

# ArubaOS-Switch Multicast and Routing Guide for YC.16.04



Part Number: 5200-3134  
Published: July 2017  
Edition: 1

## Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

## Acknowledgments

Intel<sup>®</sup>, Itanium<sup>®</sup>, Pentium<sup>®</sup>, Intel Inside<sup>®</sup>, and the Intel Inside logo are trademarks of Intel Corporation in the United States and other countries.

Microsoft<sup>®</sup> and Windows<sup>®</sup> are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Adobe<sup>®</sup> and Acrobat<sup>®</sup> are trademarks of Adobe Systems Incorporated.

Java<sup>®</sup> and Oracle<sup>®</sup> are registered trademarks of Oracle and/or its affiliates.

UNIX<sup>®</sup> is a registered trademark of The Open Group.

<b>Chapter 1 About this guide</b> .....	<b>7</b>
Applicable products.....	7
Switch prompts used in this guide.....	7
<b>Chapter 2 Multimedia Traffic Control with IP Multicast (IGMP)</b> .....	<b>8</b>
Overview.....	8
IGMP general operation and features.....	8
IGMP operating features.....	8
Enhancements.....	8
Number of IP multicast addresses allowed.....	9
How IGMP operates.....	9
Operation with or without IP addressing.....	10
Automatic fast-leave IGMP.....	11
Forced fast-leave IGMP.....	13
Configuring delayed group flush.....	14
Configuring and displaying IGMP (CLI).....	14
Configuring per-port IGMP traffic filters.....	14
Configuring the querier function.....	15
Configuring static multicast groups.....	16
Viewing IGMP configuration for VLANs.....	16
Viewing the current IGMP configuration.....	16
Viewing IGMP high-level statistics for all VLANs on the switch.....	18
Viewing IGMP historical counters for a VLAN.....	19
Viewing IGMP group address information.....	19
Viewing IGMP group information for a VLAN with a filtered address.....	20
Enabling or disabling IGMP on a VLAN.....	20
IGMP proxy forwarding.....	21
How IGMP proxy forwarding works.....	21
Configuring IGMP proxy (CLI).....	22
VLAN context command.....	23
IGMP proxy show command.....	24
Using the switch as querier.....	26
Well-known or reserved multicast addresses excluded from IP multicast filtering.....	26
IP multicast filters.....	27
Reserved addresses excluded from IP multicast filtering.....	27
IGMPv3.....	27
IGMPv3 commands.....	28
<b>Chapter 3 IP Routing Features</b> .....	<b>39</b>
Overview.....	39
IP interfaces.....	39
IP tables and caches.....	40
IP route exchange protocols.....	41
IP global parameters for routing switches.....	41
ARP age timer.....	44
IP interface parameters for routing switches.....	45
Configuring IP parameters for routing switches.....	46
Configuring ARP parameters.....	46

Configuring forwarding parameters.....	47
Configuring ICMP.....	48
Disabling ICMP messages.....	48
Disabling replies to broadcast ping requests.....	48
Disabling ICMP destination unreachable messages.....	48
Disabling ICMP redirects.....	49
Configuring static IP routes.....	49
Static route types.....	50
Other sources of routes in the routing table.....	50
Static IP route parameters.....	50
Static route states follow VLAN states.....	50
Configuring a static IP route.....	51
Viewing static route information.....	53
Configuring the default route.....	53
Configuring RIP.....	53
Overview of RIP.....	53
RIP parameters and defaults.....	54
Configuring RIP parameters.....	55
Configuring RIP redistribution.....	58
Changing the route loop prevention method.....	59
Viewing RIP information.....	60
Configuring IRDP.....	65
Enabling IRDP globally.....	66
Enabling IRDP on an individual VLAN interface.....	66
Viewing IRDP information.....	67
Configuring DHCP relay.....	68
Overview.....	68
DHCP packet forwarding.....	68
Prerequisites for DHCP relay operation.....	68
Enabling DHCP relay.....	69
Configuring an IP helper address.....	69
Verifying the DHCP relay configuration.....	69
DHCP Option 82.....	70
UDP broadcast forwarding.....	82
Overview.....	82
Subnet masking for UDP forwarding addresses.....	83
Configuring and enabling UDP broadcast forwarding.....	83
<b>Chapter 4 IP Directed Broadcasts.....</b>	<b>88</b>
Enabling forwarding of IP directed broadcasts (CLI).....	88
Introduction to feature.....	88
CLI commands.....	88
Disabling the directed broadcasts.....	92
<b>Chapter 5 RIPv2 MD5 authentication.....</b>	<b>93</b>
Introduction.....	93
Configuration commands.....	93
Show commands.....	94
Operating notes.....	95
Validation rules.....	96
Log messages.....	96
Error messages.....	96

<b>Chapter 6 RIPng</b> .....	<b>98</b>
RIPng for IPv6.....	98
Configure RIPng.....	99
Enable/Disable RIPng global.....	99
Configure a RIPng setting.....	99
Configure a default metric.....	99
Configure the administrative distance for routes.....	100
Redistribute router RIPng.....	100
Configure RIPng timers.....	101
Enable/Disable RIPng traps.....	101
VLAN Level Configuration.....	102
IPv6 RIPng.....	102
Show commands.....	103
Show IPv6 ripng general.....	103
Show IPv6 ripng interface.....	103
Show IPv6 RIPng peer.....	104
Show IPv6 RIPng redistribute.....	104
Show IPv6 RIPng traps.....	105
Show IPv6 route RIPng.....	105
Show ipv6 route summary.....	106
Debug commands.....	106
Debug IPv6 RIPng.....	106
Additional commands.....	107
VLAN VLAN-ID IPv6.....	107
Show running config.....	107
Show running-config vlan.....	108
Validation rules.....	108
Event Log.....	109
<b>Chapter 7 Websites</b> .....	<b>111</b>
<b>Chapter 8 Support and other resources</b> .....	<b>112</b>
Accessing Hewlett Packard Enterprise Support.....	112
Accessing updates.....	112
Customer self repair.....	112
Remote support.....	113
Warranty information.....	113
Regulatory information.....	113
Documentation feedback.....	114
<b>Apple’s Bonjour and Google’s Chromecast</b> .....	<b>115</b>
Overview.....	115
mDNS Gateway.....	115
Service filtering.....	116
Wireless printer service process.....	116
Wireless Printer advertising printer service.....	117
Host 2 queries for printers.....	117
iPhone 1 queries for printers.....	118
Limitations of the mDNS gateway and Chromecast.....	118
Enabling mDNS feature.....	119
Create mDNS reflection.....	119

Create or delete a mDNS profile.....	119
Set rules for mDNS profile.....	119
Set the specific mDNS profile for VLAN.....	120
Set the global mDNS profile.....	120
Show mdns.....	121
Show mDNS gateway.....	121
Show mDNS profile configuration.....	122
Show mDNS profile name.....	122
Debug mDNS.....	123
Validation rules.....	123
RMON table.....	124

This guide provides information on how to configure IGMP, PIM and routing protocols.

## Applicable products

This guide applies to these products:

Aruba 2540 Switch Series (JL354A, JL355A, JL356A, JL357A)

## Switch prompts used in this guide

Examples in this guide are representative and may not match your particular switch/environment. The following table explains the types of command prompts that may be used in examples, along with information on what each prompt indicates.

Prompt	Explanation
switch#	# indicates manager context (authority).
switch>	> indicates operator context (authority).
switch(config)#	(config) indicates the config context.
switch(vlan-x)#	(vlan-x) indicates the vlan context of config, where x represents the VLAN ID. For example: switch(vlan-128)#.
switch(eth-x)#	(eth-x) indicates the interface context of config, where x represents the interface. For example: switch(eth-48)#.
switch-Stack#	Stack indicates stacking is enabled.
switch-Stack(config)#	Stack(config) indicates the config context while stacking is enabled.
switch-Stack(stacking)#	Stack(stacking) indicates the stacking context of config while stacking is enabled.
switch-Stack(vlan-x)#	Stack(vlan-x) indicates the vlan context of config while stacking is enabled, where x represents the VLAN ID. For example: switch-Stack(vlan-128)#.
switch-Stack(eth-x/y)#	Stack(eth-x/y) indicates the interface context of config, in the form (eth-<member-in-stack>/<interface>). For example: switch(eth-1/48)#

## Overview

This chapter describes multimedia traffic control with IP multicast-Internet Group Management Protocol (IGMP). IGMP reduces unnecessary bandwidth usage on a per-port basis. For general information about IGMP, see **IGMP general operation and features** on page 8.



The use of static multicast filters is described in the chapter titled "Traffic/Security Filters" in the Access Security Guide for your HPE switch.

## IGMP general operation and features

In a network where IP multicast traffic is transmitted for various multimedia applications, you can use the switch to reduce unnecessary bandwidth usage on a per-port basis by configuring IGMP. In the factory default state (IGMP disabled), the switch simply floods all IP multicast traffic it receives on a given VLAN through all ports on that VLAN (except the port on which it received the traffic.) This can result in significant and unnecessary bandwidth usage in networks where IP multicast traffic is a factor. Enabling IGMP allows the ports to detect IGMP queries and report packets and manage IP multicast traffic through the switch.

IGMP is useful in multimedia applications such as LAN TV, desktop conferencing, and collaborative computing, where there is multipoint communication, that is, communication from one to many hosts, or communication originating from many hosts and destined for many other hosts. In such multipoint applications, IGMP is configured on the hosts, and multicast traffic is generated by one or more servers (inside or outside of the local network.) Switches in the network (that support IGMP) can then be configured to direct the multicast traffic to only the ports where needed. If multiple VLANs are configured, you can configure IGMP on a per-VLAN basis.

Enabling IGMP allows detection of IGMP queries and report packets used to manage IP multicast traffic through the switch. If no other querier is detected, the switch then also functions as the querier. If you need to disable the querier feature, you can do so using the IGMP configuration CLI commands, see **Configuring the querier function** on page 15.



IGMP configuration on the switches operates at the VLAN context level.

## IGMP operating features

### Basic operation

In the factory default configuration, IGMP is disabled. To enable IGMP

- If multiple VLANs are not configured: Configure IGMP on the default VLAN (DEFAULT\_VLAN; VID=1.)
- If multiple VLANs are configured: Configure IGMP on a per-VLAN basis for every VLAN where this feature is to be used.

### Enhancements

With the CLI, you can configure these additional options:



Auto/blocked/ forward	<p>You can use the console to configure individual ports to any of the following states:</p> <p><b>Auto</b></p> <p>(Default) Causes the switch to interpret IGMP packets and to filter IP multicast traffic based on the IGMP packet information for ports belonging to a multicast group. This means that IGMP traffic will be forwarded on a specific port only if an IGMP host or multicast router is connected to the port.</p> <p><b>Blocked</b></p> <p>Causes the switch to drop all IGMP transmissions received from a specific port, and also blocks all outgoing IP Multicast packets for that port, thus preventing IGMP traffic from moving through specific ports.</p> <p><b>Forward</b></p> <p>Causes the switch to forward all IGMP and multicast transmissions through the port.</p>
Operation with or without IP addressing	<p>This feature helps to conserve IP addresses by enabling IGMP to run on VLANs that do not have an IP address. See <b>Operation with or without IP addressing</b> on page 10. This is also referred as IGMP Snooping.</p>
Querier capability	<p>The switch performs querier function for IGMP on VLANs having an IP address when there is no other device in the VLAN acting as querier. See <b>Using the switch as querier</b> on page 26.</p>

To configure high priority settings for traffic, see “Quality of Service: managing bandwidth more effectively” in the Advanced Traffic Management Guide.



Whenever IGMP is enabled, the switch generates an Event Log message only after the querier election.

IP multicast traffic groups are identified by IP addresses in the range of 224.0.0.0 to 239.255.255.255. Also, incoming IGMP packets intended for reserved, or “well-known” multicast addresses, automatically flood through all ports (except the port on which the packets entered the switch).

For more on this topic, see **Well-known or reserved multicast addresses excluded from IP multicast filtering** on page 26.

For more information about IGMP, see **How IGMP operates** on page 9.

## Number of IP multicast addresses allowed

The number of IGMP filters (addresses) and static multicast filters available is 2,038. Additionally, 16 static multicast filters are allowed. If multiple VLANs are configured, then each filter is counted once per VLAN in which it is used.



The number of IGMP and static multicast filters available for 2540 is 502 for both IPv4 and IPv6.

## How IGMP operates

IGMP is used by IP hosts to report their multicast group memberships with neighboring multicast routers. It is an internal protocol of the IP suite. IGMP messages are encapsulated in IP datagrams, with an IP protocol number of 2. When enabled, IGMP operates in VLAN context. As a result, there is no need of a multicast routing protocol as

long as the communication between IP hosts and multicast source is on the same network. If IP hosts and the multicast source are on different network segments, multicast routing is essential.

Multicast routers use IGMP to identify the groups having members on each of their attached physical networks. A multicast router or a switch enabled with IGMP can operate in any of the two roles:

- Querier
- Non Querier

Generally, only one Querier is available per physical network. When you enable IGMP, Querier election takes place and one of the devices perform the role of a Querier. The Querier is responsible for:

- Sending out IGMP group membership queries in a timed interval
- Retrieving IGMP membership reports from an active member
- Allowing to update the group membership table

The multicast group running version 2 of IGMP uses three fundamental types of messages to communicate:

Membership Query (Query)	<p>The message sent always by the Querier to all the devices on the attached network enabled with IGMP. The Membership Query message is of two types:</p> <ul style="list-style-type: none"> <li>• General Query — Used to learn the groups which have members on the attached network.</li> <li>• Group-specific Query — Used to learn if a particular group has any members on the attached network.</li> </ul> <p>The above message types are differentiated by the Group Address. The Membership Query messages are referred as Query messages. To disable the querier feature, use the IGMP configuration CLI commands. For more information about the CLI commands, see <a href="#">Configuring the querier function</a> on page 15.</p>
Report (Join)	<p>The message sent by a host to the querier to indicate that the host wants to be or is a member of a given group indicated in the report message.</p>
Leave group	<p>The message sent by a host to all routers 224.0.0.2 is to indicate that the host has ceased to be a member of a specific multicast group.</p>

An IP multicast packet includes the multicast group (address) to which the packet belongs. When an IGMP client connected to a switch port needs to receive multicast traffic from a specific group, it joins the group by sending an IGMP report (join request) to the network. (The multicast group specified in the join request is determined by the requesting application running on the IGMP client.) When a network device with IGMP enabled receives the join request for a specific group, it forwards any IP multicast traffic it receives for that group through the port on which the join request was received. When the client is ready to leave the multicast group, it sends a Leave Group message to the network and ceases to be a group member. When the leave request is detected, the appropriate IGMP device ceases transmitting traffic for the designated multicast group through the port on which the leave request was received (as long as there are no other current members of that group on the affected port.)

Thus, IGMP identifies members of a multicast group (within a subnet) and allows IGMP-configured hosts (and routers) to join or leave multicast groups.

To display IGMP data showing active group addresses, reports, queries, querier access port, and active group address data (port, type, and access), see section "Internet Group Management Protocol (IGMP) Status" in appendix B, "Monitoring and Analyzing Switch Operation" of the Management and Configuration Guide for your switch.

## Operation with or without IP addressing

You can configure IGMP on VLANs that do not have IP addressing. The benefit of IGMP without IP addressing is a reduction in the number of IP addresses you have to use and configure. This can be significant in a network with a large number of VLANs. The limitation on IGMP without IP addressing is that the switch cannot become

Querier on any VLANs for which it has no IP address—so the network administrator must ensure that another IGMP device will act as Querier. It is also advisable to have an additional IGMP device available as a backup Querier.

**Table 1: Comparison of IGMP operation with and without IP addressing**

IGMP function available with IP addressing configured on the VLAN	Available without IP addressing?	Operating differences without an IP address
Forward multicast group traffic to any port on the VLAN that has received a join request for that multicast group.	Yes	None
Forward join requests (reports) to the Querier.	Yes	None
Configure individual ports in the VLAN to Auto (the default)/Blocked, or Forward.	Yes	None
Configure IGMP traffic forwarding to normal or high-priority forwarding.	Yes	None
Age-out IGMP group addresses when the last IGMP client on a port in the VLAN leaves the group.	Yes	Requires that another IGMP device in the VLAN has an IP address and can operate as Querier. This can be a multicast router or another switch configured for IGMP operation. (Hewlett Packard Enterprise recommends that the VLAN also include a device operating as a backup Querier in case the device operating as the primary Querier fails for any reason.)
Support Fast-Leave IGMP and Forced Fast-Leave IGMP (below.)	Yes	
Support automatic Querier election.	No	Querier operation not available.
Operate as the Querier.	No	Querier operation not available.
Available as a backup Querier.	No	Querier operation not available.

## Automatic fast-leave IGMP

IGMP fast-leave is configured for ports on a per-VLAN basis. By default, the switches send IGMP Group-Specific Query message out of the interface, upon which the Leave Group message is received to ensure that no other receivers are connected to the interface. If receivers are directly attached to the switch, there is no point in sending the membership query as the receiver wanting to leave is the only connected host.

Fast-leave processing eliminates the IGMP Group-Specific Query message. Thus, it allows the switch to immediately remove an interface from the bridge table upon receiving the Leave Group message. This processing speeds up the overall leave process and also eliminates the CPU overhead of having to generate an IGMP Group-Specific Query message.

Depending on the switch model, fast-leave is enabled or disabled in the default configuration.

With fast-leave enabled and an IGMP Group Leave being received on a noncascaded port, the following events take place:

- The switch stops forwarding multicast traffic for that group to that port.
- Does not apply to cascaded ports.

When disabled or when the port is cascaded, the regular IGMP leave time is used (up to 10 seconds when the switch is not the IGMP Querier).

On switches that do not support data-driven IGMP, unregistered multicast groups are flooded to the VLAN rather than pruned. In this scenario, fast-leave IGMP can actually increase the problem of multicast flooding by removing the IGMP group filter before the Querier has recognized the IGMP leave. The Querier will continue to transmit the multicast group during this short time, and because the group is no longer registered, the switch will then flood the multicast group to all ports.

On HPE switches that do support data-driven IGMP ("Smart" IGMP), when unregistered multicasts are received the switch automatically filters (drops) them. Thus, the sooner the IGMP leave is processed, the sooner this multicast traffic stops flowing.

Because of the multicast flooding problem mentioned above, the IGMP fast-leave feature is disabled by default on all switches that do not support data-driven IGMP (see **Operation with or without IP addressing** on page 10). The feature can be enabled on these switches via an SNMP set of this object:

```
hpSwitchIcmpPortForceLeaveState.<vid>.<port number>
```

However, Hewlett Packard Enterprise does not recommend this because it will increase the amount of multicast flooding during the period between the client's IGMP leave and the Querier's processing of that leave. For more information on this topic, see **Forced fast-leave IGMP** on page 13.

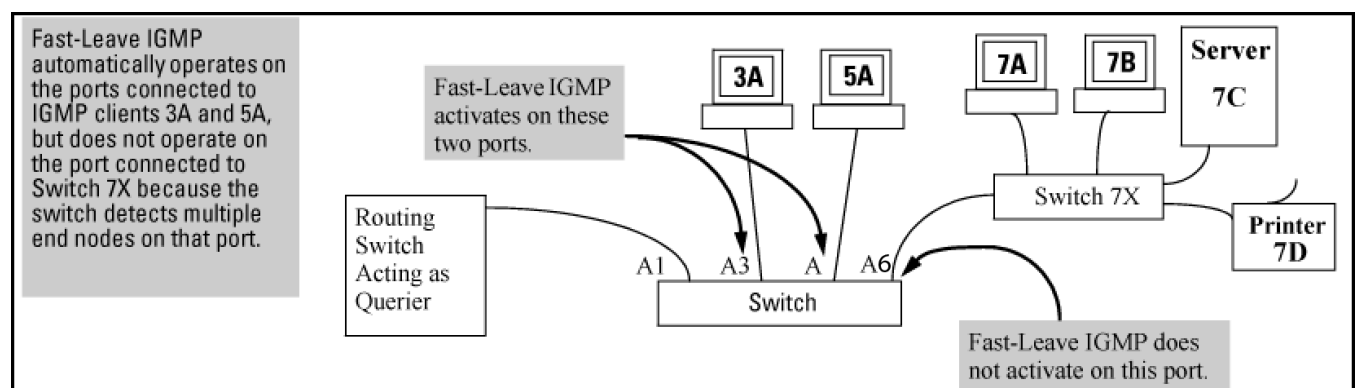
If a switch port has the following characteristics, the fast-leave operation will apply:

- Connected to only one end node.
- The end node currently belongs to a multicast group, that is, is an IGMP client.
- The end node subsequently leaves the multicast group.

Then the switch does not need to wait for the Querier status update interval, but instead immediately removes the IGMP client from its IGMP table and ceases transmitting IGMP traffic to the client. (If the switch detects multiple end nodes on the port, automatic fast-leave does not activate—regardless of whether one or more of these end nodes are IGMP clients.)

In the following figure, automatic fast-leave operates on the switch ports for IGMP clients "3A" and "5A," but not on the switch port for IGMP clients "7A" and "7B," server "7C," and printer "7D."

**Figure 1: Example of automatic fast-leave IGMP criteria**



When client "3A" running IGMP is ready to leave the multicast group, it transmits a Leave Group message. Because the switch knows that there is only one end node on port A3, it removes the client from its IGMP table and halts multicast traffic (for that group) to port A3. If the switch is not the Querier, it does not wait for the actual Querier to verify that there are no other group members on port A3. If the switch itself is the Querier, it does not query port A3 for the presence of other group members.

Fast-leave operation does not distinguish between end nodes on the same port that belong to different VLANs. Thus, for example, even if all the devices on port A6 shown in figure 1 belong to different VLANs, fast-leave does not operate on port A6.

## Default (enabled) IGMP operation solves the "delayed leave" problem

Fast-leave IGMP is enabled by default. When fast-leave is disabled and multiple IGMP clients are connected to the same port on an IGMP device (switch or router), if only one IGMP client joins a given multicast group, then later sends a Leave Group message and ceases to belong to that group, the switch automatically retains that IGMP client in its IGMP table and continues forwarding IGMP traffic to the IGMP client until the Querier triggers confirmation that no other group members exist on the same port. This delayed leave operation means that the switch continues to transmit unnecessary multicast traffic through the port until the Querier renews multicast group status.

## Configuring fast-leave IGMP

For information about fast-leave IGMP, see [Automatic fast-leave IGMP](#) on page 11.

### Syntax:

```
ip igmp fastleave <port-list>
no ip igmp fastleave <port-list>
```

Enables IGMP fast-leaves on the specified ports in the selected VLAN.

The `no` form of the command disables IGMP fast-leave on the specified ports in the selected VLAN.

Use `show running` to display the ports per-VLAN on which fast-leave is disabled.

Default: Enabled

## Forced fast-leave IGMP

When enabled, forced fast-leave IGMP speeds up the process of blocking unnecessary IGMP traffic to a switch port that is connected to multiple end nodes. (This feature does not activate on ports where the switch detects only one end node.) For example, in [Figure 1: Example of automatic fast-leave IGMP criteria](#) on page 12, even if you configured forced fast-leave on all ports in the switch, the feature would activate only on port A6 (which has multiple end nodes) when a Leave Group request arrived on that port.

When a port having multiple end node receives a Leave Group request from one end node for a given multicast group "X," forced fast-leave activates and waits a small amount of time to receive a join request from any other group "X" member on that port. If the port does not receive a join request for that group within the forced-leave interval, the switch then blocks any further group "X" traffic to the port.

## Configuring forced fast-leave IGMP

For information about forced fast-leave, see [Forced fast-leave IGMP](#) on page 13.

### Syntax:

```
vlan <vid> ip igmp forcedfastleave <port-list>
no vlan <vid> ip igmp forcedfastleave <port-list>
```

Enables IGMP forced fast-leave on the specified ports in the selected VLAN, even if they are cascaded.

The `no` form of the command disables forced fast-leave on the specified ports in the selected VLAN.

Use `show running` to display the ports per-VLAN on which forced fast-leave is enabled.

Default: Disabled

<code>show running-config</code>	Displays a nondefault IGMP forced fast-leave configuration on a VLAN. If configured, the <code>show running-config</code> output does not include forced fast-leave.
<code>forcedfastleave</code>	Can be used when there are multiple devices attached to a port.

## Configuring delayed group flush

When enabled, this feature continues to filter IGMP groups for a specified additional period after IGMP leaves have been sent. The delay in flushing the group filter prevents unregistered traffic from being forwarded by the server during the delay period. In practice, this is rarely necessary on the switches, which support data-driven IGMP. (Data-driven IGMP, which is enabled by default, prunes off any unregistered IGMP streams detected on the switch.)

### Syntax:

```
igmp delayed-flush <0-255>
```

Where leaves have been sent for IGMP groups, enables the switch to continue to flush the groups for a specified period. This command is applied globally to all IGMP-configured VLANs on the switch.

Range: 0 - 255; Default: Disabled (0)

### Syntax:

```
show igmp delayed-flush
```

Displays the current `igmp delayed-flush` setting.

## Configuring and displaying IGMP (CLI)

### Configuring per-port IGMP traffic filters

#### Syntax:

```
vlan <vid> ip igmp [auto < port-list > | blocked < port-list > | forward < port-list >]
```

Used in the VLAN context, specifies how each port handles IGMP traffic.

Default: auto.



Where a static multicast filter is configured on a port, and an IGMP filter created by this command applies to the same port, the IGMP filter overrides the static multicast filter for any inbound multicast traffic carrying the same multicast address as is configured in the static filter. See section "Filter Types and Operation" in the "Port Traffic Controls" chapter of the Management and Configuration Guide for your switch.

#### Example:

Suppose that you want to configure IGMP as follows for VLAN 1 on the 100/1000T ports on a module in slot 1:

Ports 1-2	auto	Filter multicast traffic. Forward IGMP traffic to hosts on these ports that belong to the multicast group for which the traffic is intended. (Also forward any multicast traffic through any of these ports that is connected to a multicast router.)
Ports 3-4	forward	Forward all multicast traffic through this port.
Ports 5-6	blocked	Drop all multicast traffic received from devices on these ports.

The different states of IGMP control traffic are auto, forward and blocked.

auto	(Default) Causes the switch to interpret IGMP packets and to filter IP multicast traffic based on the IGMP packet information for ports belonging to a multicast group. This means that IGMP traffic will be forwarded on a specific port only if an IGMP host or multicast router is connected to the port.
forward	Causes the switch to forward all IGMP and multicast transmissions through the port.
blocked	Causes the switch to drop all IGMP transmissions received from a specific port, and also blocks all outgoing IP Multicast packets for that port, thus preventing IGMP traffic from moving through specific ports.

For a description of the default behavior of data-driven switches, see **Automatic fast-leave IGMP** on page 11.

