGMs.
- **Interfaces** - Makes sure that all blades have the same interfaces and that they are in the same state (UP or DOWN). If inconsistencies are detected, a warning message shows.

Syntax

```
> asg_igmp -h
> asg_igmp [-i <if>] [-b <sgm_ids>]
```

Parameter	Description
`-h`	Show command syntax.
`-i` *<if>*	Source interface name.
`-b` *<sgm_ids>*	Works with SGMs and/or Chassis as specified by *<sgm_ids>*. *<sgm_ids>* can be: • No *<sgm_ids>* specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`)

Example: Show IGMP information for all interfaces and SGMs

This example shows IGMP information and multicast routes for all interfaces and SGMs. In this example, the verification detected an interface inconsistency.

```
> asg_igmp
```

```
Collecting IGMP information, may take few seconds...
+------------------------------------------------------------------------------------------+
|IGMP (All SGMs)                                                                           |
+------------------------------------------------------------------------------------------+
|Interface: eth1-01
+------------------------------------------------------------------------------------------+
|Verification:                                                                             |
|Group Verification: Passed - Information is identical on all blades                       |
|Global Properties Verification: Passed - Information is identical on all blades           |
+------------------------------------------------------------------------------------------+
|Group               |Age       |Expire                                                    |
+--------------------+----------+----------------------------------------------------------+
|225.0.90.91         |2m        |4m                                                        |
+--------------------+----------+----------------------------------------------------------+
|Flags     |IGMP Ver |Query Interval |Query Response Interval  |protocol |Advertise Address|
+----------+---------+---------------+-------------------------+---------+-----------------+
|Querier   |2        |125            |10                       |PIM      |12.12.12.10      |
+------------------------------------------------------------------------------------------+

+------------------------------------------------------------------------------------------+
|Interface: eth1-02                                                                        |
+------------------------------------------------------------------------------------------+
|Verification:                                                                             |
|Group Verification: Failed - Found inconsistency between blades                           |
| -Group 225.0.90.92: missing in blades 1_02                                               |
|Global Properties Verification: Passed - Information is identical on all blades           |
+------------------------------------------------------------------------------------------+
|Group               |Age       |Expire                                                    |
+--------------------+----------+----------------------------------------------------------+
|225.0.90.92         |2m        |3m                                                        |
+--------------------+----------+----------------------------------------------------------+
|Flags     |IGMP Ver |Query Interval |Query Response Interval  |protocol |Advertise Address|
+----------+---------+---------------+-------------------------+---------+-----------------+
|Querier   |2        |125            |10                       |PIM      |22.22.22.10      |
+------------------------------------------------------------------------------------------+

+------------------------------------------------------------------------------------------+
|Interface: eth2-01                                                                        |
+------------------------------------------------------------------------------------------+
|Verification:                                                                             |
|Group Verification: Passed - Information is identical on all blades                       |
|Global Properties Verification: Passed - Information is identical on all blades           |
+------------------------------------------------------------------------------------------+
|Group               |Age       |Expire                                                    |
+--------------------+----------+----------------------------------------------------------+
|225.0.90.90         |2m        |3m                                                        |
+--------------------+----------+----------------------------------------------------------+
|Flags     |IGMP Ver |Query Interval |Query Response Interval  |protocol |Advertise Address|
+----------+---------+---------------+-------------------------+---------+-----------------+
|Querier   |2        |125            |10                       |PIM      |2.2.2.10         |
+------------------------------------------------------------------------------------------

NOTE: Inconsistency found in interfaces configuration between blades
Inconsistent interfaces: eth1-02
```

Example: Show IGMP Information for a specified interface.

```
> asg_igmp -i bond1.3
Collecting IGMP information, may take few seconds...
+------------------------------------------------------------------------------------------+
|IGMP (All SGMs)                                                                           |
+------------------------------------------------------------------------------------------+
|Interface: bond1.3                                                                        |
+------------------------------------------------------------------------------------------+
|Verification                                                                              |
|Group Verification: Passed - Information is identical on all blades                       |
|Global Properties Verification: Passed - Information is identical on all blades           |
+------------------------------------------------------------------------------------------+
|Group               |Age       |Expire                                                    |
--------------------+----------+-----------------------------------------------------------+
|225.0.90.90         |46m       |3m                                                        |
+--------------------+----------+----------------------------------------------------------+
|Flags     |IGMP Ver |Query Interval |Query Response Interval  |protocol |Advertise Address|
+----------+---------+---------------+-------------------------+---------+-----------------+
|Querier   |2        |125            |10                       |PIM      |12.12.12.11      |
+------------------------------------------------------------------------------------------+
```

VPN Packet Tracking (bcstats)

You can run these commands to monitor the IPSEC packet flow.

To see:	Run:
Source and destination IP addresses	• `# g_tcpdump for ip proto 50` (For Site-to-Site VPN) • `# g_tcpdump for UDP port 4500` (For SecureClient and Endpoint VPN clients)

To see:	Run:
Which SGM encrypted packets are forwarded	`# bcstats vpn -v`
Which SGM holds the outbound SA	`# g_fw tab -t outbound_SPI -f` Search for MSPI in the output. MSPI is the Meta SA, and shows which SGM holds the outbound SA.

Example: g_fw tab

```
#g_fw tab -t outbound_SPI -f
using cptfmt
Formatting table’s data – this might take a while...
local host:
Date: Nov 14, 2011
12:37:15 172.16.6.171 > : (+)====================================(÷); Table_Name: outbound_sPi; : (÷);
Attributes: dynamic, id 285,
attributes: keep, sync, kbuf 6 7, expires 3600, limit 20400, hashsize 32768; product: VPN–1 & Firewall–1;
12:37:15 1172.16.6.171 >1 : (+); peer: 172.16.6.189; ,sPi: fs9baoec; CPTFMT_sep:   sPI: 1; Ic00MB1:
c5364f5e6414aad9; ,cookieR:
95a478b10f9544a6; Expires: 3540/3610; product: VPN–1 & Firewall–1;
```

The output can include Security Associations (SAs) with an MSPI of 0. These are dummy SAs and can safely be ignored.

Monitoring VPN Tunnels

Because VPN tunnels synchronize between all SGMS, use traditional tools to monitor tunnels. This gives you a better selection of monitoring tools compared to the native 61000/41000 Security System capabilities.

SmartView Monitor

You must not activate the Monitoring Blade on the 61000/41000 Security System. But, you can still use the Tunnels information in SmartView Monitor to see VPN tunnel status and details.

SNMP

- You can use the **tunnelTable** sub-tree in Check Point MIB .1.3.6.1.4.1.2620.500.9002 to see VPN status with SNMP.
- For VSX environments, search for the *SNMP Monitoring* section in the *R76 VSX Administration Guide* (http://supportcontent.checkpoint.com/documentation_download?ID=22932) for VSX related SMTP information.

CLI Tools

Use these CLI commands:

- Run `# cpstat -f all vpn` (Expert Mode) to see VPN statistics per SGM.
- Run `# vpn tu` (Expert Mode) to monitor VPN tunnels per SGM.

  Since VPN tunnels are synchronized to all SGMs, you can use run this command from the scope of one SGM.
- Run `> vpn shell tu` (gclish) to monitor VPN tunnels in the non-interactive mode.

  This command is supported for versions R76SP.20 and higher.

Note - In a VSX environment, you must run these commands from the applicable Virtual System contact.

Showing SSM Traffic Statistics (asg_traffic_stats)

Use this command to show traffic statistics, for SSM ports during a specified time period, in terms of:

- Throughput (Bits per second)
- Packet rate (packets per second)

Packet rate statistics are divided to four categories:

- Unicast
- Multicast
- Broadcast
- Total packets per second

Syntax

```
# asg_traffic_stats {<ssm_id> | <if_name>} [<delay>]
```

Parameter	Description
<ssm_id>	SSM name (1-4) Shows the traffic statistics for the specified SSM
<if_name>	The interface name: eth1-04 or eth1-Sync Shows the total traffic statistics for a specified SSM
<delay>	Length of time, in seconds, that traffic statistics are collected (Default = 5 seconds).

Example - Traffic over one interface

```
# asg_traffic_stats eth1-04
Processing traffic statistics for 5 seconds...

eth1-04 statistics
---------------------
  Incoming traffic:
  ------------------
  Throughput: 164.9 Kbps
  Packet rate: [Total: 252 pps], [Unicast: 14 pps], [Multicast: 161 pps], [Broadcast: 76 pps]

  Outgoing traffic:
  ------------------
  Throughput: 4.0 Kbps
  Packet rate: [Total: 2 pps], [Unicast: 2 pps], [Multicast: 0 pps], [Broadcast: 0 pps]
```

Example - Traffic over one SSM

```
# asg_traffic_stats 1
Processing traffic statistics for 5 seconds...

Summary on SSM1
----------------
  Incoming traffic:
  ------------------
  Throughput: 319.1 Kbps
  Packet rate: [Total: 409 pps], [Unicast: 167 pps], [Multicast: 166 pps], [Broadcast: 75 pps]

  Outgoing traffic:
  ------------------
  Throughput: 408.2 Kbps
  Packet rate: [Total: 156 pps], [Unicast: 156 pps], [Multicast: 0 pps], [Broadcast: 0 pps]
```

Showing SGM Forwarding Statistics (asg_blade_stats)

Use this command to show detailed packet forwarding statistics.

Syntax

```
> asg_blade_stats [-6] corr [[-p [-v]] [-a] | [-reset]]
> asg_blade_stats [-6] iterator
> asg_blade_stats [-6] smo
> asg_blade_stats [-6] vpn [-v]
> asg_blade_stats [-6] 6in4 [-v]
> asg_blade_stats [-6] gre [-v]
> asg_blade_stats [-6] icmp_error [-v]
> asg_blade_stats [-6] all
> asg_blade_stats -h | Help
```

Parameter	Description
`-6`	Shows only IPv6 traffic
`corr`	Shows correction layer statistics (for predefined services) for each SGM
`-p`	Shows correction layer statistics for each service (for predefined services) for each SGM Use with `corr`.
`-reset`	Resets correction layer statistics Use with `corr`.
`-a`	Shows aggregate statistics Use with `corr`.
`-v`	Shows detailed statistics (verbose)
`iterator`	Shows information about the last iterator process
`smo`	Shows statistics for SMO task, and logs for each SGM
`vpn`	Shows statistics for VPN forwarded packets
`6in4`	Shows statistics for 6in4 tunnel forwarded packets
`gre`	Shows statistics for GRE forwarded packets

Parameter	Description
`icmp_error`	Shows statistics for ICMP ERROR forwarded packets
`vs`	Show Virtual System stateless correction layer statistics (VSX mode only)
`arp_forw`	Shows statistics for ARP forwarded packets
`all`	Shows all correction layer statistics mentioned above
`help`	Shows help information

Traceroute (asg_tracert)

Use this enhanced command to show correct tracert results on the 61000/41000 Security System. The native `tracert` cannot handle `tracert` pings correctly because of the stickiness mechanism used in the 61000/41000 Security System firewall. The `asg_tracert` command supports all native `tracert` command options and parameters.

Syntax

```
> asg_tracert <ip> [<tracert_options>]
```

Parameter	Description
<*ip*>	IP address
<*tracert_options*>	Native `tracert` command options

Example

```
> asg_tracert 100.100.100.99
```

Output

```
  traceroute to 100.100.100.99 (100.100.100.99), 30 hops max, 40 byte packets
   1   (20.20.20.20)  0.722 ms  0.286 ms  0.231 ms
   2   (100.100.100.99)  1.441 ms  0.428 ms  0.395 ms
```

Monitoring Management Interfaces Link State

By default, 61000/41000 Security System monitors the link state only on data ports (ethX-YZ). The Management Monitor feature lets SNMP monitor Management ports for the SSM60 and SSM160 components. The link state is sent to all SGMs and is integrated with the Chassis High Availability mechanism. Management ports show in the `asg stat -v` output when they are enabled. (See the **Ports** > **Mgmt** line in the output example below.)

Monitored management ports are included in the Chassis grade mechanism, according to defined factors (default = 11). In addition, the `asg if` command shows the link state of Management interfaces based on the feature mechanism.

Note - For the SSM60, it is necessary to pre-configure the Base Switch to enable the SNMP server, before you enable the feature itself. See ("SSM60 snmp-server configuration" on page 112) for details. After you configure the SNMP server, run: `# set chassis high-availability mgmt-monitoring on`

```
> asg stat -v
--------------------------------------------------------------------------------
| Chassis 1                     ACTIVE                                         |
--------------------------------------------------------------------------------
| SGM ID         State          Process                   FW Policy Date       |
| 1  (local)     UP             Enforcing Security        01Sep14 20:04        |
| 2              UP             Enforcing Security        01Sep14 20:04        |
| 3              UP             Enforcing Security        01Sep14 20:04        |
| 4              UP             Enforcing Security        01Sep14 20:04        |
--------------------------------------------------------------------------------
| Chassis 2                     STANDBY                                        |
--------------------------------------------------------------------------------
| SGM ID         State          Process                   FW Policy Date       |
| 1              UP             Enforcing Security        01Sep14 20:04        |
| 2              UP             Enforcing Security        01Sep14 20:04        |
| 3              UP             Enforcing Security        01Sep14 20:04        |
| 4              UP             Enforcing Security        01Sep14 20:04        |
--------------------------------------------------------------------------------
| Chassis Parameters                                                           |
--------------------------------------------------------------------------------
| Unit                           Chassis 1     Chassis 2    Unit Weight        |
|                                                                              |
| SGMs                            4 / 4         4 / 4           6             |
| Ports                                                                        |
|    Standard                     2 / 2         2 / 2          11             |
|    Bond                         2 / 2         2 / 2          11             |
|    Mgmt                         1 / 1         1 / 1          11             |
|    Other                        0 / 0         0 / 0           6             |
| Sensors                                                                      |
|    Fans                         4 / 4         4 / 4           5             |
|    SSMs                         2 / 2         2 / 2          11             |
|    CMMs                         2 / 2         2 / 2           6             |
|    Power Supplies               5 / 5         3 / 5           6             |
|                                                                              |
| Chassis Grade                 163 / 163     163 / 163         -             |
--------------------------------------------------------------------------------
| Minimum grade gap for chassis failover:                       11             |
| Synchronization                                                              |
|     Within  chassis:           Enabled    (Default)                          |
|     Between chassis:           Enabled    (Default)                          |
|     Exception Rules:                      (Default)                          |
--------------------------------------------------------------------------------
| Chassis HA mode:               Active Up                                     |
--------------------------------------------------------------------------------
```

SSM60 snmp-server configuration

To configure SNMP on the SSM60:

On each Chassis:

1. Log in to each SSM base switch address using telnet.
2. Enter Expert mode.
3. Enter 'configure terminal' mode.
4. Run these commands:
   ```
   # snmp-server enable
   # snmp-server view myview 1.3 included
   # snmp-server group mygroup v3 auth read myview write myview notify myview
   # snmp-server system-name BI_cp
   # snmp-server user asg1 group mygroup v3 auth md5 asg1asg1
   ```
5. Exit 'configure terminal' mode.
6. Save the configuration:
   ```
   # write
   ```

Validating snmp configuration

After you configure all SSM60s, make sure the configuration is correct:

```
# mgmt_monitor snmp_verify
```

Output after successful configuration:

```
Please wait while querying the snmp—servers on all SSMs
Chassis 1:
----------
SSM1: OK
SSM2: OK

Chassis 2:
----------
SSM1: OK
SSM2: OK
```

Configuring Non-local RADIUS Users Management port factor

Management Ports are integrated as part of the Chassis HA grade mechanism. Management port factors (for all Management ports) are the same as 'Standard' or 'Other' data ports factors.

Use `set chassis high-availability factors` to change the management port factors (default = 11). For more information see `set chassis high-availability factors` ("Setting Chassis Weights (Chassis High-Availability Factors)" on page 193).

Hardware Monitoring and Control

Showing Chassis and Component States (asg stat)

Use this command to show the Chassis and hardware component state for single and dual Chassis configurations. The command shows system:

- Up-time
- CPU load: average and current
- Concurrent connections
- Health

Use Verbose mode to show SGM state, process and policy.

Syntax

```
> asg stat [-v] [-vs <vs_ids>] [-l]
```

Note - If you run this command in a VSX context, the output is for the applicable Virtual System.

Parameter	Description
`-v`	Show detailed Chassis status (verbose mode).

Parameter	Description
`-vs` *<vs_ids>*	Shows the Chassis status of Virtual Systems. *<vs_ids>* can be: - No *<vs_ids>* (default) - Shows the current Virtual System context. - One Virtual System. - A comma-separated list of Virtual Systems (1, 2, 4, 5). - A range of Virtual Systems (VS 3-5). - `all` - Shows all Virtual Systems. **Note:** This parameter is only relevant in a VSX environment. For a Chassis with more than 3 SGMs, the output uses abbreviations to make the output more compact.
`-l`	Show the meaning of the abbreviations in the output for a Chassis with more than 3 SGMs.

Chassis Status Summary

```
> asg stat
------------------------------------------------------------------------------
| VSX System Status                                                          |
------------------------------------------------------------------------------
| Up time                      | 1 day, 20:04:39 hours                       |
------------------------------------------------------------------------------
| Current CPUs load average    | N/A                                         |
| Concurrent connections       | 400                                         |
| Health                       | SGMs                         1 Inactive     |
|                              | Power Supplies               2 Down         |
|                              | Virtual Systems              6 / 6          |
------------------------------------------------------------------------------
|Chassis 1                     | STANDBY                     UP / Required   |
|                              |   SGMs                       3 / 4   (!)    |
|                              |   Ports                      2 / 2          |
|                              |   Fans                       6 / 6          |
|                              |   SSMs                       2 / 2          |
|                              |   CMMs                       2 / 2          |
|                              |   Power Supplies             3 / 5   (!)    |
------------------------------------------------------------------------------
|Chassis 2                     | ACTIVE                      UP / Required   |
|                              |   SGMs                       4 / 4          |
|                              |   Ports                      2 / 2          |
|                              |   Fans                       6 / 6          |
|                              |   SSMs                       2 / 2          |
|                              |   CMMs                       2 / 2          |
|                              |   Power Supplies             5 / 5          |
------------------------------------------------------------------------------
```

The output shows that:

- Chassis 1 is in the Standby state
- Only three out of the required four SGMs in Chassis 1 are UP
- One SGM and two power supplies in Chassis 1 do not run

Chassis Status Details

```
> asg stat -v
```

Output (Top Section)

```
------------------------------------------------------------------------------
| VSX System Status                                                          |
------------------------------------------------------------------------------
| VS ID:  0                     VS Name:  Athens                             |
------------------------------------------------------------------------------
```

```
| Chassis 1                     STANDBY                                     |
-----------------------------------------------------------------------------
| SGM ID          State          Process                   Policy Date      |
| 1  (local)      UP             Enforcing Security        09Jan14 11:30    |
| 2               UP             Enforcing Security        09Jan14 11:30    |
| 3               DOWN           Inactive                  NA               |
| 4               UP             Enforcing Security        09Jan14 11:30    |
| 5               UP             Enforcing Security        09Jan14 11:30    |
| 6               UP             Enforcing Security        09Jan14 11:30    |
-----------------------------------------------------------------------------
| Chassis 2                     ACTIVE                                      |
-----------------------------------------------------------------------------
| SGM ID          State          Process                   Policy Date      |
| 1               UP             Enforcing Security        09Jan14 11:30    |
| 2               UP             Enforcing Security        09Jan14 11:30    |
| 3               UP             Enforcing Security        09Jan14 11:30    |
| 4               UP             Enforcing Security        09Jan14 11:30    |
| 5               UP             Enforcing Security        09Jan14 11:30    |
| 6               UP             Enforcing Security        09Jan14 11:30    |
-----------------------------------------------------------------------------
```

This output shows that:

- Chassis 1 is STANDBY with 5 SGMs UP
- Chassis 2 is ACTIVE with 6 SGMs UP

Notes

- **SGM ID** is the Identifier of the SGM. **(local)** is the SGM on which you ran the command.
- **State** is the state of the SGM. This can be:
  - **UP** - The SGM is processing traffic
  - **DOWN** - The SGM is not processing traffic
  - **Detached** - No SGM has been detected in a slot.

Note - To manually change the state of an SGM, use the `asg sgm_admin` command. This command administratively changes the state to UP or DOWN. An SGM that is DOWN because of a software or hardware problem cannot be changed to UP, using this command.

- **Process** is the status of the SGM security enforcement:
  - **Enforcing Security** - UP and works properly.
  - **Inactive** - DOWN and is experiences some problem. It is not handling traffic.
  - **Initial policy** - The SGM is UP but the policy is not installed on the SGM.

Output (Bottom Section)

```
--------------------------------------------------------------------------------
| Chassis Parameters                                                           |
--------------------------------------------------------------------------------
| Unit                           Chassis 1     Chassis 2    Unit Weight        |
|                                                                              |
| SGMs                            5 / 6 (!)     6 / 6  (!)       6             |
| Ports                                                                        |
|    Standard                     0 / 0         0 / 0           11             |
|    Bond                         2 / 2         2 / 2           11             |
|    Other                        0 / 0         0 / 0            6             |
| Sensors                                                                      |
|    Fans                         9 / 9         9 / 9            5             |
|    SSMs                         2 / 2         2 / 2           11             |
|    CMMs                         2 / 2         2 / 2            6             |
|    Power Supplies               4 / 4         3 / 3            6             |
|                                                                              |
| Chassis Grade                 133 / 139     139 / 139          -             |
--------------------------------------------------------------------------------
| Minimum grade gap for chassis failover:                       11             |
| Synchronization                                                              |
|     Within  chassis:           Enabled    (Default)                          |
|     Between chassis:           Enabled    (Default)                          |
|     Exception Rules:                      (Default)                          |
--------------------------------------------------------------------------------
```

Notes

- The X/X notation shows the number of components that are UP and the components must be UP. For example, on the SGMs line, 6/6 means that 6 SGMs are UP and 6 must be UP.
- **Chassis grade** is the sum of the grades of all components. In a Dual-Chassis deployment, the Chassis with a higher grade (by at least the **Minimum grade gap**) becomes ACTIVE. The grade of each component is the Unit Weight multiplied by the number of components that are UP.

  You can configure the Unit Weight of each component to show the importance of the component in the system. To configure the Unit Weight run:

  `> set chassis high-availability factors` *<sensor_name>*

  For example, to change the weight of the SGM to 12, run:

  `> set chassis high-availability factors sgm 12`

  If you run `asg stat -v`, the output shows a higher unit weight and Chassis grade:
- **Minimum threshold for traffic processing** - The minimum grade required for the Chassis to become ACTIVE.
- **Minimum grade gap for chassis failover** - Chassis failover occurs to the Chassis with the higher grade only if its grade is greater than the other Chassis by more than the minimum gap.
- **Synchronization** - The status of synchronization:
  - **Within chassis** - between SGMs located in the same Chassis.
  - **Between chassis** - between SGMs located in different Chassis.
  - **Exception Rules** - user configured exception rules. To configure, usCR01963350e the command `g_sync_exception.`

Compact Output for Selected SGMs

```
> asg stat -v -vs 0,1,2
------------------------------------------------------------------------------
| Chassis 1                    STANDBY                                       |
------------------------------------------------------------------------------
|SGM   |1    |2    |3    |4    |  -  |  -  |  -  |  -  |  -  |  -  |  -  |  -|
------------------------------------------------------------------------------
|State | UP  | UP  |DOWN | UP  |  -  |  -  |  -  |  -  |  -  |  -  |  -  |  -|
------------------------------------------------------------------------------
| VS ID                                                                      |
------------------------------------------------------------------------------
|  0   | ES  | ES  | ES  | ES  |  -  |  -  |  -  |  -  |  -  |  -  |  -  |  -|
------------------------------------------------------------------------------
|  1   | ES  | ES  | ES  | ES  |  -  |  -  |  -  |  -  |  -  |  -  |  -  |  -|
------------------------------------------------------------------------------
|  2   | ES  | ES  | ES  | ES  |  -  |  -  |  -  |  -  |  -  |  -  |  -  |  -|
------------------------------------------------------------------------------
| Chassis 2                    ACTIVE                                        |
------------------------------------------------------------------------------
|SGM   |1 (l)|2    |3    |4    |  -  |  -  |  -  |  -  |  -  |  -  |  -  |  -|
------------------------------------------------------------------------------
|State | UP  | UP  | UP  | UP  |  -  |  -  |  -  |  -  |  -  |  -  |  -  |  -|
------------------------------------------------------------------------------
| VS ID                                                                      |
------------------------------------------------------------------------------
|  0   | ES  | ES  | ES  | ES  |  -  |  -  |  -  |  -  |  -  |  -  |  -  |  -|
------------------------------------------------------------------------------
|  1   | ES  | ES  | ES  | ES  |  -  |  -  |  -  |  -  |  -  |  -  |  -  |  -|
------------------------------------------------------------------------------
|  2   | ES  | ES  | ES  | ES  |  -  |  -  |  -  |  -  |  -  |  -  |  -  |  -|
------------------------------------------------------------------------------
```

```
| Chassis Parameters
-------------------------------------------------------------------------------
| Unit                           Chassis 1     Chassis 2    Unit Weight       |
|                                                                             |
| SGMs                            3 / 4  (!)    4 / 4            6            |
| Ports                                                                       |
|   Standard                      0 / 0         0 / 0           50            |
|   Other                         0 / 0         0 / 0            6            |
| Sensors                                                                     |
|   Fans                          6 / 6         6 / 6            5            |
|   SSMs                          2 / 2         2 / 2           11            |
|   CMMs                          2 / 2         2 / 2            6            |
|   Power Supplies                6 / 6         6 / 6            6            |
|                                                                             |
| Chassis Grade                 118 / 124     124 / 124          -            |
-------------------------------------------------------------------------------
| Minimum grade gap for chassis failover:                       11            |
| Synchronization                                                             |
|     Within  chassis:           Enabled    (Default)                         |
|     Between chassis:           Enabled    (Default)                         |
|     Exception Rules:                      (Default)                         |
| Distribution                                                                |
|     Control Blade:             Disabled   (Default)                         |
| Chassis HA mode:               Active Up                                    |
-------------------------------------------------------------------------------
```

Output State Acronyms

To see a list of the acronyms that show in the reports:

```
> asg stat -l
Legend:

      SGM States:

           ACT - ACTIVE                  DTC - DETACHED
           DWN - DOWN                    NSG - NOT IN SECURITY GROUP

      VS States:

           ES  - Enforcing Security      FSC - FullSync Client
           FSS - FullSync Server         IAC - Inactive
           IF  - Iteration Finished      IPO - Initial Policy
           IS  - Iteration Started       NPO - No Policy
           PC  - Policy Completed        PRF - Policy Ready2Finish
           PS  - Policy Started
```

Monitoring Chassis and Component Status (asg monitor)

Use this command to continuously monitor Chassis and component status. This command shows the same information as `asg stat`, but the information stays on the screen and refreshes at user-specified intervals (default = 1 second). To stop the monitor session, press **Ctrl-c**.

Note - If you run this command in a Virtual System context, you see only the output for that Virtual System. You can also specify the Virtual System as a command parameter.

Syntax

```
> asg monitor
> asg monitor [-v|-all] [-amw] [-vs <vs_ids>] <interval>
> asg monitor -l
> asg monitor -h
```

Parameter	Description
No parameters	Shows the SGM status.
`-h`	Shows the command syntax and help information.
`-amw`	Shows the Anti-Malware policy date instead of the Firewall policy date.
`-v`	Shows only Chassis component status.
`-all`	Shows both SGM and Chassis component status.
<*interval*>	Sets the data refresh interval (in seconds) for this session.
`-vs` <*vs_ids*>	Shows the component status for one or more Virtual Systems. <*vs_ids*> can be: • No <*vs_ids*> (default) - Shows the current Virtual System context. • One Virtual System. • A comma-separated list of Virtual Systems (1, 2, 4, 5). • A range of Virtual Systems (VS 3-5). • `all` - Shows all Virtual Systems. **Note:** This parameter is only relevant in a VSX environment. For a Chassis with more than 3 SGMs, the output has abbreviations to make the output more compact.
`-l`	Shows legend of column title abbreviations.
`-h`	Shows the command syntax and help information.

Examples

This example shows the SGM status with the Anti-Malware policy date.

```
> asg monitor -amw
---------------------------------------------------------------------------
| Chassis 1                     ACTIVE                                    |
---------------------------------------------------------------------------
| SGM ID         State          Process                   AMW Policy Date |
| 1              UP             Enforcing Security        10Feb14 19:56   |
| 2  (local)     UP             Enforcing Security        10Feb14 19:56   |
| 3              UP             Enforcing Security        10Feb14 19:56   |
| 4              UP             Enforcing Security        10Feb14 19:56   |
---------------------------------------------------------------------------
| Chassis 2                     STANDBY                                   |
---------------------------------------------------------------------------
| SGM ID         State          Process                   AMW Policy Date |
| 1              UP             Enforcing Security        10Feb14 19:56   |
| 2              UP             Enforcing Security        10Feb14 19:56   |
| 3              UP             Enforcing Security        10Feb14 19:56   |
| 4              UP             Enforcing Security        10Feb14 19:56   |
---------------------------------------------------------------------------
| Chassis HA mode:               Active Up                                |
---------------------------------------------------------------------------
```

This example shows the Chassis component status.

```
> asg monitor -v
-----------------------------------------------------------------------------
| Chassis Parameters                                                        |
-----------------------------------------------------------------------------
| Unit                           Chassis 1     Chassis 2    Unit Weight     |
|                                                                           |
| SGMs                            4 / 4         3 / 4  (!)       6          |
| Ports                                                                     |
|    Standard                     2 / 2         2 / 2           11          |
|    Bond                         2 / 2         2 / 2           11          |
|    Mgmt                         1 / 1         1 / 1           11          |
|    Other                        0 / 0         0 / 0            6          |
| Sensors                                                                   |
|    Fans                         4 / 6  (!)    6 / 6            5          |
|    SSMs                         2 / 2         2 / 2           11          |
|    CMMs                         2 / 2         2 / 2            6          |
|    Power Supplies               3 / 5  (!)    3 / 5  (!)       6          |
|                                                                           |
| Chassis Grade                 157 / 173     155 / 173          -          |
-----------------------------------------------------------------------------
| Minimum grade gap for chassis failover:                      200          |
| Synchronization                                                           |
|     Within  chassis:           Enabled    (Default)                       |
|     Between chassis:           Enabled    (Default)                       |
|     Exception Rules:                      (Default)                       |
-----------------------------------------------------------------------------
| Chassis HA mode:               Primary Up (Chassis 1)                     |
-----------------------------------------------------------------------------
```

This example shows the status of the SGMs and Virtual System 3.

```
> asg monitor –vs 3
--------------------------------------------------------------------------------
| Chassis 1                    ACTIVE                                          |
--------------------------------------------------------------------------------
|SGM   |1 (l)|2    |3    |4    |  -  |  -  |  -  |  -  |  -  |  -  |  -  |  -  |
--------------------------------------------------------------------------------
|State | UP  | UP  | UP  | DWN |  -  |  -  |  -  |  -  |  -  |  -  |  -  |  -  |
--------------------------------------------------------------------------------
| VS ID                                                                        |
--------------------------------------------------------------------------------
|  3   | ES  | ES  | ES  | IAC |  -  |  -  |  -  |  -  |  -  |  -  |  -  |  -  |
--------------------------------------------------------------------------------
```

Monitoring Performance (asg perf)

Use this command to continuously monitor key performance indicators and load statistics. There are different commands for IPv4 and IPv6. You can show the performance statistics for IPv4 traffic, IPv6 traffic or for all traffic.

When you run `asg perf`, the statistics display shows on the screen. The display automatically updates after a predefined interval (default = 10 seconds). To stop `asg perf` and return to the command line, press: `e`

Syntax

```
> asg perf -h
> asg perf [-b <sgm_ids>] [-vs <vs_ids>] [-k] [-v] [-vv] [-p] [-4|-6] [-c]
> asg perf [-b <sgm_ids>] [-vs <vs_ids>] [-k] [--peak_hist|--perf_hist] [-e]
[--delay <seconds>]
> asg perf [-b <sgm_ids>] [-vs <vs_ids>] [-v] [-vv [ mem [fwk|cpd|fwd|all_daemons]
| cpu [1m|1h|24h] ] ]
```

Parameter	Description
`-h`	Shows command syntax with help

Parameter	Description
`-b` *<sgm_ids>*	Works with SGMs and/or Chassis as specified by *<sgm_ids>*. *<sgm_ids>* can be: • No *<sgm_ids>* specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`)
`-vs` *<vs_ids>*	For VSX Gateways only. Shows performance for Virtual Systems as specified by *<vs_ids>*. *<vs_ids>* can be: • No *<vs_ids>* (default) - Shows the current Virtual System context. • One Virtual System. • A comma-separated list of Virtual Systems (1, 2, 4, 5). • A range of Virtual Systems (VS 3-5). • `all` - Shows all Virtual Systems. **Note:** This parameter is only relevant in a VSX environment.
`-v`	Shows statistics for each SGM.
`-vv`	For VSX Gateways only. Shows statistics for each Virtual System.
`mem`	Shows memory usage for each daemon. Use this with `-vv`. Possible values: • `fwk` (Default) • `fwd` • `cpd` • `all_daemons`
`cpu`	Shows CPU usage for a specified period of time. Use this with `-vv`. Possible values: • `1m` (default) - The last 60 seconds • `1h` - The last hour • `24h` - The last 24 hours
`-p`	Show detailed statistics and traffic distribution between these paths on the Active Chassis: • Acceleration path (Performance Pack). • Medium path (PXL). • Slow path (Firewall).

Parameter	Description
`-4|-6`	`-4` - Shows IPv4 information only. `-6` - Shows IPv6 information only. If no value is specified, the combined performance information for both IPv4 and IPv6 shows.
`-c`	Show percentages instead of absolute values.
`-k`	Show peak (maximum) system performance values.
`--peak_hist`	Creates an exportable text file that contains all data saved in the peak performance files. You must use this parameter together with -k.
`--perf_hist`	Creates exportable text files that contain all performance data saved in the history files. You must use this parameter together with -k.
`-e`	Reset peak values and delete all peaks files and system history files.
`--delay` *<seconds>*	Temporarily changes the update interval for the current `asg perf` session. Enter a delay value in seconds. Default = 10 seconds

Notes:

- The `-b` *<sgm_ids>* and `-vs` *<vs_ids>* parameters must be at the start of the command. If both parameters are used, `-b` *<sgm_ids>* must be first.
- If your 61000/41000 Security System is not configured for VSX, the VSX related commands are not available. They do not show when you run `asg perf -h.`

Summary without Parameters

```
Thu May 21 08:17:24 IDT 2015
Aggregated statistics (IPv4 Only) of SGMs: chassis_active VSs: 0
+---------------------------------------------------------------+
|Performance Summary                                            |
+----------------------------------------------+----------------+
|Name                                          |Value           |
+----------------------------------------------+----------------+
|Throughput                                    |751.6 K         |
|Packet rate                                   |733             |
|Connection rate                               |3               |
|Concurrent connections                        |142             |
|Load average                                  |2%              |
|Acceleration load (avg/min/max)               |1%/0%/4%        |
|Instances load (avg/min/max)                  |2%/0%/8%        |
|Memory usage                                  |10%             |
+----------------------------------------------+----------------+
 * Instances / Acceleration Cores: 8 / 4
 * Activated SWB: FW,IPS
```

Notes

- By default, absolute values are shown.
- Unless otherwise specified, the combined statistics for IPv4 and IPv6 are shown.
- When no SGMs are specified, performance statistics are shown for the active SGM only.

Output with Performance Summary

The `-v` parameter adds a performance summary for each SGM.

```
> asg perf -vs all -v -vv cpu 24h
Tue Oct 22 07:23:37 IST 2013
Aggregated statistics (IPv4 and IPv6) of SGMs: chassis_active Virtual Systems: 0
+-------------------------------------------------------------------------+
|Performance Summary                                                      |
+---------------------------------------------+-------------+-------------+
|Name                                         |Value        |IPv4%        |
+---------------------------------------------+-------------+-------------+
|Throughput                                   |10.2 K       |100%         |
|Packet rate                                  |11           |100%         |
|Connection rate                              |0            |N/A          |
|Concurrent connections                       |22           |100%         |
|Load average                                 |7%           |             |
|Acceleration load (avg/min/max)              |6%/6%/6%     |             |
|Instances load (avg/min/max)                 |5%/4%/9%     |             |
|Memory usage                                 |55%          |             |
+---------------------------------------------+-------------+-------------+

+-------------------------------------------------------------------------+
|Per SGM Distribution Summary                                             |
+-----+-----------+--------+--------+--------+----------+----------+------+
|SGM  |Throughput |Packet  |Conn.   |Concu.  |Accel.    |Instances |Mem.  |
|ID   |           |Rate    |Rate    |Conn    |Cores%    |Cores%    |Usage%|
+-----+-----------+--------+--------+--------+----------+----------+------+
|1_01 |10.2 K     |11      |0       |22      |6/6/6     |5/4/9     |55%   |
+-----+-----------+--------+--------+--------+----------+----------+------+
|Total|10.2 K     |11      |0       |22      |6/6/6     |5/4/9     |55%   |
+-----+-----------+--------+--------+--------+----------+----------+------+

+-----------------------------------+
|Per VS CPU Usage Summary           |
+-----+---------+---------+---------+
|VS ID|Avg. Cpu%|Min. Cpu%|Max. Cpu%|
|     |         |(SGM id) |(SGM id) |
+-----+---------+---------+---------+
| 0   |2        |1  (1_02)|2  (1_01)|
| 1   |0        |0  (1_01)|0  (1_04)|
+-----+---------+---------+---------+
* CPU stats is aggregated over the last 24hrs
```

Notes

- Make sure that resource control monitoring is enabled on all SGMs.

  To enable resource control monitoring, run: `g_fw vsx resctrl monitor enable` from the Expert mode.
- By default, absolute values are shown.
- Average, minimum and maximum values are calculated across all active SGMs.
- The SGM ID with the minimum and maximum value shows in brackets for each SGM.
- Unless otherwise specified, the combined statistics for both IPv4 and IPv6 are shown.

- When no SGMs are specified, performance statistics are shown for the active SGM only.

Virtual System Memory Summary with Performance Summary

The `-vv mem` parameter shows memory usage for each Virtual System across all active SGMs.

Example:

```
> asg perf -vs all -vv mem
Tue Jul 29 16:05:44 IDT 2014
Aggregated statistics (IPv4 Only) of SGMs: chassis_active VSs: all
+--------------------------------------------------------------+
|Performance Summary                                           |
+----------------------------------------------+---------------+
|Name                                          |Value          |
+----------------------------------------------+---------------+
|Throughput                                    |684.5 K        |
|Packet rate                                   |700            |
|Connection rate                               |3              |
|Concurrent connections                        |144            |
|Load average                                  |2%             |
|Acceleration load (avg/min/max)               |0%/0%/1%       |
|Instances load (avg/min/max)                  |2%/0%/12%      |
|Memory usage                                  |10%            |
+----------------------------------------------+---------------+
 * Instances / Acceleration Cores: 8 / 4
+-------------------------------------------------------------------------+
|Per VS Memory Summary                                                    |
+--------+-------------+-------------+-------------+-------------+---------+
| VS ID  | User Space  | Memory in   | FWK memory  | Total memory| CPU     |
|        | memory      | Kernel      |             |             | Usage % |
+--------+-------------+-------------+-------------+-------------+---------+
|   0 max|222.3M (1_01)|1.658G (1_04)|47.11M (1_04)|1.880G (1_04)|   N/A   |
|     min|215.8M (1_03)|1.213G (1_01)|45.55M (1_03)|1.249G (1_01)|   N/A   |
+--------+-------------+-------------+-------------+-------------+---------+
|   1 max|56.34M (1_02)|   0K (1_04) |31.16M (1_02)|56.34M (1_02)|   N/A   |
|     min|54.24M (1_01)|   0K (1_04) |29.52M (1_03)|54.24M (1_01)|   N/A   |
+--------+-------------+-------------+-------------+-------------+---------+
* Maximum and minimum values are calculated across all active SGMs
```

Notes:

- The SGM which uses the most user-space memory on Virtual System 1 is SGM 1_01.
- The SGM which uses the least fwk daemon memory on Virtual System 3 is SGM 1_02.
- This information is shown only if `vsxmstat` is enabled for `perfanalyze` use.
- Make sure that `vsxmstat` feature is enabled (`vsxmstat status_raw`).

Per Path Statistics

This example shows detailed performance information for each SGM and traffic distribution between different paths. It also shows VPN throughput and connections.

```
> asg perf -p -v
Tue Oct 22 07:31:31 IST 2013
Aggregated statistics (IPv4 and IPv6) of SGMs: chassis_active Virtual Systems: 0
+-------------------------------------------------------------+
|Performance Summary                                          |
+----------------------------------------+--------------------+
|Name                                    |Value               |
+----------------------------------------+--------------------+
|Throughput                              |3.3 G               |
|Packet rate                             |6.2 M               |
|Connection rate                         |0                   |
|Concurrent connections                  |3.4 K               |
|Load average                            |54%                 |
|Acceleration load (avg/min/max)         |58%/48%/68%         |
|Instances load (avg/min/max)            |3%/1%/5%            |
|Memory usage                            |18%                 |
+----------------------------------------+--------------------+
+---------------------------------------------------------------------------------------+
|Per SGM Distribution Summary                                                           |
+-------+------------+------------+-----+-----------+-------------+--------------+------+
|SGM ID |Throughput  |Packet rate |Conn.|Concurrent |Core usage   |Core Instances|Memory|
|       |            |            |Rate |Connections|avg/min/max %|avg/min/max % |Usage |
+-------+------------+------------+-----+-----------+-------------+--------------+------+
|1_01   |644.3 M     |1.2 M       |0    |520        |52/44/62     |6/3/10        |18%   |
|1_02   |526.7 M     |997.1 K     |0    |512        |61/51/68     |2/0/5         |18%   |
|1_03   |526.6 M     |997.0 K     |0    |512        |62/53/73     |2/1/3         |18%   |
|1_04   |526.7 M     |997.0 K     |0    |804        |54/48/60     |2/1/3         |18%   |
|1_05   |526.7 M     |997.1 K     |0    |512        |59/45/76     |3/1/5         |18%   |
|1_06   |526.7 M     |997.1 K     |0    |512        |61/52/70     |4/4/5         |18%   |
+-------+------------+------------+-----+-----------+-------------+--------------+------+
|Total  |3.3 G       |6.2 M       |0    |3.4 K      |58/48/68     |3/1/5         |18%   |
+-------+------------+------------+-----+-----------+-------------+--------------+------+

+-----------------------------------------------------------------------------------+
|Per Path Distribution Summary                                                      |
+-------------------------+------------+------------+------------+------------------+
|                         |Acceleration|Medium      |Firewall    |Dropped           |
+-------------------------+------------+------------+------------+------------------+
|Throughput               |3.2 G       |0           |2.1 M       |117.6 M           |
|Packet rate              |6.0 M       |0           |1.4 K       |222.8 K           |
|Connection rate          |0           |0           |0           |                  |
|Concurrent connections   |3.2 K       |0           |156         |                  |
+-------------------------+------------+------------+------------+------------------+

+----------------------------------------+--------------------+
|VPN Performance                                              |
+----------------------------------------+--------------------+
|VPN throughput                          |2.9 G               |
|VPN connections                         |3.1 K               |
+----------------------------------------+--------------------+
```

Showing Peak Values

This example shows peak values for one Virtual System.

```
> asg perf -vs 0-1 -p
Aggregated statistics (IPv4 and IPv6) of SGMs: all Virtual Systems: 0-1
+-------------------------------------------------------------------------+
|Performance Summary                                                      |
+--------------------------------------------+---------------+------------+
|Name                                        |Value          |IPv4%       |
+--------------------------------------------+---------------+------------+
|Throughput                                  |1.7 K          |100%        |
|Packet rate                                 |2              |100%        |
|Connection rate                             |0              |N/A         |
|Concurrent connections                      |20             |100%        |
|Load average                                |6%             |            |
|Acceleration load (avg/min/max)             |5%/5%/5%       |            |
|Instances load (avg/min/max)                |5%/3%/10%      |            |
|Memory usage                                |57%            |            |
+--------------------------------------------+---------------+------------+
=+------------------------------------------------------------------------+
|Per Path Distribution Summary                                            |
+------------------+------------+--------------+--------------+-----------+
|                  |Acceleration|Medium        |Firewall      |Dropped    |
+------------------+------------+--------------+--------------+-----------+
|Throughput        |0           |0             |1.7 K         |0          |
|Packet rate       |0           |0             |2             |0          |
|Connection rate   |0           |0             |0             |           |
|Concurrent conn.  |10          |0             |10            |           |
+------------------+------------+--------------+--------------+-----------+
```

Showing History and Peak Value Files

The 61000/41000 Security System periodically saves historical system performance and peak value data. New history files are created based on a predefined interval (Default = every 4 hours). New peak value files are created whenever a new peak value is detected. You can find these files at: `/var/log/asgstats`

The system saves these files until a predefined maximum number of files is reached, after which files are deleted on an oldest first basis. You can also delete all history and peak value files manually.

System performance data includes these parameters:

- Throughput
- Packet rate
- Connection rate
- Concurrent connections
- Acceleration load
- Firewall load
- Memory consumption

You can collect the data contained in the historical peak value files and save them into two comma-separated-value text files. There is one combined file for historical system performance data and another for peak values. You can export these files and analyze them in a spreadsheet or statistical analysis application. The combined files are saved at: `$FWDIR/conf/asgpeaks.conf`

To create the combined text files:

Run:

```
> asg perf -k -peak_hist
> asg perf -k -perf_hist
```

To delete the history and peak value files:

Run:

```
> asg perf -k -e
```

Configuring Alert Thresholds (chassis_alert_threshold)

Usage

Set the hardware and performance alert thresholds. You can configure alert thresholds for performance and hardware monitoring alerts. Run the alert configuration commands from gclish.

Syntax

```
set chassis alert_threshold <threshold_name> <value>
show chassis alert_threshold <threshold_name>
```

Parameter	Description
<*threshold_name*>	Threshold name as specified in the table below
<*value*>	High or low value for the applicable threshold

Example

```
> set chassis alert_threshold mem_util_threshold_perc_high 70
```

This sets the memory utilization high limit to 70% of installed memory.

Working with Alert Thresholds

These are supported alert thresholds.

Hardware Alert Thresholds

Threshold Name	Scope	Description
fans_threshold	System	Fan speed
cpus_temperature_threshold	SGM	CPU Temperature

Performance Alert Thresholds

Threshold Name	Scope	Description
`concurr_conn_threshold_high`	SGM	Concurrent connections - high limit
`concurr_conn_threshold_low_ratio`	SGM	Concurrent connections - Low limit (% of high limit)
`concurr_conn_total_threshold_high`	System	Concurrent connections - high limit
`concurr_conn_total_threshold_low_ratio`	System	Concurrent connections - Low limit (% of high limit)
`conn_rate_threshold_high`	SGM	Connection rate per second - High limit
`conn_rate_threshold_low_ratio`	SGM	Connection rate per second - Low limit (% of high limit)
`conn_rate_total_threshold_high`	System	Connection rate per second - High limit
`conn_rate_total_threshold_low_ratio`	System	Connection rate per second - Low limit (% of high limit)
`cpu_load_threshold_perc_high`	SGM	CPU load (%) - High limit
`cpu_load_threshold_perc_low_ratio`	SGM	CPU load (%) - Low limit (% of high limit)
`hd_util_threshold_perc_high`	SGM	Disk utilization (%) - High limit
`hd_util_threshold_perc_low_ratio`	SGM	Disk utilization (%) - Low limit (% of high limit)
`mem_util_threshold_perc_high`	SGM	Memory utilization (%) - High limit
`mem_util_threshold_perc_low_ratio`	SGM	Memory utilization (%) - Low limit (% of high limit)

Threshold Name	Scope	Description
packet_rate_threshold_high	SGM	Packet rate per second - High limit
packet_rate_threshold_low_ratio	SGM	Packet rate per second - Low limit (% of high limit)
packet_rate_total_threshold_high	System	Packet rate per second - High limit
packet_rate_total_threshold_low_ratio	System	Packet rate per second - Low limit (% of high limit)
throughput_threshold_high	SGM	Throughput (bps) - High limit
throughput_threshold_low_ratio	SGM	Throughput (bps) - Low limit (% of high limit)
throughput_total_threshold_high	System	Throughput (bps) - High limit
throughput_total_threshold_low_ratio	System	Throughput (bps) - Low limit (% of high limit)

Global Operating System Commands

Global operating system commands are standard Linux commands that run on all or specified SGMs. When you run a global command in the gclish shell, the operating system runs a global script, which the standard Linux command on the SGMs. When you run a command in the Expert mode, it works as a standard Linux command. To use the global command in the Expert mode, run the global command script version as shown in this table:

gclish Command	Global Command - Expert Mode
arp	g_arp
cat	g_cat
cp	g_cp
dmesg	g_dmesg
ethtool	g_ethtool
ls	g_ls
md5sum	g_md5sum
Mv	g_mv
Netstat	g_netstat
Reboot	g_reboot
tail	g_tail

gclish Command	Global Command - Expert Mode
`tcpdump`	`g_tcpdump`
`ifconfig`	`asg_ifconfig`
`top`	`g_top`

The parameters and options for the standard Linux command are available for the global command. In addition, you can use the `-b` parameter to select some or all SGMs for the global command.

Syntax

```
{<gclish_command> | <global_command>} [-b <sgm_ids>] <command_options>]
```

Parameter	Description
`-b` *<sgm_ids>*	Works with SGMs and/or Chassis as specified by <*sgm_ids*>. <*sgm_ids*> can be: • No <*sgm_ids*> specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`) Note: You can only select SGMs from one Chassis with this option.
<*gclish_command*>	Standard command in gclish
<*global_command*>	Global command as shown in the table, in Expert Mode
<*command_options*>	Standard command options for the specified command.

You can use one or more flags. However, do not use the `–l` and `–r` flags together.

g_reboot syntax

```
# g_reboot [-a]
```

g_reboot reboots all SGMs that are in the UP state. Use the `-a` option to reboot all SGMs in both the DOWN and UP states.

Global arp

This example shows the interfaces on all SGMs

```
> arp
1_01:
Address     HWtype  HWaddress          Flags Mask   Iface
192.0.2.2   ether   00:1C:7F:02:04:FE   C            Sync
172.23.9.28 ether   00:14:22:09:D2:22   C            eth1-Mgmt4
192.0.2.3   ether   00:1C:7F:03:04:FE   C            Sync
1_02:
Address     HWtype  HWaddress          Flags Mask   Iface
192.0.2.3   ether   00:1C:7F:03:04:FE   C            Sync
172.23.9.28 ether   00:14:22:09:D2:22   C            eth1-Mgmt4
192.0.2.1   ether   00:1C:7F:01:04:FE   C            Sync
```

```
1_03:
Address     HWtype  HWaddress           Flags Mask   Iface
192.0.2.1   ether   00:1C:7F:01:04:FE   C            Sync
172.23.9.28 ether   00:14:22:09:D2:22   C            eth1-Mgmt4
192.0.2.2   ether   00:1C:7F:02:04:FE   C            Sync
```

Global ls

This example runs `ls` from Expert Mode on SGMs 1_1, 1_2, and 1_3. The output shows the combined results for these SGMs.

```
# g_ls ls -b 1_1-1_3,2_1 /var/
-*- 4 blades: 1_01 1_02 1_03 -*-
CPbackup    ace    crash  lib   log   opt       run    suroot
CPsnapshot  cache  empty  lock  mail  preserve  spool  tmp
```

Global top

The global `top` command shows SGM processor activity in real time. The default output also shows a list of the most processor-intensive processes. In addition to the standard functionality of the Linux top command, global `top` adds these features for the 61000/41000 Security System:

Global `top` relies on the user configuration for the local top utility. The command uses the local SGM configuration file for configuring the output on the remote SGMs.

Syntax

```
> top [local] [-f [-o <filename>] [-n <iter>] | -s <filename>] -b <sgm_ids> [<top_params>] [-h]
```

Parameter	Description
`local`	Use the local configuration file
`-f`	Export the output to a file
`-o` *<filename>*	File and path of the output file Default: `/vat/log/gtop.`*<time>* Use with: `-f`
`-n` *<iter>*	Number of iterations Default: 1 Use with: `-f`
`-b` *<sgm_ids>*	Works with SGMs and/or Chassis as specified by *<sgm_ids>*. *<sgm_ids>* can be: • No *<sgm_ids>* specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`)

Parameter	Description
<*top_params*>	Parameters of the standard `top` command For more information, see the `top` documentation.
-s <*filename*>	Shows the content of the output file <*filename*>

Managing the g_top display

`top` uses a configuration file to manage output display. By default it copies and uses this configuration file from the local blade (usually located under `~/.toprc`). This file is copied to all SGMs and is used when `top` is run.

To manage the **`g_top`** display:

1. Run: `top`
2. Set the desired display view.
3. Save configuration (shift+w).
4. Run: `g_top`

Sending output to a file

At times, it is more convenient to send `g_top` output to a file, for example, when there are more SGMs than the screen can handle. To enable the file mode use: `-f`

Monitoring SGM Resources (asg resource)

Use this command to show SGM resource usage and thresholds for the entire NG 61000 Security System.

Syntax

```
> asg resource [-b <sgm_ids>]
> asg resource -h
```

Parameter	Description
`-b` <*sgm_ids*>	Works with SGMs and/or Chassis as specified by <*sgm_ids*>. <*sgm_ids*> can be: • No <*sgm_ids*> specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`)
`-h`	Shows usage and exits

Example

```
> asg resource
+-------------------------------------------------------------------------------------+
|Resource Table                                                                       |
+------------+--------------------------+------------+------------+-------------------+
|SGM ID      |Resource Name             |Usage       |Threshold   |Total              |
+------------+--------------------------+------------+------------+-------------------+
|1_01        |Memory                    |14%         |50%         |31.3G              |
|            |HD: /                     |22%         |80%         |19.4G              |
|            |HD: /var/log              |1%          |80%         |58.1G              |
|            |HD: /boot                 |19%         |80%         |288.6M             |
+------------+--------------------------+------------+------------+-------------------+
|1_02        |Memory                    |9%          |50%         |62.8G              |
|            |HD: /                     |23%         |80%         |19.4G              |
|            |HD: /var/log              |1%          |80%         |58.1G              |
|            |HD: /boot                 |19%         |80%         |288.6M             |
+------------+--------------------------+------------+------------+-------------------+
|1_03        |Memory                    |9%          |50%         |62.8G              |
|            |HD: /                     |23%         |80%         |19.4G              |
|            |HD: /var/log              |1%          |80%         |58.1G              |
|            |HD: /boot                 |19%         |80%         |288.6M             |
+------------+--------------------------+------------+------------+-------------------+
|2_01        |Memory                    |9%          |50%         |62.8G              |
|            |HD: /                     |23%         |80%         |19.4G              |
|            |HD: /var/log              |1%          |80%         |58.1G              |
|            |HD: /boot                 |19%         |80%         |288.6M             |
+------------+--------------------------+------------+------------+-------------------+
|2_02        |Memory                    |9%          |50%         |62.8G              |
|            |HD: /                     |23%         |80%         |19.4G              |
|            |HD: /var/log              |1%          |80%         |58.1G              |
|            |HD: /boot                 |19%         |80%         |288.6M             |
+------------+--------------------------+------------+------------+-------------------+
|2_03        |Memory                    |9%          |50%         |62.8G              |
|            |HD: /                     |23%         |80%         |19.4G              |
|            |HD: /var/log              |1%          |80%         |58.1G              |
|            |HD: /boot                 |19%         |80%         |288.6M             |
+------------+--------------------------+------------+------------+-------------------+
```

Notes

- The **SGM** column shows the SGM ID.
- The **Resource** column identifies the resource. There are four types of resources:
  - **Memory**
  - **HD** – Hard drive space (/)
  - **HD: /var/log** – Space on hard drive committed to log files
  - **HD: /boot** - Location of the kernel
- The **Usage** column shows the percentage of the resource in use.
- The **Threshold** gives an indication of the health and functionality of the component. When the value of the resource is greater than the threshold, an alert is sent. The threshold can be modified in `gclish`.
- The **Total** column is the total absolute value in units

For example, the first row shows that SGM1 on Chassis 1 has 31.3 GB of memory, 14% of which is used. An alert is sent if the usage is greater than 50%.

