



Check Point
SOFTWARE TECHNOLOGIES LTD.

15 May 2017

GAIA ADVANCED ROUTING

R80.10

Administration Guide

Classification: [Protected]



Check Point
SOFTWARE TECHNOLOGIES LTD.

INFINITY

© 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.



Check Point R80.10

For more about this release, see the R80.10 home page
<http://supportcontent.checkpoint.com/solutions?id=sk111841>.



Latest Version of this Document

Download the latest version of this document
http://supportcontent.checkpoint.com/documentation_download?ID=54803.

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 Gaia Advanced Routing R80.10 Administration Guide.



Searching in Multiple PDFs

To search for text in all the R80.10 PDF documents, download and extract the complete R80.10 documentation package

<http://downloads.checkpoint.com/dc/download.htm?ID=54846>.

Use **Shift-Control-F** in Adobe Reader or Foxit reader.

Revision History

Date	Description
15 May 2017	First release of this document

Contents

Important Information.....	3
Introduction to Gaia Advanced Routing	9
IPv6 Support.....	9
Configuring IPv6 Support - WebUI.....	9
Configuring IPv6 Support - CLI	9
VRRP Support for IPv6	10
DHCP Relay	11
DHCP Services	11
DHCP Services Initial Setup - Management Servers.....	12
Configuring DHCP Relay on the Security Gateway - WebUI	13
BOOTP/DHCP Parameters.....	13
Configuring DHCP Relay on the Security Gateway - CLI (bootp)	14
BOOTP Show Commands	14
Configuring DHCP Security Policy.....	15
IPv6 DHCP Relay	16
Configuring DHCPv6 Relay on the Security Gateway - WebUI	16
Configuring DHCPv6 Relay on the Security Gateway - CLI	18
Monitoring DHCPv6 Relay (show dhcp).....	18
Configuring DHCPv6 Security Policy.....	18
BGP	21
Configuring BGP 4-Byte AS	21
Support for IPv6 BGP (BGP-4 Multiprotocol Extensions)	22
BGP Sessions (Internal and External)	22
Preventing Private AS Numbers from Propagating	23
Peer Local AS	23
AS Override	24
BGP and ECMP	24
BGP Route Refresh.....	24
BGP Path Attributes	25
BGP Multi-Exit Discriminator.....	26
BGP Interactions with IGPs	26
Inbound BGP Route Filters.....	26
Redistributing Routes to BGP.....	27
BGP Communities	27
BGP Route Reflection.....	27
BGP Confederations	28
External BGP (EBGP) Multihop Support	29
BGP Route Dampening	29
TCP MD5 Authentication	30
Configuring BGP - WebUI	30
BGP Global Settings	31
BGP Miscellaneous Settings	33
BGP AS Peer Group Settings	34
BGP Remote Peers	35
Configuring BGP - CLI (bgp)	43
External BGP Commands	43

BGP Peer Commands	44
BGP Confederation Commands	48
BGP Route Reflection Commands.....	52
BGP Route Dampening Commands.....	53
Internal BGP Commands	54
BGP Communities Commands.....	59
BGP Show Commands	59
Rate Limiting for DoS Mitigation	60
IGMP	68
IGMP Version 3	68
Configuring IGMP - WebUI	69
Configuring IGMP - CLI (igmp)	71
Configuring Interfaces for IGMP	71
Monitoring IGMP (show igmp).....	73
IP Broadcast Helper	74
Configuring IP Broadcast helper - WebUI	74
Configuring IP Broadcast Helper - CLI (iphelper)	75
Forwarding Non-Local Packets	75
IP Broadcast Helper interfaces	75
Monitoring IP Broadcast Helper.....	76
RIP	77
RIP 2.....	77
Network Mask	77
Authentication	77
RIP 1	78
Network Mask	78
Auto Summarization	78
Virtual IP Address Support for VRRP	78
Configuring RIP - WebUI	78
RIP Global Settings.....	79
RIP Interface Options.....	79
Configuring RIP - CLI (rip).....	81
RIP Global Commands	81
RIP Interface Commands.....	82
Monitoring RIP	84
Monitoring RIP - WebUI	84
RIP Show Commands.....	84
OSPF.....	85
Types of Areas.....	85
Area Border Routers.....	86
High Availability Support for OSPF.....	86
OSPF Forced Hellos	87
Configuring OSPF - WebUI	87
Configuring Global Settings.....	88
Configuring OSPF Interfaces	89
Configuring OSPF Virtual Links	92
Configuring OSPF Areas	93
Configuring OSPF - CLI (ospf)	96
OSPF Global Settings.....	97
OSPF Interfaces.....	98
OSPF Virtual Links.....	101

OSPF Areas	103
OSPF and IPv6 OSPF Show Commands.....	104
IPv6 OSPF	110
Configuring OSPFv3 with IPv6 VRRP	110
Configuring IPv6 OSPF - WebUI	110
IPv6 OSPF Global Settings - WebUI.....	111
IPv6 OSPF Areas - WebUI	112
IPv6 OSPF Interfaces - WebUI	114
Configuring IPv6 OSPF - clish (ipv6 ospf)	116
IPv6 OSPF Global Settings - clish	116
IPv6 OSPF Areas - clish	117
IPv6 OSPF Interfaces - clish	119
IPv6 OSPF Configuration Examples - clish	121
Monitoring IPv6 OSPF	121
Monitoring IPv6 OSPF in the WebUI.....	121
Show IPv6 OSPF Commands.....	121
Route Aggregation	122
Configuring Route Aggregation - WebUI.....	122
Configuring Route Aggregation - CLI (aggregate).....	124
Routing Policy Configuration.....	126
Configuring Inbound Route Filters - WebUI	126
Add BGP Policy Window	128
Fine Tuning Policies	130
Configuring Route Redistribution - WebUI.....	132
Redistributing IPv4 Routes - WebUI.....	133
Redistributing IPv6 Routes - WebUI.....	140
Configuring Route Redistribution - CLI	142
Redistributing IPv4 Routes - CLI.....	142
Redistributing IPv6 Routes - CLI.....	154
Configuring Route Maps - CLI (routemap)	159
Set Route Map Commands.....	159
Show Routemap Commands.....	168
Routemap Protocol Commands.....	168
Supported Route Map Statements by Protocol	169
Route Map Examples	171
Routing Options.....	173
Routing Options (Apply, Reset and Reload) - WebUI	173
Equal Cost Path Splitting	173
Configuring Equal Cost Path Splitting - WebUI.....	173
Configuring Equal Cost Path Splitting - CLI (max-path-splits).....	174
Kernel Options- Kernel Routes.....	174
Configuring Kernel Routes - WebUI.....	174
Configuring Kernel Routes - CLI (kernel-routes)	174
Protocol Rank.....	174
Default Ranks	175
Configuring Protocol Rank - WebUI.....	175
Configuring Protocol Rank - CLI (protocol-rank)	175
Advanced Routing Options - Wait for Clustering	176
Configuring Wait for Clustering - WebUI	176
Configuring Wait for clustering - CLI (router-options).....	177
Advanced Routing Options - Auto Restore of Interface Routes.....	177

Trace Options	178
Trace Options - WebUI	178
Trace Options - CLI	178
Router Discovery	186
How Router Discovery Works	186
Configuring Router Discovery - WebUI	187
Configuring Router Discovery - CLI (rdisc)	188
Monitoring ICMP Router Discovery	189
IPv6 Router Discovery	190
IPv6 Router Discovery and VRRP	190
IPv6 Discovery - WebUI	190
IPv6 Discovery - clish (ipv6 rdisc6)	193
Monitoring IPv6 Router Discovery	196
Monitoring IPv6 Router Discovery in the WebUI	196
Show IPv6 Router Discovery Commands	197
Policy Based Routing	198
Configuring Policy Based Routing - WebUI	198
Action Table Parameters	199
Policy Rule Parameters	200
Configuring Policy Based Routing - CLI	200
Monitoring Policy Based Routing	202
PIM	203
Dense Mode	203
Sparse Mode	203
Source-Specific Multicast (SSM) Mode	203
PIM Dense Mode State Refresh	204
Configuring PIM - WebUI	204
PIM Global Settings	205
PIM Interfaces	205
Disabling PIM	206
Configuring PIM-DM Advanced Options (Optional)	206
Configuring PIM-SM Advanced Options (Optional)	207
Configuring PIM-SM Bootstrap and Rendezvous Point Settings	210
Configuring PIM - CLI (pim)	214
PIM Global Settings	214
PIM Interfaces	214
PIM Sparse Mode	215
Timer and Assert Rank Parameters for Dense Mode and Sparse Mode	219
Static Multicast Routes	222
Configuring Static Multicast Routes - WebUI	223
Configuring Static Multicast Routes - CLI	224
Monitoring and Troubleshooting PIM	224
PIM Trace Options	224
PIM Show and MFC Commands	226
Routing Monitor	228
Monitoring Routes - WebUI	228
Monitoring Routes - CLI (show route)	228
Monitoring the Routing Daemon - CLI (show routed)	229
Monitoring the Multicast Forwarding Cache - CLI (show mfc)	230

Regular Expressions and Character Sets	232
Regular Expression Syntax	232
Special Characters in Clish	232

Introduction to Gaia Advanced Routing

In This Section:

IPv6 Support.....	9
VRRP Support for IPv6.....	10

Dynamic Routing is fully integrated into the Gaia WebUI and command-line shell. BGP, OSPF and RIP are supported.

Dynamic Multicast Routing is supported, using PIM (Sparse Mode (SM), Dense Mode (DM) and Source-Specific Multicast (SSM)) and IGMP.

For information about Dynamic Routing features that are newly supported in Gaia updates and releases, see sk98226: *Dynamic Routing and VRRP Features on Gaia OS*
<http://supportcontent.checkpoint.com/solutions?id=sk98226>.

IPv6 Support

Gaia supports IPv6.

Before you can configure IPv6 addresses and IPv6 static routes on a Gaia Security Management Server or Security Gateway you must:

1. Enable IPv6 support for the Gaia operating system and Firewall.
2. On the Security Management Server, install and enable an IPv6 license.
3. Create IPv6 objects in SmartConsole.
4. Create IPv6 Firewall rules in SmartConsole.

Configuring IPv6 Support - WebUI

1. In the WebUI tree view, click **System Management > System Configuration**.
2. In the **IPv6 Support** area, click **On**.
3. Click **Apply**.

Configuring IPv6 Support - CLI

The IPv6-state feature configures IPv6 support.

Description Use this command to enable or disable IPv6 support.

Syntax
set ipv6-state off
set ipv6-state on
show ipv6-state

Parameters

Parameter	Description
on off	Turns IPv6 support on or off.

VRRP Support for IPv6

VRRP for IPv6 follows the same guidelines and limitations as VRRP for IPv4. If both IPv6 and IPv4 VRRP are configured, they must be symmetrical in terms of master state and fail over behavior.

To configure VRRP support for IPv6 on Gaia:

1. In clish, run this command to enable IPv6:
Hostname> set ipv6-state on
2. Configure a VIP address for each VRRP IPv6 group:
 - a) Configure a link-local VIP address for the VRRP IPv6 group. Run:
set ipv6 vrrp6 interface VALUE monitored-circuit vrid VALUE address VALUE on
Example:
Hostname> set ipv6 vrrp6 interface eth2.10 monitored-circuit vrid 11 address fe80::250:56ff:fea3:4321 on
Note - Configure a link-local address that is on the same subnet as the hosts.
 - b) Configure a link-global VIP address for the VRRP IPv6 group. Run:
set ipv6 vrrp6 interface VALUE vrid VALUE address VALUE on
For example:
Hostname> set ipv6 vrrp6 interface eth2.10 vrid 11 address 2610:18:8104:1A::5 on
3. Run this command to save the configuration:
Hostname>save configuration

To verify the configuration:

In the WebUI, go to **High Availability > IPv6 VRRP**

OR

In clish, run: show ipv6 vrrp interfaces

DHCP Relay

In This Section:

DHCP Services	11
DHCP Services Initial Setup - Management Servers.....	12
Configuring DHCP Relay on the Security Gateway - WebUI.....	13
Configuring DHCP Relay on the Security Gateway - CLI (bootp)	14
BOOTP Show Commands	14
Configuring DHCP Security Policy	15
IPv6 DHCP Relay.....	16

BOOTP/DHCP Relay extends Bootstrap Protocol (BOOTP) and Dynamic Host Configuration Protocol (DHCP) operations across multiple hops in a routed network. In standard BOOTP, all interfaces on a LAN are loaded from a single configuration server on the LAN. BOOTP Relay allows configuration requests to be forwarded to, and serviced from, configuration servers outside the LAN.

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

- **Redundancy:** Configure an interface on the Check Point system to relay client configuration requests to all the listed servers simultaneously.
- **Load Balancing:** Configure multiple interfaces on the Check Point system to relay client configuration requests to different servers.
- **Management:** Centrally manage client configurations in large enterprise environments across multiple LANs.

The Gaia implementation of BOOTP Relay is compliant with RFC 951, RFC 1542, and RFC 2131. BOOTP Relay supports Ethernet and IEEE 802 LANs with canonical MAC byte ordering, on clients that specify Bootp htype=1: 802.3 and FDDI.

When an interface configured for BOOTP Relay receives a boot request, it forwards the request to all the servers in its server list. It does this after waiting a specified length of time for a local server to answer the boot request. If a primary IP is specified, it stamps the request with that address. Otherwise, it stamps the request with the lowest numeric IP address specified for the interface.

DHCP Services

To allow the DHCP relay traffic, it is necessary to configure explicit Security Policy rules with the DHCP relay services.

Such explicit Rule Base configuration is required for these reasons:

- The DHCP relay agents and DHCP servers cannot automatically match replies with requests.
- Clients do not necessarily have a source IP when they send their initial request. The Security Policy has to allow DHCP broadcasts from *Any* source to the DHCP Server or DHCP Relay.
- The dhcp-request and dhcp-reply services use Check Point's Stateful Inspection Engine to do Stateful inspection of DHCP traffic. If you do not handle DHCP Relay traffic with these services (for example: a service of *Any* in the Security Policy or implied rules) the Security Gateway can drop the traffic.

For Gateways that are R77.20 or higher, the applicable dhcp services are the new dhcp services: *dhcp-request* and *dhcp-reply*. The procedures in this chapter are compatible with the new DHCP services. For Gateways that are older than R77.20, refer to sk98839

<http://supportcontent.checkpoint.com/solutions?id=sk98839>.

For DHCPv6, the services are *dhcp-request*, *dhcp-reply* and *dhcp-relay*.

DHCP Services Initial Setup - Management Servers

This procedure shows how to configure the DHCP services on the Security Management Server or the Multi-Domain Server.

To configure the new DHCP services on the server:

1. Connect to the command line on the Security Management Server or the Multi-Domain Server (over SSH, or console).
2. Log in to Expert mode.
3. Examine the contents of all the related table.def files. This file is in \$FWDIR/lib, and possibly other locations, such as the backwards-compatibility sub-directories. For file locations, refer to sk98339.

```
[Expert@HostName:0]# grep -E
"no_hide_services_ports|no_fold_services_ports"
/path_to_relevant/table.def
```

4. If UDP port 67 and UDP port 68 are configured in the no_hide_services_ports or the no_fold_services_ports tables, edit the related table.def file and remove these ports.

```
[Expert@HostName:0]# vi /path_to_relevant/table.def
```

Change from:

```
no_hide_services_ports = { <4500,17>, <500,17>, <259,17>, <1701,17>, ...,
<68,17>, <67,17> }
no_fold_services_ports = { <4500,17>, <500,17>, <259,17>, <1701,17>, ...,
<68,17>, <67,17> }
```

To:

```
no_hide_services_ports = { <4500,17>, <500,17>, <259,17>, <1701,17>, ...
}
no_fold_services_ports = { <4500,17>, <500,17>, <259,17>, <1701,17>, ...
}
```

Note - These table changes are only necessary if one or more VSX or ClusterXL clusters run DHCP Relay. You can skip it if DHCP Relay is only used on VRRP clusters or stand-alone Security Gateways.

5. **Install Policy.**

Configuring DHCP Relay on the Security Gateway - WebUI

Use the WebUI to enable BOOTP/DHCP Relay on each interface. If the interface is enabled for relay, you can set up a number of servers to which to forward BOOTP/DHCP requests.

To enable BOOTP/DHCP relay on an Interface

1. Open the **Advanced Routing > DHCP Relay** page of the WebUI.
2. Click **Add**.
The **Add BOOTP/DHCP Relay** window opens.
3. Select an **Interface** on which you want to enable BOOTP/DHCP.
4. Optional: Enter values for one or more of these parameters ("**BOOTP/DHCP Parameters**" on page 13):
 - **Primary Address**
 - **Wait Time**
5. Define the IPv4 address of each relay to which you want to forward BOOTP/DHCP requests ("**BOOTP/DHCP Parameters**" on page 13). For each relay:
 - a) Click **Add**
 - b) In the **Add Relay** window, enter the **IPv4 address** of the relay
 - c) Click **OK**.
6. Click **Save**.

To disable BOOTP/DHCP relay on an interface

1. Open the **Advanced Routing > DHCP Relay** page of the WebUI.
2. Select an interface.
3. Click **Delete**.

BOOTP/DHCP Parameters

Parameter	Description
Primary Address	The IP address to use as the BOOTP/DHCP router address. If you enter an IP address, all BOOTP/DHCP requests received on the interface are stamped with this gateway address. This can be useful on interfaces with multiple IP addresses (aliases).
Wait Time	The minimum time to wait (in seconds) for a local configuration server to answer the boot request before forwarding the request through the interface. This delay provides an opportunity for a local configuration server to reply before attempting to relay to a remote server. Set the wait time to a sufficient length to allow the local configuration server to respond before the request is forwarded. If no local server is present, set the time to zero (0).

Relay to Server	The IPv4 address of the BOOTP/DHCP configuration server to which to forward BOOTP/DHCP requests. You can configure relay to multiple configuration servers independently on each interface. Configuring different servers on different interfaces provides load balancing, while configuring multiple servers on a single interface provides redundancy. The server IPv4 address cannot be an address belonging to the local machine.
------------------------	---

Configuring DHCP Relay on the Security Gateway - CLI (bootp)

Use these commands to configure BOOTP properties for specific interfaces.

```
set bootp interface <if_name>
    primary ip_address wait-time <0-65535> on
    relay-to ip_address <on | off>
off
```

Arguments

Parameter	Description
primary ip_address wait-time <0-65535> on	The ip_address to stamp as the gateway address on all BOOTP requests. The wait-time value is the minimum seconds to wait before forwarding a bootp request. A client-generated bootp request includes the elapsed time after the client began to boot. The bootp relay does not forward the request until the indicated elapsed time at least equals the specified wait time. This delay lets a local configuration server reply, before it relays to a remote server.
relay-to ip_address <on off>	The server to which BOOTP requests are forwarded. You can specify more than one server.
off	Disables BOOTP on the specified interface.

Note - For ClusterXL gateways, set the `fwx_dhcp_relay_nat` kernel parameter to 0 on each cluster member.

BOOTP Show Commands

Use this group of commands to monitor and troubleshoot the BOOTP implementation.

```
show bootp
    interfaces
    interface if_name
    stats
    stats receive
    stats request
    stats reply
```

Configuring DHCP Security Policy

You configure the DHCP services on these ports:

- DHCP requests from a DHCP client are sent as UDP unicasts or broadcasts with a source port of 68 and a destination port of 67. The source IP may be 0.0.0.0 if the client does not have an IP address yet.
- DHCP replies to a client are sent as UDP unicasts or broadcasts with a source port of 67 and a destination port of 68.
- DHCP relay traffic between relay and server is sent as UDP unicasts with source port of 67 and destination port of 67.

To configure DHCP Security Policy:

1. Go to the **main SmartConsole Menu** () **Global Properties > Firewall**. If the **Accept outgoing packets originating from gateway implied rule** is enabled, then from the drop-down menu, select **Last** or **Before Last**.

2. Create a host object for the DHCP server. In the SmartConsole main view, go to **Objects > New Host**.

The **New Host** window opens.

- a) Enter the object name.
- b) Enter the IPv4 address of the DHCP server.
- c) Click **OK**.

3. Create a host object for the Global Broadcast. In the SmartConsole main view, go to **Objects > New Host**.

The **New Host** window opens.

- a) Enter the object name
- b) Enter the IPv4 Address of 255.255.255.255.
- c) Click **OK**.

4. Create a Client Network object. In the SmartConsole main view, go to **Objects > New Network**. The **New Network** window opens.

- a) Enter the object name.
- b) Enter the **Network address** and **Net mask** to which the which the DHCP clients are connected.
- c) Click **OK**.

5. Make sure that the legacy DHCP configuration does not exist:

- a) Delete/disable all security rules for DHCP traffic that use these legacy services:
 - bootp
 - dhcp-relay
 - dhcp-req-localmodule
 - dhcp-rep-localmodule
- b) Delete/disable all manual NAT rules for legacy DHCP configuration. For more about NAT rules, see sk97566 <http://supportcontent.checkpoint.com/solutions?id=sk97566>.

6. Configure the required Security Policy rules with the new DHCP services (*dhcpv6-request* and *dhcpv6-reply*).

Note - Use the DHCP-relay object which you configured on the Security Gateway ("**Configuring DHCP Relay on the Security Gateway - CLI (bootp)**" on page 14). For its value, enter the name of the Security Gateway which runs DHCP Relay.

An example for a Rule Base with the DHCP relay services:

Source IP	Destination IP	Service	Action	Description of the rule
Any	Global_Broadcast	dhcp-request	Accept	Source IP must be <i>Any</i> . A value of <i>0.0.0.0</i> does not work.
<DHCP_Relay> Client Network	DHCP_Server	dhcp-request	Accept	In some situations, the DHCP client sends some requests directly to the DHCP Server.
<DHCP_Relay>	Client_Network Global_Broadcast	dhcp-reply	Accept	The replies can be unicast or broadcast based on the DHCP client options.
DHCP_Server	Client_Network Global_Broadcast	dhcp-reply	Accept	The replies can be unicast or broadcast based on the DHCP client options. In some situations, the DHCP server sends some requests directly to the DHCP client.

7. **Install Policy** on the related Security Gateways.

IPv6 DHCP Relay

IPv6 DHCP relay is configured on the Security Gateway through the WebUI or the CLI. You can configure each interface to send requests to a different set of DHCPv6 servers.

Best Practices:

- R80.10 and newer management servers include pre-defined DHCPv6 services. These built-in services do stateful inspection of DHCPv6 traffic for enhanced security. Upgrade your older management servers before you configure DHCPv6 relay.
- In a cluster environment, You must configure DHCPv6 on each cluster member.

For more technical information on the DHCPv6 protocol and IPv6 DHCP relay, refer to RFC 3315.

Configuring DHCPv6 Relay on the Security Gateway - WebUI

To enable DHCP relay for IPv6 on an interface:

1. Open the **Advanced Routing > IPv6 DHCP Relay** page of the WebUI.
2. Click **Add**.
The **Add IPv6 DHCP Relay** window opens.
3. Select an **Interface** on which to enable DHCPv6 relay.
4. Select the settings for **Send Interface-ID**.
5. Optional: Enter a value for **Wait Time**.

6. Add an IPv6 unicast address for each relay destination to which you want to forward DHCPv6 requests.

For each relay destination:

- a) Click the **Add** button in the **Relays** section of the **Add IPv6 DHCP Relay** window.

The **Add Relay** window opens.

- b) Enter the IPv6 address of the relay destination (the DHCP server).

- c) Click **Ok**.

7. Click **Save**.

8. In the **Monitoring** tab, make sure that the DHCPv6 relay runs on the required interface(s).

DHCPv6 parameters

Parameter	Description	Values	Default Value
Interface	The name of an interface to run DHCPv6 Relay on. You can enable DHCPv6 Relay for multiple interfaces, and configure each interface independently. The interface must be directly connected to the DHCPv6 clients. It is not necessary to enable DHCPv6 Relay on the interface connected to the DHCPv6 Server.		
Wait Time	The minimum time to wait (in seconds) for a local configuration server to answer the boot request before forwarding the request through the interface. This delay provides an opportunity for a local configuration server to reply before attempting to relay to a remote server. Set the wait time to a sufficient length to let the local configuration server respond before the request is forwarded. If no local server is present, or if another server listens on the same interface, and it is the preferred server, set the time to zero (0).	0-65535.	0
Interface ID	When the DHCPv6 relay request includes the interface ID value, the server copies it from the relay request to the reply message. The interface ID identifies from which interface the DHCPv6 request was received. Use this option when the IPv6 address of the interface that runs the DHCPv6 relay is not sufficient to uniquely identify that interface.	On, Off	On

Parameter	Description	Values	Default Value
Relay to Server	The IPv6 address of the DHCP configuration server/relay to which to forward DHCP requests. You can configure relay to multiple configuration servers independently on each interface. Configuration of different servers on different interfaces provides load balancing, while configuration of multiple servers on one interface provides redundancy. The server IPv6 address cannot be an address which belongs to the local computer.	Usually a DHCP server, but can be another DHCP relay	

Configuring DHCPv6 Relay on the Security Gateway - CLI

Use these commands to configure DHCP properties for specific interfaces:

```
set ipv6 dhcp6relay interface <if_name>
    wait-time {default | <0 - 65535>}
    interface-id {on | off}
    relay-to <ipv6_address> {on | off}
    {on | off}
```

Monitoring DHCPv6 Relay (show dhcp)

Use this group of commands to monitor and troubleshoot the DHCP implementation:

```
show ipv6 dhcp6relay
    interfaces
    interface <if_name>
    stats
```

Configuring DHCPv6 Security Policy

DHCPv6 clients and servers send and receive messages through UDP. A special link-scoped multicast address is defined. This address lets DHCPv6 clients request configuration information when they do not know the IPv6 address of a relay or server:

All_DHCPv6_Relay_Agents_and_Servers (FF02::1:2 - all DHCPv6 Relays and Servers on the local link)

- The client sends DHCPv6 requests as UDP unicasts or multicasts with source port 546 and destination port 547. The related security policy service is **dhcpv6-request**.
- DHCPv6 replies to a client are sent as UDP unicasts to the client's IPv6 link-local address. They are sent with source port 547 and destination port 546. The related security policy service is **dhcpv6-reply**.
- The relay and server send DHCPv6 traffic between them as IPv6 UDP unicasts with source port 547 and destination port 547. Multiple server addresses can be specified. Each server address must be an IPv6 unicast address. Each server address can refer to a DHCPv6 server or one more DHCPv6 relay. The related security policy service is **dhcpv6-relay**.
- Unlike IPv4 BOOTP/DHCP Relay, DHCPv6 server sends its replies to the nearest relay to the server, as opposed to the nearest relay to the client.

To configure a DHCPv6 Security Policy:

1. Go to the **main SmartConsole Menu** () **Global Properties > Firewall**. If the **Accept outgoing packets originating from gateway implied rule** is enabled, then from the drop-down menu, select **Last** or **Before Last**.
2. Create a new host for the DHCP server. In the SmartConsole main view, go to **Objects > New Host**.
The **New Host** window opens.
 - a) Enter the server name.
 - b) In the **Ipv6 address** field, enter the server's address.
 - c) Click **OK**.
3. Create a new client network. In the SmartConsole main view, go to **Objects > New Network**.
The **New Network** window opens.
 - a) Enter the object name.
 - b) Enter the Network address Prefix for the network to which the DHCP clients are connected.
 - c) Click **OK**.
4. Configure the required Security Policy rules with the DHCPv6 services (*dhcpv6-request*, *dhcpv6-reply* and *dhcpv6-relay*).

Note - Use:

- The DHCPv6-Relay object which you configured on the Security Gateway.
- All_DHCPv6_Relay_Agents_and_Servers which is pre-defined in SmartConsole.
- IPv6_Link_Local_Hosts which is pre-defined in SmartConsole.

For example:

Source IP	Destination IP	Service	Notes
Ipv6_Link_Local_Hosts	All_DHCPv6_Relay_Agents_and_Servers	dhcpv6-request	Allows requests from DHCPv6 clients to a DHCPv6 relay (on a Gaia Security Gateway) which is directly connected to the client network.
<DHCPv6_Relay> IPv6_Link_Local_Hosts	IPv6_Link_Local_Hosts	dhcpv6-reply	Allows replies from a DHCPv6 relay to local DHCPv6 clients.
<DHCPv6_Relay>	DHCPv6_Server	dhcpv6-relay	Allows traffic between DHCPv6 relays and DHCPv6 servers.
DHCPv6_Server	<DHCPv6_Relay>	dhcpv6-reply	With the implied rules in their default settings (Before Last), replies from the DHCPv6 Server to the DHCPv6 relay are automatically accepted when the reply matches a request from the relay to the server. However, to make sure that such replies are accepted, we recommend to make an explicit rule to allow traffic from the DHCPv6 Server to the DHCPv6 Relay.

Source IP	Destination IP	Service	Notes
Client_Network	DHCPv6_Server	dhcpv6-request	When DHCPv6 relay is used, the DHCPv6 client can still send requests directly to the DHCPv6 server.
DHCPv6_Server	Client_Network	dhcpv6-reply	Accepts replies from DHCPv6 servers to DHCPv6 clients. When DHCPv6 relay is used, the DHCPv6 client can still receive replies directly from a remotely located DHCPv6 server through UDP unicasts.

5. **Install Policy** on the related Security Gateways.

BGP

In This Section:

Configuring BGP 4-Byte AS	21
Support for IPv6 BGP (BGP-4 Multiprotocol Extensions)	22
BGP Sessions (Internal and External)	22
BGP Path Attributes	25
BGP Multi-Exit Discriminator	26
BGP Interactions with IGPs	26
Inbound BGP Route Filters	26
Redistributing Routes to BGP	27
BGP Communities	27
BGP Route Reflection	27
BGP Confederations	28
External BGP (EBGP) Multihop Support	29
BGP Route Dampening	29
TCP MD5 Authentication	30
Configuring BGP - WebUI	30
Configuring BGP - CLI (bgp)	43

Border Gateway Protocol (BGP) is an inter-AS protocol, meaning that it can be deployed within and between autonomous systems (AS). An autonomous system is a set of routers under a single technical administration. An AS uses an interior gateway protocol and common metrics to route packets within an AS; it uses an exterior routing protocol to route packets to other ASes.



Note - This implementation supports BGP version 4, with Multiprotocol Extensions.

BGP sends update messages that consist of network number-AS path pairs. The AS path contains the string of ASes through which the specified network can be reached. An AS path has some structure in order to represent the results of aggregating dissimilar routes. These update messages are sent over TCP transport mechanism to ensure reliable delivery. BGP contrasts with IGPs, which build their own reliability on top of a datagram service.

As a path-vector routing protocol, BGP limits the distribution of router reachability information to its peer or neighbor routers.

You can run BGP over a route-based VPN by enabling BGP on a virtual tunnel interface (VTI).

Configuring BGP 4-Byte AS

BGP 4-Byte AS lets you configure 32-bit AS numbers. This feature is enabled by default and it is not possible to turn it off.

To configure as AS number:

On the gateway, run:

```
set as <asnum>
```

Possible values: {off | 1-4294967295 | 0.1-65535.65535 (dotted format)}

Example: GW-1> set as 1.14464

To monitor BGP 4-Byte AS:

On the gateway, run:

show as

Output example: Autonomous System Number 1.14464 (80000)

Support for IPv6 BGP (BGP-4 Multiprotocol Extensions)

Gaia implements BGP-4 with support for multiprotocol extensions and the exchange of IPv6 address prefixes, as described in RFCs 2545, 2858, and 3392.

You must use an IPv4 address for the router ID (BGP identifier). After the BGP session is up, prefixes can be advertised and withdrawn by sending normal UPDATE messages that include either or both of the new multiprotocol attributes MP_REACH_NLRI (used to advertise reachability of routes) and MP_UNREACH_NLRI (used to withdraw routes).

The new attributes are backward compatible. If two routers have a BGP session and only one supports the multiprotocol attributes, they can still exchange unicast IPv4 routes even though they cannot exchange IPv6 routes.

BGP IPv6 peers are not supported on ClusterXL clusters (VSX or non-VSX).

Configuring IPv6 BGP (BGP-4 Multiprotocol Extensions)

On each peer, configure the type of routes (Multiprotocol capability) to exchange between peers. Select one of these:

- IPv4 Unicast Only (the default)
- IPv6 Unicast Only
- Both IPv4 and IPv6

To establish peering, the routers must share a capability.

If your system is exchanging IPv4 routes over IPv6 or vice versa, use route map CLI commands to set nexthop to match the family of the routes being exchanged. If they do not match, the routes will not be active.

Note - To configure routing policies for BGP-4 Multiprotocol Extensions, use route map commands in the CLI. Do not use the route redistribution and inbound filters in the WebUI.

BGP Sessions (Internal and External)

BGP supports two basic types of sessions between neighbors: internal (sometimes referred to as IBGP) and external (EBGP). Internal sessions run between routers in the *same* autonomous systems, while external sessions run between routers in *different* autonomous systems. When sending routes to an external peer, the local AS number is prepended to the AS path. Routes received from an internal neighbor have, in general, the same AS path that the route had when the originating internal neighbor received the route from an external peer.

BGP sessions might include a single metric (Multi-Exit Discriminator or MED) in the path attributes. Smaller values of the metric are preferred. These values are used to break ties between routes with equal preference from the same neighbor AS.

Internal BGP sessions carry at least one metric in the path attributes that BGP calls the local preference. The size of the metric is identical to the MED. Use of these metrics is dependent on the type of internal protocol processing.

BGP implementations expect external peers to be directly attached to a shared subnet and expect those peers to advertise next hops that are host addresses on that subnet. This constraint is relaxed when the multihop option is enabled in the BGP peer template during configuration.

Type internal groups determine the immediate next hops for routes. They do this by using the next hop received with a route from a peer as a forwarding address, and use this to look up an immediate next hop in IGP routes. Type internal groups support distant peers, but they need to be informed of the IGP whose routes they are using to determine immediate next hops.

Where possible, for internal BGP group types, a single outgoing message is built for all group peers based on the common policy. A copy of the message is sent to every peer in the group, with appropriate adjustments to the next hop field to each peer. This minimizes the computational load of running large numbers of peers in these types of groups.

Preventing Private AS Numbers from Propagating

An ISP can assign private AS numbers (64512 to 65535) to a customer in order to conserve globally unique AS numbers. When an ISP does so, a BGP update from a customer network to the ISP has the private AS number in its AS_PATH attribute. When the ISP propagates its network information to other ISPs, the private AS number would normally be included. To avoid this, you can configure Gaia to remove the private AS number from BGP update messages to external peers.

To configure Gaia to remove private AS numbers from BGP updates, enable the Remove Private AS option on the configuration page for an external peer.

If you enable this option, private AS numbers are removed from BGP updates according to the following rules:

- If the AS_PATH includes both public and private AS numbers, the private AS numbers are not removed.
- If the AS_PATH contains the AS number of the destination peer, private AS numbers are not removed.
- If the AS_PATH includes confederations and all the AS numbers in the AS_PATH are private, all the private AS numbers are removed.

Peer Local AS

The Peer Local AS feature lets you configure the connection to a remote peer with a Peer Local ASN, on a per-peer basis. The Peer Local ASN replaces the Local ASN in the BGP session. Only eBGP peers are supported. It is not necessary to configure the Peer Local ASN locally.

Inbound-Peer-Local

The router adds the configured Peer Local ASN to the AS path of the routes received from the peer. Routes installed from that peer will contain the Peer Local ASN as the first entry in the AS Path.

Outbound-Local

The router adds the Local ASN to the AS Path of the routes advertised to an eBGP peer. When enabled, the Local ASN is the second ASN in the AS Path of updates sent to eBGP peers. The Peer Local ASN is always the first ASN in the AS Path if this sub-feature is enabled or not.

Dual Peering

This option enables the connection to the Local ASN or the Peer Local ASN. There can be only one active connection. If you do not enable this option, it is only possible to connect to the Peer Local ASN.

The router first tries to connect to the local ASN. If the connection is created with the Local ASN, the BGP runs as if the peer Local ASN feature is not configured. If the connection with the Local ASN fails, the router tries to connect with the Peer Local ASN

Important - Do not use this feature with an AS that already has Peer Local AS with Dual-Peering enabled. This is because the two ASes then send AS numbers in OPEN messages in a manner that does not converge.

AS Override

To prevent loops in BGP, the router examines the AS number in the AS Path. If the receiving router sees its own AS number in the AS Path of the received BGP packet, it drops the packet. The receiving router thinks that the packet originated from its own AS and returned to where it started from.

With AS Override, the router overrides the peer's AS number with the router's AS number in the outbound AS path. This helps multiple sites in the same AS accept the routes. If the Peer Local AS feature is enabled, the router uses the configured Peer Local AS to override the remote peer's AS number.

BGP and ECMP

Equal-cost multi-path routing (ECMP) is a load-balancing routing strategy.

The BGP RFC does not support ECMP routes because BGP clearly sets the route selection criteria. To overcome this issue and install ECMP route through a BGP user, enable the ECMP option.

Note - BGP ECMP is not supported for routes that are received by a mix of IBGP and EBGP.

BGP Route Refresh

Gaia supports the ability to dynamically request BGP route updates from peers and to respond to requests for BGP route updates. For example, if you change the inbound routing policy, you can request that a peer readvertise its previously advertised routes so that the routes can be checked against the new policy. This feature is often referred to as a soft reset because it provides the ability to refresh routes received from a peer without tearing down the established session.

To configure BGP route updates in the:

- CLI - Run these commands:
`set bgp external remote-as as_number peer ip_address send-route-refresh`
`set bgp internal peer ip_address send-route-refresh`
- WebUI - Click the appropriate buttons in the **Edit Peer** page, in the section **Advanced Settings > Route Refresh**.

These options work only with peers that support the same capabilities. Gaia systems can also peer with systems that do not support these options.

BGP Path Attributes

A path attribute is a list of AS numbers that a route has traversed to reach a destination. BGP uses path attributes to provide more information about each route and to help prevent routing loops in an arbitrary topology. You can also use path attributes to determine administrative preferences.

BGP collapses routes with similar path attributes into a single update for advertisement. Routes that are received in a single update are readvertised in a single update. The churn caused by the loss of a neighbor is minimized, and the initial advertisement sent during peer establishment is maximally compressed.

BGP does not read information that the kernel forms message by message. Instead, it fills the input buffer. BGP processes all complete messages in the buffer before reading again. BGP also performs multiple reads to clear all incoming data queued on the socket.



Note - This feature might cause a busy peer connection to block other protocols for prolonged intervals.

The following table displays the path attributes and their definitions

Attribute	Description
AS_PATH	Identifies the autonomous systems through which routing information carried in an UPDATE message passed. Components of this list can be AS_SETs or AS_SEQUENCES.
NEXT_HOP	Defines the IP address of the border router that should be used as the next hop to the destinations listed in the UPDATE message.
MULTI_EXIT_DISC	Discriminates among multiple exit or entry points to the same neighboring autonomous system. Used only on external links.
LOCAL_PREF	Determines which external route should be taken and is included in all IBGP UPDATE messages. The assigned BGP speaker sends this message to BGP speakers within its own autonomous system but not to neighboring autonomous systems. Higher values of a LOCAL_PREF are preferred.
ATOMIC_AGGREGATE	Specifies to a BGP speaker that a less specific route was chosen over a more specific route. The BGP speaker attaches the ATOMIC_AGGREGATE attribute to the route when it reproduces it to other BGP speakers. The BGP speaker that receives this route cannot remove the ATOMIC_AGGREGATE attribute or make any Network Layer Reachability Information (NLRI) of the route more specific. This attribute is used only for debugging purposes.

All unreachable messages are collected into a single message and are sent before reachable routes during a flash update. For these unreachable announcements, the next hop is set to the local address on the connection, no metric is sent, and the path origin is set to incomplete. On external connections, the AS path in unreachable announcements is set to the local AS. On internal connections, the AS path length is set to zero.

Routing information shared between peers in BGP has two formats: announcements and withdrawals. A route announcement indicates that a router either learned of a new network attachment or made a policy decision to prefer another route to a network destination. Route

withdrawals are sent when a router makes a new local decision that a network is no longer reachable.

BGP Multi-Exit Discriminator

Multi-exit Discriminator (MED) values are used to help external neighbors decide which of the available entry points into an AS are preferred. A lower MED value is preferred over a higher MED value and breaks the tie between two or more preferred paths.



Note - A BGP session does not accept MEDs from an external peer unless the Accept MED field is set for an external peer.

BGP Interactions with IGPs

All transit ASes must be able to carry traffic that originates from locations outside of that AS, is destined to locations outside of that AS, or both. This requires a certain degree of interaction and coordination between BGP and the Interior Gateway Protocol (IGP) that the particular AS uses. In general, traffic that originates outside of a given AS passes through both interior Gateways (Gateways that support the IGP only) and border Gateways (Gateways that support both the IGP and BGP). All interior Gateways receive information about external routes from one or more of the border Gateways of the AS that uses the IGP.

Depending on the mechanism used to propagate BGP information within a given AS, take special care to ensure consistency between BGP and the IGP, since changes in state are likely to propagate at different rates across the AS. A time window might occur between the moment when some border gateway (A) receives new BGP routing information (which was originated from another border gateway (B) within the same AS) and the moment the IGP within this AS can route transit traffic to the border gateway (B). During that time window, either incorrect routing or black holes can occur.

To minimize such routing problems, border gateway (A) should not advertise to any of its external peers a route to some set of exterior destinations associated with a given address prefix using border gateway (B) until all the interior Gateways within the AS are ready to route traffic destined to these destinations by using the correct exit border gateway (B). Interior routing should converge on the proper exit gateway before advertising routes that use the exit gateway to external peers.

If all routers in an AS are BGP speakers, no interaction is necessary between BGP and an IGP. In such cases, all routers in the AS already have full knowledge of all BGP routes. The IGP is then only used for routing within the AS, and no BGP routes are imported into the IGP. The user can perform a recursive lookup in the routing table. The first lookup uses a BGP route to establish the exit router, while the second lookup determines the IGP path to the exit router.

Inbound BGP Route Filters

BGP routes can be filtered, or redistributed by AS number or AS path regular expression, or both.

BGP stores rejected routes in the routing table with a negative preference. A negative preference prevents a route from becoming active and prevents it from being installed in the forwarding table or being redistributed to other protocols. This behavior eliminates the need to break and re-establish a session upon reconfiguration if importation policy is changed.

When you import from BGP you can add or modify the local preference, rank and nexthop. The local preference parameter assigns a BGP local preference to the imported route. The local preference is a 32-bit unsigned value, with larger values preferred. This is the preferred way to bias a routing subsystem preference for BGP routes.

Redistributing Routes to BGP

Redistributing to BGP is controlled by an AS. The same policy is applied to all firewalls in the AS. BGP metrics are 16-bit, unsigned quantities; that is, they range from 0 to 65535 inclusive, with zero being the most attractive. While BGP version 4 supports 32-bit unsigned quantities, routed does not.



Note - To define a redistribution policy, use **Advanced Routing > Route Distribution** in the WebUI, or routemaps in the CLI.

BGP Communities

BGP communities allow you to group a set of IP addresses and apply routing decisions based on the identity of the group or community.

To implement this feature, map a set of communities to certain BGP local preference values. Then you can apply a uniform BGP configuration to the community as a whole as opposed to each router within the community. The routers in the community can capture routes that match their community values.

Use community attributes to can configure your BGP speaker to set, append, or modify the community of a route that controls which routing information is accepted, preferred, or distributed to other neighbors. The following table displays some special community attributes that a BGP speaker can apply.

Community attribute	Description
NO_EXPORT (0xFFFFF01)	Not advertised outside a BGP confederation boundary. A stand-alone autonomous system that is not part of a confederation should be considered a confederation itself.
NO_ADVERTISE (0xFFFFF02)	Not advertised to other BGP peers.
NO_EXPORT_SUBCONFED(0xFFFFF03)	Not advertised to external BGP peers. This includes peers in other members' autonomous systems inside a BGP confederation.

For more about communities, see RFCs 1997 and 1998.

BGP Route Reflection

By default, all BGP peers in an Autonomous System (AS) are in a full mesh. However, if an AS has many BGP peers, the BGP configuration and hardware deployment is not easy.


To simplify configuration and deployment and avoid having to connect the peers in a full mesh, it is possible to configure:

- One BGP peer as a route reflector.

- All or some of the other BGP peers as clients of the route reflector.

The route reflector and its clients are known as a route reflection cluster. The route reflector sends the routes received from its peers to its clients.

In the example network, AS1 has five BGP-enabled Check Point routers. One of the routers is a route reflector for two clients.



Number	Description
1	Non-clients
2	IBGP
3	Route Reflector in cluster
4	Clients in cluster

Number	Description
5	AS1
6	EBGP
7	AS676

It is possible to define more than one route reflector in the AS to avoid having a single point of failure. **Best Practice** - We recommend that you not use multiple redundant reflectors unnecessarily because it increases the memory required to keep routes on the peers of redundant reflectors.

To learn more about route reflection, see RFC 2796.

BGP Confederations

An alternative to route reflection is BGP confederations. As with route reflectors, you can partition BGP speakers into clusters where each cluster is typically a topologically close set of routers. With confederations, this is accomplished by subdividing the autonomous system into multiple, smaller ASes that communicate among themselves. The internal topology is hidden from the outside world, which perceives the confederation to be one large AS.

Each distinct sub-AS within a confederation is referred to as a routing domain (RD). Routing domains are identified by using a routing domain identifier (RDI). The RDI has the same syntax as an AS number, but as it is not visible outside of the confederation, it does not need to be globally unique, although it does need to be unique within the confederation. Many confederations find it convenient to select their RDIs from the reserved AS space (ASes 64512 through 65535 [see RFC 1930]). RDIs are used as the ASes in BGP sessions between peers within the confederation.

The confederation as a whole, is referred to by a confederation identifier. This identifier is used as the AS in external BGP sessions. As far as the outside world is concerned, the confederation ID is the AS number of the single, large AS. For this reason, the confederation ID must be a globally unique, normally assigned AS number.



Note - Do not nest confederations.

For further details, refer to the confederations specification document (RFC 1965).