Depending on the privilege level, you could use one of the following commands to configure IGMP on VLAN 1 with the above settings:

```
switch(config)# vlan 1 ip igmp auto 1,2
switch(config)# vlan 1 ip igmp forward 3,4
switch(config)# vlan 1 ip igmp blocked 5,6

switch(vlan-1)# ip igmp auto 1,2
switch(vlan-1)# vlan 1 ip igmpforward 3,4
switch(vlan-1)# blocked 5,6
```

The following command displays the VLAN and per-port configuration resulting from the above commands.

```
switch> show igmp vlan 1 config
```

## Configuring the querier function

### Syntax:

```
vlan <vid> ip igmp querier
no vlan <vid> ip igmp querier
```

This command disables or re-enables the ability for the switch to become querier if necessary.

The `no` version of the command disables the querier function on the switch. The `show ip igmp config` command displays the current querier command.

Default querier capability: Enabled



It is necessary to have IP address for a switch to perform in a Querier role.

## Configuring static multicast groups

Use this command to configure a group on the switch so that multicast traffic for that group can be forwarded with a receiver host. Traffic will be flooded to all the ports in the VLAN for this group.

### Syntax:

```
ip igmp static-group <group-address>
no ip igmp static-group <group-address>
```



This command must be issued in a VLAN context.

Creates the IGMP static group with the specified *<group address>* on the selected VLAN. The no form of the command deletes the static group on the selected VLAN.

## Viewing IGMP configuration for VLANs

### Syntax:

```
show ip igmp [vlan < vid >]
```

Displays IGMP configuration for a specified VLAN or for all VLANs on the switch.

### Displaying IGMP status for a VLAN

```
switch(vlan-22)# show ip igmp vlan 22
```

```
IGMP Service Protocol Info
```

```
Total VLANs with IGMP enabled           : 2
Current count of multicast groups joined  : 2
```

```
IGMP Filter Unknown Multicast: Disabled
IGMP Filter Unknown Multicast Status: Disabled
```

```
VLAN ID : 22
VLAN Name : VLAN22
IGMP version : 2
Querier Address [this switch] : 10.255.128.2
Querier Port :
Querier UpTime : 1h 23m 55s
Querier Expiration Time : 0h 1m 49s
```

Active Group Addresses	Type	Expires	Ports	Reports	Queries
226.0.6.7	Filter	0h 4m 6s	1	2	0
226.0.6.8	Filter	0h 4m 5s	2	2	0

## Viewing the current IGMP configuration

### Syntax:

```
show ip igmp config
```

Displays IGMP configuration for all VLANs on the switch.



**Syntax:**

```
show ip igmp vlan <vid> config
```

Displays IGMP configuration for a specific VLAN on the switch, including per-port data.

For IGMP operating status, see the section "Internet Group Management Protocol (IGMP) status" in the chapter "Monitoring and Analyzing Switch Operation" of the Management and Configuration Guide for your switch.

**Example:**

Suppose that you have the following VLAN and IGMP configurations on the switch:

VLAN ID	VLAN name	IGMP enabled	Querier
1	DEFAULT_VLAN	Yes	No
22	VLAN-2	Yes	Yes
33	VLAN-3	No	Yes

You could use the CLI to display this data as follows:

**Listing of IGMP configuration for all VLANs in the switch**

```
switch(vlan-33)# show ip igmp config
```

```
IGMP Service Config
```

```
Control unknown multicast [Yes] : Yes
Forced fast leave timeout [0] : 4
Delayed flush timeout [0] : 0
Look-up Mode [mac] : mac
```

VLAN ID	VLAN Name	IGMP Enabled	Querier Allowed	IGMP Version	Querier Interval
1	DEFAULT_VLAN	Yes	No	2	125
22	VLAN22	Yes	Yes	2	125
33	VLAN33	No	Yes	2	125

```
switch(vlan-33)# show run
```

```
Running configuration:
```

```
; JL253A Configuration Editor; Created on release #WC.16.02.0000x
; Ver #0d:13.ef.7c.5f.fc.6b.fb.9f.fc.f3.ff.37.ef:09
```

```
hostname "Aruba-2930F-24G-4SFPP"
module 1 type jl253a
snmp-server community "public" unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged 1-28
  ip address dhcp-bootp
  ip igmp
  no ip igmp querier
  exit
vlan 22
  name "VLAN22"
  no ip address
```

```

ip igmp
exit
vlan 33
  name "VLAN33"
  no ip address
exit

```

The following version of the `show ip igmp` command includes the VLAN ID (*vid*) designation, and combines the above data with the IGMP per-port configuration:

### Listing of IGMP configuration for a specific VLAN

```
switch(vlan-22)# show ip igmp vlan 22 config
```

```
IGMP Service VLAN Config
```

```

VLAN ID : 22
VLAN NAME : VLAN22
IGMP Enabled [No] : Yes
Querier Allowed [Yes] : Yes
IGMP Version [2] : 2
Strict Mode : No
Last Member Query Interval (Seconds) [1] : 1
Querier Interval [125] : 125
Query Max. Response Time (Seconds) [10] : 10
Robustness Count [2] : 2

```

Port	Type	Port Mode	Forced	Fast Leave	Fast Leave
1	1000T	Auto	No		Yes
2	1000T	Auto	No		Yes
3	1000T	Blocked	No		Yes
4	1000T	Forward	No		Yes

1 IGMP configuration for the selected VLAN.

2 IGMP configuration on the individual ports in the VLAN.

## Viewing IGMP high-level statistics for all VLANs on the switch

### Syntax:

```
show ip igmp statistics
```

### Displaying statistics for IGMP joined groups

```
switch(vlan-22)#show ip igmp statistics
```

```
IGMP Service Statistics
```

```

Total VLANs with IGMP enabled : 2
Current count of multicast groups joined : 2

```

```
IGMP Joined Groups Statistics
```

VLAN ID	VLAN Name	Total	Filtered	Standard	Static
EXCLUDE	INCLUDE				

1	DEFAULT_VLAN	52	50	0	2
NA	NA				
22	VLAN22	80	75	5	0
NA	NA				
33	VLAN33	1100	1000	99	1
NA	NA				

## Viewing IGMP historical counters for a VLAN

### Syntax:

```
show ip igmp vlan <vid> counters
```

### Display of IGMP historical counters for a VLAN

```
switch(config)# show ip igmp vlan 1 counters
```

```
IGMP service Vlan counters
```

```
VLAN ID : 1
```

```
VLAN Name : DEFAULT_VLAN
```

```

General Query Rx           : 58
General Query Tx           : 58
Group Specific Query Rx    : 3
Group Specific Query Tx    : 3
V1 Member Report Rx       : 0
V2 Member Report Rx       : 2
V3 Member Report Rx       : 0
Leave Rx                    : 0
Unknown IGMP Type Rx      : 0
Unknown Pkt Rx            : 0
Forward to Routers Tx Counter : 0
Forward to Vlan Tx Counter : 0
Port Fast Leave Counter   : 0
Port Forced Fast Leave Counter : 0
Port Membership Timeout Counter : 0

```

## Viewing IGMP group address information

### Syntax:

```
show ip igmp groups
```

### Displaying IGMP groups address information

```
switch(vlan-22)# show ip igmp groups
```

```
IGMP Group Address Information
```

VLAN ID	Group Address	Expires	UpTime	Last Reporter	Type
22	226.0.6.7	0h 3m 26s	0h 14m 22s	10.255.128.1	Filter
22	226.0.6.8	0h 3m 19s	0h 13m 20s	10.255.128.3	Filter
22	239.20.255.9	0h 0m 0s	0h 0m 0s		Static

Sample configuration is as shown:

```

switch(vlan-22)# show run

Running configuration:

; JL253A Configuration Editor; Created on release #WC.16.02.0000x
; Ver #0d:13.ef.7c.5f.fc.6b.fb.9f.fc.f3.ff.37.ef:09

hostname "Aruba-2930F-24G-4SFPP"
module 1 type jl253a
snmp-server community "public" unrestricted
vlan 1
    name "DEFAULT_VLAN"
    no untagged 1-4
    untagged 5-28
    no ip address
    ip igmp
    no ip igmp querier
    exit
vlan 22
    name "VLAN22"
    untagged 1-4
    ip address 10.255.128.2 255.255.255.0
    ip igmp
    ip igmp blocked 3
    ip igmp forward 4
    ip igmp static-group 239.20.255.9
    exit
vlan 33
    name "VLAN33"
    no ip address
    exit

```

## Viewing IGMP group information for a VLAN with a filtered address

### Syntax:

```
show ip igmp vlan <vid> group <ip-addr>
```

### Group information for a VLAN with a filtered address group

```

switch(vlan-22)# show ip igmp vlan 22 group 226.0.6.7

IGMP ports and group information for group 226.0.6.7

VLAN ID: 22
Uptime: 0h 15m 32s
Last Reporter: 10.255.128.1
Type: Filter

  Port   Port Type  Port Mode Expires Access
  ----   -
  1      1000T     Auto      253   host

```

## Enabling or disabling IGMP on a VLAN

You can enable IGMP on a VLAN, along with the last-saved or default IGMP configuration (whichever was most recently set), or you can disable IGMP on a selected VLAN.

## Syntax:

```
ip igmp
no ip igmp
```

Enables IGMP on a VLAN. This command must be executed in a VLAN context.

### Enabling IGMP on VLAN 1

```
switch(vlan-1)# vlan 1 ip igmp
```

– or –

```
switch(vlan-1)# ip igmp
```

### Disabling IGMP on VLAN 1

```
switch(config)# no vlan 1 ip igmp
```

```
switch(vlan-1)# no ip igmp
```



If you disable IGMP on a VLAN and then later re-enable IGMP on that VLAN, the switch restores the last-saved IGMP configuration for that VLAN. For more information on how switch memory operates, see the chapter "Switch Memory and Configuration" in the Management and Configuration Guide for your switch.

You can also combine the `ip igmp` command with other IGMP-related commands, as described in the following sections.

## IGMP proxy forwarding

When a network has a border router connecting a PIM-SM domain to a PIM-DM domain, the routers that are completely within the PIM-DM domain have no way to discover multicast flows in the PIM-SM domain. When an IGMP join occurs on a router entirely within the PIM-DM domain for a flow that originates within the PIM-SM domain, it is never forwarded to the PIM-SM domain.

The IGMP proxy is a way to propagate IGMP joins across router boundaries. The proxy triggers the boundary router connected to a PIM-SM domain to query for multicast flows and forward them to the PIM-DM domain. IGMP needs to be configured on all VLAN interfaces on which the proxy is to be forwarded or received, and PIM-DM must be running for the traffic to be forwarded.

You can configure an IGMP proxy on a selected VLAN that will forward IP joins (reports) and IGMP leaves to the upstream border router between the two multicast domains. You must specify the VLANs on which the proxy is enabled as well as the address of the border router to which the joins are forwarded.

### How IGMP proxy forwarding works

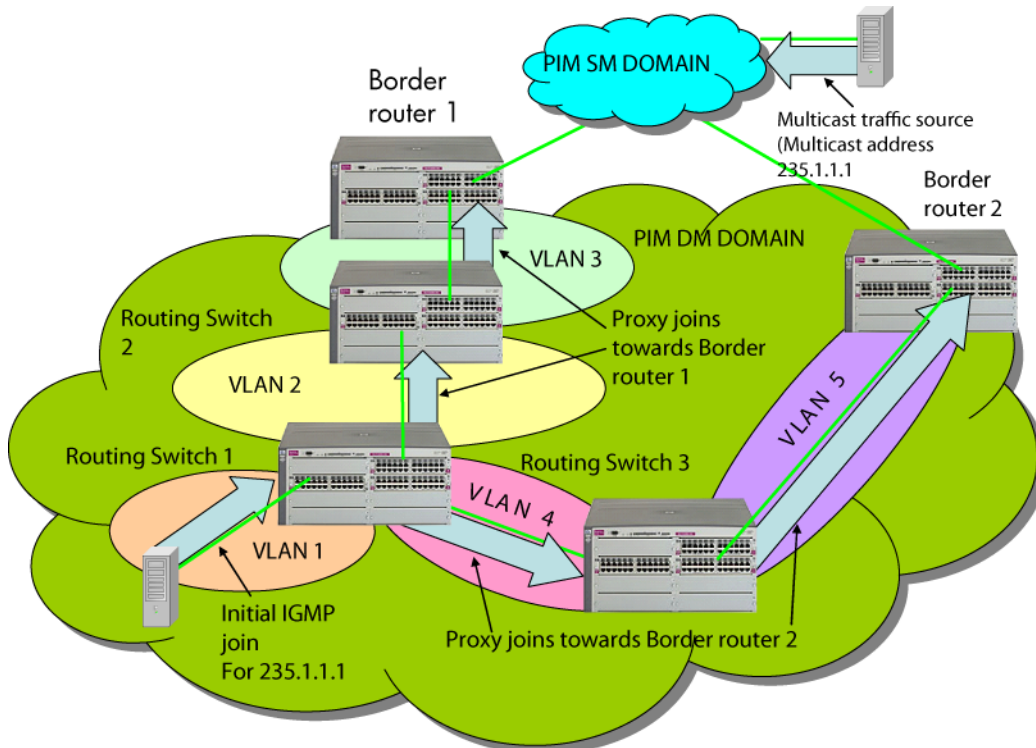
The following steps illustrate how to flood a flow from the PIM-SM domain into the PIM-DM domain when an IGMP join for that flow occurs in the PIM-DM domain. See **Figure 2: IGMP proxy example** on page 22.

#### Procedure

1. Configure Routing Switch 1 with the IGMP proxy forwarding function to forward joins toward Border Router 1; in addition, configure Routing Switch 1 to forward joins from VLAN 1 toward Border Router 2, as is VLAN 4 on Routing Switch 3.
2. Configure VLAN 2 on Routing Switch 2 to forward joins toward Border Router 1.

- When the host connected in VLAN 1 issues an IGMP join for multicast address 235.1.1.1, the join is proxied by Routing Switch 1 onto VLAN 2 and onto VLAN 4. The routing information table in Routing Switch 1 indicates that the packet to Border Router 1 and Border Router 2 is on VLAN 2 and VLAN 4, respectively.

**Figure 2: IGMP proxy example**



- Routing Switch 2 then proxies the IGMP join into VLAN 3, which is connected to Border Router 1.
- Border Router 1 uses PIM-SM to find and connect to the multicast traffic for the requested traffic. The traffic is flooded into the PIM-DM network where it is routed to the original joining host.
- Additionally, the join was proxied from Routing Switch 3 to Border Router 2. At first, both border routers will flood the traffic into the PIM-DM domain. However, PIM-DM only forwards multicasts based on the shortest reverse path back to the source of the traffic as determined by the unicast routing tables (routing FIB.) Only one multicast stream is sent to the joining host. This configuration provides a redundant in case the first fails.

## Configuring IGMP proxy (CLI)

For more information on IGMP proxy, see [IGMP general operation and features](#) on page 8.



### Adding or leaving a multicast domain

#### Syntax:

```
igmp-proxy-domain <domain-name> [< border-router-ip-address > | <mcast-range |
all>]
no igmp-proxy-domain <domain-name> [< border-router-ip-address > | <mcast-range |
all>]
```

The `no` form of the command is used to remove a multicast domain.

All VLANs associated with the domain must first be removed for this command to work. See the `no` form of `igmp-proxy` in the VLAN context command.

<code>&lt;domain-name&gt;</code>	User-defined name to associate with the PIM border router and multicast range that is being sent toward the border router.
<code>&lt;border-router-ip-addr&gt;</code>	<p>The IP address of the border router toward which IGMP proxy packets are sent. Not required for the <code>no</code> form of the command.</p> <p> The current routing FIB determines the best path toward the border router and therefore the VLAN that a proxy is sent out on</p>
<code>all</code>   <code>&lt;low-bound-ip-address  </code>	<p>The low boundary (inclusive) of the multicast address range to associate with this domain (for example, 234.0.0.1.) If <code>all</code> is selected, the multicast addresses in the range of 224.0.1.0 to 239.255.255.255 are included in this domain.</p> <p> Addresses 224.0.0.0 to 224.0.0.255 are never used, because these addresses are reserved for protocols.</p>
<code>&lt;high-bound-ip-address&gt;</code>	The high boundary (inclusive) of the multicast address range to associate with this domain (for example, 236.1.1.1.)

The following example shows the IGMP proxy border IP address (111.11.111.111) being configured.

#### IGMP proxy border IP address command

```
switch(config)# igmp-proxy-domain Bob 111.11.111.111
```

The following example shows the lower and upper boundaries of the multicast address range associated with the domain named Bob.

#### Setting the lower and upper bounds for multicasting

```
switch(config)# igmp-proxy-domain Bob 111.11.111.111 234.0.0.1
switch(config)# igmp-proxy-domain Bob 111.11.111.111 236.1.1.1
```

## VLAN context command

This command is performed when in VLAN context mode. When a query occurs on the upstream interface, an IGMP join is sent for all multicast addresses that are currently joined on the downstream interface.

#### Syntax:

```
igmp-proxy <domain-name>
no igmp-proxy <domain-name>
```

Tells the VLAN which IGMP proxy domains to use with joins on the VLAN.

The `no` version of the command with no domain name specified removes all domains associated with this VLAN.



Multiple different domains may be configured in the same VLAN context where the VLAN is considered the downstream interface. The domain name must exist prior to using this command to add the domain. If the unicast routing path to the specified IP address was through the specified VLAN, no proxy IGMP would occur, that is, a proxy is not sent back out on the VLAN that the IGMP join came in on.

If no unicast route exists to the border router, no proxy IGMP packets are sent.

## IGMP proxy show command

### Syntax:

```
show igmp-proxy {<entries | domains | vlans>}
```

Shows the currently active IGMP proxy entries, domains, or VLANs.

### Showing active IGMP proxy entries

```
switch(config)# show igmp-proxy entries
```

Total number of multicast routes: 2

Multicast Address	Border Address	VID	Multicast Domain
234.43.209.12	192.168.1.1	1	George
235.22.22.12	15.43.209.1	1	SAM
226.44.3.3	192.168.1.1	2	George

### Showing IGMP proxy domains

```
switch(config)# show igmp-proxy domains
```

Total number of multicast domains: 5

Multicast Domain	Multicast Range	Border Address	Active entries
George	225.1.1.1/234.43.209.12	192.168.1.1	2
SAM	235.0.0.0/239.1.1.1	15.43.209.1	1
Jane	236.234.1.1/236.235.1.1	192.160.1.2	0
Bill	ALL	15.43.209.1	0

### Showing active IGMP proxy VLANs

```
switch(config)# show igmp-proxy vlans
```

IGMP PROXY VLANs

VID	Multicast Domain	Active entries
1	George	1
1	Sam	1
1	Jane	0
2	George	1
4	George	0
4	Bill	0



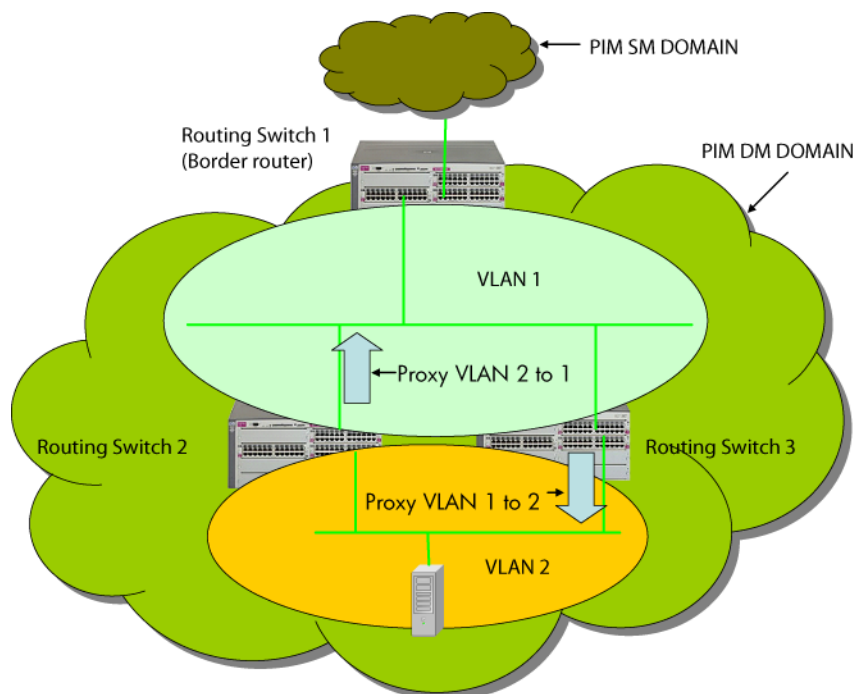
## Operating notes for IGMP proxy forwarding

- You can configure up to 12 multicast domains, which indicate a range of multicast addresses and the IP address of the PIM-SM/PIM-DM border router.
- You must give each domain a unique name, up to 20 characters.
- The domains may have overlapping multicast ranges.
- The IP address of the border router may be the same or different in each configured domain.
- Duplicate IGMP joins are automatically prevented, or leaves that would remove a flow currently joined by multiple hosts.
- Range overlap allows for redundant connectivity and the ability for multicasts to arrive from different border routers based on the shortest path back to the source of the traffic.
- The configured domain names must be associated with one or more VLANs for which the proxy joins are to be done.
- All routers in the path between the edge router receiving the initial IGMP packets and the border router have to be configured to forward IGMP using IGMP proxy.
- All upstream and downstream interfaces using IGMP proxy forwarding require IGMP and PIM to be enabled.
- You must remove all VLAN associations with the domain name before that domain name can be removed.
- The appropriate border routers must be used for each VLAN, or PIM-DM will not forward the traffic. This could occur when multiple border routers exist. It may be necessary to configure multiple overlapping domains if the multicast source address can generate the same multicast address and have different best paths to the PIM-DM domain.



Be careful to avoid configuring an IGMP forward loop, because this would leave the VLANs in a joined state forever once an initial join is sent from a host. For example, a join is issued from the host in VLAN 2 and Routing Switch 2 will proxy the join onto VLAN 1. Routing Switch 3 will then proxy the join back onto VLAN 2 and increment its internal count of the number of joins on VLAN 2. Even after the host on VLAN 2 issues a leave, the proxy join will remain and refresh itself each time a query occurs on VLAN 2. This type of loop could be created with multiple routers if an IGMP proxy is allowed to get back to the VLAN of the router that initially received the IGMP join from a host as shown.

**Figure 3:** Proxy loop scenario



## Using the switch as querier

The function of the IGMP Querier is to poll other IGMP-enabled devices in an IGMP-enabled VLAN to elicit group membership information. The switch performs this function if there is no other device in the VLAN, such as a multicasterouter, to act as Querier. Although the switch automatically ceases Querier operation in an IGMP-enabled VLAN if it detects another Querier on the VLAN, you can also use the switch's CLI to disable the Querier capability for that VLAN.



A Querier is required for proper IGMP operation. For this reason, if you disable the Querier function on a switch, ensure that there is an IGMP Querier (and, preferably, a backup Querier) available on the same VLAN.

If the switch becomes the Querier for a particular VLAN (for example, the DEFAULT\_VLAN), then subsequently detects queries transmitted from another device on the same VLAN, the switch ceases to operate as the Querier for that VLAN. If this occurs, the switch Event Log lists a pair of messages similar to these:

```
I 01/15/12 09:01:13 igmp:
DEFAULT_VLAN: Other Querier detected
I 01/15/12 09:01:13 igmp:
DEFAULT_VLAN: This switch is no longer Querier
```

In the above scenario, if the other device ceases to operate as a Querier on the default VLAN, the switch detects this change and can become the Querier as long as it is not pre-empted by some other IGMP Querier on the VLAN. In this case, the switch Event Log lists messages similar to the following to indicate that the switch has become the Querier on the VLAN:

```
I 01/15/12 09:21:55 igmp: DEFAULT_VLAN:
Querier Election in process
I 01/15/12 09:22:00 igmp: DEFAULT_VLAN:
This switch has been elected as Querier
```

## Well-known or reserved multicast addresses excluded from IP multicast filtering

Each multicast host group is identified by a single IP address in the range of 224.0.0.0 through 239.255.255.255. Specific groups of consecutive addresses in this range are termed "well-known" addresses and are reserved for predefined host groups. IGMP does not filter these addresses, so any packets the switch receives for such addresses are flooded out all ports assigned to the VLAN on which they were received (except the port on which the packets entered the VLAN.)

The following table lists the 32 well-known address groups (8192 total addresses) that IGMP does not filter on.



"X" is any value from 0 to 255.

**Table 2:** IP multicast address groups excluded from IGMP filtering

Groups of consecutive addresses in the range of 224.0.0.X to 239.0.0.X		Groups of consecutive addresses in the range of 224.128.0.X to 239.128.0.X	
224.0.0.x	232.0.0.x	224.128.0.x	232.128.0.x
225.0.0.x	233.0.0.x	225.128.0.x	233.128.0.x

Table Continued

Groups of consecutive addresses in the range of 224.0.0.X to 239.0.0.X		Groups of consecutive addresses in the range of 224.128.0.X to 239.128.0.X	
226.0.0.x	234.0.0.x	226.128.0.x	234.128.0.x
227.0.0.x	235.0.0.x	227.128.0.x	235.128.0.x
228.0.0.x	236.0.0.x	228.128.0.x	236.128.0.x
229.0.0.x	237.0.0.x	229.128.0.x	237.128.0.x
230.0.0.x	238.0.0.x	230.128.0.x	238.128.0.x
231.0.0.x	239.0.0.x	231.128.0.x	239.128.0.x



With aliasing limitation associated with MAC mode, certain non reserved multicast IP addresses are displayed as "reserved" addresses.

For example: 225.0.0.x Multicast IP address is aliased to 224.0.0.x to be displayed as "reserved".

## IP multicast filters

IP multicast addresses occur in the range from 224.0.0.0 through 239.255.255.255 (which corresponds to the Ethernet multicast address range of 01005e-000000 through 01005e-7fffff.) When a switch has a static traffic/security filter configured with a "multicast" filter type and a "multicast address" in this range, the switch will use the static filter unless IGMP learns of a multicast group destination in this range. In this case, IGMP dynamically takes over the filtering function for the multicast destination addresses, as long as the IGMP group is active. If the IGMP group subsequently deactivates, the switch returns filtering control to the static filter.

### Reserved addresses excluded from IP multicast filtering

Traffic to IP multicast groups in the IP address range of 224.0.0.0 to 224.0.0.255 will always be flooded because addresses in this range are "well known" or "reserved" addresses. Thus, if IP multicast is enabled, and there is an IP multicast group within the reserved address range, traffic to that group will be flooded instead of filtered by the switch.



In IP mode, nonreserved multicast IP addresses are not displayed as "reserved" addresses.