Searching for a Connection (asg search)

Use this command to:

- Search for a connection or a filtered list of connections.
- See which SGM handles the connection (actively or as backup), and on which Chassis.

You can run this command directly from gclish or in Interactive Mode, which lets you enter the parameters in the correct sequence. The `asg search` command also runs a consistency test between SGMs. This command supports both IPv4 and IPv6 connections.

Searching with the Command Line

Syntax

```
> asg search -help
> asg search [-v] [-vs <vs_ids>] [<source_ip> <dest_ip> <dest_port> <protocol>]
```

Parameter	Description
`-help`	Show the command syntax and help text.
Without parameters	Run in the interactive mode.
`-vs` <*vs_ids*>	Shows connections for the specified Virtual System. <*vs_ids*> can be: • No <*vs_ids*> (default) - Shows the current Virtual System context. • One Virtual System. • A comma-separated list of Virtual Systems (1, 2, 4, 5). • A range of Virtual Systems (VS 3-5). • `all` - Shows all Virtual Systems. **Note:** This parameter is only relevant in a VSX environment.
<*source_ip*>	Source IPv4 or IPv6 address.
<*dest_ip*>	Destination IPv4 or IPv6 address
<*dest_port*>	Destination port number.
<*protocol*>	IP Protocol.
<*source_port*>	Source port number.
`-v`	Shows connection indicators for • **F** - Firewall connection table • **S** - SecureXL connection table • **C** - Correction Layer table This in addition to the indicators for Active and Backup SGM.

Notes:

- You must enter the all parameters in the sequence shown in the above syntax.
- You can enter `\*` as a parameter to show all values for that parameter.
- The `-vs` parameter is only available for a 61000/41000 Security System running VSX.

Command Line Examples

One IPv4 source and destination for the TCP protocol

```
> asg search -v 192.0.2.4 192.0.2.15 \* tcp
```

```
Lookup for conn: <192.0.2.4, 192.0.2.15, *, tcp>, may take few seconds...

<192.0.2.4, 1130,  192.0.2.15, 49829, tcp> -> [2_01 A, 1_04 A]
<192.0.2.4, 36323, 192.0.2.15, 1130,  tcp> -> [2_01 A, 1_04 A]
<192.0.2.4, 1130,  192.0.2.15, 49851, tcp> -> [2_01 A, 1_04 A]
<192.0.2.4, 36308, 192.0.2.15, 1130,  tcp> -> [2_01 A, 1_04 A]
<192.0.2.4, 36299, 192.0.2.15, 1130,  tcp> -> [2_01 A, 1_04 A]
<192.0.2.4, 1130,  192.0.2.15, 49835, tcp> -> [2_01 A, 1_04 A]
<192.0.2.4, 1130,  192.0.2.15, 49856, tcp> -> [2_01 A, 1_04 A]
<192.0.2.4, 36331, 192.0.2.15, 1130,  tcp> -> [2_01 A, 1_04 A]
<192.0.2.4, 1130,  192.0.2.15, 49857, tcp> -> [2_01 A, 1_04 A]
<192.0.2.4, 1130,  192.0.2.15, 49841, tcp> -> [2_01 A, 1_04 A]
<192.0.2.4, 36315, 192.0.2.15, 1130,  tcp> -> [2_01 A, 1_04 A]
<192.0.2.4, 1130,  192.0.2.15, 49859, tcp> -> [2_01 A, 1_04 A]
<192.0.2.4, 36300, 192.0.2.15, 1130,  tcp> -> [2_01 A, 1_04 A]
<192.0.2.4, 36301, 192.0.2.15, 1130,  tcp> -> [2_01 A, 1_04 A]

Legend:
A - Active SGM
B - Backup SGM
C - Correction Layer table
F - Firewall connection table
S - SecureXL connection table
```

One IPv6 source, all destinations, source port 8080, and TCP protocol

```
> asg search 2620:0:2a03:16:2:33:0:1 \* 8080 tcp

<2620:0:2a03:16:2:33:0:1, 52117, 951::69cb:e42d:eac0:652f, 8080, tcp> -> [1_01 A, 2_01 B]
<2620:0:2a03:16:2:33:0:1, 62775, 951::69cb:e42d:eac0:652f, 8080, tcp> -> [1_01 A, 2_01 B]
<2620:0:2a03:16:2:33:0:1, 54378, 951::69cb:e42d:eac0:652f, 8080, tcp> -> [1_01 A, 2_01 B]
Legend:
A - Active SGM
B - Backup SGM
```

All sources, destinations, ports and protocols for VS0

```
> asg search -vs 0 \* \* \* \* \*.
Lookup for conn: <*, *, *, *, *>, may take few seconds...

<172.23.9.130, 18192, 172.23.9.138, 43563, tcp> -> [1_01 A]
<172.23.9.130, 32888, 172.23.9.138, 257, tcp> -> [1_01 A]
<172.23.9.130, 22, 194.29.47.14, 52120, tcp> -> [1_01 A]
<172.23.9.138, 257, 172.23.9.130, 32963, tcp> -> [1_01 A]
<172.23.9.130, 22, 194.29.47.14, 52104, tcp> -> [1_01 A]
<255.255.255.255, 67, 0.0.0.0, 68, udp> -> [1_01 A]
<172.23.9.138, 257, 172.23.9.130, 32864, tcp> -> [1_01 A]
<172.23.9.138, 257, 172.23.9.130, 32888, tcp> -> [1_01 A]
<172.23.9.138, 257, 172.23.9.130, 33465, tcp> -> [1_01 A]
<172.23.9.130, 22, 194.29.40.23, 65515, tcp> -> [1_01 A]
<172.23.9.130, 22, 194.29.47.14, 52493, tcp> -> [1_01 A]
<172.23.9.130, 18192, 172.23.9.138, 49059, tcp> -> [1_01 A]
<172.23.9.130, 18192, 172.23.9.138, 33356, tcp> -> [1_01 A]
<172.23.9.138, 33356, 172.23.9.130, 18192, tcp> -> [1_01 A]
<172.23.9.138, 43563, 172.23.9.130, 18192, tcp> -> [1_01 A]
<172.23.9.130, 32864, 172.23.9.138, 257, tcp> -> [1_01 A]
<0.0.0.0, 68, 255.255.255.255, 67, udp> -> [1_01 A]
<172.23.9.130, 32963, 172.23.9.138, 257, tcp> -> [1_01 A]
<172.23.9.130, 33465, 172.23.9.138, 257, tcp> -> [1_01 A]
<194.29.47.14, 52120, 172.23.9.130, 22, tcp> -> [1_01 A]
<194.29.47.14, 52104, 172.23.9.130, 22, tcp> -> [1_01 A]
<fe80::d840:5de7:8dbe:2345, 546, ff02::1:2, 547, udp> -> [1_01 A]
<194.29.47.14, 52493, 172.23.9.130, 22, tcp> -> [1_01 A]
<172.23.9.138, 49059, 172.23.9.130, 18192, tcp> -> [1_01 A]
<194.29.40.23, 65515, 172.23.9.130, 22, tcp> -> [1_01 A]
Legend:
A - Active SGM
B - Backup SGM
```

Searching with Interactive Mode

Interactive Mode lets you enter connection search parameters interactively in the required sequence as an alternative to the command line syntax.

To run asg search in Interactive Mode:

1. Run:

   ```
   > asg search [-vs <vs_ids>] [-v]
   ```

2. Enter these parameters in order.
   - Source IPv4 or IPv6 address
   - Destination IPv4 or IPv6 address
   - Destination port number
   - IP protocol
   - Source port number

   You can enter * to show all values for any parameter.

Interactive Mode Examples

Example 1 - One IPv4 source and destination with -v

```
> asg search -v

Please enter conn's 5 tuple:
----------------------------
Enter source IP (press enter for wildcard):
>192.0.2.4
Enter destination IP (press enter for wildcard):
>192.0.2.15
Enter destination port (press enter for wildcard):
>
Enter IP protocol ('tcp', 'udp', 'icmp' or enter for wildcard):
>tcp
Enter source port (press enter for wildcard):
>
Lookup for conn: <192.0.2.4, *, 192.0.2.15, *, tcp>, may take few seconds...
<192.0.2.4, 37408, 192.0.2.15, 1130, tcp> -> [2_01 AF, 1_04 AF]
<192.0.2.4, 1130, 192.0.2.15, 49670, tcp> -> [2_01 AF, 1_04 AF]
<192.0.2.4, 1130, 192.0.2.15, 49653, tcp> -> [2_01 AF, 1_04 AF]
<192.0.2.4, 37406, 192.0.2.15, 1130, tcp> -> [2_01 AF, 1_04 AF]
<192.0.2.4, 1130, 192.0.2.15, 49663, tcp> -> [2_01 AF, 1_04 AF]
<192.0.2.4, 1130, 192.0.2.15, 49658, tcp> -> [2_01 AF, 1_04 AF]
<192.0.2.4, 37407, 192.0.2.15, 1130, tcp> -> [2_01 AF, 1_04 AF]

Legend:
A - Active SGM
B - Backup SGM
C - Correction Layer table
F - Firewall connection table
S - SecureXL connection table
```

Example 2 - One IPv6 source with any Destination on port 8080 and TCP

```
> asg search 2620:0:2a03:16:2:33:0:1 \* 8080 tcp
Enter source IP (press enter for wildcard):
> 2620:0:2a03:16:2:33:0:1
Enter destination IP (press enter for wildcard):
>
Enter destination port (press enter for wildcard):
>8080
Enter IP protocol ('tcp', 'udp', 'icmp' or enter for wildcard):
>tcp
Enter source port (press enter for wildcard):
>
```

```
Lookup for conn: <2620:0:2a03:16:2:33:0:1, *, *, 8080, tcp>, may take few seconds...
<2620:0:2a03:16:2:33:0:1, 52117, 951::69cb:e42d:eac0:652f, 8080, tcp> -> [1_01 A, 2_01 B]
<2620:0:2a03:16:2:33:0:1, 62775, 951::69cb:e42d:eac0:652f, 8080, tcp> -> [1_01 A, 2_01 B]
<2620:0:2a03:16:2:33:0:1, 54378, 951::69cb:e42d:eac0:652f, 8080, tcp> -> [1_01 A, 2_01 B]

A - Active SGM
B - Backup SGM
```

Configuring Alerts for SGM and Chassis Events (asg alert)

Use `asg alert` an interactive wizard to configure alerts for SGM and Chassis events. Chassis events include hardware failure, recovery, and performance related events, and you can create other, general events.

An alert is sent when an event occurs, for example, when the value of a hardware resource is greater than the threshold. The alert message includes the Chassis ID, SGM ID and/or unit ID.

The wizard includes these options:

Option	Description
Full Configuration Wizard	Create a new alert
Edit Configuration	Change an existing alert
Show Configuration	Show existing alert configurations
Run Test	Run a test simulation to make sure that the alert works correctly

To create or change an alert:

1. Run:
   ```
   > asg alert
   ```
2. Select and configure these parameters as prompted by the wizard:
   - Alert type and related parameters
   - Event type
   - Alert mode

Alert Parameters

- SMS alert parameters
  - **SMS Provider URL** - Fully qualified URL to your SMS provider.
  - **HTTP proxy and port** (optional) – Necessary only if your Security Gateway requires a proxy server to reach the SMS provider.
  - **SMS rate limit** - Maximum number of SMS messages sent per hour. When there are too many messages, other messages may be combined together in one message.
  - **SMS user text** - Custom prefix for SMS messages.
- Email alert configuration
  - **SMTP server IP** - One or more SMTP servers to which the email alerts are sent.
  - **Email recipient addresses** - One or more recipient email addresses for each SMTP server.
  - **Periodic connectivity checks** - Run tests periodically to confirm connectivity with the SNMP servers. If there is no connectivity, alert messages are saved and sent in one email when connectivity is restored.
  - **Interval** - Interval, in minutes, between connectivity tests.

  - **Sender email address** - Sender email address for email alerts.
  - **Subject** - Subject header text for the email alert.
  - **Body text** - User-defined text for the alert message.

- SNMP alert parameters

  Define one or more SNMP managers to get SNMP traps sent from the Security Gateway. For each manager, configure these parameters as prompted:

  - **SNMP manager name** - Name for your SNMP manager (unique).
  - **SNMP manager IP** - Manager IP address (trap receiver).
  - **SNMP version** - SNMP version to use (v2cv3).
  - **SNMP v3 user name** - If using SNMP v3 authentication, you must configure this.
  - **SNMP v3 engine ID** - Unique SNMP v3 engine ID used by your system.
    Default = [0x80000000010203EA].
  - **SNMP v3 authentication protoco**l - MD5 or SHA.
  - **SNMP v3 authentication password** - Privacy password.
  - **SNMP v3 privacy protocol** - DES or AES.
  - **SNMP v3 privacy password** - Privacy password.
  - **SNMP user text** - Custom text for the SNMP trap messages.
  - **SNMP community string** - Community string for the SNMP manager.

Note - Some parameters do not show, based on your settings.

- Log alert parameters
  - There are no configurable parameters for log alerts.

Event types

You can select one or more event types:

- One event type.
- A comma-delimited list of more than one event type.
- `all` event types.

```
 ----------------------------------
1       | SGM State
2       | Chassis State
3       | Port State
4       | Pingable Hosts State
5       | System Monitor Daemon
6       | Route State
7       | Diagnostics
Hardware Monitor events:
8       | Fans
9       | SSM
10      | CMM
11      | Power Supplies
12      | CPU Temperature
Performance events:
13      | Concurrent Connections
14      | Connection Rate
15      | Packet Rate
16      | Throughput
17      | CPU Load
18      | Hard Drive Utilization
19      | Memory Utilization
```

Alert Modes

- **Enabled** - An alert is sent for the selected events.
- **Disabled** - No alert is sent for the selected events.
- **Monitor** - A log entry is generated instead of an alert.

Diagnostic Events

We recommend that you run the `asg diag verify` diagnostic tests on a regular basis. Alerts are sent if tests have failed. The alerts continue to show on the Message of the Day (MOTD) until the issues are resolved. Optionally, you can disable the MOTD. When the failed test has resolved, a Clear Alert message appears the next time that test runs. You can also run `asg diag verify` manually to make sure that the issue is resolved.

- The tests run by default at 01:00 each day. You can manually reset the default time.
- The daily test runs all tests, but you can exclude any tests.
  - When you manually run `asg diag verify`, all of the tests run, including those excluded from the automatic daily run.
- All failed tests show in the MOTD but you can disable this this feature.

To change the default time:

1. Open `/var/opt/CPsuite-R76/fw1/conf/asgsnmp.conf` in a text editor.
2. Change the `asg_diag_alert_wrapper`.
3. Copy this file to all other SGMs:

   > `asg_cp2blades <`*file_path*`>`

To disable the MOTD:

1. Open `/var/opt/CPsuite-R76/fw1/conf/asg_diag_config` in a text editor.
2. Set the `motd` parameter to `off`.
3. Copy this file to all other SGMs:

   > `asg_cp2blades <`*file_path*`>`
4. Run:

   ```
   > asg diag verify
   ```

   Step 4 is necessary for the change to take effect. You can also wait for the next time `asg diag verify` runs automatically.

To exclude specified tests from the daily automatic run:

1. Open `/var/opt/CPsuite-R76/fw1/conf/asg_diag_config` in a text editor.
2. Add this line to the file:

   `excluded_tests=[<`*Test1*`>][,<`*Test2*`>,...]`
3. Copy this file to all other SGMs:

   > `asg_cp2blades <`*file_path*`>`

To exclude failed test notifications in the MOTD:

1. Open `/var/opt/CPsuite-R76/fw1/conf/asg_diag_config` in a text editor.
2. Set the `failed_tests_motd` parameter to `off`.

3. Copy this file to all other SGMs:

   ```
   > asg_cp2blades <file_path>
   ```

4. Run:

   ```
   > asg diag verify
   ```

   Step 4 is necessary for this change to take effect. You can also wait for the next time `asg diag` runs automatically.

Known Limitations of asg diag Verification Tests

By default, the `asg diag resource verifier` option only shows a warning about resource mismatches between SGMs. The verification test results show as Passed in the output and no further action is taken.

You can change the default behavior with this procedure:

1. Open `$FWDIR/conf/asg_diag_config` in a text editor.
2. Search for `MismatchSeverity`.
3. Change the parameter value to one of these values:
   - `fail` - Verification test result is set to 'Failed'
   - `warn` - Verification test result is set to 'Passed' and a warning shows
   - `ignore` - Verification test result is set to 'Ignore' and no errors show

Collecting System Diagnostics (asg diag)

The `asg diag` command runs a specified set of diagnostic tests. By default, the full set of tests runs. You can optionally select the tests to run. The output shows the result of each test (Passed or Failed) and the location of the output log file.

Syntax

```
asg diag list|verify|print|except [<Test1>][,<Test2>,...]
asg diag purge [<num_logs>]
asg diag stat
```

Parameter	Description
`list`	Show the list of tests to run.
`verify`	Run tests and show a summary of the results.
`print`	Run tests and show the full output and summary of the results.
`except`	Run all except the specified tests and show a summary of the results.
`[<Test1>][,<Test2>,...]`	Comma separated list of test IDs. To see a list of test IDs, run: > `asg diag list`
`purge`	Delete the `asg diag` logs except for the newest.

Parameter	Description
[<*num_logs*>]	The number of most recent logs to keep when `asg diag` logs files. Default = 5.
`stat`	Shows the last summary of the results.

This example shows the output for the system component tests.

```
| Tests Status                                                                       |
-------------------------------------------------------------------------------------
| ID | Title              | Result     | Reason                                     |
-------------------------------------------------------------------------------------
| System Components                                                                 |
-------------------------------------------------------------------------------------
|  1 | System Health      | Passed     |                                            |
|  2 | Hardware           | Passed     |                                            |
|  3 | Resources          | Failed (!) | (1)Memory capacity                         |
|    |                    |            | (2)Primary HD capacity                     |
|    |                    |            | (3)Log HD capacity                         |
|    |                    |            | (4)Boot HD capacity                        |
|    |                    |            | (5)Primary HD capacity mismatch            |
|  4 | Software Versions  | Failed (!) |                                            |
|  5 | Software Provision | Passed     |                                            |
|  6 | CPU Type           | Failed (!) | (1)Non-compliant CPU type                  |
|  7 | Media Details      | Passed     |                                            |
|  8 | Chassis ID         | Passed     |                                            |
-------------------------------------------------------------------------------------
```

Showing the Tests

This example shows the full list of diagnostic tests. The list shows:

- Test ID
- Test name
- Command that `asg diag` runs to show the specified test results

```
> asg diag list

-----------------------------------------------------------------
| ID | Title              | Command                             |
-----------------------------------------------------------------
| System Components                                             |
-----------------------------------------------------------------
|  1 | System Health      | asg stat -v                         |
|  2 | Hardware           | asg hw_monitor -v                   |
|  3 | Resources          | asg resource                        |
|  4 | Software Versions  | asg_version verify -v               |
|  5 | Software Provision | asg_provision                       |
|  6 | CPU Type           | cpu_socket_verifier -v              |
|  7 | Media Details      | transceiver_verifier -v             |
|  8 | Chassis ID         | verify_chassis_id                   |
-----------------------------------------------------------------
| Policy and Configuration                                      |
-----------------------------------------------------------------
|  9 | Distribution Mode  | distutil verify -v                  |
| 10 | DXL Balance        | dxl stat                            |
| 11 | Policy             | asg policy verify -a                |
| 12 | AMW Policy         | asg policy verify_amw -a            |
| 13 | SWB Updates        | asg_swb_update_verifier -v          |
| 16 | Security Group     | asg security_group diag             |
| 17 | SPI Affinity       | spi_affinity_verifier -v            |
| 18 | Clock              | clock_verifier -v                   |
| 19 | Licenses           | asg_license_verifier -v             |
| 21 | LTE                | lte_verifier -v                     |
| 22 | IPS Enhancement    | asg_ips_enhance status              |
| 23 | Configuration File | config_verify -v                    |
-----------------------------------------------------------------
| Networking                                                    |
-----------------------------------------------------------------
| 24 | MAC Setting        | mac_verifier -v                     |
| 25 | ARP Consistency    | asg_arp -v                          |
| 26 | Interfaces         | interface_verifier -v               |
| 27 | Bond               | asg_bond -v                         |
| 28 | Bridge             | asg_brs_verifier -v                 |
| 29 | IPv4 Route         | asg_route                           |
| 30 | IPv6 Route         | asg_route -6                        |
| 31 | OS Route Cache     | asg_dst_route --diag                |
| 32 | Dynamic Routing    | asg_dr_verifier                     |
| 33 | Local ARP          | asg_local_arp_verifier -v           |
| 34 | Port Speed         | asg_port_speed verify               |
| 35 | SSM QoS            | asg_qos_verify                      |
| 36 | IGMP Consistency   | asg_igmp                            |
| 37 | PIM Neighbors      | asg_pim_neighbors                   |
| 38 | ACL Filter         | acl_filter_verifier                 |
-----------------------------------------------------------------
| DoS                                                           |
-----------------------------------------------------------------
| 39 | SYN Defender       | asg_synatk                          |
| 40 | F2F Quota          | asg_f2fq                            |
-----------------------------------------------------------------
| Misc                                                          |
-----------------------------------------------------------------
| 41 | Core Dumps         | core_dump_verifier -v               |
| 42 | Syslog             | asg_syslog verify                   |
| 43 | Processes          | asg_process_verifier -v             |
| 44 | Performance hogs   | asg_perf_hogs                       |
-----------------------------------------------------------------
| Run "asg diag print <TestNum>" to display test verbose output |
-----------------------------------------------------------------
```

Running all Diagnostic Tests

This example shows the summary output for all diagnostic tests. When a test fails, the reasons for failure show in the **Reason** column.

```
> asg diag verify
Duration of tests vary and may take a few minutes to complete

----------------------------------------------------------------
| Tests Status                                                                 |
-------------------------------------------------------------------------------
-
| ID | Title              | Result     | Reason                                |
-------------------------------------------------------------------------------
-
| System Components                                                            |
-------------------------------------------------------------------------------
-
|  1 | System Health      | Failed (!) | (1)Chassis 1 error                    |
|    |                    |            | (2)Chassis 2 error                    |
|  2 | Hardware           | Failed (!) | (1)Chassis fan is missing             |
|    |                    |            | (2)CMM is missing                     |
|    |                    |            | (3)CMM active/standby configuration   |
|    |                    |            | (4)SSM is missing                     |
|  3 | Resources          | Passed     | (1)Memory capacity mismatch           |
|  4 | Software Versions  | Failed (!) |                                       |
|  5 | Software Provision | Passed     |                                       |
|  6 | CPU Type           | Passed     |                                       |
|  7 | Media Details      | Passed     |                                       |
|  8 | Chassis ID         | Passed     |                                        |
--------------------------------------------------------------------------------
-
| Policy and Configuration                                                      |
--------------------------------------------------------------------------------
-
|  9 | Distribution Mode  | Passed     |                                        |
--------------------------------------------------------------------------------
-
| VSX Configuration                                                             |
--------------------------------------------------------------------------------
-
| 10 | USER KERNEL Dist   | Passed     |                                        |
--------------------------------------------------------------------------------
-
| Policy and Configuration                                                      |
--------------------------------------------------------------------------------
-
| 11 | DXL Balance        | Passed     |                                        |
| 12 | Policy             | Passed     |                                        |
| 13 | AMW Policy         | Passed     | (1)Not configured                      |
| 14 | SWB Updates        | Passed     | (1)Not configured                      |
--------------------------------------------------------------------------------
-
| VSX Configuration                                                             |
--------------------------------------------------------------------------------
-
| 15 | VSX Configuration  | Passed     |                                        |
| 16 | HW Utilization     | Failed (!) | (1)Connection capacity is too low      |
| 17 | BMAC VMAC verify   | Passed     |                                        |
--------------------------------------------------------------------------------
-
| Policy and Configuration                                                      |
--------------------------------------------------------------------------------
-
| 18 | Installation       | Passed     |                                        |
| 19 | Security Group     | Passed     |                                        |
| 20 | Cores Distribution | Passed     |                                        |
| 21 | SPI Affinity       | Passed     | (1)Not configured                      |
| 22 | Clock              | Passed     |                                        |
| 23 | Licenses           | Passed     | (1)Trial license installed             |
| 24 | Hide NAT range     | Passed     | (1)Not configured                      |
| 25 | LTE                | Passed     | (1)Not configured                      |
| 26 | IPS Enhancement    | Passed     | (1)Not configured                      |
| 27 | Configuration File | Passed     |                                        |
--------------------------------------------------------------------------------
-
| Networking                                                                    |
--------------------------------------------------------------------------------
-
| 28 | MAC Setting        | Passed     |                                        |
| 29 | ARP Consistency    | Passed     |                                        |
| 30 | Interfaces         | Failed (!) | (1)RX drop                             |
|    |                    |            | (2)Interface down                      |
| 31 | Bond               | Failed (!) |                                        |
| 32 | Bridge             | Passed     | (1)Not configured                      |
| 33 | IPv4 Route         | Passed     |                                        |
| 34 | IPv6 Route         | Passed     | (1)Not configured                      |
| 35 | OS Route Cache     | Passed     |                                        |
| 36 | Dynamic Routing    | Passed     |                                        |
| 37 | Local ARP          | Passed     | (1)Not configured                      |
| 38 | Port Speed         | Failed (!) | (1)Inconsistency between chassis and   |
|    |                    |            | conf file                              |
| 39 | SSM QoS            | Passed     |                                        |
| 40 | IGMP Consistency   | Passed     | (1)Not configured                    |
| 41 | PIM Neighbors      | Passed     | (1)Not configured                    |
| 42 | ACL Filter         | Passed     |                                      |
-------------------------------------------------------------------------------
-
| DoS                                                                          |
-------------------------------------------------------------------------------
-
| 43 | SYN Defender       | Passed     |                                      |
| 44 | F2F Quota          | Passed     |                                      |
-------------------------------------------------------------------------------
-
| Misc                                                                         |
-------------------------------------------------------------------------------
-
| 45 | Core Dumps         | Failed (!) |                                      |
| 46 | Syslog             | Passed     | (1)Log server is not configured      |
| 47 | Processes          | Passed     |                                      |
| 48 | Performance hogs   | Failed (!) |                                      |
-------------------------------------------------------------------------------
-
| Tests Summary                                                                |
-------------------------------------------------------------------------------
-
| Passed: 39/48 tests                                                          |
| Run: "asg diag list 1,2,4,7,16,30,31,38,45,48" to view a complete list of fail
|
| ed tests                                                                     |
| Output file: /var/log/verifier_sum.1-48.2015-12-30_17-07-38.txt              |
-----------------------------------------------------------------
```

Running Specified Diagnostic Tests

This example collects diagnostic information for specified tests.

```
> asg diag verify 1,2,3,4,5,30
Duration of tests vary and may take a few minutes to complete

---------------------------------------------------------------------
| Tests Status                                                      |
---------------------------------------------------------------------
| ID | Title              | Result     | Reason                     |
---------------------------------------------------------------------
| System Components                                                 |
---------------------------------------------------------------------
|  1 | System Health      | Failed (!) | (1)Chassis 1 error         |
|  2 | Hardware           | Failed (!) | (1)SSM is down             |
|  3 | Resources          | Passed     |                            |
|  4 | Software Versions  | Failed (!) |                            |
|  5 | Software Provision | Passed     |                            |
---------------------------------------------------------------------
| Networking                                                        |
---------------------------------------------------------------------
| 30 | IPv6 Route         | Passed     | (1)Not configured          |
---------------------------------------------------------------------
| Tests Summary                                                     |
---------------------------------------------------------------------
| Passed: 3/6 tests                                                 |
| Run: "asg diag list 1,2,4" to view a complete list of failed tests |
| Output file: /var/log/verifier_sum.1-5.30.2014-07-21_11-16-16.txt |
---------------------------------------------------------------------
```

Performance Hogs - asg_perf_hogs

You can run `asg_perf_hogs` by itself or as part of `asg diag`.

When you run `asg_perf_hogs` by itself, you can get the full details of all the tests it runs. When you run `asg diag`, it shows a general result of `asg_perf_hogs` in the **Misc** section of the diagnostic test output.

Syntax

```
> asg_perf_hogs
```

Output

```
-----------------------------------------------------------------
|  Status  |  Test performed                                    |
-----------------------------------------------------------------
| [PASSED] | Long running processes                             |
| [PASSED] | SecureXL status                                    |
| [PASSED] | PPACK debug flags                                  |
| [PASSED] | FW1 debug flags                                    |
| [PASSED] | Local logging                                      |
| [PASSED] | Templates disabled from rule                       |
| [PASSED] | Correction table entries                           |
| [PASSED] | Delayed notifications                              |
| [PASSED] | Routing cache entries                              |
| [PASSED] | Swap saturation                                    |
| [PASSED] | Neighbour table overflow                           |
| [PASSED] | Soft lockups                                       |
-----------------------------------------------------------------
```

When `asg diag` runs `asg_perf_hogs`, it shows a general result of in the **Misc** section of the diagnostic test output.

- If all of the `asg_perf_hogs` tests pass, `asg diag` shows **PASSED** as the result.
- If at least one of the `asg_perf_hogs` tests fails, `asg diag` shows **FAILED** as the result.

Configuration

You can configure `asg_perf_hogs` using the file: `$FWDIR/conf/performance_hogs.conf`

```
[tests]
long_running_procs=1
accel_off=1
sim_debug_flags=1
fw1_debug_flags=1
local_logging=1
templates_disabled_from_rule=1
correction_table_entries=1
routing_cache_entries=1
swap_saturation=1
delayed_notifications=1
neighbour_table_overflow=1
soft_lockups=1
[correction_table_entries]
threshold=10

[long_running_procs]
elapsed_time=60
processes_to_check=("fw ctl zdebug" "fw ctl debug" "fw ctl kdebug" "fw monitor" "sim
dbg" "tcpdump")

[routing_cache_entries]
threshold=90

[swap_saturation]
threshold=50

[neighbour_table_overflow]
timeout=3600

[soft_lockups]
timeout=3600
```

The **tests** section lets you enable and disable which tests to run.

To enable or disable a test:

In the `tests` section of `$FWDIR/conf/performance_hogs.conf`, set the parameter value:

- `1` = enable
- `0` = disable

To configure a test:

1. Find the configuration section for the test in `$FWDIR/conf/performance_hogs.conf`. If it does not exist, add the section with this format:
   [<*test_name*>]
2. Change or add the parameters for the test. See the tables below for allowed parameters.

Note - Not all the tests can be configured.

long_running_procs

`long_running_procs` - Confirms that certain processes do not run longer than the configured time.

This test runs on all VSX contexts.

Parameter	Description
`elapsed_time`	Longest time in seconds a process should run Default: 60 seconds Minimum recommended value: 30
`processes_to_check`	List of process to check Each process must be in quotes. Put a space between each test. Default: `"fw ctl zdebug" "fw ctl debug" "fw ctl kdebug" "fw monitor" "sim dbg" "tcpdump"` Example: `processes_to_check=("fw ctl zdebug" "fw ctl debug" "fw ctl kdebug" "fw monitor" "sim dbg" "tcpdump")`

Example:

```
-----------------------------------------------------------------
|  Status  |  Test performed                                    |
-----------------------------------------------------------------
| [FAILED] | Long running processes                             |
| [PASSED] | SecureXL status                                    |
| [PASSED] | PPACK debug flags                                  |
| [PASSED] | FW1 debug flags                                    |
| [PASSED] | Local logging                                      |
| [PASSED] | Templates disabled from rule                       |
| [PASSED] | Correction table entries                           |
| [PASSED] | Delayed notifications                              |
| [PASSED] | Routing cache entries                              |
| [PASSED] | Swap saturation                                    |
| [PASSED] | Neighbour table overflow                           |
-----------------------------------------------------------------
```

```
Found potential CPU hogging processes:
-----------------------------------------------------------------
Blade    PID     ELAPSED     TIME CMD
[1_01]  1484       03:48 00:00:00 tcpdump -nnni eth1-01

Found the following issues:
-----------------------------------------------------------------
[ All] The process 'tcpdump' is running for more than 60 seconds
```

accel_off

`accel_off` - Confirms that SecureXL is working.

The test runs on the current VSX context only.

This test has no configuration options.

Example

```
-----------------------------------------------------------------
|  Status  |  Test performed                                    |
-----------------------------------------------------------------
| [PASSED] | Long running processes                             |
| [FAILED] | SecureXL status                                    |
| [PASSED] | PPACK debug flags                                  |
| [PASSED] | FW1 debug flags                                    |
| [PASSED] | Local logging                                      |
| [PASSED] | Templates disabled from rule                       |
| [PASSED] | Correction table entries                           |
| [PASSED] | Delayed notifications                              |
| [PASSED] | Routing cache entries                              |
| [PASSED] | Swap saturation                                    |
| [PASSED] | Neighbour table overflow                           |
-----------------------------------------------------------------
 Found the following issues:
-----------------------------------------------------------------
[ All] SecureXL acceleration is disabled!
```

sim_debug_flags

`sim_debug_flags` - Confirms that the PPACK debug flags that are not enabled by default, stay in the not-enabled position.

This test runs on all VSX contexts.

This test has no configuration options.

Example

```
-----------------------------------------------------------------
|  Status  |  Test performed                                    |
-----------------------------------------------------------------
| [PASSED] | Long running processes                             |
| [PASSED] | SecureXL status                                    |
| [FAILED] | PPACK debug flags                                  |
| [PASSED] | FW1 debug flags                                    |
| [PASSED] | Local logging                                      |
| [PASSED] | Templates disabled from rule                       |
| [PASSED] | Correction table entries                           |
| [PASSED] | Delayed notifications                              |
| [PASSED] | Routing cache entries                              |
| [PASSED] | Swap saturation                                    |
| [PASSED] | Neighbour table overflow                           |
-----------------------------------------------------------------
 Found the following issues:
-----------------------------------------------------------------
```

```
[ All] PPACK debug flags are set: Module: vpn; ; Flags: vpnpkt
```

fw1_debug_flags

`fw1_debug_flags` - Confirms that FW1 debug flags that are not enabled by default, stay in the not-enabled position

This test runs on all VSX contexts.

This test has no configuration options.

Example

```
-----------------------------------------------------------------
|  Status  |  Test performed                                    |
-----------------------------------------------------------------
| [PASSED] | Long running processes                             |
| [PASSED] | SecureXL status                                    |
| [PASSED] | PPACK debug flags                                  |
| [FAILED] | FW1 debug flags                                    |
| [PASSED] | Local logging                                      |
| [PASSED] | Templates disabled from rule                       |
| [PASSED] | Correction table entries                           |
| [PASSED] | Delayed notifications                              |
| [PASSED] | Routing cache entries                              |
| [PASSED] | Swap saturation                                    |
| [PASSED] | Neighbour table overflow                           |
-----------------------------------------------------------------
 Found the following issues:
-----------------------------------------------------------------
[ All] FW1 debug flags are set:; Module: fw; ; Flags: error warning packet
```

local_logging

`local_logging` - Confirms that logs are written to a log server and not a local server.

This test runs on the current VSX context only.

This test has no configuration options.

Example

```
-----------------------------------------------------------------
|  Status  |  Test performed                                    |
-----------------------------------------------------------------
| [PASSED] | Long running processes                             |
| [PASSED] | SecureXL status                                    |
| [PASSED] | PPACK debug flags                                  |
| [PASSED] | FW1 debug flags                                    |
| [FAILED] | Local logging                                      |
| [PASSED] | Templates disabled from rule                       |
| [PASSED] | Correction table entries                           |
| [PASSED] | Delayed notifications                              |
| [PASSED] | Routing cache entries                              |
| [PASSED] | Swap saturation                                    |
| [PASSED] | Neighbour table overflow                           |
-----------------------------------------------------------------
 Found the following issues:
-----------------------------------------------------------------
[ All] Local logging is active: No connection with log server!
```

templates_disabled_from_rule

`templates_disabled_from_rule` - Confirms that no templates that are disabled because they mismatch the firewall rules.

This test runs regardless of the VSX context.

This test has no configuration options.

Example

```
------------------------------------------------------------------
|  Status  |  Test performed                                     |
------------------------------------------------------------------
| [PASSED] | Long running processes                              |
| [PASSED] | SecureXL status                                     |
| [PASSED] | PPACK debug flags                                   |
| [PASSED] | FW1 debug flags                                     |
| [PASSED] | Local logging                                       |
| [FAILED] | Templates disabled from rule                        |
| [PASSED] | Correction table entries                            |
| [PASSED] | Delayed notifications                               |
| [PASSED] | Routing cache entries                               |
| [PASSED] | Swap saturation                                     |
| [PASSED] | Neighbour table overflow                            |
------------------------------------------------------------------
 Found the following issues:
------------------------------------------------------------------
[ All] Templates are being disabled from rule(s): Accept Templates : disabled by
Firewall; disabled from rule #1; NAT Templates: disabled by Firewall; disabled from
rule #1
```

correction_table_entries

correction_table_entries makes sure that size ratio between corrections table and the connections table is not above the threshold.

This test runs on the current VSX Context only.

Parameter	Description
threshold	Allowed size ratio between the corrections table and the connections table Recommended range: 5-95

Example

```
------------------------------------------------------------------
|  Status  |  Test performed                                     |
------------------------------------------------------------------
| [PASSED] | Long running processes                              |
| [PASSED] | SecureXL status                                     |
| [PASSED] | PPACK debug flags                                   |
| [PASSED] | FW1 debug flags                                     |
| [PASSED] | Local logging                                       |
| [PASSED] | Templates disabled from rule                        |
| [FAILED] | Correction table entries                            |
| [PASSED] | Delayed notifications                               |
| [PASSED] | Routing cache entries                               |
| [PASSED] | Swap saturation                                     |
| [PASSED] | Neighbour table overflow                            |
------------------------------------------------------------------
 Found the following issues:
------------------------------------------------------------------
[ All] Correction table has 5 entries and is larger than 10% of connections table
(20 entries)
```

delayed_notifications

`delayed_notifications` - Confirms that delayed notifications are enabled. The output shows if delayed notifications are disabled for all services, or only for HTTP.

The test runs on all VSX contexts.

This test has no configuration options.

Example

```
-----------------------------------------------------------------
|  Status  |  Test performed                                    |
-----------------------------------------------------------------
| [PASSED] | Long running processes                             |
| [PASSED] | SecureXL status                                    |
| [PASSED] | PPACK debug flags                                  |
| [PASSED] | FW1 debug flags                                    |
| [PASSED] | Local logging                                      |
| [PASSED] | Templates disabled from rule                       |
| [PASSED] | Correction table entries                           |
| [FAILED] | Delayed notifications                              |
| [PASSED] | Routing cache entries                              |
| [PASSED] | Swap saturation                                    |
| [PASSED] | Neighbour table overflow                           |
-----------------------------------------------------------------
 Found the following issues:
-----------------------------------------------------------------
[ All] Delayed notifications for http is disabled.
```

routing_cache_entries

`routing_cache_entries` - Confirms that the IPv4 route cache capacity is not above a certain threshold.

This test runs on the current VSX context only.

Parameter	Description
`threshold`	Percent capacity of the IPv4 route cache that should not be exceeded Default: 90 Recommended range: 75-95

Example

```
-----------------------------------------------------------------
|  Status  |  Test performed                                    |
-----------------------------------------------------------------
| [PASSED] | Long running processes                             |
| [PASSED] | SecureXL status                                    |
| [PASSED] | PPACK debug flags                                  |
| [PASSED] | FW1 debug flags                                    |
| [PASSED] | Local logging                                      |
| [PASSED] | Templates disabled from rule                       |
| [PASSED] | Correction table entries                           |
| [PASSED] | Delayed notifications                              |
| [FAILED] | Routing cache entries                              |
| [PASSED] | Swap saturation                                    |
| [PASSED] | Neighbour table overflow                           |
-----------------------------------------------------------------
 Found the following issues:
-----------------------------------------------------------------
[ All] Routing cache is 93% full (983731 out of 1048576 entries).
```

swap_saturation

`swap_saturation` - Confirms that swap file usage is not above the threshold.

This test runs regardless of the VSX Context.

Parameter	Description
`threshold`	Percent usage of the swap file allowed Recommended range: 75-99

Example

```
-----------------------------------------------------------------
|  Status  |  Test performed                                    |
-----------------------------------------------------------------
| [PASSED] | Long running processes                             |
| [PASSED] | SecureXL status                                    |
| [PASSED] | PPACK debug flags                                  |
| [PASSED] | FW1 debug flags                                    |
| [PASSED] | Local logging                                      |
| [PASSED] | Templates disabled from rule                       |
| [PASSED] | Correction table entries                           |
| [PASSED] | Delayed notifications                              |
| [PASSED] | Routing cache entries                              |
| [FAILED] | Swap saturation                                    |
| [PASSED] | Neighbour table overflow                           |
-----------------------------------------------------------------
 Found the following issues:
-----------------------------------------------------------------
[ All] Swap saturation is 90%. Total swap space: 1044216 bytes, used: 950000 bytes.
```

neighbour_table_overflow

`neighbour_table_overflow` - Confirms that the ARP cache did not overflow.

This test runs regardless of the VSX context.

Parameter	Description
`timeout`	Number of seconds to look in `/var/log/messages` for ARP cache overloaded messages. Recommended range: 300-86400

To learn how to adjust the ARP cache, see sk43772 http://supportcontent.checkpoint.com/solutions?id=sk43772.

Example

```
-----------------------------------------------------------------
|  Status  |  Test performed                                    |
-----------------------------------------------------------------
| [PASSED] | Long running processes                             |
| [PASSED] | SecureXL status                                    |
| [PASSED] | PPACK debug flags                                  |
| [PASSED] | FW1 debug flags                                    |
| [PASSED] | Local logging                                      |
| [PASSED] | Templates disabled from rule                       |
| [PASSED] | Correction table entries                           |
| [PASSED] | Delayed notifications                              |
| [PASSED] | Routing cache entries                              |
| [PASSED] | Swap saturation                                    |
| [FAILED] | Neighbour table overflow                           |
-----------------------------------------------------------------
Found the following issues:
-----------------------------------------------------------------
[ All] Neighbour table overflow occurred during the last 3600 seconds. Please see
solution SK43772 for information how to configure arp cache size.
```

soft_lockups

`soft_lockups` - Confirms there are no kernel soft lockups in the timeout period.

This test runs regardless of the VSX Context.

Parameter	Description
`timeout`	Number of seconds to look back in `/var/log/messages` for kernel soft lockup messages. Default: 3600 Recommended range: 300-86400

Example

```
-----------------------------------------------------------------
|  Status  |  Test performed                                    |
-----------------------------------------------------------------
| [PASSED] | Long running processes                             |
| [PASSED] | SecureXL status                                    |
| [PASSED] | PPACK debug flags                                  |
| [PASSED] | FW1 debug flags                                    |
| [PASSED] | Local logging                                      |
| [PASSED] | Templates disabled from rule                       |
| [PASSED] | Correction table entries                           |
| [PASSED] | Delayed notifications                              |
| [PASSED] | Routing cache entries                              |
| [PASSED] | Swap saturation                                    |
| [PASSED] | Neighbour table overflow                           |
| [FAILED] | Soft lockups                                       |
-----------------------------------------------------------------
 Found the following issues:
-----------------------------------------------------------------
[1_01] Soft lockup occurred during the last 3600 seconds.
```

Troubleshooting Failures

This example shows how to use `asg diag` to troubleshoot a failed diagnostic test. In this case, the test shows that two fans are down and the CPU temperature exceeds its threshold. The output identifies the failed components.

```
> asg diag verify 2
------------------------------------------------------------------------------
| Tests Status                                                               |
------------------------------------------------------------------------------
| ID | Title               | Result | Reason                                 |
------------------------------------------------------------------------------
| System Components                                                          |
------------------------------------------------------------------------------
|  2 | Hardware            | Failed | (1)Chassis fan is down                 |
|    |                     |        | (2)Chassis fan exceeds threshold       |
|    |                     |        | (3)CPU exceeds threshold               |
------------------------------------------------------------------------------
| Tests Summary                                                              |
------------------------------------------------------------------------------
| Passed: 0/1 test                                                           |
| Run: "asg diag list 2" to view a complete list of failed tests              |
| Output file: /var/log/verifier_sum.2.2014-02-17_10-58-31.txt                |
------------------------------------------------------------------------------
> asg diag print 2
-----------------------------------------------------------------------------
| Hardware Monitor                                                          |
-----------------------------------------------------------------------------
| Sensor           | Location       | Value | Threshold | Units     | State |
-----------------------------------------------------------------------------
| Chassis 1                                                                 |
-----------------------------------------------------------------------------
| CMM              | bay 1          | 1     | 0         | <S,D>/<A>   | 1   |
| CMM              | bay 2          | 0     | 0         | <S,D>/<A>   | 1   |
| CPUtemp          | blade 1, CPU0  | 0     | 65        | Celsius     | 1   |
| CPUtemp          | blade 1, CPU1  | 0     | 65        | Celsius     | 1   |
| CPUtemp          | blade 2, CPU0  | 44    | 65        | Celsius     | 1   |
| CPUtemp          | blade 2, CPU1  | 41    | 65        | Celsius     | 1   |
| CPUtemp          | blade 3, CPU0  | 44    | 65        | Celsius     | 1   |
| CPUtemp          | blade 3, CPU1  | 40    | 65        | Celsius     | 1   |
| CPUtemp          | blade 4, CPU0  | 47    | 65        | Celsius     | 1   |
| CPUtemp          | blade 4, CPU1  | 43    | 65        | Celsius     | 1   |
| CPUtemp          | blade 5, CPU0  | 46    | 65        | Celsius     | 1   |
| CPUtemp          | blade 5, CPU1  | 42    | 65        | Celsius     | 1   |
| Fan              | bay 1, fan 1   | 0     | 11        | Speed Level | 0   |
| Fan              | bay 1, fan 2   | 0     | 11        | Speed Level | 0   |
| Fan              | bay 2, fan 1   | 15    | 11        | Speed Level | 1   |
| Fan              | bay 2, fan 2   | 15    | 11        | Speed Level | 1   |
| Fan              | bay 3, fan 1   | 15    | 11        | Speed Level | 1   |
| Fan              | bay 3, fan 2   | 15    | 11        | Speed Level | 1   |
| PowerConsumption | N/A            | 2471  | 4050      | Watts       | 1   |
| PowerUnit(AC)    | bay 1          | 0     | 0         | NA          | 1   |
| PowerUnit(AC)    | bay 2          | 0     | 0         | NA          | 1   |
| PowerUnit(AC)    | bay 3          | 0     | 0         | NA          | 1   |
| PowerUnit(AC)    | bay 4          | 0     | 0         | NA          | 0   |
| PowerUnit(AC)    | bay 5          | 0     | 0         | NA          | 0   |
| PowerUnitFan     | bay 1, fan 1   | 0     | 0         | NA          | 1   |
| PowerUnitFan     | bay 1, fan 2   | 0     | 0         | NA          | 1   |
| PowerUnitFan     | bay 2, fan 1   | 0     | 0         | NA          | 1   |
| PowerUnitFan     | bay 2, fan 2   | 0     | 0         | NA          | 1   |
| PowerUnitFan     | bay 3, fan 1   | 0     | 0         | NA          | 1   |
| PowerUnitFan     | bay 3, fan 2   | 0     | 0         | NA          | 1   |
| PowerUnitFan     | bay 4, fan 1   | 0     | 0         | NA          | 0   |
| PowerUnitFan     | bay 4, fan 2   | 0     | 0         | NA          | 0   |
| PowerUnitFan     | bay 5, fan 1   | 0     | 0         | NA          | 0   |
| PowerUnitFan     | bay 5, fan 2   | 0     | 0         | NA          | 0   |
| SSM              | bay 1          | 136   | 0         | Mbps        | 1   |
| SSM              | bay 2          | 128   | 0         | Mbps        | 1   |
-----------------------------------------------------------------------------
| Chassis 2                                                                 |
-----------------------------------------------------------------------------
| CMM              | bay 1          | 1     | 0         | <S,D>/<A>   | 1   |
| CMM              | bay 2          | 0     | 0         | <S,D>/<A>   | 1   |
| CPUtemp          | blade 1, CPU0  | 50    | 65        | Celsius     | 1   |
| CPUtemp          | blade 1, CPU1  | 64    | 65        | Celsius     | 1   |
| CPUtemp          | blade 2, CPU0  | 48    | 65        | Celsius     | 1   |
| CPUtemp          | blade 2, CPU1  | 64    | 65        | Celsius     | 1   |
| CPUtemp          | blade 3, CPU0  | 48    | 65        | Celsius     | 1   |
| CPUtemp          | blade 3, CPU1  | 64    | 65        | Celsius     | 1   |
| CPUtemp          | blade 4, CPU0  | 47    | 65        | Celsius     | 1   |
| CPUtemp          | blade 4, CPU1  | 74    | 65        | Celsius     | 1   |
| CPUtemp          | blade 5, CPU0  | 84    | 65        | Celsius     | 1   |
| CPUtemp          | blade 5, CPU1  | 71    | 65        | Celsius     | 1   |
| Fan                 | bay 1, fan 1   | 4      | 11          | Speed Level | 1    |
| Fan                 | bay 1, fan 2   | 4      | 11          | Speed Level | 1    |
| Fan                 | bay 2, fan 1   | 4      | 11          | Speed Level | 1    |
| Fan                 | bay 2, fan 2   | 4      | 11          | Speed Level | 1    |
| Fan                 | bay 3, fan 1   | 4      | 11          | Speed Level | 1    |
| Fan                 | bay 3, fan 2   | 4      | 11          | Speed Level | 1    |
|    .                                                                        |
|    .                                                                        |
-----------------------------------------------------------------------------
```

Error Types

This table shows some of the errors detected by `asg diag verify`.

Error Type	Error	Description
System health	`Chassis <X> error`	The Chassis quality grade is less than the defined threshold. We recommend that you correct this issue immediately.
Hardware	<*Component*> `is missing`	The component is not installed in the Chassis.
	<*Component*> `is down`	The component is installed in the Chassis, but is inactive.
Resources	<*Resource*> `capacity`	The specified resource capacity is not sufficient. You can change the defined resource capacity.
	<*Resource*> `exceed threshold`	The resource's usage is greater than the defined threshold.
CPU type	`Non compliant CPU type`	At least one SGM CPU type is not configured in the list of compliant CPUs. You can define the compliant CPU types.
Security group	<*Source*> `error`	The information collected from this source is different between the SGMs.
	<*Sources*> `differ`	The information collected from many sources is different.

Changing Compliance Thresholds

You can change some compliance thresholds that define a healthy, working system. In `$FWDIR/conf/asg_diag_config`, change the threshold values.

These are the resources you can control:

Resource	Description
`Memory`	RAM memory capacity in GB.
`HD: /`	Disk capacity in GB for <*disk*>`:/` partition.
`HD:/var/log`	Disk capacity in GB for the `/var/log` partition.
`HD: /boot`	Disk capacity in GB for the `/boot` partition.

Resource	Description
`Skew`	The maximum permissible clock difference, in seconds, between the SGMs and SSMs.
`Certified cpu`	Each line represents one compliant CPU type.