External BGP (EBGP) Multihop Support

Connections between BGP speakers of different ASes are referred to as External BGP (EBGP) connections. BGP enforces the rule that peer routers for EBGP connections need to be on a directly attached network. If the peer routers are multiple hops away from each other or if multiple links are between them, you can override this restriction by enabling the EBGP multihop feature. TCP connections between EBGP peers are tied to the addresses of the outgoing interfaces. Therefore, a single interface failure severs the session even if a viable path exists between the peers.

EBGP multihop support can provide redundancy so that an EBGP peer session persists even in the event of an interface failure. Using an address assigned to the loopback interface for the EBGP peering session ensures that the TCP connection stays up even if one of the links between them is down, provided the peer loopback address is reachable. In addition, you can use EBGP multihop support to balance the traffic among all links.

Use the TTL (Time to Live) parameter to limit the number of hops over which the External BGP (EBGP) multihop session is established. You can configure the TTL only if EBGP multihop is enabled. The default TTL is 64. When multihop is disabled the default TTL is 1.

When traffic comes from a router that is not directly connected and multihop is enabled, BGP uses that router as the next hop, irrespective of the advertised routes that it gets.



Important - Enabling multihop BGP connections is dangerous because BGP speakers might establish a BGP connection through a third-party AS. This can violate policy considerations and introduce forwarding loops.

BGP Route Dampening

Route dampening decreases the propagation of flapping routes. A flapping route is a route that repeatedly becomes available and then unavailable. Without route dampening, autonomous systems continually send advertisement and withdrawal messages each time the flapping route becomes available or unavailable. As the Internet grew, the number of announcements per second grew as well and caused performance problems within the routers.

Route dampening enables routers to keep a history of the flapping routes and prevent them from consuming significant network bandwidth. The routers measure how often a given route becomes available and then unavailable. When a route reaches a set threshold, that route is no longer considered valid, and is no longer propagated for a given period of time, usually about 30 minutes. If a route continues to flap even after it reaches the threshold, the time out period for that route grows in proportion to each additional flap. Once the route reaches the threshold, the route is dampened or suppressed. Suppressed routes are added back into the routing table once the penalty value decreases and falls below the reuse threshold.

Route dampening can cause connectivity to look lost to the outside world but maintained on your own network because route dampening only applies to BGP routes. Because of high load on the backbone network routers, most NSPs (MCI, Sprint, UUNet etc.) have set up route suppression.



Note - BGP route dampening is supported only for EBGP. It is not supported for IBGP.

TCP MD5 Authentication

The Internet is vulnerable to attack through its routing protocols and BGP is no exception. External sources can disrupt communications between BGP peers by breaking their TCP connection with spoofed RST packets. Internal sources, such as BGP speakers, can inject bogus routing information from any other legitimate BGP speaker. Bogus information from either external or internal sources can affect routing behavior over a wide area in the Internet.

The TCP MD5 option allows BGP to protect itself against the introduction of spoofed TCP segments into the connection stream. To spoof a connection using MD5 signed sessions, the attacker not only has to guess TCP sequence numbers, but also the password included in the MD5 digest.

Note - TCP MD5 is not supported on BGP IPv6 peers.

Configuring BGP - WebUI

This section gives per-field help for the fields in the Advanced Routing > BGP section of the Gaia WebUI.



Note - Not all fields are shown in all cases.

To configure BGP:

1. In the **Network Management > Network interfaces** page of the WebUI, configure Ethernet Interfaces and assign an IP address to the interface.
2. Open the **Advanced Routing > BGP** page of the WebUI.
3. Define **BGP Global settings**, including the Router ID.
4. **Optional**: Configure **Peer Groups**.
5. **Optional**: Define **Miscellaneous Settings**.

BGP Global Settings

Parameter	Description
Router ID	The Router ID uniquely identifies the router in the autonomous system. The BGP and OSPF protocols use the router ID. Best Practice - set the router ID rather than rely on the default setting. This prevents changes in the router ID if the interface used for the router ID goes down. Use an address on a loopback interface that is not the loopback address (127.0.0.1). Note - In a cluster, you must select a router ID and make sure that it is the same on all cluster members. • Range: Dotted-quad.([0-255].[0-255].[0-255].[0-255]). Do not use 0.0.0.0 • Default: The interface address of one of the local interfaces.
Cluster ID for Route Reflectors	The cluster ID used for route reflection. The default cluster ID is the router ID. You must override this default value if the cluster contains more than one route reflector. Typically, a single router acts as the reflector for a set, or cluster, of clients. However, for redundancy two or more routers can also be configured as reflectors for the same cluster. In this case, you must select a cluster ID to identify all reflectors serving the cluster. Gratuitous use of multiple redundant reflectors is not advised, for this situation can cause an increase in the memory required to store routes on the redundant reflectors peers. • Range: Dotted-quad ([0-255].[0-255].[0-255].[0-255]) • Default: Router ID
Local Autonomous System Number	The local autonomous system number of the router.

Change Local System Identification

Parameter	Description
Unconfigured	
Local Autonomous System Number	The local autonomous system number of the router. This setting is mutually exclusive from the Confederation and Routing Domain Identifier. The router can be configured with either the autonomous system number or the member of confederation, not both. Caution: When you change the autonomous system number, all current peer sessions are reset and all BGP routes are deleted. • Range: 1-65535 • Default: No default

Parameter	Description
Confederation	The identifier for the entire confederation system. This identifier is used as the AS in external BGP sessions. To the outside world, the confederation ID is the AS number of the single, large AS. For this reason, the confederation ID must be a globally unique, normally assigned AS number. • Range: 1-65535 • Default: No default
Number of loops permitted in AS_PATH	For the confederation: The number of times the local autonomous system can appear in an AS path for BGP-learned routes. If the number of times the local autonomous system appears in an AS path is more than the number in this field, the corresponding routes are discarded or rejected. • Range: 1-10 • Default: 1
Routing Domain Identifier	The routing domain identifier (RDI) of this router. This value is required only if BGP confederations are in use. The RDI does not have to be globally unique since it is never used outside the domain of the confederation system. However, the configured RDI must be unique within the confederation. The routing-domain identifier and autonomous system number are mutually exclusive values; that is, the router can be configured with either the autonomous system number or the member of confederation, not both. If confederations are in use, the RDI is used wherever the autonomous system would be used to communicate with peers within the confederation, including group-type confederation peers and the various internal-type peers. For correct operation of the router in confederations you must configure both the routing-domain identifier and the confederation. • Range: 1-65535 • Default: No default
Number of loops permitted in AS_PATH	For the routing domain identifier: The number of times the local autonomous system can appear in an AS path for BGP-learned routes. If the number of times the local autonomous system appears in an AS path is more than the number in this field, the corresponding routes are discarded or rejected. • Range: 1-10 • Default: 1

BGP Miscellaneous Settings

Parameter	Description
Default MED	Defines the metric (MED) used when advertising routes through BGP. If you do not specify a value, no metric is propagated. A metric specified on the neighbor configuration or in the redistribution configuration might override the metric you configure. • Range: 0-65535
Default Gateway:	A default route is generated when any BGP peer is up. This route has a higher rank than the default configured in the static routing page. If a specific BGP peer should not be considered for generating the default route, you should explicitly suppress the option in the peer-specific configuration. • Range: Dotted-quad ([0-255].[0-255].[0-255].[0-255]). • No Default.
Enable IGP Synchronization	Select this option to make internal and configured BGP peers check for a matching route from IGP protocols before installing a given route. • Default: Unselected
Enable communities	Enables communities-based policy options. • Default: Unselected
Enable ECMP	Enables Equal-Cost Multi-Path routing strategy. • Default: Off

Weighted Route Dampening Settings

Parameter	Description
Enable Weighted Route Dampening	Weighted route dampening minimizes the propagation of flapping routes across an internetwork. A route is considered to be flapping when it is repeatedly transitioning from available to unavailable or vice versa. Only routes learned through BGP are subjected to weighted route dampening. Note: BGP route dampening is only supported for External BGP (EBGP). When this option is selected, the other Route Dampening fields show.
Reuse-below metric	The value of the instability metric at which a suppressed route becomes unsuppressed if it is reachable but currently suppressed. The value assigned to the reuse-below metric must be less than the suppress-above value. • Range: 1-32 • Default: 2

Suppress-above metric	The value of the instability metric at which a route is suppressed; a route is not installed in the FIB or announced even if it is reachable during the period that it is suppressed. • Range: 2-32 • Default: 3
Max-flap metric	The upper limit of the instability. The value must be higher than one plus the suppress-above value. The metric assigned to the suppress-above, reuse-below, and max-flap metric values is a floating point number, in units of flaps. Each time a route becomes unreachable, one is added to the current instability metric. • Range: 3-64 • Default: 16
Reachable decay time	A value that determines the length of time it takes for the instability metric value to reach one half of its current value when the route is reachable. This half-life value determines the rate at which the metric value is decayed. A smaller half-life value makes a suppressed route reusable sooner than a larger value. • Range: 1-900 • Default: 300
Unreachable decay time	The rate at which the instability metric is decayed when a route is unreachable. This value must be equal to or greater than the reach-decay value. • Range: 1-2700 • Default: 900
Keep history time	The period over which the route flapping history is maintained for a given route. The size of the configuration arrays described below is directly affected by this value. • Range: 2-5400 • Default: 1800

BGP AS Peer Group Settings

Parameter	Description
Peer AS Number	The autonomous system number of the external peer group. Enter an integer from 1-65535.
Peer Group Type	One of • Unconfigured • Local Autonomous System Number • Confederation
Description	A free-text description of the peer group.

Parameter	Description
Local address	The address used on the local end of the TCP connection with the peer. For external peers that do not have multihop enabled, the local address must be on an interface that is shared with the peer or with the peer's gateway, when the gateway parameter is used. A session with an external peer opens only when an interface with a local address through which you can reach the peer or gateway address directly operates. For other types of peers, a peer session opens when an interface with the specified local address operates. In both external and other types of peers, incoming connections are recognized as matching a configured peer only if they are addressed to the configured local address. Note - If you run BGP in a cluster, you must not configure the local address. • Default: None
Out Delay	The length of time in seconds that a route must be present in the routing database before it is redistributed to BGP. This value applies to all neighbors configured in this group. The default value is zero, which means that this feature is disabled. This feature dampens route fluctuations. • Range: 0-65535 • Default: 0
Peer	Configure peers. Each peer inherits as defaults all parameters configured on a group. To change the values of a peer's parameters, select the peer and click Edit.

BGP Remote Peers

Parameter	Description
Peer	
Comment	A free-text description of the remote peer.

Advanced Settings

Multiprotocol Capabilities

Parameter	Description
IPv4 Unicast	Specifies if IPv4 unicast routes can be sent to and received from this peer. • Default: Selected
IPv6 Unicast	Specifies if IPv6 unicast routes can be sent to and received from this peer. • Default: Cleared

Peer Local AS

Parameter	Description
Enable Peer Local AS	The peer local ASN replaces the local ASN in the BGP session.
Peer Local AS	Peer local AS number.
Prepend Peer Local AS on inbound updates from peer	The router adds the configured peer local ASN to the AS path of the routes received from the peer. Routes installed from that peer will contain the peer local ASN as the first entry in the AS Path. Default: On
Prepend systemwide Local AS on outbound updates to peer	The router adds the local ASN to the AS Path of the routes advertised to an eBGP peer. When enabled, the local ASN is the second ASN in the AS Path of updates sent to eBGP peers. The peer local ASN is always the first ASN in the AS Path. Default: On
Allow peering with the Local AS	Enables the connection to the local ASN or the peer local ASN. There can be only one active connection. If you do not enable this option, it is only possible to connect to the Peer Local AS. The router first tries to connect to the local ASN. If the connection is created with the local ASN, the BGP runs as if the peer local ASN feature is not configured. If the connection with the local ASN fails, the router tries to connect with the peer local ASN. Important - Do not use this feature with an AS that already has peer local AS with Dual-Peering enabled. Default: Off

Local Address

Parameter	Description
Local Address	The address used on the local end of the TCP connection with the peer. For external peers that do not have multihop enabled, the local address must be on an interface that is shared with the peer or with the peer's gateway when the gateway parameter is used. A session with an external peer is opened only when an interface with a local address through which the peer or gateway address is directly reachable is operating. For other types of peers, a peer session is maintained when any interface with the specified local address is operating. In either case, incoming connections are recognized as matching a configured peer only if they are addressed to the configured local address. Note - If running BGP in a cluster you must not configure the local address. Default: None

Weight

Parameter	Description
Weight	The default weight associated with each route accepted from this peer. This value can be overridden by the weight specified in the import policy. • Range: 0-65535

MED

Parameter	Description
Accept MED from External Peer	MED should be accepted from this external neighbor. MEDs are always accepted from routing-type and confederation neighbors. If this parameter is not used with an external neighbor, the MED is stripped before the update is added to the routing table. If this parameter is added or deleted and routed is reconfigured, the affected peering sessions are automatically restarted. • Default: Cleared
MED Sent Out	The primary metric used on all routes sent to the specified peer. This metric overrides the default metric on any route specified by the redistribute policy. • Range: 0-4294967294 • Default: 4294967294

Next Hop and Time to Live

Parameter	Description
EGP Multihop	Multihop is used to set up EBGp peering connections with peers that are not directly connected. You can also use this option, which relies on an IGP to find the route to the peer, to set up peers to perform EBGp load balancing. You can refine the multihop session by configuring the TTL, that is, the number of hops to the EBGp peer. The TTL has a default value of 64. • Default: Cleared
Time to Live	You can use the TTL (time to live parameter) to limit the number of hops over which the EBGp multihop session is established. You can configure the TTL only if multihop is enabled. • Range: 1-255 • Default: 64

Aggregator

Parameter	Description
No Aggregator ID	Select to force this router to specify the router ID in the aggregator attribute as zero, rather than the actual router ID. This option prevents different routers in an AS from creating aggregate routes with different AS paths. • Default: Cleared

ASPATH

Parameter	Description
ASPATH prepend count	The number of times this router adds to the AS path on EBGP external or CBGP confederation sessions. Use this setting to bias the degree of preference some downstream routers have for the routes originated by this router. Some implementations prefer to select routes with shorter AS paths. This parameter has no effect when used with IBGP peers. • Range: 1-25 • Default: 1
AS Override	Overrides the peer's AS number with the router's AS number in the outbound AS path. Default: Cleared

Private AS

Parameter	Description
Remove Private AS	Remove private AS numbers from the outgoing updates to this peer. Following conditions apply when this feature is enabled: • If the AS path includes both public and private AS numbers, private AS numbers will not be removed. • If the AS path contains the AS number of the destination peer, private AS numbers will not be removed. • If the AS path contains only confederations and private AS numbers, private AS numbers will be removed. • Default: Cleared

Timers

Parameter	Description
Keep Alive Timer	An alternative way to specify a Hold Time value, in seconds, to use when negotiating the connection with this peer. The keepalive interval equals one-third the value of the holdtime. The keepalive interval is often used instead of the holdtime value, but you can specify both values, provided the value for the holdtime is three times the keepalive interval. The value must be 0, that is, no keepalives are sent, or at least 2. • Range: 0, 2-21845 • Default: 60
Hold Time	The BGP holdtime value, in seconds, to use when negotiating a connection with this peer. According to the specification, if the BGP speaker does not receive a keepalive update or notification message from its peer within the period specified by the holdtime value in the BGP Open message, the BGP connection is closed. The value must be either 0, that is, no keepalives are sent, or at least 6. • Range: 0, 6-65535 • Default: 180

Needed when Peering with Route Server

Parameter	Description
Ignore First AS Hop	Select to force this router to ignore the first AS number in the AS_PATH for routes learned from the corresponding peer. Select this option only if you are peering with a route server in so-called transparent mode, that is, when the route server is configured to redistribute routes from multiple ASes without prepending its own AS number. • Default: Cleared

Keep Alive

Parameter	Description
Keep Alive Always	Select to force this router always to send keepalives even when an update can substitute. This setting allows interoperability with routers that do not completely adhere to the protocol specifications on this point. • Default: Cleared

Routes

Parameter	Description
Accept Routes Received From the Peer	Routes received from peer routes are accepted if there is an inbound BGP route policy. If an inbound policy to accept the route does not exist, you can select All or None. • All - Specifies to accept and install routes with an invalid preference. Depending on the local BGP inbound policy the routes could become active or inactive. • None - Specifies to delete routes learned from a peer when no explicit local BGP inbound policy exists. This option is used to save memory overhead when many routes are rejected because there is no local policy. These routes can be relearned only by restarting the BGP session. • Default: All

Allows Accept TCP Sessions from Your Peer

Parameter	Description
Passive	Select to force this router to wait for the peer to issue an open. By default all explicitly configured peers are active and periodically send open messages until the peer responds. Modifying this option will reset the peer connection. • Default: Cleared

Authentication

Parameter	Description
Authentication type	The type of authentication scheme to use between given peers. In general peers must agree on the authentication configuration to form peer adjacencies. This feature guarantees that routing information is accepted only from trusted peers. If the Auth type selected is MD5 the Password field appears. When you enter a password, MD5 authentication is used with the given peer. Note - TCP MD5 authentication for BGP is not supported on 64-bit Gaia or on a 32/64-bit Gaia VSX cluster. • Options: None / MD5 • Default: None

Limit BGP Updates Send to a Peer

Parameter	Description
Throttle count	Throttles the network traffic when there are many BGP peers. Throttle count determines the number of BGP updates sent at a time. • Range: 0-65535 • Default: No default

Route Refresh

Parameter	Description
Route Refresh	Route refresh is used to either re-learn routes from the BGP peer or to refresh the routing table of the peer without tearing down the BGP session. Both peers must support the BGP route refresh capability and should have advertised this at the time peering was established. Re-learning of routes previously sent by the peer is accomplished by sending a BGP route refresh message. The peer responds to the message with the current routing table. Similarly, if a peer sends a route refresh request the current routing table is re-sent. You can also trigger a route update without having to wait for a route refresh request from the peer. Both peers must support the same address and subsequent address families. For example a request for IPv6 unicast routes from a peer that did not advertise the capability during session establishment will be ignored. Note - Clicking a Route Refresh button sends a trigger to the routing daemon. It does not change the configuration of the router.

Graceful Restart

Parameter	Description
Helper	Routes received from peer are preserved if the peer goes down till either the session is re-established (OPEN message is received from the peer after it comes back up) or the graceful restart timer expires. • Default: Cleared
Stalepath Time	Maximum time for which routes previously received from a restarting router are kept unless they are re-validated. The timer is started after the peer sends indication that it is up again. • Range: 60 - 65535 • Default: 360

Logging

Parameter	Description
Log bgp peer transitions	Select to force this router to log a message whenever a BGP peer enters or leaves the ESTABLISHED state. • Default: Cleared
Log warnings	Select to force this router to log a message whenever a warning scenario is encountered in the codepath. • Default: Cleared

Trace Options

Parameter	Description
The tracing options for BGP. The BGP implementation inherits the default values for global trace options. You can override these values on a group or neighbor basis. Log messages are saved in <code>var/log/routed.log</code>	
All	Trace all message types.
General	Trace messages related to Route and Normal.
Keepalive	Trace all the BGP keepalive messages to this peer, which are used to verify peer reachability.
Normal	Trace normal protocols occurrences. Abnormal protocol occurrences are always traced.
Open	Trace all the BGP open messages to this peer, which are used to establish a peer relationship.
Packets	Trace all the BGP packets to this peer.
Policy	Trace application of protocol and user-specified policy to routes being imported and exported.
Route	Trace routing table changes for routes installed by this peer.
State	Trace state machine transitions in the protocol.
Task	Trace system interface and processing associated with this peer.
Timer	Trace timer usage by this peer.
Update	Trace all the BGP update messages to this peer, which are used to pass network reachability information.

Configuring BGP - CLI (bgp)

External BGP Commands

Use the following commands to configure external sessions of the protocol, that is, between routers in different autonomous systems.

```
set bgp external remote-as as_number
    <on | off>
    aspath-prepend-count <1-25 | default>
    description text
    local-address ip_address <on | off>
    outdelay <0-65535>
    outdelay off
```

Parameter	Description
as_number <on off>	The autonomous system number of the external peer group. Enter an integer from 1-65535.
aspath-prepend-count <1-25> default>	The number of times this router adds to the autonomous system path on external BGP sessions. Use this option to bias the degree of preference some downstream routers have for the routes originated by this router. Some implementations prefer to select paths with shorter autonomous system paths. Default is 1.
description text	You can enter a brief text description of the group.
local-address ip_address <on off>	The address used on the local end of the TCP connection with the peer. For external peers that do not have multihop enabled, the local address must be on an interface that is shared with the peer or with the peer's gateway when the gateway parameter is used. A session with an external peer is opened only when an interface with a local address through which the peer or gateway address is directly reachable is operating. For other types of peers, a peer session is maintained when any interface with the specified local address is operating. In either case, incoming connections are recognized as matching a configured peer only if they are addressed to the configured local address. Note: If running BGP in a cluster you must not configure the local address. Default: Off
outdelay <0-65535>	The amount of time in seconds that a route must be present in the routing database before it is redistributed to BGP. The configured value applies to all peers configured in this group. This feature dampens route fluctuation. The value zero (0) disables this feature. Default: 0
outdelay off	Disables outdelay.

BGP Peer Commands

Use these commands to configure BGP peers.

Gaia supports IPv4 and IPv6 addresses for BGP peers.

BGP IPv6 peers are not supported on ClusterXL clusters (VSX or non-VSX).

```
set bgp external remote-as <as_number> peer      <ip_address>
    {on | off}
    med-out {<0-4294967294> | default}
    outgoing-interface <finterface> {on | off}
    accept-med {on | off}
    multihop {on | off}
    peer-local-as as {<1-4294967295> | <0.1-65535.65535> | off}
    peer-local-as {dual peering | inbound-peer-local | outbound-local} {on | off}
    as-override {on | off}
    ttl {1-255 | default}
    no-aggregator-id {on | off}
    holdtime {<6-65535> | default}
    keepalive {<2-21845> | default}
    ignore-first-ashop {on | off}
    send-keepalives {on | off}
    send-route-refresh {request|route-update}{ipv4 | ipv6 | All} [unicast]
    route-refresh {on | off}
    accept-routes {all | none}
    passive-tcp {on | off}
    removeprivateas {on | off}
    authtype none
    authtype md5 secret secret
    throttle-count {<0-65535> | off}
    suppress-default-originate {on | off}
    log-state-transitions {on | off}
    log-warnings {on | off}
    trace bgp_traceoption {on | off}
    capability {default | ipv4-unicast | ipv6-unicast}
    graceful-restart-helper {on | off}
    graceful-restart-helper-stalepath-time seconds
```

Parameter	Description
<code>ip_address {on off}</code>	A specified peer <ip_address> for the group.
<code>med-out {<0-4294967294> default}</code>	The multi-exit discriminator (MED) metric used as the primary metric on all routes sent to the specified peer address. This metric overrides the default metric on a metric specified by the redistribute policy. External peers use MED values to know which of the available entry points into an autonomous system is preferred. A lower MED value is preferred over a higher MED value. Default: 4294967294
<code>outgoing-interface <interface> {on off}</code>	IPv6 peer with FE80: local address only: All peer interfaces have a local address and a global address. All the peer interfaces can have the same local address, which starts with FE80:. To use the local address, you must enter the outgoing interface for the local address.

Parameter	Description
<code>accept-med {on off}</code>	Accept MED from the specified peer address. If you do not set this option, the MED is stripped from the advertisement before the update is added to the routing table.
<code>multihop {on off}</code>	Enable multihop connections with external BGP (EBGP) peers that are not directly connected. By default, external BGP peers are expected to be directly connected. You can configure the multihop session in the Time to Live (TTL) parameter, that is, the number of hops to the EBGP peer. This option can also be used to set up peers for EBGP load balancing. Default: Off
<code>peer-local-as as {<1-4294967295> <0.1-65535.65535> off}</code>	Configures the connection to a remote peer with a Peer Local ASN, on a per-peer basis. The Peer Local ASN replaces the Local ASN in the BGP session.
<code>peer-local-as {inbound-peer-local outbound-local dual-peering} {on off}</code>	Default for <code>inbound-peer-local</code> : On Default for <code>outbound-local</code> : On Default for <code>dual-peering</code> : Off
<code>as-override</code>	Overrides the peer's AS number with the router's AS number in the outbound AS path. Default: Off
<code>ttl {<1-255> default}</code>	Use the TTL (Time to Live) parameter to limit the number of hops over which the External BGP (EBGP) multihop session is created. You can configure the TTL only if EBGP multihop is enabled. The default TTL is 64. When multihop is disabled the default TTL is 1. Default: 64
<code>no-aggregator-id {on off}</code>	The router's aggregate attribute as zero (rather than the router ID value). This option prevents the creation of aggregate routes with different AS paths by different routers in an AS.
<code>holdtime {<6-65535> default}</code>	The BGP holdtime interval, in seconds, during the negotiation of a connection with the specified peer. If the BGP speaker does not receive a keepalive update or notification message from its peer within the period specified in the holdtime field of the BGP open message, the BGP connection is closed. Default: 180 seconds
<code>keepalive {<2-21945> default}</code>	The keepalive option is an alternative way to enter a holdtime value in seconds during the negotiation of a connection with the specified peer. You can use the keepalive interval instead of the holdtime interval. You can also use both intervals, but the holdtime value must be 3 times the keepalive interval value. Default: 60 seconds

Parameter	Description
<code>ignore-first-ashop</code> {on off}	Ignore the first AS number in the AS path for routes learned from the corresponding peer. Set this option only if you peer with a route server in transparent mode. In transparent mode, the route server redistributes routes from multiple other autonomous systems and does not prepend its own ASN.
<code>send-keepalives</code> {on off}	The router always sends keepalive messages even when an update message is sufficient. This option lets the router interoperate with other routers that do not strictly follow protocol specifications regarding updates.
<code>send-route-refresh</code> {request route-update}{ipv4 ipv6 All} unicast	The router dynamically requests BGP route updates from peers or responds to requests for BGP route updates.
<code>route-refresh</code> {on off}	Re-learns routes previously sent by the BGP peer or refreshes the routing table of the peer. The peer responds to the message with the current routing table. Similarly, if a peer sends a route refresh request the current routing table is re-sent. A user can also trigger a route update and not wait for a route refresh request from the peer.
<code>accept-routes</code> {all none}	An inbound BGP policy route if one is not already configured. Enter <code>all</code> to set accepting routes and install them with an invalid preference. Depending on the local inbound route policy, these routes are then made active or inactive. Enter <code>none</code> to delete routes learned from a peer. This option saves memory overhead when many routes are rejected because there is no inbound policy.
<code>passive-tcp</code> {on off}	The router waits for the specified peer to issue an open message. The router does not initiate tcp connections.
<code>removeprivateas</code> {on off}	Remove private AS numbers from BGP update messages to external peers.
<code>authtype none</code>	Do not use an authentication scheme between peers. If you use an authentication scheme, routing information is accepted only from trusted peers. Default: none
<code>authtype md5 secret</code> <code>secret</code>	Use md5 authentication between peers. In general, peers must agree on the authentication configuration to and from peer adjacencies. If you use an authentication scheme, routing information is accepted only from trusted peers. Note - TCP MD5 is not supported on BGP IPv6 peers.

Parameter	Description
<code>throttle-count {<0-65535> off}</code>	The number of BGP updates to send at one time. This option limits the number of BGP updates when there are many BGP peers. Off disables the throttle count option.
<code>suppress-default-originate {on off}</code>	Do NOT generate a default route when the peer receives a valid update from its peer.
<code>log-state-transitions {on off}</code>	The router generates a log message when a peer enters or leaves the established state.
<code>log-warnings {on off}</code>	The router generates a log message when there is a warning scenario in the codepath.
<code>trace bgp_traceoption {on off}</code>	Tracing options for the BGP implementation. Log messages are saved in the <code>var/log/routed</code> directory. Enter these words to set each trace option: <code>packets</code> — Trace all BGP packets to this peer. <code>open</code> — Trace all BGP open messages to this peer. <code>update</code> — Trace all BGP update messages to this peer. <code>keepalive</code> — Trace all keepalive messages to this peer. <code>all</code> — Trace all message types. <code>general</code> — Trace message related to Route and Normal. <code>route</code> — Trace routing table changes for routes installed by this peer. <code>normal</code> — Trace normal protocol occurrences. Abnormal protocol occurrences are always traced. <code>state</code> — Trace state machine transitions in the protocol. <code>policy</code> — Trace application of the protocol and user-specified policy to routes being imported and exported.
<code>capability {default ipv4-unicast ipv6-unicast}</code>	On each peer, configure the type of routes (Multiprotocol capability) to interchange between peers. Select one of these: - IPv4 Unicast Only (the default) - IPv6 Unicast Only - Both IPv4 and IPv6 To create peering, the routers must share a capability.
<code>graceful-restart-helper {on off}</code>	Sets the Check Point system to maintain the forwarding state advertised by peer routers even when they restart. This minimizes the negative effects caused by the restart of peer routers.
<code>graceful-restart-helper-stalepath-time seconds</code>	The maximum seconds that routes previously received from a restarting router are kept so that they can be revalidated. The timer starts after the peer sends an indication that it recovered.

BGP Confederation Commands

Use these commands to configure BGP confederations. You can configure a BGP confederation in conjunction with external BGP.

```
set bgp
  confederation identifier as_number
  confederation identifier off
  confederation aspath-loops-permitted <1-10>
  confederation aspath-loops-permitted default
  routing-domain identifier as_number
  routing-domain identifier off
  routing-domain aspath-loops-permitted <1-10>
  routing-domain aspath-loops-permitted default
  synchronization <on | off>
```

Arguments

Parameter	Description
<code>confederation identifier as_number</code>	Specifies the identifier for the entire confederation. This identifier is used as the autonomous system number in external BGP sessions. Outside the confederation, the confederation id is the autonomous system number of a single, large autonomous system. Thus the confederation id must be a globally unique, typically assigned autonomous system number.
<code>confederation identifier off</code>	Disables the confederation identifier.
<code>confederation aspath-loops-permitted <1-10></code>	Specifies the number of times the local autonomous system can appear in an autonomous system path for BGP-learned routes. If this number is higher than the number of times the local autonomous system appears in an autonomous system path, the corresponding routes are discarded or rejected.
<code>confederation aspath-loops-permitted default</code>	Specifies a value of 1.
<code>routing-domain identifier as_number</code>	Specifies the routing domain identifier (RDI) for this router. You must specify the RDI if you are using BGP confederations. The RDI does not need to be globally unique since it is used only within the domain of the confederation.
<code>routing-domain identifier off</code>	Disables the routing-domain identifier.
<code>routing-domain aspath-loops-permitted <1-10></code>	Specifies the number of times the local autonomous system can appear in an autonomous system path for BGP-learned routes. If this number is higher than the number of times the local autonomous system appears in an autonomous system path, the corresponding routes are discarded or rejected.

Parameter	Description
routing-domain aspath-loops-permitted default	Specifies a value of 1.
synchronization <on off>	Enables IGP synchronization. Set this option On to cause internal and confederation BGP peers to check for a matching route from IGP protocol before installing a BGP learned route.

Use these commands to configure BGP confederation peers.



Note - The IP address of a peer can be an IPv4 or an IPv6 address.

```
set bgp confederation member-as <as_id>
    [on|off]
    description [off|<description>]
    interface <int> [off|on]
    local-address <IP_addr> [off|on]
    med [default|<value>]
    nexthop-self [off|on]
    outdelay [off|<delay>]
    peer <IP_addr> [off|on]
    protocol [all|bgp|direct|rip|static|ospf|ospfase]
    peer <IP_addr> accept-routes [all|none]
    peer <IP_addr> authtype [none|md5 secret <passwd>]
    peer <IP_addr> capability [ipv4-unicast|ipv6-unicast] [off|on]
    peer <IP_addr> graceful-restart [off|on]
    peer <IP_addr> graceful-restart-stalepath-time [default|<time>]
    peer <IP_addr> holdtime [default|<time>]
    peer <IP_addr> ignore-first-ashop [off|on]
    peer <IP_addr> keepalive [default|<time>]
    peer <IP_addr> local-address <local_IP_addr> [off|on]
    peer <IP_addr> log-state-transitions [off|on]
    peer <IP_addr> log-warnings [off|on]
    peer <IP_addr> no-aggregator-id [off|on]
    peer <IP_addr> outgoing-interface <int> [off|on]
    peer <IP_addr> passive-tcp [off|on]
    peer <IP_addr> peer-type [none|reflector-client|no-client-reflector]
[off|on]
    peer <IP_addr> ping [off|on]
    peer <IP_addr> route-refresh [off|on]
    peer <IP_addr> send-keepalives [off|on]
    peer <IP_addr> send-route-refresh request [all|ipv4|ipv6] unicast
    peer <IP_addr> send-route-refresh route-update [all|ipv4|ipv6] unicast
    peer <IP_addr> throttle-count [off|<count>]
    peer <IP_addr> trace
[keepalive|open|packets|update|all|general|normal|policy|route|state|
task|timer] [off|on]
    peer <IP_addr> weight <weight>
    peer <IP_addr> comment <comment>
```

Arguments

Parameter	Description
<on off>	Creates (on) or removes (off) a peer group with AS id <as_id>.
description [off <description>]	Sets the peer group description to <description>, or turns off the description (off).

Parameter	Description
interface <int> [off on]	Sets a gateway interface (<int>: eth1, eth2, etc.) as the peer group interface, and turns it on or off.
local-address <IP_addr> [off on]	Sets a peer group with an IP address on the local gateway.
med [default <value>]	Sets the peer group local Multi-Exit Discriminator. The default is 0.
nexthop-self [off on]	Sets (on) or removes (off) the local gateway as the default exit gateway for the peer group.
outdelay [off <delay>]	Sets or removes the out-delay value (in seconds). Set this value to enforce rate limiting (" Rate Limiting for DoS Mitigation " on page 60).
peer <IP_addr> [off on]	Creates a peer group with the specified gateway (<IP_addr>).
protocol [all bgp direct rip static ospf ospfase]	Sets an internal peer group protocol.
peer <IP_addr> accept-routes [all none]	Accepts routes from peers only if there is an inbound BGP route policy. In the absence of a configured import policy for this peer, specify 'all' or 'none' here. The default is 'all'. - all - Accepts and installs routes with an invalid preference. Depending on the local BGP inbound policy, the routes can become active or inactive. - none - Deletes routes from a peer when no explicit local BGP inbound policy exists. Use this option to save memory overhead when many routes are rejected because there is no local policy. These routes can be re-learned only if you restart the BGP session.
peer <IP_addr> authtype [none md5 secret <passwd>]	Sets peer authentication between the local gateway and the specified peer gateway (<IP_addr>). You can set it to MD5 and specify the password (<passwd>), or you can turn it off (none).
peer <IP_addr> capability [ipv4-unicast ipv6-unicast] [off on]	Configures peer multiprotocol capabilities (ipv4-unicast or ipv6-unicast) with the specified peer (<IP_addr>). Turn these on or off, if necessary.
peer <IP_addr> graceful-restart [off on]	Turns graceful restart on and off between the local gateway and the specified peer (<IP_addr>).
peer <IP_addr> graceful-restart-stalepath-time [default <time>]	Sets graceful restart stalepath time (in seconds) with the specified peer (<IP_addr>).
peer <IP_addr> holdtime [default <time>]	Sets the maximum amount of time (in seconds) that can elapse between messages from the specified peer (<IP_addr>).

Parameter	Description
peer <IP_addr> ignore-first-ashop [off on]	Sets the router to ignore the first AS number in the AS_PATH for routes learned from the specified peer. Use this option for a route server peer in so-called transparent mode. The route server is configured to redistribute routes from multiple ASes and does not prepend its own AS number.
peer <IP_addr> keepalive [default <time>]	Sets the keepalive timer (in seconds) for the specified peer (<IP_addr>).
peer <IP_addr> local-address <local_IP_addr> [off on]	Sets a local IP address (<local_IP_addr>) for the specified peer (<IP_addr>).
peer <IP_addr> log-state-transitions [off on]	Turns logging of peer state transitions on or off for the specified peer (<IP_addr>).
peer <IP_addr> log-warnings [off on]	Turns logging of warnings on or off for the specified peer (<IP_addr>).
peer <IP_addr> no-aggregator-id [off on]	Sets the specified peer (<IP_addr>) to not aggregate AS routes (on). If set to off, the peer will create aggregate routes.
peer <IP_addr> outgoing-interface <int> [off on]	Sets a specific outgoing interface (<int>) to the specified peer (<IP_addr>).
peer <IP_addr> passive-tcp [off on]	Sets peer passive behavior. If on, the gateway does not initialize connections to the specified remote peer (<IP_addr>). The default is off.
peer <IP_addr> peer-type [none reflector-client no-client-reflector] [off on]	Sets the local gateway's peer type in the relation to the specified peer (<IP_addr>).
peer <IP_addr> ping [off on]	Sets ping capability between the local gateway and the specified peer (<IP_addr>). The default is off.
peer <IP_addr> route-refresh [off on]	Sets route refresh capability between the local gateway and the specified peer (<IP_addr>). The default is off.
peer <IP_addr> send-keepalives [off on]	Sets the gateway to always send keepalive messages to the specified peer (<IP_addr>). The default is off.

Parameter	Description
peer <IP_addr> send-route-refresh request [all ipv4 ipv6] unicast	Sets the local gateway to request BGP route updates from the specified peer (<IP_addr>).
peer <IP_addr> send-route-refresh route-update [all ipv4 ipv6] unicast	Sets the local gateway to respond to requests for BGP route updates from the specified peer (<IP_addr>).
peer <IP_addr> throttle-count [off <count>]	Sets the maximum number of BGP updates that can be sent at one time to the specified peer (<IP_addr>). The range for the <count> is 0-65535. The default is off.
peer <IP_addr> trace [keepalive open pack ets update all gener al normal policy rou te state task timer] [off on]	Sets the types of packets to trace from the specified peer (<IP_addr>).
peer <IP_addr> weight <weight>	Sets the weight for the specified peer (<IP_addr>). The value range for the <weight> is 0-65535.
peer <IP_addr> comment <comment>	Sets a comment associated with the specified peer (<IP_addr>).

BGP Route Reflection Commands

Use these commands to configure BGP route reflection. You can configure route reflection as an alternative to BGP confederations. Route reflection supports both internal and external BGP routing groups.

```
set bgp
  internal peer <ip_address> peer-type reflector-client
  internal peer <ip_address> peer-type none
  internal peer <ip_address> peer-type no-client-reflector
  cluster-id ip_address
  cluster-id off
  default-med <0-65535>
  default-med off
  default-route-gateway ip_address
  default-route-gateway off
```

Parameter	Description
internal peer <peer ip_address> peer-type reflector-client	The peer router <peer ip_address> is a reflector client of the local router.
internal peer <peer ip_address> peer-type none	The peer router <peer ip_address> is not a reflector client of the local router. This is the default .

Parameter	Description
<code>internal peer <peer ip_address></code> <code>peer-type</code> <code>no-client-reflector</code>	An advanced option.
<code>cluster-id ip_address</code>	The cluster ID used for route reflection. The cluster ID default is that of the router id. Override the default if the cluster has more than one route reflector
<code>cluster-id off</code>	Disable the cluster ID.
<code>default-med <0-65535></code>	The multi-exit discriminator (MED) metric used to advertise routes through BGP.
<code>default-med off</code>	Disable the specified MED metric.
<code>default-route-gateway ip_address</code>	The default route. This route has a higher rank than any configured default static route for this router. If you do not want a BGP peer considered for generating the default route, use the <code>peer <ip_address> suppress-default-originate</code> on command.
<code>default-route-gateway off</code>	Disables the configured default BGP route.

BGP Route Dampening Commands

Use the following commands to configure BGP route dampening. BGP route dampening maintains a history of flapping routes and prevents advertising these routes. A route is considered to be flapping when it is repeatedly transitioning from available to unavailable or vice versa.

```
set bgp dampening
  <on | off>
  suppress-above <2-32>
  suppress-above default
  reuse-below <1-32>
  reuse-below default
  max-flat <3-64>
  max-flat default
  reachable-decay <1-900>
  reachable-decay default
  unreachable-decay <1-2700>
  unreachable-decay default
  keep-history <2-5400>
  keep-history default
```

Note: BGP route dampening is only supported for External BGP (EBGP).

Parameter	Description
<code><on off></code>	Specifies whether to enable or disable BGP route dampening.
<code>suppress-above <2-32></code>	Specifies the value of the instability metric at which route suppression takes place. A route is not installed in the forwarding table or announced even if it is reachable during the period that it is suppressed.
<code>suppress-above default</code>	Specifies an instability metric value for suppressing routes of 3.
<code>reuse-below metric <1-32></code>	Specifies the value of the instability metric at which a suppressed route becomes unsuppressed if it is reachable but currently suppressed. The value assigned to the reuse-below metric must be lower than the suppress-above value.
<code>reuse-below metric default</code>	Specifies an instability metric value for announcing previously suppressed routes of 2.
<code>max-flap <3-64></code>	Specifies the upper limit of the instability metric. The value must be greater than the suppress-above value plus 1. Each time a route becomes unreachable, 1 is added to the current instability metric.
<code>max-flat default</code>	Specifies the upper limit of the instability metric as 16.
<code>reachable-decay <1-900></code>	Specifies the time for the instability metric to reach half of its value when the route is reachable. The smaller the value the sooner a suppressed route becomes reusable.
<code>reachable-decay default</code>	Specifies a value of 300.
<code>unreachable-decay <1-2700></code>	Specifies the time for the instability metric to reach half its value when the route is NOT reachable. The value must be equal to or higher than the reachable-decay value.
<code>unreachable-decay default</code>	Specifies a value of 900
<code>keep-history <2-5400></code>	Specifies the period for which route flapping history is maintained for a given route.
<code>keep-history default</code>	Specifies a value of 1800.

Internal BGP Commands

Use the following commands to configure internal BGP sessions, that is, between routers within the same autonomous system.

```
set bgp internal
  <on | off>
  description text
  med <0-65535>
  med default
  outdelay <0-65535>
  outdelay off
  nexthop-self <on | off>
```

```

local-address ip_address <on | off>
interface [all | if_name] <on | off>
protocol [all | bgp_internal_protocol] <on | off>
graceful-restart-helper <on | off>
graceful-restart-helper-stalepath-time seconds
route-refresh <on | off>
set bgp internal peer <ip_address>
peer_type <on | off>
weight <0-65535>
weight off
no-aggregator id <on | off>
holdtime <6-65535>
holdtime default
keepalive <2-21845>
keepalive default
ignore-first-ashop <on | off>
send-keepalives <on | off>
send-route-refresh [request | route-update] [ipv4|ipv6|All] [unicast]

```

Parameter	Description
<on off>	Enable or disable an internal BGP group.
description text	Optional: A brief text description of the group.
med <0-65535>	
med default	
outdelay <0-65535>	The amount of time in seconds that a route must be present in the routing database before it is redistributed to BGP. The configured value applies to all peers configured in this group. This feature dampens route fluctuation. Zero (0), which means that this feature is disabled. Default: 0
outdelay off	Disables outdelay.
nexthop-self on off>	This router sends one of its own IP addresses as the BGP next hop. Default: off

Parameter	Description
<code>local-address</code> <code>ip_address <on off></code>	The address used on the local end of the TCP connection with the peer. For external peers that do not have multihop enabled, the local address must be on an interface that is shared with the peer or with the peer's gateway when the gateway parameter is used. A session with an external peer is opened only when an interface with a local address through which the peer or gateway address is directly reachable is operating. For other types of peers, a peer session is maintained when any interface with the specified local address is operating. In either case, incoming connections are recognized as matching a configured peer only if they are addressed to the configured local address. Note: If running BGP in a cluster you must not configure the local address. Default: Off
<code>interface [all if_name] <on off></code>	Enable or disable the specified internal peer group on all interfaces or a specific interface.
<code>protocol [all bgp_internal_protocol] <on off></code>	Enable or disable all internal routing protocols on the specified internal peer group or specific internal protocols. You can enter the following specific internal protocols: <code>direct</code>, <code>rip</code>, <code>static</code>, <code>ospf</code>, and <code>ospfase</code>.
<code>peer ip_address</code> <code>peer_type <on off></code>	An internal peer address and peer type. Enter <code>reflector-client</code> to specify that the local router acts as a route reflector for the group of peers named. That is, the local router is the route reflection server, and the named peers are route reflection clients. Normally, the routing daemon readvertises, or reflects, routes it receives from one of its clients to all other IBGP peers, including the other peers in that client's group. Enter <code>no-client-reflector</code> to specify that a reflection client's routes are reflected only to internal BGP peers in other groups. Clients in the group are assumed to be direct IBGP peers of each other. Enter <code>none</code> if you do not want to specify route reflection.
<code>peer ip_address</code> <code>weight</code> <code><0-65535></code>	The weight associated with the specified peer. BGP implicitly stores any rejected routes by not mentioning them in a route filter. BGP explicitly mentions them within the routing table by using a <code>restrict</code> keyword with a negative weight. A negative weight prevents a route from becoming active, which prevents it from being installed in the forwarding table or exported to other protocols. This eliminates the need to break and reestablish a session upon reconfiguration if import route policy is changed.
<code>peer ip_address weight</code> <code>off</code>	Disables the weight associated with the specified peer.