## IGMPv3

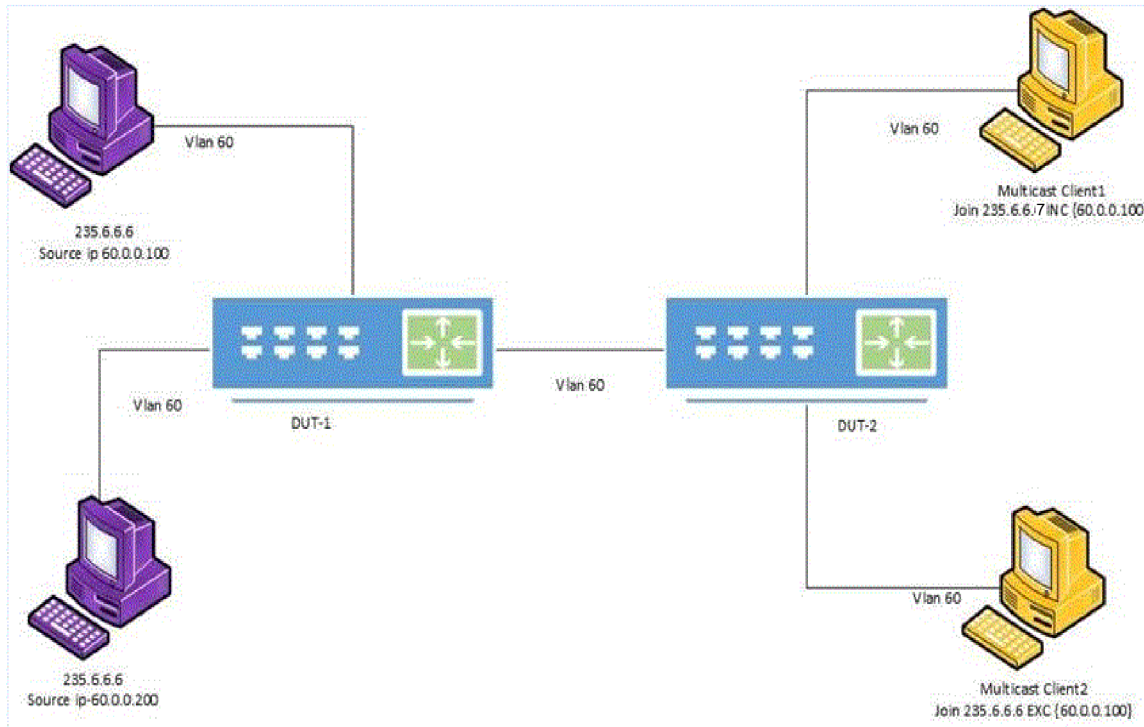
The Internet Group Management Protocol (IGMP) is used by IPv4 systems (hosts and routers) to report their IP multicast group membership to any neighboring multicast routers. This chapter is to describe version 3 of IGMP. Version 1, specified in [RFC-1112], was the first widely deployed version. Version 2, specified in [RFC-2236], added support for "low leave latency", that is, a reduction in the time it takes for a multicast router to learn that there are no longer any members of a particular group present on an attached network. Version 3 adds support for "source filtering", that is, the ability for a system to report interest in receiving packets \*only\* from specified source addresses, or from \*all but\* specified source addresses, sent to a particular multicast address. Version 3 is designed to be interoperable with Versions 1 and 2.

In the following figure, DUT-1 becomes the igmpv3 querier. Client-1 start receiving multicast traffic for group 235.6.6.6 from source 60.0.0.100 and client-2 start receiving multicast traffic for group 235.6.6.6 from source 60.0.0.200.



If multiple igmp version devices are available in the network, the igmp querier device must have the lower version of IGMP. This can be achieved by executing the `no ip igmp querier` command under the **vlan** context on other devices.

**Figure 4:** Basic topology and configuration for IGMPv3



**Table 3:** IGMPv3 configuration for Basic topology and configuration for IGMPv3

DUT-1 configurations	DUT-2 configurations
DUT-1(config)#igmp lookup-mode ip	DUT-2(config)#igmp lookup-mode ip
DUT-1(config)#vlan 60 ip address 60.0.0.1/24	DUT-2(config)#vlan 60 ip address 60.0.0.2/24
DUT-1(config)#vlan 60 ip igmp version 3	DUT-2(config)#vlan 60 ip igmp version 3
	DUT-2(config)#no vlan 60 ip igmp querier

## IGMPv3 commands

### igmp lookup-mode

To first configure IGMPv3, the igmp lookup-mode must be changed from the default mac mode to ip mode. Use the `ip igmp lookup-mode` command to set the IGMP snooping lookup mode.



IGMPv2 works both in ip mode and mac mode. Lookup-mode is applicable with IGMP disabled on all VLANs.

### Syntax

```
ip igmp lookup-mode
```

### Options

mac: Uses MAC look-up. (Default value)

ip: Uses IP look-up.

## igmp reload

This command is used to reset the IGMP state on all interfaces.

### Syntax

```
igmp reload
```

### Example output

```
IGMP application is in Error State as System Resources are exhausted. Traffic will flood.
```

```
Please disable IGMP on all VLANs or Issue the Command "igmp reload" to take it out of Error.
```

```
Refer to your product manual for information on IGMP resource consumption.
```

```
this is the output for igmp reload
```

## ip igmp

Use the **vlan** context to configure IGMPv3 on the switch.

### Syntax

```
ip igmp
```

### Options

**last-member-query-interval**: Sets the time interval that the querier waits to receive a response from members to a group-specific query message. It also specifies the amount of time between successive group-specific query messages; the default value is 1 second.

**query-max-response-time**: Sets the time interval to wait for a response to a query; the default value is 10 seconds.

**robustness**: Sets the number of times to retry a query; the default value is 2.

**version**: Sets the IGMP version to use; the default value is 2.

## ip igmp version

This command sets the IGMP version and completes igmpv3 configuration, enabling igmpv3 on the switch. Note that the default value is 2.

### Syntax

```
ip igmp version
```

```
no ip igmp version
```

### Parameters

<2-3>: The protocol version to use; the default is 2.

no: resets the version to 2.

## ip igmp last-member-query-interval

### Syntax

```
ip igmp last-member-query-interval
```

```
no ip igmp last-member-query-interval
```

### Parameters

<1-2>: The number of seconds between successive group-specific query messages; the default is 1.

The `no` version resets the value to its default value of 1 second.

## ip igmp querier

By default, IGMP querier is enabled. To disable querier functionality, use the following command:

```
switch (vlan 1)#no ip igmp querier
```

### Syntax

```
ip igmp querier
```

### Parameters

`interval`: Sets the interval in seconds between IGMP queries; the default is 125.

## ip igmp query-max-response-time

### Syntax

```
ip igmp query-max-response-time
```

```
no ip igmp query-max-response-time
```

### Parameters

<10-128>: The number of seconds to wait for a response to a query; the default value is 10.

The `no` version resets the value to its default value of 10 seconds.

## ip igmp robustness

### Syntax

```
ip igmp robustness
```

```
no ip igmp robustness
```

### Parameters

<1-8>: The number of times to retry a query; the default is 2.

The `no` version resets the value to its default value of 2.

## show ip igmp

This command is used to show IGMP information for all VLANs

### Syntax

```
show ip igmp
```

### Example output

Below is the output when version is set to 3.

```
switch(vlan-60)# show ip igmp
```

```
IGMP Service Protocol Info
```

```
Total VLANs with IGMP enabled           : 1
Current count of multicast groups joined  : 2
```

```
IGMP Filter Unknown Multicast: Disabled
IGMP Filter Unknown Multicast Status: Disabled
```

```
VLAN ID : 1
VLAN Name : DEFAULT_VLAN
IGMP version : 2
IGMP is not enabled
```

```
VLAN ID : 60
VLAN Name : VLAN60
IGMP version : 3
Querier Address : 60.0.0.1
Querier Port : 23
Querier UpTime : 0h 10m 9s
Querier Expiration Time : 0h 3m 34s
```

Active Group	Addresses	Tracking	Vers	Mode	Uptime	Expires
235.6.6.6		Filter	3	INC	0m 3s	4m 17s
235.6.6.7		Filter	3	EXC	0m 3s	4m 16s

Sample configuration is as shown:

```
switch(vlan-60)# show run
```

Running configuration:

```
; JL253A Configuration Editor; Created on release #WC.16.02.0000x
; Ver #0d:13.ef.7c.5f.fc.6b.fb.9f.fc.f3.ff.37.ef:09
```

```
hostname "Aruba-2930F-24G-4SFPP"
module 1 type jl253a
igmp lookup-mode ip
snmp-server community "public" unrestricted
vlan 1
    name "DEFAULT_VLAN"
    no untagged 1-2,23
    untagged 3-22,24-28
    ip address dhcp-bootp
    exit
vlan 60
    name "VLAN60"
    untagged 1-2,23
    ip address 60.0.0.2 255.255.255.0
    ip igmp
    no ip igmp querier
    ip igmp version 3
    exit
```

## show ip igmp vlan 1

This command is used to show IGMP information for a VLAN.

### Syntax

```
show ip igmp vlan 1
```

## Example output

Below is the output when version is set to 3.

```
switch(vlan-60)# show ip igmp vlan 60

IGMP Service Protocol Info

Total VLANs with IGMP enabled           : 1
Current count of multicast groups joined : 2

IGMP Filter Unknown Multicast: Disabled
IGMP Filter Unknown Multicast Status: Disabled

VLAN ID : 60
VLAN Name : VLAN60
IGMP version : 3
Querier Address : 60.0.0.1
Querier Port : 23
Querier UpTime : 0h 11m 44s
Querier Expiration Time : 0h 4m 5s

Active Group Addresses Tracking Vers Mode Uptime Expires
-----
235.6.6.6           Filter 3   INC 1m 38s 4m 13s
235.6.6.7           Filter 3   EXC 1m 38s 4m 19s
```

## show ip igmp vlan group

This command is used to show IGMP group information for a VLAN.

### Syntax

```
show ip igmp vlan <vid> group
```

### Example output

Below is the output when version is set to 3.

Port and source ipv4 address options are introduced under `group`. The following output captures the details of these options.

```
switch(config)# show ip igmp vlan <vid> group
  IPV4-ADDR      Show IGMP VLAN group address information.
  PORT           Show a list of all the IGMP groups on the specified port.

switch(config)# show ip igmp vlan <vid> group <ip4-addr>
  source         Show IGMP VLAN source address information.

switch(config)# show ip igmp vlan <vid> group <ip4-addr> source
  IPV4-ADDR      Specify the source IPv4 address.

switch(config)# show ipv4 igmp vlan <vid> group <ip4-addr> source <ip4-addr>

switch(vlan-60)# show ip igmp vlan 60 group 235.6.6.6

IGMP ports and group information for group 235.6.6.6
```



VLAN ID : 60    VLAN Name : VLAN60

Group Address : 235.6.6.6  
Last Reporter : 10.255.128.1  
Group Type : Filter

Port	Vers	Mode	Uptime	Expires	V1 Timer	V2 Timer	Filter Timer	Sources Forwarded	Sources Blocked
1	3	INC	2m 38s	3m 13s	-	0m 0s	-	1	0

Group Address : 235.6.6.6  
Source Address : 60.0.0.100  
Source Type : Filter

Port	Mode	Uptime	Expires	Configured Mode
1	INC	2m 38s	3m 13s	auto

### Usage errors

Error condition	Error message
Attempt to pass a nonexistent group	ipv4 address Group address is not found.

### show ip igmp vlan group source

This command is used to show IGMP group/source information for a VLAN.

#### Syntax

```
show ip igmp vlan <vid> group <ip4-addr> source <ip4-addr>
```

#### Example output

Below is the output when version is set to 3.

```
switch(vlan-60)# show ip igmp vlan 60 group 235.6.6.6 source 60.0.0.100
VLAN ID : 60    VLAN Name : VLAN60

Group Address : 235.6.6.6
Source Address : 60.0.0.100
Source Type : Filter
```

Port	Mode	Uptime	Expires	Configured Mode
1	INC	3m 31s	2m 20s	auto

### Usage errors

Error condition	Error message
Attempt to pass a nonexistent group	ipv4 address Group address is not found.

### show ip igmp vlan group port

This command is used to show IGMP group/source information for a VLAN port.

#### Syntax

```
show ip igmp vlan <vid> group <ip4-addr> port <port>
```

### Example output

Below is the output when version is set to 3.

```
switch(vlan-60)# show ip igmp vlan 60 group port 1
```

```
VLAN ID : 60      Name : VLAN60
```

```
Group Address : 235.6.6.6
Last Reporter : 10.255.128.1
Group Type    : Filter
```

Port	Vers	Mode	Uptime	Expires	Timer	Timer	Timer	Forwarded	Blocked
1	3	INC	8m 53s	3m 24s	-	0m 0s	-	1	0

```
Group Address : 235.6.6.6
Source Address : 60.0.0.100
Source Type    : Filter
```

Port	Mode	Uptime	Expires	Configured Mode
1	INC	8m 54s	3m 23s	auto

### show ip igmp vlan counters

This command is used to show IGMP counters for a VLAN.

#### Syntax

```
show ip igmp vlan <vid> counters
```

### Example output

Below is the output when version is set to 3.

```
switch(vlan-60)# show ip igmp vlan 60 counters
```

```
IGMP service Vlan counters
```

```
VLAN ID : 60      NAME : VLAN60
```

	Rx	Tx
V1 All Hosts Query	0	0
V2 All Hosts Query	0	0
V3 All Hosts Query	12	0
V1 Group Specific Query	0	0
V2 Group Specific Query	0	0
V3 Group Specific Query	8	0
Group and Source Specific Query	12	0
V3 Member Report	22	22
V2 Member Report	8	0
V1 Member Report	0	0
V2 Member Leave	0	0
Forward to Routers	0	32
Forward to VLAN	0	26

```
Errors:
```

```

Unknown IGMP Type          0
Unknown Packet             0
Malformed Packet          0
Bad Checksum               0
Martian Source             0
Packet received on IGMP-disabled Interface 0
Interface Wrong Version Query 0

```

Port Counters:

```

Fast Leave          : 4
Forced Fast Leave  : 0
Membership Timeout  : 0

```

## show ip igmp vlan statistics

This command is used to show IGMP statistics for a VLAN.

### Syntax

```
show ip igmp vlan <vid> statistics
```

### Example output

Below is the output when version is set to 3.

```
switch(vlan-60)# show ip igmp vlan 60 statistics
```

IGMP Statistics

```

VLAN ID : 60
VLAN Name : VLAN60

```

```

Number of Filtered Groups      : 2
Number of Standard Groups     : 0
Number of Static Groups       : 0
Total Multicast Groups Joined : 2

```

Mode	EXCLUDE	INCLUDE
Filtered	1	1
Standard	0	0
Total	1	1

## show ip igmp statistics

This command is used to show global IGMP statistics.

### Syntax

```
show ip igmp statistics
```

### Example output



The `show ip igmp statistics` is common for both IGMPv2 and IGMPv3. Output for the “EXCLUDE” and “INCLUDE” columns is displayed as “NA” if the version configured is IGMPv2 (as shown in the following example).

```
switch# show ip igmp statistics
```

```
IGMP Service Statistics
```

```
Total VLANs with IGMP enabled           : 1
Current count of multicast groups joined  : 2
```

```
IGMP Joined Groups Statistics
```

VLAN ID	VLAN Name	Total	Filtered	Standard	Static
EXCLUDE	INCLUDE				
1	DEFAULT_VLAN	2	2	0	0

### show ip igmp vlan config

This command is used to show the IGMP configuration for a VLAN.

#### Syntax

```
show ip igmp vlan (vid) config
```

#### Example output

Below is the output when version is set to 3.

```
switch(vlan-60)# show ip igmp vlan 60 config
```

```
IGMP Service VLAN Config
```

```
VLAN ID : 60
VLAN NAME : VLAN60
IGMP Enabled [No] : Yes
Querier Allowed [Yes] : No
IGMP Version [2] : 3
Strict Mode : No
Last Member Query Interval (Seconds) [1] : 1
Querier Interval [125] : 125
Query Max. Response Time (Seconds) [10] : 10
Robustness Count [2] : 2
```

Port	Type	Port Mode	Forced Fast Leave	Fast Leave
1	1000T	Auto	No	Yes
2	1000T	Auto	No	Yes
23	1000T	Auto	No	Yes

### show ip igmp config

This command is used to show the global IGMP configuration.

#### Syntax

```
show ip igmp config
```

## Example output

Below is the output when version is set to 3.

```
switch(vlan-60)# show ip igmp config
```

```
IGMP Service Config
```

```
Control unknown multicast [Yes] : Yes
Forced fast leave timeout [0] : 4
Delayed flush timeout [0] : 0
Look-up Mode [mac] : ip
```

VLAN ID	VLAN Name	IGMP Enabled	Querier Allowed	IGMP Version	Querier Interval
1	DEFAULT_VLAN	No	Yes	2	125
60	VLAN60	Yes	No	3	125

## show ip igmp vlan group

This command is used to show IGMP group information for a VLAN.

### Syntax

```
show ip igmp vlan <vid> group
```

### Example output

```
switch# show ip igmp vlan 60 group
```

```
IGMP ports and group information for group 235.6.6.6
```

```
VLAN ID : 60    VLAN Name : VLAN60
```

```
Group Address : 235.6.6.6
Last Reporter : 10.255.128.1
Group Type    : Filter
```

Port	Vers	Mode	Uptime	Expires	V1 Timer	V2 Timer	Filter Timer	Sources Forwarded	Sources Blocked
1	3	INC	15m 47s	2m 44s	-	0m 0s	-	1	0

```
Group Address : 235.6.6.6
Source Address : 60.0.0.100
Source Type    : Filter
```

Port	Mode	Uptime	Expires	Configured Mode
1	INC	15m 47s	2m 44s	auto

```
IGMP ports and group information for group 235.6.6.7
```

```
VLAN ID : 60    VLAN Name : VLAN60
```

```
Group Address : 235.6.6.7
Last Reporter : 10.255.128.3
Group Type    : Filter
```

V1	V2	Filter	Sources	Sources
----	----	--------	---------	---------

Port	Vers	Mode	Uptime	Expires	Timer	Timer	Timer	Forwarded	Blocked
2	3	EXC	15m 48s	2m 39s	-	0m 0s	2m 39s	0	1

Group Address : 235.6.6.7  
 Source Address : 60.0.0.100  
 Source Type : Filter

Port	Mode	Uptime	Expires	Configured Mode
2	EXC	15m 48s	2m 39s	auto

## igmp reload

This command is used to reset IGMP on all interfaces when error state is displayed.

### Syntax

```
igmp reload
```

### Example output

IGMP application is in Error State as System Resources are exhausted. Traffic will flood.  
 Please disable IGMP on all VLANs or Issue the Command "igmp reload" to take it out of Error.  
 Refer to your product manual for information on IGMP resource consumption.  
 this is the ouput for igmp reload

## Overview

The switch offers the following IP routing features:

IP Static routes	Up to 256 static routes
RIP (Router Information Protocol)	Supports RIP Version 1, Version 1 compatible with Version 2 (default), and Version 2
IRDP (ICMP Internet Router Discovery Protocol)	Advertises the IP addresses of the routing interfaces on this switch to directly attached host systems
DHCP Relay	Allows you to extend the service range of your DHCP server beyond its single local network segment
Source MAC-based ARP attack detection (ARP throttle)	ARP throttle protects the switch CPU from ARP attacks by enabling restriction of the overall number of ARP packets the CPU receives from a given client.

Throughout this chapter, the switches are referred to as "routing switches." When IP routing is enabled on your switch, it behaves just like any other IP router.

Basic IP routing configuration consists of adding IP addresses, enabling IP routing, and enabling a route exchange protocol, such as RIP.

For configuring the IP addresses, see the chapter "Configuring IP Addresses" in the Management and Configuration Guide for your switch. Use the information in this chapter if you need to change some of the IP parameters from their default values or if you want to view configuration information or statistics.

## IP interfaces

On the routing switches, IP addresses are associated with individual VLANs. By default, there is a single VLAN (Default\_VLAN) on the routing switch. In that configuration, a single IP address serves as the management access address for the entire device. If routing is enabled on the routing switch, the IP address on the single VLAN also acts as the routing interface.

Each IP address on a routing switch must be in a different subnet. You can have only one VLAN interface in a given subnet. For example, you can configure IP addresses 192.168.1.1/24 and 192.168.2.1/24 on the same routing switch, but you cannot configure 192.168.1.1/24 and 192.168.1.2/24 on the same routing switch.

You can configure multiple IP addresses on the same VLAN.

The number of IP addresses you can configure on an individual VLAN interface is 32.

You can use any of the IP addresses you configure on the routing switch for Telnet, Web management, or SNMP access, as well as for routing.



All HPE devices support configuration and display of IP address in classical subnet format (example: 192.168.1.1 255.255.255.0) and Classless Interdomain Routing (CIDR) format (example: 192.168.1.1/24.) You can use either format when configuring IP address information. IP addresses are displayed in classical subnet format only.

## IP tables and caches

### ARP cache table

The ARP cache contains entries that map IP addresses to MAC addresses. Generally, the entries are for devices that are directly attached to the routing switch.

An exception is an ARP entry for an interface-based static route that goes to a destination that is one or more router hops away. For this type of entry, the MAC address is either the destination device's MAC address or the MAC address of the router interface that answered an ARP request on behalf of the device, using proxy ARP.

### ARP cache

The ARP cache contains dynamic (learned) entries. The software places a dynamic entry in the ARP cache when the routing switch learns a device's MAC address from an ARP request or ARP reply from the device.

The software can learn an entry when the switch or routing switch receives an ARP request from another IP forwarding device or an ARP reply. Here is an example of a dynamic entry:

#### ARP cache dynamic entry

	IP Address	MAC Address	Type	Port
1	207.95.6.102	0800.5afc.ea21	Dynamic	6

Each entry contains the destination device's IP address and MAC address.

To configure other ARP parameters, see [Configuring ARP parameters](#) on page 46.

### IP route table

The IP route table contains routing paths to IP destinations.



The default gateway, which you specify when you configure the basic IP information on the switch, is used only when routing is not enabled on the switch.

### Routing paths

The IP route table can receive the routing paths from the following sources:

- Directly connected destination, which means there are no router hops to the destination
- Static IP route, which is a user-configured route
- Route learned through RIP

### Administrative distance

The IP route table contains the best path to a destination. When the software receives paths from more than one of the sources listed above, the software compares the administrative distance of each path and selects the path with the lowest administrative distance. The administrative distance is a protocol-independent value from 1 to 255.

The IP route table is displayed by entering the `show ip route` command from any context level in the console CLI. Here is an example of an entry in the IP route table:

#### IP route table entry

Destination	Gateway	VLAN	Type	Sub-Type	Metric	Dist.
10.10.10.1/32	10.10.12.1		connected		1	0



Each IP route table entry contains the destination's IP address and subnet mask and the IP address of the next-hop router interface to the destination. Each entry also indicates route type, and for OSPF routes, the subtype, and the route's IP metric (cost.) The type indicates how the IP route table received the route.

To configure a static IP route, see [Configuring a static IP route](#) on page 51.

## IP forwarding cache

The IP forwarding cache provides a fast-path mechanism for forwarding IP packets. The cache contains entries for IP destinations. When a routing switch has completed processing and addressing for a packet and is ready to forward the packet, the device checks the IP forwarding cache for an entry to the packet's destination.

- If the cache contains an entry with the destination IP address, the device uses the information in the entry to forward the packet out the ports listed in the entry. The destination IP address is the address of the packet's final destination. The port numbers are the ports through which the destination can be reached.
- If the cache does not contain an entry, the software can create an entry in the forwarding cache.

Each entry in the IP forwarding cache has an age timer. The age interval depends on the number of entries in the table. The age timer ranges from 12 seconds (full table) to 36 seconds (empty table.) Entries are aged only if they are not being used by traffic. If you have an entry that is always being used in hardware, it will never age. If there is no traffic, it will age in 12 to 36 seconds. The age timer is not configurable.



---

You cannot add static entries to the IP forwarding cache.

---

## IP route exchange protocols

The switch supports the RIP IP route exchange protocol.

This protocol provides routes to the IP route table and is disabled by default. For configuration information, see [Configuring RIP parameters](#) on page 55.


## IP global parameters for routing switches

The following table lists the IP global parameters and the page where you can find more information about each parameter.

**Table 4: IP global parameters for routing switches**

Parameter	Description	Default	See page
Address Resolution Protocol (ARP)	A standard IP mechanism that routers use to learn the MAC address of a device on the network. The router sends the IP address of a device in the ARP request and receives the device's MAC address in an ARP reply.	Enabled	<a href="#"><b>Configuring ARP parameters</b></a>
ARP age	The amount of time the device keeps a MAC address learned through ARP in the device's ARP cache. The device resets the timer to zero each time the ARP entry is refreshed and removes the entry if the timer reaches the ARP age. (Can be set using the menu interface to be as long as 1440 minutes. Go to <b>Menu &gt; Switch Configuration &gt; IP Config.</b> ) See <a href="#"><b>ARP age timer</b></a> .	Five minutes	N/A
Time to Live (TTL)	The maximum number of routers (hops) through which a packet can pass before being discarded. Each router decreases a packet's TTL by 1 before forwarding the packet. If decreasing the TTL causes the TTL to be 0, the router drops the packet instead of forwarding it.	64 hops	See the chapter "Configuring IP Addressing" in the Management and Configuration Guide.

*Table Continued*

Parameter	Description	Default	See page
Directed broadcast forwarding	<p>A directed broadcast is a packet containing all ones (or in some cases, all zeros) in the host portion of the destination IP address. When a router forwards such a broadcast, it sends a copy of the packet out each of its enabled IP interfaces.</p> <p> <b>NOTE</b></p> <p>You also can enable or disable this parameter on an individual interface basis. See <a href="#">IP interface parameters for routing switches</a>.</p>	Disabled	<a href="#">Enabling forwarding of IP directed broadcasts (CLI)</a>
ICMP Internet Router Discovery Protocol (IRDP)	<p>An IP protocol that a router can use to advertise the IP addresses of its router interfaces to directly attached hosts. You can enable or disable the protocol at the Global CLI Config level. You also can enable or disable IRDP and configure the following protocol parameters on an individual VLAN interface basis at the VLAN Interface CLI Config level.</p> <ul style="list-style-type: none"> <li>• Forwarding method (broadcast or multicast)</li> <li>• Hold time</li> <li>• Maximum advertisement interval</li> <li>• Minimum advertisement interval</li> <li>• Router preference level</li> </ul>	Disabled	<a href="#">Configuring IRDP</a>

*Table Continued*

Parameter	Description	Default	See page
Static route	An IP route you place in the IP route table.	No entries	<a href="#">Static route types</a>
Default network route	The router uses the default network route if the IP route table does not contain a route to the destination. Enter an explicit default route (0.0.0.0 0.0.0.0 or 0.0.0.0/0) as a static route in the IP route table.	None configured	<a href="#">Configuring the default route</a>

## ARP age timer

The ARP age is the amount of time the switch keeps a MAC address learned through ARP in the ARP cache. The switch resets the timer to zero each time the ARP entry is refreshed and removes the entry if the timer reaches the ARP age. For more information on ARP, see [IP tables and caches](#) on page 40.