Monitoring Hardware Components (asg hw_monitor)

Use this command to show and monitor hardware information and thresholds for monitored components:

- Security Gateway Module - CPU temperature for each socket
- Chassis fan speeds
- Security Switch Module - Throughput rates
- Power consumption for each Chassis
- Power Supply Unit: If installed or not, and PSU fan speed
- Chassis Management Module - Installed, Active or Standby

Syntax

```
> asg hw_monitor [-v] [-f <filter>]
```

Parameter	Description
`-v`	Show detailed component status report (verbose)
`-f`	Show status of one or more specified (filtered) components
<*filter*>	One or more of these component types, in a comma separated list: • `CMM` • `CPUtemp` • `Fan` • `PowerConsumption` • `PowerUnit` • `SSM`

Sample Output for the NG 61000 Security System

```
> asg hw_monitor -v
------------------------------------------------------------------------------
| Hardware Monitor                                                           |
------------------------------------------------------------------------------
| Sensor           | Location        | Value | Threshold | Units         | State|
------------------------------------------------------------------------------
| Chassis 1                                                                  |
------------------------------------------------------------------------------
| CMM               | bay 1          | 1      | 0          | <S,D>/<A>     | 1    |
| CMM               | bay 2          | 0      | 0          | <S,D>/<A>     | 1    |
| CPUtemp           | blade 1, CPU0  | 45     | 65         | Celsius       | 1    |
| CPUtemp           | blade 1, CPU1  | 39     | 65         | Celsius       | 1    |
| CPUtemp           | blade 2, CPU0  | 44     | 65         | Celsius       | 1    |
| CPUtemp           | blade 2, CPU1  | 39     | 65         | Celsius       | 1    |
| CPUtemp           | blade 3, CPU0  | 44     | 65         | Celsius       | 1    |
| CPUtemp           | blade 3, CPU1  | 38     | 65         | Celsius       | 1    |
| CPUtemp          | blade 4, CPU0  | 47     | 65         | Celsius      | 1      |
| CPUtemp          | blade 4, CPU1  | 42     | 65         | Celsius      | 1      |
| CPUtemp          | blade 5, CPU0  | 0      | 65         | Celsius      | 1      |
| CPUtemp          | blade 5, CPU1  | 0      | 65         | Celsius      | 1      |
| CPUtemp          | blade 6, CPU0  | 0      | 65         | Celsius      | 0      |
| CPUtemp          | blade 6, CPU1  | 0      | 65         | Celsius      | 0      |
| CPUtemp          | blade 7, CPU0  | 0      | 65         | Celsius      | 0      |
| CPUtemp          | blade 7, CPU1  | 0      | 65         | Celsius      | 0      |
| CPUtemp          | blade 8, CPU0  | 0      | 65         | Celsius      | 0      |
| CPUtemp          | blade 8, CPU1  | 0      | 65         | Celsius      | 0      |
| CPUtemp          | blade 9, CPU0  | 0      | 65         | Celsius      | 0      |
| CPUtemp          | blade 9, CPU1  | 0      | 65         | Celsius      | 0      |
| CPUtemp          | blade 10, CPU0 | 0      | 65         | Celsius      | 0      |
| CPUtemp          | blade 10, CPU1 | 0      | 65         | Celsius      | 0      |
| CPUtemp          | blade 11, CPU0 | 0      | 65         | Celsius      | 0      |
| CPUtemp          | blade 11, CPU1 | 0      | 65         | Celsius      | 0      |
| CPUtemp          | blade 12, CPU0 | 0      | 65         | Celsius      | 0      |
| CPUtemp          | blade 12, CPU1 | 0      | 65         | Celsius      | 0      |
| Fan              | bay 1, fan 1   | 3      | 11         | Speed Level  | 1      |
| Fan              | bay 1, fan 2   | 3      | 11         | Speed Level  | 1      |
| Fan              | bay 2, fan 1   | 3      | 11         | Speed Level  | 1      |
| Fan              | bay 2, fan 2   | 3      | 11         | Speed Level  | 1      |
| Fan              | bay 3, fan 1   | 3      | 11         | Speed Level  | 1      |
| Fan              | bay 3, fan 2   | 3      | 11         | Speed Level  | 1      |
| PowerConsumption | N/A            | 2711   | 4050       | Watts        | 1      |
| PowerUnit(AC)    | bay 1          | 0      | 0          | NA           | 1      |
| PowerUnit(AC)    | bay 2          | 0      | 0          | NA           | 1      |
| PowerUnit(AC)    | bay 3          | 0      | 0          | NA           | 1      |
| PowerUnit(AC)    | bay 4          | 0      | 0          | NA           | 0      |
| PowerUnit(AC)    | bay 5          | 0      | 0          | NA           | 0      |
| PowerUnitFan     | bay 1, fan 1   | 0      | 0          | NA           | 1      |
| PowerUnitFan     | bay 1, fan 2   | 0      | 0          | NA           | 1      |
| PowerUnitFan     | bay 2, fan 1   | 0      | 0          | NA           | 1      |
| PowerUnitFan     | bay 2, fan 2   | 0      | 0          | NA           | 1      |
| PowerUnitFan     | bay 3, fan 1   | 0      | 0          | NA           | 1      |
| PowerUnitFan     | bay 3, fan 2   | 0      | 0          | NA           | 1      |
| PowerUnitFan     | bay 4, fan 1   | 0      | 0          | NA           | 0      |
| PowerUnitFan     | bay 4, fan 2   | 0      | 0          | NA           | 0      |
| PowerUnitFan     | bay 5, fan 1   | 0      | 0          | NA           | 0      |
| PowerUnitFan     | bay 5, fan 2   | 0      | 0          | NA           | 0      |
| SSM              | bay 1          | 0      | 0          | Mbps         | 1      |
| SSM              | bay 2          | 0      | 0          | Mbps         | 1      |
---------------------------------------------------------------------------------
| Chassis 2                                                                     |
---------------------------------------------------------------------------------
| CMM              | bay 1          | 1      | 0          | <S,D>/<A>    | 1      |
| CMM              | bay 2          | 0      | 0          | <S,D>/<A>    | 1      |
| CPUtemp          | blade 1, CPU0  | 46     | 65         | Celsius      | 1      |
| CPUtemp          | blade 1, CPU1  | 46     | 65         | Celsius      | 1      |
| CPUtemp          | blade 2, CPU0  | 48     | 65         | Celsius      | 1      |
| CPUtemp          | blade 2, CPU1  | 49     | 65         | Celsius      | 1      |
| CPUtemp          | blade 3, CPU0  | 46     | 65         | Celsius      | 1      |
| CPUtemp          | blade 3, CPU1  | 47     | 65         | Celsius      | 1      |
| CPUtemp          | blade 4, CPU0  | 46     | 65         | Celsius      | 1      |
| CPUtemp          | blade 4, CPU1  | 50     | 65         | Celsius      | 1      |
| CPUtemp          | blade 5, CPU0  |        | 65         | Celsius      | 1      |
| CPUtemp          | blade 5, CPU1  |        | 65         | Celsius      | 1      |
| CPUtemp          | blade 6, CPU0  | 0      | 65         | Celsius      | 0      |
| CPUtemp          | blade 6, CPU1  | 0      | 65         | Celsius      | 0      |
| CPUtemp          | blade 7, CPU0  | 0      | 65         | Celsius      | 0      |
| CPUtemp          | blade 7, CPU1  | 0      | 65         | Celsius      | 0      |
| CPUtemp          | blade 8, CPU0  | 0      | 65         | Celsius      | 0      |
| CPUtemp          | blade 8, CPU1  | 0      | 65         | Celsius      | 0      |
| CPUtemp          | blade 9, CPU0  | 0      | 65         | Celsius      | 0      |
| CPUtemp          | blade 9, CPU1  | 0     | 65        | Celsius     | 0    |
| CPUtemp          | blade 10, CPU0 | 0     | 65        | Celsius     | 0    |
| CPUtemp          | blade 10, CPU1 | 0     | 65        | Celsius     | 0    |
| CPUtemp          | blade 11, CPU0 | 0     | 65        | Celsius     | 0    |
| CPUtemp          | blade 11, CPU1 | 0     | 65        | Celsius     | 0    |
| CPUtemp          | blade 12, CPU0 | 0     | 65        | Celsius     | 0    |
| CPUtemp          | blade 12, CPU1 | 0     | 65        | Celsius     | 0    |
| Fan              | bay 1, fan 1   | 5     | 11        | Speed Level | 1    |
| Fan              | bay 1, fan 2   | 5     | 11        | Speed Level | 1    |
| Fan              | bay 2, fan 1   | 5     | 11        | Speed Level | 1    |
| Fan              | bay 2, fan 2   | 5     | 11        | Speed Level | 1    |
| Fan              | bay 3, fan 1   | 5     | 11        | Speed Level | 1    |
| Fan              | bay 3, fan 2   | 5     | 11        | Speed Level | 1    |
| PowerConsumption | N/A            | 2711  | 4050      | Watts       | 1    |
| PowerUnit(AC)    | bay 1          | 0     | 0         | NA          | 1    |
| PowerUnit(AC)    | bay 2          | 0     | 0         | NA          | 1    |
| PowerUnit(AC)    | bay 3          | 0     | 0         | NA          | 1    |
| PowerUnit(AC)    | bay 4          | 0     | 0         | NA          | 0    |
| PowerUnit(AC)    | bay 5          | 0     | 0         | NA          | 0    |
| PowerUnitFan     | bay 1, fan 1   | 0     | 0         | NA          | 1    |
| PowerUnitFan     | bay 1, fan 2   | 0     | 0         | NA          | 1    |
| PowerUnitFan     | bay 2, fan 1   | 0     | 0         | NA          | 1    |
| PowerUnitFan     | bay 2, fan 2   | 0     | 0         | NA          | 1    |
| PowerUnitFan     | bay 3, fan 1   | 0     | 0         | NA          | 1    |
| PowerUnitFan     | bay 3, fan 2   | 0     | 0         | NA          | 1    |
| PowerUnitFan     | bay 4, fan 1   | 0     | 0         | NA          | 0    |
| PowerUnitFan     | bay 4, fan 2   | 0     | 0         | NA          | 0    |
| PowerUnitFan     | bay 5, fan 1   | 0     | 0         | NA          | 0    |
| PowerUnitFan     | bay 5, fan 2   | 0     | 0         | NA          | 0    |
| SSM              | bay 1          | 0     | 0         | Mbps        | 1    |
| SSM              | bay 2          | 0     | 0         | Mbps        | 1    |
------------------------------------------------------------------------------
```

Sample Output for a 41000 Security System

```
------------------------------------------------------------------------------
| Hardware Monitor                                                           |
------------------------------------------------------------------------------
| Sensor           | Location       | Value | Threshold | Units       | State|
------------------------------------------------------------------------------
| Chassis 1                                                                  |
------------------------------------------------------------------------------
| CMM              | bay 1          | 0     | 0         | <S,D>/<A>   | 1    |
| CMM              | bay 2          | 1     | 0         | <S,D>/<A>   | 1    |
| CPUtemp          | blade 1, CPU0  | 47    | 65        | Celsius     | 1    |
| CPUtemp          | blade 1, CPU1  | 46    | 65        | Celsius     | 1    |
| CPUtemp          | blade 2, CPU0  | 46    | 65        | Celsius     | 1    |
| CPUtemp          | blade 2, CPU1  | 44    | 65        | Celsius     | 1    |
| CPUtemp          | blade 3, CPU0  | 46    | 65        | Celsius     | 1    |
| CPUtemp          | blade 3, CPU1  | 45    | 65        | Celsius     | 1    |
| CPUtemp          | blade 4, CPU0  | 45    | 65        | Celsius     | 1    |
| CPUtemp          | blade 4, CPU1  | 46    | 65        | Celsius     | 1    |
| Fan              | bay 1, fan 1   | 4     | 11        | Speed Level | 1    |
| Fan              | bay 1, fan 2   | 4     | 11        | Speed Level | 1    |
| Fan              | bay 1, fan 3   | 4     | 11        | Speed Level | 1    |
| Fan              | bay 1, fan 4   | 4     | 11        | Speed Level | 1    |
| Fan              | bay 1, fan 5   | 4     | 11        | Speed Level | 1    |
| Fan              | bay 1, fan 6   | 4     | 11        | Speed Level | 1    |
| Fan              | bay 1, fan 7   | 4     | 11        | Speed Level | 1    |
| Fan              | bay 1, fan 8   | 4     | 11        | Speed Level | 1    |
| Fan              | bay 1, fan 9   | 4     | 11        | Speed Level | 1    |
| Fan              | bay 1, fan 10  | 4     | 11        | Speed Level | 1    |
| Fan              | bay 2, fan 1   | 4     | 11        | Speed Level | 1    |
| Fan              | bay 2, fan 2   | 4     | 11        | Speed Level | 1    |
| Fan              | bay 2, fan 3   | 4     | 11        | Speed Level | 1    |
| Fan              | bay 2, fan 4   | 4     | 11        | Speed Level | 1    |
| Fan              | bay 2, fan 5   | 4     | 11        | Speed Level | 1    |
| Fan              | bay 2, fan 6   | 4     | 11        | Speed Level | 1    |
| Fan              | bay 2, fan 7   | 4     | 11        | Speed Level | 1    |
| Fan              | bay 2, fan 8   | 4     | 11        | Speed Level | 1    |
| Fan              | bay 2, fan 9   | 4     | 11        | Speed Level | 1    |
| Fan              | bay 2, fan 10  | 4     | 11        | Speed Level | 1    |
| PowerConsumption | N/A            | 1894  | 4050      | Watts       | 1    |
| PowerUnit(AC)    | bay 1          | 0     | 0         | NA          | 1    |
| PowerUnit(AC)    | bay 2          | 0     | 0         | NA          | 1    |
| PowerUnit(AC)    | bay 3          | 0     | 0         | NA          | 1    |
| PowerUnitFan     | bay 1, fan 1   | 0     | 0         | NA          | 1    |
| PowerUnitFan     | bay 1, fan 2   | 0     | 0         | NA          | 1    |
| PowerUnitFan     | bay 2, fan 1   | 0     | 0         | NA          | 1    |
| PowerUnitFan     | bay 2, fan 2   | 0     | 0         | NA          | 1    |
| PowerUnitFan     | bay 3, fan 1   | 0     | 0         | NA          | 1    |
| PowerUnitFan     | bay 3, fan 2   | 0     | 0         | NA          | 1    |
| SSM              | bay 1          | 40    | 0         | Mbps        | 1    |
| SSM              | bay 2          | 0     | 0         | Mbps        | 1    |
------------------------------------------------------------------------------
| Chassis 2                                                                  |
------------------------------------------------------------------------------
| CMM              | bay 1          | 1     | 0         | <S,D>/<A>   | 1    |
| CMM              | bay 2          | 0     | 0         | <S,D>/<A>   | 1    |
| CPUtemp          | blade 1, CPU0  | 47    | 65        | Celsius     | 0    |
| CPUtemp          | blade 1, CPU1  | 51    | 65        | Celsius     | 0    |
| CPUtemp          | blade 2, CPU0  | 46    | 65        | Celsius     | 1    |
| CPUtemp          | blade 2, CPU1  | 56    | 65        | Celsius     | 1    |
| CPUtemp          | blade 3, CPU0  | 49    | 65        | Celsius     | 1    |
| CPUtemp          | blade 3, CPU1  | 51    | 65        | Celsius     | 1    |
| CPUtemp          | blade 4, CPU0  | 0     | 65        | Celsius     | 0    |
| CPUtemp          | blade 4, CPU1  | 0     | 65        | Celsius     | 0    |
| Fan              | bay 1, fan 1   | 3     | 11        | Speed Level | 1    |
| Fan              | bay 1, fan 2   | 3     | 11        | Speed Level | 1    |
| Fan              | bay 1, fan 3   | 3     | 11        | Speed Level | 1    |
| Fan              | bay 1, fan 4   | 3     | 11        | Speed Level | 1    |
| Fan              | bay 1, fan 5   | 3     | 11        | Speed Level | 1    |
| Fan              | bay 1, fan 6   | 3     | 11        | Speed Level | 1    |
| Fan              | bay 1, fan 7   | 3     | 11        | Speed Level | 1    |
| Fan              | bay 1, fan 8   | 3     | 11        | Speed Level | 1    |
| Fan              | bay 1, fan 9   | 3     | 11        | Speed Level | 1    |
| Fan              | bay 1, fan 10  | 3     | 11        | Speed Level | 1    |
| Fan              | bay 2, fan 1   | 3     | 11        | Speed Level | 1    |
| Fan              | bay 2, fan 2   | 3     | 11        | Speed Level | 1    |
| Fan              | bay 2, fan 3   | 3     | 11        | Speed Level | 1    |
| Fan              | bay 2, fan 4   | 3     | 11        | Speed Level | 1    |
| Fan              | bay 2, fan 5   | 3     | 11        | Speed Level | 1    |
| Fan              | bay 2, fan 6   | 3     | 11        | Speed Level | 1    |
| Fan              | bay 2, fan 7   | 3     | 11        | Speed Level | 1    |
| Fan              | bay 2, fan 8   | 3     | 11        | Speed Level | 1    |
| Fan              | bay 2, fan 9   | 3     | 11        | Speed Level | 1    |
| Fan              | bay 2, fan 10  | 3     | 11        | Speed Level | 1    |
| PowerConsumption | N/A            | 1624  | 4050      | Watts       | 1    |
| PowerUnit(AC)    | bay 1          | 0     | 0         | NA          | 1    |
| PowerUnit(AC)    | bay 2          | 0     | 0         | NA          | 1    |
| PowerUnit(AC)    | bay 3          | 0     | 0         | NA          | 0    |
| PowerUnitFan     | bay 1, fan 1   | 0     | 0         | NA          | 1    |
| PowerUnitFan     | bay 1, fan 2   | 0     | 0         | NA          | 1    |
| PowerUnitFan     | bay 2, fan 1   | 0     | 0         | NA          | 1    |
| PowerUnitFan     | bay 2, fan 2   | 0     | 0         | NA          | 1    |
| PowerUnitFan     | bay 3, fan 1   | 0     | 0         | NA          | 0    |
| PowerUnitFan     | bay 3, fan 2   | 0     | 0         | NA          | 0    |
| SSM              | bay 1          | 2     | 0         | Mbps        | 1    |
| SSM                 | bay 2            | 0       | 0            | Mbps          | 1     |
-------------------------------------------------------------------------------------------
```

Column	Meaning
Location	To identify the location, see the *61000/41000 Security System Front Panel*.
Value **Threshold** **Units**	Most components have a defined threshold value. The threshold gives an indication of the health and functionality of the component. When the value of the resource is greater than the threshold, an alert is sent ("Configuring Alerts for SGM and Chassis Events (asg alert)" on page 135).
State	**0** = Component not installed **1** = Component is installed

Chassis Control (asg_chassis_ctrl)

The Chassis Control utility lets you monitor and configure SSMs and CMMs with many different command options and parameters. Chassis Control is based on SNMP communications between the different Chassis and components.

Note - You can also configure SGMs with this utility, but we recommend that you use: `asg dxl`.

Syntax

```
> asg_chassis_ctrl <option> <parameters>
```

Options and Parameters	Description
`active_sgms`	Shows all installed SGMs.
`active_ssm`	Shows active SSMs. An SSM that is not installed or is down does not show as ACTIVE.
`get_fans_status`	Shows the health status of the Chassis fans.
`get_lb_dist` *<ssm_id>*	Shows the current distribution matrix from the specified SSM. The matrix is a table containing SGM IDs, and used to determine to which other SGMs a packet should be forwarded.
`get_ssm_firmware` *<ssm_id>*	Shows the firmware version of the specified SSM.
`get_ssm_config` *<ssm_id>*	Shows the configuration name of the specified SSM.
`get_ssm_type` *<ssm_id>*	Shows the model of the specified SSM
`get_psu_status`	Shows the current status of the PSUs.
`get_pems_status`	Shows the current status of the Chassis PEMs.
`get_cmm_status`	Shows the current status of the CMMs.
`get_cpus_temp` *<sgm_id>*	Shows temperatures of the specified SGM CPUs.

Options and Parameters	Description
`get_dist_md5sum`	Shows the md5sum of the distribution matrix for the given SSM. Comparing this checksum against the checksum on other SSM verifies that they are synchronized.
`get_ports_stat` <*ssm_id*>	Prints the port status for the specified SSM.
`get_dist_mode` <*ssm_id*>	Shows the port distribution mode for the specified SSM.
`get_dist_mask` <*ssm_id*>	Shows a summary of the distribution masks in the different modes.
`get_matrix_size` <*ssm_id*>	Shows the size, in bytes, of the SSM distribution matrix.
`get_sel_info` <*cmm_id*>	Shows data from the specified CMM event. This information is useful for troubleshooting and system forensics.
`restart_ssm` <*ssm_id*>	Restarts the specified SSM.
`restart_cmm` <*cmm_id*>	Restart the specified CMM.
`start_ssm` <*ssm_id*>	Starts the specified SSM.
`shutdown_ssm` <*smm_id*>	Shuts down the specified SSM.
`mib2_stats` <*ssm_id*> <*port_id*> `[`<*err*>`]`	Shows MIB2 statistics for the specified SSM and port. <*err*> = Error type.
`get_bmac` <*ssm_id*>	Shows SGM MAC addresses from the SSM.
`get_power_type`	Shows the Chassis input power type (AC or DC).
`get_ac_power_type`	Shows the AC power type.
`jumbo_frames enable|disable|show` <*SSM* `ID`>	Enable, disable or show Jumbo Frames on an SSM160.
`set_port_mtu` <*ssm_id*> <*port_id*> <*mtu_size*>	Sets the port MTU size for the specified SSM and Port. <*ssm_id*> - SSM identifier (1-4 or all) <*port_id*> - Port number <*mtu_size*> - This MTU size can be one of these values: • Integer value up to 12,288 • max - Maximum supported MTU size • default - System default MTU size (typically 1544)
`get_port_mtu` <*ssm_id*> <*port_id*>	Shows the MTU for the specified SSM and port.

Options and Parameters	Description
`get_port_media_details` *<ssm_id>*	Shows port information.
`get_pem_cb_status`	Shows PEM status.
`help [-v]`	Shows help messages in [-v] verbose mode.

Notes

- To see the full syntax for an option, run the command and option without parameters.
- To make sure that the Chassis Control commands work correctly, run this command on both Chassis:

```
> asg_chassis_ctrl get_cmm_status

Getting CMM(s) status
CMM #1 -> Health: 1,   Active: 1
CMM #2 -> Health: 1,   Active: 0
Active CMM firmware version: 2.83
```

Monitoring CPU Utilization (asg_cores_util)

Use this command to monitor CPU utilization on all SGMs.

Syntax

```
# asg_cores_util
```

Output

```
+---------------------+
|CPUs Utilization     |
+-----------+----+----+
|CPU \ Blade|2_3 |2_4 |
+-----------+----+----+
|cpu0       |29% |2%  |
+-----------+----+----+
|cpu1       |0%  |0%  |
+-----------+----+----+
|cpu2       |0%  |1%  |
+-----------+----+----+
|cpu3       |37% |25% |
+-----------+----+----+
|cpu4       |0%  |0%  |
+-----------+----+----+
|cpu5       |1%  |18% |
+-----------+----+----+
|cpu6       |0%  |0%  |
+-----------+----+----+
|cpu7       |0%  |0%  |
+-----------+----+----+
|cpu8       |0%  |0%  |
+-----------+----+----+
|cpu9       |0%  |1%  |
+-----------+----+----+
|cpu10      |0%  |0%  |
+-----------+----+----+
|cpu11      |0%  |0%  |
+-----------+----+----+
|cpu12      |0%  |0%  |
+-----------+----+----+
|cpu13      |1%  |1%  |
+-----------+----+----+
|cpu14      |1%  |0%  |
+-----------+----+----+
|cpu15      |0%  |0%  |
+-----------+----+----+
```

Security Monitoring

SYN Defender (sim synatk, sim6 synatk, asg synatk)

A SYN flood attack occurs when a host, typically with a forged address, sends a flood of TCP/SYN packets. Each of these packets is handled as a connection request, which causes the server to create a "half-open connection". This occurs because the gateway sends a TCP/SYN-ACK (Acknowledge) packet, and waits for a response packet, which does not arrive. These half-open connections eventually exceed the maximum available connections, which causes a denial of service condition. SYN defender protects the gateway by dropping excessive half-open connections.

You can use these commands to:

- Configure a defense against an IPv4 SYN Flood attack (`sim synatk`).
- Configure a defense against an IPv6 SYN Flood attack (`sim6 synatk`).
- Monitor the system during attacks and normal system operation (`asg synatk`).
- Simulate a SYN attack on the specified interfaces (`asg synatk state -i` <*interface_name*> `-a`)

This protection works with Performance Pack.

Syntax

```
> sim synatk [-e] [-d] [-m] [-g] [-t <threshold>] [-a] [monitor] [monitor -v]
> sim6 synatk [-e] [-d] [-m] [-g] [-t <threshold>] [-a] [monitor] [monitor -v]
> sim synatk state -i <interface_name> -a
> asg synatk [-b <sgm_ids>] [-4 | -6]
> sim6 synatk -a
```

Parameter	Description
`-e`	Enable SYN defender. This make the system engage when it recognizes an attack on an external interface. External interfaces are defined in SmartDashboard. Internal interfaces are always in monitor mode.
`-d`	Disable SYN Defender.
`-m`	Set monitor mode. SYN defender only sends a log when it recognizes an attack.
`-g`	Enforce on all interfaces.

Parameter	Description
`-t` <*threshold*>	Set the SYN Defender threshold number of half-opened connections.
`-i state` <*interface_name*>	Simulate a SYN attack on the specified interface.
`-a`	Use configuration from: `$PPKDIR/conf/synatk.conf`
`monitor`	Show the attack monitoring tool.
`monitor -v`	Show the attack monitoring tool with extra (verbose) information.
`-b` <*sgm_ids*>	Show the status for specified SGMs and Chassis. Works with SGMs and/or Chassis as specified by <*sgm_ids*>. <*sgm_ids*> can be: • No <*sgm_ids*> specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`)
`-6`	Shows the IPv6 status only.
`-4`	Shows the IPv4 status only.

SYN Defender Configuration File

The Syn Defender configuration file (default `$PPKDIR/conf/synatk.conf)` has two sections:

- Configuration fields
- Interface list

The configuration fields section consists of single lines with a field, an equal sign, and the value.

Field	Description	Default
`enabled`	• 0 - Enable SYN Defender • 1 - Disable SYN Defender	1
`enforce`	• 0 - Interfaces use monitor mode only • 1 - Enforce rules on external interfaces only • 2 - Enforce rules on internal and external interfaces	1
`global_high_threshold`	Maximum number of unestablished connections	10,000
`periodic_updates`	• 0 - Enable periodic updates of hit counters for rule enforcement • 1 - Disable periodic updates of hit counters for rule enforcement	1
`cookie_lifetime`	Maximum cookie lifetime in seconds.	10
`total_max_held_pkts`	Maximum number of cached packets. -1 means no limit.	-1

Field	Description	Default
`min_frag_sz`	Minimum size of packets that are not dropped during an attack	80
`nr_saved_pkt_on_activate`	Maximum number of packets saved to syslog when an attack starts	100
`high_threshold`	Maximum number of unestablished connections per external interface	10,000
`low_threshold`	Minimum number of unestablished connections per external interface before connections are dropped	5000
`internal_high_threshold`	Maximum number of unestablished connections per internal interface	20000
`internal_low_threshold`	Minimum number of unestablished connections per internal interface before connections are dropped	10,00
`score_alpha`	Number between 1 and 127 that represents how likely Syn Defender is to drop packets. 1 is least likely, 127 is most likely.	100
`conn_max_held_pkts`	Maximum number of held packets for a connection from before Syn Defender engages.	1
`monitor_log_interval`	Number of milliseconds between log warnings.	60,000
`grace_timeout`	Maximum number of milliseconds Syn Defender stays in grace mode.	30,000
`min_time_in_active`	Minimum number of milliseconds Syn Defender stays in active mode.	60,000
`clear_route_cache_on_activate`	• 1 - Clear the route cache when SYN Defender activates • 0 - Do not clear the route cache when SYN Defender activates	1
`revalidate_suspicious_syns`	Delete a connection and send a validation SYN+ACK packet back. This is useful to clean up spoofed connections made before SYN Defender engaged.	1

Example:

```
enabled = 1
enforce = 1
```

The interface section consists of lines in this format:

`interface <`*if_name*`> state = <`*state*`>`

Field	Description
<*if_name*>	Interface name
<*state*>	• `disabled` - Syn Defender does not protect or monitor the interface • `monitor`- Syn Defender monitors but does not protect the interface • `enforce` - Syn Defender protects the interface

Example:

```
interface eth1-01 state = enforce
interface eth2-01 state = disabled
```

Monitoring a Syn Attack - Standard Output

This example shows that there are two interfaces under attack. Interface eth2-03 was attacked 3 seconds ago and eth2-04 is recovering from an attack that ended 24 seconds ago.

```
> sim synatk monitor -b all -4
+-----------------------------------------------------------------------------+
| SYN Defender status                                                         |
+-----------------------------------------------------------------------------+
| Configuration                                                     Enforcing |
| Status                                                     Under Attack (!) |
| Non established connections                                               3 |
| Threshold                                                              1000 |
+-----------------------------------------------------------------------------+
| IF              | Topology | Enforce | State (sec)   | Non-established conns |
|                 |          |         |               | Peak      | Current   |
+-----------------------------------------------------------------------------+
| eth1-Mgmt4      | External | Prevent | Monitor       | 7         | 3         |
| eth1-01         | Internal | Detect  | Monitor       | 0         | 0         |
| eth2-01         | External | Prevent | Monitor       | 0         | 0         |
| eth2-02         | External | Prevent | Monitor       | 0         | 0         |
| eth2-03     (!) | External | Prevent | Active(   3)  | -         | -         |
| eth2-04     (!) | External | Prevent | Grace (  24)  | 0         | 0         |
+-----------------------------------------------------------------------------+
```

Output information

Column	Description
IF	Interface name
Topology	Topology as defined in SmartDashboard
Enforce	Action taken by SYN Defender: **Prevent** - Detects attacks and enforces protection **Detect** - Detects attacks, but only generates log entries. Does not enforce protection **Disabled** - Protection is disabled

Column	Description
State	Current Syn Defender state: **Disabled** - Syn Defender is disabled for this interface **Monitor** - The interface is not under attack and Syn Defender monitors connections. **Active** - The interface is under attack and Syn Defender enforces protections **Grace** - The attack on the interface ended and the normal service is restored.
non-established conns	**Peak** - The highest number of half-opened connections for this interface This can help you to configure the correct threshold. **Current** - The number of half-opened connections at this time

Monitoring a SYN Attack - Verbose Output

This example shows the verbose output.

```
> sim synatk monitor –v
+------------------------------------------------------------------------------+
| SYN Defender statistics                                                      |
+------------------------------------------------------------------------------+
| Status                                                      Under Attack (!) |
| Spoofed SYN/sec                                                       534000 |
+------------------------------------------------------------------------------+
| IF           | Topology | Defend (sec) | SYN cookie rate                     |
|              |          |              | Sent        | BAU (cps)  | Spoofed  |
+------------------------------------------------------------------------------+
| eth2-01      | External | 28           | 345345      | 40         | 95 %     |
| eth2-02      | External | 12           | 150         | 50         | 33 %     |
+------------------------------------------------------------------------------+
| Sum                                    | 345495      | 90         | 93 %     |
+------------------------------------------------------------------------------+
```

Column	Description
IF	The interface name
Topology	The interface topology as defined in SmartDashboard.
Defend	The attack duration in seconds.
Sent SYN cookie rate	Number of SYN packets received per second.
BAU	Business as usual. The number of legitimate connections handled per second.
Spoofed	The percentage of spoofed SYN packets out of all traffic.

Showing Syn Defender Status

This example shows the status of SYN Flood attack protection for all SGMs. It shows that Blade 2-02:

- There are 3 half-open connections
- Receives 10,000 spoofed syn packets per second
- Is under attack

```
> asg synatk
+-------------------------------------------------------------------------------+
| SYN Defender status                                                           |
+-------------------------------------------------------------------------------+
| Blade    | IP   | Config     | Status       | Non est. conns | Spoofed / sec |
+-------------------------------------------------------------------------------+
| 2_01     | IPv4 | Enforcing  | Normal       | 6              | 0             |
| 2_01     | IPv6 | Enforcing  | Normal       | 0              | 0             |
| 2_02     | IPv4 | Enforcing  | Normal       | 0              | 0             |
| 2_02     | IPv6 | Enforcing  | Under Attack | 3              | 10000         |
+-------------------------------------------------------------------------------+
| All                                         | 9              | 10000         |
+-------------------------------------------------------------------------------+
```

F2F Quota

- `asg f2fq`
- `fwaccel f2fq stats`
- `fwaccel6 f2fq stats`

F2F detects traffic floods and intelligently prevents performance degradation on the 61000/41000 Security System. It assigns a high priority to known, important packets from Performance Pack and drops those suspected of being part of a DDoS attack.

Two examples of known F2F flood attacks are UDP floods and fragmentation attacks. These attacks cause too much resource allocation when they try to put the packet fragments together.

Use `fwaccel` for IPv4 information and `fwaccel6` for IPv6 information.

Syntax

```
> fwaccel f2fq stats [–v]
> fwaccel f2fq -c <file>
> fwaccel f2fq -a
> fwaccel6 f2fq stats [–v]
> fwaccel6 f2fq -c <file>
> fwaccel6 f2fq -a
> asg f2fq [-b <sgm_ids> ] [-6 | -4]
```

Parameter	Description
`-v`	Shows detailed (verbose) statistics.
`-b` *<sgm_ids>*	Works with SGMs and/or Chassis as specified by *<sgm_ids>*. *<sgm_ids>* can be: • No *<sgm_ids>* specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`)

Parameter	Description
`-6`	Shows the IPv6 status only
`-4`	Shows the IPv4 status only
`-c` *<file>*	Uses the parameters in *<file>*
`-a`	Uses the parameters in `$FWDIR/conf/f2fq.conf`

Example

This example shows details of IPv4 activity for all Firewall instances.

```
> fwaccel f2fq stats -v
+----------------------------------------------------------------------------+
| DDOS Mitigation                                                            |
+----------------------------------------------------------------------------+
| Mode:                                                            Enforcing |
| Status                                                              Normal |
| Last 10 seconds drops                                                13146 |
+----------------------------------------------------------------------------+
| Instance | Reason                     | Drops / Hits                       |
+----------------------------------------------------------------------------+
| FW  0    | CONN_MISS_TCP_SYN          |           103365 / 104629          |
+----------------------------------------------------------------------------+
| FW  1    | FRAG                       |             6232 / 13816           |
|          | CONN_MISS_TCP_SYN          |           101096 / 102203          |
|          | CONN_MISS_TCP_OTHER        |            13146 / 14359           |
+----------------------------------------------------------------------------+
| FW  2    | FRAG                       |             1339 / 1339            |
|          | CONN_MISS_TCP_SYN          |           101087 / 102143          |
+----------------------------------------------------------------------------+
| All      | FRAG                       |             7571 / 15155           |
|          | CONN_MISS_TCP_SYN          |           305548 / 308975          |
|          | CONN_MISS_TCP_OTHER        |            13146 / 14359           |
+----------------------------------------------------------------------------+
```

The output shows this information:

Item	Description
Last 10 seconds drops	The number of dropped packets during the last 10 seconds.
Instance	The verbose output shows a historical aggregate of the results, for each Firewall instance.
Drops / Hits	The number of dropped packets out of the total number of packets, grouped by the attack type.

Example - asg f2fq

This output shows how the protection mitigates the DDoS attack, for each SGM.

```
> asg f2fq
+--------------------------------------------------------------------------+
| DDOS Mitigation                                                          |
+--------------------------------------------------------------------------+
| Blade    | Protocol | Config     | Status        | Last 10 sec drops     |
+--------------------------------------------------------------------------+
| 1_01 (!) | IPv4     | Enforcing  | Under Attack  | 151130                |
| 1_01     | IPv6     | Enforcing  | Normal        | 0                     |
| 1_02      | IPv4      | Enforcing  | Normal        | 0                    |
| 1_02      | IPv6      | Enforcing  | Normal        | 0                    |
| 1_03      | IPv4      | Enforcing  | Normal        | 0                    |
| 1_03      | IPv6      | Enforcing  | Normal        | 0                    |
| 1_04      | IPv4      | Enforcing  | Normal        | 0                    |
| 1_04      | IPv6      | Enforcing  | Normal        | 0                    |
+--------------------------------------------------------------------------+
```

F2F Configuration File

The F2F Configuration file (default `$FWDIR/conf/f2fq.conf`) has two sections:

- Global Options
- Packet Priority Table

The Global Options section has these options:

Option	Description	Default
`enabled`	• `1`: F2F Quota is enabled • `0`: F2f Quota is disabled	1
`enforce`	• `1`: Drop packets • `0`: Do not drop packets, log in `/var/log/messages`	1
`snapshots_interval`	Milliseconds between F2F calculations	1000
`load_threshold`	Percent capacity used of the queue load before F2F activates Range: 0 -100	80
`dynamic_prio_threshold`	Dynamic priority threshold F2F drops packets whose dynamic priorities are lower than `dynamic_prio_threshold`.	20
`print_syslog_interval`	Milliseconds between writes to `/var/log/messages`	30,000
`config_version`	Configuration file version	1
`default_priority`	Priority for a packet that does not match any rule	100

The Packet Priority table has these fields:

Field	Description
`# Interface`	The interface name. Use * for all interfaces.
`proto`	The transport layer protocol. Use * for all protocols.
`service`	Port number or port range (applicable to TCP and UDP only). Use * for all ports.
`ip`	The destination IP and subnet. Use * for all IPs.
`reason`	Reason why this packet is rejected. Use * for all reasons.

Field	Description
priority	• 0-100 - Priority for a packet that matches this rule. Packets with a higher priority have a lower chance of being dropped. • Exception - Packets that match this rule are never dropped.

Example

```
enabled = 1
enforce = 1
config_version = 1
default_priority = 100
dynamic_prio_threshold = 20
snapshots_interval = 1000
load_threshold = 80

# Interface   proto       service      ip              reason    priority
eth1-01       *           1-1024       1.1.1.0/24      *         Exception
*             TCP         *            *               FRAG      10
*             UDP         *            *               FRAG      60
```

F2F Rejection Reasons

Name	Description
FRAG	Packet is a fragment
IP_OPT	Packet has IP options
CONN_MISS_ICMP	No connection found for an ICMP packet
CONN_MISS_TCP_SYN	No connection found for a TCP syn packet
CONN_MISS_TCP_OTHER	No connection found for a TCP non-syn packet
CONN_MISS_UDP	No connection found for a UDP packet
CONN_MISS_OTHER	No connection not found for a packet of any other type
VPN_F2F	VPN connection
F2F_IS_ON_ICMP	ICMP packet set by the firewall to be rejected
F2F_IS_ON_TCP	TCP packet set by the firewall to be rejected
F2F_IS_ON_UDP	UDP packet set by the firewall to be rejected
F2F_IS_ON_OTHER	Other type of packet set by the firewall to be rejected
UNIDIR_VIOL	Unidirectional violation
SPOOF_VIOL	Possible spoof violation
TCP_STATE	Possible TCP state violation
OUT_IF	Outbound Interface is not defined or accelerated
XMT_EQ_RCV	Incoming interface is the same as the outgoing interface
ROUTING_ERR	Routing decision error

Name	Description
SANITY_CHECKS	Sanity checks failed
TEMP_CONN	Temporary connection expired
FWD_NON_PIVOT	Device cannot forward to non-pivot member
BROADCAST	Broadcast / multicast in pivot member
CLUSTER_MSG	Source address is of FWHA protocol or LS forwarding layer
PARTIAL_CONN	Partial connection
PXL_F2F	PXL connection
CLUSTER_FORWARD	Packet forwarded from another cluster member
CHAIN_FORWARD	Packet reinjection by the chain forwarding mechanism
SPORT_ALLOC_F2F	Packet rejected due to port allocation failure
GENERAL	Packet rejected for a reason not listed above

Showing the Number of Firewall and SecureXL Connections (asg_conns)

Use this command to show the number of firewall and SecureXL connections on each SGM.

Syntax

```
> asg_conns [-b <sgm_ids>]
```

Parameter	Description
<sgm_ids>	Works with SGMs and/or Chassis as specified by *<sgm_ids>*. *<sgm_ids>* can be: • No *<sgm_ids>* specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`)
-6	Show only IPv6 connections
-h	Show syntax and help information

Example

```
> asg_conns
1_01:
     #VALS     #PEAK    #SLINKS
       246      1143        246
```

```
1_02:
     #VALS      #PEAK    #SLINKS
        45        172         45
1_03:
     #VALS      #PEAK    #SLINKS
        45        212         45
1_04:
     #VALS      #PEAK    #SLINKS
       223        624        223
1_05:
     #VALS      #PEAK    #SLINKS
        45        246         45

Total (fw1 connections table): 604 connections

1_01:
There are 60 conn entries in SecureXL connections table
Total conn entries @ DB 0:  4
Total conn entries @ DB 3:  2
.
.
Total conn entries @ DB 26:  4
Total conn entries @ DB 30:  2
1_02:
There are 16 conn entries in SecureXL connections table
Total conn entries @ DB 0:  2
Total conn entries @ DB 1:  2
.
.
Total conn entries @ DB 26:  2
1_03:
There are 16 conn entries in SecureXL connections table
Total conn entries @ DB 0:  2
Total conn entries @ DB 5:  2
.
.
Total conn entries @ DB 30:  2
1_04:
There are 260 conn entries in SecureXL connections table
Total conn entries @ DB 0:  10
Total conn entries @ DB 1:  6
.
.
Total conn entries @ DB 31:  94
1_05:
There are 16 conn entries in SecureXL connections table
Total conn entries @ DB 2:  2
.
.
Total conn entries @ DB 26:  2

Total (SecureXL connections table): 368 connections
```

Packet Drop Monitoring

Use `asg_drop_monitor` in Expert mode to monitor dropped packets in real time. Drop statistics come from these modules:

- NICs
- Operating system

- CoreXL
- PSL
- Performance Pack

This command opens a monitor session and shows aggregated data from SGMs and, optionally, SSMs. To stop an open session, press **Ctl-c**.

Syntax

```
# asg_drop_monitor [-r -6 [-ssm -t <timeout>]]
# asg_drop_monitor -h
```

Parameter	Description
`-r`	Reset statistics to 0
`-ssm`	Include dropped packets from SSMs
`-t`	Maximum number of seconds to wait to report dropped packets Use with `-ssm`
`-6`	Show only IPv6 results
`-h`	Show command syntax and help information

Output

```
NICs drops (Rx):
0
IP Stack qdisc drops (Tx):
0
CoreXL queue drops (F2F):
0
CoreXL queue drops (PXL F2P)
0
PSL drops(total):
0
PSL drops(udp):
0
PSL rejects:
0
Ppak drops:

Displaying aggregated data from blades: all
Reason                Value              Reason              Value
------------------------------------------------------------------
general reason          0               PXL decision            0
fragment error          0               hl - spoof viol         0
F2F not allowed         0               hl - TCP viol           0
corrupted packet        0               hl - new conn           0
clr pkt on vpn          0               partial conn            0
encrypt failed          0               drop template           0
decrypt failed          0               outb - no conn          9
interface down          0               cluster error           0
XMT error               0               template quota          0
anti spoofing           0               Attack mitigation       0
local spoofing          0               sanity error            0
monitored spoofed       0               Conns limit. Exceed     0
Conns limit. Add fail   0
```

Monitoring System Status

Showing System Serial Numbers

These commands show and save serial numbers for 61000/41000 Security System hardware components:

- `asg_sgm_serial` - Shows serial numbers for SGMs in the UP state that belong to the security group only.
- `asg_serial_info` - Shows CMM, SSM and Chassis serial numbers.

The information is saved in the `gasginfo` archive file.

Syntax

```
# asg_sgm_serial [-a]
# asg_serial_info [-a]
```

Parameter	Description
`-a`	Apply command on all SGMs in the security group

Examples

```
# asg_sgm_serial
1_01:
 Board Serial          : AKO0769153
1_02:
 Board Serial          : AKO0585533
2_01:
 Board Serial          : AKO0462069
2_02:
 Board Serial          : AKO0447878

# asg_serial_info
chassis 1 CMM1 serial: 1163978/005
chassis 1 CMM2 serial: 1157482/001
chassis 1 SSM1 serial: 0011140011
chassis 1 SSM2 serial: 0011140012
chassis 1 serial: 1159584/016
chassis 2 CMM1 serial: 1163090/041
chassis 2 CMM2 serial: 1155519/014
chassis 2 SSM1 serial: 0311310621
chassis 2 SSM2 serial: 0311310626
chassis 2 serial: 0831232/001
```

Notes

To show CMM, SSM and Chassis serial numbers, one of the SGMs on each Chassis must be UP. For example, if no UP SGM is found on Chassis-2, the serial numbers for components for all components in the Chassis are not shown or saved.

Redirecting Alerts and Logs to External syslog server (asg_syslog)

Use `asg_syslog` to redirect alert messages and firewall logs to remote syslog servers.

This command lets you:

- Configure remote syslog servers to log all alert messages by:
  - IPv4 address
  - Hostname
- Disable/Enable firewall logs to be sent to the Log Server.
  - Log Server is configured from SmartDashboard:
  - Right-click gateway object > **Edit** > **Logs and Masters** > **Log Servers**.
- Make sure the configuration is consistent on all SGMs.
- Recover configuration on all SGMs by forcing current SGM configuration on all SGMs.

`asg_syslog` is available only from Expert mode.

Syntax

```
# asg_syslog verify|print [ -v ]|recover
```

Parameter	Description
`verify`	Verify configuration consistency on all SGMs
`print [-v]`	Print remote syslog servers configuration `-v` - Verbose mode
`recover`	Recover configuration files on all SGMs and restart syslog service

Example 1

```
# asg_syslog verify
```

Output

```
-------------------------------------------------------------------
|Service       |Path                                      |Result  |
-------------------------------------------------------------------
|CPLog         |/etc/syslog_servers_list.conf             |Passed  |
-------------------------------------------------------------------
|Alert         |/etc/syslog.conf                          |Passed  |
-------------------------------------------------------------------
```

Note - Configuration files on all SGMs are identical.

Example 2

```
# asg_syslog print
```

Output

```
---------------------------------------
|Service       |Server IP     |Status  |
----------------------------------------
|alert         |5.5.5.5       |disable |
----------------------------------------
|alert         |6.6.6.6       |enable  |
----------------------------------------
* Firewall logging is disabled
```

Syntax

Configure remote syslog servers for alerts:

```
# asg_syslog disable|enable|set|delete alert <ip>|<host_name>
```

Configure remote syslog server for firewall logs:

Syntax

```
# asg_syslog disable|enable|set[-s <status>]|delete cplog <ip>|<host_name>
```

Note - When you configure alert syslog servers, the syslog service restarts on all SGMs.

Parameter	Description
set	Set remote syslog server
-s <*status*>	Set connection status Valid values: • enable • disable
disable	Disable firewall logs and alerts to be sent to a remote syslog server defined by IP address or host name. Note: This does not remove the configuration. You can enable it again using enable.
enable	Enable firewall logs and alerts to be sent to a remote syslog server defined by IP address or host name. You can use this parameter after the remote server has been configured.
delete	Delete the remote syslog server
<*ip*>|<*host_name*>	IPv4 address or hostname of the remote syslog server.

Examples

```
# asg_syslog set alert 5.5.5.5
Writing new configuration
Updating all SGMs with new configuration
Restarting syslog service on all SGMs
syslog alert server 5.5.5.5 configured successfully
----------------------------------------
|Service       |Server IP     |Status  |
----------------------------------------
|alert         |5.5.5.5       |enable  |
----------------------------------------
Firewall logging is disabled


# asg_syslog disable alert 5.5.5.5
Updating all SGMs with new configuration
Restarting syslog service on all SGMs
syslog alert server 5.5.5.5 status changed to disable

----------------------------------------
```

```
|Service       |Server IP     |Status  |
----------------------------------------
|alert         |5.5.5.5       |disable |
----------------------------------------
* Firewall logging is disabled


#asg_syslog set cplog 6.6.6.6 -s disable
Writing new configuration
Updating all SGMs with new configuration
syslog cplog server 6.6.6.6 configured successfully

----------------------------------------
|Service       |Server IP     |Status  |
----------------------------------------
|alert         |5.5.5.5       |disable |
----------------------------------------
|cplog         |6.6.6.6       |disable |
----------------------------------------
* Firewall logging is disabled
```

Syntax

Use this command to disable or enable firewall logs to be sent to the Firewall log server (SmartView Tracker):

```
# asg_syslog disable|enable log_server
```

Parameter	Description
`disable`	Disable sending firewall logs to the log server. Log server is configured in SmartDashboard.
`enable`	Enable sending firewall logs to the log server. Log server is configured in SmartDashboard.

Example

```
# asg_syslog disable log_server
# asg_syslog print -v

--------------------------------------------------------------------------------
|Service       |Server IP     |Port          |Protocol#     |RFC version   |Status  |
--------------------------------------------------------------------------------
* Firewall logging is disabled
```

Log Server Distribution (asg_log_servers)

In SmartDashboard, you can configure multiple log servers for each gateway object. In such an environment, the gateway sends its logs to all of its configured log servers. If the gateway object is a 61000/41000 Security System (consisting of many SGMs), each SGM sends its logs to all log servers in the configuration. To reduce the load on the log servers, use `asg_log_servers` to enable log distribution (load sharing).

When enabled, each SGM sends its logs to one log server only. The 61000/41000 Security System automatically decides which log server is assigned to which SGM. This cannot be defined by the user.

Syntax

```
> asg_log_servers
```

Output

```
+-------------------------------------------------+
|             Log Servers Distribution            |
+-------------------------------------------------+
Log Servers Distribution Mode: Disabled

Available Log Servers:
* logServer
* Gaia
* LogServer2

Logs will be sent to all available servers.

Choose one of the following options:
------------------------------------
1) Configure Log Servers Distribution mode
2) Exit

>1
+-------------------------------------------------+
|             Log Servers Distribution            |
+-------------------------------------------------+

Log Servers Distribution Mode: Disabled

Choose the desired option:
--------------------------
1) Enable Log Servers Distribution mode
2) Disable Log Servers Distribution mode
3) Back
```

If log server distribution is already enabled, the command shows which log servers are assigned to each SGM:

```
+-------------------------------------------------+
|             Log Servers Distribution            |
+-------------------------------------------------+

Log Servers Distribution Mode: Enabled

Available Log Servers:
* LogServer
* Gaia
* LogServer2

Log Servers Distribution:

+--------------------------------------------------------------+
| Blade id |        Chassis 1        |        Chassis 2        |
|--------------------------------------------------------------|
|    1     |  Gaia                   |  Gaia                   |
|    2     |  LogServer2             |  LogServer2             |
|    3     |  LogServer              |  LogServer              |
|    4     |  Gaia                   |  -                      |
|    5     |  -                      |  -                      |
|    6     |  LogServer              |  -                      |
|    7     |  -                      |  Gaia                   |
|    8     |  -                      |  LogServer2             |
|    9     |  LogServer              |  LogServer              |
|    10    |  Gaia                   |  -                      |
|    11    |  LogServer2            |  -                   |
|    12    |  -                     |  -                   |
+-----------------------------------------------------------------+

("-" - Blade is not in Security Group)

Choose one of the following options:
------------------------------------
1) Configure Log Servers Distribution mode
2) Exit
```

Note - You cannot configure an SGM to send its logs to a particular log server. Distribution takes place automatically.

Configuring a Dedicated Logging Port

The 61000/41000 Security System logging mechanism lets each SGM forward logs directly to a logging server over the SSM's management ports. However, management ports can experience a high load when a large number of logs are forwarded. Load on the SSM management ports can be significantly reduced by:

- Setting up a dedicated SSM port for logging
- Assigning the dedicated logging port to each SGM

To set up a dedicated logging port:

1. Install a log server and create an object for it in SmartDashboard.
2. Connect the log server directly to a management port on the SSM.

   Important - Do not use the same port which connects to the Security Management server.
3. In `gclish`, use `set interface` to configure the port as a dedicated logging port:

Syntax

```
> set interface <if_name> ipv4-address <ip> mask-length <length>
```

Parameter	Description
<*if_name*>	The interface that connects directly to the log server
<*ip*>	IPv4 address of the logging server
<*length*>	Subnet mask length

Example

```
> set interface eth1-Mgmt2 ipv4-address 2.2.2.10 mask-length 24
```

Output

```
1_0l:
success

1_02:
success

1_03:
success

2_01:
success
```

```
  2_02:
  success

  2_03:
  success

>
```

Notes:

- For each SGM, eth1-Mgmt2 is set as a unique logging port.
- `2.2.2.0/24` is the logging server network or leads to the logs server network

To connect to the logging server:

1. Open SmartDashboard.
2. Open the Single Management Object (SMO) for the 61000/41000 Security System.
3. On the **Logs and Masters > Log Servers** page, select **Define Log Servers**.
4. Select the dedicated log server.
5. Install the policy.

Note - The SMO in SmartDashboard makes sure that return traffic from the logging server, such as ACKS, reaches the correct SGM.

Command Auditing

Command auditing:

- Notifies users about critical actions they are about to do
- Obtains confirmation for critical actions
- Creates forensic logs

If users confirm the action, they are requested to supply their names and a reason for running the command. If the command affects a critical device or a process (pnote) a second confirmation can be required.

For example, if you use administrative privileges to change the state of a SGM to DOWN the output looks like this:

```
> asg_sgm_admin –b 2_01 down
You are about to perform sgm_admin down on blades: 2_01

Are you sure? (y – yes, any other key – no) y

sgm_admin down requires auditing
Enter your full name: John Smith
Enter reason for sgm_admin down [Maintenance]:
WARNING: sgm_admin down on SGM: 2_01, User: John Smith, Reason: Maintenance
```

To see the audit logs:

Run:

```
# asg log audit
```

Example

```
# asg log audit
```

Output

```
Aug 11 14:14:21 2_01  WARNING: Chassis_admin up on chassis: 1, User: susan, Reason: Maintenance
Aug 11 16:45:15 2_01  WARNING: Reboot on blades: 1_01,1_02,1_03,1_04,1_05,2_02,2_03,2_04,2_05, User:
susan, Reason: Maintenance
Aug 18 14:28:57 2_01  WARNING: Chassis_admin down on chassis: 2, User: susan, Reason: Maintenance
Aug 18 14:31:08 2_01  WARNING: Chassis_admin up on chassis: 1, User: Peter, Reason: Maintenance
Aug 18 14:32:32 2_01  WARNING: Chassis_admin down on chassis: 2, User: O, Reason: Maintenance
Aug 20 15:38:58 2_01  WARNING: Blade_admin down on blades: 2_02,2_03,2_04,2_05, User: Paul, Reason:
Maintenance
Aug 21 10:00:05 2_01 CRITICAL: Reboot on blades: all, user: ms, Reason: Maintenance
```

Showing the 61000/41000 Security System Version (ver)

Use this command to show the 61000/41000 Security System version. For a list of official 61000/41000 Security System versions, see the R76SP.40 61000/41000 Security System home page http://supportcontent.checkpoint.com/solutions?id=sk110515.

Syntax

```
> ver
```

Output

For NG 61000 Security System:

```
1_01:
Product version Check Point 61000 R76
OS build 106
OS kernel version 2.6.18-92cpx86_64
OS edition 64-bit
```

For a 41000 Security System:

```
1_04:
Product version Check Point Gaia 41000 R76
OS build 105
OS kernel version 2.6.18-92cpx86_64
OS edition 64-bit
```

Viewing a Log File (asg log)

Use this command to see the contents of a specified log file.

Syntax

```
> asg log [-b <sgm_ids>] <log_name> [-tail [<n>]] [-f <filter>]
```

Parameter	Description
-b <*sgm_ids*>	Works with SGMs and/or Chassis as specified by <*sgm_ids*>. <*sgm_ids*> can be: • No <*sgm_ids*> specified or all shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (1_1,1_4) • A range of SGMs (1_1-1_4) • One Chassis (Chassis1 or Chassis2) • The active Chassis (chassis_active)

Parameter	Description

<*log_name*>	Enter the log file to show:
	• `audit`
	Shows the audit logs in `/var/log`
	For example:
	`/var/log/asgaudit.log.1`
	•
	• `ports`
	Shows the ports logs in `/var/log`
	For example:
	`/var/log/ports`
	• `dist_mode`
	Shows the logs for distribution mode activity.