Parameter	Description
peer ip_address aggregator id <on off>	The router's aggregate attribute as zero (rather than the router ID value). This option prevents different routers in an AS from creating aggregate routes with different AS paths Default: off
peer ip_address holdtime <6-65535>	The BGP holdtime interval, in seconds, when negotiating a connection with the specified peer. If the BGP speaker does not receive a keepalive update or notification message from its peer within the period specified in the holdtime field of the BGP open message, the BGP connection is closed.
peer ip_address holdtime default	A holdtime of 180 seconds.
peer ip_address keepalive <2-21845>	The keepalive option is an alternative way to specify a holdtime value in seconds when negotiating a connection with the specified peer. You can use the keepalive interval instead of the holdtime interval. You can also use both interval, but the holdtime value must be 3 times the keepalive interval value.
peer ip_address_keepalive default	A keepalive interval of 60 seconds.
peer ip_address ignore-first-ashop <on off>	Ignore the first autonomous system number in the autonomous system path for routes learned from the corresponding peer. Set this option only if you are peering with a route server in transparent mode, that is, when the route server is configured to redistribute routes from multiple other autonomous systems without prepending its own autonomous system number.
peer ip_address send-keepalives <on off>	This router always sends keepalive messages even when an update message is sufficient. This option allows interoperability with routers that do not strictly adhere to protocol specifications regarding update.
send-route-refresh [request route-update [ipv4 ipv6 all] [unicast]]	The router dynamically request BGP route updates from peers or respond to requests for BGP route updates.
peer ip_address accept-routes all	An inbound BGP policy route if one is not already configured. Enter all to specify accepting routes and installing them with an invalid preference. Depending on the local inbound route policy, these routes are then made active or inactive.
peer ip_address accept-routes none	An inbound BGP policy route if one is not already configured. Enter none to specify deleting routes learned from a peer. This option saves memory overhead when many routes are rejected because no inbound policy exists.
peer ip_address passive-tcp <on off>	The router waits for the specified peer to issue an open message. No tcp connections are initiated by the router. Default: off

Parameter	Description
peer ip_address authtype none	Do not use an authentication scheme between peers. Using an authentication scheme guarantees that routing information is accepted only from trusted peers.
peer ip_address authtype md5 secret secret	Use md5 authentication between peers. In general, peers must agree on the authentication configuration to and from peer adjacencies. Using an authentication scheme guarantees that routing information is accepted only from trusted peers. Note - TCP MD5 is not supported on BGP IPv6 peers.
peer ip_address throttle-count <0-65535>	The number of BGP updates to send at one time. The throttle count option limits the number of BGP updates when there are many BGP peers.
peer ip_address throttle count off	Disables the throttle count option.
peer ip_address log-state-transitions <on off>	The router generates a log message whenever a peer enters or leave the established state.
peer ip_address log-warnings <on off>	The router generates a log message whenever a warning scenario is encountered in the codepath.
peer ip_address trace bgp_traceoption <on off>	Tracing options for the BGP implementation. Log messages are saved in the <code>var/log/routed</code> directory. Enter the following words to set each trace option. Enter <code>packets</code> to trace all BGP packets to this peer. Enter <code>open</code> to trace all the BGP open messages to this peer. Enter <code>update</code> to trace all the BGP update messages to this peer. Enter <code>keepalive</code> to trace all the keepalive messages to this peer. Enter <code>all</code> to trace all the message types. Enter <code>general</code> to trace message related to Route and Normal. Enter <code>route</code> to trace routing table changes for routes installed by this peer. Enter <code>normal</code> to trace normal protocol occurrences. Abnormal protocol occurrences are always traced. Enter <code>state</code> to trace state machine transitions in the protocol. Enter <code>policy</code> to trace application of the protocol and user-specified policy to routes being imported and exported.
graceful-restart-helper <on off>	Whether the Check Point system should maintain the forwarding state advertised by peer routers even when they restart to minimize the negative effects caused by peer routers restarting.
graceful-restart-helper -stalepath-time seconds	The maximum amount of time that routes previously received from a restarting router are kept so that they can be revalidated. The timer is started after the peer sends an indication that it has recovered.

Parameter	Description
<code>route-refresh <on off></code>	Re-learns routes previously sent by the BGP peer or refreshes the routing table of the peer. The peer responds to the message with the current routing table. Similarly, if a peer sends a route refresh request the current routing table is re-sent. A user can also trigger a route update without having to wait for a route refresh request from the peer.

BGP Communities Commands

Use the following command to configure BGP communities. A BGP community is a group of destinations that share the same property. However, a community is not restricted to one network or autonomous system. Use communities to simplify the BGP inbound and route redistribution policies. Use the BGP communities commands together with inbound policy and route redistribution.

```
set bgp communities <on | off>
```

Parameter	Description
<code><on off></code>	Enable or disable BGP policy options based on communities.

BGP Show Commands

Use the following commands to monitor and troubleshoot your BGP implementation.

```
show bgp
show bgp
  groups
  memory
  errors
  paths
  stats
  peers
  peers detailed
  peer ip_address detailed
  peers established
  peer ip_address advertise
  peer ip_address received
  summary
```

Rate Limiting for DoS Mitigation

Rate Limiting is a defense against DoS (Denial of Service) attacks. A policy limits traffic coming from specified sources and services.

Rate limiting is enforced on:

- Bandwidth and packet rate
- Number of concurrent connections
- Connection rate

Rate Limiting for DoS Mitigation is scalable and can support a large number of rules. You can define policies that limit bandwidth for traffic from geographic sources that are not in your business profile. You can configure white-list bypasses.

Important: During the installation of the Access Control policy on the gateway, the rate limiting policy is not enforced.

Rate Limiting Support

Rate Limiting for DoS Mitigation is supported on:

- Gaia gateways with Performance Pack installed.
- VSX. On R77.20 and higher, you can configure different settings for the Virtual Systems or global parameters for all of them. On R77.10 and lower, you can configure only the VSX Gateway (VS0).

Configuring Rate Limiting for DoS Mitigation

To prevent Denial of Service (DoS) attacks, add rules to a policy one at a time, or in batch mode ["[Adding Rules in Batch Mode](#)" on page 64].

If this gateway is a cluster member, configure Rate Limiting for DoS Mitigation on all of the cluster members.



Note - By default, the rules are loaded only on the local gateway, unless you specify a different gateway with the `-S <server>` parameter.

Creating Quota Rules

To prevent Denial of Service (DoS) attacks, add **quota** rules to a policy.

Syntax:

```
fw samp add {-a {d | n | b}} [-l r] [-t <timeout>] [-n <name>] [-c <comment>]
[-o <originator>] [-S <IP>] source {any | range:<IP>[-<IP>] | cidr:<IP>/<netmask>
| cc:<country_code>|asn:<sys_number>} [destination {<property>:<value>}]
[source-negated {true | false}] [destination-negated {true | false}]
[service <protocol> | <protocol>-<protocol> | <protocol>/<port> |
<protocol>/<port>-<port>] [service-negated {true | false}] quota
{[new-conn-rate <seconds>] [new-conn-rate-ratio <number>] [concurrent-conns
<number>] [concurrent-conns-ratio <number>] [pkt-rate <number>]
[pkt-rate-ratio <number>] [byte-rate <number>] [byte-rate-ratio <number>]
[track {track source | track source-service}]} [flush true]
```

Parameter	Description and Values
-a	Required action on the incoming packets that match the rule. Valid values: - d - drops the packet - n - notify: logs the packet and lets it through - b - bypass: lets the packet through without checking it against the policy rules. If this is the action you set, there are no logs or limits. Bypassed packets and connections are not added to total number of packets or connections for limit enforcement of type ratio.
-l r	If given, regular logging is done on matching traffic.
-t	The number of seconds (integers) after which the rule expires. If not given, the rule does not expire.
-n	Name of the rule.
-c	Optional free-text comment for the rule. Escape spaces and backslashes with a backslash. Do not use other special characters.
-o	Name of the originator.
-s	IP address of a target gateway for policy installation. By default, the rules are loaded only on the local gateway, unless you specify a different gateway with this parameter.
source	Definition of the source, from which to limit traffic. You can add multiple properties and sources, delimited by a comma: <code>source TYPE:VALUE [,TYPE:VALUE, TYPE:VALUE, ...TYPE:VALUE]</code> Valid values: - any - The rule is applied to packets from any source. - range: <i>IP ADDRESS</i> or range: <i>IP ADDRESS- IP ADDRESS</i> - IPv4 or IPv6 addresses - cidr: <i>IP ADDRESS/NETMASK</i> - IPv4 or IPv6 address, netmask 0 to 32 for IPv4, 0 to 128 for IPv6. - cc: COUNTRY_CODE - Two-letter code defined in ISO 3166-1 alpha-2 http://www.iso.org/iso/iso-3166-1_decoding_table.html. The rule matches the country code to the addresses assigned to this country, based on the Geo IP database. - asn: AUTONOMOUS_SYSTEM_NUMBER - ASnnnn, where nnnn is a number unique to the specific organization. The rule matches the AS number of the organization to the IP addresses that are assigned to this organization, based on the Geo IP database.
destination	Definition of a specific destination, to limit inbound traffic. Properties are the same as the source.
source-negated	To ignore the action on traffic from the specified source, use <code>source-negated true</code>
destination-negated	To ignore the action on traffic from the specified destination, use <code>destination-negated true</code>

Parameter	Description and Values
<code>service</code>	Service protocols, ports, or ranges of protocols or ports. Valid values: • <code>any</code> • <i>PROTOCOL</i> - IP protocol number in the range 1 - 255 (see http://www.iana.org/assignments/protocol-numbers/protocol-numbers.xhtml) • <i>PORT</i> - TCP or UDP port number in the range 1 - 65535 (see http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xhtml)
<code>service-negated</code>	To ignore the action on traffic using the specified services, use <code>service-negated true</code>
<code>quota</code>	Indicates that the next parameters describe the Rate Limiting quota.
<code>new-conn-rate</code>	Maximum (per second) number of connections that match the rule.
<code>new-conn-rate-ratio</code>	Maximum ratio of the <code>new-conn-rate</code> value to the rate of all connections per second through the gateway, expressed in parts per 65536
<code>concurrent-conns</code>	Maximum number of concurrent active connections that match the rule.
<code>concurrent-conns-ratio</code>	Maximum ratio of the <code>concurrent-conns</code> value to the total number of active connections through the gateway, expressed in parts per 65536.
<code>pkt-rate</code>	Maximum per second number of packets that match the rule.
<code>pkt-rate-ratio</code>	Maximum ratio of the <code>pkt-rate</code> value to the rate of all connections through the gateway, expressed in parts per 65536.
<code>byte-rate</code>	Maximum total number of bytes per second in packets that match the rule.
<code>byte-rate-ratio</code>	Specifies the maximum ratio of the <code>byte-rate</code> value to the bytes per second rate of all connections through the gateway, expressed in parts per 65536.
<code>track</code>	Criteria for counting connections, packets, and bytes: • <code>track source</code> - Count for each source IP address. • <code>track source-service</code> - Count for each source IP address and service (protocol and destination port).
<code>flush</code>	These rules are not immediately applied to the gateway. They are only registered in the Suspicious Activity Monitoring policy database. To apply all the rules from the policy database immediately, use: <code>flush true</code>

Notes:

New rules apply only to new connections, not to existing connections.

Example of a rule with a range:

```
fw samp add -a d -l r -t 3600 -c rule\ for\ IPs\SE quota service any source
range:172.16.7.11-172.16.7.13 new-conn-rate 5 flush true
```

- Limits the rate of creation of new connections from the IP addresses in the range 172.16.7.11-172.16.7.13 to 5 connection per second. Drops all other attempted connections (-a d).
- Logs packets that exceed the quota set by the rule. The limit of the total number of log entries per second is set through the global parameter `sim_dos_ctl -l LOG-LIMIT`.
- Expires after one hour (3600 seconds).
- This rule is immediately compiled and loaded with other rules in the Suspicious Activity Monitoring policy database, because this rule includes `flush true`.

Example of a rule with a service specification:

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

- Logs all packets (-a n) coming from IP addresses from the country with country code SE.
- Does not let any traffic through (`byte-rate 0`), except for the packets (`service-negated true`) that match the IP protocols on the list (ICMP, IPSec, HTTPS, DNS).
- Does not expire (to cancel this rule, delete it explicitly), and is not enforced until you install the policy.

Example of a rule with ASN:

```
fw samp -a d quota source asn:AS64500,cidr:[::ffff:c0a8:1100]/120 service any
pkt-rate 0
```

- Drops all packets (-a d) with the source IP address in the IPv6 address block (`cidr:[::ffff:c0a8:1100]/120`), from the autonomous system number 64500 (`asn:AS64500`).
- Does not expire, and is not enforced until you install the policy.

Example of a whitelist rule:

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

- Ignores all other quota type of rules (other policy rules still apply) that match the traffic (-a b). Lets through all HTTP traffic (`service 6/80`) from the specified address range (`source range:172.16.8.17-172.16.9.121`).
- Does not expire, and is not enforced until you install the policy.

Example of a tracked rule:

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

- Drops (-a d) new connections for every IP address that has more than approximately 1% (655/65536) of all existing connections (`concurrent-conns-ratio 655`).
- Defines IP addresses that are assigned to a specific country (`source-negated true source cc:SE`) as exception to the rule.
- Tracks the count of connections, packets, and bytes from the source.

Adding Rules in Batch Mode

To add rules in batch mode:

1. Enter this command:

```
fw samp [-s <IP_ADDRESS>] batch << EOF
```

Note: If you include the -s parameter, all the commands in this batch apply to the specified gateway.

2. Enter one *add* or *delete* command for each line, on as many lines as necessary. Start each line with *add* or *del* parameter, and not with *fw samp*. Use the same set of parameters and values as for the individual rules. Terminate each line with a Return (ASCII 10 - Line Feed) character:

```
add -a d|n|b [-l r][-t <TIMEOUT>][-n <NAME>][-c <COMMENT>][-o <ORIGINATOR>]
quota quota {[new-conn-rate <seconds>] [new-conn-rate-ratio <number>]
[concurrent-conns <number>] [concurrent-conns-ratio <number>] [pkt-rate
<number>] [pkt-rate-ratio <number>] [byte-rate <number>]
[byte-rate-ratio <number>] [track {track source | track source-service}]}
del <UID>
```

3. To end the batch, enter: EOF.

Example:

```
fw samp -S 192.168.37.5 batch <<EOF
add -a d -l r -t 3600 -c a\ comment quota service any source
range:172.16.7.13-172.16.7.13 new-conn-rate 5
del <501f6ef0,00000000,cb38a8c0,0a0afffe>
add -a b quota source range:172.16.8.17-172.16.9.121 service 6/80
EOF
```

This batch applies two *add* commands and one *delete* command to a gateway with the IP address 192.168.37.5.



Note - A space or a backslash in comments must be each preceded by a backslash:

```
-c this\ is\ a\ comment\ with\ a\ backslash\ \
```

Deleting a Rule

To delete a rule:

1. List all the rules in the Suspicious Activity Monitoring policy database:
`fw samp get`
 The rules show in this format:
`... operation=add uid=<501f6ef0,00000000,cb38a8c0,0a0afffe> target=all
 timeout=... action=... ..`
2. Delete a rule from the list:
`fw samp del '<501f6ef0,00000000,cb38a8c0,0a0afffe>'`
3. Enter this flush-only *add* rule:
`fw samp add -t 2 quota flush true`
 This immediately deletes the rule, and times out in 2 seconds. It is a good practice to specify a short timeout period for the flush-only rules. This prevents accumulation of rules that are obsolete in the database.
 The `fw samp del` command removes a rule from the persistent database only. The deleted rule continues to be enforced until the next time a policy is compiled and loaded. To force the rule deletion immediately, you must enter a flush-only *add* rule right after the `fw samp del` command.

Configuring Rules with VSX

This feature is supported in R77.20 and higher.

Use the `fw samp` command to add rules to a Virtual System. You can run the command on the local gateway, or remotely configure the rules.

To configure rules on a local gateway:

1. Log in to the gateway CLI.
2. Set the environment for the specified Virtual System: `set virtual-system <vsid>`
3. Run the `fw samp` command ("[Creating Quota Rules](#)" on page 60).

To configure rules on a remote gateway:

1. Log in to the server CLI.
2. Run: `fw vsx showncs -vs <VSID>`
 Record the output, the name of the Virtual System that established SIC.
3. Run: `fw samp -S <VSX_IP> -s <SIC_name> [options]` ("[Creating Quota Rules](#)" on page 60)

Configuring Global Parameters

There are several global parameters that you can configure with `sim_dos ctl` command for IPv4 addresses and with `sim6_dos ctl` for IPv6 addresses. They apply to all the policy rules.



Note - `sim_dos ctl` and `sim6_dos ctl` are only available as CLI commands on the gateways. Remote command option is not available.

Use the `sim_dos ctl` or `sim6_dos ctl` command with these parameters and values.

Parameter and Values	Description
<code>-m {1 0}</code>	Turns on the monitor-only mode, when set to 1. In this mode, rules do not drop any packets, regardless of the action specified. Each rule only does logging, as specified in it.
<code>-x {1 0}</code>	When set to 1 (default), the rules are only applied to traffic that arrives on the external interfaces of the gateway. When set to 0, the rules are applied to traffic regardless of the interface on which it arrives. Note: This does not apply to other security policies on the gateway. They still get enforced.
<code>-l n</code>	Sets the limit for the number of log entries per second (the default is 100). All the entries that exceed the limit are suppressed. The number of suppressed messages shows in the following period summary.
<code>-a {1 0}</code>	Turns the quota policy rules enforcement on (1) and off (0). When the rule enforcement is turned off, no traffic is matched against the quota rules. Note: The quota rule enforcement is on automatically, when a policy with rules is loaded, and is off, when an empty policy is loaded.

The global parameters return to their default values every time the DoS in the Performance Pack module is initialized. This happens on every reboot. To keep the changes to global parameters until you decide to change them again, include the `sim_dos ctl` (or `sim6_dos ctl`) command in the `dospreload` script:

For IPv4:

```
$ cat >$PPKDIR/bin/dospreload4 <<EOF
#!/bin/bash
$PPKDIR/bin/sim_dos ctl -m 1 -x 0 -l 30
EOF
$ chmod +x $PPKDIR/bin/dospreload4
```

For IPv6:

```
$ cat >$PPKDIR/bin/dospreload6 <<EOF
#!/bin/bash
$PPKDIR/bin/sim6_dos ctl -m 1 -x 0 -l 30
EOF
$ chmod +x $PPKDIR/bin/dospreload6
```

For VSX:

Rate Limiting for VSX is supported in R77.20 and higher.

- `$PPKDIR/bin/` is shared by all the Virtual Systems
- `$FWDIR/scripts/` is specific for each Virtual System

Run these commands from Expert mode to apply monitor only mode (`-m 1`) to all Virtual Systems.

```
# cat > $PPKDIR/bin/dospreload4 << EOF
#!/bin/bash
sim_dos ctl -m 1
if test -x $FWDIR/scripts/dospreload4; then
```

```

$FWDIR/scripts/dospreload4
fi
EOF
# chmod +x $PPKDIR/bin/dospreload4

```

Run these commands from Expert mode to limit the number of log entries (-l 40) for the Virtual System with VSID 2.

```

# venv 2
Context is set to Virtual Device myVS-2 (ID 2).
# cat > $FWDIR/scripts/dospreload4 <<EOF
#!/bin/bash
sim_dos ctl -l 40
EOF
# chmod +x $FWDIR/scripts/dospreload4

```

Run similar commands to create \$PPKDIR/bin/dospreload6 and a \$FWDIR/scripts/dospreload6 script for each Virtual System.

Monitoring Events Related to DoS Mitigation

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

Command	Command Output
<pre>cat /proc/ppk/dos cat /proc/ppk6/dos (for IPv6)</pre>	Shows memory utilization, DoS policy rules, and global parameter configuration.
<pre>fw samp get -l grep '^<[0-9a-f,]*>\$' xargs sim_dos get</pre>	Shows details of active policy rules in long format. It only show rules loaded in IPv4 kernel. To see the rules in IPv6 kernel, use <code>sim6_dos get</code> command.
<pre>cat /proc/ppk/<VSID>/dos cat /proc/ppk6/<VSID>/dos (for IPv6)</pre>	VSX is supported in R77.20 and higher. Shows memory utilization, DoS policy rules, and global parameter configuration for Virtual Systems. <VSID> is the VSID for the Virtual System.

IGMP

In This Section:

IGMP Version 3	68
Configuring IGMP - WebUI	69
Configuring IGMP - CLI (igmp)	71

Internet Group Management Protocol (IGMP) allows hosts on multiaccess networks to inform locally attached routers of their group membership information. Hosts share their group membership information by multicasting IGMP host membership reports. Multicast routers listen for these host membership reports, and then exchange this information with other multicast routers.

The group membership reporting protocol includes two types of messages: host membership query and host membership report. IGMP messages are encapsulated in IP datagrams, with an IP protocol number of 2. Protocol operation requires that a designated querier router be elected on each subnet and that it periodically multicast a host membership query to the all-hosts group.

Hosts respond to a query by generating host membership reports for each multicast group to which they belong. These reports are sent to the group being reported, which allows other active members on the subnet to cancel their reports. This behavior limits the number of reports generated to one for each active group on the subnet. This exchange allows the multicast routers to maintain a database of all active host groups on each of their attached subnets. A group is declared inactive (expired) when no report is received for several query intervals.

The IGMPv2 protocol adds a leave group message and uses an unused field in the IGMPv1 host membership query message to specify a maximum response time. The leave group message allows a host to report when its membership in a multicast group terminates. Then, the IGMP querier router can send a group-directed query with a very small maximum response time to probe for any remaining active group members. This accelerated leave extension can reduce the time required to expire a group and prune the multicast distribution tree from minutes, down to several seconds.

The unicast traceroute program allows the tracing of a path from one device to another, using mechanisms that already exist in IP. Unfortunately, you cannot apply such mechanisms to IP multicast packets. The key mechanism for unicast traceroute is the ICMP TTL exceeded message that is specifically precluded as a response to multicast packets. The traceroute facility implemented within `routed` conforms to the traceroute facility for IP multicast draft specification.

Gaia supports IGMP version 1, v2 and v3. Version 2 runs by default.

IGMP Version 3

Gaia provides IGMP version 3 source filtering to support source-specific multicast (SSM), which enables the Gaia system to request traffic from specific sources via PIM join/prune messages without requiring the presence of a rendezvous point (RP). This enables the Gaia system to forward traffic from only those sources from which receivers requested traffic. IGMPv3 supports applications that explicitly signal sources from which they want to receive traffic.

With IGMP version 3, receivers (hosts) identify their membership to a multicast group in the following two modes:

- Include mode: Receivers announce membership to a group and provide a list of IP addresses (the include list) from which they want to receive traffic.
- Exclude mode: Receivers announce membership to a host group and provide a list of IP addresses (the exclude list) from which they do not want to receive traffic. To receive traffic from all sources, a host sends an empty exclude list.

The multicast group address range 232/8 (232.0.0.0 to 232.255.255.255) is reserved for use by SSM protocols and applications. The DRs of senders do not send register packets to any RPs in the SSM group range.

When SSM is enabled, all other multicast groups are treated as in normal sparse-mode.

Configuring IGMP - WebUI

IGMP is enabled by default.

To configure IGMP:

1. In the **Network Management > Network interfaces** page of the WebUI, configure Ethernet Interfaces and assign an IP address to the interface.
2. Configure a multicast routing protocol, such as PIM.
IGMP supports IP multicast groups on a network. IGMP functions only in conjunction with a multicast routing protocol to calculate a multicast distribution tree. For more information on multicast routing protocols supported by Gaia, see PIM.
3. Open the **Advanced Routing > IGMP** page of the WebUI.
4. For each interface on which you enabled a multicast routing protocol:
 - a) Select the interface and click **Add** or **Edit**.
The **Edit IGMP on Interface** window opens.
 - b) Configure the IGMP interface parameters. The parameters are optional.
 - c) **Optional:** Add a local network Multicast Group or a static multicast group. Click **Add**.
The **Add Multicast Group** window opens.
 - d) Configure the IGMP multicast group parameters.

Edit IGMP on Interface Window Parameters

Parameter	Description
Version	The version of the IGMP protocol to comply with. Note - IGMP version 2 is compatible with IGMP version 1, and version 3 is compatible with versions 2 and 1. Check Point recommends that you use version 1 only on networks that include multicast routers that are not upgraded to IGMP versions 2 or 3. IGMP version 3 is used to support source-specific multicast (SSM). Version 3 membership reports are used to request or block multicast traffic from specific sources. For example, when a host requests traffic for a multicast group from a specific source, SSM sends PIM join/prune messages towards the source. The multicast group address 232/8 is reserved for use with SSM. Version 3 is backwards compatible with versions 1 and 2. Range: 1-3. Default: 2.
Loss Robustness	Allows tuning for the expected packet loss on a subnet. If the subnet is expected to be highly lossy, then the "loss robustness" value may be increased. IGMP protocol operation is robust to (lossrobustness - 1) packet loss. Range: 1-255. Default: 2.
Query Interval	The interval (in seconds) between IGMP general queries sent by the querier router. This parameter can be used to tune the IGMP messaging overhead and has a secondary effect on the timeout of idle IP multicast groups. Range: 1-3600. Default: 125.
Query Response Interval	The maximum response time (in seconds) inserted into the periodic IGMP general queries. The query response interval may be used to tune the burstiness of IGMP messages; a larger value spreads the host IGMP reports over a larger interval, reducing burstiness. This value must always be less than the query interval. Range: 1-25. Default: 10.

Parameter	Description
Last Member Query Interval	The maximum response time (in seconds) inserted into IGMP group-specific queries. The last member query interval may be used to tune the "leave latency." A smaller value results in a reduction in the time to detect the loss of the last member of a multicast group. This value must always be less than the query interval. Range: 1-25. Default: 1.
Router Alert	Allows the "disable insertion of IP router alert" option in all IGMP messages sent on the interface. This can be useful in interoperating with broken IP implementations that may discard the packet due to the use of this option. Options: Enabled, Disabled Default: Enabled

Add Multicast Group Window

Parameter	Description
Multicast Address	The multicast address of the group
Group Type	• Local Group - Provides a mechanism to simulate the presence of local receivers for specific groups. When a multicast group is added to an interface, IGMP sends a membership report on the interface. • Static Group - Provides a mechanism to simulate the presence of local receivers on an interface. When a static group is configured on an interface that is also running a parent multicast protocol (such as PIM) IGMP informs the parent of the presence of a local receiver. In contrast to regular IGMP, no membership reports are sent on the corresponding interface. If the same multicast group is configured as both a local and a static group, local group takes precedence, that is, membership reports are sent out on the interface.

Configuring IGMP - CLI (igmp)

Use the IGMP commands to configure parameters for the internet group management protocol.

Configuring Interfaces for IGMP

Use these commands to configure IGMP for specific interfaces.

```
set igmp interface if_name
  version <1 | 2 | 3>
  last-member-query-interval <1-25>
  last-member=query-interval default
  loss-robustness <1-255>
```

```

loss-robustness default
query-interval <1-3600>
query-interval default
query-response-interval <1-25>
query-response-interval default
router-alert <on | off>
static-group address <on | off>
local-group address <on | off>

```

**Note -**

IGMP version 2 runs by default.

In a gateway cluster, run commands on every cluster member. The configuration of each cluster member must be identical.

Parameter	Description
interface if_name	The interface on which IGMP should be configured.
last-member-query-interval <1-25>	The maximum response time (in seconds) inserted into IGMP group-specific queries. You can use the last member query interval to tune the "leave latency." A smaller value results in a reduction in the time to detect the loss of the last member of a multicast group. This value must always be less than the query interval.
last-member-query-interval default	A value of 1.
loss-robustness <1-255>	Lets you tune the expected packet loss on a subnet. If you expect the subnet to be highly lossy, then you can increase the "loss robustness" value. IGMP protocol operation is robust to (lossrobustness - 1) packet loss. • Default: 2
loss-robustness default	A value of 2.
query-interval <1-3600>	The interval (in seconds) between IGMP general queries which the querier router sends. You can use this parameter to tune the IGMP messaging overhead and has a secondary effect on the timeout of idle IP multicast groups. • Default: 125
query-interval default	A value of 125.
query-response-interval <1-25>	The maximum response time (in seconds) inserted into the periodic IGMP general queries. You can use the query response interval to tune the burstiness of IGMP messages; a larger value spreads the host IGMP reports over a larger interval, which reduces burstiness. This value must always be less than the query interval. • Default: 10.
query-response-interval default	A value of 10.

Parameter	Description
<code>router-alert</code> <on off>	Lets you disable the insertion of IP router alert in all IGMP messages sent on the interface. This can be useful with broken IP implementations that may discard the packet because of the use of this option. • Default: off
<code>local-group address</code> <on off>	• A multicast group address. A local group provides a mechanism to simulate the presence of local receivers for specific groups. When a multicast group is added to an interface, IGMP sends a membership report on the interface.
<code>static-group address</code> <on off>	• A multicast group address. A static group provides a mechanism to simulate the presence of local receivers on an interface. When a static group is configured on an interface that also runs a parent multicast protocol (such as PIM), IGMP informs the parent of the presence of a local receiver. In contrast to regular IGMP, no membership reports are sent on the corresponding interface. If the same multicast group is configured as both a local and a static group, local group takes precedence, that is, membership reports are sent out on the interface.
<code>version</code> <1 2 3>	IGMP version 2 is compatible with IGMP version 1, and version 3 is compatible with versions 2 and 1. Best Practice - use version 1 only on networks that include multicast routers that are not upgraded to IGMP versions 2 or 3.

Monitoring IGMP (show igmp)

Use these commands to monitor and troubleshoot IGMP.

```
show igmp
  stats
  stats receive
  stats transmit
  stats error
  interfaces
  interfaces if_address
  groups [interface logical_interface] [local | static]
  group if_address
  if-stats
  if-stat if_address
  summary
```

IP Broadcast Helper

In This Section:

Configuring IP Broadcast helper - WebUI	74
Configuring IP Broadcast Helper - CLI (iphelper)	75
Monitoring IP Broadcast Helper	76

IP Broadcast Helper is a form of static addressing that uses directed broadcasts to forward local and all-nets broadcasts to desired destinations within the internetwork. IP Broadcast Helper allows the relaying of broadcast UDP packets on a LAN as unicasts to one or more remote servers. This is useful when you relocate servers or hosts from their original common segment and still want the service to be available.



Note - For more information, see RFC1542 section 4.

Configuring IP Broadcast helper - WebUI

To configure IP broadcast helper

1. Open the **Advanced Routing > IP Broadcast Helper** page of the WebUI.
2. Use the **Forward Non-local Packets** option to control whether to forward packets that are not locally originated by a source directly on the receiving interface.
3. In the **Configure Relays** section, click **Add**.
The **Add Relay** window opens.
4. Select the **Interface** to which you want to add support for IP helper services.
5. Add a **UDP Port** number to the helper services.
6. Add a **Relay** (server) IPv4 address for the UDP port.



Note - To enable IP Broadcast Helper in a cluster, do the configuration on each cluster member

IP Broadcast helper configuration parameters

Parameter	Description
Forward Non-local Packets	Controls whether packets will be forwarded that are not locally originated by a source directly on the receiving interface. Enable to forward packets even if the source is not directly on the receiving interface. Clear the option to require that packets are generated by a source that is directly on the receiving interface to be eligible for relay. • Default: Cleared.
Interface	The interface on which the IP helper service runs. • Default: None

Parameter	Description
UDP Port	The UDP service to be forwarded by the interface. Client UDP packets with the UDP port number are forwarded to the relay. • Range: 0-65535. • Default: None.
Relay	The IP address of the server defined for forwarding for the interface and UDP service. • Default: None

Configuring IP Broadcast Helper - CLI (iphelper)

Forwarding Non-Local Packets

Use the following commands to control whether to forward packets that are not locally originated by a source directly on the receiving interface.

```
set iphelper forward-nonlocal <on | off>
```

Parameter	Description
<code>forward-nonlocal</code> <code><on off></code>	Select on to forward packets even if the source is not directly on the receiving interface. Select off to require that packets are generated by a source that is directly on the receiving interface to be eligible for relay . • Default: off

IP Broadcast Helper interfaces

Use the following commands configure IP Broadcast Helper properties for specific interfaces.

```
set iphelper interface if_name
off
udp-port <1-65535> off
udp-port <1-65535> relay-to ip_address <on | off>
```

Parameter	Description
<code>interface <if_name></code> <code>off</code>	Disable the interface configured for iphelper
<code>udp-port <1-65535></code> <code>off</code>	Disable the UDP services configured for this interface.
<code>udp-port <1-65535></code> <code>relay-to ip_address</code> <code><on off></code>	The UDP service to be forwarded by the interface. Client UDP packets with the specified UDP port number are forwarded to the configured relay. The IP address of the server defined for forwarding for the interface and UDP service.



Note - To enable IP Broadcast Helper in a cluster, do the configuration on each cluster member.

Monitoring IP Broadcast Helper

To monitor and troubleshoot IP Broadcast Helper in the WebUI:

1. Open the **Advanced Routing > IP Broadcast Helper** page of the WebUI.
2. Click the **Monitoring** tab.



Note - The page is static. To see the latest values, click **Reload**.

To monitor and troubleshoot IP Broadcast Helper in clish:

```
show iphelper services
show iphelper stats
```

RIP

In This Section:

RIP 2	77
RIP 1	78
Virtual IP Address Support for VRRP	78
Configuring RIP - WebUI	78
Configuring RIP - CLI (rip).....	81
Monitoring RIP	84

The *Routing Information Protocol (RIP)* is one of the oldest, and still widely used, interior gateway protocols (IGP). RIP uses only the number of hops between nodes to determine the cost of a route to a destination network and does not consider network congestion or link speed. Other shortcomings of RIP are that it can create excessive network traffic if there are a large number of routes and that it has a slow convergence time and is less secure than other IGPs, such as OSPF.

Routers using RIP broadcast their routing tables on a periodic basis to other routers, whether or not the tables have changed. Each update contains paired values consisting of an IP network address and a distance to that network. The distance is expressed as an integer, the *hop count metric*. Directly connected networks have a metric of 1. Networks reachable through one other router are two hops, and so on. The maximum number of hops in a RIP network is 15 and the protocol treats anything equal to or greater than 16 as unreachable.

RIP 2

The RIP version 2 protocol adds capabilities to RIP. Some of the most notable RIP 2 enhancements follow.

Network Mask

The RIP 1 protocol assumes that all subnetworks of a given network have the same network mask. It uses this assumption to calculate the network masks for all routes received. This assumption prevents subnets with different network masks from being included in RIP packets. RIP 2 adds the ability to explicitly specify the network mask for each network in a packet.

Authentication

RIP 2 packets also can contain one of two types of authentication methods that can be used to verify the validity of the supplied routing data.

The first method is a simple password in which an authentication key of up to 16 characters is included in the packet. If this password does not match what is expected, the packet is discarded. This method provides very little security, as it is possible to learn the authentication key by watching RIP packets.

The second method uses the MD5 algorithm to create a crypto checksum of a RIP packet and an authentication key of up to 16 characters. The transmitted packet does not contain the authentication key itself; instead, it contains a crypto-checksum called the *digest*. The receiving router performs a calculation using the correct authentication key and discards the packet if the digest does not match. In addition, a sequence number is maintained to prevent the replay of older

packets. This method provides stronger assurance that routing data originated from a router with a valid authentication key.

RIP 1

Network Mask

RIP 1 derives the network mask of received networks and hosts from the network mask of the interface from which the packet was received. If a received network or host is on the same natural network as the interface over which it was received, and that network is subnetted (the specified mask is more specific than the natural network mask), then the subnet mask is applied to the destination. If bits outside the mask are set, it is assumed to be a host; otherwise, it is assumed to be a subnet.

Auto Summarization

The Check Point implementation of RIP 1 supports auto summarization; this allows the router to aggregate and redistribute nonclassful routes in RIP 1.

Virtual IP Address Support for VRRP

Gaia supports the advertising of the virtual IP address of the VRRP Virtual Router. You can configure RIP to advertise the virtual IP address rather than the actual IP address of the interface ("[Configuring RIP - WebUI](#)" on page 78). If you enable this option, RIP runs only on the master of the Virtual Router; on a failover, RIP stops running on the old master and then starts running on the new master. A traffic break might occur during the time it takes both the VRRP and RIP protocols to learn the routes again. The larger the network, the more time it would take RIP to synchronize its database and install routes again.



Note -

Gaia also provides support for BGP, OSPF, and PIM, both Sparse-Mode and Dense-Mode, to advertise the virtual IP address of the VRRP Virtual Router.

You must use Monitored Circuit mode when configuring virtual IP support for any dynamic routing protocol, including RIP.

Configuring RIP - WebUI

To configure RIP:

1. In the **Network Management > Network interfaces** page of the WebUI, configure Ethernet Interfaces and assign an IP address to the interface.
2. Open the **Advanced Routing > RIP** page of the WebUI.
3. **Optional:** In the **RIP Global Settings** section ("[RIP Global Settings](#)" on page 79):
 - a) Configure the RIP **Update Interval** and **Expire Interval**. These timers allows you to vary the frequency with which updates are sent and when routes expire.
 - b) Select **Auto Summarization** to aggregate and redistribute non-classful routes in RIP 1. Clear it to disable the option.

4. In the **RIP Interfaces** section, click **Add**.

The **Add Interface** window opens

5. Configure the **RIP Interfaces** ("**RIP Interface Options**" on page 79).

6. Click **Save**.

RIP Global Settings

Option	Description
Update Interval	The amount of time, in seconds, between regularly scheduled RIP updates. To prevent synchronization of periodic updates, RIP updates are actually sent at a time from the uniform distribution on the interval $(0.5T, 1.5T)$ where T is the Update Interval value. Note - Be careful when you set this parameter, because RIP has no protocol mechanism to detect misconfiguration. • Range: 1-65535. • Default: 30.
Expire Interval	The amount of time, in seconds, that must pass with no update for a given route before the route is considered to have timed out. This value must be 6 times the update interval before the network drops packets which contain an update. • Range: 1-65535. • Default: 180.
Auto Summarization	Automatically aggregates and redistributes non-classful RIP Version 1 into RIP. This applies only to RIP Version 1. If you do not select the Auto summarization field option, you must use route aggregation and route redistribution and do the aggregation and redistribution manually. Note - Be careful when you set this parameter, because RIP has no protocol mechanism to detect misconfiguration. • Default: Selected.

RIP Interface Options

Option	Description
Interface	The interface on which RIP is enabled.
Version	The version of RIP to run. If you enter version 2, the default is to send full version 2 packets on the RIP multicast address. • Options: 1 or 2. • Default: 1.

Option	Description
Metric	The RIP metric to add to routes that are sent with the specified interface(s). The default is zero. This is used to make other routers prefer other sources of RIP routes over this router. • Range: 0-16. • Default: 0.
Accept updates	Defines if RIP packets from other routers which use the interface are accepted or ignored. Ignoring an update may result in suboptimal routing. • Default: Selected.
Send updates	Defines if RIP packets are sent through the interface. This causes the interface to be a passive RIP listener. • Default: Selected
Virtual Address	Make RIP run only on the VRRP Virtual IP address related to this interface. If this router is not a VRRP Master then RIP does not run if this option is selected. It only runs on the VRRP Master. Note - You must use Monitored Circuit mode when you configure VRRP to work with virtual IPs, and when you configure virtual IP support for a dynamic routing protocol, including RIP. • Default: Cleared.
Transport	Selecting Multicast specifies that RIP version 2 packets should be multicast on this interface. This is the default. • Selecting Broadcast specifies that RIP version 1 packets that are compatible with version 2 should be broadcast on this interface. • Options: Broadcast/Multicast. • Default: Multicast.

Option	Description
Authentication Type	The type of authentication scheme to use for the link. This option applies to rip version 2 only. In general, routers on a given link must agree on the authentication configuration in order to form neighbor adjacencies. This is used to guarantee that routing information is accepted only from trusted routers. • None: There is no authentication scheme for the interface to accept routing information from neighboring routers. • Simple: Implement a simple authentication scheme for the interface to accept routing information from neighboring routers. Enter the Simple Password, from 1 to 16 characters. Must contain alphanumeric characters only. • MD5: Implement an authentication scheme that uses an MD5 algorithm for the interface to accept routing information from neighboring routers. Enter the password. To ensure interoperability with Cisco routers running RIP MD5 authentication, enable Cisco Compatibility. By default, RIP MD5 is set to conform to the Check Point standard, and not for Cisco compatibility. • Options: None/Simple/MD5. • Default: None.

Configuring RIP - CLI (rip)

RIP Global Commands

Use these commands to configure RIP properties that apply to all interfaces configured for RIP.

```
set rip
  auto-summary <on | off>
  update-interval <1-65535>
  update-interval default
  expire-interval <1-65535>
  expire-interval default
```

Parameter	Description
auto-summary <on off>	Automatically aggregates and redistributes non-classful RIP Version 1 into RIP. This applies only to RIP Version 1. If the Auto summarization field option is unchecked, you must do the aggregation and redistribution manually by using route aggregation and route redistribution. Note - Take care when you set this parameter, as RIP has no protocol mechanism to detect misconfiguration. Default: on

Parameter	Description
update-interval <1-65535>	The amount of time, in seconds, between regularly scheduled RIP updates. To prevent synchronization of periodic updates, RIP updates are actually sent at a time from the uniform distribution on the interval (0.5T, 1.5T) where T corresponds to the Update Interval value. Note - Take care when you set this parameter, as RIP has no protocol mechanism to detect misconfiguration.
update-interval default	A value of 30 seconds.
expire-interval <1-65535>	The amount of time, in seconds, that must pass without receiving an update for a given route before the route is considered to have timed out. This value should be 6 times the update interval in order to allow for the possibility that packets containing an update could be dropped by the network.
expire-interval default	A value of 180 seconds.

RIP Interface Commands

Use these commands to configure RIP properties that apply to a RIP interface.

```
set rip interface if_name
    <off | on>
    version <1 | 2> on
    metric <0-16>
    metric default
    accept-updates <on | off>
    send-updates <on | off>
    transport <multicast | broadcast>
    authtype none
    authtype simple password
    authtype md5 secret secret [cisco-compatibility] <on | off>
    virtual address <on | off>
```

Parameter	Description
interface if_name <off on>	Turn on or turn off RIP on the interface. Default: off
<1 2>	The version of RIP to run. If you specify version 2, the default is to send full version 2 packets on the RIP multicast address. Default: 1
metric <0-16>	The RIP metric to be added to routes that are sent using the specified interface(s). The default is zero. This is used to make other routers prefer other sources of RIP routes over this router.
metric default	A value of 0.

Parameter	Description
<code>accept-updates</code> <code><on off></code>	Whether RIP packets from other routers using the interface are accepted or ignored. Ignoring an update may result in suboptimal routing. Default: off
<code>send-updates</code> <code><on off></code>	Whether RIP packets should be sent via the interface. This causes the interface to be a passive RIP listener.
<code>transport</code> <code><multicast broadcast></code>	The transport mechanism. Selecting Multicast specifies that RIP version 2 packets should be multicast on this interface. This is the default. Note - When you use RIP 2, always select multicast. We recommend that you do not operate RIP 1 and RIP 2 together. Selecting Broadcast specifies that RIP version 1 packets that are compatible with version 2 should be broadcast on this interface.
<code>authtype none</code>	There is no authentication scheme for the interface to accept routing information from neighboring routers. This option applies to rip version 2 only. In general, routers on a given link must agree on the authentication configuration in order to form neighbor adjacencies. This is used to guarantee that routing information is accepted only from trusted routers.
<code>authtype simple</code> <code>password</code>	Implement a simple authentication scheme for the interface to accept routing information from neighboring routers. Enter the Simple Password, from 1 to 16 characters. Must contain alphanumeric characters only. This option applies to RIP version 2 only.
<code>authtype md5</code> <code>secret</code> <code>secret</code>	Implement an authentication scheme that uses an MD5 algorithm for the interface to accept routing information from neighboring routers. Enter the password.
<code>interface if_name</code> <code>virtual</code> <code><on off></code>	Make RIP run only on the VRRP Virtual IP address associated with this interface. If this router is not a VRRP Master then RIP will not run if this option is selected. It will only run on the VRRP Master. Note - You must use Monitored Circuit mode when configuring VRRP to work with virtual IPs, and when configuring virtual IP support for any dynamic routing protocol, including RIP. For more information, see <i>ICMP Router Discovery</i> (" Router Discovery " on page 186). Default: off

Parameter	Description
cisco-compatibility <on off>	To ensure interoperability with Cisco routers running RIP MD5 authentication, enable Cisco Compatibility. By default, RIP MD5 is set to conform to the Check Point standard, and not for Cisco compatibility. Default: off

Monitoring RIP

Monitoring RIP - WebUI

To monitor and troubleshoot RIP:

1. Open the **Advanced Routing > RIP** page of the WebUI.
2. Click the **Monitoring** tab.
3. In the **Information** table, click a line to see the current values.



Note - The page is static. To see the latest values, reload your browser page.

RIP Show Commands

Use these commands to monitor and troubleshoot RIP.

```
show rip
```

```
show rip
  interfaces
  interface <if_name>
  packets
  errors
  neighbors
  summary
```

OSPF

In This Section:

Types of Areas.....	85
Area Border Routers	86
High Availability Support for OSPF	86
OSPF Forced Hellos.....	87
Configuring OSPF - WebUI.....	87
Configuring OSPF - CLI (ospf).....	96

Open Shortest Path First (OSPF) is an interior gateway protocol (IGP) used to exchange routing information between routers within a single autonomous system (AS). OSPF calculates the best path based on true costs using a metric assigned by a network administrator. RIP, the oldest IGP protocol chooses the least-cost path based on hop count. OSPF is more efficient than RIP, has a quicker convergence, and provides equal-cost multipath routing where packets to a single destination can be sent using more than one interface. OSPF is suitable for complex networks with a large number of routers. It can coexist with RIP on a network.

Gaia supports OSPFv2, which supports IPv4 addressing.

You can run OSPF over a route-based VPN by enabling OSPF on a virtual tunnel interface (VTI).

Types of Areas

Routers using OSPF send packets called Link State Advertisements (LSA) to all routers in an area. *Areas* are smaller groups within the AS that you can design to limit the flooding of an LSA to all routers. LSAs do not leave the area from which they originated, thus increasing efficiency and saving network bandwidth.

You must specify at least one area in your OSPF network—the *backbone area*, which has the responsibility to propagate information between areas. The backbone area has the identifier 0.0.0.0.

You can designate other areas, depending on your network design, of the following types:

- **Normal Area** — Allows all LSAs to pass through. The backbone is always a normal area.
- **Stub Area** — Stub areas do not allow Type 5 LSAs to be propagated into or throughout the area and instead depends on default routing to external destinations. You can configure an area as a stub to reduce the number of entries in the routing table (routes external to the OSPF domain are not added to the routing table).
- **NSSA (Not So Stubby Area)** — Allows the import of external routes in a limited fashion using Type-7 LSAs. NSSA border routers translate selected Type 7 LSAs into Type 5 LSAs which can then be flooded to all Type-5 capable areas. Configure an area as an NSSA if you want to reduce the size of the routing table, but still want to allow routes that are redistributed to OSPF.

Best Practice - It is generally recommended that you limit OSPF areas to about 50 routers based on the limitations of OSPF (traffic overhead, table size, convergence, and so on).