You can increase the ARP age timeout maximum to 24 hours or more with this command:

### Syntax:

```
[no] ip arp-age <[1...1440]|infinite>
```

Allows the ARP age to be set from 1 to 1440 minutes (24 hours.)

If the option `infinite` is configured, the internal ARP age timeout is set to 99,999,999 seconds (approximately 3.2 years.) An `arp-age` value of 0 (zero) is stored in the configuration file to indicate that `infinite` has been configured. This value also displays with the `show` commands and in the menu display (**Menu > Switch Configuration > IP Config.**)

Default: 20 minutes

### Setting the ARP age time out to 1000 minutes

```
switch(config)# ip arp-age 1000
```

To view the value of ARP age timer, enter the `show ip` command. The Arp Age time value is shown in bold in the following example.

### The show ip command displaying ARP age

```
switch(config)# show ip
```

```
Internet (IP) Service
```

```
IP Routing : Disabled
```

```
Default Gateway : 15.255.120.1
```

```
Default TTL : 64
```

```
Arp Age : 1000
```

```
Domain Suffix :
```

```
DNS server :
```

```
VLAN | IP Config IP Address Subnet Mask Proxy ARP
```

-----+-----	-----+-----	-----+-----	-----+-----
DEFAULT_VLAN	Manual	15.255.111.13	255.255.248.0 No

You can also view the value of the ARP age timer in the configuration file. The ip arp-age 1000 value is shown in bold in the following example.

### The ip arp-age value in the running config file

```
switch(config)# show running-config

Running configuration:

; J9627A Configuration Editor; Created on release #XX.15.XX
; Ver #01:01:00

hostname "Switch"
savepower led
mirror-port 7
stack commander "TEST_STACK"
stack member 1 mac-address 0024A8D13A40
ip arp-age 100
vlan 1
    name "DEFAULT_VLAN"
    untagged 1-28
    ip address dhcp-bootp
    exit
vlan 222
    name "VLAN222"
    no ip address
    exit
snmp-server community "public" unrestricted
snmp-server host 16.181.51.82 community "public"
```

You can set or display the arp-age value using the menu interface (**Menu > Switch Configuration > IP Config**).

### The Menu interface displaying the ARP Age value

```
===== TELNET - MANAGER MODE =====
Switch Configuration - Internet (IP) Service

IP Routing : Disabled

Default Gateway : 15.255.120.1
Default TTL     : 64
Arp Age       : 1000

IP Config [Manual] : Manual

IP Address  : 15.255.111.11
Subnet Mask : 255.255.248.0

Actions->  Cancel  Edit  Save  Help
```

## IP interface parameters for routing switches

The following table lists the interface-level IP parameters for routing switches.

**Table 5: IP interface parameters — routing switches**

Parameter	Description	Default	More information
IP address	A Layer 3 network interface address; separate IP addresses on individual VLAN interfaces.	None configured	1
Metric	A numeric cost the router adds to RIP routes learned on the interface. This parameter applies only to RIP routes.	1 (one)	<a href="#">Configuring RIP parameters</a> on page 55
ICMP Internet Router Discovery Protocol (IRDP)	Locally overrides the global IRDP settings.	Disabled	<a href="#">Enabling IRDP globally</a> on page 66
IP helper address	The IP address of a UDP application server (such as a BootP or DHCP server) or a directed broadcast address. IP helper addresses allow the routing switch to forward requests for certain UDP applications from a client on one subnet to a server on another subnet.	None configured	<a href="#">Configuring an IP helper address</a> on page 69

<sup>1</sup>See the chapter "Configuring IP Addressing" in the Management and Configuration Guide for your switch.

## Configuring IP parameters for routing switches

The following sections describe how to configure IP parameters. Some parameters can be configured globally while others can be configured on individual VLAN interfaces. Some parameters can be configured globally and overridden for individual VLAN interfaces.



For IP configuration information when routing is not enabled, see the chapter "Configuring IP Addressing" in the Management and Configuration Guide for your routing switch.

### Configuring ARP parameters

ARP is a standard IP protocol that enables an IP routing switch to obtain the MAC address of another device's interface when the routing switch knows the IP address of the interface. ARP is enabled by default and cannot be disabled.

#### How ARP works

A routing switch needs to know a destination's MAC address when forwarding traffic, because the routing switch encapsulates the IP packet in a Layer 2 packet (MAC layer packet) and sends the Layer 2 packet to a MAC interface on a device directly attached to the routing switch. The device can be the packet's final destination or the next-hop router toward the destination.

The routing switch encapsulates IP packets in Layer 2 packets regardless of whether the ultimate destination is locally attached or is multiple router hops away. Since the routing switch's IP route table and IP forwarding cache contain IP address information but not MAC address information, the routing switch cannot forward IP packets based solely on the information in the route table or forwarding cache. The routing switch needs to know the MAC address that corresponds with the IP address of either the packet's locally attached destination or the next-hop router that leads to the destination.

For example, to forward a packet whose destination is multiple router hops away, the routing switch must send the packet to the next-hop router toward its destination, or to a default route or default network route if the IP route table does not contain a route to the packet's destination. In each case, the routing switch must encapsulate the packet and address it to the MAC address of a locally attached device, the next-hop router toward the IP packet's destination.

To obtain the MAC address required for forwarding a datagram, the routing switch does the following:

- First, the routing switch looks in the ARP cache (not the static ARP table) for an entry that lists the MAC address for the IP address. The ARP cache maps IP addresses to MAC addresses. The cache also lists the port attached to the device and, if the entry is dynamic, the age of the entry. A dynamic ARP entry enters the cache when the routing switch receives an ARP reply or receives an ARP request (which contains the sender's IP address and MAC address.) A static entry enters the ARP cache from the static ARP table (which is a separate table) when the interface for the entry comes up. To ensure the accuracy of the ARP cache, each dynamic entry has its own age timer. The timer is reset to zero each time the routing switch receives an ARP reply or ARP request containing the IP address and MAC address of the entry. If a dynamic entry reaches its maximum allowable age, the entry times out and the software removes the entry from the table. Static entries do not age-out and can be removed only by you.
- If the ARP cache does not contain an entry for the destination IP address, the routing switch broadcasts an ARP request out all its IP interfaces. The ARP request contains the IP address of the destination. If the device with the IP address is directly attached to the routing switch, the device sends an ARP response containing its MAC address. The response is a unicast packet addressed directly to the routing switch. The routing switch places the information from the ARP response into the ARP cache. ARP requests contain the IP address and MAC address of the sender, so all devices that receive the request learn the MAC address and IP address of the sender and can update their own ARP caches accordingly. Note that the ARP request broadcast is a MAC broadcast, which means the broadcast goes only to devices that are directly attached to the routing switch. A MAC broadcast is not routed to other networks. However, some routers, including routing switches, can be configured to reply to ARP requests from one network on behalf of devices on another network.



---

If the routing switch receives an ARP request packet that it is unable to deliver to the final destination because of the ARP time-out, and no ARP response is received (the routing switch knows of no route to the destination address), the routing switch sends an ICMP Host Unreachable message to the source.

---

## Configuring forwarding parameters

The following configurable parameters control the forwarding behavior of routing switches:

- Time-To-Live (TTL) threshold The configuration of this parameter is covered in the chapter "Configuring IP Addressing" in the Management and Configuration Guide for your routing switch.
- Forwarding of directed broadcasts

All these parameters are global and thus affect all IP interfaces configured on the routing switch.

To configure these parameters, use the procedures in the following sections.

### Enabling forwarding of directed broadcasts

A directed broadcast is an IP broadcast to all devices within a single directly attached network or subnet. A net-directed broadcast goes to all devices on a given network. A subnet-directed broadcast goes to all devices within a given subnet.



A less common type, the all-subnets broadcast, goes to all directly attached subnets. Forwarding for this broadcast type also is supported, but most networks use IP multicasting instead of all-subnet broadcasting.

Forwarding for all types of IP directed broadcasts is disabled by default. You can enable forwarding for all types if needed. You cannot enable forwarding for specific broadcast types.

## Configuring ICMP

You can configure the following ICMP limits:

Burst-normal	The maximum number of ICMP replies to send per second.
Reply limit	You can enable or disable ICMP reply rate limiting.

## Disabling ICMP messages

Hewlett Packard Enterprise devices are enabled to reply to ICMP echo messages and send ICMP Destination Unreachable messages by default.

You can selectively disable the following types of Internet Control Message Protocol (ICMP) messages:

Echo messages (ping messages)	The routing switch replies to IP pings from other IP devices.
Destination unreachable messages	If the routing switch receives an IP packet that it cannot deliver to its destination, the routing switch discards the packet and sends a message back to the device that sent the packet to the routing switch. The message informs the device that the destination cannot be reached by the routing switch.
Address mask replies	You can enable or disable ICMP address mask replies.

## Disabling replies to broadcast ping requests

By default, HPE devices are enabled to respond to broadcast ICMP echo packets, which are ping requests (for more information, see [Disabling ICMP messages](#) on page 48).

To disable response to broadcast ICMP echo packets (ping requests), enter the following command:

```
switch(config)# no ip icmp echo broadcast-request
```

### Syntax:

```
[no] ip icmp echo broadcast-request
```

If you need to re-enable response to ping requests, enter the following command:

```
switch(config)# ip icmp echo broadcast-request
```

## Disabling ICMP destination unreachable messages

By default, when a device receives an IP packet that the device cannot deliver, the device sends an ICMP unreachable message back to the host that sent the packet. The following types of ICMP unreachable messages are generated:



Administration	The packet was dropped by the device due to a filter or ACL configured on the device.
Fragmentation-needed	The packet has the "Don't Fragment" bit set in the IP Flag field, but the device cannot forward the packet without fragmenting it.
Host	The destination network or subnet of the packet is directly connected to the device, but the host specified in the destination IP address of the packet is not on the network.
Network	The device cannot reach the network specified in the destination IP address of the packet.
Port	The destination host does not have the destination TCP or UDP port specified in the packet. In this case, the host sends the ICMP Port Unreachable message to the device, which in turn sends the message to the host that sent the packet.
Protocol	The TCP or UDP protocol on the destination host is not running. This message is different from the Port Unreachable message, which indicates that the protocol is running on the host but the requested protocol port is unavailable.
Source-route-failure	The device received a source-routed packet but cannot locate the next-hop IP address indicated in the packet's Source-Route option.



Disabling an ICMP Unreachable message type does not change the device's ability to forward packets. Disabling ICMP Unreachable messages prevents the device from generating or forwarding the Unreachable messages.

## Disabling all ICMP unreachable messages

To disable all ICMP Unreachable messages, enter the following command:

```
switch(config)# no ip icmp unreachable
```

### Syntax:

```
[no] ip icmp unreachable
```

For more information, see [Disabling ICMP destination unreachable messages](#) on page 48.

## Disabling ICMP redirects

You can disable ICMP redirects on the routing switch only on a global basis, for all the routing switch interfaces.

To disable ICMP redirects globally, enter the following command at the global CONFIG level of the CLI:

```
switch(config)# no ip icmp redirects
```

### Syntax:

```
no ip icmp redirects
```

## Configuring static IP routes

This feature enables you to create static routes (and null routes) by adding such routes directly to the route table. This section describes how to add static and null routes to the IP route table.

## Static route types

You can configure the following types of static IP routes:

Standard	The static route consists of a destination network address or host, a corresponding network mask, and the IP address of the next-hop IP address.
Null (discard)	The null route consists of the destination network address or host, a corresponding network mask, and either the <code>reject</code> or <code>blackhole</code> keyword. Typically, the null route is configured as a backup route for discarding traffic if the primary route is unavailable. By default, when IP routing is enabled, a route for the 127.0.0.0/8 network is created to the null interface. Traffic to this interface is rejected (dropped). This route is for all traffic to the "loopback" network, with the single exception of traffic to the host address of the switch's loopback interface (127.0.0.1/32.) <b>Figure 5: Displaying the currently configured static routes</b> on page 53 shows the default null route entry in the switch's routing table.



On a single routing switch you can create one null route to a given destination. Multiple null routes to the same destination are not supported.

## Other sources of routes in the routing table

The IP route table can also receive routes from the following sources:

- Directly connected networks: One route is created per IP interface. When you add an IP interface, the routing switch automatically creates a route for the network the interface is in.
- RIP: If RIP is enabled, the routing switch can learn about routes from the advertisements other RIP routers send to the routing switch. If the RIP route has a lower administrative distance than any other routes from different sources to the same destination, the routing switch places the route in the IP route table. See **Administrative distance** on page 40.
- Default route: This is a specific static route that the routing switch uses if other routes to the destination are not available. See **Configuring the default route** on page 53.

## Static IP route parameters

When you configure a static IP route, you must specify the following parameters:

- The IP address and network mask for the route's destination network or host.
- The route's path, which can be one of the following:
  - IP address of a next-hop router.
  - "Null" interface; the routing switch drops traffic forwarded to the null interface.

The routing switch also applies default values for the route's administrative distance (**Administrative distance** on page 40). In the case of static routes, this is the value the routing switch uses to compare a static route to routes from other route sources to the same destination before placing a route in the IP route table.

The default administrative distance for static IP routes is 1, but can be configured to any value from 1 to 255.

The fixed administrative distance values ensure that the routing switch always prefers static IP routes over routes from other sources to the same destination.

## Static route states follow VLAN states

IP static routes remain in the IP route table only so long as the IP interface to the next-hop router is up. If the next-hop interface goes down, the software removes the static route from the IP route table. If the next-hop interface comes up again, the software adds the route back to the route table.

This feature allows the routing switch to adjust to changes in network topology.

The routing switch does not continue trying to use routes on unreachable paths, but instead uses routes only when their paths are reachable.

For example, the following command configures a static route to 207.95.7.0 (with a network mask of 255.255.255.0), using 207.95.6.157 as the next-hop router's IP address:

```
switch(config)# ip route 207.95.7.0/24 207.95.6.15
```

A static IP route specifies the route's destination address and the next-hop router's IP address or routing switch interface through which the routing switch can reach the destination. (The route is added to the routing switch's IP route table.)

In the above example, routing switch "A" knows that 207.95.6.157 is reachable through port A2, and assumes that local interfaces within that subnet are on the same port. Routing switch "A" deduces that IP interface 207.95.7.188 is also on port A2. The software automatically removes a static route from the route table if the next-hop VLAN used by that route becomes unavailable. When the VLAN becomes available again, the software automatically readds the route to the route table.

## Configuring a static IP route

Static route	Configure a static route to a specific network or host address
Null route	Configure a "null" route to discard IP traffic to a specific network or host address: <ul style="list-style-type: none"><li>• Discard traffic for the destination, with ICMP notification to sender</li><li>• Discard traffic for the destination, without ICMP notification to sender</li></ul>

### Syntax:

```
ip route <dest-ip-addr>/<mask-length> <next-hop-ip-addr|vlan <vlan-id>|reject|blackhole> [metric <metric>] [distance <1-255>] [tag-value <tagval>]
```

```
no ip route <dest-ip-addr>/<mask-length> <next-hop-ip-addr|vlan <vlan-id>|reject|blackhole> [metric <metric>] [distance <1-255>] [tag-value <tagval>]
```

Allows the addition and deletion of static routing table entries. A route entry is identified by a destination (IP address/mask length) and next-hop pair. The next-hop can be either a gateway IP address, a VLAN, or the keyword "reject" or "blackhole".

A gateway IP address does not have to be directly reachable on one of the local subnets. If the gateway address is not directly reachable, the route is added to the routing table as soon as a route to the gateway address is learned.

<dest-ip-addr>/<mask-bits>	The route destination and network mask length for the destination IP address. Alternatively, you can enter the mask itself. For example, you can enter either 10.0.0.0/24 or 10.0.0.0 255.255.255.0 for a route destination of 10.0.0.0 255.255.255.0.
next-hop-ip-addr	This IP address is the gateway for reaching the destination. The next-hop IP address is not required to be directly reachable on a local subnet. (If the next-hop IP address is not directly reachable, the route will be added to the routing table as soon as a route to this address is learned.)

Table Continued

reject	Specifies a null route where IP traffic for the specified destination is discarded and an ICMP error notification is returned to the sender.
blackhole	Specifies a null route where IP traffic for the specified destination is discarded and no ICMP error notification is returned to the sender.
metric	Specifies an integer value that is associated with the route. It is used to compare a static route to routes in the IP route table from other sources to the same destination.
distance	Specifies the administrative distance to associate with a static route. If not specified, this value is set to a default of 1. (Range: 1 to 255)
tag	Specifies a unique integer value for a given ECMP set (destination, metric, distance.)

The `no` form of the command deletes the specified route for the specified destination next-hop pair.

The following example configures two static routes for traffic delivery and identifies two other null routes for which traffic should be discarded instead of forwarded.

### Configuring static routes

```
switch(config)# ip route 10.10.40.0/24 10.10.10.1 1
switch(config)# ip route 10.10.50.128/27 10.10.10.1
switch(config)# ip route 10.10.20.177/32 reject2
switch(config)# ip route 10.10.30.0/24 blackhole3
```

<sup>1</sup> Configures static routes to two different network destinations using the same next-hop router IP address.

<sup>2</sup> Configures a null route to drop traffic for the device at 10.50.10.177 and return an ICMP notification to the sender.

<sup>3</sup> Configures a null route to drop traffic for the 10.50.10.0 network without any ICMP notification to the sender.

## Viewing static route information

The `show ip route` command displays the current static route configuration on the routing switch. The following figure shows the configuration resulting from the static routes configured in the previous examples.

**Figure 5:** *Displaying the currently configured static routes*

```
Switch(config)# show ip route static
```

IP Route Entries						
Destination	Gateway	VLAN	Type	Sub-Type	Metric	Dist.
10.50.10.177/32	reject		static		1	1
10.10.40.0/24	VLAN10	10	static		1	1
10.10.50.128/27	VLAN10	10	static		1	1
10.50.10.0/24	blackhole		static		1	1
127.0.0.0/8	reject		static		0	0
127.10.144.32/24	10.0.0.2	1	static		12	10
127.10.144.32/24	10.0.0.3	1	static		12	10

This reject (default null) route is included by default. Refer to "Configuring a static route" on page 1-1

An ECMP set with `ip load-sharing` set to 2 (the maximum paths allowed)

## Configuring the default route

You can also assign the default route and enter it in the routing table. The default route is used for all traffic that has a destination network not reachable through any other IP routing table entry. For example, if 208.45.228.35 is the IP address of your ISP router, all nonlocal traffic could be directed to the ISP by entering this command:

```
switch(config)# ip route 0.0.0.0/0 208.45.228.35
```

## Configuring RIP

This section describes how to configure RIP using the CLI interface.

To display RIP configuration information and statistics, see [Overview of RIP](#) on page 53. For more information on configuring RIP, see [Viewing RIP information](#) on page 60.

## Overview of RIP

Routing Information Protocol (RIP) is an IP route exchange protocol that uses a **distance vector** (a number representing distance) to measure the cost of a given route. The cost is a distance vector because the cost often is equivalent to the number of router hops between the routing switch and the destination network.

An routing switch can receive multiple paths to a destination. The software evaluates the paths, selects the best path, and saves the path in the IP route table as the route to the destination. Typically, the best path is the path with the fewest hops. A hop is another router through which packets must travel to reach the destination. If the routing switch receives an RIP update from another router that contains a path with fewer hops than the path stored in the routing switch's route table, the routing switch replaces the older route with the newer one. The routing switch then includes the new path in the updates it sends to other RIP routers, including routing switches.

RIP routers, including routing switches, also can modify a route's cost, generally by adding to it, to bias the selection of a route for a given destination. In this case, the actual number of router hops may be the same, but the route has an administratively higher cost and is thus less likely to be used than other, lower-cost routes. A RIP route can have a maximum cost of 15. Any destination with a higher cost is considered unreachable. Although limiting to larger networks, the low maximum hop count prevents endless loops in the network.

The switches support the following RIP types:

- Version 1
- V1 compatible with V2
- Version 2 (the default)



If the routing switch receives an ARP request packet that it is unable to deliver to the final destination because of the ARP timeout and no ARP response is received (the routing switch knows of no route to the destination address), the routing switch sends an ICMP Host Unreachable message to the source.

## RIP parameters and defaults

The following tables list the RIP parameters, their default values, and where to find configuration information.

### RIP global parameters

The following table lists the global RIP parameters and their default values.

**Table 6:** *RIP global parameters*

Parameter	Description	Default
<b>RIP state</b>	Routing Information Protocol V2-only.	Disabled
<b>auto-summary</b>	Enable/disable advertisement of summarized routes.	Enabled
<b>metric</b>	Default metric for imported routes.	1
<b>redistribution</b>	RIP can redistribute static, and connected routes. (RIP redistributes connected routes by default, when RIP is enabled.)	Disabled

### RIP interface parameters

The following table lists the VLAN interface RIP parameters and their default values.

**Table 7: RIP interface parameters**

Parameter	Description	Default
<b>RIP version</b>	The version of the protocol that is supported on the interface. The version can be one of the following: <ul style="list-style-type: none"><li>• Version 1 only</li><li>• Version 2 only</li><li>• Version 1 or version 2</li></ul>	V2-only
<b>metric</b>	A numeric cost the routing switch adds to RIP routes learned on the interface. This parameter applies only to RIP routes.	1
<b>IP address</b>	The routes that a routing switch learns or advertises can be controlled.	The routing switch learns and advertises all RIP routes on all RIP interfaces
<b>loop prevention</b>	The method the routing switch uses to prevent routing loops caused by advertising a route on the same interface as the one on which the routing switch learned the route: <ul style="list-style-type: none"><li>• <b>Split horizon</b><ul style="list-style-type: none"><li>— The routing switch does not advertise a route on the same interface as the one on which the routing switch learned the route.</li></ul></li><li>• <b>Poison reverse</b><ul style="list-style-type: none"><li>— The routing switch assigns a cost of 16 "infinite" or "unreachable") to a route before advertising it on the same interface as the one on which the routing switch learned the route.</li></ul></li></ul>	Poison reverse
<b>receive</b>	Define the RIP version for incoming packets	V2-only
<b>send</b>	Define the RIP version for outgoing packets	V2-only

## Configuring RIP parameters

Use the following procedures to configure RIP parameters on a systemwide and individual VLAN interface basis.

### Enabling RIP

RIP is disabled by default. To enable it, use one of the following methods. When you enable RIP, the default RIP version is RIPv2-only. You can change the RIP version on an individual interface basis to RIPv1 or RIPv1-or-v2, if needed.

To enable RIP on a routing switch, enter the following commands:

```
switch(config)# ip routing
switch(config)# router rip
switch(rip)# exit
switch(config)# write memory
```



IP routing must be enabled before enabling RIP. The first command in the preceding sequence enables IP routing.

## Enabling RIP on the routing switch and entering the RIP router context

### Syntax:

```
[no] router rip [[enable] | [disable]] [auto-summary]
```

Executed at the global configuration level to enable RIP on the routing switch and to enter the RIP router context. This enables you to proceed with assigning RIP areas and to modify RIP global parameter settings as needed. Global IP routing must be enabled before the RIP protocol can be enabled.

Default: Disabled

enable	Enables RIP routing.
disable	Disables RIP routing.

The `no` form of the command deletes all protocol-specific information from the global context and interface context. All protocol parameters are set to default values.



If you disable RIP, the switch retains all the configuration information for the disabled protocol in flash memory. If you subsequently restart RIP, the existing configuration will be applied.

The `auto-summary` form of the command enables advertisement of the summarized routes. When used with the `no` form of the command, `auto-summary` disables the advertisement of the summarized routes.

### Enter RIP router context

```
switch(config)# router rip
switch(rip)#
```

### Enable RIP routing

```
switch(config)# router rip enable
switch(rip)#
```

### Disable RIP routing

```
switch(config)# router rip disable
switch(rip)#
```



## Delete all protocol-specific information from the global and interface context

```
switch(config)# no router rip
switch(rip)#
```



Deleting all protocol-specific information from the global and interface context sets all protocol parameters to default values.



IP routing must be enabled before enabling RIP. The first command in the preceding sequence enables IP routing.

## Enabling IP RIP on a VLAN

To enable RIP on all IP addresses in a VLAN, use `ip rip` in the VLAN context. When the command is entered without specifying any IP address, it is enabled in all configured IP addresses of the VLAN.

To enable RIP on a specific IP address in a VLAN, use `ip rip [<ip-addr>|all]` in the VLAN context and enter a specific IP address. If you want RIP enabled on all IP addresses, you can specify `all` in the command instead of a specific IP address.

## Changing the RIP type on a VLAN interface

When you enable RIP on a VLAN interface, `RIPv2-only` is enabled by default. You can change the RIP type to one of the following on an individual VLAN interface basis:

- Version 1 only
- Version 2 only (the default)
- Version 1 - or - version 2

To change the RIP type supported on a VLAN interface, enter commands such as the following:

```
switch(config)# vlan 1
switch(vlan-1)# ip rip v1-only
switch(vlan-1)# exit
switch(config)# write memory
```

### Syntax:

```
ip rip {<v1-only | v1-or-v2 | v2-only>}
no ip rip {<v1-only | v1-or-v2 | v2-only>}
```

## Changing the cost of routes learned on a VLAN interface

By default, the switch interface increases the cost of an RIP route that is learned on the interface. The switch increases the cost by adding one to the route's metric before storing the route.

You can change the amount that an individual VLAN interface adds to the metric of RIP routes learned on the interface.



RIP considers a route with a metric of 16 to be unreachable. Use this metric only if you do not want the route to be used. In fact, you can prevent the switch from using a specific interface for routes learned though that interface by setting its metric to 16.

To increase the cost a VLAN interface adds to RIP routes learned on that interface, enter commands such as the following:

```
switch(config)# vlan 1
switch(vlan-1)# ip rip metric 5
```

These commands configure vlan-1 to add 5 to the cost of each route learned on the interface.

**Syntax:**

```
ip rip metric <1-16>
```

## Configuring RIP redistribution

You can configure the routing switch to redistribute connected, static, and OSPF routes into RIP. When you redistribute a route into RIP, the routing switch can use RIP to advertise the route to its RIP neighbors.

To configure redistribution, perform the following tasks:

**Procedure**

1. Configure redistribution filters to permit or deny redistribution for a route based on the destination network address or interface. (optional)
2. Enable redistribution.