`-tail [<`*n*`>]`	Show only last *n* lines of the log file for each SGM. For example, `-tail 3` shows only the last three lines of the specified log file. Default = 10 lines.
`-f <`*filter*`>`	Word or phrase use as a filter. For example, `-f debug`

Example - Audit logs

```
> asg log audit
Feb 02 17:36:12 1_01  WARNING: Blade_admin up on blades: 1_02,1_03,1_04,1_05,2_01,2_02,2_03,2_04,2_05,
User: y, Reason: y
Feb 03 08:16:17 1_01  WARNING: Blade_admin down on blades:
1_02,1_03,1_04,1_05,2_01,2_02,2_03,2_04,2_05, User: y, Reason: y
Feb 03 08:17:40 1_01  WARNING: Blade_admin up on blades: 1_02,1_03,1_04,1_05,2_01,2_02,2_03,2_04,2_05,
User: y, Reason: y
Feb 03 08:19:53 1_01  WARNING: Blade_admin down on blades:
1_02,1_03,1_04,1_05,2_01,2_02,2_03,2_04,2_05, User: y, Reason: y
Feb 03 08:22:33 1_01  WARNING: Blade_admin up on blades: 1_02,1_03,1_04,1_05,2_01,2_02,2_03,2_04,2_05,
User: y, Reason: y
Feb 03 08:23:30 1_01  WARNING: Reboot on blades: 1_02,1_03,1_04,1_05,2_01,2_02,2_03,2_04_05, User: y,
Reason: y
Feb 03 08:38:16 1_01  WARNING: Reboot on blades: 1_02,1_03,1_04,1_05,2_01,2_02,2_03,2_04,2_05, User:
y, Reason: y
Feb 03 09:21:09 1_01  WARNING: Reboot on blades: 1_02,1_03,1_04,1_05,2_01,2_02,2_03,2_04,2_05, User:
y, Reason: y
Feb 03 11:07:08 1_01  WARNING: Reboot on blades: 1_02,1_03,1_04,1_05,2_01,2_02,2_03,2_04,2_05, User:
y, Reason: y
Feb 03 11:16:56 1_01  WARNING: Reset sic on blades: all, User: y, Reason: y
Feb 03 11:33:10 1_01  WARNING: Reset sic on blades: all, User: y, Reason: y
Feb 03 11:50:08 1_01  WARNING: Reset sic on blades: all, User: y, Reason: y
Feb 03 13:32:32 1_01  WARNING: Reset sic on blades: all, User: y, Reason: y
Feb 03 14:30:26 1_01  WARNING: Reset sic on blades: all, User: kaki, Reason: pipi
Feb 03 14:48:03 1_01  WARNING: Reset sic on blades: all, User: kaki, Reason: pipi
Feb 03 15:34:11 1_01  WARNING: Reset sic on blades: all, User: y, Reason: y
Feb 03 17:55:23 1_01  WARNING: Reboot on blades: 1_02,1_03,1_04,1_05,2_01,2_02,2_03,2_04,2_05, User:
y, Reason: y
```

Example - Port logs (last 12 lines)

```
> asg log ports-tail 12
Feb  3 18:01:40 2_05 Athens-ch02-05 cmd: Chassis 2 eth2-09 link is down
Feb  3 18:01:40 2_05 Athens-ch02-05 cmd: Chassis 2 eth2-10 link is down
Feb  3 18:01:40 2_05 Athens-ch02-05 cmd: Chassis 2 eth2-11 link is down
Feb  3 18:01:40 2_05 Athens-ch02-05 cmd: Chassis 2 eth2-12 link is down
Feb  3 18:01:40 2_05 Athens-ch02-05 cmd: Chassis 2 eth2-13 link is down
Feb  3 18:01:40 2_05 Athens-ch02-05 cmd: Chassis 2 eth2-14 link is down
Feb  3 18:01:40 2_05 Athens-ch02-05 cmd: Chassis 2 eth2-15 link is down
Feb  3 18:01:40 2_05 Athens-ch02-05 cmd: Chassis 2 eth2-16 link is down
Feb  3 18:01:40 2_05 Athens-ch02-05 cmd: Chassis 2 eth2-Mgmt1 link is down
Feb  3 18:01:40 2_05 Athens-ch02-05 cmd: Chassis 2 eth2-Mgmt2 link is down
Feb  3 18:01:40 2_05 Athens-ch02-05 cmd: Chassis 2 eth2-Mgmt3 link is down
Feb  3 18:01:40 2_05 Athens-ch02-05 cmd: Chassis 2 eth2-Mgmt4 link is down
```

Example - Using a filter

```
> asg log -b 1_01,1_04 dist_mode -f bridge
Feb  2 18:10:30 1_01 Athens-ch01-01 distutil:0: initialize_environment: vs-ids-bridges = 4
Feb  2 18:10:30 1_01 Athens-ch01-01 distutil:0: initialize_environment: vs-ids-vsbridges = 4
Feb  2 18:12:31 1_01 Athens-ch01-01 distutil:0: initialize_environment: vs-ids-bridges = 4
Feb  2 18:12:31 1_01 Athens-ch01-01 distutil:0: initialize_environment: vs-ids-vsbridges = 4
Feb  2 18:14:14 1_01 Athens-ch01-01 distutil:0: initialize_environment: vs-ids-bridges = 4
Feb  2 18:14:14 1_01 Athens-ch01-01 distutil:0: initialize_environment: vs-ids-vsbridges = 4
Feb  2 18:14:30 1_01 Athens-ch01-01 distutil:0: initialize_environment: vs-ids-bridges = 4
Feb  2 18:14:30 1_01 Athens-ch01-01 distutil:0: initialize_environment: vs-ids-vsbridges = 4
Feb  2 18:16:19 1_01 Athens-ch01-01 distutil:0: initialize_environment: vs-ids-bridges = 4
```

Looking at the Audit Log File (asg_auditlog)

Use `asg_auditlog` to see the contents of the `auditlog` file. This log file contains an entry for each change made to the SGM configuration database with `gclish` or other commands. The `auditlog` file for each SGM is located in the `/var/log` directory.

The `asg_auditlog` command collects and summarizes records from the SGMs. The output shows actions that occur on different SGMs within a certain time period (default 5 seconds) on one line. These are considered to be global actions applicable to all SGMs. You can change this time period.

The log contains two types of activities:

Permanent - The activity permanently changes the configuration database on the SGM hard disk.

Transient - The activity changes the configuration database in SGM memory, which does not survive reboot.

Syntax

```
> auditlog [-b <sgm_ids>] [-d <n>] [-tail [n]] [-f <filter>]
```

Parameter	Meaning
`-b` *<sgm_ids>*	Works with SGMs and/or Chassis as specified by *<sgm_ids>*. *<sgm_ids>* can be: • No *<sgm_ids>* specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`)
`-d` *<n>*	Number of seconds between the same actions that occur on different SGMs, which show on one output line. Default = 5 seconds.

Parameter	Meaning
`-tail <n>`	Show only last `n` lines of the log file for each SGM. For example, `-tail 3` shows only the last three lines of the specified log file. Default = 10 lines.
`-f <filter>`	Word or phrase to use as an output filter. For example, `-f t` shows only transient changes.

Example - Show last lines

This example shows the last five activities, in this case, **cpstop** actions.

```
> asg_auditlog -tail 5
Feb  3 05:30:49 admin localhost p -command:cpstop t [1 Blades: 1_03]
Feb  3 05:30:49 admin localhost p -command:cpstop:description Stop\ Check\ Point\ products\ installed
[1 Blades: 1_03]
Feb  3 05:30:49 admin localhost p +command:cpstop:description Global\ extension\ for\ cpstop
1 Blades: 1_03]
Feb  3 05:30:49 admin localhost p -command:cpstop:description Global\ extension\ for\ cpstop
1 Blades: 1_03]
Feb  3 05:30:49 admin localhost p -command:cpstop:path /bin/cpstop_start [1 Blades: 1_03]
```

Notes:

- **p +** = Permanent action that added or changed an item in the configuration database.
- **p -** = Permanent action that deleted an item in the configuration database
- **t +** = Transient action that added or changed an item in the configuration database in memory only.
- **t -** = Transient action that deleted an item in the configuration database in memory only.

Example - filter

This example shows only permanent configuration save actions.

```
> asg_auditlog -f p +configurationSave
Feb  3 15:21:51 admin localhost p +configurationSave t [2 Blades: 1_01,1_02]
Feb  3 15:21:58 admin localhost p +configurationSave t [2 Blades: 1_03,1_04]
Feb  3 15:22:03 admin localhost p +configurationSave t [3 Blades: 1_01,1_02,2_02]
Feb  3 15:22:08 admin localhost p +configurationSave t [4 Blades: 2_01,2_03,2_04,2_05]
Feb  3 15:24:23 admin localhost p +configurationSave t [2 Blades: 1_03,1_04]
Feb  3 15:24:24 admin localhost p +configurationSave t [2 Blades: 1_03,1_04]
Feb  3 15:24:29 admin localhost p +configurationSave t [5 Blades: 1_03,1_04,2_03,2_04,
Feb  3 15:24:30 admin localhost p +configurationSave t [4 Blades: 2_01,2_03,2_04,2_05]
Feb  3 15:24:35 admin localhost p +configurationSave t [2 Blades: 2_01,2_02]
Feb  3 15:24:36 admin localhost p +configurationSave t [1 Blades: 2_02]
Feb  3 15:24:44 admin localhost p +configurationSave t [2 Blades: 2_01,2_03]
Feb  3 15:24:51 admin localhost p +configurationSave t [2 Blades: 2_02,2_04]
Feb  3 15:24:56 admin localhost p +configurationSave t [1 Blades: 2_05]
```

Working with the Firewall Database Configuration (asg config)

Use this command to show the newest firewall database configuration. You can also save the newest configuration to a file. The output and saved file include configuration information for all SGMs. The `asg config` command is useful to:

- Copy the firewall configuration to a different system. For example, you can use the saved configuration from an existing 61000/41000 Security System to set up the new 61000/41000 Security System.
- Quickly re-configure a system that was reverted to factory defaults. Before reverting to the factory default image, save the existing configuration. Then use it to override the factory settings.

Syntax

```
> asg config  show|save [-t] [<file_path>]
```

Parameter	Description
`show`	Show the existing database configuration
`save`	Save the current configuration to a file **Note**: If you do not include a path, the file is saved to: `/home/admin`
`-t`	Add a timestamp to the file name. (`save` only)
<*file_path*>	Name and path of the saved configuration file. If you do not enter a path, the configuration is saved to: `/home/admin`

Example

```
> asg config save -t mycongfig
```

This example saves the current configuration to: `/home/admin/myconfig`

Showing Software and Firmware versions (asg_version)

Use `asg_version` to:

- Retrieve system configuration
- Retrieve software versions:
  - Check Point software (Firewall and Performance Pack versions)
  - Firmware versions for SGMs, SSMs, and CMMs
  - Make sure that system hardware components are running approved software and firmware versions

Syntax

```
> asg_version -h
> asg_version [verify] [-v] [-i] [-b <sgm_ids>]
```

Parameter	Meaning
`-h`	Show complete command syntax

Parameter	Meaning
`verify`	Makes sure that system hardware components run approved software and firmware versions
`-i`	Show active and standby SGMs
`-b` *<sgm_ids>*	Works with SGMs and/or Chassis as specified by *<sgm_ids>*. *<sgm_ids>* can be: • No *<sgm_ids>* specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`)

Showing a List of Two SGMs

```
> asg_version -b 1_01,1_03
SGMs
======

----------
-*- 2 SGMs: 1_01 1_03 -*-
OS build 42, OS kernel version 2.6.18-92cpx86_64, OS edition 64-bit

Hardware
--------
-*- 1 blade: 1_01 -*-
BIOS: 1.30 BL: 1.52 IPMC: 1.52 FPGA: 2.40 FPGARE: 2.40
-*- 1 blade: 1_03 -*-
BIOS: 0.54 BL: 1.42 IPMC: 1.42 FPGA: 2.38 FPGARE: 2.38
OS version
----------
BIOS: 0.54 BL: 1.42 IPMC: 1.42 FPGA: 2.38 FPGARE: 2.
```

Showing Verbose Mode

```
> asg_version -v

+-----------------------------------------------------------------------------+
| Hardware Versions                                                           |
+-----------------------------------------------------------------------------+
| Component       | Type            | Configuration       | Firmware          |
+-----------------------------------------------------------------------------+
| Chassis 2                                                                   |
+-----------------------------------------------------------------------------+
| SSM1            | SSM160          | N/A                 | 2.4.C7            |
| SSM2            | N/A             | N/A                 | N/A               |
| CMM             | N/A             | N/A                 | 2.83              |
+-----------------------------------------------------------------------------+

SGMs
======
Type
----------
-*- 2 blades: 2_02 2_03 -*-
SGM220
```

```
OS version
----------
-*- 2 blades: 2_02 2_03 -*-
OS build 80, OS kernel version 2.6.18-92cpx86_64, OS edition 64-bit

FireWall-1 version
------------------
-*- 2 blades: 2_02 2_03 -*-
This is Check Point VPN-1(TM) & FireWall-1(R) 61000_R76 - Build 083
kernel: 61000_R76 - Build 083

Performance Pack version
------------------------
-*- 2 blades: 2_02 2_03 -*-
This is Check Point Performance Pack version: 61000_R76 - Build 083
Kernel version: 61000_R76 - Build 083

Hardware
--------
-*- 1 blade: 2_02 -*-
BIOS: 1.30 BL: 1.42 IPMC: 1.52 FPGA: 2.40 FPGARE: 2.40
-*- 1 blade: 2_03 -*-
BIOS: 1.30 BL: 1.52 IPMC: 1.54 FPGA: 2.40 FPGARE: 2.40

SSD
---
-*- 1 blade: 2_02 -*-
Firmware Version: 2CV102M3
-*- 1 blade: 2_03 -*-
Firmware Version: 4PC10362

Number of cores
---------------
-*- 1 blade: 2_02 -*-
8
-*- 1 blade: 2_03 -*-
12

Number of CoreXL instances
--------------------------
-*- 2 blades: 2_02 2_03 -*-
4

CPUs frequency
--------------
-*- 1 blade: 2_02 -*-
2.13GHz
-*- 1 blade: 2_03 -*-
2.4GHz
```

Showing System Messages (asg_varlog)

Use this command to show system messages written to message files, stored in the `/var/log` directory on SGMs. The output shows in chronological sequence. Each line shows the SGM that created the log entry.

Syntax

```
> asg_varlog [-b <sgm_ids>] [-tail <n>] [-f <filter>]
> asg_varlog -h
```

Parameter	Meaning
-b <*sgm_ids*>	The SGMs from which to collect /var/log/messages. Works with SGMs and/or Chassis as specified by <*sgm_ids*>. <*sgm_ids*> can be: • No <*sgm_ids*> specified or all shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (1_1,1_4) • A range of SGMs (1_1-1_4) • One Chassis (Chassis1 or Chassis2) • The active Chassis (chassis_active)
-tail <*n*>	Show only last n lines of the log file for each SGM. For example, -tail 3 shows only the last three lines of the specified log file. Default = 10 lines.
-f <*filter*>	Word or phrase to use as an output filter. For example, -f ospf shows only OSPF messages.
-h	Shows command syntax and help information.

Example

This example shows messages on Chassis1 containing the word "Restarted".

```
> asg_varlog -b chassis1 -f Restarted
Feb  5 12:40:07 1_03 Athens-ch01-03 pm[8465]: Restarted /bin/routed[8489], count=1
Feb  5 12:40:09 1_04 Athens-ch01-04 pm[8449]: Restarted /bin/routed[9995], count=1
Feb  5 12:40:09 1_04 Athens-ch01-04 pm[8449]: Restarted /opt/CPsuite-R76/fw1/bin/cmd[11291], count=1
Feb  5 12:40:09 1_04 Athens-ch01-04 pm[8449]: Restarted /usr/libexec/gexecd[11292], count=1
Feb  5 12:40:10 1_03 Athens-ch01-03 pm[8465]: Restarted /usr/libexec/gexecd[9701], count=1
Feb  5 12:40:10 1_03 Athens-ch01-03 pm[8465]: Restarted /bin/routed[11328], count=2
Feb  5 12:40:10 1_05 Athens-ch01-05 pm[8458]: Restarted /bin/routed[9734], count=1
Feb  5 12:40:10 1_05 Athens-ch01-05 pm[8458]: Restarted /usr/libexec/gexecd[11331], count=1
Feb  5 12:40:11 1_01 Athens-ch01-01 pm[8463]: Restarted /bin/routed[12253], count=3
Feb  5 12:40:11 1_04 Athens-ch01-04 pm[8449]: Restarted /bin/routed[11378], count=2
Feb  5 12:40:11 1_04 Athens-ch01-04 pm[8449]: Restarted /opt/CPsuite-R76/fw1/bin/cmd[11379], count=2
```

Monitoring Virtual Systems (cpha_vsx_util monitor)

Use this command to stop or start Virtual Systems (VS) monitoring.

The state of an SGM is not affected by an unmonitored Virtual Systems. For example, an unmonitored Virtual System in problem state (pnote) is ignored. The SGM state does change to DOWN.

A Virtual System that is not monitored is useful if you want an SGM to be UP, even if a specific Virtual System is DOWN or does not have a Policy (for example, after you unload the local policy).

Syntax

```
# cpha_vsx_util monitor start|stop <vs_ids>
# cpha_vsx_util monitor show
```

Parameter	Description
show	Show all unmonitored Virtual Systems
stop	Stop monitoring the Virtual Systems

Parameter	Description
`start`	Start monitoring the Virtual Systems
<*vs_ids*>	<*vs_ids*> can be: • No <*vs_ids*> (default) - Shows the current Virtual System context. • One Virtual System. • A comma-separated list of Virtual Systems (1, 2, 4, 5). • A range of Virtual Systems (VS 3-5). • `all` - Shows all Virtual Systems. **Note:** This parameter is only relevant in a VSX environment.

Note - When you stop Virtual System monitoring, you must run `cpha_vsx_util monitor start` to start it again. Monitoring does not start automatically after reboot.

Working with SNMP

You can use SNMP to monitor different aspects of the 61000/41000 Security System, including:

- Software versions
- Hardware status
- Key performance indicators
- Chassis high availability status

To monitor the system using SNMP

1. Upload the Check Point MIB to your third-party SNMP monitoring software.

   The SNMP MIB is located on each SGM under: `$CPDIR/lib/snmp/chkpnt.mib`

   To monitor the 61000/41000 Security System, the supported OIDs are under `iso.org.dod.internet.private.enterprise.checkpoint.products.asg` (OID 1.3.6.1.4.1.2620.1.48)
2. Enable the SNMP agent on the 61000/41000 Security System.

   In gclish, run:

   ```
   > set snmp agent on
   ```

SNMP Traps

The 61000/41000 Security System supports this SNMP trap only:

```
iso.org.dod.internet.private.enterprise.checkpoint.products.asgTrap
(OID 1.3.6.1.4.1.2620.1.2001)
```

The SNMP traps MIB is located on each SGM under: `$CPDIR/lib/snmp/chkpnt-trap.mib`

Note - The `set snmp traps` command is not supported. You must use the `asg alert` configuration wizard for this purpose.

To learn more about SNMP, see Configuring asg alerts ("Configuring Alerts for SGM and Chassis Events (asg alert)" on page 135).

SNMP in a VSX Gateway

There are two SNMP modes for a 61000/41000 Security System configured as a VSX Gateway:

- Default Mode - Monitor global SNMP data from the 61000/41000 Security System. Data is accumulated from all SGMs for all Virtual System.
- Virtual Systems Mode - Monitor each Virtual System separately.

Note - SNMP traps are supported for VS0 only.

Supported SNMP Versions

The SNMP Virtual Systems mode uses SNMP version 3 to query the Virtual Systems. You can run remote SNMP queries on each Virtual System in the VSX Gateway.

For systems that only support SNMP versions 1 and 2:

- You cannot run remote SNMP queries for each Virtual System. You can only run a remote SNMP query on VS0.
- You can use gclish to change the Virtual System context and then run a local SNMP query on it.

Enabling the SNMP Virtual System Mode

To use SNMP for each Virtual Systems:

1. Configure an SNMP V3 user:
   ```
   > add snmp usm user jon security-level authNoPriv authpass-phrase VALUE
   ```
2. Set the SNMP mode:
   ```
   > set snmp mode vs
   ```
   or
   ```
   > set snmp mode default
   ```
3. Start SNMP agent:
   ```
   > set snmp agent on
   ```

To see Virtual System throughput from a Linux host:

Run:

```
# snmpwalk -m $CPDIR/lib/snmp/chkpnt.mib -n ctxname_vsid1 -v 3 -l authNoPriv -u jon
-A mypassword 192.0.2.72 asgThroughput
```

To query Virtual System throughput, from its context:

1. Go to Expert mode.
2. Change to the applicable Virtual System:
   ```
   # vsenv <vs_ids>
   ```
3. Run:
   ```
   # snmpwalk -m $CPDIR/lib/snmp/chkpnt.mib -v 2c -c public localhost
   asgThroughput
   ```

Common SNMP MIBs

This table shows frequently used SNMP MIBs that are applicable to the 61000/41000 Security System.

Note:

<IPver_index>= 20 for IPv4 or 21 for IPv6

Name	Type	OID	Comments
System Throughput	String	1.3.6.1.4.1.2620.1.48.<IPver_index>.1	
System Connection Rate (cps)	String	1.3.6.1.4.1.2620.1.48.<IPver_index>.2	
System Packet Rate(pps)	String	1.3.6.1.4.1.2620.1.48.<IPver_index>.3	
System Concurrent conn.	String	1.3.6.1.4.1.2620.1.48.<IPver_index>.4	
System Accelerated cps	String	1.3.6.1.4.1.2620.1.48.<IPver_index>.6	
System non-accelerated cps	String	1.3.6.1.4.1.2620.1.48.<IPver_index>.7	
System Accelerated Concurrent conn.	String	1.3.6.1.4.1.2620.1.48.<IPver_index>.8	
System Non-accelerated concurrent conn.	String	1.3.6.1.4.1.2620.1.48.<IPver_index>.9	
System CPU load AVG.	String	1.3.6.1.4.1.2620.1.48.<IPver_index>.10	
System Acceleration CPU load AVG	String	1.3.6.1.4.1.2620.1.48.<IPver_index>.11	
System FW instances load AVG	String	1.3.6.1.4.1.2620.1.48.<IPver_index>.14	
System VPN Throughput	String	1.3.6.1.4.1.2620.1.48.<IPver_index>.17	

Name	Type	OID	Comments
System Path distribution (fast, medium, slow, drops).	Table	1.3.6.1.4.1.2620.1.48.<IPver_index>.24	Path Distribution of: Throughput PPS CPS Concurrent conn
Per SGM counters	Table	1.3.6.1.4.1.2620.1.48.<IPver_index>.25	Counters of: Throughput cps pps concurrent conn sxl CPU usage (avg/min/max) fw CPU usage (avg/min/max)
Performance peaks	Table	1.3.6.1.4.1.2620.1.48.<IPver_index>.26	
Sensors Per Chassis	Table	1.3.6.1.4.1.2620.1.48.22.1.1	Status Details of: Fans SSMs CPU temp CMM PSUs PSU Fans
Resources Per SGM	Table	1.3.6.1.4.1.2620.1.48.23	Memory and HD utilization
CPU Utilization Per SGM	Table	1.3.6.1.4.1.2620.1.48.29	

Working with Active/Standby High Availability

In This Section:

How Active Standby Works

Chassis Active/Standby High Availability is based on two fully synchronized Chassis for redundancy, with seamless failover. The Active Chassis handles all traffic, while the Standby Chassis is continuously synchronized with the Active Chassis. Traffic continues uninterrupted during Chassis failover.

This release supports these Active/Standby High Availability modes:

Mode ID	Description
`0`	**Active/Standby - Active Up** The currently Active Chassis stays Active unless it goes DOWN, or the Standby Chassis has a higher Chassis quality grade.
`1`	**Active/Standby - Primary Up Chassis 1** Chassis 1 always stays Active unless it goes DOWN, or the Standby Chassis has a higher Chassis quality grade.
`2`	**Active/Standby - Primary UP Chassis 2** Chassis 2 is always Active unless it goes DOWN, or the Standby Chassis has a higher Chassis quality grade.

To make sure that the most reliable Chassis is always Active, and to prevent unnecessary failover, the 61000/41000 Security System calculates a *quality grade* for each Chassis. This is based on continuous monitoring of critical components and traffic characteristics. See Setting Chassis Weights ("Setting Chassis Weights (Chassis High-Availability Factors)" on page 193) for a detailed explanation of the quality grade system.

Chassis High Availability works on the principle that the Chassis with the highest quality grade becomes the Active Chassis. A configurable minimum grade differential prevents unnecessary failover, which can cause performance degradation. Automatic failover occurs only when the Standby Chassis quality grade is greater than the Active Chassis quality grade, plus the minimum differential. See Setting the quality grade differential (on page 194) for details.

Each Chassis data port has a unique MAC address. The MAC addresses for the Chassis SGMs are the same. A Chassis failover event sends GARP/ICMv6 packets to each interface. This informs the network to use the other interfaces. See GARP Chunk Mechanism ("Working with the GARP Chunk Mechanism" on page 291) for details.

You can use `gclish` commands to configure these High Availability parameters:

- Active/Standby mode (Active UP/Primary UP)
- Chassis quality grade factors
- Failover grade difference for failover
- Failover freeze interval
- Port priority

Synchronizing Clusters on a Wide Area Network

You can install your Chassis at two different remote sites as a geographically distributed cluster. There are two limitations to this capability:

1. The synchronization network must guarantee no more than 100ms latency and no more than 5% packet loss.
2. The synchronization network can include switches and hubs. Routers cannot be installed on the synchronization network because they drop Cluster Control Protocol packets.

Configuring Active/Standby High Availability

Setting the Chassis ID

You must make sure that the Chassis IDs are different before you start to configure the software. Chassis IDs are configured on the CMM and should be <*1*> for the first Chassis and <*2*> for the second Chassis.

Note - If the 61000/41000 Security System is up and running, change the Chassis ID on the Standby Chassis. You must perform Chassis failover.

To set the Chassis ID on the NG 61000 Security System

1. Remove the top CMM from the Chassis.
2. Log in to the remaining CMM.
3. Connect the serial cable to the console port on the CMM.
4. Connect to the CMM with a terminal emulation application.
5. Make sure that the Speed (baud rate) is set to 9600.

   No IP address is necessary.
6. Log in with user name and password `admin/admin`.
7. Open `/etc/shmm.cfg` in a text editor.
8. Search for and set `SHMM_CHASSID=` to the correct Chassis ID:

   ```
   Chassis ID
   SHMM_CHASSID=<Chassis_id>
   ```
9. Remove the lower CMM, which you just reconfigured, from the Chassis.
10. Insert the top CMM into the Chassis.
11. Do steps 2 - 8 on the top CMM.
12. Remove the top CMM from the Chassis.
13. Insert both CMMs into the Chassis.
14. Attach the correct identification labels to the Chassis and CMMs.

This step is required if the Chassis has already been configured (after the First Time Configuration Wizard).

15. Remove all SGMs from the Chassis and then reinsert them.

    This step causes a hard reboot of the system.

To set the Chassis ID on the 41000 Security System

1. Remove the right CMM from the Chassis.
2. Log in to the remaining CMM.
3. Connect the serial cable to the console port on the CMM.
4. Connect to the CMM with a terminal emulation application.
5. Make sure that the Speed (baud rate) is set to 9600.

   No IP address is necessary.
6. Log in with user name and password `admin/admin`.
7. Open `/etc/shmm.cfg` in a text editor.
8. Search for and set SHMM_CHASSID= to the correct Chassis ID:
   ```
   Chassis ID
   SHMM_CHASSID=<Chassis_id>
   ```
9. Remove from the left CMM from the Chassis.
10. Insert the right CMM into the Chassis.
11. Do steps 2-8 on the right CMM.
12. Remove the right CMM from the Chassis.
13. Insert both CMMs into the Chassis.
14. Attach the correct identification labels to the Chassis and CMMs.

    This step is required if the Chassis has already been configured (after the First Time Configuration Wizard).
15. Remove all SGMs from the Chassis and then reinsert them.

    This step causes a hard reboot of the system.

Setting Chassis Weights (Chassis High-Availability Factors)

Each component in a Chassis has a quality weight factor, which sets its relative importance to overall Chassis health. For example, ports are more important than fans and are typically assigned a higher weight value. The Chassis grade is the sum of all component weight values. In a High Availability environment, the Chassis with the higher grade becomes Active and handles traffic. The grade for each component = (Unit Weight) X (Number of UP components)

To see the weight of each component:

```
> asg stat -v
```

Use `set chassis high-availability factors` to configure a component's weight.

Syntax

```
> set chassis high-availability factors SGM <sgm_factor>
> set chassis high-availability factors port other <port_other_factor> | port
standard <port_standard_factor>
> set chassis high-availability factors sensor cmm <cmm_factor> | sensor fans
<fans_factor> | sensor power_supplies <psu_factor> | sensor ssm <ssm_factor>
> set chassis high-availability factors pnote pingable_hosts <ping_factor>
```

Parameter	Description
<*sgm_factor*>	Weight factor for an SGM Valid range: Integer between 0 and 1000
<*port_other_facto*r>	High grade port factor Valid range: Integer between 0 and 1000
<*port_standard_factor*>	Standard grade port factor Valid range: Integer between 0 and 1000
<*cmm_factor*>	CMM weight factor Valid range: Integer between 0 and 100
<*fans_factor*>	Fan unit factor Valid range: Integer between 0 and 99
<*psu_factor*>	Power supply unit factor Valid range: Integer between 0 and 99
<*ssm_factor*>	SSM factor This factor applies to all SSMs. Valid range: Integer between 0 and 100
<ping_factor>	Pingable hosts factor Weight factor for pingable hosts, which shows if they are properly connected to their hosts. Valid range: Integer between 0 and 99

Examples

```
> set chassis high-availability factors sgm 100
> set chassis high-availability factors Port other 70
> set chassis high-availability factors Port standard 50
> set chassis high-availability factors sensor cmm 40
> set chassis high-availability factors sensor fans 30
> set chassis high-availability factors sensor power_supplies 20
> set chassis high-availability factors sensor ssm 45
> set chassis high-availability factors pnote pingable_hosts 99
```

Setting the Quality Grade Differential

Use the `set chassis high-availability failover` command to set the minimum quality grade differential that causes failover.

Syntax

```
> set chassis high-availability failover <trigger>
```

Parameter	Description
<*trigger*>	Minimum difference in Chassis quality grade to trigger failover Valid values: 1-1000

Setting the Failover Freeze Interval

A Chassis cannot failover a second time until the specified failover freeze interval expires. The default failover freeze interval is:

- For Primary Up - 150 seconds
- For Active Up - 30 seconds
- For VSLS - 150 seconds

If the Standby Chassis grade changes (to a value greater than the minimum quality grade gap for failover), the Standby Chassis fails over and becomes Active. The failover does not start until the freeze interval expires. This confirms that the Chassis quality grade is stable, before it becomes active. For example, a Chassis quality grade can become unstable if a fan goes UP and DOWN frequently.

Syntax

```
> set chassis high-availability freeze_interval <freeze_interval>
```

Parameter	Description
<*freeze_interval*>	Minimum time in seconds to wait until the next Chassis failover Valid range: 1-1000

Note - When you run `asg stat` after Chassis failover, the freeze time shows in the output.

Setting Port Priority

For each Chassis port, use `set chassis high-availability port priority` to set a port priority (high or standard) for each port.

Syntax

```
> set chassis high-availability port <if_name> priority <priority>
```

Parameter	Description
<*if_name*>	Interface name
<*priority*>	Port grade Valid values: • 1 - Standard priority • 2 - Other priority

Use this command together with: `set chassis high-availability factors port`

1. Set the port grade as standard or high.

   For example:
   ```
   > set chassis high-availability factors port standard 50
   ```

This sets the standard grade at 50.

2. Set the port to high grade or standard grade.

   For example:

   ```
   > set chassis high-availability port eth1-01 priority 2
   ```

   This assigns to `eth1-01` the standard port grade.

Advanced Features

Working with Link Preemption

The Link Preemption Mechanism prevents constant Chassis failover and failback when there is interface link flapping. When you enable this feature, an interface state that changes from DOWN to UP, is only included in the Chassis grade if the link state is Up for "x" seconds. The default is 10 seconds.

Configuration:

The Link Preemption Mechanism is enabled by default with a preemption time of 10 seconds.

To configure the preemption time, run:

```
> fw ctl set int fwha_ch_if_preempt_time <preemp_time>
> update_conf_file fwkern.conf fwha_ch_if_preempt_time=<preemp_time>
```

Parameter	Description
<*preemp_time*>	Link Preemption Mechanism time Default: 10 seconds

Example

```
> fw ctl set int fwha_ch_if_preempt_time 20
> update_conf_file fwkern.conf fwha_ch_if_preempt_time=20
```

To disable Link Preemption Mechanism:

Run:

```
> fw ctl set int fwha_ch_if_preempt_time 0
> update_conf_file fwkern.conf fwha_ch_if_preempt_time=0
```

To make sure the preemption time value is correct:

Run:

```
> fw ctl get int fwha_ch_if_preempt_time
```

Chassis HA – Sync Lost Mechanism

The 61000/41000 Security System uses the Check Point proprietary *Cluster Control Protocol* (CCP) to send UDP control packets between two High Availability Chassis. When a sync interface fails, it is necessary to send SYNC_LOST to the other Chassis. The SYNC_LOST mechanism handles loss of connectivity between two Chassis on the Sync network.

To prevent the two Chassis from changing their states to Active, a SYNC_LOST CCP is sent over the non-sync interface (the Data Ports and Management interfaces) to the other Chassis. This causes the two Chassis to freeze their current states until connectivity between the two Chassis is restored. During the Sync Loss, the Standby Chassis does not change its state to Active until it stops receiving SYNC_LOST packets from the other Chassis.

The 61000/41000 Security System sends SYNC_LOST messages in this manner:

- For VSX environments - All interfaces of the VS0 context only
- For non-VSX environments - All Chassis interfaces

Sync Lost Mechanism is enabled by default.

To disable Sync Lost Mechanism:

Run:

```
> fw ctl set int fwha_ch_sync_lost_mechanism_enabled 0
> update_conf_file fwkern.conf fwha_ch_sync_lost_mechanism_enabled=0
```

To enable Sync Lost Mechanism

Run:

```
> fw ctl set int fwha_ch_sync_lost_mechanism_enabled 1
> update_conf_file fwkern.conf fwha_ch_sync_lost_mechanism_enabled=1
```

To check whether the mechanism is enabled:

Run:

```
> fw ctl get int fwha_ch_sync_lost_mechanism_enabled
```

(1-enabled, 0-disabled)

Managing Connection Synchronization (asg_sync_manager)

Use the asg_sync_manager utility to manage connection synchronization for High Availability. The configuration parameters include global settings and Sync Exception rules that control connection synchronization. Global synchronization settings apply to all connections, while Sync Exception rules apply only to specified connections.

This utility also controls SecureXL delayed synchronization parameters. When a connection is created from a SecureXL template, asg_sync_manager can set the period until it synchronizes to the Firewall.

To Define the Synchronization Level:

```
> asg_sync_manager

Please choose one of the following:
-----------------------------------
1) Print sync exceptions table
2) Add new sync exceptions rule
3) Delete old sync exception rule
4) Set sync between Chassis flag on / off
5) Set sync within local Chassis on / off
6) Configure sync between Chassis blades ratio
7) Set default delay notifications
8) Enable / Disable unicast sync
e) Exit
```

Press e to return to main menu.

To show synchronization properties, run: > asg stat -v

Working with Sync Rules

Sync Exceptions are rules, contained in the *Sync Exception Table* that define how synchronization works for specified connections or connection types. A Sync Exception rule applies and the specified action occurs if the connection matches all parameters in the rule definition. Rules are examined in sequence. The first matching rule applies.

These are the parameters of a Sync Exception rule:

Parameter	Description
`Idx`	Rule sequence number. Rules are applied in sequence, starting with rule 1.
`VS`	One or more Virtual System contexts.
`Source`	Source IP address and subnet mask.
`Destination`	Destination IP address and subnet mask.
`DPort`	Destination port.
`Ipp`	IP protocol number - typically http (6) or udp (17).
`ync`	Synchronization action: 0 = No synchronization 1 = Synchronize only to the local Chassis 2 = Synchronize only to the other (remote) Chassis 3 = Synchronize both Chassis 4 = Synchronize all SGMs
`Delay`	Time that it takes for connections created from templates to synchronize.

Sync Rule Options

Option	Description
Print Sync Exceptions table	Shows the Sync exception table Each entry in this table has these parameters: **1.** <5-tuple, including wild cards> **2.** Synchronization mode (none, within Chassis only, between Chassis only, both within, between Chassis and to all SGMs) **3.** SecureXL delayed synchronization value In addition, global synchronization values are displayed.
Add new Sync exceptions rule	Add a new rule to the sync exceptions table The user can hit enter at any stage to apply the default value. Specific rules allow the use of wildcards within 5-tuple. The new rule applies to new connections.
Delete old sync exception rule	Delete a rule from the sync exceptions table

Option	Description
Set sync between Chassis flag on / off	Global system setting - Enable synchronization connections to the Standby Chassis
Set sync within local Chassis flag on / off	Global system setting - Enable synchronization connections to the Active Chassis
Configure sync between Chassis SGMs ratio	Minimal SGMs ratio between active and backup Chassis for synchronization to occur If the number of UP SGMs in Standby Chassis is significantly low, compared to Active Chassis, synchronization might overload them. Default ratio for synchronization is 70% and it can be re-configured here. After configuration, the user can also choose to restore default settings.
Set default delay notifications	Default delayed synchronization setting are divided to HTTP related services (30) and all other services (5) You can reconfigure these settings here. **Note** - When you configure service delayed synchronization in SmartDashboard it overrides these settings.
Enable / Disable unicast sync	Enable or disable unicast sync (correction layer is enabled and disabled accordingly) and returns to the legacy synchronization scheme (synchronize connections to all SGMs). If you change this setting, you musts reboot of all SGMs.

Example

This example shows how to add a Sync Exception rule for all Virtual Systems that only synchronizes HTTP traffic from 3.3.3.0/24 to 4.4.4.0/24 on the Active Chassis.

```
Enter vs range: [default: 0]
>all
Enter source IP [0.0.0.0]:
>3.3.3.0
Enter source IP mask length [0]:
>24
Enter destination IP [0.0.0.0]:
>4.4.4.0
Enter destination IP mask length [0]:
>24
Enter destination port [0]:
>80
Enter IP protocol number (for example: tcp = 6, udp = 17):
>6
Enter the sync exception rule [3 - sync to all chassis]:
0 = no sync
1 = sync only to local chassis
2 = sync only to other chassis
3 = sync to all chassis
4 = sync to all SGMs
>1
Enter delay notification [30 - http, 5 - other]:
> 30
to insert new exception to vs 0-1,2: <3.3.3.0/24, 4.4.4.0/24, 80, 6> sync rule: 1, delay: 5 ? (y/n)
>y
```

The Sync exception table shows this information:

```
+-----------------------------------------------------------------------------+
|Sync exceptions table                                                        |
+-----+-------+---------------+-----------+-----+------+----+-----+-----------+
|Idx  |VS     |Source  |Mask  |Destination|Mask |DPort |Ipp |Sync |Delay      |
+-----+-------+---------------+-----------+-----+------+----+-----+-----------+
|1    |0-1,2  |0.0.0.0 |0     |0.0.0.0    |0    |53    |17  |0    |5          |
|2    |0-1,2  |3.3.3.0 |24    |4.4.4.0    |24   |80    |6   |1    |5          |
+-----+-------+---------------+-----------+-----+------+----+-----+-----------+
*Sync: 0=no sync, 1=sync only to local Chassis,2=sync only to other Chassis,3 = sync to all Chassis
**Delay: The time it takes for connections created from templates to synchronize

+---------------------------------------------------------------------------------------+
|Sync chassis                                                                           |
+----+----------------+----------------+--------------------+--------------------+------+
|VS  |Between chassis |Within chassis  |Unicast sync        |Correction layer    |Ratio |
+----+----------------+----------------+--------------------+--------------------+------+
|0   |Enabled         |Enabled         |Enabled             |Enabled             |50    |
|1   |Enabled         |Enabled         |Enabled             |Enabled             |50    |
|2   |Enabled         |Enabled         |Enabled             |Enabled             |50    |
+----+----------------+------*---------+--------------------+--------------------+------+

+---------------------------------------------------------+
|Delay                                                    |
+---------------+--------------------+--------------------+
|VS             |http                |default             |
+---------------+--------------------+--------------------+
|0              |30                  |5                   |
|1              |30                  |5                   |
|2              |30                  |5                   |
+---------------+--------------------+--------------------+

Enter vs range: [default: 0-1,2]
```

Working with SyncXL

SyncXL is a Check Point technology that makes sure that active connections are only synchronized to one SGM each on the Active Chassis and the Standby Chassis.

When an SGM or Chassis state changes, all SGMs update their counterpart SGMs. Synchronization is triggered automatically by these events:

- **SGM Failure** – Connections with a backup connection on an SGM are synchronized to a backup SGM
- **SGM Recovery** – The newly recovered SGM can be:
  - A backup for connections that are active on other SGMs
  - Active for connections before SGM failure
- **Chassis HA failover** – When the Active Chassis fails over to the Standby Chassis, a backup entry is defined for each connection it handles.

The SyncXL mechanism can be configured using the `asg_sync_manager` command. To learn more the `asg_sync_manager` command, see Defining the Synchronization Level (asg_sync_manager).

Standby Chassis/Active SGMs ratio:

To handle load and capacity, the Standby Chassis must have at least 50% of its SGMs in the UP state, compared with the Active Chassis. For example, if there are 10 SGMs that are UP on the Active Chassis, there must be at least five UP SGMs on the Standby Chassis. SyncXL is automatically disabled if this condition is not successful. You can change the ratio parameter.

To make sure that each active connection has backups on both Chassis in a Dual Chassis system

Run:

```
# asg_sync_manager
```

To see the last connection backup operation:

Run:

```
# asg_blade_stats

Last Iterator Statistics:
---------------------------------------------
Start time:                         Thu Sep 13 10:48:18 2012
Running time:                       0 Seconds
Status:                             Finished
Reason:                             Chassis ID 2 state was changed to STANDBY
Total connections iterated          38
Connections w/ sync action          0
```

To learn more about the `asg_blade_stats` command, see Showing SGM Forwarding Statistics (asg_blade_stats) (on page 110).

Notes:

- VoIP connections are synchronized to all SGMs
- Local connections (to/from the 61000/41000 Security System pseudo IP) are not synchronized
- SyncXL does not work on the Sync interface or the Management Interface

Setting Admin DOWN on First Join

You can configure the 61000/41000 Security System to automatically set a newly installed SGM in a Security Group, to the **Admin DOWN** state. The administrator can confirm that the SGM is configured correctly before it handles traffic.

Syntax

> set chassis high-availability down_on_first_join <*first_join*>

Parameter	Description
<*first_join*>	Sets whether Admin DOWN on First Join is enabled `0` - Admin DOWN on First Join is disabled `1` - Admin DOWN on First Join is enabled

To add a new SGM to a Security Group with Admin DOWN:

1. Run:
   `> set chassis high-availability down_on_first_join 1`
2. Install the new SGM and add it to the Security Group.
3. Set the SGM to the UP state:
   > asg sgm_admin -b <*sgm_ids*> up -p

Configuring a Unique IP Address For Each Chassis (UIPC)

In Dual-Chassis deployment:

- A heavy load on the Active Chassis can prevent you from creating a network connection to the SMO and working with management tasks.
- It can be necessary to have direct access to the Standby Chassis to troubleshoot a problem, such as a DOWN SGM. You cannot use the SMO to connect to the Standby Chassis.

You can assign a unique IP address to each Chassis to help resolve these issues. This adds an extra alias IP to the management interfaces on all SGMs.

- When there is a high load on the SMO, connect using the unique IP assigned to the Standby Chassis. The SGMs on the Standby Chassis are always UP and available to run `gclish` management commands.
- To connect directly to the Standby Chassis, use the Standby Chassis unique IP address.

Notes

- Only one SGM "owns" the UIPC task.
- The UIPC feature is disabled by default.
- If the 61000/41000 Security System is not managed by a management port, you can add the unique IP to one of the data ports. The connection to the unique IP reaches a specific blade based on the distribution configuration.

Use `set chassis id` command to assign a unique IP address to a Chassis.

Syntax

```
> set chassis id <chassis_id> general unique_ip <ip>
> delete chassis id <chassis_id> general unique_ip
> show chassis id <chassis_id> general unique_ip
```

Parameter	Description
<*chassis_id*>	Chassis ID Valid values: • `1` • `2` • `all`
<*ip*>	An alias IP address on the same network as one of the SGMs interfaces

Manual configuration

UIPC is automatically enabled after you run the configuration commands. You can also manually enable or disable it.

To manually enable UIPC:

```
> g_fw ctl set int fwha_uipc_enabled 1
```

To manually disable UIPC:

```
> g_fw ctl set int fwha_uipc_enabled 0
```

Example: Add a UIPC

```
> set chassis id 1 general unique_ip 172.16.6.186
Adding alias IP: 172.16.6.186 to chassis 1
Alias IP was added successfully
```

Example: Delete a UIPC

```
> delete chassis id 1 general unique_ip
Deleting alias IP 172.16.6.186 of chassis 1
Alias IP was deleted successfully
```

VSX Layer 2 Active/Active Mode

In the VSX Active/Active mode, both Chassis in a dual Chassis deployment handle connections. Connections between both Chassis are synchronized. Active/Active High Availability supports Layer 2 topologies.

Select High Availability Active/Active L2 mode when:

- An external device or protocol sends connections to both Chassis and makes the decision as to which Chassis is Active.
- Routing to Chassis is not symmetric. Packets for some connections can be sent to both Chassis.

Working with Link Aggregation (Interface Bonds)

In This Section:

Link Aggregation binds many physical interfaces together into one virtual interface called a Bond. This provides connection redundancy and traffic load sharing for better throughput. Each physical interface in a Bond is known as a slave interface.

Configuring Link Aggregation

Use the `add bonding group` command to create new Bond and add slave interfaces to it. Use `set bonding group` to configure parameters for an existing Bond.

This section shows the full syntax for these commands. The other sections in this chapter show only the syntax for the specified activities.

Syntax

```
add bonding group <bond_id>
add bonding group <bond_id> interface <slave_interface>
set bonding group <bond_id>
   [primary <slave_interface>]
   [mii-interval <value>]
   [up-delay <value>]
   [down-delay <value>]
   [mode <value>]
   [lacp-rate  <value> ]
   [xmit-hash-policy <value>]
   [abxor-threshold <value>]
```

Parameter	Description
<bond_id>	Bond identifier, an integer between 1 and 1024.
`interface` *<slave_interface>*	Slave interface name.
`primary` *<slave_interface>*	Sets the primary slave interface. This parameter is applicable to the `active-backup` mode only.
`mii-interval` *<interval>*	Frequency (in ms) that the system polls the Media Independent Interface (MII) to get status. **Valid values** = 1-5000 ms. Default = 100 ms.

Parameter	Description
`up-delay` <*value*> `down-delay` <*value*>	Wait time (in ms) before the system confirms that a slave interface is UP or DOWN. **Valid values** = 1-5000 ms. Default = 200 ms.
`mode` <*value*>	Bond interface mode: • **active-backup** - Selects the Primary slave interface as the Active slave interface. If the Primary slave interface goes down, it fails over to a different slave interface. • **xor** - All UP slave interfaces are Active for Load Sharing. Traffic is assigned to Active interfaces based on the transmit hash policy (Layer2 or Layer3+4). • **8023AD** -Dynamically uses Active slave interfaces to share the traffic load based on the LACP protocol. This protocol uses full interface monitoring between the Security Gateway and a switch. • **abxor** - Slave interfaces are assigned to sub-groups called *bundles*. Only one bundle is active at a time. All slave interfaces in the active bundle share the traffic load. The system assigns traffic to all interfaces in the active bundle based on the defined transmit hash policy. **Note** - The Round-Robin option is not supported on the 61000/41000 Security System.
`lacp-rate` <*value*>	LACPDU packet transmission rate: **slow** - Request LACPDU every 30 seconds **fast** - Request LACPDU every 1 second. This parameter is applicable to the 8023AD mode only.
`xmit-hash-policy` <*value*>	Methodology for slave interface selection based on the TCP/IP layer. **layer2** - Use XOR of hardware MAC addresses. **layer 3+4** - Use upper layer protocol information. This parameter is applicable to the XOR and ABXOR modes only.
`abxor-threshold` <*value*>	Minimum number of slave interfaces that must be UP for a bundle to be Active. **Valid values =** 1-8 interfaces. Default = 3 interfaces. This parameter is applicable to the ABXOR mode only.

Examples:

```
> add bonding group 4 interface eth1-03
```

This command creates a new Bond (`bond4`) with one slave interface.

```
> add bonding group 4 interface eth2-03
```

This command adds another slave interface to bond4.

```
> set bonding group 4 mode xor down-delay 300 mii-interval 100
```

This command changes the `mode`, `down-delay` and `mii_interval` parameters for bond4.

Creating a New Bond and Adding Slave Interfaces

Use the `add bonding group` command to create a new Bond and to add slave interfaces to an existing Bond. You must run this command once for each slave that you add to a Bond.

Syntax

```
add bonding group <bond_id> [interface <slave_interface>]
```

Parameter	Description
<*bond_id*>	Bond identifier, an integer between 1 and 1024
<*slave_interface*>	Slave interface name

Examples

```
> add bonding group 4 interface eth1-02
```

This creates a new Bond with slave interface `eth1-02.`

```
> add bonding group 4 interface eth2-02
```

This command adds another slave interface to bond4.

Setting a Bonding Mode

Use `set bonding group` command to change the bond mode. This section shows only the options related to interface bond modes.

Syntax

```
set bonding group <bond_id> mode <bond_mode>
```

Parameter	Value
<*bond_id*>	Bond identifier, an integer between 0 and 1024.

Parameter	Value
<*bond_mode*>	Bond interface mode: • **active-backup** - Selects the Primary slave interface as the Active slave interface. If the Primary slave interface goes down, it fails over to a different slave interface. • **xor** - All UP slave interfaces are Active for Load Sharing. Traffic is assigned to Active interfaces based on the transmit hash policy (Layer2 or Layer3+4). • **8023AD** -Dynamically uses Active slave interfaces to share the traffic load based on the LACP protocol. This protocol uses full interface monitoring between the Security Gateway and a switch. • **abxor** - Slave interfaces are assigned to sub-groups called *bundles*. Only one bundle is active at a time. All slave interfaces in the active bundle share the traffic load. The system assigns traffic to all interfaces in the active bundle based on the defined transmit hash policy. **Note** - The Round-Robin option is not supported on the 61000/41000 Security System.

Example

```
> set bonding group 4 mode 8023AD
1_01:
success
1_02:
success
1_03:
success
2_01:
success
2_03:
success
```

Setting the Polling interval

Use this command to set the polling interval for a Bond. This section shows only the parameters related to the polling interval.

Syntax

```
> set bonding group <bond_id> mii-interval <interval>
```

Parameter	Description
<*bond_id*>	Bond ID
`mii-interval`<*interval*>	Frequency (in ms) that the system polls the Media Independent Interface (MII) to get status. **Valid values** = 1-5000 ms. Default = 200 ms.

Setting a Bond Interface On or Off

Use this command to turn the Bond interface on or off after you create and configure it.

Syntax

```
> set interface <bond_interface> state on
```

Parameter	Description
<bond_interface>	Slave interface name

Example

```
> set interface bond4 state on
```

Removing Slave Interfaces

Use this command to remove a slave interface from a Bond.

Syntax

```
> delete bonding group <bond_id> interface <slave_interface>
```

Parameter	Description
<bond_id>	Bond identifier, an integer between 0 and 1024.
<slave_interface>	Slave interface name

Example

```
> delete bonding group 1 interface eth1-02
```

Deleting a Bond

Use this command to delete a Bond.

⚠ **Important** - You must delete all slave interfaces in a Bond before you can delete that Bond.

Syntax

```
> delete bonding group <bond_id>
```

Parameter	Description
<bond_id>	Bond identifier, an integer between 1 and 1024.

Working with the ABXOR Bonds

R76SP.40 supports ABXOR Bonds, which provide slave interface redundancy and load sharing. An ABXOR Bond is divided into two or more sub-groups, known as *bundles*. Each bundle can have up to eight slave interfaces.

Bundles provide Active/Backup redundancy, where only one bundle is active at any given time. The system selects the active bundle based on these rules:

- The active bundle (bundle 1) has the lowest index and at least as many active (UP) slave interfaces as the *abxor threshold* defined for the Bond. The *abxor threshold* is the minimum number of active slave interfaces necessary for a bundle to become active. You define an abxor threshold for each Bond.
- If no bundle has the minimum number of active slave interfaces, the bundle with the most active slave interfaces becomes the active bundle.

For example, a bundle has four slave interfaces and the Bond has an abxor threshold of three. The active bundle must have at least three active interfaces. If no bundle has the minimum quantity of active interfaces, the bundle with the most UP interfaces becomes active.

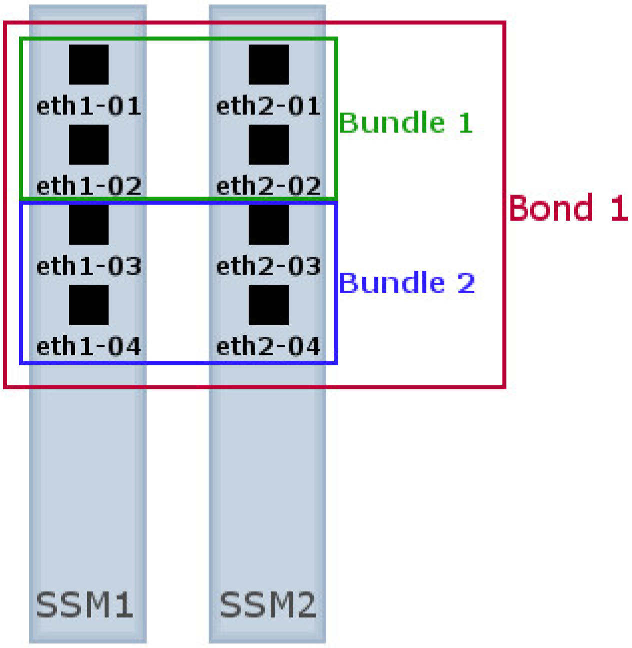

You can use abxor bonds with a different switch connected to each bundle. This provides both SSM and switch redundancy with Load Sharing. In the example above, each bundle connects to a different switch and has slave interfaces from both SSMs. If one of the switches and/or one of the SSMs fail, there is no traffic interruption.

Configuring ABXOR

To create an ABXOR Bond:

1. Create a new Bond ("Configuring Link Aggregation " on page 204).
2. Add slave interfaces ("Creating a New Bond and Adding Slave Interfaces" on page 206) to the new Bond.
3. Create the bundles.

   Run: `add bonding group` <*bond_id*> `bundle` <*bundle_id*>

   Note: `Bundle ID cannot be 0. The legal values are 1..2`
4. Assign slave interfaces to each bundle.

   Run for each slave interface:

   `add bonding group` <*bond_id*> `bundle` <*bundle_id*> `interface` <*slave_interface*>

5. Set the Bond mode to abxor:

   `set bonding group` <*bond_id*> `mode abxor`

6. Set the abxor-threshold:

   `set bonding group` <*bond_id*> `abxor-threshold` <*value*>

   The <*value*> can be from one to eight and the default value is three.

7. Set the minimum number of slave interfaces in a bond:

   ```
   set chassis high-availability bond bond1 min_slaves 2
   ```

   Notes:

   - The default value for `min_slaves` is 1.
   - In order to keep standard throughput, the number of slave interfaces has to equal the abxor-threshold.

To delete an ABXOR Bond:

Important - You must delete all slave interfaces in a Bond before you can delete that Bond.

Syntax

```
> delete bonding group <bond_id> bundle <bundle_id> interface <slave_interface>
> delete bonding group <bond_id> bundle <bundle_id>
> delete bonding group <bond_id> interface <slave_interface>
> delete bonding group <bond_id>
```

Parameter	Description
<*bond_id*>	Bond identifier, an integer between 1 and 1024.
<*slave_interface*>	Name of slave interface.
<*bundle_id*>	Bundle identifier, an integer between 1 and 2.

Working with Management Aggregation

Management Aggregation (MAGG) is a High Availability and Load Sharing solution for management interfaces. You can create bonds that link physical management interfaces together as one virtual interface.

To create a new Management Bond, run one of these commands:

- `add bonding group` <*bond_id*> `mgmt`
- `add bonding group` <*bond_id*> `mgmt interface` <*mgmt_interface_name*>

The second command creates the Management Bond and adds a slave management interface in one step.

Notes:

- Only use the `mgmt` parameter when you create a new management bond. For all other configurations, use the standard commands and parameters without the `mgmt` parameter. For more information, see Configuring L*ink A*ggregation ("Working with Link Aggregation (Interface Bonds)" on page 204).

- A bond is created for data or management, but not for both.
- We recommend that you do not mix 1G and 10G management interfaces in a bond.
- Unlike a regular bond, you cannot delete a bond slave if you configure it in a bundle.

Limitations

- You cannot include Mgmt4 interfaces.   For example, `eth0-Mgmt4`.
- You cannot configure VLANs.
- After you create the VSX object in S*martDas*hboard, you cannot configure a bond on the chassis management interface. Configure the MAGG interface after setup.
- Use XOR\HA configuration. Upon chassis failover, the LACP mode may cause up to a 1 minute outage on the management interface. Regular traffic is not affected.
- Only eth1-Mgmt1 and eth-2-Mgmt1 can be added to MAGG.