All OSPF areas must be connected to the backbone area. If you have an area that is not connected to the backbone area, you can connect it by configuring a *virtual link*, enabling the backbone area to appear contiguous despite the physical reality.



Note - If you need to connect two networks that both already have backbone areas and you do not want to reconfigure one to something other than 0.0.0.0, you can connect the two backbone areas using a virtual link.

Each router records information about its interfaces when it initializes and builds an LSA packet. The LSA contains a list of all recently seen routers and their costs. The LSA is forwarded only within the area it originated in and is flooded to all other routers in the area. The information is stored in the link-state database, which is identical on all routers in the AS.

Area Border Routers

Routers called *Area Border Routers* (ABR) have interfaces to multiple areas. ABRs compact the topological information for an area and transmit it to the backbone area. Check Point supports the implementation of ABR behavior as outlined in the Internet draft of the Internet Engineering Task Force (IETF). The definition of an ABR in the OSPF specification as outlined in RFC 2328 does not require a router with multiple attached areas to have a backbone connection. However, under this definition, any traffic destined for areas that are not connected to an ABR or that are outside the OSPF domain is dropped. According to the Internet draft, a router is considered to be an ABR if it has more than one area actively attached and one of them is the backbone area. An area is considered actively attached if the router has at least one interface in that area that is not down.

Rather than redefine an ABR, the Check Point implementation includes in its routing calculation summary LSAs from all actively attached areas if the ABR does not have an active backbone connection, which means that the backbone is actively attached and includes at least one fully adjacent neighbor. You do not need to configure this feature; it functions automatically under certain topographies.

OSPF uses the following types of routes:

- **Intra-area**—Have destinations within the same area.
- **Interarea**—Have destinations in other OSPF areas.
- **Autonomous system external (ASE)**—Have destinations external to the autonomous system (AS). These are the routes calculated from Type 5 LSAs.
- **NSSA ASE Router**—Have destinations external to AS. These are the routes calculated from Type 7 LSAs.

All routers on a link must agree on the configuration parameters of the link. All routers in an area must agree on the configuration parameters of the area. A separate copy of the SPF algorithm is run for each area. Misconfigurations prevent adjacencies from forming between neighbors, and routing black holes or loops can form.

High Availability Support for OSPF

Gaia supports the OSPF protocol in clusters configured either via VRRP or ClusterXL.

In this configuration, the cluster becomes a Virtual Router. The neighbor routers see it as a single router, where the virtual IP address of the cluster becomes the router ID. Each member of the cluster runs the OSPF process, but only the master actively exchanges routing information with

the neighbor routers. When a failover occurs, a standby member of the cluster becomes the master and begins exchanging routing information with the neighbor routers.

Gaia also supports the OSPF protocol over VPN tunnels which terminate in the VRRP or ClusterXL cluster.

VRRP

Gaia supports advertising of the virtual IP address of the VRRP Virtual Router instead of the actual interface IP address. If you enable this option, but do not enable Graceful Restart, OSPF runs only on the master of the Virtual Router. During a failover, a traffic break may occur, while the new VRRP master becomes active and learns the OSPF routes. This happens because the OSPF route database exists only on the master and is not synchronized on all members of the cluster. The larger the network, the more time it takes OSPF to synchronize its database and install routes again.

To avoid traffic loss during failovers, you can configure Graceful Restart. In this case, the master synchronizes the route table with the VRRP backup member. If the master fails, the backup member takes on a role of the new master, sends grace-LSAs to the OSPF neighbors, and establishes adjacencies with them. The new master keeps the kernel routes that were installed before the failover until it establishing full adjacency with the neighbors.



Note - You must use monitored-circuit VRRP, not VRRP v2, when configuring virtual IP support for OSPF or any other dynamic routing protocol.

ClusterXL

Gaia ClusterXL advertises the virtual IP address of the ClusterXL Virtual Router. The OSPF routes database of the master is synchronized across all members of the cluster. The OSPF task of each standby member obtains routing state and information from the master and installs the routes in the kernel as the master does. On a failover, one of the standby members becomes the new master and then continues where the old master failed. During the time that the new master resynchronizes routes database with the neighbor routers, traffic forwarding continues using the old kernel routes until OSPF routes are fully synchronized and pushed into the kernel.

OSPF Forced Hellos

When OSPF is configured with a low dead interval or too many OSPF neighbors or OSPF routes, routers can become too busy to send the OSPF hello packets on time. This can cause OSPF dead timers to expire on neighbors and cause outages. To prevent this behavior, the OSPF Forced Hellos feature was introduced. With this feature enabled, OSPF sends out hello packets at specified intervals when it processes updates or synchronizes routes. These hello packets are in addition to the regular hello packets in OSPF.

Configuring OSPF - WebUI

To configure OSPF:

1. In the **Network Management > Network interfaces** page of the WebUI, configure Ethernet Interfaces and assign an IP address to the interface.
2. Open the **Advanced Routing > OSPF** page of the WebUI.

3. Define other Global settings ("Configuring Global Settings " on page 88), including the **Router ID**.
4. Optional: Define additional OSPF areas ("Configuring OSPF Areas" on page 93) (in addition to the backbone area).
5. Optional: For each area, you can add one or more address ranges if you want to reduce the number of routing entries that the area advertises into the backbone.


Note - To prevent an address range from being advertised into the backbone, select **Restrict** for the address range
6. Configure OSPF Interfaces ("Configuring OSPF Interfaces" on page 89).
7. Configure virtual links ("Configuring OSPF Virtual Links" on page 92) for any area that does not connect directly to the backbone area.

Configuring Global Settings

The following table shows the global settings that you can specify for OSPF. Configure these settings by clicking OSPF under Configuration > Routing Configuration in the tree view and scrolling down to these fields.

Global Settings for OSPF

Parameter	Description
Router ID	The Router ID uniquely identifies the router in the autonomous system. The router ID is used by the BGP and OSPF protocols. We recommend setting the router ID rather than relying on the default setting. This prevents the router ID from changing if the interface used for the router ID goes down. Use an address on a loopback interface that is not the loopback address (127.0.0.1). Note - In a cluster, you must select a router ID and make sure that it is the same on all cluster members. • Range: Dotted-quad.([0-255].[0-255].[0-255].[0-255]). Do not use 0.0.0.0 • Default: The interface address of one of the local interfaces.
RFC1583 Compatibility	This implementation of OSPF is based on RFC2178, which fixed some looping problems in an earlier specification of OSPF. If your implementation is running in an environment with OSPF implementations based on RFC1583 or earlier, enable RFC 1583 compatibility to ensure backwards compatibility. • Default: Selected
SPF Delay	Specifies the time in seconds the system will wait to recalculate the OSPF routing table after a change in topology. • Default: 2. • Range: 1-60.

Parameter	Description
SPF Hold Time	Specifies the minimum time in seconds between recalculations of the OSPF routing table. • Default:5. • Range: 1-60.
Default ASE Route Cost	Specifies a cost for routes redistributed into OSPF as ASEs. Any cost previously assigned to a redistributed route overrides this value.
Default ASE Route Type	Specifies a route type for routes redistributed into OSPF as ASEs, unless these routes already have a type assigned. There are two types: • Type 1 external: Used for routes imported into OSPF which are from IGPs whose metrics are directly comparable to OSPF metrics. When a routing decision is being made, OSPF adds the internal cost to the AS border router to the external metric. • Type 2 external: Used for routes whose metrics are not comparable to OSPF internal metrics. In this case, only the external OSPF cost is used. In the event of ties, the least cost to an AS border router is used.

Configuring OSPF Interfaces

To configure an OSPF interface:

1. In the **Edit Interface** window, assign the appropriate **Area** to each interface by selecting the OSPF area that this interface participates in.

The OSPF interface configuration parameters are displayed showing the default settings. If you want to accept the default settings for the interface, no further action is necessary.

2. (Optional) Change any configuration parameters for the interface.



Note - The hello interval, dead interval, and authentication method must be the same for all routers on the link.

Configuration Parameters for OSPF Interfaces

Parameter	Description
Area	The drop-down list displays all of the areas configured and enabled on your platform. An entry for the backbone area is displayed even if it is disabled. An OSPF area defines a group of routers running OSPF that have the complete topology information of the given area. OSPF areas use an area border router (ABR) to exchange information about routes. Routes for a given area are summarized into the backbone area for distribution into other non-backbone areas. An ABR must have at least two interfaces in at least two different areas. For information on adding an area Configuring OSPF Areas and Global Settings.
Hello interval	Specifies the length of time in seconds between hello packets that the router sends on this interface. For a given link, this value must be the same on all routers, or adjacencies do not form. • Range: 1-65535 in seconds • Default: For broadcast interfaces, the default hello interval is 10 seconds. For point-to-point interfaces, the default hello interval is 30 seconds.
Dead interval	Specifies the number of seconds after the router stops receiving hello packets that it declares the neighbor is down. • Recommended value: Four times the hello interval. For a given link, this value must be the same on all routers, or adjacencies do not form. The value must not be 0. • Range: 1-65535 in seconds. • Default: For broadcast interfaces, the default dead interval is 40 seconds. For point-to-point interfaces, the default dead interval is 120 seconds.
Retransmit interval	Specifies the number of seconds between LSA retransmissions for this interface. This value is also used when retransmitting database description and link state request packets. Set this value well above the expected round-trip delay between any two routers on the attached network. Be conservative when setting this value to prevent necessary retransmissions. • Range: 1-65535 in seconds. • Default: 5.
OSPF cost	Specifies the weight of a given path in a route. The higher the cost you configure, the less preferred the link as an OSPF route. For example, you can assign different relative costs to two interfaces to make one more preferred as a routing path. You can explicitly override this value in route redistribution. • Range: 1-65535. • Default: 1.

Parameter	Description
Election priority	Specifies the priority for becoming the designated router (DR) on this link. When two routers attached to a network both attempt to become a designated router, the one with the highest priority wins. If there is a current DR on the link, it remains the DR regardless of the configured priority. This feature prevents the DR from changing too often and applies only to a shared-media interface, such as Ethernet. A DR is not elected on point-to-point type interfaces. A router with priority 0 is not eligible to become the DR. • Range: 0-255. • Default: 1.
Passive	Specifies that the interface does not send hello packets, which means that the link does not form any adjacencies. This mode enables the network associated with the interface to be included in the intra-area route calculation rather than redistributing the network into OSPF and having it as an ASE. In passive mode, all interface configuration information, with the exception of the associated area and the cost, is ignored. • Options: On or Off. • Default: Off.
Use Virtual Address	Makes OSPF run only on the VRRP Virtual IP address associated with this interface. If this router is not a VRRP master, then OSPF will not run if this option is On. It will only run on the VRRP master. For more information, see Configuring Monitored-Circuit VRRP. • Options: On or Off. • Default: Off

Parameter	Description
Authorization Type	Specifies which type of authentication scheme to use for a given link. In general, routers on a given link must agree on the authentication configuration to form neighbor adjacencies. This feature guarantees that routing information is accepted only from trusted routers. Options are: • None: Does not authenticate packets. This is the default option. • Simple: Uses a key of up to eight characters. Provides little protection because the key is sent in the clear, and it is possible to capture packets from the network and learn the authentication key. • MD5: Uses a key of up to 16 characters. Provides much stronger protection, as it does not include the authentication key in the packet. Instead, it provides a cryptographic hash based on the configured key. The MD5 algorithm creates a crypto checksum of an OSPF packet and an authentication key of up to 16 characters. The receiving router performs a calculation using the correct authentication key and discards the packet if the key does not match. In addition, a sequence number is maintained to prevent the replay of older packets.

Configuring OSPF Virtual Links

You must configure a virtual link for any area that does not connect directly to the backbone area. You configure the virtual link on both the ABR for the discontinuous area and another ABR that does connect to the backbone.

The virtual link acts like a point-to-point link. The routing protocol traffic that flows along the virtual link uses intra-area routing only.

To configure a virtual link:

1. Create a Normal Type area (which does not connect directly to the backbone area) and configure an interface to be in that area.
2. In the **Virtual Links** section, click **Add**.
3. In the **Add Virtual Link** window, enter the **Remote Router ID** of the remote endpoint of the virtual link.
4. Select the **Transit Area**. This is the area that connects both to the backbone and to the discontinuous area.
5. Configure the following parameters for the virtual link:
 - **Hello interval**—Length of time, in seconds, between hello packets that the router sends on the interface. For a given link, this field must be the same on all routers or adjacencies do not form.
 - **Default:** 30.
 - **Dead Interval**—Number of seconds after the router stops receiving hello packets that it declares the neighbor is down. Typically, the value of this field should be four times that of

the hello interval. For a given link, this value must be the same on all routers, or adjacencies do not form. The value must not be zero.

- **Range:** 1-65535.
 - **Default:** 120.
 - **Retransmit Interval**—Specifies the number of seconds between LSA retransmissions for adjacencies belonging to this interface. This value is also used when retransmitting database description and link state request packets. Set this value well above the expected round-trip delay between any two routers on the attached network. Be conservative when setting this value to prevent unnecessary retransmissions.
 - **Range:** 1-65535 in number of seconds.
 - **Default:** 5.
 - **Auth Type**—Type of authentication scheme to use for a given link. In general, routers on a given link must agree on the authentication configuration to form neighbor adjacencies. This feature guarantees that routing information is accepted only from trusted routers.
 - **Options:** None / Simple / MD5.
 - **Default:** None.
6. If you selected MD5 for the auth type, you must also configure the following parameters:
- **Add MD5 Key**— If the Auth type selected is MD5, the **MD5 List** appears. Click Add and specify the MD5 Key ID and its corresponding MD5 key. If you configure multiple Key IDs, the Key ID with the highest value is used to authenticate outgoing packets. All keys can be used to authenticate incoming packets.
 - **Key ID**—The Key ID is included in the outgoing OSPF packets to enable the receivers to use the appropriate MD5 secret to authenticate the packet.
 - **Range:** 0-255.
 - **Default:** None
 - **MD5 Secret**—The MD5 secret is included in encrypted form in outgoing packets to authenticate the packet. Range: 1-16 alphanumeric characters. Default: None
7. Repeat this procedure on both the ABR for the discontinuous area and an ABR that connects to the backbone area.

Configuring OSPF Areas

This table lists the parameters for areas and global settings that you use when configuring OSPF on your system. As you add areas, each is displayed with its own configuration parameters under the Areas section.

- **Area:** Choose an IPv4 address (preferred) or an integer.

OSPF Normal Type Area Configuration Parameters

Parameter	Description
Add Address Range	You can configure any area with any number of address ranges. Use these ranges to reduce the number of routing entries that a given area emits into the backbone and thus all areas. If a given IPv4 address aggregates a number of more specific IPv4 addresses within an area, you can configure an address that becomes the only IPv4 address advertised into the backbone. You must be careful when configuring an address range that covers parts of an IPv4 address not contained within the area. By definition, an address range consists of a IPv4 address and a mask length. Note: To prevent a specific IPv4 address from being advertised into the backbone, select Restrict.
Add Stub Network	OSPF can advertise reachability to IPv4 addresses that are not running OSPF using a stub network. The advertised IPv4 address appears as an OSPF internal route and can be filtered at area borders with the OSPF area ranges. The IPv4 address must be directly reachable on the router where the stub network is configured; that is, one of the router's interface addresses must fall within the IPv4 address to be included in the router-LSA. You configure stub hosts by specifying a mask length of 32. This feature also supports advertising an IPv4 address and mask that can be activated by the local address of a point-to-point interface. To advertise reachability to such an address, enter an IP address and a cost with a value other than zero.
Area Type	For descriptions of area types, see Types of Areas (on page 85). Options: Normal/Stub/NSSA.

Stub Area Parameters

The following table stub areas configuration parameters appear if you define the area as a stub area.

Parameter	Description
Cost for Default Route	Enter a cost for the default route to the stub area. • Range: 1-16777215. • Default: No default.
Import Summary Routes	Specifies if summary routes (summary link advertisements) are imported into the stub area or NSSA. Each summary link advertisement describes a route to a destination outside the area, yet still inside the AS (i.e. an inter-area route). These include routes to networks and routes to AS boundary routers. • Default: Selected.

NSSA (Not So Stubby Area) Parameters

The following table describes the configuration parameters for NSSA areas. These fields appear if you define the area as an NSSA (Not So Stubby Area). For more information on NSSA, see RFC 3101.

Parameter	Description
Translator Role	Specifies whether this NSSA border router will unconditionally translate Type-7 LSAs into Type-5 LSAs. When role is Always, Type-7 LSAs are translated into Type-5 LSAs regardless of the translator state of other NSSA border routers. When role is Candidate, this router participates in the translator election to determine if it will perform the translations duties. If this NSSA router is not a border router, then this option has no effect. • Default: Candidate.
Translator Stability Interval	Specifies how long in seconds this elected Type-7 translator will continue to perform its translator duties once it has determined that its translator status has been assumed by another NSSA border router. This field appears only if an area is defined as an NSSA with translator role as Candidate. • Default: 40 seconds.
Import Summary Routes	Specifies if summary routes (summary link advertisements) are imported into the stub area or NSSA. Each summary link advertisement describes a route to a destination outside the area, yet still inside the AS (i.e. an inter-area route). These include routes to networks and routes to AS boundary routers. • Default: On.
Cost for Default Route	Enter a cost associated with the default route to the NSSA.
Default Route Type	Specifies the route type associated with the Type-7 default route for an NSSA when routes from other protocols are redistributed into OSPF as ASEs. If a redistributed route already has a route type, this type is maintained. If summary routes are imported into an NSSA, only then a Type-7 default route is generated (otherwise a Type-3 default route is generated). This field appears only if an area is defined as an NSSA into which summary routes are imported. The route type can be either 1 or 2. A type 1 route is internal and its metric can be used directly by OSPF for comparison. A type 2 route is external and its metric cannot be used for comparison directly. • Default: 1
Redistribution	Specifies if both Type-5 and Type-7 LSAs or only Type-7 LSAs will be originated by this router. This option will have effect only if this router is an NSSA border router and this router is an AS border router. • Default: On

Parameter	Description
Type 7 Address Ranges	An NSSA border router that performs translation duties translates Type-7 LSAs to Type-5 LSAs. An NSSA border router can be configured with Type-7 address ranges. Use these ranges to reduce the number of Type-5 LSAs. Many separate Type-7 networks may fall into a single Type-7 address range. These Type-7 networks are aggregated and a single Type-5 LSA is advertised. By definition, a Type-7 address range consists of a prefix and a mask length. Note - To prevent a specific prefix from being advertised, select On in the Restrict field next to the entry for that prefix.

Configuring OSPF - CLI (ospf)

Use the following group of commands to set and view parameters for OSPF. This syntax is shown below for each set of commands.



Note - Gaia does not have CLI commands for route filtering and redistribution. You must configure inbound routing policies and redistribution of routes through the WebUI. You can configure route maps and route aggregation using CLI commands. Route map configuration done through the CLI takes precedence over route filtering and redistribution configured in the WebUI. For example if OSPF uses route maps for inbound filtering, anything configured on the WebUI page for inbound route filters for OSPF is ignored. You can still use the WebUI to configure route redistribution into OSPF.

When you do initial configuration, set the router ID. Use these commands:

```
set router-id
  default
  ip_address
```

Parameter	Description
router-id default	Selects the highest interface address when OSPF is enabled.

router-id ip_address	Specifies a specific IP address to assign as the router ID. Do not use 0.0.0.0 as the router ID address. Best Practice - Check Point recommends setting the router ID rather than relying on the default setting. Setting the router ID prevents the ID from changing if the default interface used for the router ID goes down. The Router ID uniquely identifies the router in the autonomous system. The router ID is used by the BGP and OSPF protocols. We recommend setting the router ID rather than relying on the default setting. This prevents the router ID from changing if the interface used for the router ID goes down. Use an address on a loopback interface that is not the loopback address (127.0.0.1). Note - In a cluster, you must select a router ID and make sure that it is the same on all cluster members. • Range: Dotted-quad.([0-255].[0-255].[0-255].[0-255]). Do not use 0.0.0.0 • Default: The interface address of one of the local interfaces.
----------------------	--

OSPF Global Settings

Global settings apply to all configured OSPF areas, including the backbone and stub areas.

To configure global settings:

Run the set ospf command with these options:

```
set ospf {rfc1583-compatibility {on | off} | spf-delay {default | <delay>} |
spf-holdtime {default | <holdtime>} | default-ase-cost <cost> |
default-ase-type {1 | 2} | force-hellos {on | off} | force-hellos timer
{default | <2-10>} | graceful-restart-helper {on | off} | graceful-restart
{on | off | grace-period <seconds>}}
```

Parameter	Description
rfc1583-compatibility {on off}	Ensure backward compatibility. This option is on by default.
spf-delay {default <delay>}	Specify the <delay> value, in seconds, to wait before recalculating the OSPF routing table after a change in the topology. The default is 2 seconds.
spf-holdtime {default <holdtime>}	Specify the minimum <holdtime>, in seconds, between recalculations of the OSPF routing table. The default is 5 seconds.
default-ase-cost <cost>	Specify the cost assigned to routes from other protocols that are redistributed into OSPF as autonomous systems external. If the route has a cost already specified, that cost takes precedent. Valid cost values are between 1 and 6777215.

Parameter	Description
<code>default-ase-type {1 2}</code>	Specify the type assigned to routes from other protocols that are redistributed into OSPF as autonomous systems external. If the route has a type already specified, that type takes precedent. Valid type values are 1 or 2.
<code>force-hellos</code>	In addition to OSPF regular hello packets, OSPF sends out hello packets at specified intervals when it processes updates or synchronizes routes. Default: Off
<code>force-hellos timer</code>	The time in seconds between one forced hello message to the next. Value: 2-10 Default: 5
<code>graceful-restart-helper {on off}</code>	Specify whether the Check Point system should maintain the forwarding state advertised by peer routers, even when they restart, to minimize the negative effects caused by peer routers restarting.
<code>graceful-restart {on off grace-period <seconds>}</code>	Configure Graceful Restart - turn it on, turn it off, or set the grace period to a value between 1 and 1800 seconds. The default grace period is 120 seconds.

OSPF Interfaces

Use these commands to configure a backbone and other areas, such as stub areas, for specified interfaces.

For OSPFv2 use the following commands:

```
set ospf
  area <backbone | ospf_area> range ip_prefix <on | off>
  area <backbone | ospf_area> range ip_prefix restrict <on | off>
  stub-network ip_prefix <on | off>
  stub-network ip_prefix stub-network-cost <1-677722>

set ospf interface if_name
  area <backbone | ospf_area> <on | off>
  hello-interval <1-65535>
  hello-interval default
  dead-interval <1-65535>
  dead-interval default
  retransmit-interval <1-65535>
  retransmit-interval default
  cost <1-65535>
  priority <0-255>
  passive <on | off>
  virtual-address <on | off>
  authtype none
  simple password
  md5 key authorization key id secret md5 secret
  md5 key authorization key id
```

Parameter	Description
<pre>area <backbone ospf_area> range ip_prefix <on off></pre>	Select an area from the areas already configured. Any area can be configured with any number of address ranges. These ranges are used to reduce the number of routing entries that a given area transmits to other areas. If a given prefix aggregates a number of more specific prefixes within an area, you can configure an address range that becomes the only prefix advertised to other areas. Be careful when configuring an address range that covers part of a prefix that is not contained within an area. An address range is defined by an IP prefix and a mask length. If you mark a range as restrict, it is not advertised to other areas.
<pre>area <backbone ospf_area> range ip_prefix restrict <on off></pre>	Any area can be configured with any number of address ranges. These ranges are used to reduce the number of routing entries that a given area transmits to other areas. If a given prefix aggregates a number of more specific prefixes within an area, you can configure an address range that becomes the only prefix advertised to other areas. Be careful when configuring an address range that covers part of a prefix that is not contained within an area. An address range is defined by an IP prefix and a mask length. If you mark a range as restrict, it is not advertised to other areas.
<pre>stub-network ip_prefix <on off></pre>	Specifies a stub network to which the specified interface range belongs. Configure a stub network to advertise reachability to prefixes that are not running OSPF. The advertised prefix appears as an OSPF internal route and is filtered at area borders with the OSPF area ranges. The prefix must be directly reachable on the router where the stub network is configured, that is, one of the router's interface addresses must fall within the prefix range to be included in the router-link-state advertisement. Use a mask length of 32 to configure the stub host. The local address of a point-to-point interface can activate the advertised prefix and mask. To advertise reachability to such an address, enter an IP address for the prefix and a non-zero cost for the prefix.
<pre>stub-network ip_prefix stub-network-cost <1-677722></pre>	Configure a stub network to advertise reachability to prefixes that are not running OSPF. The advertised prefix appears as an OSPF internal route and is filtered at area borders with the OSPF area ranges. The prefix must be directly reachable on the router where the stub network is configured, that is, one of the router's interface addresses must fall within the prefix range to be included in the router-link-state advertisement. Use a mask length of 32 to configure the stub host. The local address of a point-to-point interface can activate the advertised prefix and mask. To advertise reachability to such an address, enter an IP address for the prefix and a non-zero cost for the prefix.
<pre>interface if_name area <backbone ospf area> <on off></pre>	Specifies the OSPF area to which the specified interface belongs.

Parameter	Description
interface if_name hello-interval <1-65535>	Specifies the interval, in seconds, between hello packets that the router sends on the specified interface. For a given link, this value must be the same on all routers or adjacencies do not form.
interface if_name hello-interval default	Specifies the default value for the hello interval, which is 10 seconds.
interface if_name dead-interval <1-65535>	Specifies the number of seconds after which a router stops receiving hello packets that it declares the peer down. Generally, you should set this value at 4 times the value of the hello interval. Do not set the value at 0. For a given link, this value must be the same on all routers or adjacencies do not form.
interface if_name dead-interval default	Specifies the default value for the dead interval, which is 40 seconds
interface if_name retransmit-interval <1-65535>	Specifies the number of seconds between link state advertisement transmissions for adjacencies belonging to the specified interface. This value also applies to database description and link state request packets. Set this value conservatively, that is, at a significantly higher value than the expected round-trip delay between any two routers on the attached network.
interface if_name retransmit-interval default	Specifies the default for the retransmit interval, which is 5 seconds.
interface if_name cost <1-65535>	Specifies the weight of the given path in a route. The higher the cost, the less preferred the link. To use one interface over another for routing paths, assign one a higher cost.
interface if_name priority <0-255>	Specifies the priority for becoming the designated router (DR) on the specified link. When two routers attached to a network attempt to become a designated router, the one with the highest priority wins. This option prevents the DR from changing too often. The DR option applies only to a share-media interface, such as Ethernet or FDDI; a DR is not elected on a point-to-point type interface. A router with a priority of 0 is not eligible to become the DR.
interface if_name passive <on off>	Enabling this option puts the specified interface into passive mode; that is, hello packets are not sent from the interface. Putting an interface into passive mode means that no adjacencies are formed on the link. This mode enables the network associated with the specified interface to be included in intra-area route calculation rather than redistributing the network into OSPF and having it function as an autonomous system external. off
interface if_name authtype none	Specifies not to use an authentication scheme for the specified interface.

Parameter	Description
interface if_name authtype simple password	Specifies to use simple authentication for the specified interface. Enter an ASCII string that is 8 characters long. Generally, routers on a given link must agree on the authentication configuration to form peer adjacencies. Use an authentication scheme to guarantee that routing information is accepted only from trusted peers.
interface if_name authtype md5 key authorization key id secret md5 secret	Specifies to use MD5 authorization. Enter at least one key ID and its corresponding MD5 secret. If you configure multiple key IDs, the largest key ID is used for authenticating outgoing packets. All keys can be used to authenticate incoming packets. Generally, routers on a given link must agree on the authentication configuration to form peer adjacencies. Use an authentication scheme to guarantee that routing information is accepted only from trusted peers.

OSPF Virtual Links

Use these commands to configure OSPF virtual links. Configure a virtual link if the router is a border router that does not have interfaces in the backbone area. The virtual link is effectively a tunnel across an adjacent non-backbone area whose endpoint must be any of the adjacent area's border routers that has an interface in the backbone area.

For OSPFv2 use the following commands:

```
set ospf area backbone virtual-link ip_address
    transit-area ospf_area <on | off>
    transit-area ospf_area hello-interval <1-65535>
    transit-area ospf_area hello-interval default
    transit-area ospf_area dead interval <1-4294967295>
    transit-area ospf_area dead interval default
    transit-area ospf_area retransmit-interval <1-4294967295>
    transit-area ospf_area retransmit-interval default
    transit-area ospf_area authtype none
    transit-area ospf_area authtype simple password
    transit-area ospf_area authtype md5 key authorization key id secret md5 key
    transit-area ospf_area authtype md5 key authorization key id off
```

Parameter	Description
ip_address transit-area ospf_area <on off>	Specifies the IP address of the remote endpoint of the virtual link and transit area, which is a specified ospf area you configure using the set ospf area command. Configure the ospf area you are using as the transit area before you configure the virtual link. The transit area is the area shared by the border router on which you configure the virtual link and the router with an interface in the backbone area. Traffic between the endpoints of the virtual link flow through this area. The virtual link IP address functions as the router ID of the remote endpoint of the virtual link.
ip_address transit-area ospf_area hello-interval <1-65535>	Specifies the interval, in seconds, between hello packets that the router sends on the specified interface. For a given link, this value must be the same on all routers or adjacencies do not form.

Parameter	Description
ip_address transit-area ospf_area hello-interval default	Specifies an interval of 10 seconds.
ip_address transit-area ospf_area dead-interval <1-4294967295>	Specifies the number of seconds after which a router stops receiving hello packets that it declares the neighbor down. Generally, you should set this value at 4 times the value of the hello interval. Do not set the value at 0. For a given link, this value must be the same on all routers or adjacencies do not form.
ip_address transit-area ospf_area dead-interval default	Specifies a value of 40 seconds.
ip_address transit-area ospf_area retransmit-interval <1-4294967295>	Specifies the number of seconds between link state advertisement transmissions for adjacencies belonging to the specified interface. This value also applies to database description and link state request packets. Set this value conservatively, that is, at a significantly higher value than the expected round-trip delay between any two routers on the attached network.
ip_address transit-area ospf_area retransmit-interval default	Specifies a value of 5 seconds.
ip_address transit-area ospf_area authtype none	Specifies not to use an authentication scheme for the specified interface.
ip_address transit-area ospf_area authtype simple password	Specifies to use simple authentication for the specified interface. Enter an ASCII string that is 8 characters long. Generally, routers on a given link must agree on the authentication configuration to form neighbor adjacencies. Use an authentication scheme to guarantee that routing information is accepted only from trusted peers.
ip_address transit-area ospf_area authtype md5 key authorization key id secret MD5 secret	Specifies to use MD5 authorization. Enter at least one key ID and its corresponding MD5 secret. If you configure multiple key IDs, the largest key ID is used for authenticating outgoing packets. All keys can be used to authenticate incoming packets. Generally, routers on a given link must agree on the authentication configuration to form neighbor adjacencies. Use an authentication scheme to guarantee that routing information is accepted only from trusted peers.

OSPF Areas

Use the following commands to configure OSPF areas, including the backbone and stub areas.

For OSPFv2, use the following commands.

```
set ospf area backbone <on | off>
```

```
set ospf area ospf_area
  <on | off>
  stub <on | off>
  stub default-cost <1-677215>
  stub summary <on | off>
  nssa <on | off>
  nssa default-cost <1-677215>
  nssa default-metric-type <1-2>
  nssa import-summary-routes <on | off>
  nssa translator-role <always | candidate>
  nssa translator-stability-interval <1-65535>
  nssa redistribution <on | off>
  nssa range ip_addr [restrict] <on | off>
```

Parameter	Description
backbone <on off>	Specifies whether to enable or disable the backbone area. By default, the backbone area is enabled. You can disable the backbone area if the system does not have interfaces on the backbone area.
<on off>	Specifies the area ID for a new OSPF area. Best Practice - Check Point recommends that you enter the area ID as a dotted quad, but you can use any integer as the area ID. The area ID 0.0.0.0 is reserved for the backbone.
stub <on off>	Specifies the area ID for a stub area. Stub areas are areas that do not have AS external routes. Note - The backbone area cannot be a stub area.
stub default-cost <1-677215>	Specifies a default route into the stub area with the specified cost.
stub summary <on off>	Specifies the OSPF area as totally stubby, meaning that it does not have any AS external routes and its area border routers do not advertise summary routes.
nssa <on off>	Specifies the area ID for an NSSA. Note - The backbone area cannot be an NSSA area.
nssa default-cost <1-677215>	Specifies the cost associated with the default route to the NSSA.
nssa default-metric-type <1-2>	Specifies the type of metric. The default, type 1, is equivalent to the <i>Default ASE Route Type</i> on the OSPF WebUI page. A type 1 route is internal and its metric can be used directly by OSPF for comparison. A type 2 route is external and its metric cannot be used for comparison directly.

Parameter	Description
<code>nssa import-summary-routes <on off></code>	Specifies if summary routes (summary link advertisements) are imported into the NSSA.
<code>nssa translator-role <always candidate></code>	Specifies whether this NSSA border router will unconditionally translate Type-7 LSAs into Type-5 LSAs. When role is Always, Type-7 LSAs are translated into Type-5 LSAs regardless of the translator state of other NSSA border routers. When role is Candidate, this router participates in the translator election to determine if it will perform the translations duties.
<code>nssa translator-stability- interval <1-65535></code>	Specifies how long in seconds this elected Type-7 translator will continue to perform its translator duties once it has determined that its translator status has been assumed by another NSSA border router. Default: 40 seconds.
<code>nssa redistribution <on off></code>	Specifies if both Type-5 and Type-7 LSAs or only Type-7 LSAs will be originated by this NSSA border router.
<code>nssa rangeip_addr[restrict <on off></code>	Specify the range of addresses to reduce the number of Type-5 LSAs for the NSSA border router. To prevent a specific prefix from being advertised, use the restrict argument.

OSPF and IPv6 OSPF Show Commands

You can monitor and troubleshoot the OSPFv2 and OSPFv3 routing in CLI.

- For OSPFv2: `show ospf`
- For OSPFv3: `show ipv6 ospf`

To monitor and troubleshoot OSPFv2:

Run this command: `show ospf [border-routers | database [area {backbone | <area_id>} | areas [detailed] | asbr-summary-lsa [detailed] | checksum | database-summary | detailed | external-lsa [detailed] | network-lsa [detailed] | nssa-external-lsa [detailed] | opaque-lsa [detailed] | router-lsa [detailed] | summary-lsa [detailed] | type {1 | 2 | 3 | 4 | 5 | 6 | 7}] | errors {dd | hello | ip | lsack | lsr | lsu | protocol} | events | interface <int> [detailed | stats] | interfaces [detailed | stats] | neighbor <neighbor_IP> [detailed] | neighbors [detailed] | packets | routemap | summary]`

To monitor and troubleshoot OSPFv3:

Run this command: `show ipv6 ospf3 [border-routers | database [area {backbone | <area_id>} | areas [detailed] | checksum | database-summary | detailed | external-lsa [detailed] | inter-area-prefix-lsa [detailed] | inter-area-router-lsa [detailed] | intra-area-prefix-lsa [detailed] | link-lsa [detailed] | network-lsa [detailed] | router-lsa [detailed] | type {1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9}] | errors {dd | hello | ip | lsack | lsr | lsu | protocol} | events | interface <int> [detailed | stats] | interfaces [detailed | stats] | neighbor <neighbor_IP> [detailed] | neighbors [detailed] | packets | routemap | summary]`

Parameter	Description
border-routers	Shows the state of each area border router: • Router ID • OSPF area • Associated route cost
database [area {backbone <area_id>} areas [detailed] checksum database-summary detailed external-lsa [detailed] inter-area-prefix-lsa [detailed] inter-area-router-lsa [detailed] intra-area-prefix-lsa [detailed] link-lsa [detailed] network-lsa [detailed] router-lsa [detailed] nssa-external-lsa [detailed] opaque-lsa [detailed] summary-lsa [detailed] type {1 2 3 4 5 6 7 8 9}]	Show the OSPF database information: • area {backbone <area_id>} - Router-link state, network-link state, AS-border-router-link state, AS-external-link state, and summary-link state statistics for the specified OSPF area; for each interface in this OSPF area, shows the checksum, sequence number, and link count • areas [detailed] - Router-link state, network-link state, AS-border-router-link state, AS-external-link state, and summary-link state statistics for each OSPF area; for each OSPF interface, shows the checksum, sequence number, and link count • checksum - Checksum sum for each OSPF interface • database-summary - Summary of all types of LSAs • detailed - Detailed statistics on all types of LSAs • external-lsa [detailed] - Type 5 (AS-external) LSA statistics for each OSPF area • inter-area-prefix-lsa [detailed] - Type 3 (Summary) LSA statistics for each OSPF area (OSPFv3 only) • inter-area-router-lsa [detailed] - Type 4 (ASBR) LSA statistics for each OSPF area (OSPFv3 only) • intra-area-prefix-lsa [detailed] - Type 9 (Intra-Area Prefix) LSA statistics for each OSPF area (OSPFv3 only) • link-lsa [detailed] - Type 8 (Link) LSA statistics for each OSPF area (OSPFv3 only) • network-lsa [detailed] - Type 2 (Network) LSA statistics for each OSPF area • router-lsa [detailed] - Type 1 (Router) LSA statistics for each OSPF area • nssa-external-lsa [detailed] - Type 7 (NSSA) LSA statistics for each OSPF area (OSPFv2 only) • opaque-lsa [detailed] - Opaque LSA statistics for each OSPF area (OSPFv2 only) • summary-lsa [detailed] - Summary of all LSAs for each OSPF area (OSPFv2 only)

Parameter	Description
	• <code>type {1 2 3 4 5 6 7 8 9}</code> - LSA statistics for each type of LSA: • 1 - Router LSA • 2 - Network LSA • 3 - Summary LSA in OSPFv2, Inter-Area Prefix LSA in OSPFv3 • 4 - Summary ASBR LSA in OSPFv2, Inter-Area Router LSA in OSPFv3 • 5 - External LSA in OSPFv2, AS-External LSA in OSPFv3 • 7 - NSSA LSA in OSPFv2 only, not supported in OSPFv3 • 8 - Link LSA in OSPFv3 only, not supported in OSPFv2 • 9 - Intra-Area Prefix LSA in OSPFv3 only, not supported in OSPFv2
<code>errors {dd hello ip lsack lsr lsu protocol}</code>	Number of error messages sent, per type: • <code>dd</code> - Database description error messages • <code>hello</code> - Hello error messages • <code>ip</code> - IP error messages • <code>lsack</code> - Link-state acknowledgment error messages • <code>lsr</code> - Link-state request error messages • <code>lsu</code> - Link-state update error messages • <code>protocol</code> - Protocol error messages
<code>events</code>	Number of these types of events: • Interface down • Interface up • Virtual interface down • Virtual interface up • Designated Router (DR) elections • Router ID (RID) changes • Area border router (ABR) changes • AS border router (ASBR) changes • RFC1583 changes • LSA self-advertisement messages

Parameter	Description
<pre>interface <int> [detailed stats]</pre>	Shows OSPF information for the specified interface: • <int> - interface name • IP address • Area ID • State • Number of logged errors (NC) • DR Interface IP address • BDR Interface IP address • detailed - • IP Address • Area • Router ID • Network type • Cost • Authentication type • Error count • Event count • Transmit delay • State • Priority • Designated Router (DR) ID and interface IP address • Backup Designated Router (BDR) ID and interface IP address • Hello, Dead, Wait, and Retransmit timers (in seconds) • Next Hello timer • Neighbor count • Count of lost 2-way connections with neighbors • stats - • Total Errors • Hello Interval Mismatch • External Option Error • Delayed Ack Count • Dead Interval Mismatch • Lost Neighbor Count • Authentication Errors • Duplicate Router ID • Neighbor Errors • Newer Self LSA Count • Neighbor Count

Parameter	Description
<code>interfaces [detailed stats]</code>	Shows OSPF information for all interfaces: - Without command options - IP address - Area ID - State - Number of logged errors (NC) - DR Interface IP address - BDR Interface IP address <code>detailed</code> - - IP Address - Area - Router ID - Network type - Cost - Authentication type - Error count - Event count - Transmit delay - State - Priority - Designated Router (DR) ID and interface IP address - Backup Designated Router (BDR) ID and interface IP address - Hello, Dead, Wait, and Retransmit timers (in seconds) - Next Hello timer - Neighbor count - Count of lost 2-way connections with neighbors <code>stats</code> - - Total Errors - Hello Interval Mismatch - External Option Error - Delayed Ack Count - Dead Interval Mismatch - Lost Neighbor Count - Authentication Errors - Duplicate Router ID - Neighbor Errors - Newer Self LSA Count - Neighbor Count

Parameter	Description
<code>neighbor <neighbor_IP> [detailed]</code>	Shows OSPF information for the specified OSPF neighbor: • Priority • State • Number of logged errors
<code>neighbors [detailed]</code>	Shows OSPF information for each OSPF neighbor: • IP address • Priority • State • Number of logged errors
<code>packets</code>	Shows the number of received (Rx) and transmitted (Tx) OSPF packets: • Hello Rx • Hello Tx • Link State Update Rx • Link State Update Tx • Link State Ack Rx • Link State Ack Tx • Link State Request Rx • Link State Request Tx
<code>routermap</code>	Shows OSPF Import Policy and Export Policy.
<code>summary</code>	Shows detailed OSPF configuration.

IPv6 OSPF

In This Section:

Configuring OSPFv3 with IPv6 VRRP	110
Configuring IPv6 OSPF - WebUI	110
Configuring IPv6 OSPF - clish (ipv6 ospf)	116
Monitoring IPv6 OSPF	121

Open Shortest Path First (OSPF) is a link-state routing protocol that calculates forwarding tables in an IP-based internetwork. OSPF is the preferred Interior Gateway Protocol (IGP) for Check Point.

OSPF supports IPv6. OSPF for IPv6 is also referred to as OSPF version 3 (OSPFv3). OSPFv3 is defined in RFC 5340 (which makes RFC 2740 obsolete).

For IPv6, VRRP clusters are supported. ClusterXL clusters are not supported.

The IPv6 address appearing in the source of OSPF packets sent on the interface must be a link-local address, that is, an FE80::/64 address. A link-local address is automatically added to each interface when IPv6 is enabled on Gaia. The addresses are used for next hops, to advertise routes, and to send hello messages. OSPF advertises the IPv6 addresses defined by the user, but OSPF exchanges routes using the FE80 addresses. A /64 address is required by the OSPFv3 protocol. If the peer router does not use an FE80::/64 address, OSPFv3 does not work.

OSPFv2 is used with IPv4. See OSPF (on page 85).

Configuring OSPFv3 with IPv6 VRRP

To use OSPFv3 with VRRPv3, enable the Virtual Address option on the IPv6 OSPF Interface configuration page. If the configured interface is part of the VRRP master virtual router, OSPFv3 runs on the interface. When you enable this option, OSPFv3 uses the VRRPv3 virtual link-local address for the interface as the source of its control packets. This cannot be the automatically configured link-local address—that is, you must change the link-local address for the interface to something other than the default. You must configure the same link-local address on all the routers in the VRRP group.

VRRP installs the link-local address only on the master, so OSPFv3 runs only on that router. If a failover occurs, VRRPv3 installs the link-local address on the new master and OSPFv3 starts running on that system. Because OSPFv3 runs on one router at a time, there is no synchronization of OSPFv3 state between the VRRP group members.

Configuring IPv6 OSPF - WebUI

OSPFv3 has almost the same configuration parameters as OSPFv2.

To configure IPv6 OSPF:

1. Enable IPv6 ("IPv6 Support" on page 9). This automatically adds FE80::/64 link local to the interfaces.
2. Open the **Advanced Routing > IPv6 OSPF** page of the WebUI.

3. Define IPv6 OSPF Global settings, including the **Router ID** ("[IPv6 OSPF Global Settings - WebUI](#)" on page 111).
4. Optional: Define additional IPv6 OSPF areas, in addition to the backbone area ("[IPv6 OSPF Areas - WebUI](#)" on page 112).
5. Optional: For each area, you can add one or more address ranges if you want to reduce the number of routing entries that the area advertises into the backbone.



Note - To prevent an address range from being advertised into the backbone, select **Restrict** for the address range

6. Configure IPv6 OSPF Interfaces ("[IPv6 OSPF Interfaces - WebUI](#)" on page 114).

IPv6 OSPF Global Settings - WebUI

Configure these IPv6 OSPF global parameters.

Note - Graceful Restart Helper is not supported.

Parameter	Description
Router ID	A number that uniquely defines the router in a routing domain. Best Practice - Selecting an existing IP address in the unit is recommended <i>because the OSPF Router ID inherently makes it unique</i>. We recommend setting the Router ID rather than relying on the system to pick one based on an IP address. This prevents the router ID from changing if the interface used for the router ID goes down. • Range: Dotted-quad.[0-255].[0-255].[0-255].[0-255]. Do not use 0.0.0.0. • Default: A non-loopback IPv4 address assigned to the loopback interface, if available. Otherwise the system selects an IPv4 address from the list of active interfaces.
SPF Delay	The time (in seconds) the system waits before recalculating the OSPF routing table after a change in the topology. • Default: 2. • Range: 1-60.
SPF Hold Time	The minimum time (in seconds) between recalculations of the OSPF routing table. • Default: 5. • Range: 1-60.
Default ASE Route Cost	When routes from other protocols are redistributed into OSPF as ASEs, they are assigned this configured cost unless a cost has been specified for the individual route. • Range: 1-16777215. • Default: 1.

Parameter	Description
Default ASE Route Type	When routes from other protocols are redistributed into OSPF as ASEs, they are assigned this route type, unless a type has been specified for the individual route. ASEs can either be type 1 or type 2. • Type 1: Used for routes imported into OSPF which are from IGP's whose metrics are directly comparable to OSPF metrics. When a routing decision is being made, OSPF adds the internal cost to the AS border router to the external metric. • Type 2: Used for routes whose metrics are not comparable to OSPF internal metrics. In this case, only the external OSPF cost is used. In the event of ties, the least cost to an AS border router is used. • Range: Type 1, Type 2 • Default: 1.

IPv6 OSPF Areas - WebUI

Configure these IPv6 OSPF Area parameters.

The **Areas** section shows the IPv6 OSPF parameters of each area.