## Defining RIP redistribution filters

Route redistribution imports and translates different protocol routes into a specified protocol type. On the switches covered in this guide, redistribution is supported for static and directly connected routes. Redistribution of any other routing protocol into RIP is not currently supported. When you configure redistribution for RIP, you can specify that static or connected routes are imported into RIP routes.

## Configuring for redistribution

To configure for redistribution, define the redistribution tables with "restrict" redistribution filters. In the CLI, use the `restrict` command for RIP at the RIP router level.



Do not enable redistribution until you have configured the redistribution filters. Otherwise, the network might become overloaded with routes that you did not intend to redistribute.

**Example:**

To configure the switch to filter out redistribution of static or connected routes on network 10.0.0.0, enter the following commands:

```
switch(config)# router rip
switch(rip)# restrict 10.0.0.0 255.0.0.0
switch(rip)# write memory
```

The default configuration permits redistribution for all default connected routes only.

---

**Syntax:**

```
restrict {< ip-addr > < ip-mask > | < ip-addr/ < prefix length >}
```

This command prevents any routes with a destination address that is included in the range specified by the address/mask pair from being redistributed by RIP.

## Modifying default metric for redistribution

The default metric is a global parameter that specifies the cost applied to all RIP routes by default. The default value is 1. You can assign a cost from 1 to 15.

### Example:

To assign a default metric of 4 to all routes imported into RIP, enter the following commands:

```
switch(config)# router rip
switch(rip)# default-metric 4
```

### Syntax:

```
default-metric <value>
```

The <value> can be from 1 to 15. The default is 1.

## Enabling RIP route redistribution

The basic form of the `redistribute` command redistributes all routes of the selected type. For finer control over route selection and modification of route properties, you can specify the `route-map` parameter and the name of a route map.



Do not enable redistribution until you have configured the redistribution filters. Otherwise, the network might become overloaded with routes that you did not intend to redistribute.

### Syntax:

```
[no] router rip redistribute {<connected | static>} [route-map < name >]
```

Enables redistribution of the specified route type to the RIP domain.

<code>static</code>	Redistribute from manually configured routes.
<code>connected</code>	Redistribute from locally connected networks.
<code>route-map &lt;name&gt;</code>	Optionally specify the name of a route map to apply during redistribution.

The `no` form of the command disables redistribution for the specified route type.

### Example:

To enable redistribution of all connected, static, and OSPF routes into RIP, enter the following commands.

```
switch(config)# router rip
switch(rip)# redistribute connected
switch(rip)# redistribute static
switch(rip)# write memory
```

## Changing the route loop prevention method

### Syntax:

```
[no] ip rip poison-reverse
```

Entering the command without the `no` option will re-enable Poison reverse.

RIP can use the following methods to prevent routing loops:

- **Split horizon**

- the routing switch does not advertise a route on the same interface as the one on which the routing switch learned the route.

- **Poison reverse**

- the routing switch assigns a cost of 16 (“infinity” or “unreachable”) to a route before advertising it on the same interface as the one on which the routing switch learned the route. This is the default.

These loop prevention methods are configurable on an individual VLAN interface basis.



These methods are in addition to RIP's maximum valid route cost of 15.

Poison reverse is enabled by default. Disabling Poison reverse causes the routing switch to revert to Split horizon. (Poison reverse is an extension of Split horizon.) To disable Poison reverse on an interface, and thereby enable Split horizon, enter the following:

```
switch(config)# vlan 1
switch(vlan-1)# no ip rip poison-reverse
```

## Viewing RIP information

All RIP configuration and status information is shown by the CLI command `show ip rip` and options off that command.

### Viewing general RIP information

**Syntax:**

```
show ip rip
```

To display general RIP information, enter `show ip rip` at any context level. The resulting display will appear similar to the following:

#### General RIP information listing

```
switch(config)# show ip rip
```

```
RIP global parameters
```

```
RIP protocol      : enabled
Auto-summary     : enabled
Default Metric   : 1
Distance         : 120
Route changes    : 0
Queries          : 0
```

```
RIP interface information
```

IP Address	Status	Send mode	Recv mode	Metric	Auth
100.1.0.1	enabled	V2-only	V2-only	5	none
100.2.0.1	enabled	V2-only	V2-only	5	none
100.3.0.1	enabled	V2-only	V2-only	5	none
100.4.0.1	enabled	V2-only	V2-only	5	none
100.10.0.1	enabled	V2-only	V2-only	5	none
100.11.0.1	enabled	V2-only	V2-only	5	none

## RIP peer information

```
IP Address      Bad routes  Last update timeticks  
-----
```

The display is a summary of global RIP information, information about interfaces with RIP enabled, and information about RIP peers. The following fields are displayed:

RIP protocol	Status of the RIP protocol on the router. RIP must be enabled here and on the VLAN interface for RIP to be active. The default is disabled.
Auto-summary	Status of auto-summary for all interfaces running RIP. If auto-summary is enabled, subnets will be summarized to a class network when advertising outside of the given network.
Default metric	Sets the default metric for imported routes. This is the metric that will be advertised with the imported route to other RIP peers. A RIP metric is a measurement used to determine the "best" path to network: 1 is the best, 15 is the worst, 16 is unreachable.
Route changes	The number of times RIP has modified the routing switch's routing table.
Queries	The number of RIP queries that have been received by the routing switch.

*Table Continued*

RIP interface information	<p>RIP information on the VLAN interfaces on which RIP is enabled:</p> <p><b>IP address</b></p> <p>IP address of the VLAN interface running RIP.</p> <p><b>Status</b></p> <p>Status of RIP on the VLAN interface.</p> <p><b>Send mode</b></p> <p>Format of the RIP updates: RIP 1, RIP 2, or RIP 2 version 1 compatible.</p> <p><b>Recv mode</b></p> <p>The switch can process RIP 1, RIP 2, or RIP 2 version 1 compatible update messages.</p> <p><b>Metric</b></p> <p>Path "cost", a measurement used to determine the "best" RIP route path: 1 is the best, 15 is the worst, 16 is unreachable.</p> <p><b>Auth</b></p> <p>RIP messages can be required to include an authentication key if enabled on the interface.</p>
RIP peer information	<p>RIP peers are neighboring routers from which the routing switch has received RIP updates:</p> <p><b>IP address</b></p> <p>IP address of the RIP neighbor.</p> <p><b>Bad routes</b></p> <p>Number of route entries which were not processed for any reason.</p> <p><b>Last update timeticks</b></p> <p>Number of seconds that have passed since we received an update from this neighbor.</p>

## Viewing RIP interface information

To display RIP interface information, enter the `show ip rip interface` command at any context level.

### Syntax:

```
show ip rip interface [ip-addr | vlan < vlan-id >]
```

The resulting display will appear similar to the following:

### Output for the show IP RIP interface command

```
switch(config)# show ip rip interface
```

```
RIP interface information
```

IP Address	Status	Send mode	Recv mode	Metric	Auth
100.1.0.1	enabled	V2-only	V2-only	1	none
100.2.0.1	enabled	V2-only	V2-only	1	none

100.3.0.1	enabled	V2-only	V2-only	1	none
100.4.0.1	enabled	V2-only	V2-only	1	none

You can also display the information for a single RIP VLAN interface, by specifying the VLAN ID for the interface, or by specifying the IP address for the interface.

To show the RIP interface information for VLAN 1000, use the `show ip rip interface vlan <vid>` command.

### RIP interface output by VLAN

```
switch# show ip rip interface vlan 4

RIP configuration and statistics for VLAN 4

RIP interface information for 100.4.0.1

  IP Address : 100.4.0.1
  Status      : enabled

  Send Mode   : V2-only
  Recv mode   : V2-only
  Metric      : 1
  Auth        : none

  Bad packets received : 0
  Bad routes received  : 0
  Sent updates         : 0
```

For RIP interface output by VLAN field definitions, see [Viewing general RIP information](#) on page 60.

The RIP interface information also includes the following fields:

Bad packets received	Number of packets that were received on this interface and were not processed for any reason.
Bad routes received	Number of route entries that were received on this interface and were not processed for any reason.
Sent updates	Number of RIP routing updates that have been sent on this interface.

To show the RIP interface information for the interface with IP address 100.2.0.1, enter the `show ip rip interface` command:

### The show IP rip interface output by IP address

```
switch# show ip rip interface 100.2.0.1

RIP interface information for 100.2.0.1

  IP Address : 100.2.0.1
  Status      : enabled

  Send Mode   : V2-only
  Recv mode   : V2-only
  Metric      : 1
  Auth        : none

  Bad packets received : 0
```

```
Bad routes received : 0
Sent updates : 0
```

## Viewing RIP peer information

To display RIP peer information, enter the `show ip rip peer` command at any context level.

The resulting display will appear similar to the following:

### Output for the show IP rip peer command

```
switch# show ip rip peer
```

```
RIP peer information
```

IP Address	Bad routes	Last update timeticks
100.1.0.100	0	1
100.2.0.100	0	0
100.3.0.100	0	2
100.10.0.100	0	1

This display lists all neighboring routers from which the routing switch has received RIP updates. The following fields are displayed:

IP address	IP address of the RIP peer neighbor.
Bad routes	The number of route entries that were not processed for any reason.
Last update timeticks	How many seconds have passed since the routing switch received an update from this peer neighbor.

To show the RIP peer information for a specific peer with IP address 100.1.0.100, enter `show ip rip peer 100.1.0.100`.

### Output for the show IP rip peer <ip-addr> command

```
switch# show ip rip peer 100.0.1.100
```

```
RIP peer information for 100.0.1.100
```

```
IP Address : 100.1.0.100
Bad routes : 0
Last update timeticks : 2
```

This display lists information in the fields described above (IP address, Bad routes, Last update timeticks.)

## Viewing RIP redistribution information

To display RIP redistribution information, enter the `show ip rip redistribute` command at any context level:

### Output for the show IP rip redistribute command

```
switch# show ip rip redistribute
```

```
RIP redistributing
```



```
Route type Status
-----
connected enabled
static disabled
```

RIP automatically redistributes connected routes that are configured on interfaces that are running RIP and all routes that are learned via RIP. The `router rip redistribute` command, described in [Configuring for redistribution](#) on page 58, configures the routing switch to cause RIP to advertise connected routes that are not running RIP or static routes. The display shows whether RIP redistribution is enabled or disabled for connected or static routes.

## Viewing RIP redistribution filter (restrict) information

To display RIP restrict filter information, enter the `show ip rip restrict` command at any context level:

### Output for the show IP rip restrict command

```
switch# show ip rip restrict
RIP restrict list
IP Address      Mask
-----
```

The display shows if any routes identified by the IP Address and Mask fields are being restricted from redistribution. The restrict filters are configured by the `router rip restrict` command (see [Configuring for redistribution](#) on page 58).

## Configuring IRDP

The ICMP Internet Router Discovery Protocol (IRDP) is used by routing switches to advertise the IP addresses of their router interfaces to directly attached hosts. IRDP is disabled by default. You can enable the feature on a global basis or on an individual VLAN interface basis.

When IRDP is enabled, the routing switch periodically sends Router Advertisement messages out the IP interfaces on which the feature is enabled. The messages advertise the routing switch's IP addresses to directly attached hosts who listen for the messages. In addition, hosts can be configured to query the routing switch for the information by sending Router Solicitation messages.

Some types of hosts use the Router Solicitation messages to discover their default gateway. When IRDP is enabled on the routing switch, the routing switch responds to the Router Solicitation messages. Some clients interpret this response to mean that the routing switch is the default gateway. If another router is actually the default gateway for these clients, leave IRDP disabled on the routing switch.

IRDP uses the following parameters. If you enable IRDP on individual VLAN interfaces, you can configure these parameters on an individual VLAN interface basis.

Packet type	The routing switch can send Router Advertisement messages as IP broadcasts or as IP multicasts addressed to IP multicast group 224.0.0.1. The default packet type is IP broadcast.
Hold time	Each Router Advertisement message contains a hold time value. This value specifies the maximum amount of time the host should consider an advertisement to be valid until a newer advertisement arrives. When a new advertisement arrives, the hold time is reset. The hold time is always longer than the maximum advertisement interval. Therefore, if the hold time for an advertisement expires, the host can reasonably conclude that the router interface that sent the advertisement is no longer available. The default hold time is three times the maximum message interval.
Maximum message interval and minimum message interval	When IRDP is enabled, the routing switch sends the Router Advertisement messages every 450-600 seconds by default. The time within this interval that the routing switch selects is random for each message and is not affected by traffic loads or other network factors. The random interval minimizes the probability that a host will receive Router Advertisement messages from other routers at the same time. The interval on each IRDP-enabled routing switch interface is independent of the interval on other IRDP-enabled interfaces. The default maximum message interval is 600 seconds. The default minimum message interval is 450 seconds.
Preference	If a host receives multiple Router Advertisement messages from different routers, the host selects the router that send the message with the highest preference as the default gateway. The preference can be a number from -4294967296 to 4294967295. The default is 0.

## Enabling IRDP globally

Enter the following command:

```
switch(config)# ip irdp
```

This command enables IRDP on the IP interfaces on all ports. Each port uses the default values for the IRDP parameters.

## Enabling IRDP on an individual VLAN interface

To enable IRDP on an individual VLAN interface and configure IRDP parameters, enter commands such as the following:

```
switch(config)# vlan 1
switch(vlan-1)# ip irdp maxadvertinterval 400
```

This example shows how to enable IRDP on a specific interface (VLAN 1) and change the maximum advertisement interval for Router Advertisement messages to 400 seconds.

### Syntax:

```
[no] ip irdp [broadcast | multicast] [holdtime < seconds >] [maxadvertinterval < seconds >] [minadvertinterval < seconds >] [preference < number >]
```

[broadcast   multicast]	<p>Specifies the packet type the routing switch uses to send the Router Advertisement:</p> <p><b>broadcast</b></p> <p>The routing switch sends Router Advertisements as IP broadcasts.</p> <p><b>multicast</b></p> <p>The routing switch sends Router Advertisements as multicast packets addressed to IP multicast group 224.0.0.1. This is the default.</p>
holdtime <seconds>	<p>Specifies how long a host that receives a Router Advertisement from the routing switch should consider the advertisement to be valid. When a host receives a new Router Advertisement message from the routing switch, the host resets the hold time for the routing switch to the hold time specified in the new advertisement. If the hold time of an advertisement expires, the host discards the advertisement, concluding that the router interface that sent the advertisement is no longer available. The value must be greater than the value of the maxadvertinterval parameter and cannot be greater than 9000. The default is three times the value of the maxadvertinterval parameter.</p>
maxadvertinterval	<p>Specifies the maximum amount of time the routing switch waits between sending Router Advertisements. You can specify a value from 1 to the current value of the holdtime parameter. The default is 600 seconds.</p>
minadvertinterval	<p>Specifies the minimum amount of time the routing switch can wait between sending Router Advertisements. The default is three-fourths (0.75) the value of the maxadvertinterval parameter. If you change the maxadvertinterval parameter, the software automatically adjusts the minadvertinterval parameter to be three-fourths the new value of the maxadvertinterval parameter. If you want to override the automatically configured value, you can specify an interval from 1 to the current value of the maxadvertinterval parameter.</p>
preference <number>	<p>Specifies the IRDP preference level of this routing switch. If a host receives Router Advertisements from multiple routers, the host selects the router interface that sent the message with the highest preference as the host's default gateway. The valid range is -4294967296 to 4294967295. The default is 0.</p>

## Viewing IRDP information

To display IRDP information, enter `show ip irdp` from any CLI level.

### Output for show ip irdp

```
switch# show ip irdp
Status and Counters - ICMP Router Discovery Protocol

Global Status : Disabled

VLAN Name      Status   Advertising Min int Max int Holdtime Preference
              Address (sec)  (sec)      (sec)

```

DEFAULT_VLAN	Enabled	multicast	450	600	1800	0
VLAN20	Enabled	multicast	450	600	1800	0
VLAN30	Enabled	multicast	450	600	1800	0

## Configuring DHCP relay

### Overview

The Dynamic Host Configuration Protocol (DHCP) is used for configuring hosts with IP address and other configuration parameters without user intervention. The protocol is composed of three components:

- DHCP client
- DHCP server
- DHCP relay agent

The DHCP client sends broadcast request packets to the network; the DHCP servers respond with broadcast packets that offer IP parameters, such as an IP address for the client. After the client chooses the IP parameters, communication between the client and server is by unicast packets.

HPE routing switches provide the DHCP relay agent to enable communication from a DHCP server to DHCP clients on subnets other than the one the server resides on. The DHCP relay agent transfers DHCP messages from DHCP clients located on a subnet without a DHCP server to other subnets. It also relays answers from DHCP servers to DHCP clients.

The DHCP relay agent is transparent to both the client and the server. Neither side is aware of the communications that pass through the DHCP relay agent. As DHCP clients broadcast requests, the DHCP relay agent receives the packets and forwards them to the DHCP server. During this process, the DHCP relay agent increases the hop count by one before forwarding the DHCP message to the server. A DHCP server includes the hop count from the DHCP request that it receives in the response that it returns to the client.

### DHCP packet forwarding

The DHCP relay agent on the routing switch forwards DHCP client packets to all DHCP servers that are configured in the table administrated for each VLAN.

#### Unicast forwarding

The packets are forwarded using unicast forwarding if the IP address of the DHCP server is a specific host address. The DHCP relay agent sets the destination IP address of the packet to the IP address of the DHCP server and forwards the message.

#### Broadcast forwarding

The packets are forwarded using broadcast forwarding if the IP address of the DHCP server is a subnet address or IP broadcast address (255.255.255.255.) The DHCP relay agent sets the DHCP server IP address to broadcast IP address and is forwarded to all VLANs with configured IP interfaces (except the source VLAN.)

### Prerequisites for DHCP relay operation

For the DHCP relay agent to work on the switch, you must complete the following steps:

#### Procedure

1. Enable DHCP relay on the routing switch (the default setting.)
2. Ensure that a DHCP server is servicing the routing switch.
3. Enable IP routing on the routing switch.

4. Ensure that there is a route from the DHCP server to the routing switch and back.
5. Configure one or more IP helper addresses for specified VLANs to forward DHCP requests to DHCP servers on other subnets.

## Enabling DHCP relay

The DHCP relay function is enabled by default on a routing switch. However, if DHCP has been disabled, you can re-enable it by entering the following command at the global configuration level:

```
switch(config)# dhcp-relay
```

To disable the DHCP relay function, enter the `no` form of the command:

```
switch(config)# no dhcp-relay
```

## Configuring an IP helper address

To add the IP address of a DHCP server for a specified VLAN on a routing switch, enter the `ip helper-address` command at the VLAN configuration level as in the following example:

```
switch(config)# vlan 1
switch(vlan-1)# ip helper-address <ip-addr>
```

To remove the DHCP server helper address, enter the `no` form of the command:

```
switch(vlan-1)# no ip helper-address <ip-addr>
```

## Operating notes

- You can configure up to 4000 IP helper addresses on a routing switch. The helper addresses are shared between the DHCP relay agent and UDP forwarder feature.
- A maximum of 16 IP helper addresses is supported in each VLAN.

## Verifying the DHCP relay configuration

### Viewing the DHCP relay setting

Use the `show config` command (or `show running` for the running-config file) to display the current DHCP relay setting.



The DHCP relay and hop count increment settings appear in the `show config` command output only if the nondefault values are configured.

### Displaying startup configuration with DHCP relay disabled

```
switch# show config
Startup configuration:
; J9726A Configuration Editor; Created on release #xx.15.xx
hostname "switch"
cdp run
module 1 type J9726A
ip default-gateway 18.30.240.1
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged A1
```

```
ip address 18.30.240.180 255.255.248.0
no untagged A2-A24
exit
no dhcp-relay
```

- Non-Default DHCP Relay setting

## Viewing DHCP helper addresses

To display the list of currently configured IP Helper addresses for a specified VLAN on the switch, enter the `show ip helper-address vlan` command.

### Syntax:

```
show ip helper-address [vlan < vlan-id >]
```

Displays the IP helper addresses of DHCP servers configured for all static VLANs in the switch or on a specified VLAN, regardless of whether the DHCP relay feature is enabled. The `vlan <vlan-id>` parameter specifies a VLAN ID number.

The following command lists the currently configured IP Helper addresses for VLAN 1.

### Displaying IP helper addresses

```
switch(config)# show ip helper-address vlan 1

IP Helper Addresses

IP Helper Address
-----
10.28.227.97
10.29.227.53
```

## DHCP Option 82

Option 82 is called the relay agent information option and is inserted by the DHCP relay agent when forwarding client-originated DHCP packets to a DHCP server. Servers recognizing the relay agent information option may use the information to implement IP address or other parameter assignment policies. The DHCP server echoes the option back verbatim to the relay agent in server-to-client replies, and the relay agent strips the option before forwarding the reply to the client.

The relay agent information option is organized as a single DHCP option that contains one or more suboptions that convey information known by the relay agent. The initial suboptions are defined for a relay agent that is colocated in a public circuit access unit. These include a circuit ID for the incoming circuit and a remote ID that provides a trusted identifier for the remote high-speed modem.

The routing switch can operate as a DHCP relay agent to enable communication between a client and a DHCP server on a different subnet. Without Option 82, DHCP operation modifies client IP address request packets to the extent needed to forward the packets to a DHCP server. Option 82 enhances this operation by enabling the routing switch to append an Option 82 field to such client requests. This field includes two suboptions for identifying the routing switch (by MAC address or IP address) and the routing switch port the client is using to access the network. A DHCP server with Option 82 capability can read the appended field and use this data as criteria for selecting the IP addressing it will return to the client through the usual DHCP server response packet. This operation provides several advantages over DHCP without Option 82:

- An Option 82 DHCP server can use a relay agent's identity and client source port information to administer IP addressing policies based on client and relay agent location within the network, regardless of whether the relay agent is the client's primary relay agent or a secondary agent.
- A routing switch operating as a primary Option 82 relay agent for DHCP clients requesting an IP address can enhance network access protection by blocking attempts to use an invalid Option 82 field to imitate an

authorized client, or by blocking attempts to use response packets with missing or invalid Option 82 suboptions to imitate valid response packets from an authorized DHCP server.

- An Option 82 relay agent can also eliminate unnecessary broadcast traffic by forwarding an Option 82 DHCP server response only to the port on which the requesting client is connected, instead of broadcasting the DHCP response to all ports on the VLAN.



The routing switch's DHCP relay information (Option 82) feature can be used in networks where the DHCP servers are compliant with RFC 3046 Option 82 operation. DHCP servers that are not compliant with Option 82 operation ignore Option 82 fields. For information on configuring an Option 82 DHCP server, see the documentation provided with the server application.

Some client applications can append an Option 82 field to their DHCP requests; see the documentation provided for your client application.

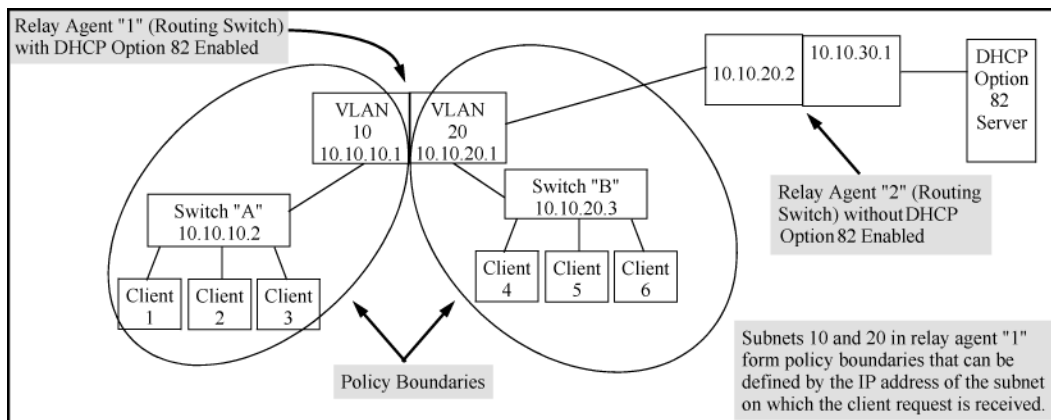
It is not necessary for all relay agents on the path between a DHCP client and the server to support Option 82, and a relay agent without Option 82 should forward DHCP packets regardless of whether they include Option 82 fields. However, Option 82 relay agents should be positioned at the DHCP policy boundaries in a network to provide maximum support and security for the IP addressing policies configured in the server.

### Option 82 server support

To apply DHCP Option 82, the routing switch must operate in conjunction with a server that supports Option 82. (DHCP servers that do not support Option 82 typically ignore Option 82 fields.) Also, the routing switch applies Option 82 functionality only to client request packets being **routed** to a DHCP server. DHCP relay with Option 82 does not apply to **switched** (nonrouted) client requests.

For information on configuring policies on a server running DHCP Option 82, see the documentation provided for that application.

**Figure 6:** Example of a DHCP Option 82 application



### General DHCP Option 82 requirements and operation Requirements

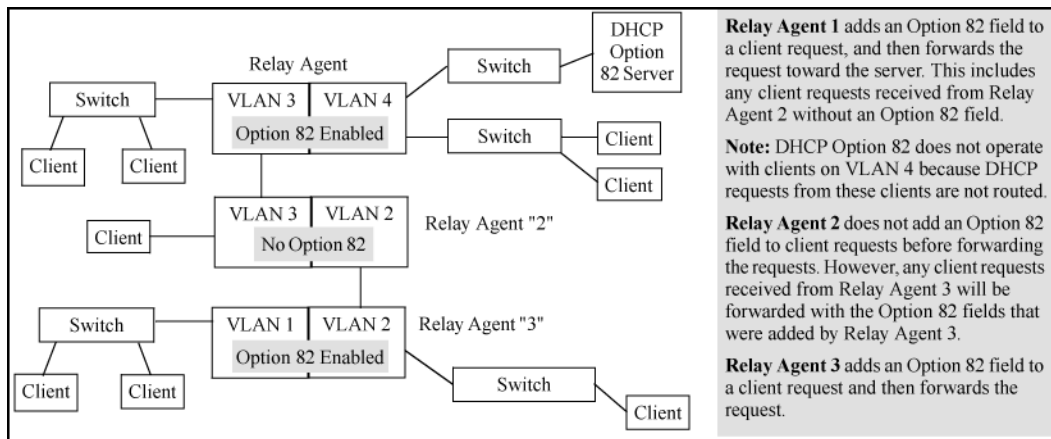
DHCP Option 82 operation is configured at the global config level and requires the following:

- IP routing enabled on the switch
- DHCP-relay option 82 enabled (global command level)
- Routing switch access to an Option 82 DHCP server on a different subnet than the clients requesting DHCP Option 82 support
- One IP helper address configured on each VLAN supporting DHCP clients

## General DHCP-relay operation with Option 82

Typically, the first (primary) Option 82 relay agent to receive a client's DHCP request packet appends an Option 82 field to the packet and forwards it toward the DHCP server identified by the IP helper address configured on the VLAN in which the client packet was received. Other, upstream relay agents used to forward the packet may append their own Option 82 fields, replace the Option 82 fields they find in the packet, forward the packet without adding another field, or drop the packet. (Intermediate next-hop routing switches without Option 82 capability can be used to forward—route—client request packets with Option 82 fields.) Response packets from an Option 82 server are routed back to the primary relay agent (routing switch) and include an IP addressing assignment for the requesting client and an exact copy of the Option 82 data the server received with the client request. The relay agent strips off the Option 82 data and forwards the response packet out the port indicated in the response as the Circuit ID (client access port.) Under certain validation conditions described later in this section, a relay agent detecting invalid Option 82 data in a response packet may drop the packet.