Unique IP address per Chassis (UIPC) is not supported when Management Aggregations are enabled. See sk107955 https://supportcenter.checkpoint.com/supportcenter/portal?eventSubmit_doGoviewsolutiondetails=&solutionid=sk107955

Example

This example creates a management bond with two slaves:

```
> add bonding group 7 mgmt
> add bonding group 7 mgmt interface eth1-Mgmt3
> add bonding group 7 mgmt interface eth2-Mgmt3
> set bonding group 7 mode xor
> set interface magg7 state on
> set interface magg7 ipv4-address X.X.X.X mask-length X

> show bonding group 7
1_01:
Bond Configuration
    xmit-hash-policy layer2
    down-delay 200
    primary Not configured
    lacp-rate Not configured
    mode xor
    up-delay 200
    mii-interval 100
    abxor-threshold 3
    type mgmt
    Bond Interfaces
        eth1-Mgmt3
        eth2-Mgmt3

1_02:
Bond Configuration
    xmit-hash-policy layer2
    down-delay 200
    primary Not configured
    lacp-rate Not configured
    mode xor
    up-delay 200
    mii-interval 100
    abxor-threshold 3
    type mgmt
    Bond Interfaces
        eth1-Mgmt3
        eth2-Mgmt3
```

Example:

This example creates a new interface name:

```
> set interface magg4 state on
> set interface magg4 ipv4-address X.X.X.X mask-length X
```

Working with Sync Bonds

The *Sync Interface* is a special Bond used for these and other synchronization and control tasks:

- To send CCP packets to other SGMs
- To share policies and configuration files amongst SGMs
- Connection state synchronization
- Packet forwarding
- Firewall synchronization
- Daemon synchronization
- Monitor and control commands (gclish, asg) that get information from many SGMs

A Sync Interface has one or two slave interfaces, based on the number of SSMs in each Chassis. The system automatically creates the slave interfaces based on this algorithm:

Notes:

- A Chassis with one SSM always uses `eth1-Sync`. `eth2-Sync`s not assigned.
- A chassis with two or more SSMs always uses these slave interfaces:
  - eth1-Sync
  - eth2-Sync
  - Sync ports on SSM3 and SSM4 are not used when there are more than two SSMs
- The system automatically creates the Sync Bond during installation and assigns these IP addresses:
  - SGM1_1 - 192.0.2.1
  - SGM1_12 - 192.0.2.14
  - SGM2_1 - 192.0.2.15
  - SGM2_12 - 192.0.2.28

  No manual configuration is necessary.
- The system automatically assigns the Sync slave to port 8 on SSM1 and SSM2.
- The system sets the Sync slave port speed to 10 Gb by default. We recommend that your do not change this parameter.
- Sync Bonds support both LR and SR transceivers.
- The Sync Bond uses the XOR mode ("Setting a Bonding Mode " on page 206).
- The default network (192.0.2.X) is defined by the applicable RFC as a private network for documentation. It is unlikely to cause collisions with user networks.

Limitations:

- LACP is not supported.
- VLANs are not supported for Sync slave interfaces.
- A Sync Bond can have up to two slave interfaces.
- Sync ports cannot be changed to data ports. This is true for both Single and Dual Chassis systems.

Sync Lost

Sync Lost is a Check Point feature that makes sure that both Chassis do not become Active if the Sync network fails. The system sends special SYNC_LOST packets to the other Chassis over the data and management interfaces. This action prevents a state change on both Chassis until the Sync network is restored.

The Sync Lost mechanism is enabled by default.

Connecting Physical Cables

Single Chassis Systems

It is not necessary to connect Sync ports in a single Chassis system, because communication between SGMs are handled internally by the Chassis infrastructure.

Dual Chassis system with a Cross cable

- `eth1-Sync` in chassis1 connects to `eth1-Sync` in chassis2.
- `eth2-Sync` in chassis1 is connects to `eth2-Sync` in chassis2.

Dual Chassis system with the Sync Bond on the switch

- Configure all ports in the switch in the same VLAN broadcast domain.
- Each switch configures a bond for its chassis.

Dual chassis system without a Sync Bond on a the switch

- `eth1-Sync` communicates over VLAN X.
- `eth2-Sync` communicates over VLAN Y.
- Configure switches for VLAN access on each related port.

  Configure "Link State Tracking" (Cisco), or an equivalent mechanism, so that Sync port peers go down on both Chassis after a failure on one Chassis. If you do not do this, and the Sync port fails on one Chassis, the related peer Sync port stays UP.

Working with VSX

In This Section:

Provisioning VSX

Create VSX objects with one of these procedures:

- Create new Security Gateways, Virtual Systems and other virtual objects in SmartDashboard.
- Run `vsx_util reconfigure` from the management server.

Notes:

- The SMO reboots automatically when you create a new Virtual System.
- Before you start one of these procedures, make sure that the SMO is the only SGM in the security group. After successful configuration, you can add more SGMs to the security group.

Configuring 64 Bit Virtual System Support

You can configure the 61000/41000 Security System to run `fwk` as a 64 bit process. This lets VSX Virtual Systems use more than 4GB of RAM, which significantly increases the concurrent connection capacity for each Virtual System.

Use the `vs_bits` command to configure `fwk` to run in the 64 or 32 bit mode. The system automatically reboots when you run the command.

Syntax:

```
vs_bits [-stat | 32 | 64 ]
```

Parameter	Description
`stat`	Shows the current `fwk` mode.
32	Run `fwk` in the 32 bit mode.
64	Run `fwk` in the 64 bit mode.

Examples:

This example changes the `fwk` mode to 64 bits.

```
vs_bits  64
```

This example shows all the `fwk` mode for all SGMs.

```
# vs_bits -stat
```

```
1_01:
64
1_02:
64
1_03:
64
1_04:
64
2_01:
64
2_02:
64
2_03:
64
2_04:
64
```

Known limitation:

This feature only works on a 64 bit operating system.

Creating a new VSX Gateway

This section shows you how to create a new VSX Gateway with the **VSX Gateway Wizard**. After you finish the VSX Gateway Wizard, you can configure the VSX Gateway definition with SmartDashboard. For example, you can add or delete interfaces, or configure existing interfaces to support VLANs.

Before starting, you must make sure that the SMO is the only SGM in the group.

To start the VSX Gateway W**izard**:

1. Open SmartDashboard. If you are using Multi-Domain Security Management, open SmartDashboard from the Domain Management Server.
2. From the **Network Objects** tree, right-click on **Check Point** and select **VSX** > **Gateway**.
   The **General Properties** page opens.
3. Do the instructions on the screen.

Configuring VSX Gateway General Properties

The **General Properties** page contains basic identification properties for VSX Gateways.

- **VSX Gateway Name**: Unique, alphanumeric for the VSX Gateway. The name cannot contain spaces or special characters except the underscore.
- **VSX Gateway IP Address**: Management interface IP address.
- **VSX Gateway Version**: Select the VSX version installed on the VSX Gateway from the drop-down list.

Selecting Virtual Systems Creation Templates

The **Creation Templates** page lets you provision predefined, default topology and routing definitions to Virtual Systems. This makes sure Virtual Systems are consistent and makes the definition process faster. You always have the option to override the default creation template when you create or change a Virtual System.

The Creation Templates are:

- **Shared Interface** - Not supported for the 61000/41000 Security System.
- **Separate Interfaces:** Virtual Systems use their own separate internal and external interfaces. This template creates a Dedicated Management Interface (DMI) by default.
- **Custom Configuration:** Define Virtual System, Virtual Switch, and Interface configurations.

For this example, choose **Custom configuration**.

Establishing SIC Trust

Troubleshooting SIC Trust Initialization Problems

If SIC trust did not successfully connect, click **Check SIC Status**. The most common reasons for an unsuccessful connection are:

- Entering an incorrect activation key
- Connectivity problems between the management server and the VSX Gateway

To troubleshoot and resolve SIC initialization problems:

- Re-enter and re-confirm the activation key.
- Confirm that the IP address defined in **General Properties** is correct.
- Ping the management server to verify connectivity. Resolve connectivity issues.
- From the VSX Gateway command line, use `cpconfig` to re-initialize SIC. When this has finished, click **Reset** in the wizard and re-enter the activation key.

For more about resolving SIC initialization, see sk65385 http://supportcontent.checkpoint.com/solutions?id=sk65385.

Initializing SIC Trust

Initialize Secure Internal Communication (SIC) trust between the VSX Gateway and the management server. The gateway and server cannot communicate without trust.

To you create a VSX Gateway:

1. Enter and confirm the activation key from the installation wizard setup program.
2. Click **Initialize**.

   If you entered the correct activation key, the **Trust state** changes to **Trust established**.

Note: To reset SIC trust for a VSX Gateway or Virtual System, you must use a console connection. Do not use an SSH client.

Defining Physical Interfaces

In the **VSX Gateway Interfaces** window, you can define physical interfaces as VLAN trunks. The page shows the interfaces currently defined on the VSX Gateway.

To define an interface as a VLAN trunk, select **VLAN Trunk** for the interface.

You can define VLAN trunks another time. For this example, choose **Next**.

Virtual Network Device Configuration

If you chose the **Custom Configuration** option, the **Virtual Network Device Configuration** window opens.

The options in this window are not supported for the 61000/41000 Security System.

Click **Next**.

VSX Gateway Management

In the **VSX Gateway Management** window, define security policy rules that protect the VSX Gateway. This policy is installed automatically on the new VSX Gateway.

Note - This policy applies only to traffic destined for the VSX Gateway. Traffic destined for Virtual Systems, other virtual devices, external networks, and internal networks, is not affected by this policy.

The security policy consists of predefined rules for these services:

- **UDP -** SNMP requests
- **TCP -** SSH traffic
- **ICMP -** Echo-request (ping)
- **TCP -** HTTPS traffic

To modify and configure the Gateway Security Policy you can:

- Select **Allow** to pass traffic on the selected services, or clear this option to block traffic. By default, all services are blocked.

  For example, to ping the gateway from the management server, **Allow** ICMP echo-request traffic.
- Click the arrow and select a **Source Object** from the list or **New Source Object** for a new source.

  The default value is *Any

You can modify security policy rules that protect the VSX Gateway at any time.

To complete the Virtual System wizard:

1. Click **Next**.
2. Click **Finish**.

   It can take several minutes to complete.

If this ends unsuccessfully, click **View Report** to see the error messages.

After the VSX wizard has finished successfully, other SGMs can be added to security group.

Virtual System

After you create a Virtual System on a 61000/41000 Security System, we recommend that you limit the maximum number of concurrent connections to no more than 500,000.

To limit the maximum number of concurrent connections:

1. In SmartDashboard, double-click Virtual System.
2. Click **Optimizations** > **Calculate the maximum limit for concurrent connections.**

3. Select **Manually**.
4. Enter `500000`.
5. Click **OK**.

Reconfigure (vsx_util reconfigure)

Use `vsx_util reconfigure` on the management server to restore a VSX configuration to a newly installed gateway.

Syntax

```
> vsx_util reconfigure
```

Input

- VSX Gateway name
- SIC activation key assigned to the Security Management Server or Domain Management Server
- Retype to confirm the SIC activation key

Notes

- This command is also useful for restoring a gateway or cluster member after a system failure.
- Run the command and follow the instructions on the screen.
- A new gateway must have the same hardware specifications and configuration as its replacement and other cluster members. Most importantly, it must have the same number of interfaces (or more) and the same management IP address.
- The new or replacement machine must be a new installation. You cannot use a machine with a previous VSX configuration.

Working with VSLS

VSLS is a Virtual System Load Sharing solution for the 61000/41000 Security System that uses both Chassis to handle traffic. Each Virtual System works as an independent cluster. For each Virtual System, one Chassis is Active and the other Chassis becomes the Standby. The selection of the Active Chassis is based on interface availability, SGM availability, and Virtual System stability.

A Virtual System in the DOWN state fails over to the Standby Virtual System in the other Chassis. By default, a Virtual System in the DOWN state **does not** put the SGM in the DOWN state. Because of this, there is no effect on other Virtual System states.

The SGM continues to receive traffic from the SSM. This behavior is different from Chassis High Availability, where a Virtual System in the DOWN state causes the SGM to go DOWN.

Notes:

- If VS0 goes DOWN, its related SGM also goes DOWN.
- Run this command (in gclish) to change the VSLS behavior so that a Virtual System in the DOWN state causes the SGM to go DOWN :

  ```
  > g_update_conf_file fwkern.conf fwha_mbs_vsls_only_vs0_decide_state=0
  ```

  Reboot the Chassis.

  This behavior is now the same as for standard Chassis High Availability.

- When an SGM contains a DOWN Virtual System, the SMO and Chassis Monitor tasks move to a different valid SGM. Because these tasks can move to a different SGM, connections to the Virtual Systems can become disconnected.
- We recommend that you work with UIPC. This is because the UIPC task does not move to a different SGM.

Activating Chassis VSLS

To use Chassis VSLS features, you must first activate the **Chassis VSLS High Availability mode**.

To activate Chassis VSLS:

Run: `> set chassis high-availability mode 4`

Note - This command can cause Chassis failover.

Selecting the Active Chassis for a Virtual System

VSLS dynamically assigns an Active Chassis to each Virtual System based on criteria in this order of priority:

1. **Availability of functional interfaces for the Virtual System**

   VSLS selects the Chassis with the most connected interfaces to be the Active Chassis.
2. **Availability of UP SGMs**

   If both Chassis have the same number of connected interfaces, VSLS uses this ratio to select the Active Chassis:

   ```
   SGM Ratio = Fewest_UP_SGMS/Most_UP_SGMS
   ```

   If the SGM Ratio is less than the predefined threshold (default=50%), VSLS selects the Chassis with the most available SGMs. If the SGM Ratio is greater or equal to the threshold, VSLS does not select an Active Chassis based on SGM availability.

   Example:
   Chassis 1 has two UP SGMs and Chassis 2 has five UP SGMS. The ratio is 2/5 (40%), which is less that the default threshold of 50%. VSLS selects Chassis 2 as the Primary Chassis.
3. **Virtual System with a problem**

   When a Virtual System fails, VSLS automatically fails over to the related Virtual System on the other Chassis, which becomes the Active Chassis.
4. **Primary Chassis**

   If none of the above criteria causes VSLS to select an Active Chassis, the Primary Chassis automatically becomes the Active Chassis.

To change the SGM threshold value:

Run:

```
> set chassis vsls sgm_ratio <percent_value>
```

Virtual System Failover

With VSLS, a Virtual System can failover to the Standby Chassis independently of the other Virtual Systems. When VSLS selects a different Chassis for a Virtual System based on the selection criteria, only that Virtual System fails over. There is no effect on the other Virtual Systems.

Virtual System failover works the same way as a regular layer2/layer3 failover. The Virtual System sends GARP/NDS packets in layer 3 and MAC learning packets in Layer 2.

Example:

For VS1, Chassis 2 is both the Active and the Primary Chassis. If an interface used by VS1 on Chassis 2 is disconnected, VS1 fails over to Chassis 1 based on the dynamic selection procedure. When the port is reconnected, VS1 fails back to Chassis2.

SGM Failover

When an SGM fails, it no longer receives traffic. When a single Virtual System fails on an SGM, this Virtual System can do a Virtual System Chassis Failover. If a Virtual System Chassis failover does not occur, the failed Virtual System on the SGM continues to receive traffic.

Configuring the VSLS Primary Chassis

When you create a new Virtual System, VSLS automatically assigns a Primary Chassis based on the system default. You can change the default Primary Chassis when it is necessary. When you change the default Primary Chassis, it changes for all Virtual Systems that do not have a manually defined Primary Chassis. This can cause Virtual Systems to failover to a different Active Chassis.

You can manually define the Primary Chassis for specified Virtual Systems. Manually defined Virtual Systems do not change their Primary Chassis when you change the default Primary Chassis.

To change the system default Primary Chassis:

1. Change the context to VS0:
   ```
   > set virtual-system 0
   ```
2. Run:
   ```
   > set chassis vsls system primary_chassis <option>
   ```
   <*option*> is an integer between 0 and 2:

   0 - Automatic (VSLS automatically assigns the Primary Chassis)

   1 - Define Chassis 1 as the default Primary Chassis

   2 - Define Chassis 2 as the default Primary Chassis

To manually define a Primary Chassis for a Virtual System:

1. Go to the Virtual System context to be changed.
   ```
   > set virtual-system <vsid>
   ```
2. Run:
   ```
   > set chassis vsls vs primary_chassis <option>
   ```
   <*option*> is an integer between 0 and 2:

   0 - Use the system default Primary Chassis

   1 - Define Chassis 1 as the Primary Chassis

   2 - Define Chassis 2 as the Primary Chassis

To show the Primary Chassis for all Virtual Systems:

Run:

```
> show configuration vsls
```

```
-------------------------------------------------
| Default Mode:         Automatic               |
| Virtual Systems:      10                      |
-------------------------------------------------
| VS   | VS-Name         | Chassis 1 | Chassis 2 |
-------------------------------------------------
| 0    | 61000-VSLS      | Default   |           |
| 1    | VS1             | Manual    |           |
| 2    | VS2             | Default   |           |
| 3    | VS3             |           | Default   |
| 4    | VS4             | Default   |           |
| 5    | VS5             |           | Default   |
| 6    | VS6             | Default   |           |
| 7    | VS7             |           | Default   |
| 8    | VS8             | Default   |           |
| 9    | VS9             |           | Manual    |
-------------------------------------------------
| Total:                 | 6         | 4         |
-------------------------------------------------
```

This example shows that:

- The default Primary Chassis mode is Automatic (0).
- The deployment has 10 Virtual Systems including VS0.
- VS1 and VS9 have manually assigned Primary Chassis (Chassis 1 and Chassis 2 respectively)
- All others use the default Primary Chassis, which are assigned to different chassis to effectively distribute the traffic load.
- Chassis 1 is configured as the Primary Chassis for VS0, VS1, VS2, VS4, VS6, and VS8.
- Chassis 2 is configured as the Primary Chassis for VS3, VS5, VS7, and VS9.

Monitoring VSLS

Using asg stat

Use the `asg stat` without arguments to see general VSX and system information. You can run this command from gclish or the Expert mode.

```
> asg stat
--------------------------------------------------------------------------
|                        VSX System Status - 61000                        |
--------------------------------------------------------------------------
| Chassis Mode               | VSLS                                       |
| Up time                    | 21:36:26 hours                             |
| SGMs                       | 5 / 8                                      |
| Virtual Systems            | 10                                         |
| Version                    | R76SP.10_VSLS (Build Number 20)            |
--------------------------------------------------------------------------
|                            | Chassis 1              | Chassis 2         |
--------------------------------------------------------------------------
|                            | UP / Required          | UP / Required     |
| SGMs                       |  4 / 4                 |  1 / 4  (!)       |
| Ports                      |  3 / 3                 |  3 / 3            |
| Bonds                      |  2 / 2                 |  2 / 2            |
| Fans                       |  3 / 4  (!)            |  6 / 6            |
| SSMs                       |  2 / 2                 |  2 / 2            |
| CMMs                       |  1 / 2  (!)            |  1 / 2  (!)       |
| Power Supplies             |  4 / 5  (!)            |  3 / 5  (!)       |
--------------------------------------------------------------------------
```

The output shows that:

- System is running in VSLS mode.
- System has 10 Virtual Systems configured, including VS0.
- System has eight SGMs in Security Group.
- System has five SGMs in UP state.
- All SGMs on Chassis 1 are UP.
- Only one SGM on Chassis 2 is UP.

Using asg stat vs all

Use the `asg stat vs all` command to see which Virtual Systems are Active on each Chassis and their health status. You can run this command from `gclish` or the Expert Mode.

```
> asg stat vs all

Output: VSLS
-----------------------------------------------------------------------
| VSX System Status - 61000                                           |
-----------------------------------------------------------------------
| Chassis Mode                | VSLS                                  |
| Up time                     | 4 days, 16:05:08 hours                |
| SGMs                        | 1 / 3 (!)                             |
| Virtual Systems             | 4                                     |
| Version                     | R76SP.40 (Build Number 2)             |
-----------------------------------------------------------------------
| VSID   | VS Type & Name     | Chassis 1   | Chassis 2   | Health    |
-----------------------------------------------------------------------
| 0      | V 61000-VSLS       | DOWN    (P) | ACTIVE      | Problem   |
| 1      | S VS1              | DOWN        | ACTIVE  (P) | Problem   |
| 2      | S VS2              | DOWN    (P) | ACTIVE      | Problem   |
-----------------------------------------------------------------------
| Active Virtual Systems      | 0           | 3           |           |
-----------------------------------------------------------------------
| Errors:                                                             |
| VSID's not on Primary chassis: 0 2                                  |
-----------------------------------------------------------------------
| Synchronization                                                     |
|     Within  chassis:           Enabled    (Default)                 |
|     Between chassis:           Disabled   (Auto)                    |
|           Reason: Chassis states doesn't allow Sync between chassis |
|     Exception Rules:                      (Default)                 |
-----------------------------------------------------------------------

(P) - VS Primary Chassis
```

```
Output: Virtual System HA


------------------------------------------------------------------------

| VSX System Status - 61000                                            |

------------------------------------------------------------------------

| Chassis Mode                | Active Up                              |

| Up time                     | 4 days, 12:04:35 hours                 |

| SGMs                        | 19/24 (!)                              |
| Virtual Systems             | 103                                  |
| Version                     | R76SP.40 (Build Number 53)           |
----------------------------------------------------------------------
| VSID    | VS Type & Name    | Chassis 1    | Chassis 2    | Health   |
----------------------------------------------------------------------
| 0       | V na-core-gw      | STANDBY      | ACTIVE       | Problem  |
| 1       | S MEX-T2-VS       | STANDBY      | ACTIVE       | Problem  |
| 2       | S EQU-T2-VS       | STANDBY      | ACTIVE       | Problem  |
| 3       | S CAN-T2-VS       | STANDBY      | ACTIVE       | Problem  |
| 4       | S SHR-T2-VS       | STANDBY      | ACTIVE       | Problem  |
| 5       | S EQU-T3-VS       | STANDBY      | ACTIVE       | Problem  |
| 6       | S EXTRANET-VS     | STANDBY      | ACTIVE       | Problem  |
| 7       | S MEX-T3-VS       | STANDBY      | ACTIVE       | Problem  |
| 8       | S SHR-T3-VS       | STANDBY      | ACTIVE       | Problem  |
| 9       | S FUSION-VS       | STANDBY      | ACTIVE       | Problem  |
| 10      | S IPT-VS          | STANDBY      | ACTIVE       | Problem  |
| 11      | S JP-T3-VS        | STANDBY      | ACTIVE       | Problem  |
| 12      | S AB-T3-VS        | STANDBY      | ACTIVE       | Problem  |
| 13      | S MGMT-VS         | STANDBY      | ACTIVE       | Problem  |
| 14      | S VENDOR-VS       | STANDBY      | ACTIVE       | Problem  |
| 15      | S ALM-T3-VS       | STANDBY      | ACTIVE       | Problem  |
| 16      | B VSB_1_TAP       | STANDBY      | ACTIVE       | Problem  |
| 19      | B VSB_2_Access    | STANDBY      | ACTIVE       | Problem  |
| 22      | S VS_Barakeo      | STANDBY      | ACTIVE       | Problem  |
| 25      | B VSB_Packet_Bro  | STANDBY      | ACTIVE       | Problem  |
| 26      | B VSB_Packet_Bro  | STANDBY      | ACTIVE       | Problem  |
| 27      | S VS027           | STANDBY      | ACTIVE       | Problem  |
| 28      | S VS028           | STANDBY      | ACTIVE       | Problem  |
| 29      | S VS029           | STANDBY      | ACTIVE       | Problem  |
| 30      | S VS030           | STANDBY      | ACTIVE       | Problem  |
| 31      | S VS031           | STANDBY      | ACTIVE       | Problem  |
| 32      | S VS032           | STANDBY      | ACTIVE       | Problem  |
| 33      | S VS033           | STANDBY      | ACTIVE       | Problem  |
| 34      | S VS034           | STANDBY      | ACTIVE       | Problem  |
| 35      | S VS035           | STANDBY      | ACTIVE       | Problem  |
| 36      | S VS036           | STANDBY      | ACTIVE       | Problem  |
| 37      | S VS037           | STANDBY      | ACTIVE       | Problem  |
| 38      | S VS038           | STANDBY      | ACTIVE       | Problem  |
| 39      | S VS039           | STANDBY      | ACTIVE       | Problem  |
| 40      | S VS040           | STANDBY      | ACTIVE       | Problem  |
| 41     | S VS041            | STANDBY     | ACTIVE      | Problem    |
| 42     | S VS042            | STANDBY     | ACTIVE      | Problem    |
| 43     | S VS043            | STANDBY     | ACTIVE      | Problem    |
| 44     | S VS044            | STANDBY     | ACTIVE      | Problem    |
| 45     | S VS045            | STANDBY     | ACTIVE      | Problem    |
| 46     | S VS046            | STANDBY     | ACTIVE      | Problem    |
| 47     | S VS047            | STANDBY     | ACTIVE      | Problem    |
| 48     | S VS048            | STANDBY     | ACTIVE      | Problem    |
| 49     | S VS049            | STANDBY     | ACTIVE      | Problem    |
| 50     | S VS050            | STANDBY     | ACTIVE      | Problem    |
| 51     | S VS051            | STANDBY     | ACTIVE      | Problem    |
| 52     | S VS052            | STANDBY     | ACTIVE      | Problem    |
| 53     | S VS053            | STANDBY     | ACTIVE      | Problem    |
| 54     | S VS054            | STANDBY     | ACTIVE      | Problem    |
| 55     | S VS055            | STANDBY     | ACTIVE      | Problem    |
| 56     | S VS056            | STANDBY     | ACTIVE      | Problem    |
| 57     | S VS057            | STANDBY     | ACTIVE      | Problem    |
| 58     | S VS058            | STANDBY     | ACTIVE      | Problem    |
| 59     | S VS059            | STANDBY     | ACTIVE      | Problem    |
| 60     | S VS060            | STANDBY     | ACTIVE      | Problem    |
| 61     | S VS061            | STANDBY     | ACTIVE      | Problem    |
| 62     | S VS062            | STANDBY     | ACTIVE      | Problem    |
| 63     | S VS063            | STANDBY     | ACTIVE      | Problem    |
| 64     | S VS064            | STANDBY     | ACTIVE      | Problem    |
| 65     | S VS065            | STANDBY     | ACTIVE      | Problem    |
| 66     | S VS066            | STANDBY     | ACTIVE      | Problem    |
| 67     | S VS067            | STANDBY     | ACTIVE      | Problem    |
| 68     | S VS068            | STANDBY     | ACTIVE      | Problem    |
| 69     | S VS069            | STANDBY     | ACTIVE      | Problem    |
| 70     | S VS070            | STANDBY     | ACTIVE      | Problem    |
| 71     | S VS071            | STANDBY     | ACTIVE      | Problem    |
| 72     | S VS072            | STANDBY     | ACTIVE      | Problem    |
| 73     | S VS073            | STANDBY     | ACTIVE      | Problem    |
| 74     | S VS074            | STANDBY     | ACTIVE      | Problem    |
| 75     | S VS075            | STANDBY     | ACTIVE      | Problem    |
| 76     | S VS076            | STANDBY     | ACTIVE      | Problem    |
| 77     | S VS077            | STANDBY     | ACTIVE      | Problem    |
| 78     | S VS078            | STANDBY     | ACTIVE      | Problem    |
| 79     | S VS079            | STANDBY     | ACTIVE      | Problem    |
| 80     | S VS080            | STANDBY     | ACTIVE      | Problem    |
| 81     | S VS081            | STANDBY     | ACTIVE      | Problem    |
| 82     | S VS082            | STANDBY     | ACTIVE      | Problem    |
| 83     | S VS083            | STANDBY     | ACTIVE      | Problem    |
| 84     | S VS084            | STANDBY     | ACTIVE      | Problem    |
| 85     | S VS085            | STANDBY     | ACTIVE      | Problem    |
| 86     | S VS086            | STANDBY     | ACTIVE      | Problem    |
| 87     | S VS087            | STANDBY     | ACTIVE      | Problem    |
| 88     | S VS088            | STANDBY     | ACTIVE      | Problem    |
| 89     | S VS089            | STANDBY     | ACTIVE      | Problem    |
| 90     | S VS090            | STANDBY     | ACTIVE      | Problem    |
| 91     | S VS091            | STANDBY     | ACTIVE      | Problem    |
| 92     | S VS092            | STANDBY     | ACTIVE      | Problem    |
| 93     | S VS093            | STANDBY     | ACTIVE      | Problem    |
| 94     | S VS094            | STANDBY     | ACTIVE      | Problem    |
| 95     | S VS095            | STANDBY     | ACTIVE      | Problem    |
| 96     | S VS096            | STANDBY     | ACTIVE      | Problem    |
| 97     | S VS097            | STANDBY     | ACTIVE      | Problem    |
| 98     | S VS098            | STANDBY     | ACTIVE      | Problem    |
| 99     | S VS099            | STANDBY     | ACTIVE      | Problem    |
| 100    | S VS100            | STANDBY     | ACTIVE      | Problem    |
| 102    | B VSB_Packet_Bro   | STANDBY     | ACTIVE      | Problem    |
| 251    | S PROV_VS1         | STANDBY     | ACTIVE      | Problem    |
------------------------------------------------------------------------
| Active Virtual Systems      | 0           | 97          |            |
------------------------------------------------------------------------
| Synchronization                                                      |
|     Within  chassis:           Enabled    (Default)                  |
|     Between chassis:           Enabled    (Default)                  |
|     Exception Rules:                      (Default)                  |
------------------------------------------------------------------------
```

Using asg stat vs

The `asg stat vs` command shows status information, SGM states, and problems for a specified Virtual System. You can run this command from `gclish` or the Expert mode. Select the Virtual System context before you run this command.

In `gclish`, run:

```
> set virtual_system <context>
> asg stat vs
```

In the Expert mode, run:

```
# vsenv <context>
# asg stat vs
--------------------------------------------------------------------------
|                          VSX System Status - 61000                     |
--------------------------------------------------------------------------
| VS ID                    1                                             |
| VS Name                  VS1                                           |
| Chassis Mode             VSLS                                          |
| FW Policy Date           09Jun14 19:12                                 |
--------------------------------------------------------------------------
| Chassis 1  (Primary)     STANDBY                                       |
--------------------------------------------------------------------------
| SGM ID      State        Process             Health                    |
| 1           DOWN         Inactive            fwk                       |
| 2 (local)   UP           Enforcing Security  OK                        |
| 3           UP           Enforcing Security  OK                        |
| 4           UP           Enforcing Security  OK                        |
--------------------------------------------------------------------------
| Chassis 2                ACTIVE                                        |
--------------------------------------------------------------------------
| SGM ID      State        Process             Health                    |
| 1           UP           Enforcing Security  OK                        |
| 2           UP           Enforcing Security  OK                        |
| 3           UP           Enforcing Security  OK                        |
| 4           UP           Enforcing Security  OK                        |
--------------------------------------------------------------------------
| Active Chassis: 2                                                      |
| Primary chassis has a problem. Secondary chassis health is better.     |
--------------------------------------------------------------------------
|                          Chassis 1           Chassis 2                 |
| Ports                     1 / 1               1 / 1                    |
| Bonds                     0 / 0               0 / 0                    |
| FWKs                      3 / 4               4 / 4                    |
| SGMs                      4 / 4               4 / 4                    |
--------------------------------------------------------------------------
```

This example shows that:

- VS1 on Chassis1- SGM1 is DOWN.
- The Primary Chassis for this Virtual System (Chassis 1) has a problem with the Firewall, but is otherwise working properly.
- VS1 failed over to Chassis 2, which is not the defined Primary Chassis for this Virtual System.
- All other SGMs are working properly.

SGM health status:

- **OK** - This SGM does not have problems.
- **SGM** - The SGM has a problem.
- **fwk** - The Firewall kernel has a problem.
- **Policy** - The policy date for this SGM is different from the Firewall policy date.
- **Interface** - The number of interface on this SGM is different from the related SGM on the other Chassis.
- **Problem** - This SGM has one or more problems.
- **Pnote** - This SGM has a problem that generated a pnote.

The bottom section shows the Active Chassis and the reason why the Primary Chassis is not Active, if applicable. Possible reasons:

- Primary Chassis health is good.
- Primary Chassis has a problem. Secondary Chassis health is better.
- Primary Chassis is above Active SGM threshold.
- Primary Chassis is below Active SGM threshold.
- Both Chassis have fwk problems. Continue using the Primary Chassis.
- Both Chassis have fwk problems. Primary Chassis health is better.
- Both Chassis have fwk problems. Secondary Chassis health is better.
- Both Chassis have interface problems. Continue using the Primary Chassis.
- Both Chassis have interface problems. Primary Chassis health is better.
- Both Chassis have interface problems. Secondary Chassis health is better.
- Both Chassis have problems. Continue using the Primary Chassis.
- Both Chassis have problems. Secondary Chassis health is better.

Using SNMP

SNMP information for VSLS is located under `iso.org.dod.internet.private.enterprise.checkpoint.products.asg.asgVSX.asgVslsInfo` (OID 1.3.6.1.4.1.2620.1.48.30.20)

VSLS SNMP monitors:

- SGM ratio threshold value
- System Primary Chassis
- Active Chassis for each Virtual System
- Primary Chassis for each Virtual System
- Number of configured interfaces for each Virtual System
- Number of UP interfaces for each Virtual System
- Number of working FWK instances for each Virtual System
- Total number of FWK instances for each Virtual System

SNMP for VSLS supports these modes:

- Default - SNMP collects data from all SGMs for all Virtual Systems
- Virtual Systems - SNMP monitors each Virtual System separately

Monitoring and Logging in VSX

VSX Functionality

The VSX commands run only on a VSX machine:

- stat - Print information about the VSX environment
- verify - Verify integrity and correctness of the configurations on all the blades

- logs - Collect VSX related logs
- hw_utilization - Hardware utilization
- mstat - Print VSX memory related information

Monitoring Hardware Utilization for VSX (hw_utilization)

Use the `hw_utilization` command monitor system CPU configuration connection capacity and CoreXL status on VS0. This set of tests runs as part of the `asg diag verify` utility for VSX environments only. The results show in the **VSX Configuration** section.

You can also run `hw_utilization`, in the Expert mode, as an independent command.

Syntax:

hw_utilization *<parameter>*

Parameter	Description
`cpu`	Shows alerts for CPU configuration issues.
`conn`	Shows alerts connection capacity issues.
`wizard`	Shows recommendations for optimum CPU distribution between Multi-Queue and `fwk` instances. Also shows recommendations for the optimal number of CoreXL instances for each Virtual System.
`configure`	Changes the default parameter values for this command.
`set_suppress`	Toggle display of alerts that show if CoreXL is enabled for VS0.

Examples:

```
# hw_utilization cpu
CPU utilization:
================
FWK cores: 0 1 10 11 2 20 21 22 23 24 25 26 27 28 29 3 31 39 4 5 6 7 8 9
MQ cores: 12 13 14 15 16 17 18 19 32 33 34 35 36 37 38 39
No overlapping CPU/s
Unused CPU ID/s: 30
Overlapping CPU ID/s:39
```

```
# hw_utilization conn
Connection capacity utilization:
=================================
+------+---------+-----------------+-----------------+---------------+-------------+
| VSID | Type    | Name            | [SGM_ID]        | [SGM_ID]      | Conn. limit |
|      |         |                 | Max Conn.       | Max Conn.     |             |
|      |         |                 | Number          | Peak          |             |
+------+---------+-----------------+-----------------+---------------+-------------+
| 0    | VSX     | Guru-T3-127     | [1_05] 572      | [1_02] 9312   | 31800       |
| 1    | VS      | vs1-T3          | [1_02] 4900     | [1_08] 95     | 49800       |
| 3    | VS      | vs2-T3          | [1_11] 8        | [1_03] 540    | 1999900     |
| 4    | VS      | vs3-T3          | [2_03] 9        | [1_02] 530    | 999900      |
| 5    | VS      | vs4-T3          | [1_02] 19502    | [1_02] 0      | 24900       |
| 7    | VSB     | vsb1-T3         | [1_03] 350      | [1_05] 0      | 49800       |
+------+---------------------------+-----------------+----------------+------------+
All virtual devices are above the minimum connection capacity limit (24000)
**Concurrent connections amount almost exceeds connection limit**
Virtual devices 1 5 are close to their connection limit (less than 25000 new connections can be opened)
```

```
# hw_utilization wizard
How much traffic is accelerated (in percentage)?
```

```
Give the traffic distribution for each configured VS (in percentage).

According to the given information a recommended CPU tuning for the system is presented. For example:
How much traffic is accelerated (in percentage)?
40
How much traffic is distributed to the following 7 VSs (in percentage):

vs1-T3:  10
vs2-T3:  10
vs3-T3:  10
vs4-T3:  10
vs5-T3:  10
vsb1-T3: 40

 Recommended optimization:
=========================
4 cores for Multi-queue
8 cores for VSs
Instances per VSs:
1 instance for Gruffalo-T3-127
1 instance for vs1-T3
1 instance for vs2-T3
1 instance for vs3-T3
1 instance for vs4-T3
1 instance for vs5-T3
4 instances for vsb1-T3
Please note that the number of recommended assigned FWK instances (10) is higher by 2 than the number
of CPUs that are available for FWK (8)
Hence, there will be maximum 2 CPUs that will run more than 1 fwk instance simultaneously
**The instances per VS recommendation assumes that all VSs are handling traffic simultaneously
Connections Capacity
According to pre-defined values: MIN_CONN_LIMIT and CONN_DIFF_FROM_LIMIT the following checks will be
performed and the user will be alerted:
1.     connections limit < MIN_CONN_LIMIT
2.     connections number + CONN_DIFF_FROM_LIMIT >= connections limit
```

Monitoring VSX Memory Resources

Use vsxmstat to enable memory monitoring for the VSX Gateway. This command shows an overview of the memory that the system and each virtual device is using. These are the global memory resources that are shown:

- **Memory Total** - Total physical memory on the VSX Gateway.
- **Memory Free** - Available physical memory.
- **Swap Total** - Total of swap memory.
- **Swap Free** - Available swap memory.
- **Swap-in rate** - Total memory swaps per second.

The virtual devices are listed according to the VSIDs. Run vsx stat -v to show the VSID for the virtual devices.

You must be in Expert mode to run vsxmstat.

Managing vsxmstat

Use the vsxmstat command to enable or disable memory information collection on the VSX Gateway.

Syntax

```
# vsxmstat {enable_raw|disable_raw|status_raw}
```

Parameter	Description
`enable_raw`	Enables memory resource monitoring for perfanalyze use.
`disable_raw`	Disables memory resource monitoring for perfanalyze use.
`status_raw`	Shows if memory resource monitoring is enabled or disabled for perfanalyze use.

Example

```
# vsxmstat disable_raw
```

Output

```
VSX memory resource control is disabled for perfanalyze use
```

Memory Resources for Each Virtual Device

Use `vsxmstat` to show memory usage for each virtual device. You can use the `-vs` parameter to show specified virtual devices only.

Use these parameters to show more data:

- `unit` - Change the memory measurement unit shown in the command output.
- `sort` - Sort the results according to the virtual devices that use the most memory. Limit the display to the specified number of results.

Syntax

```
# vsxmstat [-vs <vsid>] [unit <unit>] [sort <top>]
```

Parameter	Description
`-vs`	Shows the memory usage of the specified virtual devices.
<vsid>	virtual device identification. To show multiple devices: • Put a space between each VSID: `-vs 1 3 5` • List a range of VSIDs: `-vs 1-4` **Note**: You can combine VSID ranges together with single VSIDs
`unit`	Change the memory measurement unit shown in the command output.
<unit>	The memory measurement unit. The default value is megabytes. Use with the `unit` parameter. The values are: • B - bytes • K, KB - kilobytes • M, MB - megabytes (default) • G, GB - gigabytes

Parameter	Description
`sort`	Sort the results according to the virtual devices that use the most memory.
<*top*>	Maximum number of virtual devices to show. Only those virtual devices that use the most memory are shown. Use with the `sort` parameter. Use `all` to show all virtual devices.

Example

```
# vsxmstat -vs  0 1 3 5-8 unit MB sort 5
# vsxmstat sort 5
```

Output (Both examples show the same results)

```
VSX Memory Status
=================
Memory Total: 997.22 MB
Memory Free: 232.56 MB
Swap Total: 2047.34 MB
Swap Free: 2047.16 MB
Swap-in rate: 0.00 MB
 VSID | Memory Consumption
======+====================
    0 |         133.50 MB
    8 |          92.41 MB
    3 |          43.81 MB
    6 |          42.47 MB
    1 |          42.47 MB
```

Configuring Swap-in Sample Rate

The swap-in rate measures how much memory per second that the system swaps-in from the disk. You can configure how frequently the system calculates the swap-in rate. For example, a sample rate of 5 means that the system calculates the swap-in rate at five minute intervals.

Syntax

```
# vsxmstat swap <minutes>
```

Parameter	Description
<*minutes*>	Number of minutes that the system measures memory swaps to determine the swap-in rate. Only integers are valid values. The default swap-in sample rate is `10`.

Example

```
# vsxmstat swap 5
```

Output

```
Swap-in sample rate was changed successfully to 5 minutes.
```

Comments

Swap-in sample rate is a system wide Linux setting. When you change the value for memory monitoring, all the swap-in rates are calculated according to the new value.

When you enable the monitoring memory resources feature, the swap-in rate setting is saved. When you disable the feature, the system restores the saved setting.

Using Debug Mode

Use the debug parameter to show more data about the memory that the VSX Gateway uses. You cannot use the `-vs`, `unit` and `sort` parameters in debug mode. The memory is shown in kilobytes.

Syntax

```
# vsxmstat debug
```

Output

```
VSX Memory Status
=================
Memory Total: 1021152.00 KB
Memory Free: 324788.00 KB
Swap Total: 2096472.00 KB
Swap Free: 2096404.00 KB
Swap-in rate: 375.34 KB

 VSID |      Private_Clean |      Private_Dirty |    DispatcherGConn
======+====================+====================+====================+
    0 |        13544.00 KB |       144268.00 KB |            0.00 KB |
    1 |         1740.00 KB |        46276.00 KB |            0.00 KB |
    2 |         1720.00 KB |        46868.00 KB |            0.00 KB |
    3 |         1720.00 KB |        46644.00 KB |            0.00 KB |
    4 |         1712.00 KB |        45144.00 KB |            0.00 KB |
    5 |         1712.00 KB |        45836.00 KB |            0.00 KB |
    6 |         1720.00 KB |        45000.00 KB |            0.00 KB |
    7 |         1720.00 KB |        45044.00 KB |            0.00 KB |
```

Comments

By default the debug parameter shows these memory fields:

Field	Description
Private_Clean	Clean private pages. (`/proc/[`*pid*`]/smaps`)
Private_Dirty	Dirty private pages. (`/proc/[`*pid*`]/smaps`)
DispatcherHTab	Hash table for each Virtual System
DispatcherGConn	Global connections for each Virtual System
SecureXL	SecureXL memory each Virtual System uses

Monitoring VSX Configuration (vsx stat)

This tool runs only on VSX.

VSX stat (bin vsx_stat):

Configuration	Description
policy	Print policies on Virtual Systems
sw_blades	Print Software Blades on Virtual Systems
processes	Print Virtual Systems processes
topology	Print topology on Virtual Systems
routes	print routes on Virtual Systems
interfaces	print interfaces and distributions on Virtual Systems
cores	print core allocations on Virtual Systems
conn_vmems	print connections and virtual memory on Virtual Systems

Usage:

Option	Description
--version	Show program's version number and exit
-h, --help	Show this help message and exit
-p, --policy	Print policies on Virtual Systems
-v, --v	Legacy print policies on Virtual Systems
-s, --sw_blades	Print Software Blades on Virtual Systems
-c, --processes	Print processes on Virtual Systems
-t, --topology	Print topology on Virtual Systems
-r, --routes	Print routes on Virtual Systems
-i, --interfaces	Print interfaces and distributions on Virtual Systems
-u, --cores	Print core allocations on Virtual Systems
-n, --conn_vmems	Print connections and virtual memory on Virtual Systems
-a, --all	Print all the information on Virtual Systems

```
Policy Table
=============
+----+-----------+-----------------+--------------+--------------+
| ID | Type & Name | Security Policy | Installed at    SIC State |
+----+-------------+-----------------+---------------+-----------+
| 0  | S Cost     | NA-Core-GW_VSX | 2016-04-03 17:30:54 | Trust |
| 1  | S MEX-T2-VS  | MEX-T2       | 2016-04-03 17:36:47 | Trust |
| 2  | S EQU-T2-VS  | EQU-T2       | 2016-04-10 10:47:38 | Trust |
| 3  | S CAN-T2-VS  | CAN-T2       | 2016-04-10 10:38:26 | Trust |
| 4  | S SHR-T2-VS  | SHARED-T2    | 2016-04-03 17:46:34 | Trust |
| 5  | S EQU-T3-VS  | EQU-T3       | 2016-04-03 17:33:46 | Trust |
| 6  | S EXTRANET-VS | EXTRANET    | 2016-04-03 17:54:09 | Trust |
| 7  | S MEX-T3-VS  | MEX-T3       | 2016-04-03 17:37:56 | Trust |
| 8  | S SHR-T3-VS  | SHARED-T3    | 2016-04-03 17:47:40 | Trust |
| 9  | S FUSION-VS  | FUSION       | 2016-04-03 17:51:28 | Trust |
| 10 | S IPT-VS     | IPT          | 2016-04-03 17:35:55 | Trust |
| 11 | S JP-T3-VS   | JP-T3        | 2016-04-03 17:42:27 | Trust |
| 12 | S AB-T3-VS   | AB-T3        | 2016-04-03 17:45:34 | Trust |
| 13 | S MGMT-VS    | MGMT         | 2016-04-03 17:50:08 | Trust |
| 14 | S VENDOR-VS  | VENDOR       | 2016-04-03 17:49:02 | Trust |
| 15 | S ALM-T3-VS  | ALM-T3       | 2016-04-03 17:43:23 | Trust |
| 16 | B VSB_1_TAP  | <Not Applicable> |                 | Trust |
| 17 | W VSW1       | <Not Applicable> |                 | Trust |
| 18 | W VSW2       | <Not Applicable> |                 | Trust |
| 19 | B VSB_2_Access | <Not Applicable> |               | Trust |
| 20 | W VSW_Share_identities | <Not Applicable> |       | Trust |
| 21 | W VSW_Barakeo1 | <Not Applicable> |               | Trust |
| 22 | S VS_Barakeo | Big_Policia1 | 2016-04-10 10:43:31 | Trust |
| 23 | W DMZ_VSW      | <Not Applicable> |               | Trust |
| 24 | W VPN_DMZ_VSW  | <Not Applicable> |               | Trust |
| 25 | B VSB_Packet_Brocker | <Not Applicable> |         | Trust |
| 26 | S VS026     | Standard      | 2016-04-03 17:53:16 | Trust |
| 27 | S VS027     | Standard      | 2016-04-03 17:53:36 | Trust |
| 28 | S VS028     | Standard      | 2016-04-03 17:53:56 | Trust |
+----+----------------+---------------+------------------+-------+


Software Blades Table
======================
+-----+----------------+-------------------------------------------+
| ID | Type & Name     | Software Blades                          |
+----+----------------+--------------------------------------------+
| 0  | S Cost          | FW                                       |
| 1 | S MEX-T2-VS | FW VPN URLF AV APPI IPS IDENTITYSERVER ANTI_BOT |
| 2 | S EQU-T2-VS | FW VPN URLF AV APPI IPS IDENTITYSERVER ANTI_BOT |
| 3 | S CAN-T2-VS | FW VPN URLF AV APPI IPS IDENTITYSERVER ANTI_BOT |
| 4 | S SHR-T2-VS | FW VPN                                     |
| 5 | S EQU-T3-VS | FW VPN URLF AV APPI IPS IDENTITYSERVER ANTI_BOT |
| 6 | S EXTRANET-VS | FW VPN APPI IDENTITYSERVER                 |
| 7  | S MEX-T3-VS | FW URLF APPI IDENTITYSERVER                 |
| 8  | S SHR-T3-VS | FW AV ANTI_BOT                              |
| 9  | S FUSION-VS | FW                                          |
| 10 | S IPT-VS    | FW                                          |
| 11 | S JP-T3-VS  | FW                                          |
| 12 | S AB-T3-VS  | FW IDENTITYSERVER                           |
| 13 | S MGMT-VS   | FW VPN IPS                                  |
| 14 | S VENDOR-VS | FW                                          |
| 15 | S ALM-T3-VS | FW VPN                                      |
| 16 | B VSB_1_TAP | FW URLF AV APPI IPS SSL_INSPECT ANTI_BOT    |
| 19 | B VSB_2_Access | FW URLF AV APPI SSL_INSPECT ANTI_BOT     |
| 22 | S VS_Barakeo   | FW IPS                                   |
| 25 | B VSB_Packet_Brocker | FW URLF AV APPI IPS ANTI_BOT       |
| 26 | S VS026        | FW IPS                                   |
| 27 | S VS027        | FW IPS                                   |
+----+----------------+-------------------------------------------+


Processes Table
================
+-----+------------------------+-------+-------+-------+--------+
|  ID | Type & Name            |  FWK  |  FWD  |  CPD  | ROUTED |
+-----+------------------------+-------+-------+-------+--------+
|  0  | S Cost                 |  9095 | 23132 | 18466 |  8153  |
|  1  | S MEX-T2-VS            |  6899 |  6912 |  6901 | 20694  |
|  2  | S EQU-T2-VS            |  5313 |  5308 |  5306 |  1220  |
|  3  | S CAN-T2-VS            | 31757 | 31818 | 30765 |  7084  |
|  4  | S SHR-T2-VS            | 27556 | 27567 | 27563 | 21253  |
|  5  | S EQU-T3-VS            | 15792 | 14992 | 14889 |  8155  |
|  6  | S EXTRANET-VS          | 10289 | 10298 | 10288 | 31735  |
|  7  | S MEX-T3-VS            | 12681 | 11924 | 11915 | 25467  |
|  8  | S SHR-T3-VS            | 28849 | 28876 | 28833 | 21163  |
|  9  | S FUSION-VS            | 29471 | 29470 | 29440 | 15220  |
|  10 | S IPT-VS               |  7530 |  7545 |  7536 | 15307  |
|  11 | S JP-T3-VS             | 15126 | 15149 | 15127 | 20674  |
|  12 | S AB-T3-VS             |  4063 |  4045 |  4036 |  7480  |
|  13 | S MGMT-VS              | 12592 | 13018 | 13003 |  1001  |
|  14 | S VENDOR-VS            | 18563 | 18576 | 18567 | 20665  |
|  15 | S ALM-T3-VS            | 23961 | 22049 | 22009 |  7972  |
|  16 | B VSB_1_TAP            | 28033 | 28055 | 28040 |   -    |
|  17 | W VSW1                 | 13577 | 13560 | 13540 |   -    |
|  18 | W VSW2                 | 13638 | 13767 | 13704 |   -    |
|  19 | B VSB_2_Access         | 25870 | 25876 | 25868 |   -    |
|  20 | W VSW_Share_identities | 13670 | 13763 | 13691 |   -    |
|  21 | W VSW_Barakeo1         | 13630 | 13685 | 13639 |   -    |
|  22 | S VS_Barakeo           | 32648 | 32656 | 32637 | 12980  |
|  23 | W DMZ_VSW              |  5952 |  5948 |  5938 |   -    |
|  24 | W VPN_DMZ_VSW          | 13516 | 13538 | 13528 |   -    |
|  25 | B VSB_Packet_Brocker   |  9757 |  9773 |  9756 |   -    |
|  26 | S VS026                | 14382 | 14306 | 14111 | 29379  |
|  27 | S VS027                | 14485 | 14465 | 14440 | 29381  |
|  28 | S VS028                | 14308 | 14612 | 14463 | 29384  |
|  29 | S VS029                |  7669 |  7733 |  7671 | 25901  |
|  30 | S VS030                |  9471 |  9454 |  9447 | 28496  |
|  31 | S VS031                | 14680 | 14709 | 14691 |  5603  |
|  32 | S VS032                | 14584 | 14697 | 14519 | 29397  |
|  33 | S VS033                | 14502 | 14609 | 14523 |  5616  |
|  34 | S VS034                | 16692 | 18837 | 16815 |  5630  |
|  35 | S VS035                | 21622 | 22052 | 22038 |  5641  |
+-----+------------------------+-------+-------+-------+--------+

Topology Table
===============
+------+------------------+------+--------------------+-----------+
| VSID | Type & Name      | VSID | Type & Name        | Interface |
+------+------------ -----+------+--------------------+-----------+
|  10  | S IPT-VS         |  23  | W DMZ_VSW          | wrpj640   |
+------+------------------+------+--------------------+-----------+
| 100  | S kaki_pipi      |  17  | W VSW1             | wrpj6400  |
+------+------------------+------+--------------------+-----------+
|  12  | S AB-T3-VS       |  17  | W VSW1             | wrpj768   |
|      | S AB-T3-VS      | 20   | W VSW_Share_identities | wrpj769 |
|      | S AB-T3-VS      |  23   | W DMZ_VSW             | wrpj770 |
+------+-----------------+------+---------------------+-----------+
|  17  | W VSW1          | 100  | S kaki_pipi         | wrp6400   |
|      | W VSW1          |  12  | S AB-T3-VS          | wrp768    |
|      | W VSW1          |  3   | S CAN-T2-VS         | wrp193    |
+------+-----------------+------+---------------------+-----------+
|  2   | S EQU-T2-VS      |  23  | W DMZ_VSW            | wrpj130   |
|      | S EQU-T2-VS      |  24  | W VPN_DMZ_VSW        | wrpj131   |
+------+-----------------+------+---------------------+-----------+
|  20  | W VSW_Share_identities |  12  | S AB-T3-VS   | wrp769    |
|      | W VSW_Share_identities |  3   | S CAN-T2-VS  | wrp192    |
+------+------------------------+------+--------------+-----------+
|  23  | W DMZ_VSW              |  10  | S IPT-VS     | wrp640    |
|      | W DMZ_VSW              |  12  | S AB-T3-VS   | wrp770    |
|      | W DMZ_VSW              |  2   | S EQU-T2-VS  | wrp130    |
|      | W DMZ_VSW              |  3   | S CAN-T2-VS  | wrp194    |
|      | W DMZ_VSW              |  4   | S SHR-T2-VS  | wrp257    |
|      | W DMZ_VSW              |  5   | S EQU-T3-VS  | wrp322    |
|      | W DMZ_VSW              |  6   | S EXTRANET-VS | wrp384   |
|      | W DMZ_VSW              |  7   | S MEX-T3-VS   | wrp448   |
+------+------------------------+------+---------------+----------+
|  24  | W VPN_DMZ_VSW          |  2   | S EQU-T2-VS   | wrp131   |
|      | W VPN_DMZ_VSW          |  4   | S SHR-T2-VS   | wrp256   |
|      | W VPN_DMZ_VSW          |  5   | S EQU-T3-VS   | wrp321   |
+------+------------------------+------+---------------+----------+
|  3   | S CAN-T2-VS            |  17  | W VSW1        | wrpj193  |
|      | S CAN-T2-VS      | 20 | W VSW_Share_identities | wrpj192 |
|      | S CAN-T2-VS      | 23 | W DMZ_VSW              | wrpj194 |
+------+------------------+----+------------------------+---------+
|  4   | S SHR-T2-VS      | 23 | W DMZ_VSW              | wrpj257 |
|      | S SHR-T2-VS      | 24 | W VPN_DMZ_VSW          | wrpj256 |
+------+------------------+----+------------------------+---------+
|  5   | S EQU-T3-VS      | 23 | W DMZ_VSW              | wrpj322 |
|      | S EQU-T3-VS      | 24 | W VPN_DMZ_VSW          | wrpj321 |
+------+------------------+----+------------------------+---------+
|  6   | S EXTRANET-VS    | 23 | W DMZ_VSW              | wrpj384 |
+------+------------------+----+------------------------+---------+
|  7   | S MEX-T3-VS      | 23 | W DMZ_VSW              | wrpj448 |
+------+------------------+----+------------------------+---------+


Routes Table
=============
+----+-------------+----------------+----------------+------------+
| ID | Type & Name | Destination    |    Gateway     | Interface  |
+----+-------------+----------------+----------------+------------+
|  1 | S MEX-T2-VS |   120.100.1.96 |    11.1.1.3    | bond2.120  |
|    |             |   120.100.1.97 |    11.1.1.3    | bond2.120  |
|    |             |   120.100.1.98 |    11.1.1.3    | bond2.120  |
|    |             |   120.100.1.99 |    11.1.1.3    | bond2.120  |
|    |            |   120.100.1.100 |    11.1.1.3    | bond2.120  |
|    |             |   120.100.1.88 |    11.1.1.3    | bond2.120  |
|    |             |   120.100.1.91 |    11.1.1.3    | bond2.120  |
|    |             |   120.100.1.90 |    11.1.1.3    | bond2.120  |
|    |             |   120.100.1.93 |    11.1.1.3    | bond2.120  |
|    |             |   120.100.1.92 |    11.1.1.3    | bond2.120  |
|    |             |   120.100.1.95 |    11.1.1.3    | bond2.120  |
|    |             |   120.100.1.94 |    11.1.1.3    | bond2.120  |
|    |             |   120.100.1.87 |    11.1.1.3    | bond2.120  |
|    |             |   120.100.1.86 |    11.1.1.3    | bond2.120  |
|    |             |    5.5.5.5     | 10.133.252.27  | bond3.48   |
|    |             |   120.100.1.58 |    11.1.1.3    | bond2.120  |
|     |            |    120.100.1.59   |    11.1.1.3 | bond2.120  |
|     |            |    81.81.81.1 |     11.1.1.4    | bond2.120  |
|     |            |   120.100.1.60 |    11.1.1.3    | bond2.120  |
|     |            |    120.100.1.48 |   11.1.1.3    | bond2.120  |
|     |            |   120.100.1.42  |   11.1.1.3    | bond2.120  |
|     |            |   120.100.1.43  |   11.1.1.3    | bond2.120  |
|     |            |   120.100.1.40  |   11.1.1.3    | bond2.120  |
|     |            |   120.100.1.41  |   11.1.1.3    | bond2.120  |
|     |            |   120.100.1.46  |   11.1.1.3    | bond2.120  |
|     |            |   120.100.1.47  |    11.1.1.3   | bond2.120  |
|     |            |   120.100.1.44  |    11.1.1.3   | bond2.120  |
+-----+------------+-----------------+---------------+------------+


Interfaces Table
=================
+----+-------------+-----------+----------+---------+--------------+
| ID | Type & Name | Interface | Address  | Netmask | Distribution |
+----+-------------+-----------+----------+---------+--------------+
|    |             |           |          |         |              |
| 1 | S MEX-T2-VS | bond1.2303 | 10.133.242.215 | 28 | policy-internal |
|   |             |            |       -         |    -    |           |
|   |             | bond2.2025 | 81.81.81.97 |  30   | policy-external |
|   |             |            |        -         |    -    |          |
|   |             | bond2.2024 |  81.81.81.93 |  30 | policy-internal |
|   |             |            |        -        |    -    |          |
|   |             | bond2.2021 |  81.81.81.81 |  30 | policy-internal |
|   |             |            |     -        |  -  |                 |
|   |             | bond2.2020 |  81.81.81.77 |  30 | policy-internal |
|   |             |            |     -        |  -  |                 |
|   |             | bond2.2023 |  81.81.81.89 |  30 | policy-external |
|   |             |            |      -       |  -  |                 |
|   |             | bond2.2022 | 81.81.81.85  |  30 | policy-internal |
|   |             |            |       -      |  -  |                 |
|   |            | bond1.252  | 10.133.242.100 | 28 | policy-external |
|   |             |            |       -        |  -   |              |
|   |             | bond2.120  |     11.1.1.1  |  8 | policy-internal |
|   |             |            |      -        |  - |                 |
|   |             | bond2.2011 |   81.81.81.41 | 30 | policy-external |
|   |             |            |      -        |  - |                 |
|   |             | bond2.2015 | 81.81.81.57  |  30 | policy-internal |
|   |             |            |     -        |  -  |                 |
|   |             | bond6.562 | 171.171.251.11 | 24 | policy-internal |
|   |             |           |       -        |  -  |                |
|   |             | bond1.120 |  111.1.1.1    | 8  | policy-external |
|   |             |           |      -        |  -    |               |
|   |             | bond2.2010 | 81.81.81.37  |  30 | policy-internal |
|   |             |            |      -       |  -  |                 |
|   |             | bond2.2012 |  81.81.81.45 |  30 | policy-internal |
|   |             |            |     -        |  -  |                 |
|   |             | bond2.2013 | 81.81.81.49  |  30 | policy-internal |
|   |             |            |      -       |  -  |                 |
|   |             | bond2.2014 |  81.81.81.53 |  30 | policy-internal |
|   |             |            |      -       |  -  |                 |
|   |             | bond2.2018 | 81.81.81.69  |  30 | policy-external |
|   |             |            |      -       |  -  |                 |
|   |             | bond2.2019 | 81.81.81.73  |  30 | policy-internal |
|   |             |            |      -       |  -  |                 |
|   |             | bond3.120  |   1.1.1.1    | 8  | policy-internal |
|   |             |            |      -       |  -  |                 |
|   |             | bond2.2016 |  81.81.81.61 |  30 | policy-internal |
|   |             |            |      -       |  -  |                 |
|   |             | bond2.2017 | 81.81.81.65  |  30 | policy-external |
|   |             |            |      -       |  -  |                 |
|   |             | bond6.3000 |   31.0.0.1   |  8  | policy-internal |
|   |             |            |      -       |  -  |                 |
|   |             | bond4.120  |   21.0.0.1   |  8  | policy-external |
|   |             |            |      -       |  -  |                 |
|   |             | bond2.2007 | 81.81.81.25  |  30 | policy-internal |
|   |             |            |      -       |  -  |                 |
|   |             | bond2.2006 | 81.81.81.21  |  30 | policy-internal |
|   |             |            |      -       |  -  |                 |
+---+-------------+------------+--------------+-----+-----------------+


Core Allocations Table
=======================
+---+--------------+------------+-------+-----------------------------+
| ID | Type & Name | CoreXL IPv(4/6) | Type  | CPUs                    |
+----+-------------+-----------------+-------+------------------------+
| 0 | S Cost | -/-              |       |                               |
|   |   |    | P FWK | 0 1 2 3 4 5 6 7 8 9 10 11 20 21 22 23 24 25 26 27
28 29 30 31 |
|   |   |    | P FWD | 0 1 2 3 4 5 6 7 8 9 10 11 20 21 22 23 24 25 26 27
28 29 30 31 |
|   |   |    | P CPD | 0 1 2 3 4 5 6 7 8 9 10 11 20 21 22 23 24 25 26 27
28 29 30 31 |
| 1 | S MEX-T2-VS | 8/0 |                                                  |
|   |   |    | P FWK | 0 1 2 3 4 5 6 7 8 9 10 11 20 21 22 23 24 25 26 27
28 29 30 31 |
|   |   |    | P FWD | 0 1 2 3 4 5 6 7 8 9 10 11 20 21 22 23 24 25 26 27
28 29 30 31 |
|   |   |    | P CPD | 0 1 2 3 4 5 6 7 8 9 10 11 20 21 22 23 24 25 26 27
28 29 30 31 |
| 2 | S EQU-T2-VS | 8/0 |   |                                             |
|   |   |    | P FWK | 0 1 2 3 4 5 6 7 8 9 10 11 20 21 22 23 24 25 26 27
28 29 30 31 |
|   |   |    | P FWD | 0 1 2 3 4 5 6 7 8 9 10 11 20 21 22 23 24 25 26 27
28 29 30 31 |
|   |   |    | P CPD | 0 1 2 3 4 5 6 7 8 9 10 11 20 21 22 23 24 25 26 27
28 29 30 31 |
| 3 | S CAN-T2-VS | 8/0 |   |                                             |
|   |   |    | P FWK | 0 1 2 3 4 5 6 7 8 9 10 11 20 21 22 23 24 25 26 27
28 29 30 31 |
|   |   |    | P FWD | 0 1 2 3 4 5 6 7 8 9 10 11 20 21 22 23 24 25 26 27
28 29 30 31 |
|   |   |   | P CPD | 0 1 2 3 4 5 6 7 8 9 10 11 20 21 22 23 24 25 26 27
28 29 30 31 |
| 4 | S SHR-T2-VS | 8/0 |   |                                          |
|   |   |   | P FWK | 0 1 2 3 4 5 6 7 8 9 10 11 20 21 22 23 24 25 26 27
28 29 30 31 |
|   |   |   | P FWD | 0 1 2 3 4 5 6 7 8 9 10 11 20 21 22 23 24 25 26 27
28 29 30 31 |
|   |   |   | P CPD | 0 1 2 3 4 5 6 7 8 9 10 11 20 21 22 23 24 25 26 27
28 29 30 31 |
| 5 | S EQU-T3-VS | 3/0 |   |                                          |
|   |   |   | P FWK | 0 1 2 3 4 5 6 7 8 9 10 11 20 21 22 23 24 25 26 27
28 29 30 31 |
|   |   |   | P FWD | 0 1 2 3 4 5 6 7 8 9 10 11 20 21 22 23 24 25 26 27
28 29 30 31 |
|   |   |   | P CPD | 0 1 2 3 4 5 6 7 8 9 10 11 20 21 22 23 24 25 26 27
28 29 30 31 |
| 6 | S EXTRANET-VS | 4/0 |   |                                        |
|   |   |   | P FWK | 0 1 2 3 4 5 6 7 8 9 10 11 20 21 22 23 24 25 26 27
28 29 30 31 |
|   |   |   | P FWD | 0 1 2 3 4 5 6 7 8 9 10 11 20 21 22 23 24 25 26 27
28 29 30 31 |
|   |   |   | P CPD | 0 1 2 3 4 5 6 7 8 9 10 11 20 21 22 23 24 25 26 27
28 29 30 31 |
+---+---+---+-------+--------------------------------------------------+


Connections and Virtual memory Table
=====================================
+-----+------------------------+-------------+---------------------+
|  ID | Type & Name            | Virtual Mem | Connections         |
+-----+------------------------+-------------+---------------------+
|  0  | S Cost                 | 789/62994   | 3973/7043/16900     |
|  1  | S MEX-T2-VS            | 3042/62994  | 1427/151559/4999900 |
|  2  | S EQU-T2-VS            | 2376/62994  | 4576/189922/1999900 |
|  3  | S CAN-T2-VS            | 2245/62994  | 141/193912/999900   |
|  4  | S SHR-T2-VS            | 1864/62994  | 10/982133/999900    |
|  5  | S EQU-T3-VS            | 1330/62994  | 13/521/999900       |
|  6  | S EXTRANET-VS          | 1319/62994  | 310/1038/999900     |
|  7  | S MEX-T3-VS            | 2249/62994  | 13/406/1999900      |
|  8  | S SHR-T3-VS            | 989/62994   | 18/38/1499900       |
|  9  | S FUSION-VS            | 719/62994   | 15/17/299900        |
|  10 | S IPT-VS               | 1706/62994  | 19/509/4999900      |
|  11 | S JP-T3-VS             | 743/62994   | 33/47/19900         |
|  12 | S AB-T3-VS             | 750/62994   | 15/39/14900         |
|  13 | S MGMT-VS              | 801/62994   | 0/0/149900          |
|  14 | S VENDOR-VS            | 741/62994   | 0/0/14900           |
|  15 | S ALM-T3-VS            | 1483/62994  | 0/0/499900          |
|  16 | B VSB_1_TAP            | 1878/62994  | 2835/6372/999900    |
|  17 | W VSW1                 | 723/62994   | 0/0/900             |
|  18 | W VSW2                 | 659/62994   | 0/0/900             |
|  19 | B VSB_2_Access         | 1129/62994  | 0/0/14900           |
|  20 | W VSW_Share_identities | 659/62994   | 0/0/900             |
|  21 | W VSW_Barakeo1         | 723/62994   | 0/0/900             |
|  22 | S VS_Barakeo           | 747/62994   | 0/0/14900           |
|  23 | W DMZ_VSW              | 659/62994   | 0/0/900             |
|  24 | W VPN_DMZ_VSW          | 659/62994   | 0/0/900             |
|  25 | B VSB_Packet_Brocker   | 2148/62994  | 24797/40281/99900   |
|  26 | S VS026                | 670/62994   | 0/0/14900           |
|  27 | S VS027                | 734/62994   | 0/0/14900           |
|  28 | S VS028                | 670/62994   | 0/0/14900           |
|  29 | S VS029                | 734/62994   | 0/0/14900           |
|  30 | S VS030                | 734/62994   | 0/0/14900           |
|  31 | S VS031                | 734/62994   | 0/0/14900           |
|  32 | S VS032                | 734/62994   | 0/0/14900           |
|  33 | S VS033                | 734/62994   | 0/0/14900           |
|  34 | S VS034                | 734/62994   | 0/0/14900           |
|  35 | S VS035                | 670/62994   | 0/0/14900           |
|  36 | S VS036                | 734/62994   | 0/0/14900           |
|  37 | S VS037                | 734/62994   | 0/0/14900           |
|  38 | S VS038                | 670/62994   | 0/0/14900           |
+-----+------------------------+-------------+---------------------+
```

VSX Legacy Bridge Mode

VSX Legacy Bridge Mode lets Virtual Systems in Bridge Mode ignore tagged packets.

Use `fw -i k ctl set int fw_vsx_legacy_bridge_mode` <*mode*> to manage VSX Legacy Bridge Mode.

Syntax

```
> fw -i k ctl set int fw_vsx_legacy_bridge_mode <mode>
```

Parameter	Description
<*mode*>	Sets the VSX Legacy Bridge Mode Valid values: • `0` (Default) - Disable VSX Legacy Bridge Mode • `1` - Enable VSX Legacy Bridge Mode

Working with LTE Features

In This Section:

The 61000/41000 Security System includes features that support advanced LTE telecommunication. Most of these features are configured with SmartDashboard or on the management server. See the R76 LTE Release Notes http://downloads.checkpoint.com/dc/download.htm?ID=29339 for detailed information and configuration procedures. Configuration procedures for SGMs are included in this section for your convenience.

These LTE features include:

- LTE S1 VPN
- Firewall GX support
- GTPv2 support
- GTP CoreXL support
- GTP Signaling rate limit
- SCTP support
- Diameter inspection
- Third-Party Syslog
- MSS adjustment
- CGNAT
- Stateless NAT46 translation
- NAT 64
- Large Scale VPN

Enabling LTE Support

LTE configuration includes hundreds or thousands of eNodeB VPN peers. Each eNodeB has its own IPSec tunnel to the 61000/41000 Security System. eNodeB encrypts GTP traffic from mobile clients behind the eNodeB.

You must enable LTE support to use LTE features and S1 VPN.

To enable LTE support for all SGMs:

On the 61000/41000 Security System, run:

```
> asg_lte_config enable
> reboot -b all
```

Note: Hyper-threading must be disabled for LTE.

To disable hyper-threading:

1. Enter #*g_cpconfig ht disable*
2. Reboot.

Limitations:

- Connections are synchronized to all SGMs, not just the standby SGM.
- You must not enable SPI distribution.

VPN Sticky SA

By default, the VPN Sticky *Security Association (SA)* feature is enabled. This feature makes sure that the 61000/41000 Security System has only one outgoing SA to remote peers. This is a requirement for some network device manufacturers to minimize security vulnerabilities.

⚠ **Important** - Make sure that SPI distribution and Sticky SA are not enabled. at the same time.

Configuring VPN Sticky SA

1. To disable VPN Sticky SA, run this command in the Expert mode:
   ```
   # g_update_conf_file $FWDIR/modules/fwkern.conf
   fwha_vpn_sticky_tunnel_enabled=0
   ```
2. To re-enable VPN Sticky SA, run this command in the Expert mode:
   ```
   # g_update_conf_file $FWDIR/modules/fwkern.conf
   fwha_vpn_sticky_tunnel_enabled=1
   ```
3. Reboot all SGMs:
   ```
   # reboot -b all
   ```

Verification

To see the VPN Sticky SA status, run this command in the Expert mode:

```
# g_fw ctl get int fwha_vpn_sticky_tunnel_enabled
```

```
-*- 12 blades: 1_01 1_02 1_03 1_04 1_05 1_06 2_01 2_02 2_03 2_04 2_05 2_06 -*-
fwha_vpn_sticky_tunnel_enabled = 0
```

Notes:

- Only outbound sticky SA connections are synchronized.
- Connections are not synchronized to all SGMS.

  To synchronize connections to all SGMs, run:
  ```
  # asg_lte_config enable
  ```

Configuring SCTP Acceleration on SGMs

To enable SCTP acceleration:

Run:

```
> sim feature sctp on
```

To disable SCTP acceleration:

Run:

```
> sim feature sctp off
```

Notes:

- You must configure SCTP in SmartDashboard before you can use this feature. See the R76 LTE Release Notes http://downloads.checkpoint.com/dc/download.htm?ID=29339 for detailed information and configuration procedures.
- If SCTP acceleration is activated and SCTP inspection is deactivated, the Performance Pack accelerates all SCTP packet types.

Configuring SCTP NAT on SGMs

SCTP NAT overrides the currently defined NAT policy. When this feature is not activated, SCTP connections do not use NAT.

To activate SCTP NAT:

Run:

```
> fw ctl set int fwx_enable_sctp_nat 1
```

To deactivate SCTP NAT:

Run:

```
> fw ctl set int fwx_enable_sctp_nat 0
```

System Optimization

In This Section:

Firewall Connections Table Size for VSX Gateway

You configure the Firewall connections table for VSX Gateway, Virtual Systems and other VSX virtual devices in SmartDashboard.

To configure the Firewall connections table:

1. Open the **Virtual Device** object in SmartDashboard.
2. Select the applicable virtual device.
3. Select **Optimizations** in the navigation tree.
4. On the Optimizations page, select Manually in **Calculate the maximum limit for concurrent connections**.
5. Enter or select a value.

Using the Fast Accelerator (sim fastaccel)

The Fast Accelerator lets you define trusted connections to allow bypassing of Medium Pass inspection (of Application Control, URL Filtering, Anti-Bot, anti-Virus, and Threat Emulation connections).

Those connections are handled in the regular way by SecureXL, while bypassing Medium Pass inspection, which requires forwarding to the Firewall.

This feature significantly improves throughput for these trusted high volume connections and reduces CPU consumption.

Syntax

```
sim fastaccel add|delete <source_ip> <source_port> <dest_ip> <dest_port> <protocol>
sim fastaccel conns
sim fastaccel -h
```

Parameter	Description
`Add`	Add a new trusted connection
`Delete`	Delete a trusted connection
`Conns`	Show all trusted connections
<*source_ip*>	Connection - Source IP address and optional subnet
<*source_port*>	Connection - Source port
<*dest_ip*>	Connection - Destination IP address and optional subnet
<*dest_port*>	Connection - Destination port
<*protocol*>	IP Protocol - (TCP=6, UDP=17)
`-h`	Show help information

Notes

- You can use the @ character as a wildcard to represent a valid parameter value.
- Enter the subnet in the `/nn` format. If you do not define a subnet, connection is defined as one, specified IP address.

Examples

To add a new trusted connection:

```
sim fastaccel add 1.1.1.1 80 2.2.2.0/24 19186 6
sim fastaccel add 192.168.0.0/16 @ @ 16 17
```

In the second example, the source connection is all addresses in the 192.168.0.0/16 subnet from all valid ports. The destination connection is defined as all IP addresses to port 16 over UDP.

To delete a trusted connection:

```
sim fastaccel delete 1.1.1.1 80 2.2.2.0/24 19186 6
```

To show trusted connections:

```
# sim fastaccel conns

### Source           SPort  Destination      DPort  PR
-------------------  -----  ---------------  -----  ---
  1)        1.1.1.1     80          2.2.2.0  19186   6
  2)    192.168.2.0     80      192.168.0.0     16  17
```

Known Limitation

Once you create a connection, you cannot disable or enable it. You can only delete the connection and then add the connection later.

Fastaccel does not support IPv6.

Reserved Connections

Normally, when the connection table limit is reached, no more connections are allowed, even ones critical for operating and managing the gateway. The reserved connections feature allows the gateway to process these critical connections, even after the connections table limit is reached. There is a user defined amount of space that is reserved in the connections table for these critical connections. If the Rule Base allows these connections, they are allowed even if no other connections can be accepted.

For example, when the connections table limit is reached, the administrator may not be able to install a new policy that increases the connections limit or open other necessary connections, such as SSH to the gateway.

Enforcing the reserved connections limit

By default, the number of reserved connections is limited to 2000. The actual limit of the connections table is increased by this amount.

Before a new connection is recorded, the system makes sure that there is sufficient space in the connections table. If connections table limit is reached, the connection is recorded if it satisfies these conditions:

- The limit is below the limit sum of connections table limit and reserved connections limit.
- Connection matches one of the rules in the reserved connections table

If not, the connection is not recorded.

In VSX, Reserved Connections are supported for VS0 only.

Syntax

```
# asg_reserved_conns
Please choose one of the following:
-----------------------------------
1) Print reserved connections table
2) Add new reserved connection rule
3) Delete reserved connection rule
4) Exit
>
```

To show the reserved connections table:

Enter: 1

Output

```
Idx  Source           Mask  Destination      Mask  DPort  Ipp    Interface
---  ---------------  ----  ---------------  ----  -----  -----  ------------
 1)          0.0.0.0     0          0.0.0.0     0   1129      6         Sync
 2)          0.0.0.0     0          0.0.0.0     0   1130      6         Sync
 3)          0.0.0.0     0          0.0.0.0     0   4444      6         Sync
 4)          0.0.0.0     0          0.0.0.0     0     22      6         Sync
 5)          0.0.0.0     0          0.0.0.0     0   8888      6         Sync
 6)          0.0.0.0     0          0.0.0.0     0   2010      6         Sync
 7)          0.0.0.0     0          0.0.0.0     0   1131      6         Sync
 8)          0.0.0.0     0          0.0.0.0     0   1132      6         Sync
 9)          0.0.0.0     0          0.0.0.0     0    256      6         Sync
10)          0.0.0.0     0          0.0.0.0     0      0      1         Sync
11)          0.0.0.0     0          0.0.0.0     0   8116     17         Sync
```

```
12)          0.0.0.0     0          0.0.0.0     0      0      1      eth1-CIN
13)          0.0.0.0     0          0.0.0.0     0     22      6      eth1-CIN
14)          0.0.0.0     0          0.0.0.0     0     23      6      eth1-CIN
15)          0.0.0.0     0          0.0.0.0     0    161     17      eth1-CIN
16)          0.0.0.0     0          0.0.0.0     0    623     17      eth1-CIN
17)          0.0.0.0     0          0.0.0.0     0      0      1      eth2-CIN
18)          0.0.0.0     0          0.0.0.0     0     22      6      eth2-CIN
19)          0.0.0.0     0          0.0.0.0     0     23      6      eth2-CIN
20)          0.0.0.0     0          0.0.0.0     0    161     17      eth2-CIN
21)          0.0.0.0     0          0.0.0.0     0    623     17      eth2-CIN
22)          0.0.0.0     0          0.0.0.0     0     22      6           Any
23)          0.0.0.0     0          0.0.0.0     0    256      6           Any
24)          0.0.0.0     0          0.0.0.0     0  18191      6           Any
25)          0.0.0.0     0          0.0.0.0     0  18192      6           Any
Press enter to continue
```

Field	Description
Idx	Rule number
Source	Source IP If the IP is 0.0.0.0, all IPs are allowed.
Mask	Subnet mask for the Source
Destination	Destination IP If the IP is 0.0.0.0, all IPs are allowed.
Mask	Subnet mask for the Destination
DPort	TCP/UDP Port This is ignored with non-TCP/UDP traffic.
Ipp	IP protocol number
Interface	Interface for this rule

To add a reserved connection rule:

1. Enter: 2
2. Follow the directions on the screen.

```
Enter source IP [0.0.0.0]:
>10.10.10.10
Enter source IP mask length [0]:
>24
Enter destination IP [0.0.0.0]:
>20.20.20.0
Enter destination IP mask length [0]:
>24
Enter destination port [0]:
>0
Enter IP protocol number (for example: tcp = 6, udp = 17):
>6
Enter interface number [0 = Any]:
0: Any
1: eth1-Mgmt4
2: eth2-Mgmt4
3: BPEth0
```

```
4: BPEth1
5: eth1-Mgmt1
6: eth1-CIN
7: eth1-01
8: eth2-Mgmt1
9: eth2-CIN
10: eth2-01
11: Sync
>0
OK to insert new reserved conn rule: <10.10.10.10/24, 20.20.20.0/24, 0, 6, Any> ?
(y/n)
>y
entry inserted, rule will apply when new connection will be opened
Press enter to continue
```

To make sure that the feature is configured correctly:

1. Confirm that the value of the kernel global parameter `fwconn_reserved_conn_active` is set to: `1`
2. Run `asg_reserved_conns` and enter: `1`
3. Run `fw tab -t reserved_conns_table` and confirm that the table contains the entries for the rules above.
4. Confirm that the contents of `$FWDIR/bin/reserved_conns_table` has rules of this feature.

To debug the feature:

1. Set the kernel global parameter `fwreserved_conns_debug` to: `1`
2. Use the `CONN` kernel debug flag to see reserved connections related debugs.

To troubleshoot the feature:

1. Run:
   ```
   # fw tab -t reserved_conns_table
   ```
2. Confirm that the table contains the entries for the rules in this feature.
3. Confirm that the contents of `$FWDIR/bin/reserved_conns_table` has rules of this feature.

   ⚠ Important - Do not make changes to this file.

4. Delete all current rules from the kernel and reload the rules from `$FWDIR/bin/reserved_conns_tab`:
   ```
   # asg_reserved_conns -f
   ```
   It is useful if there were changes in network interface names or if `$FWDIR/bin/reserved_conns_table` was edited directly.

Configuration

The feature works after installation without additional configuration.

The rules are stored in:

```
$FWDIR/bin/reserved_conns_table
```

The feature uses these kernel global variables:

Variable	Description
`fwconn_reserved_conn_active`	Enables or disables the feature Valid values: • 1 - Enabled • Any other integer: Disabled
`fwconn_reserved_limit`	Maximum allowed number of entries in `$FWDIR/bin/reserved_conns_table` Default: 2000

Policy Acceleration – SecureXL Keep Connections

To allow flow acceleration while a policy is pushed to the system:

In SmartDashboard, under the Gateway's properties, select **Other > Connection Persistence > Keep all connections**.

Note - This is enabled if SecureXL, and only the Firewall Software Blade are enabled.

Legacy Mode

To allow **Keep all connections** while disabling SecureXL keep connections:

In `$FWDIR/boot/modules/fwkern.conf,` set `cphwd_policy_accel` to: `0`

Verification

After policy installation, delete the old policy templates.

To make sure the templates of the old policy were deleted:

1. Run:
   ```
   # g_fwaccel stats
   ```
2. Save the old value of the `Policy deleted tmpl` statistics.
3. Install the policy.
4. Run:
   ```
   # g_fwaccel stats
   ```
5. Confirm that the templates were deleted.

VPN Performance Enhancements

These VPN performance enhancements are included in this release:

- **SPI Based Traffic Distribution for SSM160** - Uses all SGMs to handle VPN traffic based on the SPI instead of the IP address
- **SPI affinity** - Better traffic assignment to SGM CPU cores
- **VPN Templates** - Accelerates the session rate by adding VPN Templates to the SecureXL technology

SPI Distribution on SSM160 (asg dxl spi)

By default, the SSM160 distributes traffic to SGMs based on the IP address in the packet header. This methodology can be inefficient when working with a small number of remote peers in a Site-To-Site VPN topology. This is because the SSM160 only sees the VPN tunnel IP address and causes distribution only to some SGMs.

To resolve this issue, you can enable SPI distribution for VPN traffic. Run this command in `gclish` mode.

```
# set distribution spi mode on|off
```

⚠ **Important** - You must not enable SPI distribution for the LTE mode ("Working with LTE Features" on page 244) or when working with 3rd party VPN peers.

When you enable SPI distribution, you must also run:

```
# g_update_conf_file fwha_vpn_sticky_tunnel_enabled=0
```

When you disable SPI distribution in LTE mode or with a 3rd party peer, you must also run:

```
# g_update_conf_file fwha_vpn_sticky_tunnel_enabled=1
```

Note – SPI distribution mode is disabled by default.

SPI Affinity (asg_spi_affinity)

The `asg_spi_affinity` command helps you improve VPN performance with more efficient traffic assignment to SGMs and SGM cores. Typically, most VPN traffic goes to the same tunnel IP addresses. Because traffic is usually assigned to SGMs based on the destination IP address, VPN traffic is frequently assigned to the same SGMs. The solution is to assign VPN traffic to SGMs based on the SPI field in the packet header as an alternative to the IP address.

A related issue occurs with Multi-core VLAN traffic, where traffic is assigned to CPU cores based on IP addresses. As with VPN traffic, `asg_spi_affinity` can also assign VLAN traffic to CPU cores based on the SPI field.

Run this command in Expert mode.

Syntax

```
# asg_spi_affinity mode|vlan <ssm_id> on|off
# asg_spi_affinity verify
```

Parameter	Description
`mode`	Configure VPN affinity for specified SSM.
`vlan`	Configure VLAN affinity for the specified SSM interfaces.
`verify`	Show SPI affinity status.

Parameter	Description
<*ssm_id*>	SSM ID Valid values: - Integer between `1` and 4 - `all` - All SSMs
`on|off`	Enable or disable SPI affinity. You must enable vlan and mode (VPN) affinity separately.

Notes:

- When some SSM interfaces are not configured as VLANs, we recommend that you enable VLAN affinity only if most traffic passes through VLAN interfaces.
- SPI affinity can affect the distribution of clear packets. We recommend that you use SPI affinity only if most of the inbound traffic is VPN traffic.

Examples

`# asg_spi_affinity mode 1 on` - Enable VPN affinity for SSM 1

`# asg_spi_affinity mode 2 off` - Disable VPN affinity for SSM 2

`# asg_spi_affinity vlan all on` - Enable VLAN affinity for all SSM interfaces

`# asg_spi_affinity vlan all off` - Disable VLAN affinity for all SSM interfaces

VPN Templates

VPN templates accelerate the session rate, particularly for short connections (HTTP, DNS). These templates, which are part of the SecureXL template set, let you create new connections in the acceleration layer. They only send a notification to the Firewall layer if the connection is too long or if an F2F attack is detected. VPN templates are enabled by default.

To enable VPN templates:

Change `cphwd_offload_vpn_templates` to: 1

To disable VPN templates:

1. Run:
   ```
   > update_conf_file fwkern.conf cphwd_offload_vpn_templates=0
   ```
2. Reboot all SGMs.

Using Third Party VPN Peers with Many External Interfaces

When you use third-party VPN peers and have multiple external interfaces on the 61000/41000 Security System, you must configure the SGMs and the Management Server.

To configure the 61000/41000 Security System:

1. Run this command on the SMO:
   ```
   # g_update_conf_file $FWDIR/modules/vpnkern.conf ipsec_use_p1_src_ip=1
   ```
2. Reboot all SGMs.

To configure the Management Server:

1. Open `/opt/CPR76CMP-R77/lib/vpn_table.def` in a text editor.
2. Add this line to the configuration file:
   `dynamic_ipsec_source_address = dynamic sync keep expires EX_INFINITE;`
3. In SmartDashboard, install policy.

SCTP Acceleration

To enable SCTP Acceleration:

1. In SmartDashboard, create SCTP as **Other** using IP protocol 132
2. Enable **Accept Replies** in the **Advanced** tab of the SCTP service.
3. On the 61000/41000 Security System, **c**onnect to the SMO in Expert mode:
   ```
   > shell
   ```
4. Open: `$FWDIR/boot/modules/fwkern.conf` for editing. If the file does not exist, create it.
5. Add `sxl_accel_proto_list=132` to the file.
6. Open `$PPKDIR/boot/modules/simkern.conf` for editing. If the file does not exist, create it.
7. Add `sim_accel_non_tcpudp_proto=1` to the file.
8. Copy the file to all SGMs:
   ```
   # g_cp2blades $FWDIR/boot/modules/fwkern.conf
   # g_cp2blades $PPKDIR/boot/modules/simkern.conf
   ```
9. Reboot all SGMs:
   ```
   reboot -b all
   ```

Configuring DNS Session Rate

To improve the DNS session rate, the 61000/41000 Security System includes these enhancements:

- **Delayed Connection -** When a DNS connection matches a SecureXL template, the 61000/41000 Security System firewall is not immediately notified. The notification is delayed using the global parameter: `cphwd_udp_selective_delay_ha`. After a delay is set, the connection is handled fully by the acceleration device.

  Note - If the connection is not fully handled (and closed) by the acceleration device during the set delay period, the firewall is notified in the usual manner.

- **Delete on Reply -** After the DNS reply is received, the connection is immediately deleted from the gateway instead of being kept for an additional 40 seconds (the UDP connection default timeout).

To improve the DNS session rate:

Run:

```
> fw ctl set int cphwd_udp_selective_delay_ha <delay_secs>
> fwaccel off
> fwaccel on
```

To make sure that DNS connections are delayed by the set value:

1. Open a number of DNS connections from the same client to the same server.
2. Run:

```
> fwaccel templates
```

```
Source          SPort Destination     DPort PR  Flags     Conns  Open   LCT  DLY
--------------- ----- --------------- ----- -- --------- ------ ------ ---- ---
    10.33.87.12     *   192.168.15.31    53 17 .........     25      0    2  30
```

The number under **DLY** should match <*delay_secs*>.

Note – The default value for this parameter is 30 seconds. The maximum value is 60

To **Enable** or **Delete on Reply**:

1. Make sure SmartDashboard is disconnected.
2. Open the Check Point Database Tool at: C:\Program Files (x86)\Check Point\SmartConsole\R77.30\PROGRAM\GuiDBedit.exe.
3. Go to **Services>domain-udp**.
4. In **domain-udp**, change the value of delete_on_reply to "true".
5. Save changes using the **File** menu.
6. Exit Database Tool.
7. Launch SmartDashboard.
8. Install policy.

To make the enhancements **Permanent**:

Update `fwkern.conf`:

```
> update_conf_file fwkern.conf cphwd_udp_selective_delay_ha=<delay>
```

To turn off the enhancements:

To turn off **Delayed Connection** and **Delete on Reply**:

- Run:
  `fw ctl set int cphwd_udp_selective_delay_ha 0`
  or
- Remove all services from: `cphwd_delayed_udp_ports`

Note - This disables both enhancements.

Extending Session Rate Enhancements to other UDP Services

Change the value of `cphwd_delayed_udp_ports` in `fwkern.conf` to extend the benefits of these two DNS session rate enhancements to other services. For example, to add UDP service 100 to the list, run:

```
> update_conf_file fwkern.conf cphwd_delayed_udp_ports=53,100,0,0,0,0,0,0
```

Notes:

- The number of services is limited to 8.
- The command must contain 8 values. If you configure less than 8 services, enter 0 for the others.

- This is the only way to extend the DNS session rate enhancements to other UDP services.

The fw ctl set int command is not supported.

- The configuration takes effect only after reboot.

Accelerated Drop Enhancement

Use Accelerated Drop Enhancement to enforce drop rules in SecureXL on new or accelerated connections, without policy installation.

To configure Accelerated Drop Enhancement:

1. Log in to Expert mode.
2. Edit `$PPKDIR/conf/sim_drop_rules.conf`
3. Run `asg_sim_dropcfg` on the local SGM.

Limitations:

- Accelerated Drop Enhancement does not support IPv6.
- Accelerated Drop Enhancement and the `sim template quota exclude` list (`sim tmplquota -f`) cannot be enabled at the same time.
- Accelerated Drop Enhancement enforces rules only if SecureXL is ON. For example, it does not enforce rules during policy installation.
- Accelerated Drop Enhancement is not supported for VSX environments.

Configuration File

Add the drop rules in this file only for the local SGM. Each line must contain one rule, and each rule must contain one or more parameters.

Parameter	Description
`src` <*Source IP*> [<`Subnet`>]	Subnet of the source is optional
`dst` <*Destination IP*> [<`Subnet`>]	Subnet of the destination is optional
`dport` <*Destination port*>	Valid port number
`proto` <*IP protocol*>	An integer that represents a protocol, according to the IANA standards (http://www.iana.org/assignments/protocol-numbers/protocol-numbers.xhtml).

Example:

```
src 1.1.1.0/24 dst 2.2.0.0/16 dport 53 proto 17
```

Control Commands

Use this syntax to send commands to the local SGM in command-line mode.

```
asg_sim_dropcfg [enforce [-before | -ext | -nolog] | disable] [status] [conf [-comp]]
[stats] [fix]
```

Command Parameters:

Parameter	Description
`enforce`	Apply configuration to SecureXL to start rule enforcement.
`enforce -before`	Test packets against drop rules, and then against a connection or a template. Use this option to apply drop rules to a new or an existing connection.
`enforce -ext`	Enforce drop rules only on external interfaces. Default is Enforce rules on all interfaces.
`enforce -nolog`	Disable automatic log sent to SmartConsole.
`disable`	Disable enforcement of rules.
`status`	Show configuration file and SecureXL configuration status.
`conf`	Show configuration file settings.
`conf -comp`	Compare configuration files between SGMs.
`stats`	Show drop counters for each SGM.
`fix`	Set a consistent configuration across SGMs. If this fails, disable Accelerated Drop Enhancement. Use this option for error recovery.

Examples:

To enforce drop rules in the configuration file on external interfaces of new and existing connections:

```
# asg_sim_dropcfg enforce -before -ext
```

To disable enforcement:

```
# asg_sim_dropcfg disable
```

Configuring Hyper-Threading

Hyper-threading lets a compatible operating system run more than one process run simultaneously on a CPU core. A Hyper-threading processor adds one or more logical processors, which the operating system sees as independent processors.

To enable Hyper-threading, run `g_cpconfig` in Expert mode.

Syntax

```
# g_cpconfig ht stat
# g_cpconfig ht enable
# g_cpconfig ht disable
# g_cpconfig ht show stat
```

Parameter	Description
`Stat`	Shows whether hyper-threading is enabled for the 61000/41000 Security System

Parameter	Description
`Enable`	Enable Hyper-threading
`disable`	Disable Hyper-threading
`show stat`	Shows the hyper-threading status for all SGMs

Notes:

Hyper-threading is enabled by default on the SGM260.

You must reboot all SGMs after you enable or disable hyper-threading.

Configuring CoreXL (g_cpconfig)

Use `g_cpconfig` to configure CoreXL the 61000/41000 Security System. The number of instances for the VSX Gateway is limited to the physical number of CPU cores on the 61000/41000 Security System.

Note – If you run this command in a Virtual System, the output applies to VS0.

Syntax

```
> g_cpconfig corexl stat
> g_cpconfig corexl enable <n> [-6 <k>]
> g_cpconfig corexl disable
> g_cpconfig corexl instances <n> [-6 <k>]
> g_cpconfig corexl show instances
> g_cpconfig corexl show stat
```

Parameter	Description
`stat`	Show current status and number of instances on all SGMs.
`enable <n> [-6 <k>]`	Enable CoreXL *<n>* - Number of IPv4 Firewall instances -6 *<k>* - Number of IPv6 Firewall instances. Valid values: 2 - 32 Default - 16
`disable`	Disable CoreXL.
`instances <n> [-6 <k>]`	Change the number instances *<n>* - Number of IPv4 Firewall instances -6 *<k>* - Number of IPv6 Firewall instances Valid values: 2 - 32 Default - 16
`show instances`	Show the number of instances on each blade

Parameter	Description
`show stat`	Show the status on each blade

Example - Enabling Cores

```
> g_cpconfig corexl enable 8 -6 8
```

```
-*- 5 blades: 1_01 1_02 2_01 2_02 2_04 -*-
rx_num for ixgbe interfaces was set to: 16

CoreXL was successfully enabled with 8 IPv4 and 8 IPv6 firewall instances.

Important: This change will take effect after rebooting all blades.
```

Example - Showing CoreXL status for each SGM

```
> g_cpconfig corexl show stat
```

```
blade 1_01 corexl is enabled
blade 1_02 corexl is enabled
blade 1_03 corexl is enabled
```

CoreXL configuration on a VSX system

When you change the number of CoreXL instances in a Security Gateway environment, all CPUs not assigned to CoreXL are assigned to Performance Pack. When you change the number of CoreXL instances in a VSX Gateway environment, you only change the number of user-mode threads. This has no effect on Performance Pack affinity. The number of CPUs assigned to Performance Pack does not change.

This example shows a system with 12 CPUs and 3 Virtual Systems where:

- Each Virtual Systems has 1 CoreXL instance.
- CPUs 0-7 are assigned to Firewall packet inspection.
- CPUs 8-11 are assigned to Performance Pack.

```
> g_cpconfig corexl instances 3
```

- The number of CoreXL instances (user-mode threads) changes from 1 to 3. Each Virtual System still has one CoreXL instance.
- CPUs 0-7 are still assigned to Firewall packet inspection.
- CPUs 8-11 are still assigned to Performance Pack.

VSX Affinity Commands (fw ctl affinity-s -d)

This section shows you how to use the `fw ctl affinity` command to set affinities in a VSX environment. When you run this command, the system automatically creates or updates the affinity configuration files. All affinity configurations are kept after reboot.

You can define specified processes as affinity exceptions. Affinity commands do not apply these processes. To define an exception, add the process name to the `$FWDIR/conf/vsaffinity_exception.conf` file. You cannot add kernel threads as affinity exceptions.

⚠ **Important** - Do not add Check Point processes to the exception list. This can cause system instability.

Affinity Priorities

When a CPU core has more than one affinity, the affinity is applied based on these priorities:

1. Firewall instance
2. Process
3. Virtual System

Setting Affinities

Use `fw ctl affinity-s -d` to set these CPU affinities:

- Firewall instance
- Processes
- Virtual System

You can set Firewall instance affinity to one or more CPUs on each Virtual System individually.

Syntax

```
> fw ctl affinity-s -d
> fw ctl affinity-s -d [-vsid <vs_ids>] -cpu <cpu_id>
> fw ctl affinity-s -d -pname <process> [-vsid <ranges>] -cpu <cpu_id>
> fw ctl affinity-s -d -inst <instance_id> -cpu <cpu_id>
```

Parameter	Description
`-s -d`	Set affinity for a VSX environment.
`-vsid` *<vs_ids>*	*<vs_ids>* can be: • No *<vs_ids>* (default) - Shows the current Virtual System context. • One Virtual System. • A comma-separated list of Virtual Systems (1, 2, 4, 5). • A range of Virtual Systems (VS 3-5). • `all` - Shows all Virtual Systems. **Note:** This parameter is only relevant in a VSX environment.
`-cpu` *<cpu_id>*	One or more CPU cores. You can define a range from which the system selects the instances. The format for a range is: *<from_cpu_id>-<to_cpu_id>*.
`-pname` *<process>*	Configure affinity for the specified process.
`-inst` *<instance_id>*	One or more Firewall instances. You can define a range from which the system selects the instances. The format for a range is: *<from_instance_id>-<to_instance_id>*.

Setting affinities for all SGMs from the SMO:

From `gclish`, run:

```
> fw ctl affinity-s -d <options>
```

From Expert mode, run:

```
# g_fw ctl affinity-s -d <options>
```

To set affinities for a specified SGM:

Run:

```
> blade <sgm_id>
> fw ctl affinity-s -d <options>
```

Setting Firewall instance affinity with ranges

This example creates two Firewall instance affinities for the Virtual System on context 1. One affinity is assigned to instance 0 and the other is automatically assigned from the range of instances 2-4. These instances are automatically assigned to CPU cores in the range of 0-2.

```
> vsenv 1
> fw ctl affinity-s -d -inst 0 2-4 -cpu 0-2

VDevice 0: CPU 0 1 2 - set successfully

Note: If there were previously configured processes/FWK instances, this operation
has overridden them and deleted their configuration files
Athens-ch01-02:0>
```

Setting VSX processes affinity (-pname)

Set the affinity of processes to one or more CPUs. You can use `-vsid` to set the affinity for a process to Virtual Systems in any context. If you do not use `-vsid`, the affinity of the current context is set.

```
> fw ctl affinity-s -d -pname cpd -vsid 0-1 -cpu 0 2

VDevice 0-1 : CPU 0 2 - set successfully
```

Virtual System affinity (-vsid)

Use `-vsid` to define an affinity for specified Virtual Systems. This example sets the affinity for Virtual System contexts 0 and 1 to CPU cores 0 and 2. If you do not use `-vsid`, this command sets the affinity for the current VSX context.

```
> fw ctl affinity-s -d -vsid 0-1 -cpu 0 2
VDevice 0-1 : CPU 0 2 - set successfully
```

Setting Affinity for all Virtual Systems (fw ctl affinity -s -d -fwkall)

Use `fw ctl affinity -s -d -fwkall` to assign the specified number of CPU cores to all Virtual Systems at one time.

Effect on Multi-queue settings for ixgbe interfaces

The use of this command to change the number of cores assigned to Virtual Systems, changes the number of cores available for **ixgbe** interface **rx queues**. Conversely, when you change the number of cores assigned to **ixgbe** interface queues, you also change the number of cores assigned to Virtual Systems.

For example, if your SGMs have 16 cores, and you assign 9 cores to Virtual Systems, the remaining 7 cores are available to the **ixgbe** interfaces.

Syntax

```
> fw ctl affinity -s -d -fwkall <cores>
```

Parameter	Description
`-s -d`	Set affinity for a VSX environment.
`-fwkall <cores>`	Defines the number of cores assigned to all Virtual Systems.

Example

This example assigns three cores to Firewall instances for all Virtual Systems.

```
> fw ctl affinity -s -d -fwkall 3
VDevice 0-2 : CPU 0 1 2 - set successfully
```

Note - You can run this command from the VS0 context only.

Monitoring Process Affinity (fw ctl affinity -l -x)

You can monitor the affinity of processes and Virtual Systems on a VSX Gateway. You can use the -vsid parameter to show the affinity for a process to the specified Virtual Systems.

Syntax

```
> fw ctl affinity -l -x [-vsid <vsid>] [-flags {e|h|k|n|t|o}]
```

Parameter	Description
`<vsid>`	Shows the affinity for processes for these Virtual System IDs. Use a dash to set a range of Virtual Systems.
e	Do not show processes that are affinity exceptions. You define affinity exceptions in: `$FWDIR/conf/vsaffinity_exception.conf`
h	Show CPU affinity mask in hexadecimal format.
k	Do not show kernel threads.
n	Show the process name instead of `/proc/<pid>/cmdline`
t	Show information about process threads.
o	Print the list to a file.

Example

```
> fw ctl affinity -l -x -vsid 1 -flags tn

-----------------------------------------------------------------------
|PID     |VSID |         CPU              |SRC|V|KT |EXC| NAME
-----------------------------------------------------------------------
|   4756 |   0 |                      all |   | |   |   | pm
|   4773 |   0 |                      all |   | |   |   | confd
|   4774 |   0 |                      all |   | |   |   | searchd
|   5008 |   0 |                      all |   | |   |   | |---searchd
|   4780 |   0 |                       all |   |  |    |  | httpd2
|   4781 |   0 |                       all |   |  |    |  | monitord
|  24700 |   0 |                       0 1 | P |  |    |  | |---cpd
|  24704 |   0 |                       0 1 | P |  |    |  | |---cpd
|  24705 |   0 |                       0 1 | P |  |    |  | |---cpd
|  22800 |   0 |                       all |   |  |    |  | mpdaemon
|  24523 |   0 |                       all |   |  |    |  | fwk_forker
|  24525 |   0 |                       all |   |  |    |  | fwk_wd
|  24573 |   0 |                   1 3 4 6 | P |  |    |  | fw
|  24667 |   0 |                   1 3 4 6 | P |  |    |  | |---fw
|  24668 |   0 |                   1 3 4 6 | P |  |    |  | |---fw
|  24670 |   0 |                   1 3 4 6 | P |  |    |  | |---fw
|  24671 |   0 |                   1 3 4 6 | P |  |    |  | |---fw
|  25412 |   0 |                   1 3 4 6 | P |  |    |  | |---fw
|  24642 |   0 |               2 3 4 5 6 7 | P |  |    |  | fwk0_dev
|  24643 |   0 |               2 3 4 5 6 7 | P |  |    |  | |---fwk0_0
|  30186 |   0 |                       all |   |  |    |  | clishd
------------------------------------------------------------------------
```

System Under Load

The System Under Load (SUL) feature delays SGM failover for a specified time (default=10 seconds) during periods of high system CPU utilization. This helps to prevent unnecessary SGM failovers caused by CCP packet transmission delays.

The system automatically turns on System Under Load when at least one SGM has kernel CPU usage above the specified threshold (default = 80%). SUL turns off automatically when no SGM has high CPU utilization for at least 10 seconds or when SUL is active for more than three minutes.

Logs

A log entry is generated for every System Under Load state change (ON/OFF). Only the SMO sends messages to the log server. This example shows System Under Load logs in SmartView Tracker.

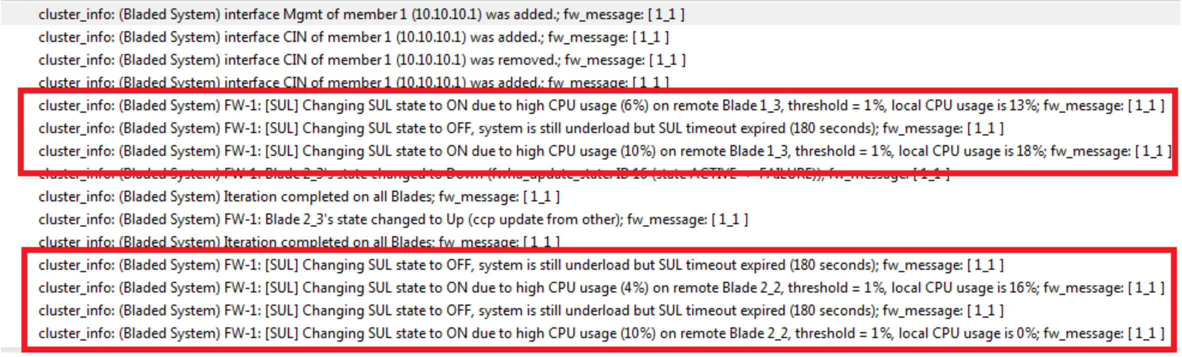

```
cluster_info: (Bladed System) interface Mgmt of member 1 (10.10.10.1) was added.; fw_message: [ 1_1 ]
cluster_info: (Bladed System) interface CIN of member 1 (10.10.10.1) was added.; fw_message: [ 1_1 ]
cluster_info: (Bladed System) interface CIN of member 1 (10.10.10.1) was removed.; fw_message: [ 1_1 ]
cluster_info: (Bladed System) interface CIN of member 1 (10.10.10.1) was added.; fw_message: [ 1_1 ]
cluster_info: (Bladed System) FW-1: [SUL] Changing SUL state to ON due to high CPU usage (6%) on remote Blade 1_3, threshold = 1%, local CPU usage is 13%; fw_message: [ 1_1 ]
cluster_info: (Bladed System) FW-1: [SUL] Changing SUL state to OFF, system is still underload but SUL timeout expired (180 seconds); fw_message: [ 1_1 ]
cluster_info: (Bladed System) FW-1: [SUL] Changing SUL state to ON due to high CPU usage (10%) on remote Blade 1_3, threshold = 1%, local CPU usage is 18%; fw_message: [ 1_1 ]
cluster_info: (Bladed System) FW-1: Blade 2_3's state changed to Down (fwha_update_state: ID 16 (state ACTIVE -> FAILURE)); fw_message: [ 1_1 ]
cluster_info: (Bladed System) Iteration completed on all Blades; fw_message: [ 1_1 ]
cluster_info: (Bladed System) FW-1: Blade 2_3's state changed to Up (ccp update from other); fw_message: [ 1_1 ]
cluster_info: (Bladed System) Iteration completed on all Blades; fw_message: [ 1_1 ]
cluster_info: (Bladed System) FW-1: [SUL] Changing SUL state to OFF, system is still underload but SUL timeout expired (180 seconds); fw_message: [ 1_1 ]
cluster_info: (Bladed System) FW-1: [SUL] Changing SUL state to ON due to high CPU usage (4%) on remote Blade 2_2, threshold = 1%, local CPU usage is 16%; fw_message: [ 1_1 ]
cluster_info: (Bladed System) FW-1: [SUL] Changing SUL state to OFF, system is still underload but SUL timeout expired (180 seconds); fw_message: [ 1_1 ]
cluster_info: (Bladed System) FW-1: [SUL] Changing SUL state to ON due to high CPU usage (10%) on remote Blade 2_2, threshold = 1%, local CPU usage is 0%; fw_message: [ 1_1 ]
```

System Under Load log entries are typically a symptom of intensive CPU activity. See Hardware Monitoring and Control (on page 113) to learn how to resolve these issues.

Working with Jumbo Frames

The 61000/41000 Security System supports Jumbo Frames with a total size of up to 9,146 bytes for the SSM60 and 12,200 bytes for the SSM160.

Note: Carefully calculate the MTU. For example: IPSEC or GRE traffic adds bytes to the header, and this leaves fewer bytes for the data payload.

The next topics explain how to configure Jumbo Frames:

- Configure Jumbo Frames on the SGM.
- Configure Jumbo Frames on VSX ("Configuring Jumbo Frames on VSX" on page 266).
- Run tests on the SSMs, SGMs, and SGM interfaces before you use the system for Jumbo Frames production traffic.
  - Confirming Jumbo Frames Configuration on SSM160 (on page 268)
  - Confirming Jumbo Frames Configuration on SSM60 (on page 267)
  - Confirming Jumbo Frames on SGMs and SGM Interfaces (on page 267)

Configuring Jumbo Frames on Security Gateway

Configuring SGMs (set interface)

Use `set interface` <*if_name*> `mtu` <*size*> to configure Jumbo Frames for each applicable interface on an SGM.

To enable Jumbo Frames, you must set the MTU on at least one interface to more than 1500. In a Dual Chassis environment, this enables Jumbo Frames on both Chassis.

To enable Jumbo Frames on the SSM60, you must also do the procedure in SSM60.

Note – This command can take several seconds to work.

Syntax

```
> set interface <if_name> mtu <size>
```

Parameter	Description
<*if_name*>	Interface name as defined in the operating system
<*size*>	MTU size Allowed values: • `68 - 9,124` for **SSM60** • `68 - 12,288` for **SSM160**

Example

```
> set interface eth1-01 mtu 9000
```

Output

```
1_02:
Note: MTU changes are propagated to the SSMs. Use "asg_jumbo_conf show" to validate
changes
```

Configuring Jumbo Frames on SSM60

Configure Jumbo Frames for each SSM60 in the Chassis. In a Dual Chassis system, do this procedure for both Chassis.

1. Connect to the SSM using telnet.

   The default password is `admin`.