Add/Edit Area

Parameter	Description
Area	For the name of the area, choose an IPv4 address (preferred) or an integer.
Use Stub Area Type	A Stub Area is an area in which there are no Autonomous System External (ASE) routes. ASE routes are routes to destinations external to the AS. Note: The backbone area cannot be a stub area. NSSA Areas are not supported. • Range: Selected (Area Type is Stub), Deselected (Area Type is Normal) • Default: Deselected

Stub Area Parameters

These parameters show if you define the area as a stub area.

Parameter	Description
Cost for Default Route	The cost for the default route to the stub area. • Range: 1-16777215. • Default: No default.

Parameter	Description
Cost for Default Route	The cost for the default route to the stub area. • Range: 1-16777215. • Default: No default.
Totally Stubby	An area in which there are no ASE routes or summary routes. • Default: Areas are not initially totally stubby.

Add/Edit Address Range

Parameter	Description
Address Range	An area can be configured with any number of address ranges. Address ranges are used to reduce the number of routing entries that a given area will emit into the backbone (and hence all) areas. An address range is defined by a prefix and a mask length. If a given prefix aggregates a number of more specific prefixes within an area, then an address range can be configured and will be the only prefix advertised into the backbone. Be careful when configuring an address range that covers parts of a prefix that are not contained within the area.
Restrict	Prevent an address from being advertised into the backbone. • Range: Selected (Restrict), Deselected (Do not restrict)

Add/Edit Stub Network

Parameter	Description
Address Range	OSPF can advertise reachability to prefixes that are not running OSPF by configuring a stub network. The advertised prefix shows as an OSPF internal route and can be filtered at area borders with the OSPF area ranges. The prefix must be directly reachable on the router where the stub network is configured (that is, one of the routers interface addresses must fall in the prefix) in order to be included in the router-LSA. Stub hosts are configured by specifying a mask length of 128. An address range is defined by a prefix and a mask length. A prefix and mask can be advertised. That can be activated by the local address of a point-to-point interface. To advertise reachability to such an address, enter an IP address for the prefix and a non-zero cost for the prefix.
Cost	The cost associated with the stub network. The higher the cost, the less preferred the prefix as an OSPF route. • Range: 1-65535 • Default: 1

IPv6 OSPF Interfaces - WebUI

To configure an IPv6 OSPF interface:

1. In the **Add/Edit Interface** window, assign the appropriate **Area** to each interface by selecting the OSPF area that this interface participates in.

The OSPF interface configuration parameters are displayed showing the default settings. If you want to accept the default settings for the interface, no further action is necessary.

2. (Optional) Change any configuration parameters for the interface.

Note - The hello interval and dead interval must be the same for all routers on the link. Authentication is not supported for IPv6 OSPF interfaces.

Add/Edit Area

Parameter	Description
Interface	The interfaces for OSPF configuration. An interface must have an area associated with it in order to become active in OSPF.
Area	The OSPF area to which the interface belongs. An OSPF area defines a group of routers running OSPF that have the complete topology information of the area. OSPF areas use an area border router to exchange information about routes. Routes for a given area are summarized into the backbone area for distribution into other non-backbone areas. An area border router (ABR) is one that has at least two interfaces in at least two different areas. One of those areas must be the backbone or the router must have a virtual link configured. OSPF forces a hub and spoke topology of areas, with the backbone area always being the hub. • Range: All areas currently configured. • Default: None.
Hello interval	The time, in seconds, between hello packets that the router sends on the interface. For a given link, this must be the same on all routers or adjacencies will not be created. • Range: 1-65535 (seconds). • Default: For broadcast interfaces, the default is 10 seconds. For point-to-point interfaces, the default is 30 seconds.
Dead interval	The number of seconds after a router stops receiving hello packets that it declares the neighbor is down. Typically, the value of this field should be four times the size of the hello interval. For a given link, this field must be the same on all routers or adjacencies will not be created. The value must not be zero. • Range: 1-65535 (seconds). • Default: For broadcast interfaces, the default is 40 seconds. For point-to-point interfaces, the default is 120 seconds.

Parameter	Description
Retransmit interval	The number of seconds between LSA retransmissions, for adjacencies belonging to this interface. Also used when retransmitting Database Description and Link State Request Packets. This should be well over the expected round-trip delay between any two routers on the attached network. The setting of this value should be conservative or needless retransmissions will result. • Range: 1-65535 (seconds). • Default: 5.
OSPF Cost	The weight of a given path in a route. The higher the cost, the less preferred the link. You may explicitly override this setting in route redistribution. To use one interface over another for routing paths, give one a higher cost. • Range: 1-65535. • Default: 1.
Election Priority	The priority for becoming the designated router (DR) on this link. When two routers attached to a network both attempt to become a designated router, the one with the highest priority wins. If there is currently an elected DR on the link, it will remain the DR regardless of the configured priority. This feature prevents the DR from changing too often. This field is only relevant on a shared-media interface (Ethernet), as a DR is not elected on point-to-point type interfaces. A router with priority 0 is not eligible to become the designated router. • Range: 0-255. • Default: 1.
Passive	An interface in passive mode does not send hello packets out of the given interface. This means no adjacencies are formed on the link. The purpose of passive mode is to make it possible for the network associated with the interface to be included in the intra-area route calculation. In non passive mode, the network is redistributed into OSPF and is an ASE. In passive mode, all interface configuration information is ignored, with the exception of the associated area and the cost. • Default: Not selected.
Virtual Address	Directs OSPFv3 to use the VRRPv3 virtual link-local address as the source of its control packets. When enabled, OSPFv3 runs on the interface only while the local router is the master with respect to a VRRPv3 state machine on the interface. Note: The VRRPv3 state machine must back-up an IPv6 link-local address that is not auto-configured on the interface. • Default: Unselected.

Configuring IPv6 OSPF - clish (ipv6 ospf)

The commands for OSPFv3 are similar to those for OSPFv2, except that instead of `ospf` you type `ipv6 ospf3`.

- To configure OSPFv3 use `set ipv6 ospf` clish commands. There are no `add ipv6 ospf` commands.
- To disable a configuration, use the `off` setting in the `set` command. There are no `delete ipv6 ospf` commands.

To show the configuration, use `show ipv6 ospf` commands.

To work with OSPFv3, you must enable IPv6 on Gaia ("IPv6 Support" on page 9). This automatically adds FE80::/64 link local to the interfaces.

When you do initial configuration, set the router ID. Use these commands:

```
set router-id
  default
  ip_address
```

Parameter	Description
<code>router-id default</code>	Selects the highest interface address when OSPF is enabled.
<code>router-id ip_address</code>	Specifies a specific IP address to assign as the router ID. Do not use 0.0.0.0 as the router ID address. Best Practice - Check Point recommends setting the router ID rather than relying on the default setting. Setting the router ID prevents the ID from changing if the default interface used for the router ID goes down. The Router ID uniquely identifies the router in the autonomous system. The router ID is used by the BGP and OSPF protocols. We recommend setting the router ID rather than relying on the default setting. This prevents the router ID from changing if the interface used for the router ID goes down. Use an address on a loopback interface that is not the loopback address (127.0.0.1). Note - In a cluster, you must select a router ID and make sure that it is the same on all cluster members. • Range: Dotted-quad.([0-255].[0-255].[0-255].[0-255]). Do not use 0.0.0.0 • Default: The interface address of one of the local interfaces.

IPv6 OSPF Global Settings - clish

```
set ipv6 ospf3
  default-ase-cost <1-16777215>
  default-ase-type <1 | 2>
  spf-delay <1-60>
  spf-delay default
  spf-holdtime <1-60>
  spf-holdtime default
```

Note - The `graceful-restart-helper` parameter is not supported.

Parameter	Description
<code>default-ase-cost <1-16777215></code>	When routes from other protocols are redistributed into OSPF as ASEs, they are assigned this configured cost unless a cost has been specified for the individual route. • Default: 1
<code>default-ase-type <1 2></code>	When routes from other protocols are redistributed into OSPF as ASEs, they are assigned this route type, unless a type has been specified for the individual route. ASEs can either be type 1 or type 2. • Type 1: Used for routes imported into OSPF which are from IGPs whose metrics are directly comparable to OSPF metrics. When a routing decision is being made, OSPF adds the internal cost to the AS border router to the external metric. • Type 2: Used for routes whose metrics are not comparable to OSPF internal metrics. In this case, only the external OSPF cost is used. In the event of ties, the least cost to an AS border router is used. • Default: 1.
<code>spf-delay <1-60></code>	The time (in seconds) the system waits before recalculating the OSPF routing table after a change in the topology. • Default: 2.
<code>spf-holdtime <1-60></code>	The minimum time (in seconds) between recalculations of the OSPF routing table. • Default: 5.

IPv6 OSPF Areas - clish

Use the following commands to configure OSPFv3 (IPv6 OSPF) areas, including the backbone and stub areas.

NSSA is not available for OSPFv3

```
set ipv6 ospf3 area ospf_area <on | off>
set ipv6 ospf3 area ospf_area
  stub <on | off>
  stub default-cost <1-677215>
  stub summary <on | off>
  range VALUE <on | off>
  range VALUE restrict VALUE
  stub-network VALUE <on | off>
  stub-network VALUE stub-network-cost <1-65535>
```

Parameter	Description
<code>ospf_area <on off></code>	For the name of the area, choose an IPv4 address (preferred) or an integer.
<code>stub <on off></code>	A Stub Area is an area in which there are no Autonomous System External (ASE) routes. ASE routes are routes to destinations external to the AS. Note: The backbone area cannot be a stub area. NSSA Areas are not supported. Default: Off
<code>stub default-cost <1-677215></code>	The cost for the default route to the stub area. Default: No default.
<code>stub summary <on off></code>	An area in which there are no ASE routes or summary routes. Default: Off
<code>range VALUE <off on></code>	An area can be configured with any number of address ranges. Address ranges are used to reduce the number of routing entries that a given area will emit into the backbone (and hence all) areas. An address range is defined by a prefix and a mask length. If a given prefix aggregates a number of more specific prefixes within an area, then an address range can be configured and will be the only prefix advertised into the backbone. Be careful when configuring an address range that covers parts of a prefix that are not contained within the area.
<code>range VALUE restrict <on off></code>	Prevent an address from being advertised into the backbone.
<code>stub-network VALUE <on off></code>	OSPF can advertise reachability to prefixes that are not running OSPF by configuring a stub network. The advertised prefix shows as an OSPF internal route and can be filtered at area borders with the OSPF area ranges. The prefix must be directly reachable on the router where the stub network is configured (that is, one of the routers interface addresses must fall in the prefix) in order to be included in the router-LSA. Stub hosts are configured by specifying a mask length of 128. An address range is defined by a prefix and a mask length. A prefix and mask can be advertised. That can be activated by the local address of a point-to-point interface. To advertise reachability to such an address, enter an IP address for the prefix and a non-zero cost for the prefix.
<code>stub-network VALUE stub-network-cost <1-65535></code>	The cost associated with the stub network. The higher the cost, the less preferred the prefix as an OSPF route. Default: 1

IPv6 OSPF Interfaces - clish

Note - The hello interval and dead interval must be the same for all routers on the link. Authentication is not supported for IPv6 OSPF interfaces.

```
set ipv6 ospf3 interface interface_name
  area ospf_area <on|off>
  cost <1-65535>
  dead-interval <1-65535>
  hello-interval <1-65535>
  passive <on | off>
  priority <0-255>
  retransmit-interval <1-65535>
  virtual-address <on|off>
```

Parameter	Description
interface <i>interface_name</i>	The interfaces for OSPF configuration. An interface must have an area associated with it in order to become active in OSPF.
area <i>ospf_area</i> <off on>	The OSPF area to which the interface belongs. An OSPF area defines a group of routers running OSPF that have the complete topology information of the area. OSPF areas use an area border router to exchange information about routes. Routes for a given area are summarized into the backbone area for distribution into other non-backbone areas. An area border router (ABR) is one that has at least two interfaces in at least two different areas. One of those areas must be the backbone or the router must have a virtual link configured. OSPF forces a hub and spoke topology of areas, with the backbone area always being the hub. For the name of the area, choose an IPv4 address (preferred) or an integer. • Range: All areas currently configured. • Default: None.
cost <1-65535>	The weight of a given path in a route. The higher the cost, the less preferred the link. You may explicitly override this setting in route redistribution. To use one interface over another for routing paths, give one a higher cost. • Default: 1.
dead-interval <1-65535>	The number of seconds after a router stops receiving hello packets that it declares the neighbor is down. Typically, the value of this field should be four times the size of the hello interval. For a given link, this field must be the same on all routers or adjacencies will not be created. The value must not be zero. • Default: For broadcast interfaces, the default is 40 seconds. For point-to-point interfaces, the default is 120 seconds.

Parameter	Description
hello-interval <1-65535>	The time, in seconds, between hello packets that the router sends on the interface. For a given link, this must be the same on all routers or adjacencies will not be created. • Default: For broadcast interfaces, the default is 10 seconds. For point-to-point interfaces, the default is 30 seconds.
passive <on off>	An interface in passive mode does not send hello packets out of the given interface. This means no adjacencies are formed on the link. The purpose of passive mode is to make it possible for the network associated with the interface to be included in the intra-area route calculation. In non passive mode, the network is redistributed into OSPF and is an ASE. In passive mode, all interface configuration information is ignored, with the exception of the associated area and the cost. • Default: Off
priority <0-255>	The priority for becoming the designated router (DR) on this link. When two routers attached to a network both attempt to become a designated router, the one with the highest priority wins. If there is currently an elected DR on the link, it will remain the DR regardless of the configured priority. This feature prevents the DR from changing too often. This field is only relevant on a shared-media interface (Ethernet), as a DR is not elected on point-to-point type interfaces. A router with priority 0 is not eligible to become the designated router. • Default: 1.
retransmit-interval <1-65535>	The number of seconds between LSA retransmissions, for adjacencies belonging to this interface. Also used when retransmitting Database Description and Link State Request Packets. This should be well over the expected round-trip delay between any two routers on the attached network. The setting of this value should be conservative or needless retransmissions will result. • Default: 5.
virtual-address <on off>	Directs OSPFv3 to use the VRRPv3 virtual link-local address as the source of its control packets. When enabled, OSPFv3 runs on the interface only while the local router is the master with respect to a VRRPv3 state machine on the interface. Note: The VRRPv3 state machine must back-up an IPv6 link-local address that is not auto-configured on the interface. • Default: Off.

IPv6 OSPF Configuration Examples - clish



Note - If OSPF is used without the virtual-address option in a VRRP cluster, you must make sure that each router selects a different router-id. To see and correct the router-id:

```
show router-id
set router-id ipv4-address
```

Example of OSPFv3 configuration:

```
set ipv6 ospf3 interface eth1 area backbone on
```

To change default OSPFv3 parameters:

```
set ipv6 ospf3 interface eth1 cost 2
set ipv6 ospf3 interface eth1 priority 2
set ipv6 ospf3 interface eth1 hello-interval 20
set ipv6 ospf3 interface eth1 dead-interval 80
set ipv6 ospf3 interface eth1 retransmit-interval 20
```

To configure OSPFv3 to run on the VRRP Virtual Address:

```
set ipv6 ospf3 interface eth1 virtual-address on
```

Monitoring IPv6 OSPF

You can monitor IPv6 OSPF in the WebUI and in the clish CLI.

Monitoring IPv6 OSPF in the WebUI

1. In the WebUI, go to the **Advanced Routing > IPv6 OSPF** page.
2. Click the **Monitoring** tab.

Show IPv6 OSPF Commands

See OSPF and IPv6 OSPF Show Commands (on page 104).

Route Aggregation

In This Section:

Configuring Route Aggregation - WebUI	122
Configuring Route Aggregation - CLI (aggregate)	124

Route aggregation is a method of generating a more general route, given the presence of a specific route. Use this method to reduce the number of routes advertised for a given protocol. For example, if a router has many stub interface routes subnetted from a class C and is running RIPv2 on another interface, you can use interface routes to create an aggregate route of the class C that can then be redistributed into RIP. This action reduces the number of routes advertised through RIP. Be careful when aggregating if there are "holes" in the route that is aggregated.

The interface that originates the aggregate routes does not use aggregate routes to forward packets. Only the router receiving data can use the Aggregate routes to forward traffic. A router that receives a packet that does not match one of the component routes that generated an aggregate route should respond with an ICMP network unreachable message. This action prevents packets for unknown component routes from following a default route into another network, where they are continually forwarded back to the border router until their TTLs expire.

Create an aggregate route by specifying the network address and mask length and then providing a set of contributing routes. Define a contributing route by specifying a source, such as, a routing protocol, a static route, an interface route, and a route filter, which is either a prefix or the keyword all. An aggregate route can have many contributing routes, but at least one of the routes must be present to generate an aggregate route.

Configuring Route Aggregation - WebUI

To create aggregate routes

1. Go to the **Advanced Routing > Route Aggregation** page of the WebUI.
2. Click **Add**.
The **Add Aggregate Route** window opens.
3. Enter the **IPv4 address** for the new contributing route.
4. Enter the **Subnet mask**.
The IPv4 address and subnet mask correspond to a single routing table entry.
5. In the **Contributing Protocol** section, click **Add**.
The **Add Contribution Setting** window opens.
6. Select the **Protocol** for the new aggregate route.
7. Optional: Select **Contribute All Routes**.
8. Optional: To specify a prefix, fill in the **Address** and **Subnet mask**.
9. Select the **Match Type** (None, Refines or Exact).
10. Click **Save**.

Add Aggregate Route Window

Parameter	Description
IPv4 address	The IPv4 address of a route that activates the aggregate route, if contributed by the protocol. • Range: Dotted-quad ([0-255].[0-255].[0-255].[0-255]) • Default: No default.
Subnet mask	The mask length of the route that activates the aggregate route, if contributed by the protocol. • Range: 0-32. • Default: No default.
Rank	The routing system uses rank when there are routes from different protocols to the same destination. For each route, the route from the protocol with the lowest rank is used. • Range: 0-255. • Default: Each routing protocol has a default rank value. Aggregate routes have a default rank of 130. The default value for OSPF is 10, and the default value for RIP is 100.
Weight	A second tie breaker after the rank. It selects routes going to the same destination. The route with the highest weight is an active route and is installed in the kernel forwarding table and redistributed to other routing protocols. • Range: 0-65535. • Default: 0.
AS Path Truncate	When selected, the AS path is truncated to the longest common AS path. • Options: Select / Clear • Default: Clear. Build an AS path consisting of SETs and SEQUENCEs of all contributing AS paths.
Contributing Protocol	Contributing protocols whose routes activate the aggregate routes.

Add Contribution Setting window

Parameter	Description
Protocol	A contributing protocol whose routes activate the aggregate route. • Range: All routing protocols, as well as interface routes, aggregate routes, and static routes. Select one of the following: All / Direct / Static / Aggregate / OSPF2 / OSPF2ASE/ RIP /BGP. • Default: No default
Contribute All Routes	When selected, lets any routes contributed by the protocol activate the aggregate route. • Options: Select / Clear • Default: Clear

Parameter	Description
Address	The IP address of the aggregate route being created. • Range: Dotted-quad ([0-255].[0-255].[0-255].[0-255]) • Default: No default
Subnet mask	The mask length of the aggregate route being created. • Range: 0-32 • Default: No default
Match Type	The routes that are filtered for the Address and Subnet mask . These are the ways to compare other routes against it: • None - matches any route that equals the specified route or is more specific than the specified route. • Exact - matches a route only if it equals the From Address and Subnet mask of the specified route. • Refines - matches a route only if it is more specific than the specified route. • Options: None, Exact, Refines. • Default: None.

Configuring Route Aggregation - CLI (aggregate)

Create aggregate routes using these commands.

```
set aggregate ip_prefix
  contributing protocol <protocol> contributing-route
    <all | ip_prefix> <on | off>
    <ip_prefix> exact on
    <ip_prefix> refines on
  off
  contributing protocol <protocol> off
  rank default
  rank <0-255>
  weight default
  aspath-truncate <on | off>
```

Parameter	Description
contributing protocol <protocol> contributing-route <all ip_prefix <on off>	The IP address and mask length of the new aggregate route and the contributing protocol or interface route. To specify a protocol, enter direct, static, ospf2, ospf2ase, bgp, rip, igrp, rip, or aggregate. To specify a contributing route, enter all to contribute all the routes for a specific protocol or enter the IP address and mask length to contribute a specific route.
<ip_prefix> exact on	The contributing route is limited to the specified IP address and mask length only. You cannot enable both exact on and refines on at the same time. If you enable refines on when exact on is enabled, exact on is automatically disabled.

Parameter	Description
<code><ip_prefix> refines on</code>	The contributing route is based on addresses with a greater value than the specified mask length of the specified IP address. You cannot enable both <code>exact on</code> and <code>refines on</code> at the same time. If you enable <code>refines on</code> when <code>exact on</code> is enabled, <code>exact on</code> is automatically disabled.
<code>rank default</code>	The rank to assign to the aggregate route when routes from different protocols to the same destination are present. For each route, the route from the protocol with the lowest rank is used. Each routing protocol has a different default rank value. Aggregate routes have a default rank of 130.
<code>rank <0-255></code>	The rank to assign to the aggregate route when routes from different protocols to the same destination are present. For each route, the route from the protocol with the lowest rank is used. Each routing protocol has a different default rank value. Each routing protocol has a different default rank value. Aggregate routes have a default rank of 130.
<code>weight default</code>	A value that breaks a tie if select routes going to the same destination have the same rank value. The route with the highest weight is the active route. The active route is installed in the kernel forwarding table and redistributed to the other routing protocols. • Range: 0-65535. • Default: 0
<code>weight <0-65535></code>	A value that breaks a tie if select routes going to the same destination have the same rank value. The route with the highest weight is the active route. The active route is installed in the kernel forwarding table and redistributed to the other routing protocols. • Default: 0
<code>aspath-truncate <on off></code>	Specifies that the autonomous system (AS) path be truncated to the longest common AS path. The default behavior is to build an AS path that consists of sets and sequences of all contributing AS paths. • Default: off

Routing Policy Configuration

In This Section:

Configuring Inbound Route Filters - WebUI	126
Configuring Inbound Route Filters - CLI	132
Configuring Route Redistribution - WebUI	132
Configuring Route Redistribution - CLI	142
Configuring Route Maps - CLI (routemap)	159

You can configure routing policy for RIP, OSPFv2 and BGP in these ways:

Routing Policy Configuration	Description	Configured Using
Inbound Route filters	Define filters for routes accepted by a given routing protocol. Inbound Route filters are similar to route maps for an import policy.	WebUI
Route Redistribution	Redistribute routes learned from one routing protocol into another routing protocol. It is also useful for advertising static routes, such as the default route, or aggregate routes. Route Redistribution is similar to route maps for an export policy.	WebUI clish
Routemaps	Control which routes are accepted and announced. Used to configure inbound route filters, outbound route filters, and to redistribute routes from one protocol to another. Route maps offer more configuration options than the WebUI options. However, they are not functionally equivalent. Routemaps assigned to a protocol for import or export override corresponding filters and route redistribution rules.	clish

Configuring Inbound Route Filters - WebUI

Inbound Route Filters let you define which external to a routing protocol routes are accepted by that protocol. You can define Inbound Route Filters through the WebUI or through the CLI.

By default, all routes, external to RIP, OSPFv2 (IPv4), and OSPFv3 (IPv6), are accepted by these protocols. To narrow down the selection of accepted routes, you can edit the default policies and configure new policies.

By default, BGP does not accept any routes. You must configure explicit policies for BGP to accept routes.

When you configure Inbound Route Filters, to specify precision with which the network addresses

are matched, use the same Match Type criteria rules as for route redistribution:

- The prefix and the mask length are matched exactly.
- The prefix is matched exactly, and the mask length is greater than the one specified.
For example, if the network address 10.0.0.0/8 is specified in the filter, then any route with the prefix 10 and the mask length greater than 8 is matched, but those with the mask length of exactly 8 are not matched.
- The prefix is matched exactly, and the mask length is equal to or greater than the one specified.
For example, if the network address 10.0.0.0/8 is specified in the filter, then any route with the prefix 10 and the mask length equal to or greater than 8 is matched.
- The prefix is matched exactly, and the mask length is within the specified range of masks. The mask range values must be equal to or greater than the network mask.
For example, if the network address 10.0.0.0/8 and the mask range 16 to 8 are specified in the filter, then any route with the prefix 10 and the mask length between 8 and 16 is matched.



Note - routemap import configuration overrides the Inbound Route Filters configuration.

To change a default Inbound Router Filter policy :

1. In the WebUI, go to the **Advanced Routing > Inbound Route Filters** page.
2. Select a default Inbound Route Filter policy
 - **OSPFv2** (IPv4 routes)
 - **RIP**
 - **OSPFv3** (IPv6 routes)
3. Click **Edit**.
4. In the window that opens, do one of these:
 - Select **Action > Accept** (default) and enter the **Rank** from 0 to 255 (the default value for RIP is 10, for OSPFv3 - 150, and for OSPFv2 - 200)
 - Select **Action > Restrict**
5. Click **Save**.

To configure an Inbound Route Filter for an individual route:

1. In the WebUI, go to the **Advanced Routing > Inbound Route Filters** page.
2. Click **Add** and select one of these:
 - **Add Individual IPv4 Route**
 - **Add Individual IPv6 Route**
3. In the **Add IPv4/IPv6 Route** window that opens, set the route parameters:
 - **Import To** - select the routing protocol for which you want to add the inbound route filter
 - **Route** - enter the network prefix and mask
 - **Match Type** - set the precision with which you want the network addresses in the routes to be matched against the filter criteria
 - **Action** - choose to **Accept** or to **Restrict** the routes that match the filter criteria
 - **Rank** - assign a rank to the route

4. Click **Save**.
5. **Add Route** window opens ("[Fine Tuning Policies](#)" on page 130).

To configure a policy for RIP routes:

1. Go to the **Advanced Routing > Inbound Route Filters** page of the WebUI.
2. In the **Inbound Route Protocols and BGP Policies** section, select **RIP Routes**.
3. Click **Edit**.
4. In the **Configure RIP All Routes** window, select the **Action**:
 - **Options**: Accept or Restrict
 - **Default**: Accept
5. If you selected **Accept**, change the **Rank**:
 - **Range**: 0-255
 - **Default**: 100
6. You can fine tune the policy for RIP routes. In the **Individual Routes** section click **Add**.
The **Add Route** window opens ("[Fine Tuning Policies](#)" on page 130).

To configure a policy for BGP routes:

1. Go to the **Advanced Routing > Inbound Route Filters** page of the WebUI.
2. In the **Inbound Route Protocols and BGP Policies** section, click **Add BGP Policy**.
The **Add BGP Policy** window opens ("[Add BGP Policy Window](#)" on page 128).
3. You can fine tune the policy for BGP routes. In the **Individual Routes** section click **Add**.
The **Add Route** window opens ("[Fine Tuning Policies](#)" on page 130).



Note - For BGP, no routes are accepted from a peer by default. You must configure an explicit Inbound BGP Route Filter to accept a route from a peer.

Add BGP Policy Window

Parameter	Description
BGP Type: Based on AS_PATH Regular Expression (1-511)	An autonomous system can control BGP importation. BGP supports propagation control through the use of AS-PATH regular expressions. BGP version 4 supports the propagation of any destination along a contiguous network mask.
BGP Type: Based on Autonomous System Number (512-1024)	An autonomous system can control BGP importation. BGP can accept routes from different BGP peers based on the peer AS number.
Import ID	The order in which the import lists are applied to each route. • Range for BGP Type based on AS_PATH Regular Expression: 1-511 • Range for BGP Type based on Autonomous System Number: 512-1024 • Default: No default

Parameter	Description
AS Number	Autonomous system number of the peer AS. • Range: 0-65535
AS-PATH Regular Expression	The following definitions describe how to create regular expressions. AS-PATH operators are one of the following: • aspath_term (m n) A regular expression followed by (m n), where m and n are both non-negative integers and m is less than or equal to n. This expression means that there are at least m, and at most, n repetitions. • aspath_term m A regular expression followed by m, where m is a positive integer and means exactly m repetitions. • aspath_term (m) A regular expression followed by m, where m is a positive integer. This expression means that there are exactly m repetitions. • aspath_term * A regular expression followed by *, which means zero or more repetitions. • aspath_term + A regular expression followed by +, which means one or more repetitions. • aspath_term ? A regular expression followed by ?, which means zero or one repetition. • aspath_term aspath_term Match either the AS term on the left or the AS term on the right of the pipe.
Origin	The completeness of AS-PATH information. • Any - • IGP - A route was learned from an interior routing protocol and is probably complete. • EGP - The route was learned from an exterior routing protocol that does <i>not</i> support AS-PATHs, and the path is probably incomplete. • Incomplete - The path information is incomplete. • Options: Any / IGP / EGP / Incomplete • Default: No default

Parameter	Description
Weight	BGP stores any routes that are rejected by not mentioning them in a route filter. BGP explicitly mentions these rejected routes in the routing table and assigns them a restrict keyword with a negative weight. A negative weight prevents a route from becoming active, which means that it is not installed in the forwarding table or exported to other protocols. This feature eliminates the need to break and re-establish a session upon reconfiguration if importation policy is changed. • Range: 0-65535 • Default: No default
Local Pref.	The BGP local preference to the imported route. Check Point recommends that you configure this value to bias the preference of routed for BGP routes. Note: Do not use the local preference parameter when importing BGP. The local preference value is sent automatically when redistributing external BGP routes to an internal BGP route. The local preference parameter is ignored if used on internal BGP import statements. • Range: 0-65535. Larger values are preferred • Default: No default
All Routes: Action	Whether the routing protocol should accept or restrict the All Routes route, equivalent to 0.0.0.0/0, from the given AS-Path or AS. If set to Accept, you can specify a Rank for all routes. • Options: Accept / Restrict • Default: Restrict
All Routes: Rank	If All Routes: Action is set to Accept, you can specify a Rank for all routes. • Range: 0 - 65535 • Default: no default.

Fine Tuning Policies

To fine tune your OSPF, RIP or BGP Policy:

1. Specify which routes should be filtered by:
 - IP address
 - Subnet mask
 - Match type
 - Optional: Parameters that depend on the match type. For routes that match a filter, you can select Accept or Restrict. If the route is accepted, you can specify its rank.
2. Specify what actions to perform on a route if it matches the route filter.

Do these steps by configuring the parameters in the **Add Route** window.

Add Route Window

Parameter	Description
Protocol	The protocol for which you want to create the inbound route filter.
Address Subnet mask	A baseline route that specifies a route filter. This route is the <i>specified route</i> in the context of a single route filter.
Matchtype	The routes that are filtered for the From Address and Subnet mask. These are the ways to compare other routes against it: • Normal - matches any route that equals the specified route or is more specific than the specified route. • Exact - matches a route only if it equals the From Address and Subnet mask of the specified route. • Refines - matches a route only if it is more specific than the specified route. • Range - matches any route whose Ip prefix equals the specified route's From Address and whose Subnet Mask falls within the specified Subnet Mask length range. • Options: Normal, Exact, Refines, Range. • Default: Normal.
Action	What to do with the routes that match the filter that is defined by the From Address, Subnet mask and Matchtype. • Options: Accept, Restrict. • Default: Accept.
Weight	BGP stores any routes that are rejected by not mentioning them in a route filter. BGP explicitly mentions these rejected routes in the routing table and assigns them a restrict keyword with a negative weight. A negative weight prevents a route from becoming active, which means that it is not installed in the forwarding table or exported to other protocols. This feature eliminates the need to break and re-establish a session upon reconfiguration if importation policy is changed. • Range: 0-65535 • Default: No default
Local Pref	The BGP local preference to the imported route. Check Point recommends that you configure this value to bias the preference of routed for BGP routes. Note: Do not use the local preference parameter when importing BGP. The local preference value is sent automatically when redistributing external BGP routes to an internal BGP route. The local preference parameter is ignored if used on internal BGP import statements. • Range: 0-65535. Larger values are preferred • Default: No default

Configuring Route Redistribution - WebUI

Route redistribution lets a router propagate routes between routing protocols - IPv4 or IPv6. Route redistribution is also useful for advertising the default route, static routes, or aggregate routes.



Note - Static routes take precedence over dynamic routes of any kind, native or redistributed.

To Configure Route Redistribution

1. In the WebUI, go to the **Advanced Routing > Route Redistribution** page.
2. In the **Route Redistributions** section:
 - To add a redistributed route, click **Add Redistribution From**.
 - To edit a redistributed route, select it and click **Edit**.
3. Configure the route source, the destination routing protocol, and the corresponding redistribution settings.

Routes from these sources can be redistributed into IPv4 and IPv6 routing protocols:

Interface

Redistribute interface routes.

Parameter	Description
To Protocol	The destination routing protocol.
Interface	The interface from which to distribute the routes. You can select all , or one of the configured interfaces.
Metric	The cost of the redistributed routes in the destination routing protocol.

Static

Redistribute static routes.

Parameter	Description
To Protocol	The destination routing protocol.
Static Route	The static route to be redistributed into the destination routing protocol.
Metric	The cost of the redistributed routes in the destination routing protocol. Note - this parameter is mandatory when configuring redistribution into RIP.

Redistributing IPv4 Routes - WebUI

Configure route sources for redistribution into IPv4 routing protocols.

Aggregate

Redistribute aggregate routes.

Parameter	Description
To Protocol	The destination routing protocol.
Aggregate Route	The aggregate route to be redistributed into the destination protocol.
Metric	The cost of the redistributed routes in the destination routing protocol: • RIP - a number in the range 1-16 • OSPFv2 - a number in the range 1-16777215 • BGP - a number in the range 1-4294967295 Note - this parameter is mandatory when configuring redistribution into RIP.

RIP

Redistribute RIP routes.

Parameter	Description
To Protocol	The destination routing protocol.
Route: All	Choose which RIP routes to redistribute into the destination routing protocol. When All is selected, all active RIP routes get redistributed. When cleared (default), only the RIP routes that match the Address Range and Match Type criteria get redistributed.
Address Range	The network route to be redistributed into the destination routing protocol. Specify the IPv4 prefix and the mask length.
Match Type	The precision with which the Address Range routes are matched: • Normal (default)- matches routes for the IP addresses and IP subnets in the specified range, including the range itself • Exact - matches the routes for the specified address range exactly as it appears • Refines - matches routes for the IP addresses and IP subnets in the specified range, but not the range itself • Range - matches routes for the IP subnets in the specified range, including the range itself

Parameter	Description
Action	To redistribute the routes that match the specified criteria or not: • Accept (default) - to redistribute • Restrict - not to redistribute
Metric	The cost of redistributed routes in the destination routing protocol: • OSPFv2 - a number in the range 1-16777215 • BGP - a number in the range 1-4294967295

OSPFv2

Redistribute OSPFv2 routes.

Parameter	Description
To Protocol	The destination routing protocol.
Route: All	Choose which OSPFv2 routes to redistribute into the destination routing protocol. When All is selected, all active OSPFv2 routes get redistributed. When cleared (default), only the OSPFv2 routes that match the Address Range and Match Type criteria get redistributed.
Address Range	The network route to be redistributed into the destination routing protocol. Specify the IPv4 prefix and the mask length.
Match Type	The precision with which the Address Range routes are matched: • Normal (default)- matches routes for the IP addresses and IP subnets in the specified range, including the range itself • Exact - matches the routes for the specified address range exactly as it appears • Refines - matches routes for the IP addresses and IP subnets in the specified range, but not the range itself • Range - matches routes for the IP subnets in the specified range, including the range itself
Action	To redistribute the routes that match the specified criteria or not: • Accept (default) - to redistribute • Restrict - not to redistribute
Metric	The cost of the redistributed routes in the destination routing protocol: • RIP - a number in the range 1-16 • BGP - a number in the range 1-4294967295 Note - this parameter is mandatory when configuring redistribution into RIP.

Parameter	Description
RIP Tag (when redistributing into RIP only)	RIP route tag (optional) - a value in the range 1-65535.

OSPFv2 External

Redistribute OSPFv2 External routes.

Parameter	Description
To Protocol	The destination routing protocol.
Route: All	Choose which OSPFv2 External routes to redistribute into the destination routing protocol. When All is selected, all active OSPFv2 routes get redistributed. When cleared (default), only the OSPFv2 routes that match the Address Range and Match Type criteria get redistributed.
Route: Address Range	The network route to be redistributed into the destination routing protocol. Specify the IPv4 prefix and the mask length.
Route: Match Type	The precision with which the Address Range routes are matched: • Normal (default)- matches routes for the IP addresses and IP subnets in the specified range, including the range itself • Exact - matches the routes for the specified address range exactly as it appears • Refines - matches routes for the IP addresses and IP subnets in the specified range, but not the range itself • Range - matches routes for the IP subnets in the specified range, including the range itself
Action	To redistribute the routes that match the specified criteria or not: • Accept (default) - to redistribute • Restrict - not to redistribute
Metric	The cost of the redistributed routes in the destination routing protocol: • RIP - a number in the range 1-16 • BGP - a number in the range 1-4294967295 Note - this parameter is mandatory when configuring redistribution into RIP.
RIP Tag (when redistributing into RIP only)	RIP route tag (optional) - a value in the range 1-65535.

BGP Based on AS-Path

Redistribute BGP routes, based on the AS path criteria.

Parameter	Description
To Protocol	The destination routing protocol.
From BGP AS-Path	AS Path, defined through a regular expression. An AS-Path regular expression consists of one or more <i>term-operator</i> pairs. Each <i>term</i> identifies one or more of the Autonomous Systems, and each <i>operator</i> specifies the matching criteria for the corresponding term. An AS-Path Regular Expression <i>term</i> can be defined through single AS numbers and groups of AS numbers. These are the AS-Path Regular Expression <i>operators</i>: • . - Matches any single character • \ - Matches the character right after the backslash (to enter \ in CLI, type another \ immediately after the first one) • ^ - Matches the characters or a null string at the beginning of an AS path • \$ - Matches the characters or a null string at the end of an AS path • ? - Matches zero or one occurrence of the preceding term (to enter ? in CLI, hit <i>Ctrl-V</i>, then type ?) • * - Matches zero or more occurrences of the preceding term • + - Matches one or more occurrences of the preceding term • - Matches one of the terms on either side of the " " character • _ - Matches a non-alphanumeric character (" ", "{", "}", "^", "\$") or a whitespace • [] - Matches a term in a set of terms (separated by whitespaces) or a range of terms (separated by a hyphen) • () - Matches a group of terms as a single term • {m, n} - Matches at least m and at most n repetitions of the preceding term • {m} - Matches exactly m repetitions of the preceding term • {m, } - Matches m or more repetitions of the preceding term

Parameter	Description
Origin	The origin and the completeness of the AS-Path information: • Any - Any route • IGP - The route was learned from an interior routing protocol and the AS-Path information is probably complete • EGP - The route was learned from an exterior routing protocol that does <i>not</i> support AS-Path, and the AS-Path information is probably incomplete • Incomplete - The AS-Path information is incomplete
Route: All	Choose which BGP Based on AS-Path routes to redistribute into the destination routing protocol. When All is selected, all active BGP Based on AS-Path routes get redistributed. When cleared (default), only the BGP Based on AS-Path routes that match the Address Range and Match Type criteria get redistributed.
Route: Address Range	The network route to be redistributed into the destination routing protocol. Specify the IPv4 prefix and the mask length.
Route: Match Type	The precision with which the Address Range routes are matched: • Normal (default)- matches routes for the IP addresses and IP subnets in the specified range, including the range itself • Exact - matches the routes for the specified address range exactly as it appears • Refines - matches routes for the IP addresses and IP subnets in the specified range, but not the range itself • Range - matches routes for the IP subnets in the specified range, including the range itself
Action	To redistribute the routes that match the specified criteria or not: • Accept (default) - to redistribute • Restrict - not to redistribute
Metric	The cost of the redistributed routes in the destination routing protocol: • RIP - a number in the range 1-16 • OSPFv2 - a number in the range 1-16777215 • BGP - a number in the range 1-4294967295 Note - this parameter is mandatory when configuring redistribution into RIP.
RIP Tag (when redistributing into RIP only)	RIP route tag (optional) - a value in the range 1-65535.

Parameter	Description
Automatic Tag (when redistributing into OSPFv2 only)	OSPFv2 route tag (optional) - when selected, uses the value specified in Arbitrary Tag to generate the route tag value.
Arbitrary Tag (when redistributing into OSPFv2 only)	Used only when Automatic Tag is selected. Valid values: 0-4095.
Manual Tag (when redistributing into OSPFv2 only)	OSPFv2 route tag (optional). Valid values: 1-2147483647

BGP Based on AS

Redistribute BGP routes from a specific BGP Autonomous System (AS).

Parameter	Description
To Protocol	The destination routing protocol.
From BGP AS	The source BGP AS, from which routes get redistributed.
Routes: All	Choose which BGP Based on AS routes to redistribute into the destination routing protocol. When All is selected, all active BGP Based on AS routes get redistributed. When cleared (default), only the BGP Based on AS routes that match the Address Range and Match Type criteria get redistributed.
Route: Address Range	The network route to be redistributed into the destination routing protocol. Specify the IPv4 prefix and the mask length.
Route: Match Type	The precision with which the Address Range routes are matched: • Normal (default)- matches routes for the IP addresses and IP subnets in the specified range, including the range itself • Exact - matches the routes for the specified address range exactly as it appears • Refines - matches routes for the IP addresses and IP subnets in the specified range, but not the range itself • Range - matches routes for the IP subnets in the specified range, including the range itself
Action	To redistribute the routes that match the specified criteria or not: • Accept (default) - to redistribute • Restrict - not to redistribute

Parameter	Description
Metric	The cost of the redistributed routes in the destination routing protocol: • RIP - a number in the range 1-16 • OSPFv2 - a number in the range 1-16777215 • BGP - a number in the range 1-4294967295 Note - this parameter is mandatory when configuring redistribution into RIP.
RIP Tag (when redistributing into RIP only)	RIP route tag (optional) - a value in the range 1-65535.
Automatic Tag (when redistributing into OSPFv2 only)	OSPFv2 route tag (optional) - when selected, uses the value specified in Arbitrary Tag to generate the route tag value.
Arbitrary Tag (when redistributing into OSPFv2 only)	Used only when Automatic Tag is selected. Valid values: 0-4095.
Manual Tag (when redistributing into OSPFv2 only)	OSPFv2 route tag (optional). Valid values: 1-2147483647

BGP Default Origin

Redistribute default routes into a specific BGP Autonomous System (AS).

Parameter	Description
To Protocol	The destination BGP AS ID.
Metric	The cost of the redistributed routes in the destination AS - a number in the range 1-4294967295

For redistribution into BGP, you can configure redistribution settings for each Autonomous System.

To configure BGP redistribution settings for each Autonomous System:

1. In the **BGP Redistribution Settings** section, select a BGP AS and click **Edit**.
2. Configure the parameters:
 - **MED** (Multi-Exit Discriminator) - The cost of using this route (0 - 2147483647)
 - **Local Preference** - Local BGP route preference value when routes are redistributed into BGP (0 - 2147483647: the higher the local preference, the more preferred is the route)
3. Click **Save**.

Redistributing IPv6 Routes - WebUI

Configure route sources for redistribution into IPv6 routing protocols.

RIPng

Redistribute RIPng (IPv6) routes.

Parameter	Description
To Protocol	The destination routing protocol. Note - Only distribution into interior routing protocols is possible.
Route: All	Choose which RIPng routes to redistribute into the destination routing protocol. When All is selected, all active RIPng routes get redistributed. When cleared (default), only the RIPng routes that match the Address Range and Match Type criteria get redistributed.
Address Range	The network route to be redistributed into the destination routing protocol. Specify the IPv6 prefix and the mask length.
Match Type	The precision with which the Address Range routes are matched: • Normal (default)- matches routes for the IP addresses and IP subnets in the specified range, including the range itself • Exact - matches the routes for the specified address range exactly as it appears • Refines - matches routes for the IP addresses and IP subnets in the specified range, but not the range itself
Action	To redistribute the routes that match the specified criteria or not: • Accept (default) - to redistribute • Restrict - not to redistribute
Metric	The cost of the redistributed routes in the destination routing protocol: • OSPFv3 - a number in the range 1-16777215

OSPFv3

Redistribute OSPFv3 routes.

Parameter	Description
To Protocol	The destination routing protocol.
Route: All	Choose which OSPFv3 routes to redistribute into the destination routing protocol. When All is selected, all active OSPFv3 routes get redistributed. When cleared (default), only the OSPFv3 routes that match the Address Range and Match Type criteria get redistributed.

Parameter	Description
Address Range	The network route to be redistributed into the destination routing protocol. Specify the IPv6 prefix and the mask length.
Match Type	The precision with which the Address Range routes are matched: • Normal (default)- matches routes for the IP addresses and IP subnets in the specified range, including the range itself • Exact - matches the routes for the specified address range exactly as it appears • Refines - matches routes for the IP addresses and IP subnets in the specified range, but not the range itself
Action	To redistribute the routes that match the specified criteria or not: • Accept (default) - to redistribute • Restrict - not to redistribute
Metric	The cost of the redistributed routes in the destination routing protocol: • RIPng - a number in the range 1-16 Note - this parameter is mandatory when configuring redistribution into RIPng.

OSPFv3 External

Redistribute OSPFv3 External routes.

Parameter	Description
To Protocol	The destination routing protocol.
Route: All	Choose which OSPFv3 External routes to redistribute into the destination routing protocol. When All is selected, all active OSPFv3 routes get redistributed. When cleared (default), only the OSPFv3 routes that match the Address Range and Match Type criteria get redistributed.
Route: Address Range	The network route to be redistributed into the destination routing protocol. Specify the IPv6 prefix and the mask length.
Route: Match Type	The precision with which the Address Range routes are matched: • Normal (default)- matches routes for the IP addresses and IP subnets in the specified range, including the range itself • Exact - matches the routes for the specified address range exactly as it appears • Refines - matches routes for the IP addresses and IP subnets in the specified range, but not the range itself

Parameter	Description
Action	To redistribute the routes that match the specified criteria or not: • Accept (default) - to redistribute • Restrict - not to redistribute
Metric	The cost of the redistributed routes in the destination routing protocol: • RIPng - a number in the range 1-16 Note - this parameter is mandatory when configuring redistribution into RIPng.