**Figure 7:** Example of DHCP Option 82 operation in a network with a noncompliant relay agent



## Option 82 field content

The remote ID and circuit ID subfields comprise the Option 82 field a relay agent appends to client requests. A DHCP server configured to apply a different IP addressing policy to different areas of a network uses the values in these subfields to determine which DHCP policy to apply to a given client request.

### Remote ID

Remote ID is a configurable subfield that identifies a policy area that comprises either the routing switch as a whole (by using the routing switch MAC address) or an individual VLAN configured on the routing switch (by using the IP address of the VLAN receiving the client request.)

- Use the IP address option if the server will apply different IP addressing policies to DHCP client requests from ports in different VLANs on the same routing switch.
- Use the Management VLAN option if a management VLAN is configured and you want all DHCP clients on the routing switch to use the same IP address. (This is useful if you are applying the same IP addressing policy to DHCP client requests from ports in different VLANs on the same routing switch.) Configuring this option means the management VLAN's IP address appears in the remote ID subfield of all DHCP requests originating with clients connected to the routing switch, regardless of the VLAN on which the requests originate.
- Use the MAC address option if, on a given routing switch, it does not matter to the DHCP server which VLAN is the source of a client request (that is, use the MAC address option if the IP addressing policies supported by the target DHCP server do not distinguish between client requests from ports in different VLANs in the same routing switch.)

To view the MAC address for a given routing switch, execute the `show system-information` command in the CLI.



## Using the CLI to view the switch MAC address

```
switch(config)# show system information

Status and Counters - General System Information

System Name       : HP Switch
System Contact    :
System Location   :

MAC Age Time (sec) : 300

Time Zone         : 0
Daylight Time Rule : None

Software revision : xx.15.xx      Base MAC Addr   : 0026f1-152e10
ROM Version       : xx.15.xx      Serial Number    : CN9458Q011
Allow V1 Modules  : Yes

Up Time          : 68 mins        Memory - Total   : 58,720,256
CPU Util (%)     : 5              Free            : 39,500,456

IP Mgmt - Pkts Rx : 28,959        Packet - Total   : 3022
                Pkts Tx : 1340    Buffers - Free   : 2902
                                                Lowest          : 2742
                                                Missed          : 0
```

## Circuit ID

Circuit ID is a nonconfigurable subfield that identifies the port number of the physical port through which the routing switch received a given DHCP client request and is necessary to identify if you want to configure an Option 82 DHCP server to use the Circuit ID to select a DHCP policy to assign to clients connected to the port. This number is the identity of the inbound port. On HPE fixed-port switches, the port number used for the circuit ID is always the same as the physical port number shown on the front of the switch. On HPE chassis switches, where a dedicated, sequential block of internal port numbers are reserved for each slot, regardless of whether a slot is occupied, the circuit ID for a given port is the sequential index number for that port position in the slot. (To view the index number assignments for ports in the routing switch, use the `walkmib ifname` command.)

For example, the Circuit ID for port 11 on an HPE switch is “11”.

## Using walkmib to determine the Circuit ID for a port on an HPE chassis

```
switch(config)# walkmib ifname
ifName.1 = 1
ifName.2 = 2
ifName.3 = 3
ifName.4 = 4
ifName.5 = 5
ifName.6 = 6
ifName.7 = 7
ifName.8 = 8
ifName.9 = 9
ifName.10 = 10
ifName.11 = 11
ifName.12 = 12
```

For example, suppose that you want port 10 on a given relay agent to support no more than five DHCP clients simultaneously. You can configure the server to allow only five IP addressing assignments at any one time for the circuit ID (port) and remote ID (MAC address) corresponding to port 10 on the selected relay agent.

Similarly, if you want to define specific ranges of addresses for clients on different ports in the same VLAN, you can configure the server with the range of IP addresses allowed for each circuit ID (port) associated with the remote ID (IP address) for the selected VLAN.

## Forwarding policies

DHCP Option 82 on HPE switches offers four forwarding policies, with an optional validation of server responses for three of the policy types (*append*, *replace*, or *drop*.)

## Configuration options for managing DHCP client request packets


Option 82 configuration	DHCP client request packet inbound to the routing switch	
	Packet has no Option 82 field	Packet includes an Option 82 field
Append	Append an Option 82 field	<p>Append</p> <p>allows the most detail in defining DHCP policy boundaries. For example, where the path from a client to the DHCP Option 82 server includes multiple relay agents with Option 82 capability, each relay agent can define a DHCP policy boundary and append its own Option 82 field to the client request packet. The server can then determine in detail the agent hops the packet took, and can be configured with a policy appropriate for any policy boundary on the path. Note:</p> <hr/> <p> In networks with multiple relay agents between a client and an Option 82 server, <i>append</i> can be used only if the server supports multiple Option 82 fields in a client request. If the server supports only one Option 82 field in a request, consider using the <i>keep</i> option.</p>
Keep	Append an Option 82 field	<p>If the relay agent receives a client request that already has one or more Option 82 fields, <i>keep</i> causes the relay agent to retain such fields and forward the request without adding another Option 82 field. But if the incoming client request does not already have any Option 82 fields, the relay agent appends an Option 82 field before forwarding the request. Some applications for <i>keep</i> include:</p> <ul style="list-style-type: none"> <li>• The DHCP server does not support multiple Option 82 packets in a client request, and there are multiple Option 82 relay agents in the path to the server.</li> <li>• The unusual case where DHCP clients in the network add their own Option 82 fields to their request packets, and you do not want any additional fields added by relay agents.</li> </ul> <p>This policy does not include the <i>validate</i> option (described in the next section) and allows forwarding of all server response packets arriving inbound on the routing switch (except those without a primary relay agent identifier.)</p>


Table Continued

Option 82 configuration		
	Packet has no Option 82 field	Packet includes an Option 82 field
Replace	Append an Option 82 field	Replace replaces any existing Option 82 fields from downstream relay agents (and/or the originating client) with an Option 82 field for the current relay agent. Some applications for <code>replace</code> include: <ul style="list-style-type: none"> <li>The relay agent is located at a point in the network that is a DHCP policy boundary, and you want to replace any Option 82 fields appended by downstream device with an Option 82 field from the relay agent at the boundary. (This eliminates downstream Option 82 fields you do not want the server to use when determining which IP addressing policy to apply to a client request.)</li> <li>In applications where the routing switch is the primary relay agent for clients that may append their own Option 82 field, you can use <code>replace</code> to delete these fields if you do not want them included in client requests reaching the server.</li> </ul>
Drop	Append an Option 82 field	Drop causes the routing switch to drop an inbound client request with an Option 82 field already appended. If no Option 82 fields are present, <code>drop</code> causes the routing switch to add an Option 82 field and forward the request. As a general guideline, configure <code>drop</code> on relay agents at the edge of a network, where an inbound client request with an appended Option 82 field may be unauthorized, a security risk, or for some other reason, should not be allowed.

### Multiple Option 82 relay agents in a client request path

Where the client is one router hop away from the DHCP server, only the Option 82 field from the first (and only) relay agent is used to determine the policy boundary for the server response. Where there are multiple Option 82 router hops between the client and the server, you can use different configuration options on different relay agents to achieve the results you want. This includes configuring the relay agents so that the client request arrives at the server with either one Option 82 field or multiple fields. (Using multiple Option 82 fields assumes that the server supports multiple fields and is configured to assign IP addressing policies based on the content of multiple fields.)


**Figure 8:** Example configured to allow only the primary relay agent to contribute an Option 82 field



The above combination allows for detection and dropping of client requests with spurious Option 82 fields. If none are found, the drop policy on the first relay agent adds an Option 82 field, which is then kept unchanged over the


next two relay agent hops ("B" and "C".) The server can then enforce an IP addressing policy based on the Option 82 field generated by the edge relay agent ("A".) In this example, the DHCP policy boundary is at relay agent 1.

**Figure 9:** Example configured to allow multiple relay agents to contribute an Option 82 field



This is an enhancement of the previous example. In this case, each hop for an accepted client request adds a new Option 82 field to the request. A DHCP server capable of using multiple Option 82 fields can be configured to use this approach to keep a more detailed control over leased IP addresses. In this example, the primary DHCP policy boundary is at relay agent "A," but more global policy boundaries can exist at relay agents "B" and "C."

**Figure 10:** Example allowing only an upstream relay agent to contribute an Option 82 field



Like the first example, above, this configuration drops client requests with spurious Option 82 fields from clients on the edge relay agent. However, in this case, only the Option 82 field from the last relay agent is retained for use by the DHCP server. In this case the DHCP policy boundary is at relay agent "C." In the previous two examples the boundary was with relay "A."

## Validation of server response packets

A valid Option 82 server response to a client request packet includes a copy of the Option 82 fields the server received with the request. With validation disabled, most variations of Option 82 information are allowed, and the corresponding server response packets are forwarded.

Server response validation is an option you can specify when configuring Option 82 DHCP for `append`, `replace`, or `drop` operation. Enabling validation on the routing switch can enhance protection against DHCP server responses that are either from untrusted sources or are carrying invalid Option 82 information.

With validation enabled, the relay agent applies stricter rules to variations in the Option 82 fields of incoming server responses to determine whether to forward the response to a downstream device or to drop the response due to invalid (or missing) Option 82 information. The following table describes relay agent management of DHCP server responses with optional validation enabled and disabled.

**Table 8:** Relay agent management of DHCP server response packets.

Response packet content	Option 82 configuration	Validation enabled on the relay agent	Validation disabled (the default)
Valid DHCP server response packet without an Option 82 field.	append , replace, or drop <sup>1</sup>	Drop the server response packet.	Forward server response packet to a downstream device.
	keep <sup>2</sup>	Forward server response packet to a downstream device.	Forward server response packet to a downstream device.
The server response packet carries data indicating a given routing switch is the primary relay agent for the original client request, but the associated Option 82 field in the response contains a remote ID and circuit ID combination that did not originate with the given relay agent.	append	Drop the server response packet.	Forward server response packet to a downstream device.
	replace or drop <sup>1</sup>	Drop the server response packet.	Drop the server response packet.
	keep <sup>2</sup>	Forward server response packet to a downstream device.	Forward server response packet to a downstream device.
The server response packet carries data indicating a given routing switch is the primary relay agent for the original client request, but the associated Option 82 field in the response contains a <b>Remote ID</b> that did not originate with the relay agent.	append	Drop the server response packet.	Forward server response packet to a downstream device.
	replace or drop <sup>1</sup>	Drop the server response packet.	Drop the server response packet.

*Table Continued*

Response packet content	Option 82 configuration	Validation enabled on the relay agent	Validation disabled (the default)
	keep <sup>2</sup>	Forward server response packet to a downstream device.	Forward server response packet to a downstream device.
All other server response packets <sup>3</sup>	append , keep <sup>2</sup> , replace, or drop <sup>1</sup>	Forward server response packet to a downstream device.	Forward server response packet to a downstream device.

<sup>1</sup>Drop is the recommended choice because it protects against an unauthorized client inserting its own Option 82 field for an incoming request.

<sup>2</sup> A routing switch with DHCP Option 82 enabled with the `keep` option forwards all DHCP server response packets except those that are not valid for either Option 82 DHCP operation (compliant with RFC 3046) or DHCP operation without Option 82 support (compliant with RFC 2131.)

<sup>3</sup> A routing switch with DHCP Option 82 enabled drops an inbound server response packet if the packet does not have any device identified as the primary relay agent (`giaddr=null`; see RFC 2131.)

## Multinetted VLANs

On a multinetted VLAN, each interface can form an Option 82 policy boundary within that VLAN if the routing switch is configured to use IP for the remote ID suboption. That is, if the routing switch is configured with IP as the remote ID option and a DHCP client request packet is received on a multinetted VLAN, the IP address used in the Option 82 field will identify the subnet on which the packet was received instead of the IP address for the VLAN. This enables an Option 82 DHCP server to support more narrowly defined DHCP policy boundaries instead of defining the boundaries at the VLAN or whole routing switch levels. If the MAC address option (the default) is configured instead, the routing switch MAC address will be used regardless of which subnet was the source of the client request. (The MAC address is the same for all VLANs configured on the routing switch.)

All request packets from DHCP clients in the different subnets in the VLAN must be able to reach any DHCP server identified by the IP helper addresses configured on that VLAN.

## Configuring Option 82

For information on Option 82, see the sections beginning with **DHCP Option 82** on page 70.

To configure DHCP Option 82 on a routing switch, enter the `dhcp-relay option 82` command.

### Syntax:

```
dhcp-relay option 82 <append[validate]|replace[validate]|drop[validate]|keep> [ip|mac|mgmt-vlan]
```

append	Configures the switch to append an Option 82 field to the client DHCP packet. If the client packet has existing Option 82 field(s) assigned by another device, the new field is appended to the existing fields. The appended Option 82 field includes the switch Circuit ID (inbound port number*) associated with the client DHCP packet, and the switch Remote ID. The default switch remote ID is the MAC address of the switch on which the packet was received from the client. To use the incoming VLAN's IP address or the Management VLAN IP address (if configured) for the remote ID instead of the switch MAC address, use the <code>ip</code> or <code>mgmt-vlan</code> option (below).
replace	Configures the switch to replace existing Option 82 fields in an inbound client DHCP packet with an Option 82 field for the switch. The replacement Option 82 field includes the switch circuit ID (inbound port number*) associated with the client DHCP packet and the switch remote ID. The default switch remote ID is the MAC address of the switch on which the packet was received from the client. To use the incoming VLAN's IP address or the Management VLAN IP address (if configured) for the remote ID instead of the switch MAC address, use the <code>ip</code> or <code>mgmt-vlan</code> option (below).
drop	Configures the routing switch to unconditionally drop any client DHCP packet received with existing Option 82 fields. This means that such packets will not be forwarded. Use this option where access to the routing switch by untrusted clients is possible. If the routing switch receives a client DHCP packet without an Option 82 field, it adds an Option 82 field to the client and forwards the packet. The added Option 82 field includes the switch circuit ID (inbound port number*) associated with the client DHCP packet and the switch remote ID. The default switch remote ID is the MAC address of the switch on which the packet was received from the client. To use the incoming VLAN's IP address or the Management VLAN IP address (if configured) for the remote ID instead of the switch MAC address, use the <code>ip</code> or <code>mgmt-vlan</code> option (below).
keep	For any client DHCP packet received with existing Option 82 fields, configures the routing switch to forward the packet as-is, without replacing or adding to the existing Option 82 fields.

*Table Continued*

[validate]	Operates when the routing switch is configured with append, replace, or drop as a forwarding policy. With <code>validate</code> enabled, the routing switch applies stricter rules to an incoming Option 82 server response to determine whether to forward or drop the response. For more information, see <a href="#">Validation of server response packets</a> .
[ip   mac   mgmt-vlan]	<p>Specifies the remote ID suboption that the switch uses in Option 82 fields added or appended to DHCP client packets. The type of remote ID defines DHCP policy areas in the client requests sent to the DHCP server. If a remote ID suboption is not configured, the routing switch defaults to the <code>mac</code> option. See <a href="#">Option 82 field content</a> on page 72.</p> <ul style="list-style-type: none"> <li>• <code>ip</code>: Specifies the IP address of the VLAN on which the client DHCP packet enters the switch.</li> <li>• <code>mac</code>: Specifies the routing switch's MAC address. (The MAC address used is the same MAC address that is assigned to all VLANs configured on the routing switch.) This is the default setting.</li> <li>• <code>mgmt-vlan</code>: Specifies the IP address of the (optional) management VLAN configured on the routing switch. Requires that a management VLAN is already configured on the switch. If the management VLAN is multinetted, the primary IP address configured for the management VLAN is used for the remote ID.</li> </ul> <p>If you enter the <code>dhcp-relay option 82</code> command without specifying either <code>ip</code> or <code>mac</code>, the MAC address of the switch on which the packet was received from the client is configured as the remote ID. For information about the remote ID values used in the Option 82 field appended to client requests, see <a href="#">Option 82 field content</a> on page 72.</p>

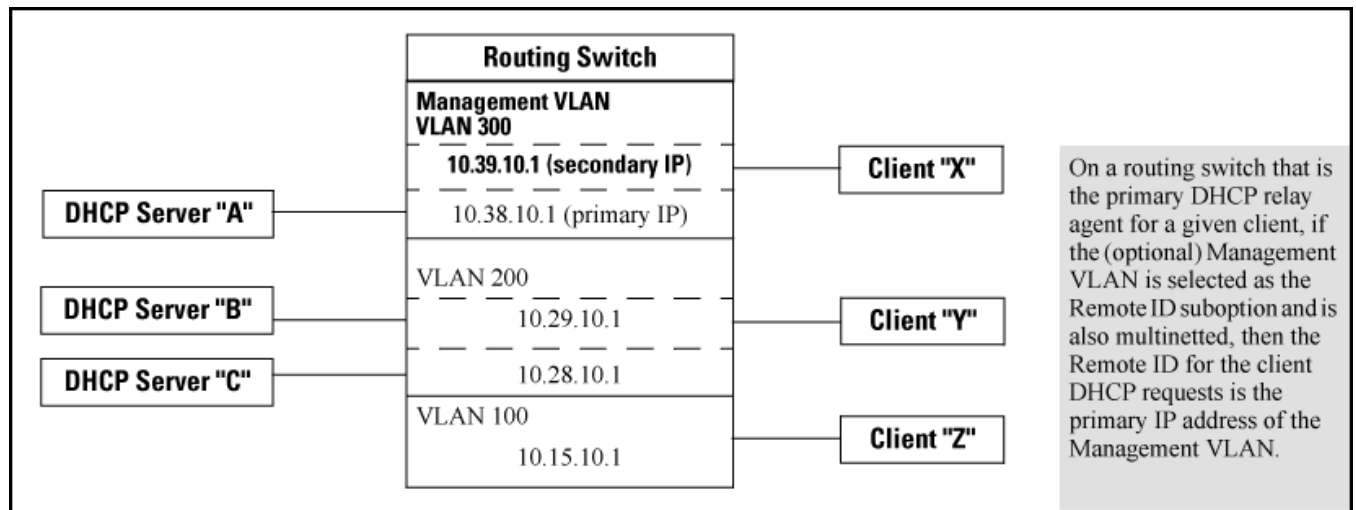
## Example of Option 82 configuration

In the routing switch shown below, option 82 has been configured with `mgmt-vlan` for the remote ID.

```
switch(config)# dhcp-relay option 82 append mgmt-vlan
```

The resulting effect on DHCP operation for clients X, Y, and Z is shown in the following table.

**Figure 11:** DHCP Option 82 when using the management VLAN as the remote ID suboption





**Table 9:** DHCP operation for the topology in DHCP Option 82 when using the management VLAN as the remote ID suboption

Client	Remote ID	giaddr	DHCP server	
X	10.38.10.1	10.39.10.1	A only	If a DHCP client is in the management VLAN, its DHCP requests can go only to a DHCP server that is also in the management VLAN. Routing to other VLANs is not allowed.
Y	10.38.10.1	10.29.10.1	B or C	Clients outside of the management VLAN can send DHCP requests only to DHCP servers outside of the management VLAN. Routing to the management VLAN is not allowed.
Z	10.38.10.1	10.15.10.1	B or C	

### Operating notes

- This implementation of DHCP relay with Option 82 complies with the following RFCs:
  - RFC 2131
  - RFC 3046
- Moving a client to a different port allows the client to continue operating as long as the port is a member of the same VLAN as the port through which the client received its IP address. However, rebooting the client after it moves to a different port can alter the IP addressing policy the client receives if the DHCP server is configured to provide different policies to clients accessing the network through different ports.
- The IP address of the primary DHCP relay agent receiving a client request packet is automatically added to the packet, and is identified as the giaddr (gateway interface address.) (That is, the giaddr is the IP address of the VLAN on which the request packet was received from the client.) For more information, see RFC 2131 and RFC 3046.
- DHCP request packets from multiple DHCP clients on the same relay agent port will be routed to the same DHCP servers. When using 802.1X on a switch, a port's VLAN membership may be changed by a RADIUS server responding to a client authentication request. In this case the DHCP servers accessible from the port may change if the VLAN assigned by the RADIUS server has different DHCP helper addresses than the VLAN used by unauthenticated clients.
- Where multiple DHCP servers are assigned to a VLAN, a DHCP client request cannot be directed to a specific server. Thus, where a given VLAN is configured for multiple DHCP servers, all these servers should be configured with the same IP addressing policy.
- Where routing switch "A" is configured to insert its MAC address as the remote ID in the Option 82 fields appended to DHCP client requests, and upstream DHCP servers use that MAC address as a policy boundary for assigning an IP addressing policy, then replacing switch "A" makes it necessary to reconfigure the upstream DHCP servers to recognize the MAC address of the replacement switch. This does not apply in the case where an upstream relay agent "A" is configured with `option 82 replace`, which removes the Option 82 field originally inserted by switch "A."

- Relay agents without Option 82 can exist in the path between Option 82 relay agents and an Option 82 server. The agents without Option 82 forward client requests and server responses without any effect on Option 82 fields in the packets.
- If the routing switch cannot add an Option 82 field to a client's DHCP request because the message size exceeds the MTU size, the request is forwarded to the DHCP server without Option 82 data and an error message is logged in the switch's Event Log.
- Because routing is not allowed between the Management VLAN and other VLANs, a DHCP server must be available in the management VLAN if clients in the management VLAN require a DHCP server.
- If the Management VLAN IP address configuration changes after `mgmt-vlan` has been configured as the remote ID suboption, the routing switch dynamically adjusts to the new IP addressing for all future DHCP requests.
- The Management VLAN and all other VLANs on the routing switch use the same MAC address.

## UDP broadcast forwarding

### Overview

Some applications rely on client requests sent as limited IP broadcasts addressed to a UDP application port. If a server for the application receives such a broadcast, the server can reply to the client. Since typical router behavior, by default, does not allow broadcast forwarding, a client's UDP broadcast requests cannot reach a target server on a different subnet unless the router is configured to forward client UDP broadcasts to that server.

A switch with routing enabled includes optional per-VLAN UDP broadcast forwarding that allows up to 256 server and/or subnet entries on the switch (16 entries per-VLAN.) If an entry for a particular UDP port number is configured on a VLAN, and an inbound UDP broadcast packet with that port number is received on the VLAN, the switch routes the packet to the appropriate subnet. (Each entry can designate either a single device or a single subnet. The switch ignores any entry that designates multiple subnets.)



The number of UDP broadcast forwarding entries supported is affected by the number of IP helper addresses configured to support DHCP relay. See **Operating notes for UDP broadcast forwarding** on page 86.

A UDP forwarding entry includes the desired UDP port number and can be either an IP unicast address or an IP subnet broadcast address for the subnet the server is in. Thus, an incoming UDP packet carrying the configured port number will be:

- Forwarded to a specific host if a unicast server address is configured for that port number.
- Broadcast on the appropriate destination subnet if a subnet address is configured for that port number.

A UDP forwarding entry for a particular UDP port number is always configured in a specific VLAN and applies only to client UDP broadcast requests received inbound on that VLAN. If the VLAN includes multiple subnets, the entry applies to client broadcasts with that port number from any subnet in the VLAN.

For example, VLAN 1 (15.75.10.1) is configured to forward inbound UDP packets as shown in the following table.

**Table 10: Example of a UDP packet-forwarding environment**

Interface	IP address	Subnet mask	Forwarding address	UDP port	Notes
VLAN 1	15.75.10.1	255.255.255.0	15.75.11.43	1188	Unicast address for forwarding inbound UDP packets with UDP port 1188 to a specific device on VLAN 2.
			15.75.11.255	1812	Broadcast address for forwarding inbound UDP packets with UDP port 1812 to any device in the 15.75.11.0 network.
			15.75.12.255	1813	Broadcast address for forwarding inbound UDP packets with UDP port 1813 to any device in the 15.75.12.0 network.
VLAN 2	15.75.11.1	255.255.255.0	None	N/A	Destination VLAN for UDP 1188 broadcasts from clients on VLAN 1. The device identified in the unicast forwarding address configured in VLAN 1 must be on this VLAN. Also the destination VLAN for UDP 1812 from clients on VLAN 1.
VLAN 3	15.75.12.1	255.255.255.0	None	N/A	Destination VLAN for UDP 1813 broadcasts from clients on VLAN 1.



If an IP server or subnet entry is invalid, a switch will not try to forward UDP packets to the configured device or subnet address.

## Subnet masking for UDP forwarding addresses

The subnet mask for a UDP forwarding address is the same as the mask applied to the subnet on which the inbound UDP broadcast packet is received. To forward inbound UDP broadcast packets as limited broadcasts to other subnets, use the broadcast address that covers the subnet you want to reach. For example, if VLAN 1 has an IP address of 15.75.10.1/24 (15.75.10.1 255.255.255.0), you can configure the following unicast and limited broadcast addresses for UDP packet forwarding to subnet 15.75.11.0:

Forwarding destination type	IP address
UDP unicast to a single device in the 15.75.11.0 subnet	15.75.11.X
UDP broadcast to subnet 15.75.11.0	15.75.11.255

## Configuring and enabling UDP broadcast forwarding

To configure and enable UDP broadcast forwarding on the switch:

### Procedure

1. Enable routing.
2. Globally enable UDP broadcast forwarding.
3. On a per-VLAN basis, configure a forwarding address and UDP port type for each type of incoming UDP broadcast you want routed to other VLANs.

## Globally enabling UDP broadcast forwarding

### Syntax:

```
ip udp-bacst-forward  
no ip udp-bcast-forward
```

Enables or disables UDP broadcast forwarding on the routing switch. Routing must be enabled before executing this command.

Using the `no` form of this command disables any `ip forward protocol udp` commands configured in VLANs on the switch.

Default: Disabled

## Configuring UDP broadcast forwarding on individual VLANs

This command routes an inbound UDP broadcast packet received from a client on the VLAN to the unicast or broadcast address configured for the UDP port type.