2. Go to Enable mode:
   ```
   > en
   ```
3. Go to the Configuration terminal:
   ```
   # conf t
   ```
4. Configure all the downlink interfaces:
   ```
   # interface range 1/2/1-1/14/1
   ```
5. Configure the MTU:
   ```
   # packet-size-limit 9146
   ```
6. Configure the required front panel ports:
   ```
   # interface range 1/2/1-1/14/1
   ```
   Interfaces 1/15/1 – 1/15/5 = SSM ports 1-5.
7. Set the required MTU:
   ```
   # packet-size-limit 9146
   ```
8. Close the Configuration terminal and save the configuration:
   ```
   # end
   # write
   ```

Example

```
# telnet 198.51.100.32
Trying 198.51.100.32...
Connected to 198.51.100.32.
Escape character is '^]'.

User Access Verification
Password:
> en
# conf t
# interface range 1/2/1-1/14/1
# packet-size-limit 9146
# interface range 1/15/1-1/15/5
# packet-size-limit 9146
# end
# write
```

Configuring Jumbo Frames on VSX

Configure MTU in the interface's properties, as described in the Configuring Jumbo Frames section of *VSX Administration Guides:* R76 (http://supportcontent.checkpoint.com/documentation_download?ID=22932) R77 http://supportcontent.checkpoint.com/documentation_download?id=24802 R80 http://supportcontent.checkpoint.com/documentation_download?id=46537.

1. Open SmartDashboard and connect to the Security Management Server or <mds>.
2. Open the 61000/41000 Chassis object.
3. Click **Topology**.
4. Edit the relevant interface.
5. On the **General** tab, set the MTU.

   Valid values:

   - SSM160 - 68 to 12,288 bytes
   - SSM60 - 68 to 9,124 bytes
6. Click **OK**.

7. Install the policy on the 61000/41000 object.

Confirming Jumbo Frames on SGMs and SGM Interfaces

To confirm configuration on SGMs and SGM Interfaces (asg_jumbo_conf show):

Use `asg_jumbo_conf` show to:

- Make sure that Jumbo Frames are enabled on the SGMs
- See the configured MTU values on SGM interfaces configured for Jumbo Frames

Syntax

```
# asg_jumbo_conf show [-v]
```

Parameter	Description
-v	Detailed report (verbose)

Example

```
# asg_jumbo_conf show -v
```

Output

```
Jumbo frames are enabled on SGMs (SSM1 max MTU: 12288 SSM2 max MTU: 12288 )
Retrieving SSMs Jumbo frames configuration
Chassis1

SSMs:
Jumbo frames are enabled on SSM1
Jumbo frames are enabled on SSM2
Interfaces MTU configuration:
interface:BPEth0:mtu 12288
interface:BPEth1:mtu 12288
The MTU of all the interfaces which are not in the list is 1500
```

Confirming Jumbo Frames Configuration on SSM60

To run the validation test on the SSM60:

1. Connect to the SSM with telnet.

   The default password is `admin`.
2. Use **Enable Mode**:
   ```
   > en
   ```
3. Show the running configuration:
   ```
   # show run
   ```
4. Make sure that all applicable interfaces (downlinks and front panel ports) show the required packet size limit.
   ```
   # telnet 198.51.100.32
   Trying 198.51.100.32...
   Connected to 198.51.100.32.
   Escape character is '^]'.
   ```

```
User Access Verification
Password:
FI_cp>en
#show run
.
.
.
!
interface 1/2/1
flow-control disable
packet-size-limit 9146
!
```

Confirming Jumbo Frames Configuration on SSM160

To run the validation test on the SSM160:

1. Show the Jumbo Frames configuration on the specified SSM:

   `> asg_chassis_ctrl jumbo_frames show <`*`ssm_id`*`>`
2. Show the configured MTU on the specified port:

   `> asg_chassis_ctrl get_port_mtu <`*`ssm_id`*`> <`*`port_id`*`>`

Example

```
> asg_chassis_ctrl jumbo_frames show 1
Jumbo frames are enabled on SSM1
> asg_chassis_ctrl get_port_mtu 1 1
MTU of port 1 on SSM1 is 1544
```

Disabling Jumbo Frames

Use `set interface` to disable Jumbo Frames and change the MTU of each interface to 1500 or lower.

To disable Jumbo Frames on a Security Gateway:

Syntax

```
set interface <interface> mtu {1500..1}
```

Example

```
> set interface eth1-01 mtu 1500
```

Output

```
1_02:
```

Note - MTU changes are propagated to the SSMs. Use `asg_jumbo_conf show` to validate changes.

To disable Jumbo Frames on a VSX:

1. Connect to the Security Management server with the SmartDashboard.
2. Open the 61000/41000 Chassis Object.
3. Open **Topology**.

4. Edit the interface.
5. On the **General** tab, set the MTU.

TCP MSS Adjustment

TCP MSS Adjustment allows MSS (Maximum Segment Size) clamping of TCP traffic. This enables the configuration of the MSS that is part of OPTIONS in the TCP header.

This feature provides a method to prevent fragmentation when the MTU value on the communication path is lower than the MSS value.

Syntax

```
> fw ctl set int clamp_mss|mss_value <num>
```

Parameter	Description
clamp_mss <*num*>	Enable or Disable MSS Adjustment: • 0 - Disable (default) • 1 - Enable
mss_value <*num*>	MSS Value If this is set to `0,` the MSS value is based on the interface's MTU.

Notes:

- If you want the modified parameters including state (ON/OFF), to be persistent, use `g_update_conf_`file in Expert mode to add them to: $FWDIR/boot/modules/fwkern.conf
- Verification - You can use Packet Sniffers to make sure that MSS is clamped when the feature is enabled according to the configuration.
- MSS value is applied on all interfaces, including Management.

Debugging

1. Enable SIM debug:
   ```
   > sim dbg -m pkt + pkt
   ```
2. Start fw debugging:
   ```
   > fw ctl zdebug + packet
   ```
3. Look for output that contains the string: `MSS`

Working with Session Control (asg_session_control)

Based on a predefined set of rules, use `asg_session_control` to set the rate at which new communication sessions are opened. `asg_session_control` is also known as **Session Rate Throttling**. You can only run `asg_session_control` from Expert mode.

Create session control rules in the `$FWDIR/conf/control_rules` file.

Note - Session rate control is disabled by default.

Syntax

```
# asg_session_control apply|disable|stats|verify
```

Parameter	Description
No parameters	Shows command syntax and helpful information
`apply`	Applies session rate rules to all SGMs
`disable`	Disables session rate rules for all SGMs
`stats`	Shows all session rate rules and dropped traffic statistics
`verify`	Makes sure that the session rate rules are the same on all SGMs

Defining Session Control Rules

You define session rate rules in the `$FWDIR/conf/control_rules` file. Use one line for each rule.

Each rule must contain the `limit` parameter. The other parameters are optional.

⚠ **Important** - Define rules as specifically as possible, so that more than one rule cannot apply to the same traffic. Overlapping rules can cause unpredictable results. We recommend that you explicitly define all parameters in each rule.

Rule Syntax

```
[src <ip>/<mask>] [dst <ip>/<mask>] [dport <port>] [proto <protocol_id>] [limit
<rate>] [limit_ongoing 0|1]
```

Parameter	Description
`src` *<ip>/<mask>*	Source IP address and net mask
`dst` *<<ip>/<mask>*	Destination IP address and net mask
`dport` *<port>*	Destination port
`proto` *<protocol_id>*	Protocol code, typically 6 (TCP) or 17 (UDP) To learn more about protocol codes, see IANA protocol codes http://www.iana.org/assignments/protocol-numbers/protocol-numbers.xhtml.
`limit` *<rate>*	Maximum number of new connections allowed per second
`limit_ongoing 0|1`	• `0` - Do not limit the number of packets on an established connection • `1` - Limit the number of packets on an established connection

Rule Examples

```
src * dst 1.1.1.0/24 dport 67 proto 17 limit 20 limit_ongoing 1
```

This rule defines a limit of 20 new connections per second for traffic going from all sources to:

- Network 1.1.1.0/24
- Port 67
- Using protocol 17 (UDP)
- Including current connections

```
dst 1.1.1.1/32 dport 80 proto 6 limit 13
```

This rule defines a limit of 13 new connections per second for traffic going from all sources to:

- Network 1.1.1.1/32
- Port 80
- Using protocol 6 (TCP)

Notes

- New connections above the specified limit are dropped.
- If you do not include a parameter, the rule applies to all values for that parameter. For example, if you do not include the `src` parameter, the rule applies to all servers.
- The * character as a parameter value explicitly says that a rule applies to all values.

Enabling and Disabling Session Control

To enable Session Control:

1. Define Session Control rules.
2. Run:

   ```
   # asg_session_control apply
   ```

To disable Session Control:

Run:

```
# asg_session_control disable
```

Output:

```
-*- 2 blades: 1_01 1_02 -*-
Resetting session rate entries
Session rate entries configured successfully
```

Applying Session Control Rules

To apply Session Control rules:

Run:

```
# asg_session_control apply
```

Output

```
-*- 2 blades: 1_01 1_02 -*-
Rule ID Source            Destination        DPort PR  Limit Ongoing
------- ----------------- ------------------ ----- --- ----- -------
      1                 *         1.1.1.0/24    67  17    20       1
      2                 *         2.2.2.2/32    80   6    13       0
```

The output shows the Session Control rules that applied.

Showing Session Control Statistics

To show Session Control statistics:

Run:

```
# asg_session_control stats
```

Output

```
1_01:
Rule ID Source             Destination        DPort PR  Limit Drops         Attempts
------- ------------------ ------------------ ----- --- ----- ------------- -------------
       1                 *         1.1.1.0/24    67  17    20             3 19
       2                 *         2.2.2.2/32    80   6    13             0 12

1_02:
Rule ID Source             Destination        DPort PR  Limit Drops         Attempts
------- ------------------ ------------------ ----- --- ----- ------------- -------------
       1                 *         1.1.1.0/24    67  17    20             0 19
       2                 *         2.2.2.2/32    80   6    13             2 13
```

The output shows the session control rules for each SGM and the connections dropped by each rule.

Acceleration Not Disabled Because of Traceroute Rule (asg_tmpl_special_svcs)

This feature safely prevents security policy rules with the Traceroute service from disabling acceleration for all subsequent rules.

Syntax

```
> asg_tmpl_special_svcs on|off
```

Parameter	Description
`on`	Acceleration is not disabled because of Traceroute rules
`off`	Acceleration is disable because of Traceroute rules

Example

```
> asg_tmpl_special_svcs on
```

- This feature requires a patch on the Management side. To get it, contact Check Point support.
- For this feature to work correctly, the Traceroute service object in SmartDashboard must remain with default settings and not customized.

Improving Inbound HTTPS Performance

You can improve the performance of inbound HTTPS traffic from outside the organization.

To improve the performance of inbound HTTPS:

Run:

```
> fw ctl set int choose_active_streaming 0
```

To restore the default HTTPS performance settings:

Run:

```
> fw ctl set int choose_active_streaming 1
```

Supported SSL Ciphers

These SSL ciphers are supported on internal HTTPS servers when the parameter `choose_active_streaming` is set to `0`:

- RSA+AES
- RSA+RC4
- RSA+3DES

You must update the list of supported SSL ciphers on the protected HTTPS servers.

61000/41000 Security System Concepts

In This Section:

Single Management Object and Policies

Single Management Object is a Check Point technology that manages the 61000/41000 Security System as one large Security Gateway with one management IP address. All management tasks are handled by one SGM (the SMO Master), which updates all other SGMs. All management tasks, such as Security Gateway configuration, policy installation, remote connections and logging are handled by the SMO master. The active SGM with the lowest ID number is automatically assigned to be the SMO.

Use this command to identify the SMO and see how tasks are distributed on the SGMs:

```
> asg stat -i tasks
```

```
Chassis ID: 1
-------------
Task (Task ID)          SGM ID

General    (1)          3
LACP       (2)          4
CH Monitor (3)          5

Chassis ID: 2
-------------
Task (Task ID)          SGM ID

SMO        (0)          2(local)
DR Manager (4)          2(local)
General    (1)          3
LACP       (2)          4
CH Monitor (3)          5
```

Installing and Uninstalling Policies

To install a policy on the 61000/41000 Security System, select **Policy** > **Install** in SmartDashboard. The installation procedure includes these steps:

1. The Security Management server installs the policy on the SMO Master.
2. The SMO copies the policy to all SGMs.

3. Each SGM installs the policy locally.

During the installation, each SGM sends and receives policy status updates to/from the other SGMs. This is because the SGMs must install their policies in a synchronized manner. Policy installation has these stages:

- **Policy Started -** Policy installation started on the SGM.
- **Policy Ready2Finish** - Policy installation is completed, but the SGM is waiting for other SGMs to reach the same stage.
- **Policy Completed** - The policy is synchronized with the other SGMs.
- **Enforcing Security -** The SGM enforces the new policy.

Note - When installing the 61000/41000 Security System, SGMs enforce an initial policy where only the implied rules necessary for management are enforced.

To uninstall a policy:

Open a serial connection to the 61000/41000 Security System and run:

```
> asg policy unload
```

Notes:

- You cannot uninstall policies with SmartDashboard.
- To learn more about the working with policies, see `asg policy` ("Working with Policies (asg policy)" on page 275).

Working with Policies (asg policy)

`Use asg policy to do these policy-related actions:`

Action	Description
`verify`	Make sure that the correct policies are installed on all SGMs.
`verify_amw`	Makes sure that the correct Anti-malware policies are installed on all SGMs.
`unload`	Uninstall the policy from SGMs.

Syntax

```
> asg policy -h
> asg policy verify|verify_amw [-vs <vs_ids>] [-a] [-vs] [-v]
> asg policy unload [--disable_pnotes] [-a]
> asg policy unload --ip_forward
```

Parameter	Description
`-h`	Show syntax and help information.

Parameter	Description
`-vs <vs_ids>`	Shows verification results for each Virtual System. <*vs_ids*> can be: - No <*vs_ids*> (default) - Shows the current Virtual System context. - One Virtual System. - A comma-separated list of Virtual Systems (1, 2, 4, 5). - A range of Virtual Systems (VS 3-5). - `all` - Shows all Virtual Systems. **Note:** This parameter is only relevant in a VSX environment.
`-v`	Shows detailed verification results for SGMs in each Virtual System.
`-a`	Run the verification on both UP and DOWN SGMs.
`--disable_pnotes`	SGMs stay in the UP state without an installed policy.
`--ip_forward`	Enable IP forwarding.

Example - Detailed Virtual System Output

```
> asg policy verify -vs all -v
+------------------------------------------------------------------------------+
|Policy Verification                                          |
+-------+-------+-------------------+---------------+-----------------+--------+
|VS     |SGM    |Policy Name        |Policy Date    |Policy Signature |Status  |
+-------+-------+-------------------+---------------+-----------------+--------+
|0      |1_01   |Standard           |26Nov12 21:11  |996eee5e6        |Success |
|       |1_03   |Standard           |26Nov12 21:11  |996eee5e6        |Success |
|       |1_04   |Standard           |26Nov12 21:11  |996eee5e6        |Success |
|       |1_05   |Standard           |26Nov12 21:11  |996eee5e6        |Success |
|       |1_06   |Standard           |26Nov12 21:11  |996eee5e6        |Success |
|       |1_11   |Standard           |26Nov12 21:11  |996eee5e6        |Success |
|       |1_12   |Standard           |26Nov12 21:11  |996eee5e6        |Success |
+-------+-------+-------------------+---------------+-----------------+--------+
|1      |1_01   |Standard           |27Nov12 13:03  |836fa2ec1        |Success |
|       |1_03   |Standard           |27Nov12 13:03  |836fa2ec1        |Success |
|       |1_04   |Standard           |27Nov12 13:03  |836fa2ec1        |Success |
|       |1_05   |Standard           |27Nov12 13:03  |836fa2ec1        |Success |
|       |1_06   |Standard           |27Nov12 13:03  |836fa2ec1        |Success |
|       |1_11   |Standard           |27Nov12 13:03  |836fa2ec1        |Success |
|       |1_12   |Standard           |27Nov12 13:03  |836fa2ec1        |Success |
+-------+-------+-------------------+---------------+-----------------+--------+
|2      |1_01   |Standard           |26Nov12 21:11  |10eef9ced        |Success |
|       |1_03   |Standard           |26Nov12 21:11  |10eef9ced        |Success |
|       |1_04   |Standard           |26Nov12 21:11  |10eef9ced        |Success |
|       |1_05   |Standard           |26Nov12 21:11  |10eef9ced        |Success |
|       |1_06   |Standard           |26Nov12 21:11  |10eef9ced        |Success |
|       |1_11   |Standard           |26Nov12 21:11  |10eef9ced        |Success |
|       |1_12   |Standard           |26Nov12 21:11  |10eef9ced        |Success |
+-------+-------+-------------------+---------------+-----------------+--------+

+------------------------------------------------------------------------------+
|Summary                                                                       |
+------------------------------------------------------------------------------+
|Policy Verification completed successfully                                    |
+------------------------------------------------------------------------------+
```

Example - Uninstall Policy

```
> asg policy unload
You are about to perform unload policy on blades: all
All SGMs will be in DOWN state, beside local SGM. It is recommended to run the procedure
via serial connection

Are you sure? (Y - yes, any other key - no) y

Unload policy requires auditing
Enter your full name: ploni
Enter reason for unload policy [Maintenance]:
WARNING: Unload policy on blades: all, User: ploni, Reason: Maintenance
+-------------------------------+
|Unload policy                  |
+---------------+---------------+
|SGM            |Status         |
+---------------+---------------+
|1_3            |Success        |
+---------------+---------------+
|1_2            |Success        |
+---------------+---------------+
|1_1            |Success        |
+---------------+---------------+
|2_3            |Success        |
+---------------+---------------+
|2_2            |Success        |
+---------------+---------------+
|2_1            |Success        |
+---------------+---------------+

+--------------------------------------------------------------------------------+
|Summary                                                                         |
+--------------------------------------------------------------------------------+
|Unload policy completed successfully                                            |
+--------------------------------------------------------------------------------+
```

Note - We recommend that you run this command over a serial connection.

SGM Policy Management

Because the 61000/41000 Security System works as one large Security Gateway, all SGMs are configured with the same policy. When you install a policy from the management server, it first installs the policy on the SMO. The SMO copies the policy and SGM configuration to all SGMs in the UP state. When an SGM enters the UP state, it automatically gets the currently installed policy and configuration from the SMO. If there is no SMO (when there is only one SGM in the UP state), that SGM uses its local policy and configuration.

If there are problems with the policy or configuration on an SGM, you can manually copy the information from a different SGM.

An SGM configuration has these components:

- Firewall policy, which includes the Rulebase.
- Set of configuration files defined in the `/etc/xfer_files_list` file. This file contains the location of all related configuration files. It also defines the action to take if the copied file is different from the one on the local SGM.

Synchronizing Policy and Configuration between SGMs

The `asg_blade_config pull_config` command manually synchronizes Policies, and optionally, configuration files from a specified source SGM to the target SGM. The target SGM is the SGM on which you run this command.

To manually synchronize SGMs:

1. Run: `asg_blade_config pull_config`
2. Reboot the target CMA or run these commands:
   - cpstart
   - asg sgm_admin up

Note - You can run `asg stat -i all_sync_ips` to get a list of all SGM synchronization IP addresses.

Understanding the Configuration File List

The `xfer_file_list` file contains pointers to the related configuration files on an SGM. Each record defines the path to a configuration file, followed by the action to take if the imported file is different from the local file. This table shows an example of the record structure.

Context	File name and path	Action
`global_context`	`$FWDIR/modules/fwkern.conf`	`/bin/false`

The context field defines the type of configuration file:

- `global_context` - Security Gateway configuration file
- `all_vs_context` - Virtual Systems configuration file

The action field defines that action to be taken when the imported (copied) file is different that the local file:

- `/bin/true` - Reboot is required
- `/bin/false` - No reboot is required
- String enclosed in double quotes - Name of a "callback script" that selects the applicable action.

Example of a configuration file list:

```
global_context $PPKDIR/boot/modules/sim_aff.conf "sim affinityload"
global_context $PPKDIR/boot/modules/simkern.conf /bin/false
global_context $FWDIR/modules/fwkern.conf /bin/false
all_vs_context $FWDIR/conf/fwauthd.conf /bin/false
all_vs_context $FWDIR/conf/discntd.if /bin/false
global_context /var/opt/fw.boot/ha_boot.conf /bin/false
all_vs_context $FWDIR/conf/sync_exceptions_tab "g_sync_exception -f"
all_vs_context $FWDIR/bin/reserved_conns_tab "g_reserved_conns -f"
global_context /config/active  /usr/bin/confd_clone /config/db/cloned_db
global_context /tmp/sms_rate_limit.tmp /bin/true
global_context /tmp/sms_history.tmp /bin/true
global_context /home/admin/.ssh/known_hosts /bin/true
global_context /etc/passwd /bin/true
global_context /etc/shadow /bin/true


all_vs_context $FWDIR/bin/iproute.load /bin/true
all_vs_context $FWDIR/conf/gre_loader.conf /bin/true
global_context $FWDIR/conf/fwha_ch_uptime /bin/true
global_context $FWDIR/modules/mq_aff.conf "mq_affinity -s"
global_context $FWDIR/conf/pingable_hosts.conf "pingable_hosts local on"
all_vs_context $FWDIR/conf/pingable_hosts.ips /bin/true
global_context $FWDIR/conf/alert.conf /bin/true
all_vs_context $FWDIR/conf/asg_log_servers.conf "log_servers_util refresh"
global_context $FWDIR/modules/vlan_mq.conf "vlan_perf_enhancement -c"
global_context $FWDIR/conf/fw_global_params.conf "cpha_blade_config
fw_global_params_changed"
global_context $FWDIR/boot/mq.conf "cpmq reconfigure"
global_context /etc/modprobe.conf asg_update_modprobe_conf /tmp/modprobe.conf.new
global_context $FWDIR/boot/modules/vpnkern.conf /bin/false
global_context /etc/ssm_port_speed.conf /bin/asg_update_port_speed
/tmp/ssm_port_speed.conf.new
all_vs_context $FWDIR/conf/selective_template_exclude.conf /bin/true
global_context /etc/syslog_servers_list.conf asg_syslog_helper
global_context $FWDIR/conf/vsaffinity_exception.conf /bin/false
all_vs_context $FWDIR/conf/manual.affinity.conf "check_smo_affinity_files manual"
```

```
global_context $FWDIR/conf/fwkall.affinity.conf "check_smo_affinity_files fwdir"
$FWDIR/tmp/
all_vs_context $CPDIR/conf/*.affinity.conf "check_smo_affinity_files cpdir"
$CPDIR/tmp/
global_context $FWDIR/conf/resctrl "$FWDIR/bin/fw vsx resctrl load_configuration"
```

MAC Addresses and Bit Conventions

MAC addresses are divided into these types:

- **BMAC** - A MAC address assigned to all interfaces with the "BPEthX" naming convention. This is unique for each member. It does not rely on the interface index number.
- **VMAC** - A MAC address assigned to all interfaces with "ethX-YZ" naming convention. This is unique for each Chassis. It does not rely on the interface index number.
- **SMAC** - A MAC address assigned to Sync interfaces. This is unique for each member. It does not rely on the interface index number.

Bit Conventions

BMAC

Bit range	Description
1	Distinguishes between VMAC and other MAC address This is used to prevent possible collisions with VMAC space. Possible values: • `0` - BMAC or SMAC • `1` - VMAC
2-8	Member ID (starting from 1) This is limited to 127 members.
9-13	Always zero
14	Distinguishes between BMAC and SMAC address This is used to prevent possible collisions with SMAC space. Possible values: • `0` - BMAC • `1` - SMAC
15-16	Absolute interface number This is taken from the interface name. When the BPEthX format is used, X is the interface number. This is limited to four interfaces.

VMAC

Bit range	Description
1	Distinguishes between VMAC and other MAC address This is used to prevent possible collisions with VMAC space. Possible values: • `0` - BMAC or SMAC • `1` - VMAC
2-3	Chassis ID Limited to 4 Chassis
4-8	Switch number Limited to 32 switches
9-16	Port number Limited to 256 for each switch

SMAC

Bit range	Description
1	Distinguishes between VMAC and other MAC address This is used to prevent possible collisions with VMAC space. Possible values: • `0` - BMAC or SMAC • `1` - VMAC
2-8	Member ID (starting from 1) This is limited to 127 members.
9-13	Always zero
14	Distinguishes between BMAC and SMAC address This is used to prevent possible collisions with SMAC space. Possible values: • `0` - BMAC • `1` - SMAC
15	Always zero
16	Sync interface Possible values: • `0` - Sync1 • `1` - Sync2

MAC Address Resolver (asg_mac_resolver)

Use `asg_mac_resolver` to make sure that all types of MAC address (BMAC, VMAC, and SMAC) are correct. From the given MAC address, `asg_mac_resolver` determines the:

- MAC type
- Chassis ID
- SGM ID
- Assigned interface

Syntax

```
> asg_mac_resolver <mac_addr>
```

Example

```
> asg_mac_resolver 00:1C:7F:01:00:FE
```

Output

```
[00:1C:7F:01:00:FE, BMAC] [Chassis ID: 1] [SGM ID: 1] [Interface: BPEth0]
```

Notes

- The specified MAC Address comes from the BPEth0, on SGM 1 on Chassis 1.
- 00:1C:7F:**01:00:FE** is the Magic MAC attribute, which is identified by **FE**.
- The index is 16 bits (2 Bytes) identified by **01:00** 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16.

Security Group (asg security_group)

To be part of the Security Gateway, an SGM must belong to the Security Group. SGMs are added to the Security group using the `asg security_group` command. SGMs in the security group:

- Are selected during the initial installation procedure (after running: `#setup`)
- Are automatically installed once installation of the first SGM has completed
- Can be changed by using the `asg security_group` command

Syntax

```
> asg security_group
```

Example

```
> asg security_group
```

Output

```
+-------------------------------------+
|       Security Group Utility        |
+------------------------------------ +

Current Security Group:

+-------------------------------------+
| Chassis |  Security Gateway Modules  |
|-------------------------------------|
|    1    | 1,2,3                      |
 |---------------------------------------|
 |    2    | 1,2,3                        |
 +---------------------------------------+

 Choose one of the following options:
------------------------------------
 1) Add SGMs to Security Group
 2) Remove SGMs from Security Group
 3) Exit
```

Notes

Select which SGMs must be added or removed from the security group. Note that:

- An SGM added to the security group automatically joins the Single Management Object of the Security Gateway and then reboots
- Before you remove an SGM from the security gateway, make sure that its state is DOWN.
- To optimize connection distribution amongst the SGMs, keep the security group updated with the actual number of SGMs in the appliance.

Important - Run `asg security_group verify` to make sure that the security group is correctly configured.

Working with the Distribution Mode

The Distribution Mode is the way that an SSM assigns incoming traffic to SGMs. These are the supported Distribution Modes:

Mode	Description	Applies to
User	Packets are assigned to an SGM based on the packet destination.	One SSM
Network	Packets are assigned to an SGM based on the packet source.	One SSM
General	Packets are assigned to an SGM based on both the packet source and destination.	All SSMs in the 61000/41000 Security System
Per-Port	Each SSM data interface is configured separately as **User mode** or **Network mode**.	SSM data interface

Note – The User and Network modes always work together and are known collectively as the User/Network mode.

By default, the 61000/41000 Security System automatically configures the Distribution Mode. You can manually assign the General mode as necessary. There can be some scenarios where you must manually assign the General mode. The system does not automatically assign the General mode, with these exceptions:

- For Security Gateway deployments, the General mode is automatically assigned if there is at least one Bridge Mode interface.
- For VSX environments, the General mode is automatically assigned if there is at least one Virtual System configured in the Bridge mode.

Automatic Distribution Configuration (Auto-Topology)

By default, the 61000/41000 Security System automatically configures the Distribution Mode. The best Distribution Mode is selected based on the Gateway topology as defined in SmartDashboard.

The Distribution Mode is automatically based on these interface types:

- Physical interfaces, except for management and synchronization interfaces
- VLAN
- Bond
- VLAN over Bond

These examples show how the distribution Mode can be automatically configured for each interface.

Physical Interfaces

Physical Interface	Topology	SSM	Distribution Mode
eth1-01	Internal	1	User
eth1-02	Internal		
eth2-01	External	2	Network
eth2-02	External		

In this example, all ports on each SSM are either Internal or External. The Distribution Mode for the two SSMs is automatically configured as **User** or **Network**.

Physical interfaces

Interface	Topology	SSM	Port	Distribution Mode
eth1-01	Internal	1	1	User
eth1-02	External	1	2	Network
eth2-01	External	2	1	Network
eth2-02	External	2	2	Network

On at least one of the SSMs, some ports are Internal and others are External. The Distribution Mode for the SSMs is automatically configured as **Per Port**.

Physical and VLAN interfaces

Interface	Topology	SSM	Port	VLAN	Distribution Mode
eth1-01	External	1	1	NA	Network
eth1-01.100	Internal	1	1	100	User
eth1-01.200	External	1	1	200	Network

Interface	Topology	SSM	Port	VLAN	Distribution Mode
eth1-01.300	Internal	1	1	300	User

Three VLANs are defined on one SSM port. On at least one of the SSMs, some VLANs are Internal and others are External. Therefore, the SSM Distribution Mode is automatically configured as Per-Port.

Note - Automatic physical and VLAN interface configuration is not supported for the SSM60. For an SSM60, the Distribution Mode of all the VLANs on each port must be the same as the Distribution Mode of the port.

VSX Virtual Systems

Interface	Topology	Distribution Mode
eth1-01	External	N/A
wrpj64	Internal	Network
wrpj128	Internal	Network
wrpj192	Internal	User

Because a Virtual Switch does not have topology, the Distribution Mode is calculated based on the topologies of the WARP interfaces connected to the Virtual Systems, as shown. In this example, the Distribution Mode is calculated as **Network**.

Bond interfaces

Interface	Topology	Slaves	SSM	Port	Distribution Mode
bond1	Internal	eth1-01	1	1	User
		eth2-01	2	1	User
bond2	External	eth1-02	1	2	Network
		eth2-02	2	2	Network

In this example, both interfaces on each Bond are configured with the same distribution mode. Both bond interfaces are configured with one port for SSM1 and one port for SSM2. On both SSMs, one port is Internal and the other is External. The SSM Distribution Mode is automatically configured as **Per-Port**.

VLAN over Bond Interfaces

Interface	Topology	Slaves	SSM	Port	VLAN	Distribution Mode
bond1.100	Internal	eth1-01	1	1	100	User
		eth2-01	2	1	100	User
bond1.200	External	eth1-01	1	1	200	Network
		eth2-01	2	1	200	Network

The automatic distribute mode configuration is based on the VLAN topology. In this example, both interfaces on each VLAN are configured with the same distribution mode. Both Bond interfaces are configured on port 1 for each SSM. The SSM Distribution Mode is automatically configured as **Per-Port**.

Note - Automatic VLAN over Bond Interfaces configuration is not supported for the SSM60. For an SSM60 the Distribution Mode of all the VLANs must be the same.

SSM60 VLAN Legacy Support

The SSM60 does not support the new VLAN scheme used by the SSM160. SSM60 users should continue to use the legacy VLAN scheme.

To activate the legacy VLAN scheme on SSM60:

1. Run these commands in Expert model:
   ```
   # dbset chassis:id:1:SSM1:legacy_vlan on
   # dbset chassis:id:1:SSM2:legacy_vlan on
   # dbset chassis:id:2:SSM1:legacy_vlan on
   # dbset chassis:id:2:SSM2:legacy_vlan on
   ```
2. Reboot the SGMs.

   You can reboot the SGMs one Chassis at a time to keep connectivity during this procedure.

Manual Distribution Configuration (Manual-General)

In some deployments, you must manually configure a Distribution Mode of general. In other cases, you may want to force the system to work in General Mode.

When the Distribution Mode is manually configured (Manual-General Mode), the Distribution Mode of each SSM is General. In this configuration, the topology of the interfaces is irrelevant.

Note - We do not recommend that you manually change the Distribution mode of a Virtual System. This can cause performance degradation.

Setting and Showing the Distribution Configuration

Use these gclish commands to set and show the distribution configuration.

Syntax

```
> set distribution configuration auto-topology|manual-general
> show distribution configuration
```

Note - If the system is a VSX system, configure the command below on VS-0 only. It applies immediately across all VS's.

To change the distribution mode to manual-general:

```
> set distribution configuration manual-general
1_01:
configuration update completed successfully
```

```
1_02:
configuration update completed successfully


1_03:
configuration update completed successfully
```

To show the distribution:

```
> show distribution configuration
1_01:
manual-general


1_02:
manual-general


1_03:
manual-general
```

Configuring the Interface Distribution Mode (set distribution interface)

Use these commands to:

- Set the Distribution Mode for an interface when the system is not working in the **General Mode**.
- Show the Distribution Mode and whether it is assigned by Auto-Topology, or is manually configured.

Note - When working with Virtual Systems, you must move to the applicable virtual system context before you can change the Distribution Mode.

To move to the Virtual System:

```
> set virtual-system <vs_ids>
```

Parameter	Description
<*if_name*>	Interface name as assigned by the operation system
<*vs_ids*>	Virtual System context

Syntax

```
> set distribution interface <if_name> configuration user|network|policy
> show distribution interface <if_name> configuration
```

Parameter	Description
<*if_name*>	Interface name as assigned by the operation system
user	Manually assign the user Distribution Mode

Parameter	Description
network	Manually assign the network Distribution Mode
policy	Use Auto-Topology to automatically assign the Distribution Mode according to the policy

Example:

This example shows how to:

- Manually change the Distribution Mode for interface `eth1-01` from `policy` to `network`.
- Change the Distribution Mode on interface `eth1-01` from `network` to `policy:`

```
> set distribution interface eth1-01 configuration network
1_01:
configuration update completed successfully
1_02:
configuration update completed successfully
1_03:
configuration update completed successfully
> set distribution interface eth1-01 configuration policy
1_01:
configuration update completed successfully
1_02:
configuration update completed successfully
1_03:
configuration update completed successfully
```

Showing Distribution Status

Use this command to show a summary or detailed status report of the Distribution mode.

Syntax

```
> show distribution status [verbose]
```

Parameter	Description
verbose	Shows a detailed report for all SGMs and SSMs

Example

```
> show distribution status verbose
```

Output

```
Topic:                                    Configuration:
distribution mode                         user-network
policy mode                               on
ssm 1 mode                                user
ssm 2 mode                                network
ipv6 mode                                 off
```

```
spi mode                                  off
40g mode                                  off
matrix size                               2048
interface eth1-01 mode                    policy-internal
interface eth2-01 mode                    policy-external
```

Field	Description
distribution mode	Currently configured Distribution mode
policy mode	Auto-Topology assignment • On - Auto-topology • On - Manual override • Off - Manual-General
ssm mode	Distribution Mode assignment for each SSM
ipv6 mode	Shows if IPv6 is enabled for this system (on/off)
spi mode	Shows if SPI affinity is enabled for this system (on/off)
40g mode	Shows if QSFP ports are working at 40GbE (On) or at 4 x 10GbE (Off)
matrix size	Size of the Distribution matrix The Distribution matrix is a table that contains SGM IDs that are used for traffic assignment.
interface	Shows the Distribution mode assignment for each interface

Running a Verification Test (show distribution verification)

Use this command to run a verification test of the Distribution Mode configuration. This test compares the SGM and SSM configuration with the actual results. You can see a summary or a detailed (verbose) report of the test results.

Syntax

```
> show distribution verification [verbose]
```

Parameter	Description
`verbose`	Shows a detailed report for all SGMs and SSMs

Example

Note – This example shows only a small sample of the data. The checksums are truncated to fit on the page.

```
> show distribution verification verbose
Test:                                               Configuration Verification Result
chassis 1 blade 1 dxl-general-mode                  off           off          Passed
chassis 1 blade 1 dxl-md5sum                        5be67561a...  5be675611... Passed
chassis 1 blade 1 dxl-size                          2048          2048         Passed
chassis 1 blade 2 dxl-general-mode                  off           off          Passed
chassis 1 blade 2 dxl-md5sum                    5be67561a...  5be675611... Passed
chassis 1 blade 2 dxl-size                      2048          2048         Passed
chassis 1 blade 3 dxl-general-mode              off           off          Passed
chassis 1 blade 3 dxl-md5sum                    5be67561a...  5be675611... Passed
chassis 1 blade 3 dxl-size                      2048          2048         Passed
chassis 1 ssm 1 ipv6-mode                       off           off          Passed
chassis 1 ssm 1 mask ipv4 general destination   0000001f      0000001f     Passed
chassis 1 ssm 1 mask ipv4 general source        0000001f      0000001f     Passed
chassis 1 ssm 1 mask ipv4 user-network destination  000007ff  000007ff     Passed
chassis 1 ssm 1 mask ipv4 user-network source   000007ff      000007ff     Passed

Summary:
verification passed successfully
```

NAT and the Correction Layer on a Security Gateway

For optimal system performance, one SGM handles all traffic for a session. With NAT, packets sent from the client to the server can be distributed to a different SGM than packets from the same session sent from the server to the client. The system Correction Layer then must forward the packet to the correct SGM.

Correctly configuring Distribution Modes keeps corrections situations to a minimum and optimizes system performance. To achieve optimal distribution between SGMs on the gateway:

- **When not using NAT rules:** Set the General Distribution Mode.
- **When using NAT rules:** Set the hidden network(s) to User Mode, and the destination network(s) to Network Mode.

NAT and the Correction Layer on a VSX Gateway

In a VSX Gateway, the guidelines in NAT and the Correction Layer on Security Gateway apply to each Virtual System individually. In particular, an entire session should be handled by the same SGM by a given Virtual System. When a Virtual Switch ("Junction") connects several Virtual Systems, the same session may be handled by one Virtual System on one SGM, and by another Virtual System on a different SGM.

When a packet reaches a Virtual System from a Junction, the system VSX Stateless Correction Layer rechecks the distribution according to the Warp interface's Distribution Mode. It can decide to forward the packet to a different SGM.

In addition, on each Virtual System the system Correction Layer, which is stateful, can forward session's packets, similar to Security Gateway.

All forwarding operations have a performance impact. Therefore, the Distribution Mode configuration should minimize forwarding operations.

To achieve optimal distribution between SGMs on the VSX Gateway:

1. If you do not use NAT rules on any Virtual System, set the General Distribution Mode.
2. If you use NAT rules on at least one Virtual System, set the hidden network(s) to User Mode, and the destination network(s) to Network Mode.
3. On the remaining Virtual Systems that do not use NAT rules, set internal network(s) to User Mode, and the external network(s) to Network Mode.

Hybrid System

A 61000/41000 Security System *Hybrid System* is a deployment with SGMs that have different quantities of physical CPU cores. In a Hybrid System, the total number of CoreXL and Performance Pack instances that can run on one SGM is equal to the number of physical CPU cores. All SGMs **must** have the same number of CoreXL instances. The number of Performance Pack instances can be different.

Note - While it is possible to mix SGM220 and SGM260 units in the same environment, we do not recommend this configuration.

For example, a Hybrid System can contain these SGMs:

SGM	Physical CPU Cores	CoreXL Instances	Performance Pack Instances
1_01	12	10	2
1_03	20	10	10
1_04	40	10	20

How this works:

When an SGM boots, the 61000/41000 Security System makes sure that the number of CoreXL instances on the SGM matches the number defined for all other SGMs. Typically, this information comes from the SMO.

If the SGM has too many CoreXL instances, the system automatically reassigns these instances as Performance Pack instances. If the SGM has insufficient CPU cores, the SGM stays in the **DOWN** state. You must manually change the number of CoreXL instances and then reboot the SGM.

To see the number of CoreXL instances defined for ALL SGMs:

Run:

```
> asg_cores_stats
```

To manually change the number of CoreXL instances for ALL SGMs:

Run:

```
> cpconfig corexl instances <num_instances>
```

Parameter	Description
<*num_instances*>	Number of CoreXL instances for all SGMs

Important Notes:

- There is always at least one CPU core configured as a CoreXL instance and one as a Performance Pack instance.
- The maximum number of Performance Pack instances on an SGM is the lesser of **Physical cores -1** or **16**.
- The maximum number of CoreXL instances on an SGM is **Physical cores -1**.
- If manual Performance Pack core configuration for one SGM causes an invalid configuration on a different SGM, it automatically goes back to the default Performance Pack configuration.

- It is possible to have overlapping CoreXL and Performance Pack instances, where the number of instances is greater than the number of physical cores. We do not recommend this configuration.

Working with the GARP Chunk Mechanism

When Proxy ARP is enabled, the Firewall responds to ARP requests for hosts other than itself. When Chassis failover occurs, the new Active Chassis sends GARPs with its own (new) MAC address to update the network ARP tables.

To prevent network congestion during Chassis failover, GARP requests/responses are sent in user defined groups called "chunks". Each chunk contains a predefined number of GARP messages based on these parameters:

- The number of GARP messages in each chunk
- **HTU** (High Availability Time Unit) - Time interval, after which a chunk is sent.
- The chunk mechanism is iterating on the proxy ARP IPs, and each time sends GARPs only for some of them until it completes the full list.

In each HA Time Unit (HTU=0.1s) - a chunk of the GARP list is sent.

When the iteration sends the full list, it waits N HTU and sends the list again.

Configuration:

In each HTU (=0.1 second) - a chunk of the GARP list is sent.

For example, to send 10 GARPs each second, set `fwha_refresh_arps_chunk` to: `1`

```
# fw ctl set int fwha_refresh_arps_chunk 1
```

To send 50 GARPs per second, set `fwha_refresh_arps_chunk` to: `5`

```
# fw ctl set int fwha_refresh_arps_chunk 5
```

Whenever the iteration is finished sending GARPs for the entire list, it waits N HTU and re-sends the GARPS again. The time between the iterations can be configured with:

```
fwha_periodic_send_garps_interval1 = (1 HTU) /* should not be changed, send
immediately after failover */
fwha_periodic_send_garps_interval2 = (10 HTU) /* 01 seconds */
fwha_periodic_send_garps_interval3 = (20 HTU) /* 02 seconds */
fwha_periodic_send_garps_interval4 = (50 HTU) /* 05 seconds */
fwha_periodic_send_garps_interval5 = (100 HTU) /* 10 seconds */
```

In the above (default) configuration, after the iteration sends the list:

- Wait 1 second and start send again.
- Wait 2 seconds and start send again.
- Wait 5 seconds and start send again.
- Wait 10 seconds and start send again.

To change the interval:

Run:

```
# fw ctl set int fwha_periodic_send_garps_interval<1-5> 1
```

To apply the intervals:

Run:

```
# fw ctl set int fwha_periodic_send_garps_apply_intervals 1
```

Verification:

To manually send garp messages:

On the Chassis monitor blade, run:

```
> fw ctl set int test_arp_refresh 1
```

This causes garp messages to be sent (same as was failover).

Debug:

```
> fw ctl zdebug -m cluster + ch_conf | grep fw_refresh_arp_proxy_on_failover
```

Port Forwarding on Management Servers

Initiating traffic from an SGM which is not the SMO through the management interface (for example eth1-mgmt4) only works with specific services:

- On UDP: RADIUS, TACACS, SYSLOG, DNS, NTP
- On TCP: CRL, URLF proxy, URLF no proxy, LDAP, TACACS, CPD, SMTP, SSH

To add new services to the list:

1. Edit `$FWDIR/conf/fw_global_params.conf.`
2. example for number of ports type: mgmt_forwarding_tcp_ports_list_string 55201,55200,55202
3. Run `g_cp2blades $FWDIR/conf/fw_global_params.conf.`
4. Run `cpha_blade_config fw_global_params_changed` to apply the string.

Advanced Hardware Configuration

In This Section:

Chassis Management Module (CMM) CLI

The Chassis Management Module (CMM) monitors and controls all hardware components in the Chassis. The CMM communicates with a dedicated SGM using SNMP. If a hardware sensor reports a problem, the CMM automatically takes action or sends a report. CMMs also have a Command Line Interface.

To connect to the active CMM:

1. Connect to the serial port on the front panel of the CMM.
2. In your terminal emulation program, set the baud rate to 9600.
3. Enter `admin` for the user name and password.
4. Open a telnet or SSH session from one of the SGMs.
5. Ping these addresses:
   - 198.51.100.33
   - 198.51.100.233
6. Telnet or SSH from the SGM to the active CMM.
7. Enter `admin` for the user name and password.

To connect to the standby CMM:

1. Connect to the active CMM.
2. At the command prompt, run: `ifconfig`
3. Record the IP Address for the USB interface.

`Telnet or SSH from the active CMM to the standby CMM with the IP from the table below.`

Active CMM IP	Standby CMM IP
192.168.1.131	192.168.1.130
192.168.1.131	192.168.1.130
192.168.1.2	192.168.1.3
192.168.1.3	192.168.1.2

Logging CMM Diagnostic Information

How to log CMM diagnostic information:

1. Log into the active CMM.
2. Run:
   ```
   > /etc/summary
   ```
   This command can take several minutes to run.
3. Run:
   ```
   > cat /tmp/debug.log
   > cat /etc/shmm.cfg
   > clia fruinfo 20 0
   > clia fruinfo 20 1
   > clia fruinfo 20 2
   > clia fruinfo 20 3
   > clia fruinfo 20 4
   > clia fruinfo 20 5
   > clia fruinfo 20 6
   > clia fruinfo 20 7
   > clia fruinfo 20 8
   > clia fruinfo 20 9
   ```
4. On the NG 61000 Security System, run:
   ```
   > clia fruinfo 20 10
   > clia fruinfo 20 11
   > clia fruinfo 20 12
   > clia fruinfo 20 13
   > clia fruinfo 20 14
   > clia fruinfo 20 15
   > clia fruinfo 20 16
   ```
5. On the 61000/41000 Security System, run:
   ```
   > clia fruinfo y 10
   > clia fruinfo y 12
   > clia fruinfo y 82
   > clia fruinfo y 84
   > clia fruinfo y 86
   > clia fruinfo y 88
   > clia fruinfo y 8a
   > clia fruinfo y 8c
   > clia fruinfo y 8e
   > clia fruinfo y 90
   > clia fruinfo y 92
   > clia fruinfo y 94
   > clia fruinfo y 96
   > clia fruinfo y 98
   > clia fruinfo y 9a
   > clia fruinfo y 9c
   ```

The logs are stored in `/tmp/debug.log` on the CMM.

Changing the CMM Administrator Password

To change the CMM Administrator Password:

1. In Expert mode, run:
   ```
   # passwd admin
   ```
2. Enter and confirm the new password.

Changing the Chassis Configuration

To change the Chassis configuration:

Edit: `/etc/shmm.cfg`

Chassis Management Module (CMM) CLI Commands

Use the CMM CLI commands to monitor and manage the CMM.

Some commands use SGM/SSM IDs or Slot IDs. Use these tables to find the correct SGM ID or Slot ID.

NG 61000 Security System slot information

Physical slot	Slot ID	SGM/SSM
1	9a	SGM1
2	96	SGM2
3	92	SGM3
4	8e	SGM4
5	8a	SGM5
6	86	SGM6
7	82	SSM1
8	84	SSM2
9	88	SGM7
10	8c	SGM8
11	90	SGM9
12	94	SGM10
13	98	SGM11
14	9c	SGM12

41000 Security System slot information

Physical slot	Slot ID	SGM/SSM
Upper most slot	8C	SGM1
	8A	SGM2
	88	SGM3
	86	SGM4
	84	SSM2
Lowest Slot	82	SSM1

clia alarm

Use this command to:

- Shows the current alarms on the CMM
- Reset the alarms

Syntax

```
> clia alarm [0]
```

Parameter	Description
`0`	Reset the alarms

clia board

Use this command to make sure the boards are recognized.

Syntax

```
> clia board
```

clia boardreset

Use this command to reset a board.

Syntax

```
> clia boardreset <slot_num>
```

Parameter	Description
<*slot_num*>	Slot number of the board to reset

clia fru

Use this command to see information for an SGM or SSM.

Syntax

```
> clia fru <sgm_id>|<ssm_id>
```

Parameter	Description
<sgm_id>|<ssm_id>	ID of an SGM or SSM

clia help

Use this command to see a list of available commands.

Syntax

```
> clia help
```

clia reboot

Use this command to reboot the CMM.

The Chassis fails over to the standby CMM.

Syntax

```
> clia reboot
```

clia shelf pd

Use this command to see power consumption information for all boards.

Syntax

```
> clia shelf pd
```

clia sel

Use this command to retrieves event logs.

Syntax

```
> clia sel
```

i2c_test

Use this command to:

- Test the I2C connection
- See all devices connected to the CMM using I2C

Syntax

```
> i2c_test
```

Security Switch Module (SSM) CLI

The Security Switch Module (SSM):

- Distributes network traffic to the Security Gateway Modules (SGMs)
- Forwards traffic from the SGMs to the network
- Shares the load amongst the SGMs

The SSMs and SGMs communicate automatically through SNMP requests. You can also connect directly to the SSM and run commands.

You can connect to the SSM CLI through:

- A serial port on the front panel of the SSM.
- A telnet session from one of the SGMs

SSM60 CLI

1. Connect to a serial port on the front panel of the SSM.

   The SSM60 has two serial ports, one for the fabric switch (data ports) and one for the base switch (management ports).

   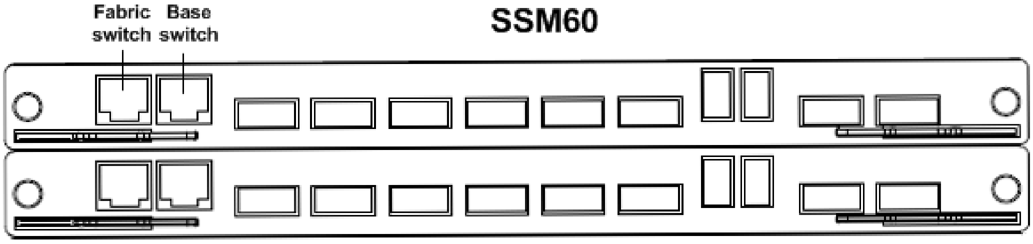

2. In your terminal emulation program, set the baud rate to 9600.
3. Enter `admin` for the password.
4. Give read-write permissions to the system:

   ```
   > enable
   ```

5. Enter ? for a list of available commands and usage.

Note - Load balancing commands are run on the fabric switch only.

6. Open a telnet session from one of the SGMs.
7. Ping these address to make sure you have connectivity to the SSMs:

SSM	Switch	IP Address
1	Base	198.51.100.31
	Fabric	198.51.100.32
2	Base	198.51.100.231
	Fabric	198.51.100.232

8. Telnet from the SGM to the SSM
9. Enter `admin` for the password.
10. Give read-write permissions to the system:

    ```
    > enable
    ```

11. Enter ? for a list of available commands and usage.

When connected, you can use these troubleshooting commands:

To	Run:
View the current configuration	`# show running-config`

To	Run:
View current ports status	`# show interface`
View interface statistics	`# show interface <`*interface_id*`> statistics [extended]`
View SSM logs	`# show log buffer`

Modify the group of SGMs amongst which the load is distributed	`# configure terminal`
	`(config)# load-balance mtx-bucket {<`*sgm_id1*`>,<`*sgm_id2*`>...}`
	`(config)# load-balance apply`
	Note: The command does not work if you have an odd number of SGMs in the group. For example, do not run:
	`#load-balance mtx-bucket 1,2,3`
	Run:
	`#load-balance mtx-bucket 1,2,3,1,2,3`

SSM160 CLI

The SSM (Security Switch Module) is the networking module of the gateway. The SSM transmits traffic to and from the SGM and performs the load distribution among the SGMs.

The SSM includes two modules:

- Fabric switch - includes the Data ports
- Base switch - includes the Management ports

Usually the SSM communicates with the SGM through SNMP. Sometimes you can need to connect directly to the SSM.

Configuration

You can connect to the SSM CLI:

- With a serial console to the CLI port on the SSM front panel (baud rate 9600).
- From one of the SGMs with SSH.
  You can get the SSM IPs in clish/gclish:
  - `show Chassis id 1|2|all module SSM{1|2} ip`
  - The password for the SSM is `admin`.

To see the current configuration:

Run:

```
# show running-config [<feature_name>]
```

Because the full configuration is very long, we recommended that you specify the feature that you are interested in. For example, run `show running-config load-balance` to see the Load Balance configuration. Press **tab** to see a full list of the features.

To see the current port status:

Run:

```
# show port
```

To see detailed port information (speed, administrative state, link state, etc.):

Run:

```
# show port <port_id>
```

To see interface statistics:

```
# show port <port_id> statistics
```

Example

```
# show port 1/3/1 statistics
```

Output

```
===============================================================================
 Port Statistics
===============================================================================
                                                  Input                 Output
-------------------------------------------------------------------------------
Unicast Packets                                    5003                   7106
Multicast Packets                                568409                   1880
Broadcast Packets                                122151                   1972
Flow Control                                          0                      0
Discards                                             16                      0
Errors                                                0                      0
-------------------------------------------------------------------------------
-
Total                                            695563                  10958
===============================================================================

===============================================================================
 Ethernet Statistics in Packets
===============================================================================
RX CRC Errors                     0         TX Collisions                    0
RX Undersize                      0
-------------------------------------------------------------------------------
                                                  Input                 Output
-------------------------------------------------------------------------------
Fragments                                             0                      0
Oversize                                              0                      0
Jabbers                                               0                      0
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Packets                                                       Input and Output
-------------------------------------------------------------------------------
Octets                                                                71085491
Packets                                                                 706521
Packets of 64 Octets                                                      2290
Packets of 65 to 127 Octets                                             689951
Packets of 128 to 255 Octets                                              4122
Packets of 256 to 511 Octets                                              6009
Packets of 512 to 1023 Octets                                              258
Packets of 1024 to 1518 Octets                                             994
Packets of 1519 or more Octets                                               0
-------------------------------------------------------------------------------
-
Total                                            695563                  10958
```

```
===============================================================================

===============================================================================
 Rates in Bytes per Second
===============================================================================
                                                  Input                 Output
Rate for last 10 sec                               1477                     25
Rate for last 60 sec                               1435                     50
===============================================================================
```

Pay special intention to "Discards" and "Errors" fields which might show a problem if they constantly increase.

To view the SSM logs:

Run:

```
# unhide private
```

The default password is: `private`

```
# show private shell
# tail /var/log/messages
```

To change the load distribution on SGM groups:

Run:

```
# configure terminal
(config)# load-balance mtx-bucket 1 buckets
[<SGM_ID1><SGM_ID2>:<SGM_ID3><SGM_ID4>…]
(config)# commit
(config)# exit
#load-balance apply
```

Note - You need to provide a full list of the SGMs when you use this command. Otherwise, traffic might be dropped on the SSM.

To set port modes for 40G ports (4X10G or 1X40G):

1. Run:
   ```
   # unhide private
   ```
   The default password is: `private`
2. Run:
   ```
   # show private shell
   ```
3. To set 1X40G mode, run:
   ```
   # /batm/binux/bin/ub_util -s ahub4_40G yes
   ```
4. To set 4X10G mode, run:
   ```
   # /batm/binux/bin/ub_util -s ahub4_40G
   # exit
   # config terminal
   (config)# system reload
   ```

Note - This procedure requires you to reload the SSM. It is recommended that you do one SSM at a time.

To see the current version information:

Run:

```
# show version
```

To log out from current session:

Run:

```
# logout
```

To change the SSM160 admin password:

1. Log in using SSH or a serial console to an SGM on the Chassis.
2. In Expert mode, log in to one of the SSMs in the Chassis:
   `ssh admin@ssm<ssm_id>`
3. Enter admin password when prompted.
4. Run these commands:
   ```
   # conf t
   # system security user admin
   # password
   ```
5. Enter the new password.
6. Run these commands:
   ```
   # commit
   # end
   # logout
   ```

Notes

- This procedure should be done separately on each SSM in the system.
- This procedure does not cause any traffic interruption.