Configuring Route Redistribution - CLI

You can configure IPv4 and IPv6 route redistribution through CLI.

Redistributing IPv4 Routes - CLI

To configure IPv4 route redistribution:

```
set route-redistribution to <destination_protocol>
<destination_protocol_parameters> from <source_protocol>
<source_protocol_parameters>
```

These are the destination protocol options:

- BGP
- OSPFv2
- RIP

Each destination protocol has a set of destination protocol parameters.

These are the source protocol options:

- BGP, specific AS
- BGP, specific AS path
- OSPFv2
- OSPFv2 External
- RIP

You can also redistribute:

- Aggregate routes
- Default routes
- Routes from specific interfaces
- Static routes

Destination protocol - BGP:

```
bgp-as <peer_group> [community-append <community_ID> [as <AS_number> | off]]
| [community-match <community_ID> [as <AS_number> | off]] | [localpref
<preference>] | [med <MED>]
```

Parameter	Description
<code>bgp-as <peer_group></code>	Destination BGP Autonomous System (AS).
<code>community-append <community_ID></code>	The ID of the community appended to the route. Valid community ID values are 1 through 65535.
<code>community-match <community_ID></code>	The ID of the community to route. Valid community ID values are 1 through 65535.
<code>as <AS_number> off</code>	Turn the specified community-AS association on (as <AS_number>) or off (off)
<code>localpref <preference></code>	Assign a BGP local preference to routes redistributed into the destination BGP AS. Valid preference values are 0 through 2147483647. The default is none. If selected, it leaves the BGP local preference value of the source BGP AS routes.
<code>med <MED></code>	Assign a Multi Exit Discriminator (MED). The routes with the lowest MED are preferred.

Examples:

- Redistribute all the static routes into the BGP AS 11.

```
set route-redistribution to bgp-as 11 from static-route all on
```
- Assign the BGP local preference of 100 to routes redistributed into BGP AS 11.

```
set route-redistribution to bgp-as 11 localpref 100
```

Destination protocol - OSPFv2:

```
ospf2
```

There are no protocol options for redistribution into OSPFv2.

Example:

Redistribute all routes from the interface eth0 into OSPFv2, and assign the metric of 50 to them.

```
set route-redistribution to ospf2 from interface eth0 metric 50 on
```

Destination protocol - RIP:

```
rip
```

There are no protocol options for redistribution into RIP.

Example:

Redistribute all external OSPFv2 routes for the network 192.168.0.0/16 into RIP, and assign the metric of 3 to them.

```
set route-redistribution to rip from ospf2ase network 192.168.0.0/16
metric 3 on
```

Source - aggregate routes:

```
aggregate {all | <IPv4_address/mask>} {metric <cost> on | on | off}
```

Parameter	Description
<code>aggregate {all <IPv4_address/mask>}</code>	IPv4 aggregate routes to redistribute into the destination routing protocol. Valid options are: <code><IPv4_address/mask></code> - Aggregate routes to the specified network <code>all</code> - All aggregate routes
<code>metric <cost></code>	The cost of the redistributed routes in the destination routing protocol: - RIP - a number in the range 1-16 - OSPFv2 - a number in the range 1-16777215 - BGP - a number in the range 1-4294967295 Note - this parameter is mandatory when configuring redistribution into RIP.

Example:

Redistribute aggregate routes for the network 1.2.0.0/16 into RIP, and assign the cost of 2 to them.

```
set route-redistribution to rip from aggregate 1.2.0.0/16 metric 2 on
```

Source protocol - BGP, specific AS:

```
bgp-as-number <AS_number> {all {metric <cost> on | on | off} | network
<IPv4_address/mask> {action {accept | restrict} | match-type {exact | normal
| range between <lower_IP> and <upper_IP> | refines} on | metric <cost> on
| on | off} | ospf-automatic-tag {on | off} | ospf-automatic-tag-value
{default | <value>} | ospf-manual-tag-value {default | <value>} | riptag
{default | <value>}}
```

Parameter	Description
<code>bgp-as-number <AS_number></code>	BGP Autonomous System number from which to redistribute routes into the destination routing protocol.

Parameter	Description
<pre>all {metric <cost> on on off}</pre>	Redistribute all the BGP routes from the specified BGP AS. The options for redistribution of all routes are: • <code>metric <cost> on</code> - Assign a cost to the redistributed routes and turn the redistribution on • <code>on</code> - Turn the redistribution on • <code>off</code> - Turn the redistribution off Make sure to assign a route redistribution cost that is in a valid for the destination protocol range: • RIP - 1-16 • OSPFv2 - 1-16777215 • BGP - 1-4294967295 Note - Route cost assignment is mandatory when configuring redistribution into RIP.
<pre>network <IPv4_address/mask> {action <action> match-type <precision> on metric <cost> on on off}</pre>	Redistribute the BGP routes to the specified network in the specified BGP AS, unless a more specific route redistribution rule exists. Specify an <code><action></code> to accept or reject the network routes: • <code>accept</code> - Redistribute the routes • <code>restrict</code> - Do not redistribute the routes Specify the <code><precision></code> with which the routes are matched: • <code>exact</code> - Matches only routes to the network • <code>normal</code> - Matches routes within the network, including the network route itself (default) • <code>range</code> - Matches routes that fall within the specified range in the network (the lower limit for the mask length must not be lower than the network mask length) • <code>refines</code> - Matches routes within the network, not including the network route itself Assign a cost to redistributed routes - <code><cost></code>. Make sure to assign a value that is in a valid for the destination protocol range: • RIP - 1-16 • OSPFv2 - 1-16777215 • Note - Route cost assignment is mandatory when configuring redistribution into RIP.
<pre>ospf-automatic-tag {on off}</pre>	Generate tags for external OSPF routes, based on the BGP AS (optional): • <code>on</code> - Turn the tag generation on • <code>off</code> - Turn the tag generation off

Parameter	Description
<code>ospf-automatic-tag-value {default <value>}</code>	Assign a specific value to the automatic tag for the external OSPF routes (optional). Valid tag values are 1 through 4095.
<code>ospf-manual-tag-value {default <value>}</code>	Assign a tag value for the external OSPF routes (optional). This tag overrides all automatic tag configurations. Valid tag values are 1 through 2147483647.
<code>riptag <RIP_tag></code>	Assign a tag value to a BGP route redistributed into RIP. Valid tag values are 1 through 65535.

Examples:

- Redistribute all routes for the IP addresses in the Autonomous System 100 into OSPF (valid for OSPFv2 only), and assign the cost of 99 to them.

```
set route-redistribution to ospf2 from bgp-as-number 100 all metric 99 on
```
- Redistribute all routes for the IP addresses in the Autonomous System 100 into OSPF (valid for OSPFv2 only), except for routes for the network 192.168.0.0/16.

```
set route-redistribution to ospf2 from bgp-as-number 100 network 192.168.0.0/16 action restrict
set route-redistribution to ospf2 from bgp-as-number 100 all on
```

Source protocol - BGP, specific AS path:

```
bgp-aspath <AS-Path_regexp> origin {any | IGP | EGP | incomplete} {all {metric <cost> on | on | off} | network <IPv4_address/mask> {action {accept | restrict} | match-type {exact | normal | range between <lower_IP> and <upper_IP> | refines} on | metric <cost> on | on | off} | ospf-automatic-tag {on | off} | ospf-automatic-tag-value {default | <value>} | ospf-manual-tag-value {default | <value>} | riptag {default | <value>}}
```

Parameter	Description
<AS-Path_regexp>	AS Path, defined through a regular expression. An AS-Path regular expression consists of one or more <i>term-operator</i> pairs. Each <i>term</i> identifies one or more of the Autonomous Systems, and each <i>operator</i> specifies the matching criteria for the corresponding term. An AS-Path Regular Expression <i>term</i> can be defined through single AS numbers and groups of AS numbers. These are the AS-Path Regular Expression <i>operators</i>: • <code>.</code> - Matches any single character • <code>\</code> - Matches the character right after the backslash (to enter <code>\</code> in CLI, type another <code>\</code> immediately after the first one) • <code>^</code> - Matches the characters or a null string at the beginning of an AS path • <code>\$</code> - Matches the characters or a null string at the end of an AS path • <code>?</code> - Matches zero or one occurrence of the preceding term (to enter <code>?</code> in CLI, hit <i>Ctrl-V</i>, then type <code>?</code>) • <code>*</code> - Matches zero or more occurrences of the preceding term • <code>+</code> - Matches one or more occurrences of the preceding term • <code> </code> - Matches one of the terms on either side of the <code> </code> character • <code>_</code> - Matches a non-alphanumeric character (<code>"</code>, <code>,</code>, <code>{</code>, <code>}</code>, <code>^</code>, <code>\$</code>) or a whitespace • <code>[]</code> - Matches a term in a set of terms (separated by whitespaces) or a range of terms (separated by a hyphen) • <code>()</code> - Matches a group of terms as a single term • <code>{m, n}</code> - Matches at least <i>m</i> and at most <i>n</i> repetitions of the preceding term • <code>{m}</code> - Matches exactly <i>m</i> repetitions of the preceding term • <code>{m, }</code> - Matches <i>m</i> or more repetitions of the preceding term
origin {any IGP EGP incomplete}	The origin of the route: • <code>any</code> - Any origin - IGP or EGP • <code>IGP</code> - The route was learned from an interior routing protocol (IGP), and the AS-Path information is probably complete • <code>EGP</code> - The route was learned from an exterior routing protocol (EGP) that does <i>not</i> support AS-Path, and the AS-Path information is probably incomplete • <code>incomplete</code> - The AS-Path information is incomplete Routes from IGPs are preferred over routes from EGPs, and routes from EGPs are preferred over routes with incomplete AS paths.

Parameter	Description
all {metric <cost> on on off}	Redistribute all the BGP routes that match this BGP AS Path. The options for redistribution of all routes are: - metric <cost> on - Assign a cost to the redistributed routes and turn the redistribution on - on - Turn the redistribution on - off - Turn the redistribution off Make sure to assign a value that is in a valid for the destination protocol range: - RIP - 1-16 - OSPFv2 - 1-16777215 Note - Route cost assignment is mandatory when configuring redistribution into RIP.
network <IPv4_address/mask> {action <action> match-type <precision> on metric <cost> on on off}	Redistribute the BGP routes to the specified network that match this BGP AS Path, unless a more specific route redistribution rule exists. Specify an <action> to accept or reject the network routes: - accept - Redistribute the routes - restrict - Do not redistribute the routes Specify the <precision> with which the routes are matched: - exact - Matches only routes to the network - normal - Matches routes within the network, including the network route itself (default) - range - Matches routes that fall within the specified range in the network (the lower limit for the mask length must not be lower than the network mask length) - refines - Matches routes within the network, not including the network route itself Assign a cost to redistributed routes - <cost>. Make sure to assign a value that is in a valid for the destination protocol range: - RIP - 1-16 - OSPFv2 - 1-16777215 Note - Route cost assignment is mandatory when configuring redistribution into RIP.
ospf-automatic-tag {on off}	Generate tags for external OSPF routes, based on the BGP AS (optional): - on - Turn the tag generation on - off - Turn the tag generation off
ospf-automatic-tag-value {default <value>}	Assign a specific value to the automatic tag for the external OSPF routes (optional). Valid values are 1 through 4095.

Parameter	Description
<code>ospf-manual-tag-value {default <value>}</code>	Assign a tag value for the external OSPF routes (optional). This tag overrides all automatic tag configurations. Valid tag values are 1 through 2147483647.
<code>riptag <RIP_tag></code>	Assign a tag value to a BGP route redistributed into RIP. Valid tag values are 1 through 65535.

Examples:

- Redistribute all routes that are learned from BGP AS 100 into RIP, and assign the RIP tag value of 99 to them.

```
set route-redistribution to rip from bgp-aspath ^100_ origin any riptag 99
```
- Redistribute routes for network 192.168.0.0/16 that are originated in BGP AS 100 and learned through any interior routing protocol, into RIP, and assign the cost of 10 to them.

```
set route-redistribution to rip from bgp-aspath _100$ origin IGP network 192.168.0.0/16 metric 10 on
```

Source protocol - default route:

The destination protocol must be BGP.

```
default-origin all {metric <cost> on | on | off}
```

Parameter	Description
<code>all {metric <cost> on on off}</code>	Redistribute the default route. The options for redistribution of the default route are: • <code>metric <cost> on</code> - Assign a cost in the range of 1-4294967295 to the redistributed default route and turn the redistribution on • <code>on</code> - Turn the redistribution on • <code>off</code> - Turn the redistribution off

Example:

Redistribute all default routes into BGP AS 100, and assign the cost of 10 to them.

```
set route-redistribution to bgp-as 100 from default-origin all metric 10 on
```

Source - specific interface:

```
interface {all | <interface1> | <interface2> | ..} {metric <cost> on | on | off}
```

Parameter	Description
<pre>{all <interface1> <interface2> ..} {metric <cost> on on off}</pre>	Redistribute routes from the specified interfaces. The interface options are: • all - all interfaces • <interface#> - specific interface The options for redistribution are: • metric <cost> on - Assign a cost to the redistributed routes and turn the redistribution on • on - Turn the redistribution on • off - Turn the redistribution off Make sure to assign a route redistribution cost that is in a valid for the destination protocol range: • RIP - 1-16 • OSPFv2 - 1-16777215 • BGP - 1-4294967295

Example:

Redistribute all IPv4 routes from the interface Eth2 into OSPFv2.
 set route-redistribution to ospf2 from interface eth2 on

Source protocol - OSPFv2:

```
ospf2 {all {metric <cost> on | on | off} | network <IPv4_address/mask> {action
{accept | restrict} | match-type {exact | normal | range between <lower_IP>
and <upper_IP> | refines} on | metric <cost> on | on | off}}
```

Parameter	Description
<pre>all {metric <cost> on on off}</pre>	Redistribute all the OSPFv2 routes. The options for redistribution of all routes are: • metric <cost> on - Assign a cost to the redistributed routes and turn the redistribution on • on - Turn the redistribution on • off - Turn the redistribution off Make sure to assign a route redistribution cost that is in a valid for the destination protocol range: • RIP - 1-16 • BGP - 1-4294967295 Note - Route cost assignment is mandatory when configuring redistribution into RIP.

Parameter	Description
<code>network <IPv4_address/mask> {action <action> match-type <precision> on metric <cost> on on off}</code>	Redistribute the OSPFv2 routes to the specified network, unless a more specific route redistribution rule exists. Specify an <action> to accept or reject the network routes: • accept - Redistribute the routes • restrict - Do not redistribute the routes Specify the <precision> with which the routes are matched: • exact - Matches only routes to the network • normal - Matches routes within the network, including the network route itself (default) • range - Matches routes that fall within the specified range in the network (the lower limit for the mask length must not be lower than the network mask length) • refines - Matches routes within the network, not including the network route itself Assign a cost to redistributed routes - <cost>. Make sure to assign a value that is in a valid for the destination protocol range: • RIP - 1-16 • BGP - 1-4294967295 Note - Route cost assignment is mandatory when configuring redistribution into RIP.

Example:

Redistribute all routes for the network 192.168.0.0/16 from OSPF into RIP, and assign the cost of 2 to them.

```
set route-redistribution to rip from ospf2 network 192.168.0.0/16 metric 2 on
```

Source protocol - OSPFv2 External:

```
ospf2ase {all {metric <cost> on | on | off} | network <IPv4_address/mask>  
{action {accept | restrict} | match-type {exact | normal | range between  
<lower_IP> and <upper_IP> | refines} on | metric <cost> on | on | off} | riptag  
{default | <value>}}
```

Parameter	Description
<code>all {metric <cost> on on off}</code>	Redistribute all the OSPFv2 External routes. The options for redistribution of all routes are: • <code>metric <cost> on</code> - Assign a cost to the redistributed routes and turn the redistribution on • <code>on</code> - Turn the redistribution on • <code>off</code> - Turn the redistribution off Make sure to assign a value that is in a valid for the destination protocol range: • RIP - 1-16 • BGP - 1-4294967295 Note - Route cost assignment is mandatory when configuring redistribution into RIP.
<code>network <IPv4_address/mask> {action <action> match-type <precision> on metric <cost> on on off}</code>	Redistribute the OSPFv2 External routes to the specified network, unless a more specific route redistribution rule exists. Specify an <code><action></code> to accept or reject the network routes: • <code>accept</code> - Redistribute the routes • <code>restrict</code> - Do not redistribute the routes Specify the <code><precision></code> with which the routes are matched: • <code>exact</code> - Matches only routes to the network • <code>normal</code> - Matches routes within the network, including the network route itself (default) • <code>range</code> - Matches routes that fall within the specified range in the network (the lower limit for the mask length must not be lower than the network mask length) • <code>refines</code> - Matches routes within the network, not including the network route itself Assign a cost to redistributed routes - <code><cost></code>. Make sure to assign a value that is in a valid for the destination protocol range: • RIP - 1-16 • BGP - 1-4294967295 Note - Route cost assignment is mandatory when configuring redistribution into RIP.
<code>riptag <RIP_tag></code>	Assign a tag value to an OSPFv2 External route redistributed into RIP. Valid tag values are 1 through 65535.

Examples:

- Redistribute all OSPF external routes into RIP, and assign the cost of 2 to them.
set route-redistribution to rip from ospf2ase all metric 2 on
- Redistribute all OSPF external routes into RIP, and assign the RIP tag value of 20 to them.
set route-redistribution to rip from ospf2ase riptag 20

Source protocol - RIP:

```
rip {all {metric <cost> on | on | off} | network <IPv4_address/mask> {action
{accept | restrict} | match-type {exact | normal | range between <lower_IP>
and <upper_IP> | refines} on | metric <cost> on | on | off}}
```

Parameter	Description
all {metric <cost> on on off}	Redistribute all the RIP routes. The options for redistribution of all routes are: • metric <cost> on - Assign a cost to the redistributed routes and turn the redistribution on • on - Turn the redistribution on • off - Turn the redistribution off Make sure to assign a value that is in a valid for the destination protocol range: • OSPFv2 - 1-16777215 • BGP - 1-4294967295
network <IPv4_address/mask> {action <action> match-type <precision> on metric <cost> on on off}	Redistribute the RIP routes to the specified network, unless a more specific route redistribution rule exists. Specify an <action> to accept or reject the network routes: • accept - Redistribute the routes • restrict - Do not redistribute the routes Specify the <precision> with which the routes are matched: • exact - Matches only routes to the network • normal - Matches routes within the network, including the network route itself (default) • range - Matches routes that fall within the specified range in the network (the lower limit for the mask length must not be lower than the network mask length) • refines - Matches routes within the network, not including the network route itself Assign a cost to redistributed routes - <cost>. Make sure to assign a value that is in a valid for the destination protocol range: • OSPFv2 - 1-16777215 • BGP - 1-4294967295

Example:

Redistribute all RIP routes for the network 192.168.0.0/16 that fall in the range between 192.168.0.0/18 and 192.168.0.0/24 into OSPF (valid for OSPFv2 only).

```
set route-redistribution to ospf2 from rip network 192.168.0.0/16
match-type range between 18 and 24 on
```

Source - static routes:

```
static-route {all | default} {metric <cost> on | on | off}
```

Parameter	Description
<code>static {all default}</code>	IPv4 static routes to redistribute into the destination routing protocol. Valid options are: <code>default</code> - Default static route <code>all</code> - All static routes
<code>metric <cost> on</code>	Assign a cost to redistributed routes - <code><cost></code> . Make sure to assign a value that is in a valid for the destination protocol range: - RIP - 1-16 - OSPFv2 - 1-16777215 - BGP - 1-4294967295 Note - Route cost assignment is mandatory when configuring redistribution into RIP.

Example:

Redistribute all IPv4 static routes into BGP AS 100, and assign the cost of 1 to them.

```
set route-redistribution to bgp-as 100 from static-route all metric 1 on
```

Redistributing IPv6 Routes - CLI

To configure IPv6 route redistribution:

```
set ipv6 route-redistribution to <destination_protocol> {from
<source_protocol> <source_protocol_parameters> | off}
```

These are the destination protocol options:

- OSPFv3
- RIPng

Each destination protocol has a set of destination protocol parameters.

These are the source protocol options:

- BGP
- OSPFv3
- OSPFv3 External
- RIPng

You can also redistribute:

- Routes from specific interfaces
- Static routes

Source protocol - OSPFv3:

```
ospf3 {all {metric <cost> on | on | off} | network <IPv6_address/mask> {action {accept | restrict} | match-type {exact | normal | range between <lower_IP> and <upper_IP> | refines} on | metric <cost> on | on | off}}
```

Parameter	Description
<code>all {metric <cost> on on off}</code>	Redistribute all the OSPFv3 routes. The options for redistribution of all routes are: • <code>metric <cost> on</code> - Assign a cost in the range of 1-16 to the redistributed routes and turn the redistribution on • <code>on</code> - Turn the redistribution on • <code>off</code> - Turn the redistribution off Note - Route cost assignment is mandatory when configuring redistribution into RIP.
<code>network <IPv6_address> {action <action> match-type <precision> on metric <cost> on on off}</code>	Redistribute the OSPFv3 routes to the specified network, unless a more specific route redistribution rule exists. Specify an <code><action></code> to accept or reject the network routes: • <code>accept</code> - Redistribute the routes • <code>restrict</code> - Do not redistribute the routes Specify the <code><precision></code> with which the routes are matched: • <code>exact</code> - Matches only routes to the network • <code>normal</code> - Matches routes within the network, including the network route itself (default) • <code>refines</code> - Matches routes within the network, not including the network route itself Assign a cost in the range of 1-16 to the redistributed routes. Note - Route cost assignment is mandatory when configuring redistribution into RIP.

Example:

Redistribute all routes for the network `fd14:8502:5b14:456a::/64` from OSPFv3 into RIPng, and assign the cost of 999 to them.

```
set ipv6 route-redistribution to ripng from ospf3 network
fd14:8502:5b14:456a::/64 metric 999 on
```

Source protocol - OSPFv3 External:

```
ospf3ase {all {metric <cost> on | on | off} | network <IPv6_address/mask> {action {accept | restrict} | match-type {exact | normal | range between <lower_IP> and <upper_IP> | refines} on | metric <cost> on | on | off}}
```

Parameter	Description
<code>all {metric <cost> on on off}</code>	Redistribute all the OSPFv3 External routes. The options for redistribution of all routes are: • <code>metric <cost> on</code> - Assign a cost in the range of 1-16 to the redistributed routes and turn the redistribution on • <code>on</code> - Turn the redistribution on • <code>off</code> - Turn the redistribution off Note - Route cost assignment is mandatory when configuring redistribution into RIP.
<code>network <IPv6_address/mask> {action <action> match-type <precision> on metric <cost> on on off}</code>	Redistribute the OSPFv3 External routes to the specified network, unless a more specific route redistribution rule exists. Specify an <code><action></code> to accept or reject the network routes: • <code>accept</code> - Redistribute the routes • <code>restrict</code> - Do not redistribute the routes Specify the <code><precision></code> with which the routes are matched: • <code>exact</code> - Matches only routes to the network • <code>normal</code> - Matches routes within the network, including the network route itself (default) • <code>refines</code> - Matches routes within the network, not including the network route itself Assign a cost in the range of 1-16 to the redistributed routes. Note - Route cost assignment is mandatory when configuring redistribution into RIP.

Examples:

- Redistribute all OSPFv3 External routes into RIPng, and assign the cost of 22 to them.
`set ipv6 route-redistribution to ripng from ospf3ase all metric 22 on`
- Do not redistribute routes for the network fd14:8502:5b14:456a::/64 from OSPFv3 External into RIPng.
`set ipv6 route-redistribution to ripng from ospf3ase network fd14:8502:5b14:456a::/64 action restrict`

Source protocol - RIPng:

```
rip {all {metric <cost> on | on | off} | network <IPv6_address/mask> {action {accept | restrict} | match-type {exact | normal | range between <lower_IP> and <upper_IP> | refines} on | metric <cost> on | on | off}}
```

Parameter	Description
<code>all {metric <cost> on on off}</code>	Redistribute all the RIPv6 routes. The options for redistribution of all routes are: • <code>metric <cost> on</code> - Assign a cost in the range of 1-16777215 to the redistributed routes and turn the redistribution on • <code>on</code> - Turn the redistribution on • <code>off</code> - Turn the redistribution off
<code>network <IPv6_address/mask> {action <action> match-type <precision> on metric <cost> on on off}</code>	Redistribute the RIPv6 routes to the specified network, unless a more specific route redistribution rule exists. Specify an <code><action></code> to accept or reject the network routes: • <code>accept</code> - Redistribute the routes • <code>restrict</code> - Do not redistribute the routes Specify the <code><precision></code> with which the routes are matched: • <code>exact</code> - Matches only routes to the network • <code>normal</code> - Matches routes within the network, including the network route itself (default) • <code>refines</code> - Matches routes within the network, not including the network route itself Assign a cost in the range of 1-16777215 to the redistributed routes.

Example:

Redistribute all RIPv6 routes for the network fd14:8502:5b14:456a::/64, including the routes for the addresses within the network into OSPFv3.

```
set ipv6 route-redistribution to ospf3 from rip network
fd14:8502:5b14:456a::/64 match-type normal on
```

Source - specific interface:

```
interface {all | <interface1> | <interface2> | ..} {metric <cost> on | on  
| off}
```

Parameter	Description
<code>{all <interface1> <interface2> ..} {metric <cost> on on off}</code>	Redistribute routes from the specified interfaces. The interface options are: • <code>all</code> - all interfaces • <code><interface#></code> - specific interface The options for redistribution are: • <code>metric <cost> on</code> - Assign a cost to the redistributed routes and turn the redistribution on • <code>on</code> - Turn the redistribution on • <code>off</code> - Turn the redistribution off Make sure to assign a route redistribution cost that is in a valid for the destination protocol range: • RIPng - 1-16 • OSPFv3 - 1-16777215

Example:

Redistribute all IPv6 routes from the interface Eth0 into OSPFv3 and assign the cost of 10 to them.

```
set ipv6 route-redistribution to ospf3 from interface eth0 metric 10 on
```

Source - static routes:

```
static-route {all | default} {metric <cost> on | on | off}
```

Parameter	Description
<code>static-route {all default}</code>	IPv6 static routes to redistribute into the destination routing protocol. Valid options are: • <code>default</code> - Default static route • <code>all</code> - All static routes
<code>metric <cost> on</code>	Assign a cost to redistributed routes - <code><cost></code>. Make sure to assign a value that is in a valid for the destination protocol range: • RIPng - 1-16 • OSPFv3 - 1-16777215 Note - Route cost assignment is mandatory when configuring redistribution into RIP.

Example:

Redistribute the default IPv6 static route into OSPFv3, and assign the cost of 15000 to it.

```
set route-redistribution to ospf3 from static-route default metric 15000 on
```

Configuring Route Maps - CLI (routemap)

Route maps support both IPv4 and IPv6 protocols, which includes RIP, RIPv6, BGP, OSPFv2, and OSPFv3. You can only define BGP-4 Multiprotocol Extensions policy with route maps. For the other protocols, you can use route maps or the Route Redistribution and Inbound Route Filters features that you configure on the WebUI.

Each route map includes a list of criteria and statements. You can apply route maps to inbound, outbound, or redistribution routes. Routes are compared to the match criteria, and all the actions defined in the criteria are applied to those routes which match all the criteria. You can set the match criteria in any order. If you do not define match criteria in a route map, the route map matches all routes.

You define route maps, then assign them to protocols for export or import policy for that protocol. Route maps override WebUI based configuration.

To create a route map, use CLI commands to define a set of criteria that must be matched for the command to run. If the criteria are matched, then the system runs the actions you define. A route map is identified by a name and a number, an Allow or Restrict clause, and a collection of match and set statements.

There can be more than one instance of a route map (same name, different ID). The lowest numbered instance of a route map is checked first. Route map processing stops when all the criteria of a route map instance are matched, or all the instances of a route map are exhausted. If the criteria are matched, the actions in the section are run.

Routing protocols can use more than one route map when you set clear preference values for each. The applicable route map with lowest preference value is checked first.

Set Route Map Commands

To set a route map:

```
set routemap <rm_name> id {default | <1-65535>}
    {off|on}
    allow
    inactive
    restrict
```

Parameter	Description
routemap <rm_name>	The name of the route map.
id {default <1-65535>}	Route map ID: - You can enter the keyword <code>default</code>. Default value is 10, or - Enter a value in the range of 1-65535.
{off on}	<code>on</code> to create a route map, <code>off</code> to delete a route map.
allow	Allow routes that match the route map.
inactive	Temporarily disable a route map. To activate the route map, use the <code>allow</code> or <code>restrict</code> arguments.
restrict	Do not allow routes that match the route map.

To configure actions for a routemap:

Note - Some statements have an effect on some protocols only ("Supported Route Map Statements by Protocol" on page 169).

The same parameter cannot appear both as a match and as an action statement in a route map. These include Community, Metric, and Nexthop.

```
set routemap <rm_name> id {default | {1-65535}} action
    aspath-prepend-count {1-25}
    community {append | replace | delete} {on|off}
    community {1-65535} as {1-65535} {on|off}
    community no-export {on|off}
    community no-advertise {on|off}
    community no-export-subconfed {on|off}
    community none {on|off}
    localpref {0-65535}
    metric {add | subtract} {1-4294967295}
    metric igp {add | subtract} {1-4294967295}
    metric value {1-4294967295}
    nexthop ip {ipv4_address}
    nexthop ipv6 {ipv6_address}
    precedence {0-65535}
    preference {0-65535}
    route-type {type-1|type-2}
    remove <action_name>
    ospfautomatictag {0-4095}
    ospfmanualtag {1-4294967295}
    riptag {1-65535}
```

Parameter	Description
routemap <rm_name>	Name of route map.
id {default {1-65535}}	Route map ID: - You can enter the keyword default. Default value is 10, or - Enter a value in the range of 1-65535.
aspath-prepend-count {1-25}	The number of times the local AS number is added to the beginning of the AS path. The number is added when routes which match the route map are advertised through BGP. Only applicable to BGP and export of routes with export-routemap. The action is otherwise ignored.
community {append replace delete} {on off}	Operates on a BGP community string. You can form a community string with these community action statements: append, replace, or delete. The default operation is append. BGP only. If the action is set to replace, the Communities attribute is replaced with this Community ID and AS number for matching BGP routes. If the action is set to append, then the given Community ID and AS number are added to the Communities attribute. See also: set routemap VALUE id VALUE action community append set routemap VALUE id VALUE action community replace

Parameter	Description
<code>community {1-65535} as {1-65535} {on off}</code>	A BGP community value.
<code>community no-export {on off}</code>	This also applies to stand-alone Autonomous Systems which are not part of a Confederation. Only Applicable to BGP and export of routes with the <code>export - routemap</code> command. The action is otherwise ignored.
<code>community no-advertise {on off}</code>	This action only applies to BGP, and only to export of routes through the <code>export - routemap</code> command. The action is otherwise ignored.
<code>community no-export-subconfed {on off}</code>	This action only applies to BGP and export of routes through the <code>export - routemap</code> command. The action is otherwise ignored.
<code>community none {on off}</code>	In action statement, this statement makes sense only if used with <code>replace</code>. This deletes all communities associated with a route so that the route has no communities associated with it. If it is used with <code>append</code> or <code>delete</code>, there is no operation. The CLI returns an error if you turn <code>none</code> on and other community values are already defined, or if <code>none</code> is defined and you add another community value.
<code>localpref {0-65535}</code>	Set the local preference for BGP routes which match the route map. This action only applies to the BGP protocol, and only to route maps with the <code>import - routemap</code> command.
<code>metric {add subtract} {1-4294967295}</code>	Increments the metric of matching routes by the given amount. This action only applies to: - Export of OSPF/OSPFv3 to BGP with <code>export - routemap</code> - Export of RIP/RIPng routes to BGP with <code>export - routemap</code> - Import of RIP/RIPng routes from other routers with <code>import - routemap</code> This action otherwise has no effect. For example, in the import of OSPF/OSPFv3 routes from other routers. In export of routes to BGP, the range of the resulting values is 0-4294967295. In import of routes from RIP, the range of the resulting values is 1-65535. If the value is 16 or higher, the route is treated as unreachable and it is not installed in the kernel routing table. In import of routes from OSPF, the range of the resulting values is 0-65535.
<code>metric igp {add subtract} {1-4294967295}</code>	This action only applies in export of OSPF/OSPFv3 or RIP/RIPng routes to BGP with <code>export - routemap</code> . This action has no effect otherwise. The range of values for the resulting metric is 0-4294967295.

Parameter	Description
metric value {0-4294967295}	Sets the metric of matching routes to the given value. For RIP the metric is <i>metric</i>, for OSPF the metric is <i>cost</i>, and for BGP the metric is <i>MED</i>. The actual range is different for each protocol. In export to RIP or IPv6 RIP (RIPng), this action sets the RIP metric. The range of values is 1 - 65535, where 16 or larger shows that the route is unreachable. In export to OSPF or IPv6 OSPF (OSPFv3), this action sets the OSPF route cost. The range of values is 0 - 65535. In export to BGP, this action affects the BGP MED. The range of values is 0 - 4294967295. This action only applies to route maps with the <code>export - routemap</code> command. It has no effect otherwise.
nexthop ip <ipv4_address>	Sets the IPv4 next hop address for routes that match this Route Map ID. This action only applies to import or export of BGP routes.
precedence {0-65535}	Sets the priority of the routes which match the route map ID. Use this setting to give priority to routes of one protocol over the other. The route with the lower value has priority. The other routes are shown as inactive and are not installed in the kernel. This action only applies to the <code>import - routemap</code> command. See also: <code>set protocol-rank</code> <code>set static-route VALUE rank</code>
preference {0-65535}	Applies only to routes received through BGP advertisements. This is equivalent to the <code>bgp weight</code> (in Cisco terms) of the route. However, unlike Cisco, the route with the lower value has precedence. This value only applies to the local router.
route-type {type-1 type-2}	Type of OSPF external route. The metric type of AS External route is set to the specified value. Only applies when used with <code>export - routemap</code> to export routes to other routers through OSPF or OSPFv3.

Parameter	Description
<code>remove <action_name></code>	Remove the specified action from the route map. For community, it removes all community statements. Permitted values: • aspath-prepend-count • community • localpref • metric • nexthop • ospfautomatictag • ospfmanualtag • precedence • preference • riptag • route-type
<code>ospfautomatictag {0-4095}</code>	Creates an automatic tag for OSPF external routes that match the route map. This action only applies to export of non-OSPF routes into OSPF with <code>export - routemap</code>. See RFC 1403 for more information on OSPF tags.
<code>ospfmanualtag {1 - 4294967295}</code>	Creates a manual tag for OSPF external routes that match the route map. This action only applies to export of non-OSPF routes into OSPF with <code>export - routemap</code>. See RFC 1403 for more information on OSPF tags.
<code>riptag {1-65535}</code>	Creates a RIP tag for external routes that match the given Route Map ID. This action only applies to export of non-RIP routes into RIP with <code>export - routemap</code>. See RFC 2453 for more information on RIP route tags.

To configure the route map criteria:



Note - Some statements have an effect on some protocols only ("[Supported Route Map Statements by Protocol](#)" on page 169).

The same parameter cannot show both as a match and as an action statement in a route map. These include Community, Metric, and Nexthop.

```
set routemap <rm_name> id {default | {1-65535}} match
as {1-65535} {on|off}
aspath-regex {{regular_expression} | empty} origin {any | egp | igp | incomplete}
community {1-65535} as {1 - 65535} {on|off}
community exact {on|off}
community no-export {on|off}
community no-advertise {on|off}
community no-export-subconfed {on|off}
community none {on|off}
ifaddress <interface_address> {on|off}
interface <interface_name> {on|off}
metric value {0-4294967295}
```

```

neighbor <ip_address> {on|off}
network <network>/<mask_length> all [restrict]
network <network>/<mask_length> exact [restrict]
network <network>/<mask_length> refines [restrict]
network <network>/<mask_length> between <mask_length> and
<mask_length>[restrict]
network <network>/<mask_length> off
nexthop <ip_address> {on|off}
protocol <protocol_name>
route-type {type-1 | type-2 | inter-area | intra-area} {on|off}
remove <match_condition_name>
tag <tag_name>

```

Parameter	Description
as <1-65535> {on off}	Configures the route map to only match routes which are received from or advertised to the given BGP AS number. Applies to BGP only. You can configure multiple AS match conditions for a given route map ID. A match occurs when one of the AS conditions matches a given route. This match condition applies to both the <code>import - routemap</code> and <code>export - routemap</code> commands, but only for BGP.
aspath-regex {<regular-expression> empty} origin {any egp igp incomplete}	aspath-regex Match the specified aspath regular expression. For BGP only. The regular expression can have only digits, metacharacters ("Regular Expression Syntax" on page 232) and clish special characters ("Special Characters in Clish" on page 232). origin Configures the route origin type for the match condition. This match condition applies to both <code>import - routemap</code> and <code>export - routemap</code>, but only for BGP.
community <1-65535> as <1-65535> {on off}	Configures the route map to match the given BGP Community ID and Community AS number. This match condition applies to both <code>import - routemap</code> and <code>export - routemap</code>, but only for BGP. If you configure multiple community match conditions under a given route map ID, then a given BGP route must match each condition to match the route map.

Parameter	Description
<code>community exact {on off}</code>	Matches the communities in the route to the communities in the route map. If there is no exact clause, the route can have other community values related to it in addition to the ones in the route map. You can have multiple community statements in a route map to form a community string. The Community String(s) to be matched must be added to this Route Map ID with this command: <pre>set routemap <Route Map Name> id <Route Map ID> match\ community <Community-ID> as <AS-Number> on</pre> This match condition applies only to BGP, both for <code>import - routemap</code> and <code>export - routemap</code>.
<code>community no-export {on off}</code>	This option also applies to stand-alone Autonomous Systems which are not part of a Confederation. This match condition applies only to BGP, both for <code>import - routemap</code> and <code>export - routemap</code>.
<code>community no-advertise {on off}</code>	This match condition applies only to BGP, both for <code>import - routemap</code> and <code>export - routemap</code>.
<code>community no-export - subconfed {on off}</code>	For Confederations, routes with this value are not advertised to peers with a different Routing Domain Identifier. This match condition applies only to BGP, both for <code>import - routemap</code> and <code>export - routemap</code>.
<code>community none {on off}</code>	Matches an empty community string, namely, a route with no communities related to it. The CLI returns an error if you turn <i>none</i> on and other community values are already defined, or if <i>none</i> is defined and you add some other community value. This match condition applies only to BGP, both for <code>import - routemap</code> and <code>export - routemap</code>.
<code>ifaddress <interface_address> {on off}</code>	Configures the route map to match the specified interface address. Value: an IPv4 or IPv6 interface address For example: 192.168.2.14 The given IP address is matched against the address of the interface which received the route. There can be multiple match conditions for interface address under the same route map ID. This type of match condition applies to route maps which are used with <code>import - routemap</code> and to route maps which are used with <code>export - routemap</code>.

Parameter	Description
<code>interface</code> <code><interface_name></code> <code>{on off}</code>	Match the route if the nexthop lies on the specified interface name. There can be multiple interface statements.
<code>metric value</code> <code><0-4294967295></code>	Configures the route map to match routes with a specified metric. This match condition applies RIP, BGP, and OSPF. For RIP and IPv6 RIP (RIPng), this matches the RIP metric. The valid range of values is 1-65535. If the value is 16 or higher, the route is unreachable. For OSPF and IPv6 OSPF (OSPFv3), this value matches the route cost. The valid range of values is 0 - 65535. For BGP, this value matches the MED. The valid range of values is 0 - 4294967295.
<code>neighbor <ip_address></code> <code>{on off}</code>	Configures the Route Map to match routes from the given neighbor. Value: an IPv4 or IPv6 network address For example: 192.168.2.14 This match rule only applies to BGP or RIP and only with <code>import - routemap</code>. You can configure multiple neighbor match conditions for a given route map ID.
<code>network</code> <code><network>/<mask_length> all [restrict]</code>	Configures the route map to match all routes which are equal to or contained within the specified IPv4 or IPv6 subnet. The subnet is expressed by CIDR notation. Examples: 192.168.2.0/24, fc00:1:2:3::/64 If <i>restrict</i> is specified, it prevents import or export of matching routes.
<code>network</code> <code><network>/<mask_length> exact [restrict]</code>	Configures the route map to only match routes with the same prefix and mask length as the given IPv4 or IPv6 subnet. The subnet is expressed by CIDR notation. Examples: 192.168.2.0/24, fc00:1:2:3::/64 If <i>restrict</i> is specified, it prevents import or export of matching routes.
<code>network</code> <code><network>/<mask_length> refines [restrict]</code>	Configures the route map to match routes which are contained within the given IPv4 or IPv6 subnet. Routes which match the given subnet exactly (namely, which have the same mask length) are excluded from the match. The subnet is expressed by CIDR notation. Examples: 192.168.2.0/24, fc00:1:2:3::/64 If <i>restrict</i> is specified, it prevents import or export of matching routes.

Parameter	Description
<code>network</code> <code><network>/<mask_length> between</code> <code><mask_length> and</code> <code><mask_length></code> <code>[restrict]</code>	Configures the Route Map to match routes that are within the given IPv4 or IPv6 subnet and which have a mask length that is between the given range of values. The subnet is expressed by CIDR notation Examples: 192.168.2.0/24, fc00:1:2:3::/64 If <i>restrict</i> is specified, it prevents import or export of matching routes.
<code>network</code> <code><network>/<mask_length> off</code>	Deletes the network match statement.
<code>nexthop <ip_addr></code> <code>{on off}</code>	Configures the route map to match routes with a given IPv4 or IPv6 next hop gateway address. For example: 192.168.2.14 You can configure multiple next hop match conditions for a given route map ID. A match occurs if a next hop value matches a given route. This match only applies to BGP and only with <code>export - routemap</code>.
<code>protocol</code> <code><protocol_name></code>	Configures the route map to match routes of the given protocol type. Use this for route redistribution between protocols. <code><protocol_name></code>: • static - static routes • direct - interface routes • aggregate - aggregate routes • rip - RIP routes • ripng - RIPng (IPv6 RIP) routes • ospf2 - OSPFv2 routes • ospf2ase - OSPFv2 external routes • ospf3 - OSPFv3 (IPv6 OSPF) routes • ospf3ase - OSPFv3 external routes • bgp - BGP routes • kernel - kernel (injected) routes This match condition only applies to <code>export - routemap</code> and is ignored for <code>import-routemap</code>. You can define one protocol match condition for a given route map ID. If you do not define a match condition, the default behavior is to export only RIP routes to the RIP protocol, only BGP routes to the BGP protocol, and so on.

Parameter	Description
<code>route-type {type-1 type-2 inter-area intra-area} {on off}</code>	Configures the route map to match the given route type. This match condition applies only to <code>export - routemap</code>, in the export of routes to other routers through OSPF or OSPFv3. If you define a route type of inter-area or intra-area, set the protocol match condition to <code>ospf2</code>. If you define a route type of type-1 or type-2, set the protocol match condition to <code>ospf2ase</code>. In the export of OSPF ASE routes to other protocols, if the metric match condition is set but the route type match condition is not set, it will try to match the metric value for both type-1 and type-2 routes. There can be multiple route type match conditions. Best Practice - Do not use this match condition simultaneously with the <code>route-type</code> action.
<code>remove <match_condition_name></code>	Removes the specified match condition from the routemap. If a match condition has multiple match statements (such as network, neighbor), this argument removes all of them. Permitted values: • as • aspath-regex • community • community-regex • ifaddress • interface • metric • neighbor • network • nexthop • protocol • route-type • tag
<code>tag <tag_name></code>	Matches OSPF external routes with the given tag value. Value: 1 - 4294967295 You can add multiple tag match conditions to a given route map ID to broaden the range of matched tag values. Currently, you can only use this feature to export OSPF routes to BGP.

Show Routemap Commands

```
show routemap rm_name {all | <id VALUE>}
show routemaps
```

Routemap Protocol Commands

To assign routemaps to protocols:

The preference value specifies which order the protocol will use each routemap.

```
set {ospf | ipv6 ospfv3 | rip | ipv6 ripng}
```

```
export-routemap <rm_name> preference VALUE on
import-routemap <rm_name> preference VALUE on
```

To turn a routemap off:

```
set {ospf | ipv6 ospfv3 | rip | ipv6 ripng}
    export-routemap <rm_name> off
    import-routemap <rm_name> off
```

To view routemaps assigned to protocols:

```
show {ospf | ipv6 ospfv3 | rip | ipv6 ripng | bgp) routemap
```

To set BGP routemaps for export and import policies:

```
set bgp external remote-as <1-65535> export-routemap <rm_name> off
    preference <1-65535> [family {inet | inet6 | inet-and-inet6}] on
```

```
set bgp external remote-as {1-65535} import-routemap <rm_name> off
    preference <1-65535> [family {inet | inet6 | inet-and-inet6}] on
```

```
set bgp internal export-routemap <rm_name> off
    preference <1-65535> [family {inet | inet6 | inet-and-inet6}] on
```

```
set bgp internal import-routemap <rm_name> off
    preference <1-65535> [family {inet | inet6 | inet-and-inet6}] on
```

Note - You cannot use route maps in BGP confederations. To configure route filters and redistribution for BGP confederations, use the Inbound Route Filters and Route Redistribution pages in the WebUI.

Supported Route Map Statements by Protocol

Some statements affect only a particular protocol, for example, matching the Autonomous System Number is applicable only to BGP. If such a condition is in a routemap used by OSPF, the match condition is ignored. Any non-applicable match conditions or actions are ignored and processing is done as if they do not exist. A log message is generated in /var/log/messages for any such statements.



Note - The same parameter cannot appear both as a match and action statement in a routemap. These include Community, Metric, and Nexthop.