### Syntax:

```
[no] ip forward-protocol udp < ip-address > {< port-number | port-name >}
```

Used in a VLAN context to configure or remove a server or broadcast address and its associated UDP port number. You can configure a maximum of 16 `forward-protocol udp` assignments in a given VLAN. The switch allows a total of 256 `forward-protocol udp` assignments across all VLANs.

You can configure UDP broadcast forwarding addresses regardless of whether UDP broadcast forwarding is globally enabled on the switch. However, the feature does not operate unless globally enabled.

<ip-address>

This can be either of the following:

- The unicast address of a destination server on another subnet. For example: 15.75.10.43.
- The broadcast address of the subnet on which a destination server operates. For example, the following address directs broadcasts to All hosts in the 15.75.11.0 subnet: 15.75.11.255.



The subnet mask for a forwarded UDP packet is the same as the subnet mask for the VLAN (or subnet on a multinetted VLAN) on which the UDP broadcast packet was received from a client.

<udp-port-#>

Any UDP port number corresponding to a UDP application supported on a device at the specified unicast address or in the subnet at the specified broadcast address. For more information on UDP port numbers, see [TCP/UDP port number ranges](#) on page 86.

<port-name>

Allows use of common names for certain well-known UDP port numbers. You can type in the specific name instead of having to recall the corresponding number:

**dns**

Domain name service (53)

**netbios-ns**

NetBIOS name service (137)

**netbios-dgm**

NetBIOS datagram service (138)

**radius**

Remote authentication dial-in user service (1812)

**radius-old**

Remote authentication dial-in user service (1645)

**rip**

Routing information protocol (520)

**snmp**

Simple network management protocol (161)

**snmp-trap**

Simple network management protocol (162)

**tftp**

Trivial file transfer protocol (69)

**timep**

Time protocol (37)

**Example:**

The following command configures the routing switch to forward UDP broadcasts from a client on VLAN 1 for a time protocol server:

```
switch(vlan-1)# ip forward-protocol udp 15.75.11.155 timep
```

## Viewing the current IP forward-protocol configuration

### Syntax:

```
show ip forward-protocol [vlan < vid >]
```

Displays the current status of UDP broadcast forwarding and lists the UDP forwarding addresses configured on all static VLANs in the switch or on a specific VLAN.

The following example shows the global display showing UDP broadcast forwarding status and configured forwarding addresses for inbound UDP broadcast traffic for all VLANs configured on the routing switch.

### Displaying global IP forward-protocol status and configuration

```
switch(config)# show ip forward-protocol
```

```
IP Forwarder Addresses
```

```
    UDP Broadcast Forwarding: Disabled
```

```
VLAN: 1
```

```
IP Forward Addresses  UDP Port
```

```
-----
```

```
15.75.11.43           37
15.75.11.255          53
15.75.12.255          1813
```

```
VLAN: 2
```

```
IP Forward Addresses  UDP Port
```

```
-----
```

```
15.75.12.255          1812
```

The following example shows the display of UDP broadcast forwarding status and the configured forwarding addresses for inbound UDP broadcast traffic on VLAN 1.

### Displaying IP forward-protocol status and per-VLAN configuration

```
switch(config)# show ip forward-protocol vlan 1
```

```
IP Forwarder Addresses
```

```
    UDP Broadcast Forwarding: Disabled
```

```
IP Forward Addresses  UDP Port
```

```
-----
```

```
15.75.11.43           37
15.75.11.255          53
15.75.12.255          1813
```

## Operating notes for UDP broadcast forwarding

### Maximum number of entries

The number of UDP broadcast entries and IP helper addresses combined can be up to 16 per VLAN, with an overall maximum of 2048 on the switch. (IP helper addresses are used with the switch's DHCP relay operation.)

For example, if VLAN 1 has two IP helper addresses configured, you can add up to 14 UDP forwarding entries in the same VLAN.

### TCP/UDP port number ranges

There are three ranges:

- Well-known ports: 0 to 1023
- Registered ports: 1024 to 49151
- Dynamic and/or private ports: 49152 to 65535

For more information, including a listing of UDP/TCP port numbers, go to the **Internet Assigned Numbers Authority (IANA)** website at: [www.iana.org](http://www.iana.org).

Click on:

### Protocol Number Assignment Services

P (Under "Directory of General Assigned Numbers" heading)

### Port Numbers

## Messages related to UDP broadcast forwarding

Message	Meaning
udp-bcast-forward: IP Routing support must be enabled first.	Appears in the CLI if an attempt to enable UDP broadcast forwarding has been made without IP routing being enabled first. Enable IP routing, then enable UDP broadcast forwarding.
UDP broadcast forwarder feature enabled	UDP broadcast forwarding has been globally enabled on the router. Appears in the Event Log and, if configured, in SNMP traps.
UDP broadcast forwarder feature disabled	UDP broadcast forwarding has been globally disabled on the routing switch. This action does not prevent you from configuring UDP broadcast forwarding addresses, but does prevent UDP broadcast forwarding operation. Appears in the Event Log and, if configured, in SNMP traps.
UDP broadcast forwarder must be disabled first.	Appears in the CLI if you attempt to disable routing while UDP forwarding is enabled on the switch.

## Enabling forwarding of IP directed broadcasts (CLI)

To enable forwarding of IP directed broadcasts, enter the following CLI command:

### Syntax:

```
no ip directed-broadcast
```

```
switch(config)# ip directed-broadcast
```

HPE software makes the forwarding decision based on the routing switch's knowledge of the destination network prefix. Routers cannot determine that a message is unicast or directed broadcast apart from the destination network prefix. The decision to forward or not forward the message is by definition only possible in the last-hop router.

### Introduction to feature

Wake-on-LAN (WOL) is an Ethernet computer networking standard that allows a computer to be turned on or awakened by a network message. The message is sent by a program executed on the same local area network. Messages can also be initiated from another network by using subnet directed broadcasts or a WOL gateway service. WOL is implemented using specially designed packet called magic packet. WOL is enabled on the switch by using a `ip directed-broadcast` command with an IPv4 configuration, which can be used to specify an access list name, thus avoiding unnecessary administrative overhead.

IP directed-broadcasts would only be forwarded if permitted by the associated access-list. An `implicit deny` at the end of an access list drops all IP directed-broadcasts that are not authorized according to the access-list entries.



---

IP routing must be enabled on the switch for this feature to work.

---

### CLI commands

The optional association of access-list with IP directed-broadcast allows user to filter directed broadcast traffic alone based on access-list entry rule. The feature's CLI includes an optional parameter to specify access-list name along with the already existing `ip directed-broadcast` command. The access-list rule specified is applied globally on the switch and is not specific to any vlan's alone. There is an Implicit Deny at the end of an access list that will drop all IP Directed Broadcasts that do not match any of the access list entries.

### Configuration commands

Enable IP directed broadcast forwarding for Wake-on-LAN support. An optional ACL can also be applied to control what packets are forwarded.

### Syntax

```
Switch(config)# ip directed-broadcast [access-group <ACL-ID>]
```



## access-group

Apply the specified access control list.

## access-list-name-str

ASCII string specifying an ACL

---

## Example configuration

```
Switch(config)# ip directed-broadcast [access-group] <wol-acl>
```

---

## <wol-acl> entries

```
ip access-list extended <wol-acl>
10 permit ip 192.168.1.1 255.255.255.0 182.168.1.1 55.255.255.0
20 deny ip 172.168.1.1 255.255.255.0 162.168.1.1 255.255.255.0
Exit
```

---


## Example running configuration

```
; J9573A Configuration Editor; Created on release #xx.15.18.0000x
; Ver #06:7c.fd.ff.ff.3f.ef:57
hostname "switch"
module 1 type j9573x
ip access-list extended "wol-acl"
....10 permit ip 192.168.1.10 0.0.0.0 182.168.1.1 0.0.0.255
....exit

ip directed-broadcast access-group "wol-acl"
ip routing
snmp-server community "public" unrestricted
oobm
....ip address dhcp-bootp
    exit
vlan 1
....name "DEFAULT_VLAN"
....no untagged 1,23-24
....untagged 2-22,25-26
....ip address dhcp-bootp
....exit
vlan 10
....name "VLAN10"
....untagged 1
....ip address 192.168.1.1 255.255.255.0
....exit
vlan 20
....name "VLAN20"
....untagged 23-24
```

```
....ip address 182.168.1.1 255.255.255.0
....exit
```

**Figure 12: Configuration diagram**



- If specified ACL ID is nonexistent, it is not possible to associate with IP Directed Broadcast. An error will be shown to the user.
- It is not allowed to delete an ACL which is associated with IP Directed Broadcast and on attempt, an error message will be shown to user.
- The same ACL *wol-acl* can be applied to any other interface like VLAN, port, and tunnel.

## Show commands

IP directed broadcast hit counts for the associated access-list with can be displayed using the `show` command.

## Show statistics

Show IPV4 ACL Statistics.

## Syntax

```
show statistics aclv4 <acl-id>
```

## Options

```
port <port>
vlan <vlan-id> vlan
ip-directed-broadcast
```



---

Please note that the existing help text of all other parameters listed other than newly added `ip-directed-broadcast` will remain the same.

---

## Syntax

```
show statistics aclv4 <acl-name-str>
```

### ip-directed-broadcast

Show Statistics for the IP Directed Broadcast ACL.

```
switch # show statistics aclv4 wol-acl ip-directed-broadcast
HitCounts for ip-directed-broadcast ACL wol-acl
Total
(      0 )      10 permit ip 192.168.1.1 255.255.255.0 182.168.1.1 55.255.255.0
(      0 )      20 deny ip 172.168.1.1 255.255.255.0 162.168.1.1 255.255.255.0
```

## Clear command

The hit count statistics for ACL on IP directed broadcast can be cleared using clear command.

## Syntax

```
clear statistics aclv4 <acl-id>
```

## Options

```
port <port>
vlan <vlan-id> vlan
<ip-directed-broadcast>
```

Reset IPV4 Statistics.



---

Please note that the existing help text of all other parameters listed other than newly added `ip-directed-broadcast` will remain the same.

---

## Syntax

```
clear statistics aclv4 <acl-name-str>
```

ip-directed-broadcast Clear Statistics for the IP Directed Broadcast ACL.

## show access-list command

The existing “show access-list” command will have the following modification to support ip- directed-broadcast.

## Syntax

```
show access-list
```

## Options

```
<ACL-ID> [config]
<config>
<ip-directed-broadcast>
ports <<PORT-LIST>>
```

```
<radius>
<resources>
```

Show Access Control List Information.



Please note that the existing help of all other parameters listed other than newly added ip-directed-broadcast will remain the same.

### Show ACL's applied to IP Directed Broadcast traffic

```
show access-list <ip-directed-broadcast>
```

```
Switch # show access-list ip-directed-broadcast
```

```
Access Lists for IP Directed Broadcast
```

```
IPv4 : wol-acl Type: Extended
```

If user uses already existing `show access-list <ACL_NAME-STR>` command, the status of ACL on IP Directed Broadcast will be shown applied as in this example below.

```
switch # sh access-list wol-acl
```

```
Access Control Lists
```

```
.....Name: wol-acl
```

```
.....Type: Extended
```

```
.....Applied: Yes
```

```
.....SEQ Entry
```

```
-----
```

```
10 .Action: permit
```

```
.....Src IP: 192.168.1.1 Mask: 255.255.255.0 Port(s):
```

```
.....Dst IP: 182.168.1.1 Mask: 55.255.255.0 Port(s):
```

```
.....Proto : IP
```

```
.....TOS : - Precedence: -
```

```
20 Action: deny
```

```
.....Src IP: 172.168.1.1 Mask: 255.255.255.0 Port(s):
```

```
.....Dst IP: 162.168.1.1 Mask: 255.255.255.0 Port(s):
```

```
.....Proto : IP
```

```
.....TOS : - Precedence: -
```

### Disabling the directed broadcasts


```
switch(config)# no ip directed-broadcast
```

## Introduction

Routing Information Protocol version 2 includes authentication types `simple`, `none`, and MD5 (MD5 message-digest algorithm.)

Both `simple` and `none` authentication types are vulnerable to passive attacks currently widespread in the Internet. Clear text passwords, currently specified for use with Routing Internet Protocol version 2 (RIPv2), are no longer considered sufficient to provide security. Keyed MD5 is the standard authentication algorithm for RIPv2. It provides a greatly enhanced probability that a system being attacked will detect and ignore hostile messages.

**Figure 13:** MD5 use case diagram



## Configuration commands

Configure MD5 authentication for RIPv2 and MD5 key chain for RIPv2 interfaces by using the following commands.

### Syntax:

```
ip rip authentication-type <none|text|md5>
no ip rip authentication-type <none|text|md5>
```

Configure RIP authentication type as None, text and md5 on an interface.

When `no` is specified, the command disables RIP on the interface.

This command can be followed by a RIP configuration command. This is a VLAN context command that can be entered in a VLAN context or following the `vlan enable/disable/configure RIP` command on the VLAN interface.

#### **none**

Do not use authentication.

#### **text**

Use simple password.

#### **MD5**

Use MD5 authentication.

#### **Using MD5**

```
switch(vlan-10)# ip rip authentication-type md5
```

#### **Using none**

```
switch(vlan-10)# ip rip authentication-type none
```

#### **Syntax:**

```
ip rip md5-auth-key-chain <keychain-name>  
no ip rip md5-auth-key-chain <keychain-name>
```

Configure Routing Internet Protocol (RIP) md5 key chain on the VLAN interface.

When `no` is specified, the command disables RIP on the interface. The command can be followed by a RIP configuration command. This is a VLAN context command that can be entered in a VLAN context or following the `vlan <vlan-id>` command.

```
no ip rip md5-auth-key-chain
```

No authentication for RIP interfaces is the default configuration.

`md5-auth-key-chain`: Set the RIP MD5 authentication key chain (maximum 32 characters).

#### **Using MD5-auth-key-chain**

```
switch(vlan-10)# ip rip md5-auth-key-chain <abc>
```



`simple` and `none` authentication is supported on all RIP interfaces. MD5 authentication is supported for RIPv2 interfaces. With MD5 authentication, MD5-keyed digest is put into the packet instead secret password. This mechanism better protects RIPv2 routing message from any eavesdropping than simple or none.

## **Show commands**

#### **Syntax:**

```
show ip rip
```

Once MD5 authentication is configured, the command will show the authentication type as MD5 for the configured RIPv2 interface.

### IP RIP interface under VLAN context

```
switch(vlan-1)# show ip rip interface
```

RIP interface information

IP Address	Status	Send mode	Recv mode	Metric	Auth
30.0.0.1	Enabled	V2-only	V2-only	1	MD5

#### Syntax:

```
show key-chain <key-name>
```

Show association with RIPv2 interface if any.

### Show key-chain

```
DUT1(vlan-30)# show key-chain abc
```

Chain - test2

Key	Accept	Start	GMT	Accept	Stop	GMT	Send	Start	GMT	Send	Stop	GMT
---	+	-----		-----			-----			-----		
1		Bootup		Infinite			Bootup			Infinite		

OSPF Interface References

Interface

-----

OSPF Virtual Link References

Area/Virtual Link

-----

RIP Interface References

Interface

-----

30.0.0.1

## Operating notes

- If the authentication type MD5 is configured without a `md5-auth-key-chain`, the authentication will fail.
- If the `md5-auth-key-chain` is configured but authentication type not set to MD5, the authentication will fail.
- If the authentication type MD5 is configured and the `md5-auth-key-chain` is already configured, the MD5 authentication will begin working.
- If the `md5-auth-key-chain` is configured but the authentication type set to MD5, the MD5 authentication will begin working.
- When the MD5 authentication is working and you remove the `md5-auth-key-chain`, the MD5 authentication will fail.
- When the MD5 authentication is working and you change authentication type to other than MD5, the MD5 authentication will fail.
- When the MD5 authentication is working and you remove the used key-chain from global configuration, the request to remove the used key chain will fail.

- When the MD5 authentication is working and you remove the key from key-chain in the global configuration, the MD5 authentication will not work.



Hewlett Packard Enterprise recommends using a single key in the key-chain.

- Only RIPv2 supports MD5 authentication.

## Validation rules

Validation	Error/Warning/Prompt
If RIP interface is running in v1-only and v1-or-v2 mode.	Only RIPv2 interfaces support MD5 authentication.
When key chain does not exist.	Chain %s is not found.
When key chain exist without any key.	Chain %s has no keys configured.
When key chain exist without any key string.	Chain %s has no keystring configured.
If md5-auth-key-chain name length is < 0 or > 32.	Invalid length.

## Log messages

Event	Message
RMON_RIP_NO_VALID_KEY	SEND: No valid key found in key ring; no update is sent from interface %s.
RMON_RIP_AUTHENTICATION_FAILED	RECV: Authentication failed for packet received from IP address %s.
RMON_RIP_AUTHENTICATION_FAILED	RECV: Packet received on interface %s has its MD5 key expired.

## Error messages

### Configuring MD5 authentication for RIP v1 or RIP v1-or-v2 interface

```
switch(vlan-10)# ip rip v1-only
switch(vlan-10)# ip rip authentication-type md5

Only RIPv2 interfaces support MD5 authentication.
```

```
switch(vlan-10)# ip rip v1-or-v2
switch(vlan-10)# ip rip authentication-type md5

Only RIPv2 interfaces support MD5 authentication
```



---

### Configuring MD5 key chain for RIP v2 interface without key chain or key or keystring

```
switch(vlan-10)# ip rip md5-auth-key-chain rip-md5-chain  
Chain rip-md5-chain is not found.
```

```
switch(vlan-10)# ip rip md5-auth-key-chain rip-md5-chain  
Chain rip-md5-chain has no keys configured.
```

```
switch(vlan-10)# ip rip md5-auth-key-chain rip-md5-chain  
Chain rip-md5-chain has no keystring configured.
```

## RIPng for IPv6

While the mechanisms of RIP remain unchanged, RIPng for IPv6 has been added to include support for IPv6 addressing and prefixes, different packet formats, packet lengths, and no authentication on HPE switches.

RIPng is for IPv6 only just as RIPv2 is for IPv4 only. RIPv2 and RIPng must be regarded as two independent protocols with no interaction between them.

RIPng is specified by RFC 2080 and RFC 2081



RIPng and RIPv2 can be supported on the same interface/VLAN.

### Supported features

- RIPng global enable/disable  
Enables/Disables RIPng protocol in the config context.
- Split horizon  
Prevents the formation of loops in routing. A router is not allowed the advertisement of routes back to the interface where it was initially learned. Enabled by default. Split Horizon is a nonconfigurable feature.
- Poison-Reverse  
Optimizes the transmission of routing information and improves the time-to-reach network convergence. Enabled by default and can be disabled per VLAN interface.
- Redistribute connected/staticroutes  
RIPng protocol advertises routes learned from static and connected networks. to its peers.
- Metric configuration for imported routes  
Updates the metric for imported routes based on the value configured.
  - Router ripng default-metric — for routes imported from protocols other than RIPng
  - `vlan <id> ipv6 ripng`  
metric — for routes received from other RIPng peer
- Configuration of RIPng timers: update, timeout, and garbage collect.
  - Update timer defines interval between update messages.
  - Timeout timer defines route aging time.
  - The garbage-collect timer defines the time interval when the metric of a route is 16 to the time when it is deleted from the routing table.
- Administrative distances: The default value can be modified and the value is applied to all routes learned through RIPng.
- RIPng will listen only to RIPng packets sent to the multicast address FF02::9. All packets sent out will be addressed to FF02::9 and the source IP will be the link local IPv6 address of the VLAN.
- Route maps — Route maps are applied in the redistribution process to control route prefixes or to modify the attributes of the routes. Route-maps can be used in RIPng redistribution to apply route policy configurations.
- RIPng notifications/traps — Traps are generated as the result of finding an unusual condition while parsing an RIPng packet or a processing a timer event. Disabled by default.

### Limitations

Limits imposed on RIPng are as follows:

IPv6 loop back addresses cannot be redistributed into RIPng.

Number of interfaces/VLANs on which RIPng may be run:	128
Total number of routes supported:	5000
Maximum number of IPv6 addresses per Vlan:	32
Maximum number of IPv6 Vlans:	512
Maximum number of IPv6 addresses:	2048



---

Starting from 16.01 onwards, the redistribution of OSPFv3 external routes (E1/E2/N1/N2) into RIPng is not supported.

---

## Configure RIPng

From within the configuration context, use the following commands to configure, enable, disable a RIPng setting.

### Enable/Disable RIPng global

#### Syntax

```
router ripng enable | disable
```

#### Description

From within the configuration context, enable RIPng globally or disable RIPng globally.

### Configure a RIPng setting

#### Syntax

```
router ripng  
no router ripng
```

#### Description

From within the configuration context, configure a RIPng setting or enter RIPng context.

Use the `no` argument to remove all RIPng configurations.

### Configure a default metric

#### Syntax

```
router ripng default-metric 1-15
```

#### Description

Configure a default metric for routes that are imported from protocols other than RIPng.

The default value is 1.

## Configure the administrative distance for routes

### Syntax

```
router ripng distance 1-255
```

### Description

Configure the administrative distance for routes that are learned via RIPng.  
The default value is 120.

## Redistribute router RIPng

### Syntax

```
router ripng redistribute  
no router ripng redistribute
```

### Description

Redistribute connected/static/other protocols routes.  
Use `no` to disable redistribution of the specified protocol.

### Options

#### connected

Redistribute locally connected networks.

#### ospf3

Redistribute OSPFv3 routes.

#### static

Redistribute manually configured routes.

#### include-all

Include blackhole and reject routes.



---

Include-all option is only for static routes.

---

#### route-map

Redistribute a route map.



---

Route-map option comes only after we specify the protocol (static/connected).

---

## Usage

```
redistribute connected route-map NAME  
no redistribute connected route-map NAME
```

```
ospf3 route-map NAME  
no redistribute ospf3 route-map NAME
```

```
redistribute static include-all route-map NAME  
no redistribute static include-all route-map NAME
```

## Configure RIPng timers

### Syntax

```
router ripng timers
```

### Description

Configure RIPng timers.

### Options

#### garbage-collect

Set the garbage-collect interval for the route.

The default value is 120 seconds.

#### timeout

Set the interval for the route timeout.

The default value is 180 seconds.

#### update

Set the interval for the update timer.

The default value is 30 seconds.



---

HPE does not recommend changing the default values.

---

### Usage

```
router ripng timers garbage-collect 5-65535
```

```
router ripng timers timeout 5-65535
```

```
router ripng timers update 5-65535
```

## Enable/Disable RIPng traps

### Syntax

```
router ripng trap
```

## Description

Enable/Disable RIPng traps.

## Options

Traps are generated as the result of finding an unusual condition while parsing an RIPng packet or a processing a timer event. If more than one type of unusual condition is encountered while parsing the packet or processing an event, only the first one will generate a trap.

```
interface-state-change
```

Send a trap when the state of an interface changes.

```
interface-config-error
```

Send a trap when a configuration conflict occurs for an interface.

```
interface-receive-bad-packet
```

Send a trap when an invalid packet is received on an interface.

```
all
```

Enable all the RIPng traps.

## Usage

```
trap TRAP-NAME | all  
no trap TRAP-NAME | all
```

# VLAN Level Configuration

This is a VLAN context command. It can be entered in VLAN context as shown or following the `vlan VLAN-ID` command.

## IPv6 RIPng

### Syntax

```
ipv6 ripng  
no ipv6 ripng
```

### Description

Enables/disables/configures the RIPng protocol for IPv6 on the interface.

The argument `no` disables or disconfigures RIPng on the interface.

### Options

#### **enable**

Enable RIPng on the VLAN.

#### **metric**

Set the metric for the interface.

#### **poison-reverse**

Enable/Disable poison reverse.

## Show commands

If RIPng is not configured on the switch, any show commands related to RIPng are executed, the following output is displayed.

```
switch (config)# show ipv6 ripng
RIPng Configuration Information
RIPng protocol : Disabled
```

### Show IPv6 ripng general

#### Syntax

```
show ipv6 ripng general
```

#### Description

Displays RIPng global parameters only as shown below.

#### RIPng global parameters

```
switch(config)# show ipv6 ripng general
RIPng global parameters
RIPng protocol : Enabled
Default metric : 1
Administrative distance : 120
Route changes : 1090
Queries : 134457
Update time : 30
Timeout : 180
Garbage-collect time : 120
switch(config)#
```

### Show IPv6 ripng interface

#### Syntax

```
show ipv6 ripng interface
```

#### Description

Displays basic config, interface, and peer information as shown below.

#### Options

##### VLAN

Specify the VLAN of the interface requesting detailed information.

##### VLAN-ID

Enter a VLAN identifier or a VLAN name.

#### Usage

```
show ipv6 ripng interface vlan VLAN-ID
```

## RIPng interface information

```
switch(config)# show ipv6 ripng
RIPng global parameters
RIPng protocol : Enabled
Default metric : 1
Administrative distance : 120
Route changes : 2090
Queries : 134877
Update time : 30
Timeout : 180
Garbage-collect time : 120
```

```
RIPng interface information
VLAN          Status      Metric
-----
10            Enabled    1
20            Enabled    1
```

```
RIPng peer information
IPv6 Address                               Bad packets  Last update
-----
fe80::200:eff:feda:98b6%vlan10            0            27
```

## Show IPv6 RIPng peer

### Syntax

```
show ipv6 ripng peer
```

### Description

Shows the peers learned through RIPng.

## RIPng peer information

```
switch (config)# show ipv6 ripng peer
RIPng peer information
IPv6 Address                               Bad packets  Last update
-----
fe80::ab23:ccff:fef4:fc40 0            30
```



Since RIPng does not have an active peering mechanism, this command shows only those RIPng peers from which a route was taken and added to the routing table. For example, if two peers advertise the same route(s) with the same metric only one of them will be shown as peer.

## Show IPv6 RIPng redistribute

### Syntax

```
show ipv6 ripng redistribute
```

### Description



List the protocols that are being redistributed into RIPng.

### RIPng redistributing without route-maps

```
switch (config)# show ipv6 ripng redistribute
RIPng redistributing
Route type Route map Options
-----
Connected
```

### RIPng redistribute with route-maps

```
RIPng redistributing
Route type Route map Options
-----
Connected map2
static map1 Include blackhole and reject
```

## Show IPv6 RIPng traps

### Syntax

```
show ipv6 ripng traps
```

### Description

Display the enabled RIPng traps.

### RIPng Traps : Enabled

```
switch(config)#show ipv6 ripng traps
RIPng Traps : Enabled
RIPng Traps Enabled
-----
Interface State Change
Interface Configuration Error
Interface Bad Packet Receive Error
```

## Show IPv6 route RIPng

### Syntax

```
show ipv6 route ripng
```

### Description

Show the IPv6 routing table. The output can be restricted to a specific destination or type of route.

### Options

#### IPv6-ADDR

The destination IPv6 address for which to display the routes.