Example

```
# ssh ssm2
admin@ssm2's password:
BATM T-HUB4
admin connected from 198.51.100.215 using ssh on T-HUB4
T-HUB4#conf t
Entering configuration mode terminal
T-HUB4(config)#system security user admin
T-HUB4(config-user-admin)#password
(<MD5 digest string>): *****
T-HUB4(config-user-admin)#commit
Commit complete.
T-HUB4(config-user-admin)#end
T-HUB4#log
Connection to ssm2 closed.
```

Each port ID on the SGM maps to a port on the SSM. The table below maps SSM port IDs to SGM port IDs.

Note - This table relates to SSM1. For SSM2 replace eth1-X with eth2-X.

SGM	SSM
eth1-01	1/3/1
eth1-02	1/3/2
eth1-03	1/3/3

SGM	SSM
eth1-04	1/3/4
eth1-05	1/3/5
eth1-06	1/3/6
eth1-07	1/3/7
eth1-Sync	1/3/8
eth1-09	1/1/1
eth1-10	1/1/2
eth1-11	1/1/3
eth1-12	1/1/4
eth1-13	1/1/5
eth1-14	1/1/6
eth1-15	1/1/7
eth1-16	1/1/8
eth1-Mgmt1	1/5/1
eth1-Mgmt2	1/5/2
eth1-Mgmt3	1/5/3
eth1-Mgmt4	1/5/4

Verification

To make sure that you have connectivity to the SSMs from the SGMs, ping all the SSM modules IPs.

You can also make sure that SNMP connectivity is available:

```
> asg_chassis_ctrl get_ssm_firmware all
```

Adding/Removing SSMs After Initial Setup

If you add or remove SSMs after the initial installation, the system can show an incorrect number of installed SSMs or show some SSMs in the DOWN state. Use `asg_ssm_amount` to define the correct number of SSMs in the Chassis.

Procedure

	Operation	Command	Comment
1	Physically pull out all the SGMs, except the SMO		
2	Physically install the additional SSMs		
3	Connect to the individual SMO with a console cable		
4	Update the number of SSMs	`# asg_ssm_amount`	
5	Reboot the individual SMO	`# reboot`	

6	When the SGM is UP, make sure it matches the number of SSMs	Example:
		`# ccutil active_ssm`
		`SSM1 ACTIVE`
		`SSM2 ACTIVE`
		`SSM3 ACTIVE`
		`SSM4 ACTIVE`
		`# asg stat -v`
		`# ifconfig`
			See the `SSMs Unit` output in the `asg stat -v` command
			In the `ifconfig` output, make sure the system has `eth3-XX, eth4-XX` ports

7	Add the remaining disconnected SGMs		

Syntax

```
asg_ssm_amount <ssm_quantity>
```

- For the NG 61000 Security System, *<ssm_quantity>* can be 2 or 4.
- For the 41000 Security System, *<ssm_quantity>* can be 1 or 2.

Notes:

- You must run this command if you add or remove SSMs in your Chassis.
- Run this command in Expert mode.
- Make sure that only SGM is turned on when you run this command.
- Reboot the system after you run this command.

Examples:

```
[expert@gw:0] # asg_ssm_amount 1
[expert@gw:0] # asg_ssm_amount 2
[expert@gw:0] # asg_ssm_amount 4
```

Security Gateway Modules

The Security Gateway Modules (SGMs) in the Chassis work together as a single, high performance Security Gateway or VSX Gateway. Adding a Security Gateway Module scales the performance of the system. A Security Gateway Module can be added and removed without losing connections. If an SGM is removed or fails, traffic is distributed to the other active SGMs.

These SGM versions are available:

- SGM220 (Not supported in a 4-SSM configuration or the 41000 Security System.)
- SGM220T (for NEBS only - Not supported for the 41000 Security System)
- SGM260 (Supports 4-SSM configuration)

The SGM260 has more powerful CPUs and uses a more advanced technology. It also has a different front panel layout and different LEDs.

Identifying SGMs in the Chassis (asg_detection)

Use this command to flash the LEDs of a SGM. This lets you identify a specified SGM.

Syntax

```
# asg_detection [ -b <sgm_ids> ] [ -t <time> | off ]
```

Parameter	Description
`-b` <*sgm_ids*>	Works with SGMs and/or Chassis as specified by <*sgm_ids*>. <*sgm_ids*> can be: • No <*sgm_ids*> specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`) Default: Local SGM
`-t` <*time*>	Time in seconds the LEDs flash Default: 60 seconds
`-t off`	Stops LED flashes if they continue after the time in `-t` <*time*>

SGM260 LEDs

Item		LED	Status	Description
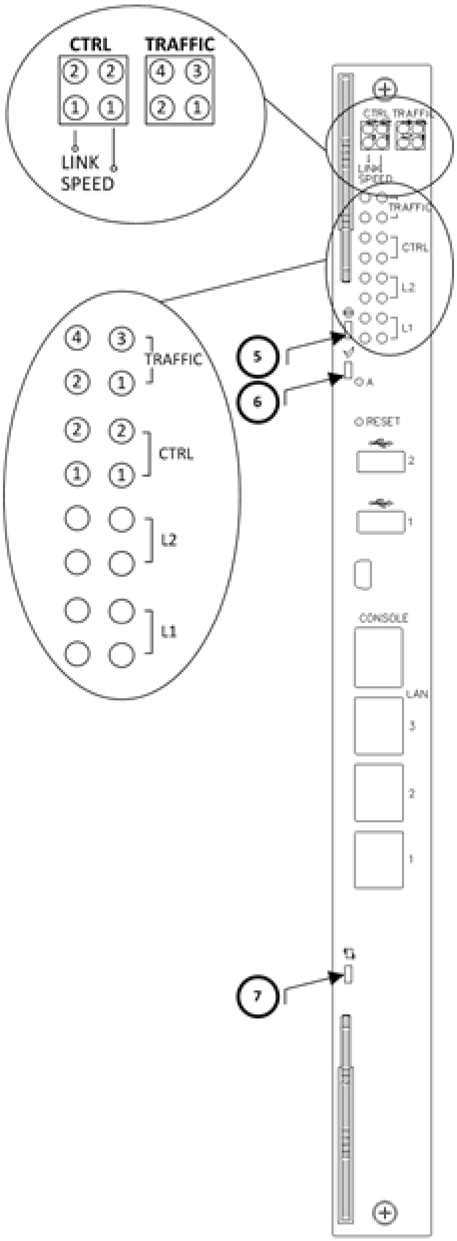	5	Out of service ⛔	Red	SGM out of service
			Off (Normal)	SGM hardware is normal
	6	Health ✓	Green (Normal)	SGM core operating system is active
			Green blinking	SGM core operating system is partially active
			Off	SGM operating system is in standby mode
	7	Hot-swap ⇅	Blue	SGM can be safely removed
			Blue blinking	SGM is going to standby mode. Do not remove
			Off (Normal)	SGM is active. Do not remove
	CLTR LINK 1 CTRL LINK 2	SSM1 and SSM2 management ports	Yellow	Link enabled
			Yellow blinking	Link is active
			Off	Link is disabled
	CTRL SPEED 1 CTRL SPEED 2	SSM1 and SSM2 management ports	Yellow	10 Gbps
			Green	1 Gbps
			Off	100 Mbps
	Traffic	1 2 3 4	On	Data and sync traffic in SSM1, SSM2, SSM3, SSM4
	L2		Off	Not used

Item		LED	Status	Description
	L1		Red. Lower Right ○ ○ ○ 🔴	Installation started
			Red blink, in sequence ◑ ◑ ◑ ◑	Installation in progress
			Red. All 🔴 🔴 🔴 🔴	Installation failure
			Yellow. Left 🟡 ○ 🟡 ○	Installation completed
			Green. Right ○ 🟢 ○ 🟢	SGM is being configured. (Using First Time Configuration Wizard or adding a new SGM into a Chassis)
			Off	SGM is configured and ready

SGM220 LEDs

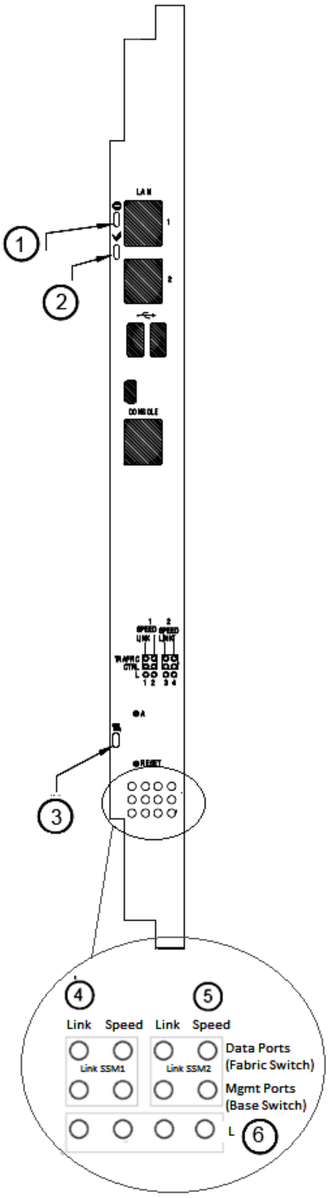

Item	LED	Status	Description
1	Out of service ⛔	Red	SGM out of service
		Off (Normal)	SGM hardware is normal
2	Health ✓	Green (Normal)	SGM core operating system is active
		Green blinking	SGM core operating system is partially active
		Off	SGM operating system is in standby mode
3	Hot-swap ⇅	Blue	SGM can be safely removed
		Blue blinking	SGM is going to standby mode. Do not remove
		Off (Normal)	SGM is active. Do not remove
4	Link	Yellow	Link enabled
		Yellow blinking	Link is active
		Off	Link is disabled
5	Data port speed	Yellow	10 Gbps
		Green	1 Gbps
		Off	100 Mbps
	Management port speed	Yellow	1 Gbps
		Green	100 Mbps
		Off	10 Mbps
6	L	LEDs 2 and 4 - Green	SGM is being configured. (Using First Time Wizard or adding a new SGM into a Chassis)
		All LEDs - Off	SGM is configured and ready

Security Switch Module LEDs

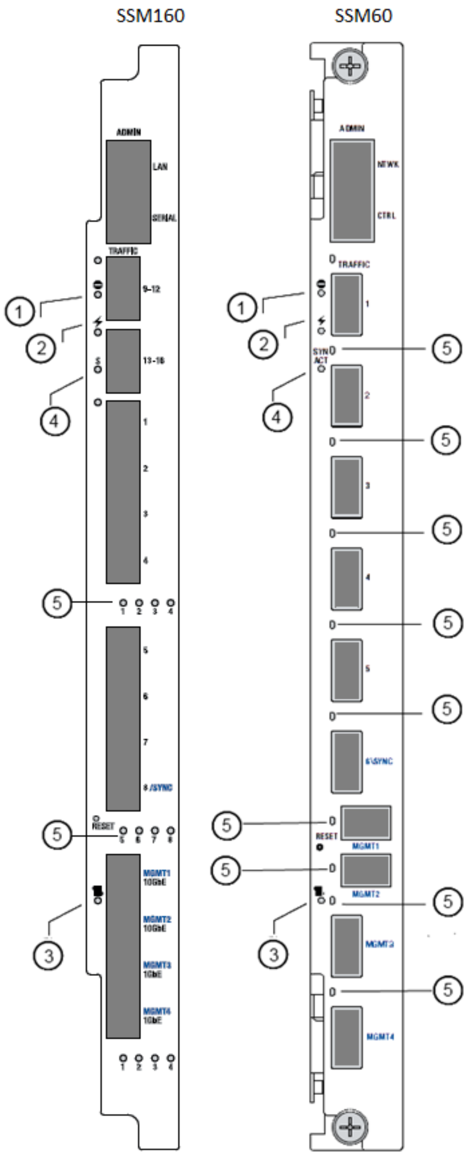

Item	LED	Status	Description
1	Out of service	Red	SSM out of service
		Off (Normal)	SSM hardware is normal
2	Power	On (Normal)	Power on
		Off	Power off
3	Hot-swap	Blue	SSM can be safely removed
		Blue blinking	SSM is going to standby mode. Do not remove
		Off (Normal)	SSM is active. Do not remove
4	SYN ACT	On (Normal)	Normal operation
		Off	N/A
5	Link	On	Link enabled
		Yellow blinking	Link is active
		Off	Link is disabled

Software Blades Support

In This Section:

Software Blades Update Verification

Use `asg_swb_update_verifier` to make sure that the signatures are up to date for these products:

- Anti-virus
- Anti-bot
- Application control
- URL filtering

Syntax

```
> asg_swb_update_verifier [-v] [-b <sgm_ids> [-m <product>] [-n [-p <ip>:<port>]]
] [-u <product>]
```

Parameter	Description
`-v`	Verbose - Shows detailed output
`-b` *<sgm_ids>*	Works with SGMs and/or Chassis as specified by <*sgm_ids*>. *<sgm_ids>* can be: • No <*sgm_ids*> specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`)

Parameter	Description
`-m <product>`	Force a manual update for SGMs specified with -b Valid values: • `all` - All products on the SGM • `Anti-Bot` • `Anti-Virus` • `APPI` • `URLF`
`-n`	Force an update download from the internet Use with `-m`.
`-p <ip>:<port>`	Force an update download from the internet and use a specific HTTP proxy. Use with `-m`. • *<ip>* - IP of the HTTP proxy • *<port>* - TCP port to use on the HTTP proxy
`-u <product>`	Force a database update for a specific product Valid values: • `all` - All products on the SGM • `Anti-Bot` • `Anti-Virus` • `APPI` • `URLF`

Example

```
> asg_swb_update_verifier
```

Output

```
+-------------------------------------------------------------------------------+
| product    | sgm  | status            | DB version | next update check        |
+-------------------------------------------------------------------------------+
| APPI       | 2_01 | failed            | 14061202   | Thu Jun 12 10:32:55 2014 |
| APPI       | 2_02 | failed            | 14061202   | Thu Jun 12 10:32:41 2014 |
| Anti-Bot   | 2_01 | up-to-date        | 1405220911 | Thu Jun 12 09:28:34 2014 |
| Anti-Bot   | 2_02 | up-to-date        | 1405220911 | Thu Jun 12 09:28:45 2014 |
| Anti-Virus | 2_01 | up-to-date        | 1406121233 | Thu Jun 12 09:28:12 2014 |
| Anti-Virus | 2_02 | new               | 1406121234 | Thu Jun 12 09:28:10 2014 |
| URLF       | 2_01 | not-installed     | N/A        | N/A                      |
| URLF       | 2_02 | not-installed     | N/A        | N/A                      |
+-------------------------------------------------------------------------------+

Report:
------------------------------ APPI ---------------------------------------------
DB versions verification                                               [   OK   ]
statuses verification                                                  [ FAILED ]

------------------------------ URLF ---------------------------------------------
DB versions verification                                               [   OK   ]
statuses verification                                                  [   OK   ]

------------------------------ Anti-Bot -----------------------------------------
DB versions verification                                               [   OK   ]
statuses verification                                                  [   OK   ]
```

```
------------------------------ Anti-Virus --------------------------------------
DB versions verification                                              [   OK   ]
statuses verification                                                 [   OK   ]
```

Field	Description
product	Name of the Product
sgm	SGM ID
status	Update status
DB version	Product database version
next update check	Date and time for the next automatic update
DB versions verification	• **OK** - The database version is correct • **FAILED** - The database version is incorrect
statuses verification	• **OK** - The update installed correctly or no update is needed • **FAILED** - The update did not install correctly

Threat Emulation

R76SP.40 supports the Threat Emulation installed on the 61000/41000 Security System platform. The Threat Emulation and Threat Prevention software blade is supported on a Security Management Server which has the latest Jumbo Hotfix installed.

To learn how to install Threat Emulation on the 61000/41000 Security System, see sk111405 http://supportcontent.checkpoint.com/solutions?id=sk111405. To learn how to work with Threat Emulation, see the *R77 versions Threat Emulation Administration Guide* https://sc1.checkpoint.com/documents/R77/CP_R77_ThreatPrevention_WebAdmin/html_frameset.htm.

IPS Bypass Under Load

Bypass Under Load allows the administrator to define a gateway resource load level at which IPS inspection suspends temporarily until the gateway's resources return to satisfactory levels.

IPS inspection can make a difference in connectivity and performance. Usually, the time it takes to inspect packets is not noticeable. However, under heavy loads it can be a critical issue.

You have the option to temporarily stop IPS inspection on a gateway if it experiences heavy load.

IPS Cluster Failover Management

You can configure how IPS is managed during a cluster failover. This occurs when one member of a cluster takes over for a different member, to supply High Availability.

To configure failover behavior for a cluster:

In Expert mode, run:

```
# asg_ips_failover_behavior connectivity|security
```

Parameter	Description
`connectivity`	Prefer connectivity - Close connections for which IPS inspection cannot be guaranteed
`security`	Prefer security - Keep connections alive even if IPS inspections cannot be guaranteed

Optimizing IPS (asg_ips_enhance)

Description

R76SP.40 supports HyperSpect optimization for IPS on systems that use the SGM260. HyperSpect uses adaptive traffic inspection to focus on the most important parts of each connection. This can give up to a 50% improvement for IPS inspection in real-life traffic scenarios.

Run the `asg_ips_enhance` command from the Expert mode to:

- Enable or disable HyperSpect
- Show HyperSpect status and enforce consistency across SGMs
- Synchronize the configuration

Syntax

```
asg_ips_enhance  [enable |disable] [status] [sync]
```

Parameter	Description
`enable`	Enable HyperSpect on all SGMs
`disable`	Disable HyperSpect on all SGMs
`status`	Show HyperSpect status and consistency for all Security Gateway Modules
`sync`	Synchronize the HyperSpect configuration file across all Security Gateway Modules

Examples

```
# asg_ips_enhance enable
```

Enables HyperSpect on all Security Gateway Modules

Replacing Hardware Components

In This Section:

Replacing the CMM

Install the replacement CMM that you received in the Return Merchandise Authorization (RMA). These steps are for CMM installation on a Standby Chassis in a Dual Chassis environment.

Before you begin:

1. Make sure you have a supported Chassis type.

   The supported Chassis types for the NG 61000 Security System are:

   - DC Chassis
   - AC Telkoor: The AC Chassis has two rows of three Telkoor power supplies in each row.
   - AC Lambda: The AC Chassis has one row of five Lambda power supplies

   The supported Chassis types for the 61000 NG Security Systems are:

   - DC Chassis
   - AC Lambda: The AC Chassis has one row of four Lambda power supplies

   The supported Chassis types for the 41000 Security System are:

   - AC Telkoor: Three Telkoor power supplies
   - DC Chassis
2. Get the label from the CMM box.

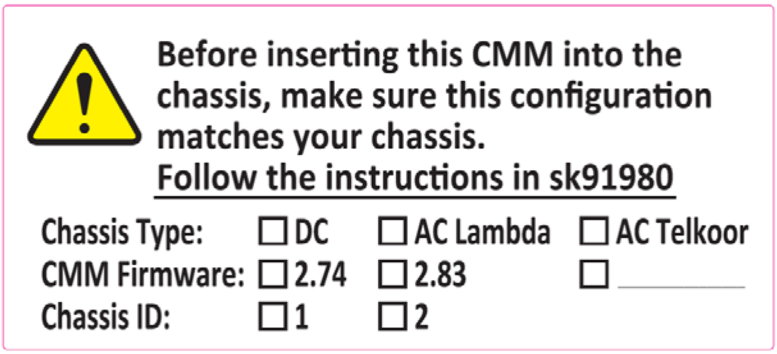

Before inserting this CMM into the chassis, make sure this configuration matches your chassis.
Follow the instructions in sk91980

Chassis Type:	□ DC	□ AC Lambda	□ AC Telkoor
CMM Firmware:	□ 2.74	□ 2.83	□ ________
Chassis ID:	□ 1	□ 2	

To replace the CMM:

1. Install the replacement CMM to the Standby Chassis.
2. Make sure that all CMMs in the environment have the same firmware version:

```
> asg_version -i
+-------------------------------------------------------------------+
| Hardware Versions                                                 |
+-------------------------------------------------------------------+
| Component       | Type       | Configuration      | Firmware       |
+-------------------------------------------------------------------+
| Chassis 1                                                         |
+-------------------------------------------------------------------+
| SSM1            | SSM160    | N/A                  | 2.4.C9         |
| SSM2            | SSM160    | N/A                  | 2.4.C9         |
| CMM(active)     | N/A       | N/A                  | 2.83           |
| CMM(standby)    | N/A       | N/A                  | 2.83           |
+-------------------------------------------------------------------+
+-------------------------------------------------------------------+
| Hardware Versions                                                 |
+-------------------------------------------------------------------+
| Component       | Type      | Configuration        | Firmware       |
+-------------------------------------------------------------------+
| Chassis 2                                                         |
+-------------------------------------------------------------------+
| SSM1            | SSM160    | N/A                  | 2.4.C9         |
| SSM2            | SSM160    | N/A                  | 2.4.C9         |
| CMM(active)     | N/A       | N/A                  | 2.83           |
| CMM(standby)    | N/A       | N/A                  | 2.83           |
+-------------------------------------------------------------------+
```

The output must be the same as the box label.

- If the firmware versions are not the same, upgrade the CMM Firmware.
- If the Chassis IDs are not the same, change the RMA CMM Chassis ID ("Setting the Chassis ID" on page 192).
- If the Chassis Types are not the same, run the next procedure.

To fix incorrect Chassis Type:

1. Put the Chassis in Standby state: `> asg chassis_admin -c <Chassis_id> down`
2. Remove all CMMs from the Chassis.
3. Insert the replacement CMM in the Chassis.
4. Open a console connection to the CMM:

   a) Connect one end of a serial cable to the serial port on the CMM front panel.

   b) Connect the other end of the serial cable to a computer.

   c) Open a console window. Use the default serial connection parameters: `9600`, 8, N, 1
5. Start the installation: `# install.sh`
6. For the NG 61000 Security System, select the applicable Chassis type.

   The menu can be different based on the CMM firmware. This menu shows for firmware 2.74.

   ```
   -----------------------------------------------------------------
   |               Select one of following options.               |
   |        1: Press 1 for 13U chassis (Telkoor PSU).             |
   |        2: Press 2 for 14U chassis (Telkoor PSU).             |
   |        3: Press 3 for 14U chassis (Lambda PSU).              |
   |        Q: Press Q for to skip.                              |
   -----------------------------------------------------------------
   ```

   - If the Chassis type is AC Telkoor PSU or a DC Chassis, enter: 2
   - If the Chassis type is AC Lambda, enter: 3

7. Insert the second CMM.
8. For the 41000 Security System: When the option to upgrade EEprom shows, select option `1`.

   ```
   -----------------------------------------------------
   | EEprom upgrading                                  |
   | 1: Press 1 for EEProm upgrading.                  |
   | 2. Press 2 to skip.                               |
   -----------------------------------------------------
   ```

   Note - In the 61000 Security System, there is no need to update EEprom.
9. Return the Chassis to the Standby state: `> asg chassis_admin -c <chassis_id> up`

Adding or Replacing an SGM

This section describes the procedure for doing an operating system upgrade on a new or replacement SGM.

There are two methods to update operating system versions:

- Create a snapshot image from one of the standby SGMs and revert the new SGM to this snapshot.
- Install from the distribution media. Please contact Check Point support for more information.

New or Replacement SGM Procedure Using Snapshot

Use this procedure to make sure that the current environment, including latest hotfixes, is installed on a new or replacement SGM. You can use this if an SGM is sent for service as an RMA.

This procedure has these basic steps:

1. Create image snapshot for the existing configuration and export it.
2. Import the snapshot to the new or replacement SGM.
3. Add the new or replacement SGM to the security group.
4. Make sure the new or replacement SGM works correctly.

To create and export a snapshot of the existing configuration:

Note - In a Dual Chassis configuration, we recommended that you create a snapshot on the Standby Chassis.

1. Switch to an SGM on the standby Chassis:

   `# blade <standby_Chassis_id>_<sgm_id>`
2. Set the global mode to off:

   `> set global-mode off`

   This makes sure that the new snapshot image is created only on this SGM
3. Create a new image snapshot:

   `> add snapshot <snapshot_name> desc <snapshot_desc>`
4. Monitor the creation process progress:

   `> show snapshots`
5. Insert a removable disk to the USB port of the SGM and mount it to: `/mnt/usb`

   To learn how to mount a USB drive, see Mounting and Dismounting a USB Disk (on page 318).

6. When creation process is done, export the snapshot to a tar file under `/mnt/usb`:

   ```
   > set snapshot export <snapshot_name_without_.tar> path /mnt/usb
   ```

7. Monitor the export process progress:

   ```
   > show snapshots
   ```

8. Un-mount `/usb/mnt`:

   ```
   # umount /mnt/usb
   ```

9. Remove the USB drive from the SGM.

Example

```
> set global-mode off
> add snapshot rma_62 desc rma
Taking snapshot. You can continue working normally.
You can use the command ‘show snapshots’ to monitor creation progress,
> show sna
snapshot - show snapshot data
snapshots — list of local snapshots
> show snapshots
Restore points:
---------------
armdilo62_2
Restore point now under creation:
riua_62 (19%)

Creation of an additional restore point will need 2.624G
Amount of space available for restore points is il.41G
test-chO2—03> show snapshots
Restore points:
----------------
rma_62
armdi 1062_2

Creation of an additional restore point will need 2.624G
Amount of space available for restore points is 41.53G
test-chO2—03> set snapshot export rma_62 path /mnt/usb/
Exporting snapshot. You can continue working normally.
You can use the command ‘show snapshots’ to monitor exporting progress.

# blade 2_3
Moving to blade 2_3
This system is for authorized use only.
Last login: Wed Jun 20 08:43:28 2012 from test—chO2—03
CLINFRO771 This gclish instance cannot run “set” operations. To allow running “
set” operations, run “set config—lock omm Override”
> shell
# cd /mnt/usb
# ls
rzna_62.tar

> exit
Connection to 192.0.2.17 closed.
# umount /uint/usb
```

To import the snapshot to the new or replacement SGM:

1. Choose the Standby Chassis and insert the new or replacement SGM in a slot that is not part of the security group.

   If all the slots are taken, reconfigure the security group and remove one of the SGM from it:

   ```
   # asg security_group
   ```

2. Insert the removable disk to the USB port of the RMA and mount it to: `/mnt/usb`

   To learn how to mount a USB drive, see Mounting and Dismounting a USB Disk (on page 318).
3. Connect to the SGM using a console connection.
4. Import the snapshot file:

   ```
   > set snapshot import <filename_without_.tar> path /mnt/usb/
   ```
5. Monitor the import progress:

   ```
   > show snapshots
   ```
6. Dismount `/mnt/usb` and remove the removable disk:

   ```
   # umount /mnt/usb
   ```
7. Revert the RMA to the snapshot image:

   ```
   > set snapshot revert <snapshotname>
   ```
8. The revert procedure can take a long time and includes reboot. When the reboot starts, continue to the next step.

To add the new or replacement SGM to the security group

Update the security group to include the new or replacement SGM:

```
# asg security_group
```

To make sure the SGM works correctly:

1. Make sure that that the new or replacement SGM is up and enforces the latest policy:

   ```
   > asg monitor
   ```
2. Make sure that all the SGMs have the same operating system version:

   ```
   # asg_version
   ```

Installing a New SGM Using a CD/DVD

To install an SGM:

1. Install the new SGM into an unoccupied slot in the standby Chassis.
2. If necessary, reconfigure the security group to include the new SGM.
3. Connect to the new SGM with a Console connection.
4. Remove the SGM boot sector:

   ```
   # eraseboot
   ```
5. Insert the CD.
6. Reboot the SGM.

Mounting and Dismounting a USB Disk

To Mount a USB device:

1. Insert the removable disk into the USB port.

2. Find the USB file system name for the USB in message log file:

   ```
   > shell run tail /var/log/messages
   ```

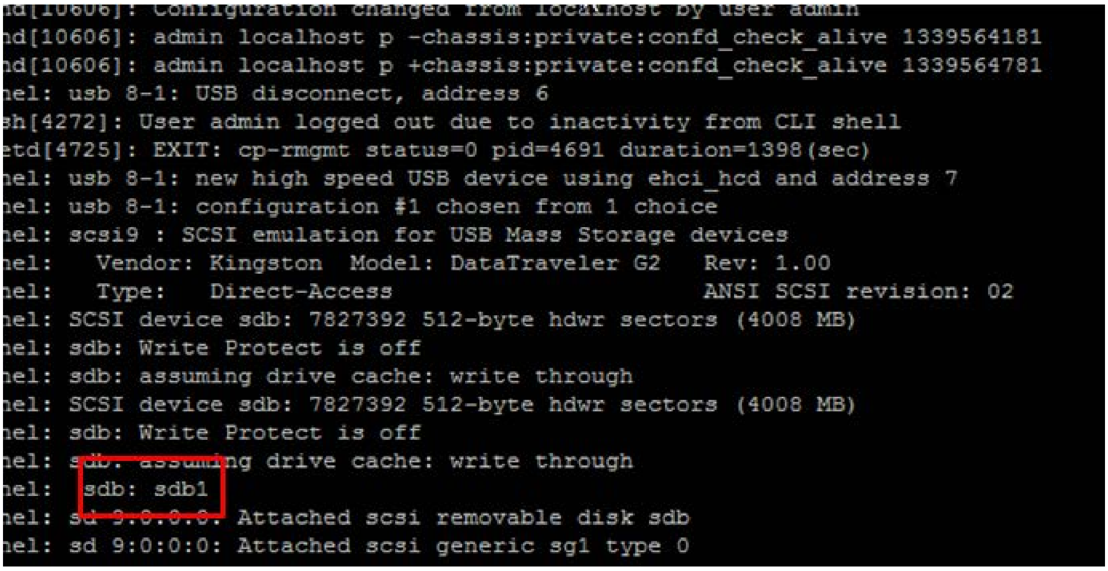

```
nd[10606]: Configuration changed from localhost by user admin
nd[10606]: admin localhost p -chassis:private:confd_check_alive 1339564181
nd[10606]: admin localhost p +chassis:private:confd_check_alive 1339564781
nel: usb 8-1: USB disconnect, address 6
sh[4272]: User admin logged out due to inactivity from CLI shell
etd[4725]: EXIT: cp-rmgmt status=0 pid=4691 duration=1398(sec)
nel: usb 8-1: new high speed USB device using ehci_hcd and address 7
nel: usb 8-1: configuration #1 chosen from 1 choice
nel: scsi9 : SCSI emulation for USB Mass Storage devices
nel:   Vendor: Kingston  Model: DataTraveler G2   Rev: 1.00
nel:   Type:   Direct-Access                      ANSI SCSI revision: 02
nel: SCSI device sdb: 7827392 512-byte hdwr sectors (4008 MB)
nel: sdb: Write Protect is off
nel: sdb: assuming drive cache: write through
nel: SCSI device sdb: 7827392 512-byte hdwr sectors (4008 MB)
nel: sdb: Write Protect is off
nel: sdb: assuming drive cache: write through
nel:  sdb: sdb1
nel: sd 9:0:0:0: Attached scsi removable disk sdb
nel: sd 9:0:0:0: Attached scsi generic sg1 type 0
```

3. If necessary, create: `/mnt/usb`
4. Mount the USB file system to your usb directory:

   ```
   > mount /dev/sdb1 /mnt/usb
   ```

To dismount the USB Disk:

1. Run:

   ```
   > umount /mnt/usb
   ```

2. Remove the USB disk.

Troubleshooting

In This Section:

Collecting System Information (asg_info)

`asg_info` is a command that collects information from the systems that generate data files and command line output.

The information is collected from these areas:

- Log files
- Configuration files
- System status
- System diagnostics

The information is sent to a compressed folder at: `/var/log/asg_info.<`*hostname*`>.<`*date*`>.tar.` By default, information is collected from all SGMs/VSs.

Commands

`asg_info` executes commands with this granularity:

- SGMs
  - All SGMs
  - Single SGM for each Chassis
  - Selective SGM
- VSX
  - Per VS
  - VS0 only
  - Selective VS
- CMM

Files

`asg_info` collects a predefined list of files from the SGM and VS folders. A global file is located in the Global folder.

Examples:

1. `latest_policy.policy.tgz` is collected as a global file, and is located in `\global\VS0\var\CPbackup\asg_backup\`

2. `dist_mode.log` is collected from the SGM and VS folders, and is located in `\SGM_1_01\VS1\var\log\dist_mode.log\`
3. `start_mbs.log` is collected from the SGM folder and not from the VS folder, and is located in `\SGM_1_01\VS0\var\log\start_mbs.log\`

Syntax

```
> asg_info [-b <sgm_ids>] [--vs <vs_ids>]  <collect_flags> [options]
    asg_info [-b <sgm_ids>] [--vs <vs_ids>]  [--user_conf <xml_filename>]
[options]
```

Parameter	Description
-b <*sgm_ids*>	Works with SGMs and/or Chassis as specified by <*sgm_ids*>. <*sgm_ids*> can be: • No <*sgm_ids*> specified or `all` shows all SGMs and Chassis • One SGM • A comma-separated list of SGMs (`1_1,1_4`) • A range of SGMs (`1_1-1_4`) • One Chassis (`Chassis1` or `Chassis2`) • The active Chassis (`chassis_active`) List of SGMs, default: all UP SGMs
-vs <*vs_ids*>	<*vs_ids*> can be: • No <*vs_ids*> (default) - Shows the current Virtual System context. • One Virtual System. • A comma-separated list of Virtual Systems (1, 2, 4, 5). • A range of Virtual Systems (VS 3-5). • `all` - Shows all Virtual Systems. **Note:** This parameter is only relevant in a VSX environment.
Collect Flags	
--all	Collect all log files and commands output
-q	Collect major log files and commands output
-f	Collect comprehensive log files and commands output
-c	Collect core dump information
-i	Collect cpinfo output
-m	Collect CMM log files
-s	Collect setup information
-a	Collect archive files
-h	Display usage message

Parameter	Description
--user_conf	Add xml configuration file with files and commands
Options	
--list	Display all the files and commands to be collected without collecting them in practice
-h	Display this help message and exit
-v	Display verbose output
-u	Upload `asg_info` output file to C*heck* Point User Center
-t	Upload `asg_info` output file using SFTP only - default is https and sftp
-uk	Upload result file using cp_uploader-k
-e	Semicolon separated list of email addresses for upload notifications

Configuration Files:

- **Default**
  `$FWDIR/conf/asg_info_config.xml`
  Files and commands are defined automatically
- **User defined**
  The user can define files and commands, following the same standard

The user can configure any command and/or file for collection, and it is used with --user_conf option.

Note - `asg_info` can run the user-defined file or the default file. They cannot be run together.

User-defined XML configuration file example:

```
<configurations>
     <collect_file_list>
     <upgrade_wizard>
          <collect_mode>-f</collect_mode>
          <path>/var/log/upgrade_wizard.log*</path>
          <per_vs>0</per_vs>
          <per_sgm>1</per_sgm>
          <delete_after_collect>0</delete_after_collect>
     </upgrade_wizard>
     <active_cmm_debug>
          <collect_mode>-m</collect_mode>
          <path>/var/log/active_cmm_debug.log</path>
```

```
           <per_vs>0</per_vs>
           <per_sgm>1</per_sgm>
           <delete_after_collect>1</delete_after_collect>
      </active_cmm_debug>
      </collect_file_list>
<cmd_list>
      <asg_stat_vs>
<mode>-f</mode>
           <pre_command></pre_command>
           <command>asg stat vs</command>
           <ipv6>0</ipv6>
           <esx>1</esx>
           <per_chassis>1</per_chassis>
           <per_vs>1</per_vs>
           <per_sgm>0</per_sgm>
           <vsx_only>1</vsx_only>
           <dest_file_name>asg_info</dest_file_name>
      </asg_stat_vs>
      <asg_if>
           <mode>-f</mode>
           <pre_command>g_all</pre_command>
           <command>asg if</command>
           <ipv6>0</ipv6>
           <esx>1</esx>
           <per_chassis>0</per_chassis>
           <per_vs>1</per_vs>
           <per_sgm>0</per_sgm>
           <vsx_only>0</vsx_only>
           <dest_file_name>asg_info</dest_file_name>
      </asg_if>
</cmd_list>
</configurations>
```

Verifiers

MAC Verification (mac_verifier)

Each MAC address contains information about the Chassis ID, SGM ID and interfaces. Use this command to make sure that the virtual MACS on physical and bond interfaces are the same for all SGMs on each Chassis. Run this command in Expert mode.

Syntax

```
# mac_verifier [-l] [-v]
# mac_verifier -h
```

Parameter	Description
`-l`	Shows MAC address consistency on the Active Chassis
`-v`	Shows information for each interface MAC
`-h`	Help screen

Example

```
# mac_verifier
Starting mac address verification on local chassis... (Chassis 1)
No inconsistency found on local chassis

Starting mac address verification on remote chassis... (Chassis 2)
MAC address inconsistency found on interface eth2-11
```

L2 Bridge Verifier (asg_br_verifier)

Use `asg_br_verifier` to make sure that:

- There are no bridge configuration problems.
- The `fdb_shadow` tables are the same.

Syntax

```
> asg_br_verifier
> asg_br_verifier -v
```

Parameter	Description
`-v`	Verbose mode

Example

```
> asg_br_verifier
```

Output

```
================================================================================

Number of entries in fdb_shadow table:

-*- 10 blades: 1_01 1_02 1_03 1_04 1_05 2_01 2_02 2_03 2_04 2_05 -*-
11

Status: OK

================================================================================
```

In this example there is a misconfiguration.

Example

```
> asg_br_verifier -v
```

Output

```
================================================================================

Number of entries in fdb_shadow table:

-*- 9 blades: 1_01 1_03 1_04 1_05 2_01 2_02 2_03 2_04 2_05 -*-
11
-*- 1 blade: 1_02 -*-
0


Status: number of entries is different

================================================================================

Collecting table info from all SGMs. This may take a while.

Table entries in fdb_shadow table:

-*- 9 blades: 1_01 1_03 1_04 1_05 2_01 2_02 2_03 2_04 2_05 -*-
address="00:00:00:00:00:00" Interface="eth1-07"
address="00:10:AA:7D:08:81" Interface="eth2-07"
address="00:1E:9B:56:08:81" Interface="eth1-07"
address="00:23:FA:4E:08:81" Interface="eth1-07"
address="00:49:DC:58:08:81" Interface="eth2-07"
address="00:7E:60:77:08:81" Interface="eth1-07"
address="00:80:EA:55:08:81" Interface="eth1-07"
address="00:8D:86:52:08:81" Interface="eth2-07"
address="00:9E:8C:7F:08:81" Interface="eth1-07"
address="00:E5:DB:78:08:81" Interface="eth2-07"
address="00:E5:F7:78:08:81" Interface="eth2-07"
-*- 1 blade: 1_02 -*-
fdb_shadow table is empty
Status: Table entries in fdb_shadow table is different between SGMs

================================================================================
```

Port Connectivity Verification (asg_pingable_hosts)

The Port Connectivity Verification feature makes sure that 61000/41000 Security System ports are connected to their hosts. When this feature is enabled, the system automatically adds a predefined value (default=50) to the Chassis Grade.

- When Port Connectivity Verification detects a host connectivity error, this value is subtracted from the Chassis Grade. The system continuously runs connectivity tests at a predefined interval (Default = 4 seconds). You can change the interval with the `asg_pingable_hosts enable -i` <*interval*> command.

Notes and Limitations

- Port Connectivity Verification is not supported for VSX
- Port Connectivity Verification only supports IPv4 addresses
- A port is considered to be down when all connected hosts fail to respond to pings

Syntax

```
# asg_pingable_hosts --help
# asg_pingable_hosts status
# asg_pingable_hosts load_ips
# asg_pingable_hosts disable
# asg_pingable_hosts enable [-i <interval>] [-monitor]
```

Parameter	Description
`--help`	Show commands and syntax

Parameter	Description
`status`	Show Port Connectivity Verification status and parameters
`load_ips`	Load IPS
`disable`	Disable Port Connectivity Verification
`enable`	Enable Port Connectivity Verification and configure options
`-i` <*interval*>	Enter a verification interval in seconds (Default = 4)
`-monitor`	Enable monitor only mode, which does not change the Chassis grade if connectivity verification detects an error.

Notes:

- `asg stat` shows the Pingable Posts and verification results in the bottom row for each Chassis.

```
> asg stat
---------------------------------------------------------------------------------
| System Status - 61000                                                         |
---------------------------------------------------------------------------------
| Up time                      | 7 days, 01:56:22 hours                         |
---------------------------------------------------------------------------------
| Current CPUs load average    | 4 %                                            |
| Concurrent connections       | 0                                              |
| Health                       | Pingable Hosts              1 Down             |
---------------------------------------------------------------------------------
| Chassis 1                    | ACTIVE                     UP / Required       |
|                              |   SGMs                      3 / 3              |
|                              |   Ports                     0 / 0              |
|                              |   Fans                      4 / 4              |
|                              |   SSMs                      2 / 2              |
|                              |   CMMs                      2 / 2              |
|                              |   Power Supplies            6 / 6              |
|                              |   Pingable Hosts            1 / 1              |
---------------------------------------------------------------------------------
| Chassis 2                    | ACTIVE                     UP / Required       |
|                              |   SGMs                      3 / 3              |
|                              |   Ports                     0 / 0              |
|                              |   Fans                      4 / 4              |
|                              |   SSMs                      2 / 2              |
|                              |   CMMs                      2 / 2              |
|                              |   Power Supplies            6 / 6              |
|                              |   Pingable Hosts            0 / 1   (!)        |



---------------------------------------------------------------------------------
```

- The **UP/Required** column shows the verification status, not the number of pingable hosts up or required. The status means:
  - 1 / 1 = OK
  - 0 / 1 when one of the pingable hosts on the list fails to reply
- Port Connectivity log files are stored at `/var/log/pingable_hosts`
- The default Port Connectivity Verification value added to the Chassis Score is 50. To change this value, run
  `> set chassis high-availability factors pnote pingable_hosts` <*factor*>

Working with Pingable Hosts

Before you can use Port Connectivity Verification, you must first define your interfaces and host IPv4 addresses in the `$FWDIR/conf/pingable_hosts.ips` configuration file. When this task is completed, you import the definitions to your SGMs and the enable Port Connectivity Verification.

Port Connectivity Verification is disabled by default.

To define interfaces and host IP addresses:

1. On an SGM, open `$FWDIR/conf/pingable_hosts.ips` in a text editor.
2. Enter the interface and host IPv4 address with this syntax:

   <*if_name*>;ipv4;<*host_ip1*>,<*host_ip2*>...

   Example: `eth0-01;ipv4;192.168.2.41,192.168.2.88,192.168.2.123`

   Each line contains one port definition, which can include one interface and many host IP addresses separated by commas. Do not put other data in this file.
3. Run:

   ```
   # pingable_hosts load_ips
   ```

Example:

```
# pingable_hosts load_ips

New IPs loaded successfully

Ports and IPs:
-------------------
eth0-1;ipv4;192.168.2.88,192.168/2.123
eth1-01;ipv4;10.2.2.1,10.10.2.2,10.30.2.3

Pingable hosts is DISABLED
```

To enable Port Connectivity Verification:

Run:

```
# pingable_hosts enable
```

Example:

```
# pingable_hosts enable
1_01:
1_02:
1_03:
No additional settings, using default values:
enable=1 interval=4 monitor=0
```

This action updates the Chassis Grade.

To disable Port Connectivity Verification:

Run:

```
# pingable_hosts disable
```

This updates the Chassis Grade.

Verifying VSX Gateway Configuration (bin vsx_verify)

Use this command to make sure that all SGMs have the same VSX Configuration: Interfaces, Routes, and Virtual Systems configuration:

- md5sum similarity on configuration files that must be identical between SGMs.
- Similarity in configuration files that must be identical but not necessarily written that way (like `/config/active`). The command uses db_cleanup report to do this.
- vsx stat among SGMs.
- vmacs/bmacs similarity.

Note - `bin vsx_verify` replaces the old verifier in asg diag and runs on a VSX system only.

Usage

Output when there is an inconsistency in the configuration:

The differences are compared in two ways:

- The return value of the command run on the SGMs with `gexec_inner_command`
- The output of the commands
  Example of difference in the command output:
  Difference between blade: 1_01 and blade: 2_01 found.
  ==
  --- 1_01
  +++ 2_01
  -73b4c20e598d6b495de7515ad4ea2fdc   /opt/CPsuite-R76/fw1/conf/fwha_vsx_conf_id.conf
  +b21dfa3feab817c3640bbb984346cdf1   /opt/CPsuite-R76/fw1/conf/fwha_vsx_conf_id.conf

When a command fails, the output contains:

Command "asg xxx" failed to run on blade "2_01"

Syntax

```
> asg vsx_verify [-a|-c|-v]
```

Parameter	Description
`-a`	Include SGMs in the administrative DOWN state
`-c`	Compare these items: • Database configuration between SGMs • Operating system and database configuration on each SGM
`-v`	Include Virtual Systems Configuration Verification table

Example

```
> bin vsx_verify –v
```

Output

```
+--------------------------------------------------------------------------------+
|Chassis 1 SGMs:                                                                 |
|1_01 1_02 1_03                                                                  |
+--------------------------------------------------------------------------------+
+--------------------------------------------------------------------------------+
|Chassis 2 SGMs:                                                                 |
|2_01* 2_02 2_03                                                                 |
+-------------------------------------------------------------------------------+

+-------------------------------------------------------------------+
|VSX Global Configuration Verification                              |
+------+---------------------------------+------------------+-------+
|SGM   |VSX Configuration Signature      |Virtual Systems   |State  |
|      |VSX Configuration ID             |Installed\Allowed |       |
+------+---------------------------------+------------------+-------+
|all   |8ef02b3e73386afd6e044c78e466ea82 |5\25              |UP     |
|      |9                                |                  |       |
+------+---------------------------------+------------------+-------+

+--------------------------------------------------------------------------+
|Virtual Systems Configuration Verification                                |
+----+-----+-----------+---------------+----------------+---------+--------+
|VS  |SGM  |VS Name    |VS Type        |Policy Name     |SIC State|Status  |
+----+-----+-----------+---------------+----------------+---------+--------+
|0   |all  |VSX_OBJ    |VSX Gateway    |Standard        |Trust    |Success |
+----+-----+-----------+---------------+----------------+---------+--------+
|1   |all  |VSW-INT    |Virtual Switch |<Default Policy>|Trust    |Success |
+----+-----+-----------+---------------+----------------+---------+--------+
|2   |all  |VSW-INT    |Virtual Switch |<Not Applicable>|Trust    |Success |
+----+-----+-----------+---------------+----------------+---------+--------+
|3   |all  |VS-1       |Virtual System |Standard        |Trust    |Success |
+----+-----+-----------+---------------+----------------+---------+--------+
|4   |all  |VS-2       |Virtual System |Standard        |Trust    |Success |
+----+-----+-----------+---------------+----------------+---------+--------+
Comparing Routes DB & OS. This procedure may take some time...
Press 'y' to skip this procedure...
Comparing..

+--------------------------------------------------------------------------+
|Summary                                                                   |
+--------------------------------------------------------------------------+
|VSX Configuration Verification completed successfully                     |
+--------------------------------------------------------------------------+

All logs collected to /var/log/vsx_verify.1360846320.log
```

Example

```
> asg vsx_verify -v -a
```

Output

```
+--------------------------------------------------------------------------+
|Chassis 1 SGMs:                                                           |
|1_01* 1_02 1_03 1_04                                                      |
+--------------------------------------------------------------------------+
+--------------------------------------------------------------------------+
|Chassis 2 SGMs:                                                           |
|2_01 2_02 2_03 2_04                                                       |
+--------------------------------------------------------------------------+

+-------------------------------------------------------------------+
|VSX Global Configuration Verification                              |
+------+---------------------------------+------------------+-------+
|SGM   |VSX Configuration Signature      |Virtual Systems   |State  |
|      |VSX Configuration ID             |Installed\Allowed |       |
+------+---------------------------------+------------------+-------+
|1_01  |8ef02b3e73386afd6e044c78e466ea82 |5\25              |UP     |
|      |9                                |                  |       |
+------+---------------------------------+------------------+-------+
|1_02  |8ef02b3e73386afd6e044c78e466ea82 |5\25              |UP     |
|      |9                                |                  |       |
+------+---------------------------------+------------------+-------+
|1_03  |8ef02b3e73386afd6e044c78e466ea82 |5\25              |UP     |
|      |9                                |                  |       |
+------+---------------------------------+------------------+-------+
|1_04  |8ef02b3e73386afd6e044c78e466ea82 |5\25              |DOWN   |
|      |9                                |                  |       |
+------+---------------------------------+------------------+-------+
|2_01  |8ef02b3e73386afd6e044c78e466ea82 |5\25              |UP     |
|      |9                                |                  |       |
+------+---------------------------------+------------------+-------+
|2_02  |8ef02b3e73386afd6e044c78e466ea82 |5\25              |UP     |
|      |9                                |                  |       |
+------+---------------------------------+------------------+-------+
|2_03  |8ef02b3e73386afd6e044c78e466ea82 |5\25              |UP     |
|      |9                                 |                  |       |
+------+----------------------------------+------------------+-------+
|2_04  |8ef02b3e73386afd6e044c78e466ea82 |5\25              |UP     |
|      |9                                 |                  |       |
+------+----------------------------------+------------------+-------+

+-------------------------------------------------------------------------+
|Virtual Systems Configuration Verification                               |
+----+-----+-----------+--------------+----------------+---------+--------+
|VS  |SGM  |VS Name    |VS Type       |Policy Name     |SIC State|Status  |
+----+-----+-----------+--------------+----------------+---------+--------+
|0   |all  |VSX_OBJ    |VSX Gateway   |Standard        |Trust    |Success |
+----+-----+-----------+--------------+----------------+---------+--------+
|1   |all  |VSW-INT    |Virtual Switch |<Default Policy>|Trust    |Success |
+----+-----+-----------+--------------+----------------+---------+--------+
|2   |all  |VSW-INT    |Virtual Switch |<Not Applicable>|Trust    |Success |
+----+-----+-----------+--------------+----------------+---------+--------+
|3   |all  |VS-1       |Virtual System |Standard        |Trust    |Success |
+----+-----+-----------+--------------+----------------+---------+--------+
|4   |all  |VS-2       |Virtual System |Standard        |Trust    |Success |
+----+-----+-----------+--------------+----------------+---------+--------+
Comparing Routes DB & OS. This procedure may take some time...
Press 'y' to skip this procedure...
Comparing..

+-------------------------------------------------------------------------+
|Summary                                                                  |
+-------------------------------------------------------------------------+
|VSX Configuration Verification completed with the following errors:     |
|1. [1_02:1] eth1-06 operating system   address doesn't match             |
|2. [1_02:1] eth1-06 DB address doesn't match                             |
|3. [1_01:1] Found inconsistency between addresses in operating system  ,DB and NCS ofeth1-06 |
|                                                                         |
+-------------------------------------------------------------------------+
All logs collected to /var/log/vsx_verify.1360886320.log
```

Resetting SIC (g_cpconfig sic init)

Use this command to reset Secure Internal Communication (SIC) between the gateway and the Security Management server. For example, if you replace the management server you must reset the SIC.

Important - This procedure causes downtime for the system and traffic outage because all SGMs are rebooted.

Resetting SIC on a Security Gateway or VSX Gateway (VS0)

The procedure to reset SIC on a Security Gateway or VSX Gateway (VS0) has these basic steps:

1. Initialize SIC on the gateway.
2. Initialize SIC in SmartDashboard.
3. Make sure that Trust is established on the gateway.

To initialize SIC on the Gateway:

1. Use a serial console to connect to the gateway.
2. Enter Expert mode.
3. Find out which SGM is the SMO:
   ```
   > asg stat –i tasks
   ```
4. Run:
   ```
   # g_cpconfig sic init <activation_key>
   ```

Notes:

- The SIC Reset procedure lasts about 3 to 5 minutes.
- During the SIC reset procedure, on a Security Gateway, all SGMs other than the SMO reboot.
- On a VSX Gateway: Do the next steps immediately when this procedure is done.

To initializing SIC in SmartDashboard:

1. On the gateway object, open the **General Properties** > **Communication** window.
2. Click **Reset**.
3. Enter the same activation key used when you initialized SIC on the gateway.
4. Click **Initialize**.
5. On a VSX Gateway:
   a) Install the policy on the VSX Gateway.
   b) At the serial console connection to the gateway, press `c` to complete the procedure.

Note - At this stage, all SGMs except the SMO, reboot.

To make sure that Trust is established on the Gateway:

Run:

```
# g_cpconfig sic state

-*- 6 blades: 1_01 1_02 1_03 2_01 2_02 2_03 -*-
Trust State: Trust established
```

Reset SIC for non-VS0 Virtual Systems

To reset SIC on Virtual Systems that are not VS0 (a non-VSX object):

1. Log into the SMO with a SSH client.
2. Go to Expert mode.
3. Go to the applicable context ID:
   ```
   # vsenv <vsid>
   ```
4. Initialize SIC:
   ```
   # g_cpconfig sic init
   ```
5. Revoke the VSID certificate defined in the management server.

   See Part II of sk34098 http://supportcontent.checkpoint.com/solutions?id=sk34098 for the detailed procedure.
6. In SmartDashboard, open and save the Virtual System object.

   This pushes the configuration to the management server and re-establishes SIC trust with the SMO.
7. Install a policy on the Virtual System.

Troubleshooting SIC reset

SIC reset requires 3-5 minutes. If SIC reset was interrupted (for example by loss of network connectivity), run `g_cpconfig sic state` to get the SIC state. If the SIC State is:

SIC state	Do this
Trust established	Repeat the SIC reset procedure
Initialized but Trust was not established	1. Reboot all SGMs. 2. In **SmartDashboard > General Properties > Communication**, initialize SIC. 3. Install the policy.

SIC Cleanup

To resolve other SIC issues, do a SIC cleanup. There are two ways to do a SIC cleanup:

Run:

```
# asg_blade_config reset_sic -reboot_all <activation_key>
```

OR

1. Shutdown all SGMs (but not the SMO) using `ccutil` in Expert mode.
2. Shutdown all SGMs (but not the SMO) using `ccutil` in Expert mode.
3. Connect to the SMO using a serial console.
4. Initialize SIC in **SmartDashboard > General Properties > Communication**.
5. Install policy on the SMO.
6. Turn on all SGMs.

Debug files

These are the 61000/41000 Security System debug files:

Feature	Debug File
FWK	`$FWDIR/log/fwk.elg.*`
Policy	`$FWDIR/log/cpha_policy.log.*`
SGM Configuration / Pull Configuration	`$FWDIR/log/blade_config.*`
Alerts	`/var/log/send_alert.*`
Distribution	`$FWDIR/log/dist_mode.log.*`
Installation – OS	`/var/log/anaconda`
Installation – 61000/41000 Security System	`/var/log/start_mbs.log`
Installation – 61000/41000 Security System	`/var/log/mbs.log`

Feature	Debug File
Dynamic Routing	/var/log/routed.log
CPD	$CPDIR/log/cpd.elg
FWD	$FWDIR/log/fwd.elg
General	/var/log/messages*
Log servers	/var/log/log_servers*
Pingable hosts	/var/log/pingable_hosts*
Clish auditing	/var/log/auditlog*
Command auditing	/var/log/asgaudit.log*
VPND	$FWDIR/log/vpnd.elg*
Reboot logs	/var/log/blade_reboot_log

Index