RIP

- Import Match conditions: Neighbor, Network, Interface, Ifaddress, Metric, Neighbor, Nexthop.
- Import Actions: Precedence, Metric Add/Subtract
- Export Match conditions when exporting from RIP - Interface, Ifaddress, Metric, Network, Nexthop
- Export Match Conditions when redistributing using Protocol match: According to the protocol from which route is being redistributed.
- Export Actions when exporting from RIP - Metric Add/Subtract
- Export Actions when redistributing - Metric Set

OSPFv2

- Import Match conditions: **Network** (Route Prefix)
- Import Actions: **Precedence**
- Export Match conditions when other protocols redistribute OSPF routes: **Network, Interface, Ifaddress, Metric, Route-type, Nexthop**
- Export Match conditions when OSPF redistributes routes from other protocols: Conditions supported by that protocol
- Export Actions when redistributing to AS External: **Metric, Route-type**

BGP

When you do initial configuration, set the router ID. You can also use the following commands to change the router ID.

```
set router-id default
set router-id ip_address
```

Command	Description
default	Selects the highest interface address when OSPF is enabled.
ip_address	The Router ID uniquely identifies the router in the autonomous system. The router ID is used by the BGP and OSPF protocols. We recommend setting the router ID rather than relying on the default setting. This prevents the router ID from changing if the interface used for the router ID goes down. Use an address on a loopback interface that is not the loopback address (127.0.0.1). Note - In a cluster, you must select a router ID and make sure that it is the same on all cluster members. • Range - Dotted-quad. 9[0-255].[0-255].[0-255]. Do not use 0.0.0.0 • Default - The interface address of one of the local interfaces.

Use the following group of commands to set and view parameters for BGP.

```
set as as_number
set as off
```

Parameter	Description
as as_number	The local autonomous system number of the router. This number is mutually exclusive from the confederation and routing domain identifier. The router can be configured with either the autonomous system number or confederation number, not both. Caution: When you change the autonomous system number, all current peer sessions are reset and all BGP routes are deleted.
as off	Disables the configured local autonomous system number.

Redistributing Static, Interface, or Aggregate Routes

When redistributing **static** routes into BGP, OSPFv2 or RIP the following match conditions are supported:

- Network Prefix,
- Nexthop
- Interface
- Ifaddress
- Protocol (proto = static)

When redistributing **interface/direct** routes into BGP, OSPFv2 or RIP the following match conditions are supported:

- Network Prefix
- Interface
- Ifaddress
- Protocol (proto = direct)

When redistributing **aggregate** routes into BGP, OSPFv2 or RIP the following match conditions are supported:

- Network Prefix
- Protocol (proto = aggregate)

Route Map Examples

Example 1

Redistribute interface route for eth3c0 into ospf, and set the ospf route-type to AS type-2 with cost 20.

```
set routemap direct-to-ospf id 10 on
set routemap direct-to-ospf id 10 match interface eth3c0
set routemap direct-to-ospf id 10 match protocol direct
set routemap direct-to-ospf id 10 action route-type type-2
set routemap direct-to-ospf id 10 action metric value 20
```

```
set ospf export-routemap direct-to-ospf preference 1 on
```

Example 2

Do not accept routes from RIP neighbor 192.0.2.3, accept routes from neighbor 192.0.2.4 as is, and for all other routes increment the metric by 2.

```
set routemap rip-in id 10 on
set routemap rip-in id 10 restrict
set routemap rip-in id 10 match neighbor 192.0.2.3
```

```
set routemap rip-in id 15 on
set routemap rip-in id 15 match neighbor 192.0.2.4
```

```
set routemap rip-in id 20 on
set routemap rip-in id 20 action metric add 2
```

```
set rip import-routemap rip-in preference 1 on
```

Example 3

Redistribute all static routes into BGP AS group 400. Set the MED value to 100, prepend our AS number to the aspath 4 times. If the route belongs to the prefix 192.0.2.0/8, do not redistribute. Send all BGP routes whose aspath matches the regular expression (100 200+) and set the MED value to 200.

```
set routemap static-to-bgp id 10 on
set routemap static-to-bgp id 10 restrict
set routemap static-to-bgp id 10 match protocol static
set routemap static-to-bgp id 10 match network 192.0.2.0/8 all
```

```
set routemap static-to-bgp id 15 on
set routemap static-to-bgp id 15 match protocol static
set routemap static-to-bgp id 15 action metric 100
set routemap static-to-bgp id 15 action aspath-prepend-count 4
```

```
set routemap bgp-out id 10 on
set routemap bgp-out id 10 match aspath-regex "(100 200+)" origin any
set routemap bgp-out id 10 action metric 200
```

```
set bgp external remote-as 400 export-routemap bgp-out preference 1 family inet on
set bgp external remote-as 400 export-routemap static-to-bgp preference 2 family inet on
```



Note - There is no need for a match protocol statement for routes belonging to the same protocol.

Routing Options

In This Section:

Routing Options (Apply, Reset and Reload) - WebUI	173
Equal Cost Path Splitting	173
Kernel Options- Kernel Routes.....	174
Protocol Rank	174
Advanced Routing Options - Wait for Clustering	176
Advanced Routing Options - Auto Restore of Interface Routes	177
Trace Options	178

This chapter describes routing options that apply to all dynamic routing protocols.

Routing Options (Apply, Reset and Reload) - WebUI

In the **Advanced Routing > Routing Options** page of the WebUI, clicking these buttons has this effect:

- **Apply** - Save changes in this page.
- **Reload** - Discard unsaved changes. This is the same as navigating away from the page, discarding changes, and returning to the page.
- **Reset** - Restart the routed routing daemon on the Gaia appliance or computer.

Equal Cost Path Splitting

You can configure the maximum number of equal-cost paths that will be used when there is more than one equal-cost path to a destination. You can specify a value for the maximum number of equal-cost paths that will be used when there is more than one equal-cost path to a destination. Only OSPF routes and Static routes are able to use more than one "next hop"

- Range: 1 to 8
- Default: 8

The "next hop" algorithm that is used for forwarding when there is more than one "next hop" to a destination is Source/destination hash: A hash function is performed on the source and destination IP address of each packet that is forwarded to a multipath destination. This result is used to determine which next hop to use.



Important - Changing this option causes all routes to be reinstalled.

Configuring Equal Cost Path Splitting - WebUI

To configure equal cost path splitting using the WebUI:

1. In the tree view, click **Advanced Routing > Routing Options**.
2. In the **Equal Cost Multipath** section, select the **Maximum Paths**.
3. Click **Apply**.

Configuring Equal Cost Path Splitting - CLI (max-path-splits)

To configure equal cost path splitting using the CLI:

1. Run: `set max-path-splits <1-8>`
For example: `set max-path-splits 2`
2. Run: `save config`

Kernel Options- Kernel Routes

Route Injection Mechanism (RIM) enables a Security Gateway to use dynamic routing protocols to propagate the encryption domain of a VPN peer Security Gateway to the internal network and then initiate back connections. When a VPN tunnel is created, RIM updates the local routing table of the Security Gateway to include the encryption domain of the VPN peer.

In Gaia, the Route Injection Mechanism adds routes directly to the kernel. For the routes to remain in the Kernel, you must configure this option.

For more about configuring RIM, see the *R80.10 Site to Site VPN Administration Guide* http://supportcontent.checkpoint.com/documentation_download?ID=53104.

Configuring Kernel Routes - WebUI

To set kernel routes using the WebUI:

1. In the tree view, click **Advanced Routing > Routing Options**.
2. In the **Kernel Options** area, select the **Kernel Routes** option.
3. Click **Apply**.

Configuring Kernel Routes - CLI (kernel-routes)

To set kernel routes using the CLI:

1. Run: `set kernel-routes on`.
2. Run: `save config`.

Protocol Rank

The *protocol rank* is the value that the routing daemon uses to order routes from different protocols to the same destination. It is an arbitrarily assigned value used to determine the order of routes to the same destination. Each route has only one rank associated with it, even though rank can be set at many places in the configuration. The route derives its rank from the most specific route match among all configurations.

The *active route* is the route installed into the kernel forwarding table by the routing daemon. In the case where the same route is contributed by more than one protocol, the one with the lowest rank becomes the active route.

Rank cannot be used to control the selection of routes within a dynamic interior gateway protocol (IGP); this is accomplished automatically by the protocol and is based on the protocol metric. Instead, rank is used to select routes from the same external gateway protocol (EGP) learned from different peers or autonomous systems.

Some protocols—BGP and aggregates—allow for routes with the same rank. To choose the active route in these cases, a separate tie breaker is used. This tie breaker is called *LocalPref* for BGP and *weight* for aggregates.

Default Ranks

A *default rank* is assigned to each protocol. Rank values range from 0 to 255, with the lowest number indicating the most preferred route.

The default rank values are:

Preference of	Default
Interface routes	0
Static routes	60
OSPF routes	10
RIP routes	100
BGP routes	170
OSPF AS external routes	150
IPv6 OSPF Routes	10
IPv6 OSPF AS external routes	150
Aggregate routes	130

These numbers do not generally need to be changed from their defaults. Use caution when modifying the default route ranks. Rank affects the route selection process, so unexpected consequences may occur throughout the network. Such a change should be planned carefully and take into account both the protocols being used and the location of the router in the network.

Configuring Protocol Rank - WebUI

To set route rank:

1. Open the **Advanced Routing > Routing Options** page of the WebUI.
2. In the **Protocol Rank** section, enter the route rank for each protocol.
3. Click **Apply**.

Configuring Protocol Rank - CLI (protocol-rank)

Rank is used by the routing system when there are routes from different protocols to the same destination. For each route, the route from the protocol with lowest rank number is used.

Syntax

```
set protocol-rank protocol
    bgp rank <0-255>
    bgp rank default
    rip rank <0-255>
```

```

rip rank default
ospf rank <0–255>
ospf rank default
ospfase rank <0–255>
ospfase rank default
ospf3 rank <1–255>
ospf3 rank default
ospf3ase rank <1–255>
ospf3ase rank default

```

Parameter	Description
rank <0–255>	The protocol rank value.
ospf rank default	The default rank value for OSPF is 10.
rip rank default	The default rank value for RIP is 100.
bgp rank default	The default rank value for BGP is 170.
ospfase rank default	The default rank value for OSPF ASE routes is 150.
ospf3 rank default	The default rank value for IPv6 OSPF (OSPFv3) is 10
ospf3ase rank default	The default rank value for IPv6 OSPF (OSPFv3) ASE routes is 150

Advanced Routing Options - Wait for Clustering

In a clustering environment, *Wait for Clustering* has this effect:

WebUI	CLI	The routed routing daemon
Selected	on	For ClusterXL clusters, use this setting. - Does not start the routing protocols if the cluster state is down. - Turns on the routing protocols after the cluster goes up.
Cleared	off	Ignores the state of the cluster. The state of the routing protocols does not depend on the state of the cluster. This is the default.



Important - Changing the setting of this option restarts the routed routing daemon. This option is for ClusterXL clusters only.

Turn on this option with ClusterXL.

Configuring Wait for Clustering - WebUI

To set the Wait for Clustering routing option:

1. In the tree view, click **Advanced Routing > Routing Options**.
2. In the **Router Options** area, select **Wait for Clustering**.
3. Click **Apply**.

Configuring Wait for clustering - CLI (router-options)

To turn on Wait for Clustering:

1. Run: `set router-options wait-for-clustering on`
2. Run: `save config`

To turn off Wait for Clustering:

1. Run: `set router-options wait-for-clustering off`
2. Run: `save config`

To show the state of the Wait for Clustering option:

`show router-options`

Advanced Routing Options - Auto Restore of Interface Routes

An interface route may be automatically deleted in error from the router kernel when it becomes reachable from another interface. The command `show route` in clish shows the route, but the Expert mode command `ip route` does not show the route. A scenario where this can happen is when the same route is learned via OSPF and BGP.

You can avoid losing the interface routes by enabling the Automatic Restoration of Interface Routes option. By default, this option is Off.

Note: If the interface route was deleted and the option was not enabled, bring the interface DOWN and then bring it UP again using clish, the WebUI or, in Expert mode, using `ifconfig`.

To Automatically Restore Routes - WebUI:

1. In the Gaia WebUI, go to the **Advanced Routing > Routing Options** page.
2. In the **Advanced Router Options** area, select **Auto Restore of Iface Routes**.
3. Click **Apply**.

To turn on Auto Restore of Interface Routes - clish:

1. Run: `set router-options Auto-restore-iface-routes on`
2. Run: `save config`

To turn off Auto Restore of Interface Routes - clish:

1. Run: `set router-options Auto-restore-iface-routes off`
2. Run: `save config`

To show the state of the Auto Restore of Interface Routes option - clish:

`show router-options`

Trace Options

The routing system can optionally log information about errors and events. Logging is configured for each protocol or globally. Logging is not generally turned on during normal operations, as it can decrease performance. Log messages are saved in `/var/log/routed.log`

Trace Options - WebUI

To Enable Trace options:

1. In the tree view, click **Advanced Routing > Routing Options**.
2. Click the **Configuration** tab.
3. In the **Trace Options** section, configure:
 - **Maximum Trace File Size** - enter a value between 1 to 2047 MB (the default is 1)
Note - When the trace file reaches the specified size, it is renamed to `file.0`, then `file.1`, `file.2`.
 - **Number of Trace Files** - enter a number between 1 to 4294967295 (the default is 10)
 - **Filter Visible Tables Below** - select trace options tables - **Show All, Global, BGP, Bootp/DHCP Relay, ICMP, IGMP, IP Broadcast Helper, Kernel, MFC, OSPF, Policy Based Routing, PIM, RIP, Router Discovery, VRRP, IPv6 OSPF, IPv6 Router Discovery, IPv6 VRRP**
 - In each trace options table, select options (to select multiple options, hold down Shift and click the options)
 - Click **Add**.
4. Click **Apply** at the top of the page

If you want to enable tracing of a specific routing option for all protocols, in the **Global** options table, select that option. If you want to enable tracing of all routing options for all protocols, in the **Global** options table, select **All**.

For an explanation of each trace option, see the *Trace Options - CLI* (on page 178).

You can see the most recent trace log messages in the `/var/log/routed.log` file.

To monitor an Option:

1. In the **Monitoring** tab of the **Advanced Routing > Routing Options** view, configure the **Number of lines** that you want to show at the end (the "tail") of the log file -minimum is 5, maximum is 100.
2. Click **Get Tail**.

The log messages show.

Trace Options - CLI

Before you configure the routing trace options, you can configure the log file size and the maximum number of log files.

To configure the log file options:

Run this command: `set tracefile {maxnum {<num> | default} | size {<size_MB> | default}}`

Command parameters:

Parameter	Description
maxnum {<num> default}	Maximum number of log files - a number between 1 and 4294967295. The default is 10.
size {<size_MB> default}	Maximum size of the log file, in MB - a number between 1 and 4095. The default is 1MB.

To configure the routing trace options:

Run this command: `set trace {global | BGP | bootp | icmp | igmp | iphelper | kernel | mfc | ospf | ospf3 | pbr | pim | rip | router-discovery | router-discovery6 | vrrp | vrrp6} {<protocol-specific event option>} {on | off}`

These event options are common among all protocol option tables:

Parameter	Description
all	Trace all the routing events.
general	Trace the events related to normal and route options.
normal	Trace all the normal protocol occurrences. Abnormal protocol occurrences are always traced.
policy	Trace the application of protocol- and user-specified policy to imported and exported routes.
route	Trace the routing table changes.
state	Trace the state machine transitions in the protocols.
task	Trace the system interface and processing events.
timer	Trace the timer usage events.
cluster	Trace the cluster-specific events.
{on off}	Turn the options on or off.

For options that are unique to each protocol, see relevant protocol sections below.

Global Trace options

To configure global routing trace options, run the command: `set trace global {adv | parse | all | general | normal | policy | route | state | task | timer | cluster} {on | off}`

To add several options, run the command once for every option.

These event options are specific to global options table:

Parameter	Description
adv	Trace the allocation of and freeing of policy blocks.
parse	Trace the lexical analyzer and parser events.

BGP Trace Options

To configure BGP trace options, run the command: `set trace bgp {keepalive | open | packets | update | all | general | normal | policy | route | state | task | timer | cluster} {on | off}`

To add several options, run the command once for every option.

These event options are specific to BGP options table:

Parameter	Description
keepalive	Trace all the BGP keepalive messages to this peer. These messages are used to verify peer reachability.
open	Trace all the BGP open messages to this peer. These messages are used to establish a peer connection.
packets	Trace all the BGP events.
update	Trace all the BGP update messages to this peer. These messages are used to pass network reachability information.

Bootp/DHCP Relay

To turn Bootp/DHCP Relay trace options on or off:

To configure Bootp/DHCP Relay trace options, run the command: `set trace bootp {packets | request | response | all | general | normal | policy | route | state | task | timer | cluster} {on | off}`

To add several options, run the command once for every option.

These event options are specific to Bootp/DHCP Relay options table:

Trace option	Description
packets	Trace all sent and received bootp packets.
request	Trace all bootp requests.
response	Trace all bootp responses.

ICMP Trace Options

To configure ICMP trace options, run the command: `set trace icmp {error | info | router-discovery | packets | all | general | normal | policy | route | state | task | timer | cluster} {on | off}`

To add several options, run the command once for every option.

These event options are specific to ICMP options table:

Parameter	Description
error	Trace ICMP error packets: - time exceeded - parameter problem - unreachable - source quench
info	Trace ICMP informational packets: - mask request/response - info request/response - echo request/response - time stamp request/response
router-discovery	Trace ICMP router discovery packets.
packets	Trace all ICMP packets.

IGMP Trace Options

To configure IGMP trace options, run the command: `set trace igmp {group | leave | mtrace | query | report | request | packets | all | general | normal | policy | route | state | task | timer | cluster} {on | off}`

To add several options, run the command once for every option.

These event options are specific to IGMP options table:

Parameter	Description
group	Trace multicast group add, delete, refresh and accelerated leave events.
leave	Trace IGMP "leave group" messages.
mtrace	Trace IGMP multicast traceroute events.
query	Trace IGMP membership query packets (both general and group-specific).
report	Trace IGMP membership report packets (both IGMPv1 and IGMPv2).
request	Trace IGMP multicast traceroute request packets.
packets	Trace all IGMP packets.

IP Broadcast Helper Trace Options

To configure IP Broadcast Helper trace options, run the command: `set trace iphelper {packets | all | general | normal | policy | route | state | task | timer | cluster} {on | off}`

To add several options, run the command once for every option.

These event options are specific to IP Broadcast Helper options table:

Parameter	Description
packets	Trace all IP Broadcast Helper packets.

Kernel Trace Options

To configure Kernel trace options, run the command: `set trace kernel {iflist | interface | packets | remnants | request | routes | all | general | normal | policy | route | state | task | timer | cluster} {on | off}`

To add several options, run the command once for every option.

These event options are specific to Kernel options table:

Parameter	Description
iflist	Trace iflist, the interface list scan.
interface	Trace interface status kernel messages.
packets	Trace kernel packets.
remnants	Trace kernel routes at the time when the routing daemon starts.
request	Trace add, delete, or change route requests in the kernel forwarding table.
routes	Trace routes that are exchanged with the kernel, including add, delete, or change messages and add, delete, or change messages received from other processes.

MFC Trace Options

To configure MFC trace options, run the command: `set trace mfc {alerts | cache | interface | mcastdist | packets | resolve | wrongif | all | general | normal | policy | route | state | task | timer | cluster} {on | off}`

To add several options, run the command once for every option.

These event options are specific to MFC options table:

Parameter	Description
alerts	Trace multicast protocol alert callback events.
cache	Trace cache maintenance log details: • addition or deletion of orphan entries (entries with no route to source) • addition or deletion of normal entries • cache state aging and refresh entries

Parameter	Description
interface	Trace log changes requested by external routed modules (IGMP and multicast routing protocols) that affect the forwarding dependencies on an interface: • addition or deletion of a forwarding interface due to routing changes • changing of the parent (reverse path forwarding) interface due to routing changes
mcastdist	Trace generic kernel multicast distribution and PIM register encapsulation and decapsulation entries.
packets	Trace all MFC packets.
resolve	Trace normal kernel and PIM register external resolve requests.
wrongif	Trace kernel multicast incoming physical interface and PIM register violation notifications.

OSPFv2 (IPv4) Trace Options

To configure OSPFv2 trace options, run the command: `set trace ospf {ack | dd | dr | hello | lsa | packets | request | spf | trap | update | all | general | normal | policy | route | state | task | timer | cluster} {on | off}`

To add several options, run the command once for every option.

These event options are specific to OSPF options table:

Parameter	Description
ack	Trace link-state acknowledgment packets.
dd	Trace database description packets.
dr	Trace designated router packets.
hello	Trace hello packets.
lsa	Trace link-state advertisement packets.
packets	Trace all OSPF packets.
request	Trace link-state request packets.
spf	Trace shortest-path-first (SPF) calculation events.
trap	Traces OSPF trap packets.
update	Trace link-state updates packets.

PIM Trace Options

To configure PIM trace options, run the command: `set trace pim {assert | bootstrap | crp | graft | hello | join | mfc | mrt | packets | refresh | register | rp | trap | all | general | normal | policy | route | state | task | timer | cluster} {on | off}`

To add several options, run the command once for every option.

These event options are specific to PIM options table and apply to dense-mode and sparse-mode implementations:

Parameter	Description
<code>assert</code>	Trace PIM assert messages.
<code>hello</code>	Trace PIM router hello messages.
<code>join</code>	Trace PIM join/prune messages.
<code>mfc</code>	Trace calls to or from the multicast forwarding cache
<code>mrt</code>	Trace PIM multicast routing table events.
<code>packets</code>	Trace all PIM packets.
<code>trap</code>	Trace PIM trap messages.
<code>all</code>	Trace all PIM events and packets.

These event options are specific to PIM options table and apply to sparse-mode implementation only:

Parameter	Description
<code>bootstrap</code>	Trace bootstrap messages.
<code>crp</code>	Trace candidate-RP-advertisements.
<code>rp</code>	Trace RP-specific events, including RP set-specific and bootstrap-specific events.
<code>register</code>	Trace register and register-stop packets.

This event option is specific to PIM options table and applies to dense-mode implementation only:

Parameter	Description
<code>graft</code>	Trace graft and graft acknowledgment packets.

RIP Trace Options

To configure RIP trace options, run the command: `set trace rip {packets | all | general | normal | policy | route | state | task | timer | cluster} {on | off}`

To add several options, run the command once for every option.

This event option is specific to RIP options table:

Parameter	Description
packets	Trace all RIP packets.

ICMP Router Discovery Trace Options

To configure ICMP Router Discovery trace options, run the command: `set trace router-discovery {all | general | normal | policy | route | state | task | timer | cluster} {on | off}`

To add several options, run the command once for every option.

There are no event options that are specific to ICMP Router Discovery options table. Use options that are common for all protocols. See Trace Options - CLI (on page 178).

VRRP Trace Options

To configure VRRP trace options, run the command: `set trace vrrp {advertise | all | general | normal | policy | route | state | task | timer | cluster} {on | off}`

To add several options, run the command once for every option.

This event option is specific to VRRP options table:

Parameter	Description
advertise	Trace all VRRP packets.

Router Discovery

In This Section:

How Router Discovery Works	186
Configuring Router Discovery - WebUI	187
Configuring Router Discovery - CLI (rdisc)	188
Monitoring ICMP Router Discovery	189

The ICMP Router Discovery protocol is an IETF standard protocol that allows hosts running an ICMP router discovery client to learn dynamically about the presence of a viable default router on a LAN. It is intended to be used instead of having hosts *wiretap* routing protocols such as RIP. It is used in place of, or in addition to, statically configured default routes in hosts.



Note - Only the *server* portion of the Router Discovery Protocol is supported.

Gaia implements only the ICMP router discovery server portion, which means that a Check Point router can advertise itself as a candidate default router, but it will not adopt a default router using the router discovery protocol.

The ICMP Router Discovery Service provides a mechanism for hosts attached to a multicast or broadcast network to discover the IP addresses of their neighboring routers. This section describes how you can configure a router to advertise its addresses by using ICMP Router Discovery.

How Router Discovery Works

The router discovery server runs on routers and announces their existence to hosts. It does this by periodically multicasting or broadcasting a *router advertisement* to each interface on which it is enabled. These advertisements contain a list of all the router addresses on a given interface and their preference for use as a default router.

Initially, these router advertisements occur every few seconds. They then fall back to every few minutes. In addition, a host can send a *router solicitation*, to which the router responds with a unicast router advertisement. However, if a multicast or broadcast advertisement is due in a moment, the router does not respond with a unicast advertisement.

Each router advertisement contains an *advertisement lifetime* field indicating the length of time that the advertised addresses are valid. This lifetime is configured such that another router advertisement is sent before the lifetime expires. A lifetime of zero (0) indicates that one or more addresses are no longer valid.

On systems that support IP multicasting, the router advertisements are sent by default to the all-hosts multicast address 224.0.0.1. However, you can specify the use of broadcast. All IP addresses configured on the physical interface are included in the router advertisement when:

- Router advertisements are sent to the all-hosts multicast address, *or*
- An interface is configured for the limited-broadcast address 255.255.255.255.

When the router advertisements are sent to a net or subnet broadcast, only the address associated with that net or subnet is included.

Configuring Router Discovery - WebUI

To enable router discovery services:

1. Open the **Advanced Routing > Router Discovery** page of the WebUI.
2. Click **Add**.
The **Add Interface** window opens.
3. Select the **Interface** on which to enable Router Discovery.
4. **Optional:** Enter values for the **Router Discover Configuration parameters**.
 - **Enable Router Discovery**
 - **Min. Advertise Interval**
 - **Max. Advertise Interval**
 - **Advertisement Lifetime**
5. **Optional:** For each IP address on the interface, define the **Router Discover Configuration parameters**:
 - **Advertise**
 - **Eligibility**
 - **Preference**
6. Click **OK**.
7. Click **Save**.

To disable router discovery service on an interface:

1. Open the **Advanced Routing > Router Discovery** page of the WebUI.
2. Select an Interface and click **Edit**.
3. Clear **Enable Router Discovery**.
4. Click **Save**.

Router Discover Configuration parameters

Parameter	Description
Interface	The interface on which Router Discovery occurs.
Enable Router Discovery	Whether ICMP router discovery is running on the interface. After you enable ICMP router discovery, configuration options for the interface appear. • Default: Unselected
Min. Advertise Interval	The minimum time (in seconds) allowed between sending unsolicited broadcast or multicast ICMP Router Advertisements on the interface. • Range: Between 3 seconds and the value in the Max advertise interval. • Default: 0.75 times the value in the Max advertise interval.

Parameter	Description
Max. Advertise Interval	The maximum time (in seconds) allowed between sending unsolicited broadcast or multicast ICMP Router advertisements on the interface. • Range: 4-1800 • Default: 600
Advertisement Lifetime	The lifetime (in seconds) of the advertisements sent from the interface. • Range: Max. Advertise Interval-9000 • Default: 3 x Max. Advertise Interval
Advertise	Whether the address should be advertised in the Router Advertisement packets. This applies to each address on the interface and not to the interface itself. • Default: Selected
Eligibility	You can make an IP address ineligible as a default router address. A router address that is not to be used as a default router has a Preference of 0. • Options: Eligible/ineligible • Default: Eligible.
Preference	The level of preference of the IP address as a default router address, relative to other router addresses on the same subnet. The minimum value corresponds to Ineligible and indicates that the address is not to be used as a default router. • Range: 0 (Ineligible)-2147483648 (2^{31}) • Default is 0

Configuring Router Discovery - CLI (rdisc)

Use the following commands to configure router discovery properties for specific interfaces.

```
set rdisc interface if_name
    <on | off>
    min-adv-interval <3-1800>
    min-adv-interval default
    max-adv-interval <4-1800>
    max-adv-interval default
    adv-lifetime integer
    adv-lifetime default
    advertise ip_address <on | off>
    advertise ip_address preference ineligible
    advertise ip_address preference integer
```

Parameter	Description
<code><on off></code>	Whether to run ICMP router discovery on the interface.
<code>min-adv-interval</code> <code><3-1800></code>	The minimum time (in seconds) allowed between sending unsolicited broadcast or multicast ICMP router advertisements on the interface.
<code>min-adv-interval</code> <code>default</code>	A value of 450 seconds.
<code>max-adv-interval</code> <code><4-1800></code>	The maximum time (in seconds) allowed between sending unsolicited broadcast or multicast ICMP router advertisements on the interface.
<code>max-adv-interval</code> <code>default</code>	A value of 600 seconds.
<code>adv-lifetime integer</code>	The lifetime (in seconds) of the advertisements sent from the interface. An integer value between the configured value for the maximum advertisement interval and 9000.
<code>adv-lifetime default</code>	A value of 1800 or 3 times the maximum advertisement interval.
<code>advertise ip_address</code> <code><on off></code>	Whether to advertise the specified IP address that is associated with the interface should be advertised in router advertisement packets.
<code>advertise ip_address</code> <code>preference ineligible</code>	Do not use the specified IP address as a default router.
<code>advertise ip_address</code> <code>preference integer</code>	The preferability of the specified IP address as a default router address relative to other router addresses on the same subnet.

Monitoring ICMP Router Discovery

Use the following commands to monitor and troubleshoot your ICMP router discovery implementation.

```
show rdisc
  interfaces
  interface if_name
  stats
  summary
```

IPv6 Router Discovery

In This Section:

IPv6 Router Discovery and VRRP	190
IPv6 Discovery - WebUI	190
IPv6 Discovery - clish (ipv6 rdisc6)	193
Monitoring IPv6 Router Discovery	196

ICMPv6 Router Discovery Protocol is an IETF standard protocol. It lets hosts running an ICMPv6 router discovery client:

- Dynamically find other IPv6 nodes.
- Find available routers and Domain Name System (DNS) servers.
- Learn prefixes and configuration parameters related to address configuration.
- Autoconfigure addresses and make relationships between link layer addresses and IPv6 addresses of other nodes.
- Find out if a neighbor is reachable and maintaining paths to other active neighbor nodes.
- Find duplicated addresses.

Gaia acts as an ICMPv6 router discovery server. It can advertise itself as a candidate default router, but it will not make a router its default router using the IPv6 Router Discovery protocol.

Note - IPv6 Router Discovery and ClusterXL cannot be enabled at the same time. We recommend the VRRP clustering solution with IPv6 Router Discovery.

IPv6 Router Discovery and VRRP

To support VRRP for IPv6 interfaces, only the router in a VRRP master state sends router discovery advertisements. The master sends the advertisements with the Virtual IP address as the source address and the Virtual MAC address as the MAC address. Routers in VRRP backup status do not send router discovery advertisements. When VRRP failover occurs, the new master begins to send out router discovery advertisements.

IPv6 Discovery - WebUI

1. Open the **Advanced Routing > IPv6 Router Discovery** page of the WebUI.
2. Click **Add**.
The **Add Interface** window opens.
3. Select an interface on which to enable IPv6 Router Discovery.
4. **Optional:** Configure the **Router Discovery** parameters.
5. **Optional:** Configure how address prefix information is advertised, using the **Advertise Addresses** parameters.

Router Discovery Parameters

Parameter	Description
Interface	The interface on which IPv6 Router Discovery runs.
Min. Advertise interval	The minimum time (in seconds) permitted between sending unsolicited multicast ICMPv6 Router Advertisements on the interface. Unsolicited Router Advertisements are not strictly periodic. The interval between two advertisements is randomized to decrease the probability of synchronization with the advertisements from other routers on the same links. When an unsolicited advertisement is sent, the timer is reset to a random value between the Max. Advertise interval and the Min. Advertise interval. • Range: Between 3 seconds and the value of the Max. Advertise interval. • Default: 1/3 of the Max. Advertise interval.
Max. Advertise interval	The maximum time (in seconds) permitted between sending unsolicited multicast ICMPv6 Router advertisements on the interface. • Range: 4-1800. • Default: 600.
Advertisement Lifetime	The length of time (in seconds) during which a host that receives information from a Check Point router thinks of it as a valid router. This value is refreshed when the host sees a router advertisement. If a host does not see a router advertisement for a longer period than this time, the host thinks of the router as "dead" and stops using it. A value of zero means that the router must not be used as a default router. The value is placed in the <i>Router Lifetime</i> field of the Router Advertisements packet. • Range: zero, or between Max. Advertise interval and 9000. • Default: 3 * Max. Advertise interval.
Reachable timer	The time (in seconds) a node assumes a neighbor is reachable after it received a reachability confirmation. This value is used by the Neighbor Unreachability Detection. The value zero means unspecified (by this router). The reachable time is placed in the <i>Reachable Time</i> field in the router advertisement packet. • Range: 0-3,600,000. • Default: 0.

Parameter	Description
Retransmission Timer	The time (in seconds) between retransmitted Neighbor Solicitation messages if the node does not receive a response. This value is used by address resolution and Neighbor Unreachability Detection. The value zero means unspecified (by this router). This value is placed in the <code>Retrans Timer</code> field in the Router Advertisement packet. • Range: number. • Default: 0.
Hop limit	Nodes use this value in the Hop count field of the IP header for outgoing IP packets. The value zero means unspecified (by this router). The default value is placed in the <code>Cur Hop Limit</code> field in the Router Advertisement packet. • Range: 0-255 • Default: 64
Managed Config	Specify if hosts do stateful autoconfiguration to get addresses. The <code>Managed Config</code> flag is placed in the <code>managed address configuration flag</code> field in the router advertisement packet. • Default: Not enabled.
Other Config Flag	Specify if hosts do stateful autoconfiguration to get more information (without addresses). The <code>Other Config Flag</code> is placed in the <code>Other stateful configuration flag</code> field in the router advertisement packet. • Default: Not enabled.
Send MTU	If enabled, router advertisement packets include MTU options. • Default: Not enabled.

Advertise Addresses Parameters

Parameter	Description
Address	Routers can use IPv6 Router Discovery to communicate address prefixes so that hosts can configure their own IPv6 addresses automatically. Check Point routers automatically configure these prefixes based on their own IPv6 address on the interface which runs IPv6 Router Discovery. The <code>address</code> field is set to be the interface address, and the prefix length sent is the mask length of the router's interface address. Therefore, hosts configure themselves to have the same prefix or mask length as the router. For example, if the router has the interface address <code>2001:db8::1/32</code>, hosts automatically configure themselves to have an address with prefix <code>2001:db8::/32</code>.

Parameter	Description
Enable On-Link	Configure if this address prefix is available on the link. This is necessary because it is possible to have multiple prefix combinations on the same subnet in IPv6. • Default: Enabled.
Enable Autonomous Address Configuration	If enabled, this prefix can be used for autonomous address configuration. • Default: Enabled.
Valid Lifetime	The length of time in seconds (relative to when the packet is sent) that the prefix is valid for on-link determination. The designated value of all 1s (0xffffffff) represents infinity. This value is placed in the Valid Lifetime field in the Prefix Information option. • Range: integer. • Default: 2592000 seconds (30 days)
Preferred Lifetime	The length of time in seconds (from the time the packet is sent) that addresses generated from the prefix through stateless address autoconfiguration stay preferred. That means that the node can use the prefix in existing connections, but it is not valid for new connections. The designated value of all 1s (0xffffffff) represents infinity. This value is placed in the Preferred Lifetime field in the Prefix Information option. • Range: integer. • Default: 604800 seconds (7 days).

IPv6 Discovery - clish (ipv6 rdisc6)

Use these commands to configure IPv6 router discovery properties for a named interface:

```
set ipv6 rdisc6 interface if_name
  <on | off>
  min-adv-interval <3-1800>
  min-adv-interval default
  max-adv-interval <4-1800>
  max-adv-interval default
  hop-limit <0-255>
  hop-limit default
  managed-config <on | off>
  other-config <on | off>
  reachable-time <0-3600000>
  reachable-time default
  retransmit-timer integer
  retransmit-timer default
  router-lifetime integer
  router-lifetime default
  send-mtu <on | off>
```

Use these Advertise Address commands to configure how address prefix information is advertised:

```
set ipv6 rdisc6 interface if_name
  address ip6_address autonomous <on | off>
  address ip6_address on-link <on | off>
```

```

address ip6_address prefix-pref-lifetime integer
address ip6_address prefix-pref-lifetime default
address ip6_address prefix-valid-lifetime integer
address ip6_address prefix-valid-lifetime default

```

Parameter	Description
interface <i>if_name</i>	The interface on which IPv6 Router Discovery is running.
<on off>	Whether to run ICMPv6 router discovery on a specified interface.
min-adv-interval <3-1800>	The minimum time (in seconds) allowed between sending unsolicited multicast ICMPv6 Router Advertisements on the interface. Unsolicited Router Advertisements are not strictly periodic. The interval between two advertisements is randomized to decrease the probability of synchronization with the advertisements from other routers on the same links. When an unsolicited advertisement is sent, the timer is reset to a random value between the Max. Advertise interval and the Min. Advertise interval.
min-adv-interval default	1/3 of the max-adv-interval.
max-adv-interval <4-1800>	The maximum time (in seconds) allowed between sending unsolicited multicast ICMPv6 Router advertisements on the interface.
max-adv-interval default	600 seconds
hop-limit <0-255>	Nodes use this value in the Hop count field of the IP header for outgoing IP packets. The value zero means unspecified (by this router). The default value is placed in the Cur Hop Limit field in the Router Advertisement packet.
hop-limit default	64
managed-config <on off>	Specify if hosts do stateful autoconfiguration to get addresses. The Managed Config flag is placed in the managed address configuration flag field in the router advertisement packet. Default: Off
other-config <on off>	Specify if hosts do stateful autoconfiguration to get more information (without addresses). The Other Config Flag is placed in the Other stateful configuration flag field in the router advertisement packet. Default: Off
reachable-time <0-3600000>	The time (in seconds) a node assumes a neighbor is reachable after having received a reachability confirmation. This value is used by the Neighbor Unreachability Detection. The value zero means unspecified (by this router). The reachable time is placed in the Reachable Time field in the router advertisement packet.

Parameter	Description
reachable-time default	Zero (0) seconds.
retransmit-timer <i>integer</i>	The time (in seconds) between retransmitted Neighbor Solicitation messages if the node does not receive a response. This value is used by address resolution and Neighbor Unreachability Detection. The value zero means unspecified (by this router). This value is placed in the Retrans Timer field in the Router Advertisement packet.
retransmit-timer default	Zero (0) seconds.
router-lifetime <i>integer</i>	The length of time (in seconds) that a host that is receiving information from a Check Point router thinks of it as a valid router. This value is refreshed when the host sees a router advertisement. If a host does not see a router advertisement for more than this time, the host thinks of the router as "dead" and stops using it. A value of zero means that the router is not to be used as a default router. The value is placed in the <i>Router Lifetime</i> field of the Router Advertisements packet. Range: zero, or between Max adv interval and 9000.
router-lifetime default	3 * max-adv-interval
send-mtu <on off>	If enabled, router advertisement packets include MTU options. Default: Off

Advertise Addresses Parameters

Parameter	Description
address <i>ip6_address</i>	Routers can use IPv6 Router Discovery to communicate address prefixes for hosts to configure their own IPv6 addresses automatically. Check Point routers automatically configure these prefixes based on their own IPv6 address on the interface running IPv6 Router Discovery. The address field is set to be the interface address, and the prefix length sent is the mask length of the router's interface address. Therefore, hosts configure themselves to be the same prefix / mask length as the router. For example, if the router has the interface address 2001:db8::1/32, hosts will automatically configure themselves to have an address with prefix 2001:db8::/32.
autonomous <on off>	If enabled, this prefix can be used for autonomous address configuration. Default: On

Parameter	Description
on-link <on off>	Configure if this address prefix is available on the link. This is necessary because it is possible to have multiple prefix combinations on the same subnet in IPv6. Default: On
prefix-valid-lifetime <i>integer</i>	The length of time in seconds (relative to the time the packet is sent) that the prefix is valid for on-link determination. The designated value of all 1s (0xffffffff) represents infinity. This value is placed in the Valid Lifetime field in the Prefix Information option. Range: Integer.
prefix-valid-lifetime default	2592000 seconds (30 days)
prefix-pref-lifetime <i>integer</i>	The length of time in seconds (from the time the packet is sent) that addresses generated from the prefix through stateless address autoconfiguration stay preferred. That means that the node can use the prefix in existing connections, but it is not valid for new connections. The designated value of all 1s (0xffffffff) represents infinity. This value is placed in the Preferred Lifetime field in the Prefix Information option. Range: Integer
prefix-pref-lifetime default	604800 seconds (7 days).

Monitoring IPv6 Router Discovery

You can monitor IPv6 Router Discovery in the WebUI and in the clish CLI.

Monitoring IPv6 Router Discovery in the WebUI

1. In the WebUI, go to the **Advanced Routing > IPv6 Router Discovery** page.
2. Click the **Monitoring** tab.

The page shows:

- A **Summary**.
- Settings for the **Interfaces**.
- **Statistics** per interface.

Show IPv6 Router Discovery Commands

Use these commands to monitor IPv6 Router Discovery:

```
show ipv6 rdisc6
  summary
  interface if_name
  interfaces
  stats
```

Policy Based Routing

In This Section:

Configuring Policy Based Routing - WebUI.....	198
Configuring Policy Based Routing - CLI	200
Monitoring Policy Based Routing.....	202

In addition to dynamic and static routing, you can use *Policy Based Routing* (PBR) to control traffic. PBR Policy Rules have priority over static and dynamic routes in the routing table. When a packet arrives at a Gaia Security Gateway, the gateway goes through the PBR Rules in the order of their set priority, and looks for a match. If the match exists, the gateway forwards the packet according to the rule. If there is no match in the PBR Policy, the gateway forwards the packet according to static or dynamic routes in the routing table.

To configure Policy Based Routing:

1. Create *Action Tables* - Sets of static routes to destination networks.
2. Configure *Policy Rules* - For each set of matching criteria, define the priority and the routing action.

You can configure Policy Based Routing in Check Point Gaia WebUI or in CLI.

Configuring Policy Based Routing - WebUI

To add static routes in an Action Table:

1. In the Gaia WebUI, go to **Advanced Routing > Policy Based Routing**.
2. In the **Action Tables** section, click **Add**.
The **Add Policy Table with Static Route** window opens.
3. Define the route parameters -
 - **Table Name** - Name of the Policy Table
Note - **Table ID** is assigned by the system.
 - **Default Route** (optional) - Make this the default route
Note - If selected, the **Destination** address and **Subnet mask** fields do not show.
 - **Destination** - Destination IPv4 address
 - **Subnet mask** - Destination IPv4 subnet mask
 - **Next Hop Type** -
 - **Normal** - Accept and forward packets
 - **Reject** - Drop packets and send *Unreachable* message to the sender
 - **Black Hole** - Drop packets without a notification to the sender
4. Configure the next hop (if **Normal** is selected for the **Next Hop Type**) - click **Add Gateway** and select one of these:
 - **IP Address** - Enter the **Gateway Address** and select a **Priority**
 - **Network Interfaces** - Select a **Gateway Interface** and a **Priority**

Note - You can configure several next hops.

5. Click **Save**.

To delete an Action Table:

1. In the **Action Tables** section of the **Policy Based Routing** page, select a static route table.
2. Click **Delete**.

To add a Policy Rule:

1. In the **Policy Rules** section of the **Policy Based Routing** page, click **Add**.
2. The **Add Policy Rule** window opens.
3. Set the **Priority** of the rule - an integer between 1 and 32765.
4. Set the routing **Action** for the traffic that matches the specified criteria -
 - **Prohibit** - Drop the packet and send a *Prohibit* message to the sender
 - **Unreachable** - Drop the packet and send an *Unreachable* message to the sender
 - **Table** - Forward the packet according to the routes in the selected Action Table
5. Configure one or more of the **Match** criteria -
 - **Interface** - Interface on which the traffic arrived at the gateway
 - **Source** - IPv4 address of the source
 - **Subnet mask** - Subnet mask of the source address
 - **Destination** - IPv4 address of the destination
 - **Subnet mask** - Subnet mask of the destination address
 - **Service Port** - Service port - enter a number between 1 and 65535, or select a predefined port from the drop-down menu
 - **Protocol** - Protocol - enter a number between 1 and 255, or select a predefined protocol from the drop-down menu
6. Click **Save**.

To Delete a Policy Rule:

1. In the **Policy Rules** section of the **Policy Based Routing** page, select a rule.
2. Click **Delete**.

Action Table Parameters

Parameter	Description
Table Name	The name of the table.
Table ID	A numerical ID for the table. Assigned by the system.
Default route	The default static route in the system routing table.
Destination	The destination of the route.
Subnet mask	Subnet mask for the destination of the route.

Parameter	Description
Table Name	The name of the table.
Next Hop Type	Choose one of: • Normal: Accept and forward packets. • Reject: Drop packets and send <i>unreachable</i> messages. • Black Hole: Drop packets but don't send <i>unreachable</i> messages.
Gateway IP address	Next hop gateway IPv4 address.
Gateway Interface	Security Gateway interface that leads to the next hop gateway.
Gateway Priority	The preference of the particular route. • Range: 1-8

Policy Rule Parameters

Parameter	Description
Priority	You can define many Policy Rules. Traffic is matched to all the rules, one rule at a time, according to the priority that is configured for the rule.
Action	The action to take if the traffic is matched
Prohibit	Send a Prohibit message to the sending host.
Unreachable	Send an Unreachable message to the sending host.
Table	Do the actions defined in an Action Table.
Match	
Interface	Match by: Interface through which the packets enter the Security Gateway from the source host.
Source, subnet mask	Match by: Source IPv4 address and subnet mask.
Destination, Subnet mask	Match by: Destination IPv4 address and subnet mask.

Configuring Policy Based Routing - CLI

To create an Action Table:

```
Run this command: set pbr table <table_name> static-route {default |
<destination_ip/mask>} nexthop {{gateway {address <nexthop_ip> | logical
<interface>}} {priority <route_priority_value> | on | off}} | reject | blackhole}
```

Parameter	Description
<code>table <table_name></code>	Name of the PBR Policy Table.
<code>static-route {default <destination_ip/mask>}</code>	Route to - <code>default</code> - Default route <code><destination_ip/mask></code> - Destination IPv4 address and mask.
<code>gateway {address <nexthop_ip> logical <interface>}</code>	Next hop - <code>address <nexthop_ip></code> - IPv4 address of the next hop gateway <code>logical <interface></code> - Exit interface that leads to the next hop gateway
<code>priority <route_priority_value> on off</code>	Control the route - <code>priority <route_priority_value></code> - Set priority of the route - a value from 1 to 8 <code>on</code> - Turn on the route <code>off</code> - Turn off the route
<code>reject</code>	Drop packets and send <i>Unreachable</i> messages to the sender
<code>blackhole</code>	Drop packets and do not send any notifications to the sender

Note - You can add multiple routes to the same table. To do that, run `set pbr table` command with the same *table_name*.

Example:

Create an Action Table named *PBRtable1*, with a route to the network 192.0.2.0/24 out of the interface *Ethernet 0* and a route to the network 192.0.3.0/24 through the next-hop gateway with the IP address 192.168.1.1.

```
set pbr table PBRtable1 static-route 192.0.2.0/24 nexthop gateway logical eth0 on
```

```
set pbr table PBRtable1 static-route 192.0.3.0/24 nexthop gateway address 192.168.1.1 on
```

To configure a Policy Rule:

Run this command: `set pbr rule priority <priority_value> {action {prohibit | table <PBR_Table> | unreachable} | match {from <source_IP/mask> | interface <interface> | port {<port_num> | off} | protocol {<num> | tcp | udp | icmp | off} | to <destination_IP/mask>} | off}`

Parameter	Description
<code>priority <priority_value></code>	Unique integer value between 1 and 5000. The gateway checks all Policy Rules, one at a time, in order of priority. The highest priority is 1.

Parameter	Description
<pre>action {prohibit table <PBR_Table> unreachable}</pre>	If the packet matches the specified parameters, select a routing action: prohibit - Drop the packet and send a <i>Prohibit</i> message to the sender table <PBR_Table> - Forward the packet according to the specified Action Table - <i><PBR_Table></i> unreachable - Drop the packet and send an <i>Unreachable</i> message to the sender
<pre>match {from <source_IP/mask> interface <interface> port {<port_num> off} protocol {<num> tcp udp icmp off} to <destination_IP/mask>}</pre>	Configure the traffic matching criteria: from <source_IP/mask> - IPv4 address and the subnet mask of the source interface <interface> - Incoming interface port <service_port> - Service port number, and integer between 1 and 65535 protocol {<num> tcp udp icmp off} - Protocol, an integer between 1 and 255, or one of predefined protocols - TCP, UDP, and ICMP to <destination_IP/mask> - Destination IPv4 address and the subnet mask
off	Delete the Policy Rule.

Example:

Create a Policy Rule that forwards all packets with the destination address 192.0.2.1/32 that arrive on the interface *Ethernet 2* according to the PBR Table **PBRtable1**, and assign to it the priority of 100.

```
set pbr rule priority 100 match to 192.0.2.1/32 interface eth2
set pbr rule priority 100 action table PBRtable1
```

Monitoring Policy Based Routing

To monitor Policy Based Routing - WebUI

1. In the Gaia WebUI, go to **Advanced Routing > Policy Based Routing**.
2. Click the **Monitoring** tab.

To monitor Policy Based Routing - CLI

Run these commands:

```
show pbr tables
show pbr rules
show pbr summary
```

PIM

In This Section:

Dense Mode	203
Sparse Mode	203
Source-Specific Multicast (SSM) Mode	203
PIM Dense Mode State Refresh	204
Configuring PIM - WebUI	204
Configuring PIM - CLI (pim)	214
Static Multicast Routes	222
Monitoring and Troubleshooting PIM	224

Protocol-Independent Multicast (PIM) can forward multicast packets with a unicast protocol. PIM efficiently routes multicast traffic for groups that span wide area (and inter-domain) networks. It works with all existing unicast routing protocols. PIM supports three modes: **dense**, **sparse** and Source-Specific Multicast (**SSM**). You can enable only one mode of PIM at a time.

Note - Due to an OS limitation, PIM supports only 31 interfaces. If more are configured, PIM runs only on the first 31 interfaces.

Dense Mode

Dense mode is most useful when:

- Senders and receivers are in close proximity to one another.
- There are few senders and many receivers.
- The volume of multicast traffic is high.
- The stream of multicast traffic is constant.

Sparse Mode

Sparse mode is most useful when:

- There are few receivers in a group.
- Senders and receivers are separated by WAN links.
- The type of traffic is intermittent.

Source-Specific Multicast (SSM) Mode

Source-Specific Multicast (SSM) is most useful when:

- Most multicast traffic is from well-known sources.
- It is desirable to avoid the overhead of shared tree and rendezvous point processing associated with sparse mode.

SSM is a version of PIM sparse-mode. It is used in conjunction with IGMP version 3 to request or block multicast traffic from specific sources. For example, when a host requests traffic for a multicast group from a specific source, SSM sends PIM join/prune messages towards the source.

The multicast group range 232/8 is reserved for SSM. When SSM is enabled, sparse-mode accepts only IGMPv3 reports for groups that fall within this range. Sparse-mode ignores IGMP v1 and v2 reports in this range. In addition, only shortest-path-tree (SPT) join/prune messages for these groups are accepted from neighboring routers. All other multicast groups are processed as in native sparse mode.

SSM does not need a rendezvous-point (RP). The presence of an RP for any of the SSM groups does not have any influence on the processing of join/prune messages.

PIM Dense Mode State Refresh

The State Refresh option can be used in conjunction with dense mode to eliminate the periodic flood-and-prune of multicast data with no active receivers. All PIM routers must have State Refresh enabled to take advantage of this feature.

PIM dense mode builds multicast distribution trees that operate on a flood and prune principle. Multicast packets from a source are flooded throughout a PIM dense mode network. PIM routers that receive multicast packets and have no directly connected multicast group members or PIM neighbors send a prune message back up the source-based distribution tree toward the source of the packets. As a result, subsequent multicast packets are not flooded to pruned branches of the distribution tree. However, the pruned state in PIM dense mode times out approximately every three minutes and the entire PIM dense mode network is reflooded with multicast packets and prune messages. This reflooding of unwanted traffic throughout the PIM dense mode network consumes network bandwidth unnecessarily.

Use the PIM Dense Mode State Refresh feature to keep the pruned state in PIM dense mode from timing out by periodically forwarding a control message down the distribution tree. The control message refreshes the prune state on the outgoing interfaces of each router in the tree. This saves network bandwidth by greatly reducing the reflooding of unwanted multicast traffic to pruned branches of the PIM dense mode network.



Note - You must enable state refresh on all the PIM routers in the distribution tree to take advantage of this feature.

Configuring PIM - WebUI

To configure PIM using the WebUI:

1. Open the **Advanced Routing > PIM** page of the WebUI.
2. Configure the PIM Global Settings ("**PIM Global Settings**" on page 205). In the **PIM Global Settings** section, select the **PIM Protocol**. One of:
 - **Sparse Mode (SM)**
 - **Dense Mode (DM)**. Enable **State Refresh**, if appropriate.
 - **Source Specific Multicast (SSM)**
3. In the **PIM Interfaces** section, click **Add**.
The **Add Interface** window opens.
4. In the **Add Interface** ("**PIM Interfaces**" on page 205) window

- a) Select the **Interface** on which you want to run PIM.
- b) Optional: For each interface that is running PIM, enter the **Local Address**.
- c) Optional: To configure this interface to use the VRRP virtual IP address, click **Use Virtual address**.
- d) Optional: Enter a new **DR Priority** (Designated Router priority).

5. Click **Save**.

PIM Global Settings

Parameter	Description
PIM Protocol	The PIM mode to use. One of: • Sparse Mode (SM) • Dense Mode (DM) • Source-Specific Multicast (SSM)
State Refresh	In Dense Mode, use state refresh messages to delay timing out prune state of multicast traffic that has no active receivers. This helps suppress the flood-and-prune cycle inherent to dense mode.

PIM Interfaces

Parameter	Description
Interface	Select the interface on which to enable PIM
Local Address	Use the specified local address for all advertisements sent on the interface. This is useful on interfaces with multiple IP addresses (aliases). The local address must match one of the addresses configured on the interface. If a local address is specified with the virtual address option enabled, the local address must match a virtual address on the interface. Note - Each router must have at least one interface address with a subnet prefix shared by all neighboring PIM routers. If any neighboring routers select advertisement addresses that are not on a shared subnet, all messages from those neighbors are rejected. This is also true when the virtual address option is enabled. • Range: Dotted-quad ([0-255].[0-255].[0-255].[0-255]). • Default: Selects one of the IP addresses configured on the interface. If the virtual address option is enabled, PIM uses the virtual address configured on the interface after the router changes to master state with respect to VRRP.

Parameter	Description
Use Virtual Address	Use the VRRP virtual IP address on this interface. If enabled, PIM runs on this interface only after the router becomes a VRRP master after a failover. When you enable virtual IP support for VRRP on a PIM interface, it creates the neighbor relationship with the virtual IP, if the router is a VRRP master. The master in the VRRP pair sends hello messages that include the virtual IP as the source address and processes PIM control messages from routers that neighbor the VRRP pair. Note - You cannot configure a local address or a virtual address when ClusterXL is enabled. • Range: Enabled, cleared • Default: Cleared
DR Priority	The dr-priority advertised in the PIM hello messages sent on the interface. This is used for DR selection on a LAN. The router with the highest priority is selected as the designated router. To break a tie, the DR is selected on the basis of the highest IP address. If even one router does not advertise a dr-priority, the DR election is based on the IP address. • Range: 0-4294967295 ($2^{32} - 1$). • Default: 1. Note - To make sure that a PIM neighbor supports DR Priority, run this command: <code>show pim neighbor <ip_address></code> For neighbors that advertise a DR selection priority value, this message shows in the summary: DRPriorityCapable Yes.

Disabling PIM

You can disable PIM on one or more interfaces configured on the WebUI platform.

1. Open the **Advanced Routing > PIM** page of the WebUI.
2. In the **PIM Interfaces** section, select the interface on which to disable PIM and click **Delete**.
To disable PIM entirely, delete all PIM interfaces.

Configuring PIM-DM Advanced Options (Optional)

To configure PIM Dense Mode Advanced Options:

1. Open the **Advanced Routing > PIM** page of the WebUI.
2. Select the **PIM Protocol: Dense Mode (DM)**.
3. Click **Edit Settings**. The **Advanced Options** window opens.
4. In the **General Timers** section ("**PIM Advanced Options- General Timers**" on page 208), enter values for the:
 - **Hello interval** (in seconds).

- **Data Interval** (in seconds).
 - **Assert Interval** (in seconds).
 - **Assert-rate Limit**.
 - **Join Prune Interval** (in seconds).
 - **Join Prune Delay Interval** (in seconds).
 - **Join Prune Suppress Interval** (in seconds).
5. In the **Assert Ranks** section ("[PIM Advanced Options- Assert Ranks](#)" on page 209), enter values for the routing protocol(s) you use.
 6. Click **Save**.

Configuring PIM-SM Advanced Options (Optional)

To configure PIM Simple Mode Advanced options:

1. Open the **Advanced Routing > PIM** page of the WebUI.
2. Select the **PIM Protocol**. One of:
 - **Sparse Mode (SM)**
 - **Source Specific Multicast (SSM)**
3. Click **Edit Settings**.
The **Advanced Options** window opens.
4. In the **Sparse Mode Timers** ("[PIM Advanced Options- Sparse Mode Timers](#)" on page 210) section, enter a value for the
 - a) **Register Suppression Interval** (in seconds).
 - b) **CRP Advertise Interval** (in seconds)
5. In the **Shortest Path First Threshold** section, click **Add**.
The **Shortest Path First Threshold: Add Multicast Group** window opens.
6. In the **SPT: Add Multicast Group** window ("[PIM Advanced Options- Shortest First Path Threshold - Add Multicast Group](#)" on page 210), enter a value for the:
 - a) **Multicast Group** to which the SPT threshold applies.
 - b) **Subnet mask** for the group multicast address.
 - c) **Threshold to Switch** (in kilobits per second).
 - d) Click **OK**.
7. In the **General Timers** section ("[PIM Advanced Options- General Timers](#)" on page 208), enter values for the:
 - **Hello interval** (in seconds).
 - **Data Interval** (in seconds).
 - **Assert Interval** (in seconds).
 - **Assert-rate Limit**.
 - **Join Prune Interval** (in seconds).
 - **Join Prune Delay Interval** (in seconds).
8. In the **Assert Ranks** section ("[PIM Advanced Options- Assert Ranks](#)" on page 209), enter values for the routing protocol(s) you use.
9. Click **Save**.

PIM Advanced Options- General Timers

Parameter	Description
Hello interval	Interval between which PIM hellos are sent on a multicast-capable interface. Hello messages are addressed to the All-PIM-Routers multicast group (224.0.0.13) so that PIM routers may discover neighbors on a multi-access network. • Range: 1-21845 seconds • Default: 30
Data Interval	The life-time of a new PIM forwarding entry. Subsequently, the life of the entry is extended in different ways based on the location of this router in the network. For example, in some cases the receipt of PIM control messages (e.g. periodic join/prune messages) extends the life of the entry and in others the presence of local senders of multicast traffic prevents the deletion of the entry. • Range: 11-3600 seconds • Default: 210
Assert Interval	If an assert battle on an upstream interface results in the selection of a PIM neighbor other than the unicast reverse-path-forwarding (RPF) neighbor towards the source of the data traffic (for which the assert battle was generated) as the designated forwarder on that interface, the winner is used as the upstream neighbor for all subsequent join/prune messages. This change is timed-out after expiry of the assert interval (measured in seconds). • Range: 1-3600 seconds • Default: 180
Assert-rate Limit.	The number of asserts to send per second. The router generates the Asserts when data from a source is detected on an interface other than the incoming interface (based on the unicast routing table) towards the source. These messages are rate-limited and the router must not originate more than a fixed number of assert messages per second. If the limit is set to 0, rate-limiting is disabled. • Range: 0, 10-10000 • Default: 10
Join Prune Interval	Interval between sending Join/Prune messages. • Range: 1-3600 seconds • Default: 60
Join Prune Delay Interval	The maximum interval from the time when the unicast Reverse Path Forwarding (RPF) neighbor (towards a source or the RP) changes, and a triggered Join/Prune message is sent. • Range: 1-3600 seconds • Default: 5

Parameter	Description
Join Prune Suppress Interval	Mean interval from the receipt of a Join/Prune with a higher Holdtime (with ties broken by higher network layer address) and allowing duplicate Join/Prunes to be sent again. Set this interval to 1.25 times the join/prune interval. • Range: 2-3600 seconds • Default: 75

PIM Advanced Options- Assert Ranks

Parameter	Description
Direct OSPF Kernel Static RIP BGP	Compares the cost of protocols to find which router will forward multicast data packets on a multi-access LAN. These values are used in assert messages sent out on a LAN when a router detects data packets on an interface other than the incoming interface towards the source of the data. These values must be the same for all routers on a multi-access LAN that run the same protocol. Therefore, the default values were specially configured to match those of other implementations. • Range: 0-255 • Defaults: • Direct: 0 • OSPF: 10 • Kernel: 40 • Static: 60 • RIP: 100 • BGP: 170

PIM Advanced Options- State Refresh Parameters (DM)

Parameter	Description
State Refresh Interval	For Dense Mode, the interval at which state refresh messages are sent for multicast traffic originated by directly-connected sources. • Range: 1-255 • Default: 60
State Refresh TTL	For Dense Mode, the time-to-live (TTL) placed in the state refresh messages originated for multicast traffic from directly-connected sources. You can use this value to limit the forwarding of state refresh messages in the network. In the absence of user configuration, it is derived from the multicast data. • Range: 1-255 • Default: None

PIM Advanced Options- Sparse Mode Timers

Parameter	Description
Register Suppression Interval	The mean interval between receipt of a register-stop and the time when registers can be sent again. A lower value means more frequent register bursts at the rendezvous point. A higher value means a longer join latency for new receivers. • Range: 60-3600 seconds • Default: 60 seconds
CRP Advertise Interval	The interval between which candidate-rendezvous point routers send candidate-rendezvous point advertisements to the elected bootstrap router. • Range: 1-3600 seconds • Default: 60 seconds

PIM Advanced Options- Shortest First Path Threshold - Add Multicast Group

Parameter	Description
Multicast Group	The multicast group address to apply to the shortest path tree (spt) threshold. • Range: Dotted-quad ([224-239].[0-255].[0-255].[0-255]) • Default: None
Subnet mask	Mask length. • Range: 1-32 • Default: None
Use Infinity	Specifies that there is no spt switch.
Use Integer	The data rate threshold in Kbits/sec to start the spt switchover. When the data rate for a sparse-mode group is higher than the shortest-path-tree threshold at the last-hop router, an (S,G) entry is created and a join/prune message is sent toward the source. Setting this option builds a shortest-path tree from the source S to the last-hop router. • Range: 0-1000000 Kbits/s, or infinity (for no switchover) • Default: None

Configuring PIM-SM Bootstrap and Rendezvous Point Settings

To configure this router as a Bootstrap router, Candidate Rendezvous Point and Static Rendezvous Point:

1. Open the **Advanced Routing > PIM** page of the WebUI.
2. Select the **PIM Protocol**. One of:
 - **Sparse Mode (SM)**

- **Source Specific Multicast (SSM)**
3. In the **Bootstrap and Rendezvous Point Settings** section, click **Edit Settings**.
The **Bootstrap and Rendezvous Point Settings** window opens.
 4. To enable the router as a Bootstrap Router ("PIM Bootstrap Settings" on page 211):
 - a) Select **Enable Bootstrap Router**.
 - b) Optional: Enter the **Local Address** of the bootstrap router.
 - c) Optional: Enter the **Local Preference**.
 5. To enable the router as a Candidate Rendezvous Point ("PIM Candidate Rendezvous Point" on page 212):
 - a) Select **Enable Candidate RP**.
 - b) Optional: Enter the **Local Address** of the Candidate Rendezvous Point router.
 - c) Optional: Enter the **Local Preference**.
 - d) Optional: Click **Add** to Configure a **Multicast Group** and **Subnet mask** for which this router is designated as the candidate rendezvous point.
 6. To enable a Static Rendezvous Point ("PIM Static Rendezvous Point" on page 213):
 - a) Select **Enable Static RP**.
 - b) Optional: Click **Add** to enter the **Static Rendezvous Point** IP address.
 - c) Optional: Click **Add** to configure a **Multicast Group** and **Subnet mask** for which this router is designated as the static rendezvous point.
 7. Click **Save**.

PIM Bootstrap Settings

Parameter	Description
Enable Bootstrap Router	If enabled, this router is a candidate bootstrap router (C-BSR). All candidate RPs (C-RPs) send C-RP-Advertisements to the selected bootstrap router (BSR). The BSR then disseminates this information in bootstrap messages across the PIM domain. To prevent a single point of failure, configure more than one router in a domain as a candidate BSR. • Default: • Cleared

Parameter	Description
Local Address	Address used for the C-BSR state machine and the bootstrap messages. If PIM clustering is enabled, then this address must be configured and must match that of one of the PIM interfaces. If PIM clustering is not enabled, this address can be that of the PIM interfaces or an address configured on the loopback interface. If an address from the loopback interface is used, do not select an address in the 127/8 address range. • Range: Address of PIM interface or a non 127.0.0.0/8 loopback address. • Default: The IP address of one of the interfaces on which PIM is enabled. The default does not apply if PIM clustering is enabled.
Local Preference	The priority advertised in C-BSR messages. The candidate bootstrap router with the highest priority value is selected as the bootstrap router for the domain. The C-RP with the lowest priority has the highest preference. The highest priority value is 0. • Range: 0-255 • Default: 0

PIM Candidate Rendezvous Point

Parameter	Description
Enable Candidate RP	Specifies that the platform is a candidate rendezvous point router.
Local Address	Address used for the C-RP state machine and in the C-RP-Advertisements sent to the elected bootstrap router. If PIM clustering is enabled, then this address must be configured and must match that of one of the PIM interfaces. If clustering is not enabled, this address can either be that of one of the PIM interfaces or an address configured on the loopback interface. If an address from the loopback interface is used, take care not select an address in the 127/8 address range. • Range: Address of PIM interface or a non 127.0.0.0/8 loopback address. • Default: Selects the IP address of one of the interfaces on which PIM is enabled. The default does not apply if PIM clustering is enabled.
Local Preference	The priority of this C-RP. All PIM routers select the same RP for a multicast group address from the list of C-RPs received in the bootstrap messages from the elected BSR. The lower the Local Preference of the C-RP, the higher the priority. • Range: 0-255 • Default: 0

Parameter	Description
Multicast Group	The multicast group(s) for which this rendezvous point is responsible. Enter the group multicast IP address and mask length. Multicast group address • Range: Dotted-quad ([224-239].[0-255].[0-255].[0-255]) • Default: 224.0.0.0/4
Subnet mask	Mask length: • Range: 1-32 • Default: None

PIM Static Rendezvous Point

Parameter	Description
Enable Static RP	Enables or disables the static rendezvous point.
Static Rendezvous Point	A static rendezvous point. If an associated multicast group and prefix is not configured, the static-rp is considered to be responsible for all multicast groups (224.0.0.0/4). This needs to be consistent with the RP information at other routers in a multicast domain irrespective of the RP-dissemination mechanism (bootstrap or autoRP) used. Note - The static RP overrides the RP information received from other RP-dissemination mechanisms, such as bootstrap routers. • Range: Any IP address • Default: None
Multicast Group	The multicast group(s) for which this rendezvous point is responsible. Enter the group multicast IP address and mask length. Multicast group address • Range: Dotted-quad ([224-239].[0-255].[0-255].[0-255]) • Default: 224.0.0.0/4
Subnet mask	Mask length: • Range: 1-32 • Default: None

Configuring PIM - CLI (pim)

Use the commands in this section to configure PIM via the CLI.

PIM Global Settings

```
set pim mode
    dense
    sparse
    ssm
```

```
set pim state-refresh <on | off>
```

Parameter	Description
<dense sparse ssm>	The PIM mode to use. One of: - Sparse Mode (SM) - Dense Mode (DM) - Source-Specific Multicast (SSM)
state-refresh <on off>	In Dense Mode, use state refresh messages to delay timing out prune state of multicast traffic that has no active receivers. This helps suppress the flood-and-prune cycle inherent to dense mode.

PIM Interfaces

After you set PIM to run dense mode, sparse mode or SSM, use the following commands to configure PIM for specific interfaces.

```
set pim interface if_name
    <on | off>
    virtual-address <on | off>
    local-address ip_address
    dr-priority <0-4294967295>
    dr-priority default
```

Parameter	Description
interface if_name <on off>	Whether to enable or disable PIM on a specified interface.
virtual-address <on off>	Use the VRRP virtual IP address on this interface. If enabled, PIM runs on this interface only after the router transitions to become a VRRP master after a failover. When you enable virtual IP support for VRRP on a PIM interface, it establishes the neighbor relationship by using the virtual IP if the router is a VRRP master. The master in the VRRP pair sends hello messages that include the virtual IP as the source address and processes PIM control messages from routers that neighbor the VRRP pair. Note - You cannot configure a local address or a virtual address when ClusterXL is enabled. Default: Off

Parameter	Description
local-address ip_address	Use the specified local address for all advertisements sent on the interface. This is useful on interfaces with multiple IP addresses (aliases). The local address must match one of the addresses configured on the interface. If a local address is specified while the virtual address option enabled, the local address must match a virtual address on the interface. Note - Each router must have at least one interface address with a subnet prefix shared by all neighboring PIM routers. If any neighboring routers choose advertisement addresses that do not appear to be on a shared subnet all messages from those neighbors will be rejected. This holds true even when the virtual address option is enabled. • Range: Dotted-quad ([0-255].[0-255].[0-255].[0-255]). • Default: Selects one of the IP addresses configured on the interface. If the virtual address option is enabled, PIM will use the virtual address configured on the interface after the router transitions to master state with respect to VRRP.
dr-priority <0-4294967295>	The dr-priority advertised in the PIM hello messages sent on the interface. This is used for DR election on a LAN. The router with the highest priority is elected the designated router. To break a tie, the DR is selected on the basis of the highest IP address. If even one router does not advertise a dr-priority, the DR election is based on the IP address. • Range: 0-4294967295 ($2^{32} - 1$). • Default: 1. • Note - To verify whether a PIM neighbor supports DR Priority, use the following command: <code>show pim neighbor <ip_address></code> For neighbors that advertise a DR election priority value, the following message appears in the summary: DRPriorityCapable Yes.
dr-priority default	A value of 1.

PIM Sparse Mode

Use the following commands to configure parameters for sparse mode PIM only.

```
set pim
  bootstrap-candidate <on | off>
  bootstrap-candidate local-address ip_address
  bootstrap-candidate priority <0-255>
  bootstrap-candidate priority default
  candidate-rp <on | off>
  candidate-rp local-address ip_address
  candidate-rp priority <0-255>
  candidate-rp priority default
  candidate-rp multicast group mcast_ip_prefix <on | off>
  static-rp off
  static-rp rp-address ip_address < on | off>
```

```

static-rp rp-address ip_address multicast-group <mcast_address>/<mask_length>
<on | off>
register-suppress-interval <60-3600>
register-suppress-interval default
candidate-rp advertise-interval <1-3600>
candidate rp-advertise-interval default
spt-threshold multicast <mcast_address>/<mask_length> threshold <0-1000000>
<on | off>
spt-threshold multicast <mcast_address>/<mask_length> threshold infinity
<on | off>

```

Bootstrap and Rendezvous Point Settings

Parameter	Description
bootstrap-candidate <on off>	If enabled, this router is a candidate bootstrap router (C-BSR). All candidate RPs (C-RPs) send C-RP-Advertisements to the elected bootstrap router (BSR). The BSR then disseminates this information in bootstrap messages across the PIM domain. To avoid a single point of failure, configure more than one router in a domain as a candidate BSR. • Default: • off
bootstrap-candidate local-address ip_address	Address used for the C-BSR state machine and the bootstrap messages. If PIM clustering is enabled, then this address must be configured and must match that of one of the PIM interfaces. If PIM clustering is not enabled, this address can either be that of the PIM interfaces or an address configured on the loopback interface. If an address from the loopback interface is used, take care not select an address in the 127/8 address range. • Range: Address of PIM interface or a non 127.0.0.0/8 loopback address. Default: The IP address of one of the interfaces on which PIM is enabled. The default does not apply if PIM clustering is enabled.
bootstrap-candidate priority <0-255>	The priority advertised in C-BSR messages. The candidate bootstrap router with the highest priority value is elected as the bootstrap router for the domain. The C-RP with the lowest priority has the highest preference. The highest priority value is 0. • Range: 0-255 • Default: 0
bootstrap-candidate priority default	Specifies a value of 0.

Candidate Rendezvous Point

Parameter	Description
candidate-rp <on off>	Specifies that the platform is a candidate rendezvous point router. • Default: off
candidate-rp local-address ip_address	Address used for the C-RP state machine and in the C-RP-Advertisements sent to the elected bootstrap router. If PIM clustering is enabled, then this address must be configured and must match that of one of the PIM interfaces. If clustering is not enabled, this address can either be that of one of the PIM interfaces or an address configured on the loopback interface. If an address from the loopback interface is used, take care not select an address in the 127/8 address range. • Range: Address of PIM interface or a non 127.0.0.0/8 loopback address. • Default: Selects the IP address of one of the interfaces on which PIM is enabled. The default does not apply if PIM clustering is enabled.
candidate-rp priority <0-255>	The priority of this C-RP. All PIM routers select the same RP for a multicast group address from the list of C-RPs received in the bootstrap messages from the elected BSR. The lower the Local Preference of the C-RP, the higher the priority. • Range: 0-255 • Default: 0
candidate-rp priority default	A value of 0.

Candidate Rendezvous Point

Parameter	Description
candidate-rp multicast-group <mcast_address>/<mask_length> <on off>	The multicast address advertised in the candidate rendezvous point advertisements. Enter the group multicast IP address and mask length. If you do not specify a group multicast address, the candidate rendezvous point advertises itself as the rendezvous point for all multicast groups. Multicast group address • Range: Dotted-quad ([224-239].[0-255].[0-255].[0-255]). • Default: None. Mask length: • Range: 1-32 • Default: None

Static Rendezvous Point

Parameter	Description
<code>static-rp off</code>	Disables the static rendezvous point option.
<code>static-rp rp-address ip_address <on off></code>	A static rendezvous point. If an associated multicast group and prefix is not configured, the static-rp is considered to be responsible for all multicast groups (224.0.0.0/4). This needs to be consistent with the RP information at other routers in a multicast domain irrespective of the RP-dissemination mechanism (bootstrap or autoRP) used. Note - The static RP overrides the RP information received from other RP-dissemination mechanisms, such as bootstrap routers. • Range: Any IP address • Default: None
<code>static-rp rp-address ip_address multicast-group <mcast_address>/<mask_length> <on off></code>	Mask length: • Range: 1-32 • Default: None

Sparse Mode Timers

Parameter	Description
<code>register-suppress-interval <60-3600></code>	The mean interval between receiving a register-stop and allowing registers to be sent again. A lower value means more frequent register bursts at the rendezvous point. A higher value means a longer join latency for new receivers. • Range: 60-3600 seconds • Default: 60 seconds
<code>register-suppress-interval default</code>	A value of 60 seconds.
<code>candidate-rp advertise-interval <1-3600></code>	The interval between which candidate-rendezvous point routers send candidate-rendezvous point advertisements to the elected bootstrap router. • Range: 1-3600 seconds • Default: 60 seconds
<code>candidate-rp advertise-interval default</code>	A value of 60 seconds.

Shortest Path First Threshold

Parameter	Description
<code>spt-threshold</code> <code>multicast</code> <code><mcast_address>/<mask_length> threshold</code> <code><0-1000000></code>	The interval between which candidate-rendezvous point routers send candidate-rendezvous point advertisements to the elected bootstrap router. • Range: 1-3600 seconds • Default: 60 seconds Mask length. • Range: 1-32 • Default: None The data rate threshold in Kbits/sec to trigger the spt switchover. When the data rate for a sparse-mode group exceeds the shortest-path-tree threshold at the last-hop router, an (S,G) entry is created and a join/prune message is sent toward the source. Setting this option builds a shortest-path tree from the source S to the last-hop router. • Range: 0-1000000 Kbits/s, or infinity (for no switchover) • Default: None
<code>spt-threshold</code> <code>multicast</code> <code><mcast_address>/<mask_length> threshold</code> <code>infinity <on off></code>	Specifies that there is no spt switchover. Multicast group address • Range: Dotted-quad ([224-239].[0-255].[0-255].[0-255]) • Default: 224.0.0.0/4 Mask length: • Range: 1-32 • Default: None

Timer and Assert Rank Parameters for Dense Mode and Sparse Mode

Use these commands to change or restore default values for timers and assert ranks.

```
set pim
  hello-interval <1-21845>
  hello-interval default
  data-interval <11-3600>
  data-interval default
  assert-interval <1-3600>
  assert-interval default
  assert-limit <10-100000>
  assert-limit default
  assert-limit <0>
  jp-interval <1-3600>
  jp-interval default
  jp-delay-interval <1-3600>
  jp-delay-interval default
  jp-suppress-interval <2-3600>
  jp-suppress-interval default
```

```

assert-rank protocol protocol name rank <0-255>
assert-rank protocol protocol name rank default
state-refresh-interval <1 - 255>
state-refresh-ttl <1 - 255>

```

General Timers

Parameter	Description
hello interval <1-21845>	Interval between which PIM hellos are sent on a multicast-capable interface. Hello messages are addressed to the All-PIM-Routers multicast group (224.0.0.13) so that PIM routers may discover neighbors on a multi-access network. • Range: 1-21845 seconds • Default: 30
hello interval default	A value of 30.
data-interval <11-3600>	The life-time of a new PIM forwarding entry. Subsequently the life of the entry will be extended in different ways depending on the location of this router in the network. For example, in some cases the receipt of PIM control messages (e.g. periodic join/prune messages) will extend the life of the entry and in others the presence of local senders of multicast traffic will prevent the entry from being deleted. • Range: 11-3600 seconds • Default: 210
data-interval default	A value of 210.
assert-interval <1-3600>	If an assert battle on an upstream interface results in a PIM neighbor other than the unicast reverse-path-forwarding (RPF) neighbor towards the source of the data traffic (for which the assert battle was generated) getting elected as the designated forwarder on that interface, the winner is used as the upstream neighbor for all subsequent join/prune messages. This change is timed-out after expiry of the assert interval (measured in seconds). • Range: 1-3600 seconds • Default: 180
assert-interval default	A value of 180.
assert-limit <10-10000>	The number of asserts to send per second. Asserts are generated by the router when data from a source is detected on an interface other than the incoming interface (based on the unicast routing table) towards the source. These messages are rate-limited and the router should not originate more than a fixed number of assert messages per second. If the limit is set to 0, rate-limiting is disabled. • Range: 0, 10-10000 • Default: 10

Parameter	Description
<code>assert-limit default</code>	A value of 10.
<code>assert-limit <0></code>	Disables the limit placed on the number of asserts that can be sent per second.
<code>jp-interval <1-3600></code>	Interval between sending Join/Prune messages. • Range: 1-3600 seconds • Default: 60
<code>jp-interval default</code>	A value of 60.
<code>jp-delay-interval <1-3600></code>	The maximum interval from the time when the unicast Reverse Path Forwarding (RPF) neighbor (towards a source or the RP) changes, and a triggered Join/Prune message is sent. • Range: 1-3600 seconds • Default: 5
<code>jp-delay-interval default</code>	A value of 5.
<code>jp-suppress-interval <2-3600></code>	Mean interval from receiving a Join/Prune with a higher Holdtime (with ties broken by higher network layer address) and allowing duplicate Join/Prunes to be sent again. Set this interval to 1.25 times the join/prune interval. • Range: 2-3600 seconds • Default: 75
<code>jp-suppress-interval default</code>	A value of 75.

Assert Ranks

<code>assert-rank protocol protocol name rank <0-255></code>	Used to compare the cost of protocols in order to determine which router will forward multicast data packets on a multi-access LAN. These values are used in assert messages sent out on a LAN when a router detects data packets on an interface other than the incoming interface towards the source of the data. These values must be the same for all routers on a multi-access LAN that are running the same protocol. Hence, the default values have been specifically configured to match that of other implementations. • Range: 0-255
--	---

assert-rank protocol protocol name rank <0-255>	Used to compare the cost of protocols in order to determine which router will forward multicast data packets on a multi-access LAN. These values are used in assert messages sent out on a LAN when a router detects data packets on an interface other than the incoming interface towards the source of the data. These values must be the same for all routers on a multi-access LAN that are running the same protocol. Hence, the default values have been specifically configured to match that of other implementations. • Range: 0-255
assert-rank protocol protocol name rank default	Default assert-rank values for supported protocols that match other implementations. • Defaults: • Direct: 0 • OSPF: 10 • Kernel: 40 • Static: 60 • RIP: 100 • BGP: 170

State Refresh Parameters- Dense Mode

state-refresh-interval <1 - 255>	For Dense Mode, the interval at which state refresh messages are sent for multicast traffic originated by directly-connected sources. • Range: 1-255 • Default: 60
state-refresh-ttl <1 - 255>	For Dense Mode, the time-to-live (TTL) placed in the state refresh messages originated for multicast traffic from directly-connected sources. This value can be used to limit the forwarding of state refresh messages in the network. In the absence of user configuration it is derived from the multicast data. • Range: 1-255 • Default: None

Static Multicast Routes

PIM expects packets to arrive on the reverse-path forwarding (RPF) interface - the interface used to reach the source of the multicast data. PIM also checks the RPF to learn which interface it should use to send join/prune messages. By default, PIM checks the unicast routing table to identify the RPF interface. Static multicast routes provide an alternative route table to use for the RPF check. If a static multicast route and a unicast route are available for a specific destination, PIM uses the static multicast route.

Static multicast routes let PIM be independent of unicast routing. This lets you deploy topologies in which multicast and unicast traffic flow over different paths. For example, in order to balance the traffic load, you can separate the HTTP traffic from the stock quotes traffic. You simply

configure a static multicast route to the source network that specifies a next hop gateway address different from the next hop address (for the same source) in the unicast routing table.



Note - PIM is the only protocol that uses static multicast routes.

You can configure static multicast routes through the WebUI or through the CLI.

Configuring Static Multicast Routes - WebUI

In WebUI, you can add, delete, and edit static multicast routes.

To add a static multicast route:

1. Open the **Advanced Routing > Static Multicast Routes** page of the WebUI.
2. Click **Add**.
The **Add Destination Route** window opens
3. Configure the multicast sender parameters:
 - **Destination** - the IP address
 - **Subnet mask** - the network mask
4. Configure the gateway:
 - a) Click **Add Gateway**.
 - b) In the window that opens, enter the **IPv4 Address** and the **Priority (optional)**.
 - c) Click **OK**.
5. Click **Save**.

To delete a static multicast route:

1. In the **Advanced Routing > Static Multicast Routes** page of the WebUI, select a route.
2. Click **Delete**.

To edit a static multicast route:

1. In the **Advanced Routing > Static Multicast Routes** page of the WebUI, select a route.
2. Click **Edit**.
The **Edit Destination Route** window opens.
3. For the selected Destination route, add, edit, or delete a gateway:
 - To add a gateway - click **Add Gateway**, enter the **IPv4 address** and the **Priority (optional)** value in the window that opens, and click **OK**
 - To edit a gateway - select a gateway, click **Edit**, change the **IPv4 address** and the **Priority (optional)** values in the window that opens, and click **OK**
 - To delete a gateway - select a gateway and click **Delete**
4. Click **Save**.

Configuring Static Multicast Routes - CLI

In CLI, you can add, delete, and verify static multicast routes.

To add or delete a static multicast route:

Run this command: `set static-mroute {<sender_IP>/<mask> | default} nexthop gateway address <gateway_IP> [priority <priority_value>] {on | off}`

Parameter	Description
<sender_IP>/<mask>	The address and the network mask of the multicast sender.
default	The default gateway to use for FPF lookups.
<gateway_IP>	The IP address of the gateway router.
priority <priority_value>	Valid values are 1 through 8. Defines the order of priority in which the nexthop gateway is selected, when multiple nexthop gateways are configured. Lower priority value takes higher preference. If not priority is defined, the default value of 1 is assigned. If multiple gateways are configured with the same priority value, the one with the lower IP address takes higher preference.
on	Installs the specified multicast static route in the routing table.
off	Removes the specified multicast static route from the routing table.

To verify static multicast routes:

Run this command: `show static-mroute`

To show static multicast routes configuration:

Run this command: `show configuration static-mroute`

Monitoring and Troubleshooting PIM

PIM Trace Options

To log information about PIM errors and events using the WebUI:

1. In the tree view, click **Advanced Routing > Routing Options**.
2. In the **Configuration** tab, select **Filter Visible Tables Below** and select **PIM**.
3. In the option variables area, do one of:
 - Double-click an option.
 - Select an option (to select multiple options, use Shift-Click) and click **Activate**.
4. Click **Apply** at the top of the page

To log information about PIM errors and events using the CLI:

```
set trace pim
```

```
  assert <on | off>
  bootstrap <on | off>
  crp <on | off>
  graft <on | off>
  hello <on | off>
  join <on | off>
  mfc <on | off>
  mrt <on | off>
  packets <on | off>
  rp <on | off>
  register <on | off>
  trap <on | off>
  traceoptions <on | off>
```

These trace options apply both to dense-mode and sparse-mode implementations:

Parameter	Description
assert	Trace PIM assert messages.
hello	Trace PIM router hello messages.
join	Trace PIM join/prune messages.
mfc	Trace calls to or from the multicast forwarding cache
mrt	Trace PIM multicast routing table events.
packets	Trace all PIM packets.
trap	Trace PIM trap messages.
all	Trace all PIM events and packets.

The following trace options apply to sparse-mode implementations only:

Parameter	Description
bootstrap	Trace bootstrap messages.
crp	Trace candidate-RP-advertisements.
rp	Trace RP-specific events, including RP set-specific and bootstrap-specific events.
register	Trace register and register-stop packets.

The following trace option applies to dense-mode implementations only:

Parameter	Description
graft	Trace graft and graft acknowledgment packets.

PIM Show and MFC Commands

Use these commands to monitor and troubleshoot PIM.

These commands apply to both dense-mode and sparse-mode PIM:

```
show pim
  interfaces
  interfaces if_address
  neighbors
  neighbor ip_address
  memory
  timers
  stats
  summary
```

These commands apply only to sparse-mode PIM:

```
show pim
  bootstrap
  candidate-rp
  joins
  rps
  sarse-mode-stats
  group-rp-mapping <mcast_address>
```

Dense Mode and Sparse Mode Parameters

Parameter	Description
interface	The interfaces that are running PIM, their status, and their mode. This command also displays the interface and its DR priority and the number of PIM neighbors on the interface.
neighbors	The IP address of each PIM neighbor and the interface on which the neighbor is present. This command also displays the neighbor's DR priority, generation ID, holdtime and the time the neighbor is set to expire based on the holdtime received in the most recent hello message.
neighbor ip_address	Use this command to verify whether a PIM neighbor supports DR Priority. For neighbors that advertise a DR election priority value, the following message appears in the summary: DRPriorityCapable Yes.
stats	The number of different types of PIM packets received and transmitted and any associated errors.
memory	
timers	
summary	

Sparse Mode Parameters

Parameter	Description
<code>bootstrap</code>	The IP address and state of the Bootstrap router.
<code>candidate-rp</code>	The state of the Candidate Rendezvous Point state machine.
<code>joins</code>	PIM's view of the join-prune (*, G and S, G) state, including RP for the group, incoming, and outgoing interface(s), interaction with the multicast forwarding cache and the presence of local members. To view the equivalent information for dense-mode PIM, use the <code>show mfc cache</code> command.
<code>rps</code>	The active RP-set, including the RP addresses, their type (or source of information about them) and the groups for which they are configured to act as RP.
<code>group-rp-mapping <group-address></code>	The RP selected for a particular group based on information from the active RP-set.
<code>sparse-mode stats</code>	Error statistics for multicast forwarding cache (MFC); Bootstrap Router (BSR) messages; Candidate Rendezvous Point (CRP) advertisements; and the Internet Group Management Protocol (IGMP).

MFC Commands and Trace Options

To log multicast errors and events, use the `show mfc` commands or the MFC Trace options in the WebUI or the CLI ("[MFC Trace Options](#)" on page 182).

```
show mfc
  cache
  interface
```

Parameter	Description
<code>cache</code>	Multicast source and group forwarding state by prefix.
<code>interface</code>	Multicast source and group forwarding state by interface.

Routing Monitor

In This Section:

Monitoring Routes - WebUI	228
Monitoring Routes - CLI (show route)	228
Monitoring the Routing Daemon - CLI (show routed)	229
Monitoring the Multicast Forwarding Cache - CLI (show mfc)	230

Use the routing monitor to show summary information about routes on the Gaia system.

Monitoring Routes - WebUI

To show the routes on the Gaia system using the WebUI:

1. In the tree view, click **Advanced Routing > Routing Monitor**.
2. Optional: In the **Filter Protocols** column, select the protocol whose routes you want to see. Shift-click to select multiple items.

Monitoring Routes - CLI (show route)

Use these commands to show information about active, inactive or all (both active and inactive) routes on your system for BGP, OSPF and RIP protocols:

```
show route
  rip
  bgp {aspath | communities | detailed | metrics | suppressed}
  inactive {bgp | rip}
  all {bgp | rip | ospf}
```

Use these commands to show information about active, inactive, or all routes on your system for OSPF:

```
show route
  ospf
  inactive ospf
  all ospf
```

Use these commands to show information about active, inactive and all aggregate routes on your system:

```
show route
  aggregate
  inactive aggregate
  all aggregate
```

Use these commands to show additional information about routes on your system:

```
show route
  all
  all direct
  all static
  all kernel
  direct
  inactive
  inactive direct
```

```

inactive static
inactive kernel
kernel
static
summary
destination ip_address
exact ip_prefix
less-specific ip_prefix
more-specific ip_prefix

```

Monitoring the Routing Daemon - CLI (show routed)

Use the following commands to view general information recorded by the Gaia routing daemon (routed).

```

show routed
  memory
  resources
  krt
  version

```

Parameter	Description
memory	Show the memory usage of the routing daemon, for each routing protocol running on the system. • Total memory usage. • Memory used by each routing protocol. • MFC- memory used for the multicast forwarding cache (MFC). • Core - memory used by routed for internal purposes.
resources	Show the following system information: • Total uptime • Total user time • Total system time • Page faults • Page reclaims • File system writes • File system reads • Message writes • Message reads • Signals received • Total swaps • Voluntary context switches • Involuntary context switches

Parameter	Description
krt	Show statistical information about the messages sent and received on the raw sockets between the kernel and routed. • KRT interface message count • KRT interface message length • KRT route message count (rx) • KRT route message length (rx) • KRT route message count (tx) • KRT route message length (tx) • KRT route adds • KRT route changes • KRT route deletes
version	Shows the following system information: • routed version • System start time • Current time • System uptime

Monitoring the Multicast Forwarding Cache - CLI (show mfc)

Use the following commands to see information about multicast forwarding cache (MFC) on your system.

```
show mfc
  cache
  summary
  interface
  orpans
  stats
```

Parameter	Description
cache	Shows MFC state information.

Parameter	Description
summary	Shows the following MFC state information: • Number of interfaces enabled • Number of cache entries • Kernel forwarding entry limit • Number of kernel forwarding entries • Cache entry average lifetime • Prune average lifetime • Cache age cycle • Data rate update interval • Multicast protocol (instance)
interface	Shows MFC interface state information.
orphans	Shows MFC orphan state information.
stats	Shows various information about the following MFC properties: • Resolve task summary • Resolve requests • RPF failure notifications • MFC maintenance

Regular Expressions and Character Sets

In This Appendix

Regular Expression Syntax	232
Special Characters in Clish	232

Regular Expression Syntax

This table shows the Check Point implementation of standard regular expression metacharacters.

Metacharacter	Name	Description
\	Backslash	escape metacharacters non-printable characters character types
[]	Square Brackets	character class definition
{ }	Parenthesis	sub-pattern, to use metacharacters on the enclosed string
{min[,max]}	Curly Brackets	min/max quantifier {n} - exactly n occurrences {n,m} - from n to m occurrences {n,} - at least n occurrences
.	Dot	match any character
?	Question Mark	zero or one occurrences (equals {0,1})
*	Asterisk	zero or more occurrences of preceding character
+	Plus Sign	one or more occurrences (equals {1,})
	Vertical Bar	alternative
^	Circumflex	anchor pattern to beginning of buffer (usually a word)
\$	Dollar	anchor pattern to end of buffer (usually a word)
-	hyphen	range in character class

Special Characters in Clish

- for the '?' character, type ctrl-v and then '?'
- for the '\' character, type '\\'