## Usage

```
show ipv6 route IPv6-ADDR static | connected | ripng | ospfv3
```

### IPv6 route entries

```
switch (config)# show ipv6 route
IPv6 Route Entries
T (Type):
S: Static C: Connected

Destination/      Gateway                               T   ST  Distance
Metric
-----
::1/128
lo0                C   NA  0
1
```

## Show ipv6 route summary

### Syntax

```
show ipv6 route summary
```

### Description

Show the summary of IPv6 routing table.

### IPv6 route summary

```
switch(config)#show ipv6 route summary
IPv6 Route Table Summary
Protocol Active Routes
-----
Connected      5
Ripng          5002
```

## Debug commands

### Debug IPv6 RIPng

#### Syntax

```
debug ipv6 ripng
```

#### Description

Enable debug messages for RIPng.

#### Options

##### database

Show RIPng database changes.

## events

Show RIPng events.

## trigger

Show RIPng trigger messages.

## Usage

```
debug ipv6 ripng database | events | trigger
```

## Additional commands

Following CLI commands are enhanced to accommodate RIPng.

### VLAN VLAN-ID IPv6

This is a VLAN context command.

#### Syntax

```
vlan VLAN-ID ipv6 ripng
```

#### Description

Enables/Disables/Configures RIPng protocol for IPv6 on the interface. The command `no ipv6 ripng enable` disables or disconfigures RIPng on the interface. This command can be followed by a RIPng configuration command.

### Show running config

#### Syntax

```
show running-config router { rip | ripng }
```

#### Description

Show the running configuration for layer 3 routing protocols.

#### Show running-config router ripng

```
router ripng
enable
default-metric 3
distance 95
redistribute connected
redistribute static
redistribute ospf3
```

#### Options

```
show running-config router rip
```

Show the running configuration for RIP.

```
show running-config router ripng
```

Show the running configuration for RIPng.

## Show running-config vlan

### Syntax

```
show running-config vlan VLAN-ID
```

### Description

Shows the IPv6 ripng vlan configuration along with other vlan specific configuration.

### show running-config vlan

```
switch (config)#show running-config vlan 15
vlan 15
name "VLAN15"
tagged Trk10
no ip address
ipv6 enable
ipv6 address 3005::10/64
ipv6 ripng enable
```

## Validation rules

Validation	Error/Warning/Prompt
Attempt to enable IPv6 RIPng before enabling ipv6 unicast routing.	IPv6 unicast routing must be enabled first.
Attempt to enable IPv6 RIPng before enabling IPv6 on any interface.	IPv6 must be enabled on at least one interface.
Attempt to configure RIPng on vlan without having assigned an IPv6 address for the vlan or the IPv6 status on the vlan is disabled.	IPv6 should be enabled before configuring RIPng.
Attempt to disable IPv6 on a vlan when RIPng is configured on that vlan.	To disable IPv6, RIPng configuration must be removed from this interface.
Attempt to disable ipv6 unicast routing when RIPng is configured on the switch.	RIPng must be disabled first.
Attempt to configure route map when redistribution is already configured.	Redistribution of routes without route-map must be disabled first.
Attempt to configure redistribution when route map is already configured.	Redistribution of routes with route-map must be disabled first.
Attempt to include blackhole or reject static routes when redistribution of static routes is already configured.	Redistribution of static routes must be disabled first.

*Table Continued*

Validation	Error/Warning/Prompt
Attempt to configure redistribution of static routes only when redistribution of blackhole or reject static routes is already configured.	Redistribution of blackhole/reject routes must be disabled first.
Attempt to configure garbage-collect time greater than time out.	Garbage-collect timer must be shorter than time out.
Attempt to configure garbage-collect time less than update-time.	Garbage-collect timer must be longer than update.
Attempt to configure timeout lesser than garbage-collect time.	Timeout must be longer than garbage-collect.
Attempt to configure update time greater than garbage-collect time.	Update timer must be shorter than garbage-collect.
User inputs vlan in show command but RIPng is not configured for that vlan.	RIPng is not configured on this interface.

## Event Log

Event	Message
RIPng has been configured on the device with the CLI command <code>router ripng enable</code> .	RIPng is enabled.
RIPng has been unconfigured on the device with the CLI command <code>router ripng disable</code> .	RIPng is disabled.
RIPng has been unconfigured on the device with the CLI command <code>no router ripng</code> . All the existing configuration of RIPng is deleted.	RIPng is disabled.
An incoming RIPng packet has been rejected because the source address is not IPv6.	Bad packet – protocol is not IPv6.
An incoming RIPng packet has been rejected because the source address is not link-local.	Bad packet – source address must be link-local.
An incoming RIPng packet has been rejected because the version number is invalid.	Bad packet – version must be 1.
An incoming RIPng packet has been rejected because the interface it was received on is marked to restrict RIPng updates.	Bad packet – received packet dropped on an interface that is marked to restrict RIPng updates.
An incoming RIPng packet has been rejected because it was sent by the switch itself.	Bad packet – originator and receiver are the same.

*Table Continued*

Event	Message
An incoming RIPng packet has been rejected because the reserved header field was not set to zero.	Bad packet – reserved field must be zero.
An incoming RIPng packet has been rejected because the hop limit is not 255.	Bad packet – hop limit must be 255.
An incoming RIPng packet has been rejected because the source port is not valid.	Bad packet – source port must be 521.

### Networking Websites

Hewlett Packard Enterprise Networking Information Library

[www.hpe.com/networking/resourcefinder](http://www.hpe.com/networking/resourcefinder)

Hewlett Packard Enterprise Networking Software

[www.hpe.com/networking/software](http://www.hpe.com/networking/software)

Hewlett Packard Enterprise Networking website

[www.hpe.com/info/networking](http://www.hpe.com/info/networking)

Hewlett Packard Enterprise My Networking website

[www.hpe.com/networking/support](http://www.hpe.com/networking/support)

Hewlett Packard Enterprise My Networking Portal

[www.hpe.com/networking/mynetworking](http://www.hpe.com/networking/mynetworking)

Hewlett Packard Enterprise Networking Warranty

[www.hpe.com/networking/warranty](http://www.hpe.com/networking/warranty)

### General websites

Hewlett Packard Enterprise Information Library

[www.hpe.com/info/EIL](http://www.hpe.com/info/EIL)

For additional websites, see [Support and other resources](#).

## Accessing Hewlett Packard Enterprise Support

- For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:  
<http://www.hpe.com/assistance>
- To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:  
<http://www.hpe.com/support/hpesc>

### Information to collect

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

## Accessing updates

- Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.
- To download product updates:

### Hewlett Packard Enterprise Support Center

[www.hpe.com/support/hpesc](http://www.hpe.com/support/hpesc)

### Hewlett Packard Enterprise Support Center: Software downloads

[www.hpe.com/support/downloads](http://www.hpe.com/support/downloads)

### Software Depot

[www.hpe.com/support/softwaredepot](http://www.hpe.com/support/softwaredepot)

- To subscribe to eNewsletters and alerts:  
[www.hpe.com/support/e-updates](http://www.hpe.com/support/e-updates)
- To view and update your entitlements, and to link your contracts and warranties with your profile, go to the Hewlett Packard Enterprise Support Center **More Information on Access to Support Materials** page:  
[www.hpe.com/support/AccessToSupportMaterials](http://www.hpe.com/support/AccessToSupportMaterials)



Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HPE Passport set up with relevant entitlements.

---

## Customer self repair

Hewlett Packard Enterprise customer self repair (CSR) programs allow you to repair your product. If a CSR part needs to be replaced, it will be shipped directly to you so that you can install it at your convenience. Some parts



do not qualify for CSR. Your Hewlett Packard Enterprise authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider or go to the CSR website:

<http://www.hpe.com/support/selfrepair>

## Remote support

Remote support is available with supported devices as part of your warranty or contractual support agreement. It provides intelligent event diagnosis, and automatic, secure submission of hardware event notifications to Hewlett Packard Enterprise, which will initiate a fast and accurate resolution based on your product's service level. Hewlett Packard Enterprise strongly recommends that you register your device for remote support.

If your product includes additional remote support details, use search to locate that information.

### Remote support and Proactive Care information

#### HPE Get Connected

[www.hpe.com/services/getconnected](http://www.hpe.com/services/getconnected)

#### HPE Proactive Care services

[www.hpe.com/services/proactivecare](http://www.hpe.com/services/proactivecare)

#### HPE Proactive Care service: Supported products list

[www.hpe.com/services/proactivecaresupportedproducts](http://www.hpe.com/services/proactivecaresupportedproducts)

#### HPE Proactive Care advanced service: Supported products list

[www.hpe.com/services/proactivecareadvancedsupportedproducts](http://www.hpe.com/services/proactivecareadvancedsupportedproducts)

### Proactive Care customer information

#### Proactive Care central

[www.hpe.com/services/proactivecarecentral](http://www.hpe.com/services/proactivecarecentral)

#### Proactive Care service activation

[www.hpe.com/services/proactivecarecentralgetstarted](http://www.hpe.com/services/proactivecarecentralgetstarted)

## Warranty information

To view the warranty for your product or to view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products* reference document, go to the Enterprise Safety and Compliance website:

[www.hpe.com/support/Safety-Compliance-EnterpriseProducts](http://www.hpe.com/support/Safety-Compliance-EnterpriseProducts)

### Additional warranty information

#### HPE ProLiant and x86 Servers and Options

[www.hpe.com/support/ProLiantServers-Warranties](http://www.hpe.com/support/ProLiantServers-Warranties)

#### HPE Enterprise Servers

[www.hpe.com/support/EnterpriseServers-Warranties](http://www.hpe.com/support/EnterpriseServers-Warranties)

#### HPE Storage Products

[www.hpe.com/support/Storage-Warranties](http://www.hpe.com/support/Storage-Warranties)

#### HPE Networking Products

[www.hpe.com/support/Networking-Warranties](http://www.hpe.com/support/Networking-Warranties)

## Regulatory information

To view the regulatory information for your product, view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at the Hewlett Packard Enterprise Support Center:

**[www.hpe.com/support/Safety-Compliance-EnterpriseProducts](http://www.hpe.com/support/Safety-Compliance-EnterpriseProducts)**

### **Additional regulatory information**

Hewlett Packard Enterprise is committed to providing our customers with information about the chemical substances in our products as needed to comply with legal requirements such as REACH (Regulation EC No 1907/2006 of the European Parliament and the Council). A chemical information report for this product can be found at:

**[www.hpe.com/info/reach](http://www.hpe.com/info/reach)**

For Hewlett Packard Enterprise product environmental and safety information and compliance data, including RoHS and REACH, see:

**[www.hpe.com/info/ecodata](http://www.hpe.com/info/ecodata)**

For Hewlett Packard Enterprise environmental information, including company programs, product recycling, and energy efficiency, see:

**[www.hpe.com/info/environment](http://www.hpe.com/info/environment)**

## **Documentation feedback**

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (**[docsfeedback@hpe.com](mailto:docsfeedback@hpe.com)**). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.

## Overview

The HPE mDNS Gateway and Google Chromecast solution adds support for Apple's Bonjour and Google's Chromecast discovery from a HPE switch. The solution uses mDNS protocol for discovery and is responsible for handling mDNS packets.

### Bonjour

HPE's mDNS Gateway solution supports Apple's Bonjour protocol to the switch.

Bonjour is Apple's implementation of a suite of zero-configuration networking protocols and is supported by both Mac OS X devices (such as laptops and desktops), and Apple iOS devices (such as iPhones and iPads).

Bonjour's zero-configuration network services benefits include:

- No longer having to assign IP addresses or host names to access network services on Mac OS X and Apple iOS devices
- Applications can leverage Bonjour to automatically detect required services.
- Interacts with other applications to allow for automatic connection of devices.
- Communication and data exchange is possible without user configuration.

### Google's Chromecast

Chromecast is a digital media player developed by Google. The device is a HDMI dongle that plays audio and video content on a high-definition screen by directly streaming it via Wi-Fi from the Internet or a local network. The media is selected, by users, to play on devices by enabling Chromecast mobile and web applications. Casting a tab for sites that are not Google Cast-enabled. mirrors most Google Chrome browser content running on the device (MAC OSX and Windows).

Chromecast uses a simple multicast protocol for discovery and launch. This protocol enables users to mirror their devices on a second screen.

### HPE mDNS protocol

HPE supports mDNS protocol implemented as a server. mDNS is the primary method of discovering a Chromecast that supports the v2 API. While SSDP/DIAL support is still present and used by some applications (such as "You Tube"), existing applications have to migrate to the new SDK using the new protocol.

## mDNS Gateway

The mDNS gateway, running on a switch, will listen for Bonjour responses and Bonjour queries and forward them to different subnets. Its main function is to forward Bonjour traffic by retransmitting the traffic between reflection enabled VLANs. The switches are configured interfaces in the VLANs for which they are performing packet reflection.



---

The mDNS gateway in a switch acts as an application layer gateway between subnets. An IP interface is required on each of the network that it is reflecting between.

---

## Service filtering

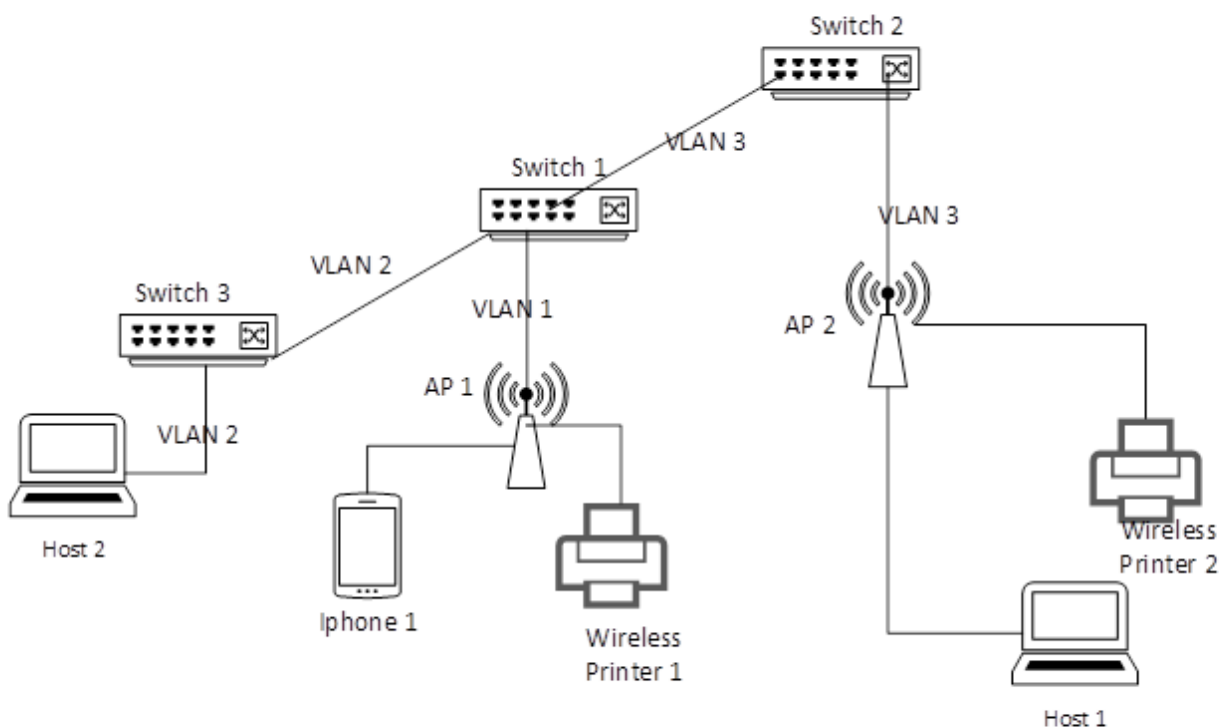
The mDNS profiles feature is responsible for applying filter profiles to mDNS resource records in mDNS response/query packets. The mDNS response/query can be filtered to give better control of the services. Service filtering allows network administrators to manipulate both the responses sent to and coming from clients in order to allow or deny mDNS services. This mechanism prevents clients from being aware of both specified services and announce specific services. These filters can be outbound from the switch to clients or inbound from clients to the switch. Profiles can be applied per-VLAN.

There is a global default which allows or denies traffic that does not match any rule. After a match is found other filter rules are ignored.



Service filtering cannot block the connection between devices. For example, if the client knows the remote device's IP address, they can still establish a connection without utilizing the mDNS protocol. Service filtering functions to keep names and addresses out of mDNS responses.

**Figure 14:** mDNS query and response assessment



- Switch 1 — Reflection enabled on VLAN 2 and VLAN 3
- Global Filters — set to permit both inbound and outbound mDNS traffic on Switch 1, 2 and 3.
- Specific Filter — Switch 1 – VLAN 3 – Deny –outbound – service type – wireless printer.
- Specific Filter — Switch 1 – VLAN 2 – Permit – inbound – instance name – Host 2.

## Wireless printer service process

Process overview of service for a wireless printer:

## Procedure

1. Wireless Printer 1 sends an mDNS response advertising printer service in Switch 1 on VLAN 1.
2. Switch 1 has no inbound filter in VLAN 1. The global filter set to **permit all**.
3. Switch 1 checks the outbound filter in VLAN 1. As there is no specific outbound filter, the global status is **permit all**. It will flood the packet in VLAN 1 except the source port.
4. iPhone 1 in VLAN 1 receives the service announcement.
5. Switch 1 checks the reflection status. Reflection is enabled on VLAN 2 and 3.
6. Switch 1 checks the outbound filter in VLAN 2. As there is no specific outbound filter, it will forward the service announcement in VLAN 2.
7. Default action **permit all**.
8. Switch 1 checks the outbound filter in VLAN 3. The outbound filter is set to **deny wireless printer** therefore the packet will not be forwarded to VLAN 3.
9. Switch 3 receives the service advertisement in VLAN 2. It will flood the packet in VLAN 2 except the source port.
10. Host 2 in Switch 3 receives the service announcement.

## Wireless Printer advertising printer service

The following procedure depicts an advertising service process for a wireless printer in the form of an example.

### Procedure

1. Wireless Printer 2 sends an mDNS response advertising printer service in VLAN 3.
2. Switch 2 does not have any inbound filter in VLAN 3, so it receives the wireless printer service announcement.
3. Switch 2 checks the outbound filter in VLAN 3. There is no specific outbound filter on VLAN 3, so it floods the service announcement in VLAN 3 (except at the source port.)
4. Switch 2 checks the reflection status. Since switch 2 is not enabled, switch 2 does not forward.
5. As there is no inbound filter in VLAN 3 of switch 1, it receives the service announcement on VLAN 3. When switch 1 checks the outbound filter in VLAN 3, there is `deny operation for service type wireless printer` error message. Therefore switch 1 will not flood the packet in VLAN 3.
6. Switch 1 checks the reflection status. The reflection is enabled on VLAN 2 and 3 however VLAN 3 is incoming so the reflection will not function. In VLAN 2 it checks the outbound filter. There is no outbound filter in VLAN 2 so switch 1 forwards the service announcement in VLAN 2.
7. Switch 3 does not have any inbound filter therefore. It receives service announcements in VLAN 2.
8. Switch 2 checks the outbound filter in VLAN 2. As there is no specific outbound filter, the global action is to `permit all` so switch 2 floods the packet in VLAN 2 (except the source port.)
9. Host 2 receives the switch 2 print service announcement.

## Host 2 queries for printers

The following procedure depicts a service process for mDNS queries for a wireless printer in the form of an example.

### Procedure

1. Host 2 sends an mDNS query for printers.
2. There is no inbound filter in VLAN 2 of Switch 3 therefore it receives the query.
3. Switch 3 checks the outbound filter in VLAN 2. As there is no specific outbound filter the default action is **permit all**.
4. Switch 3 floods the query in VLAN 2 (except the source port.)
5. Switch 1 receives the query and check the inbound filters. Permit for the instance name, Host 2, allows the packet on VLAN 2.
6. Switch 1 checks the outbound filter for VLAN 2. As there is no specific filter and global filter is **permit all**, it will flood the packet in VLAN 2 (except the source port.)

7. Switch 1 checks the reflection status. Reflection is enabled on VLAN 2 and VLAN 3. Since VLAN 2 is an incoming VLAN, it will not pass the reflection on VLAN 2.
8. Switch 1 checks the outbound filters on VLAN 3. There is no rule to deny Host two query and the global filter is set to **permit all** so it will forward the packet to VLAN 3.
9. Switch 2 receives the service and checks for any inbound and outbound filters in VLAN 3.
10. There is no specific inbound and outbound filter in VLAN 3 therefore it will flood the query in VLAN 3 (except the source port.)
11. Reflection is not enabled in Switch 2 therefore it will not pass any further reflection.
12. Wireless printer 2 responses to the query and switch 2 does not have any inbound and outbound filters therefore it will flood the response to VLAN 3 (except the source port.)
13. Switch 1 receives the packet as there are no inbound filters in VLAN 3. VLAN 3 has an outbound filter set to deny wireless printer service. The service will not flood VLAN 3.
14. Switch 1 checks the reflection status which is enabled in VLAN 2 and 3. Since the incoming VLAN is 3, the packet will not forward to VLAN 3.
15. Switch 1 checks the outbound filter in VLAN 2. As there is no specific filter, it will forward the response to VLAN 2.
16. Switch 3 receives the response on VLAN 2 as there is no inbound filter to deny this service.
17. Switch 3 does not have any outbound filters in VLAN 2, so it will flood the response in VLAN 2 (except the source port.)
18. Host 2 receives the Wireless Printer 2 service response.

## iPhone 1 queries for printers

The following depicts a service process for iPhone queries for a wireless printer in the form of an example.

1. iPhone 1 sends an mDNS query for printers in switch 1 on VLAN 1.
2. Switch 1 checks the inbound filter in VLAN 1. As there is no specific filters, it receives the query.
3. Switch 1 checks the outbound filter in VLAN 1. As there is no specific filter therefore it flood the packet in VLAN 1 (except the source port.)
4. Switch 1 checks the reflection status. The reflection is enabled on VLAN 2 and 3.
5. Switch 1 checks the outbound filters on VLAN 2 and 3. In VLAN 3 the outbound filter is set to deny wireless printer therefore it will not reflect the packet to VLAN 3. There is no specific outbound filter in VLAN 2 so it will forward the packet to VLAN 2.
6. In switch 1, wireless printer 1 receives the iPhone 1 query and sends a response. Switch 1 checks the inbound filter, outbound filter and floods the response to VLAN 1 (except the source port.)
7. Switch 3 receives the iPhone 1 query and floods the packet in VLAN 2. As there is no specific inbound and outbound filters in switch 3, there is no associated printers in switch 3. There will not be any further response.

## Limitations of the mDNS gateway and Chromecast

The following are limitations of the mDNS gateway and Chromecast features:

- IPv6 is not supported.
- In distributed environment enable gateway in one switch to avoid loops.
- Chromecast v1 (DIAL over SSDP) is not supported.
- Custom filters are not supported. For example:

```
rule <name> service *tv*
rule <name> instance *ipad*
```

- mDNS commands are not available from the web and the menu.

- If the user configures both permit and deny for same service/instance and assign that to same VLAN then it is not valid configuration. System will not behave properly.
- If the user has detected the Chromecast device via a permit profile VLAN and is doing a transition to deny profile, VLAN will need to clean the cache memory. Otherwise the system might get connected with already discovered device. It will not try to discover it again. This is an expected behavior.

## Enabling mDNS feature

This command is supported In the config context with manager permissions.

### Syntax

```
mdns enable
no mdns enable
```

### Description

Enable or disables mDNS gateway support on switch.

The default value is disabled.

## Create mDNS reflection

This command is supported in the config context.

### Syntax

```
mdns gateway vlan VLAN-LIST
no mdns gateway vlan VLAN-LIST
```

### Description

Configures the VLAN reflection for mDNS traffic. If the VLAN is not set, the mDNS traffic will not flood to different subnets, it will only flood to the incoming VLAN.

### Options

#### gateway

Enable VLAN for mDNS gateway.

## Create or delete a mDNS profile

This command will be supported on config context in manager mode. This is a context command. Separate context is created for this.

### Syntax

```
mdns profile PROFILE-NAME
no mdns profile PROFILE-NAME
```

### Description

Create or delete an mDNS profile.

## Set rules for mDNS profile

This command is supported in the mDNS profile context.

## Syntax

```
rule rule-id instance | service NAME action permit | deny  
no rule rule-id instance | service NAME action permit | deny
```

## Description

Sets rules for each mDNS profile. You can configure specific rule to permit or deny the mDNS packet.

## Options

### rule

Create or delete a rule for mDNS profile.

### instance

Instance name of the client.

### service

Service name of the client.

### action

Specify the action for mDNS traffic.

### permit

Permit the packet upon successful match.

### deny

Deny the packet upon successful match.

## Set the specific mDNS profile for VLAN

This command is supported in the mDNS profile context.

## Syntax

```
vlan VLAN-LIST  
no vlan VLAN-LIST
```

## Description

Used to set the mDNS profile for a particular VLAN. Based on the rule, the filter permits or denies traffic.

## Options

*VLAN-LIST*

## Set the global mDNS profile

This command is supported in the configure context in manager mode.

## Syntax

```
mdns default filter in | out action permit | deny
```

## Description

Used to set the default action for all VLANs. If there is no specific rule for a particular VLAN, the default action will be applied. By default, the global action is set to deny for both inbound and outbound traffic.

## Options



**filter**

Specify the mDNS filter on this VLAN.

**in**

Match inbound traffic.

**out**

Match outbound traffic.

**default**

Set the action of the mDNS default filter

## Show mdns

**Syntax**

```
show mdns
```

**Description**

Display the status of the mDNS feature.

**Options****mDNS**

Display the status of the mDNS feature

**Example show mDNS**

```
show mDNS
mDNS Configuration
mDNS: Enabled
```

## Show mDNS gateway

**Syntax**

```
show mdns gateway
```

**Description**

Display the reflection VLAN list of the mDNS gateway.

**Options****gateway**

mDNS gateway

**Example**

```
show mDNS gateway

mDNS Gateway Configuration
Gateway VLAN List: 1-10,12
```

# Show mDNS profile configuration

## Syntax

```
show mdns profile
```

## Description

Display mDNS profile configuration information.

## Options

### profile

mDNS profile information

## Example

```
mDNS profile configuration
Profile Name: Students
VLANs       : 1-3,25

Rules:
ID  Instance          Service          Action
---  -
1   ANY                AppleTV         Deny
2   MyComputer        ANY             Permit

Profile Name: Professors
VLANs       : 3-6,10

Rules:
ID  Instance          Service          Action
---  -
1   ANY                AppleTV         Deny
2   MyComputer        ANY             Permit
```

# Show mDNS profile name

## Syntax

```
show mdns profile PROFILE-NAME
```

## Description

Display mDNS profile name information.

## Options

### PROFILE-NAME

Specify the profile name.

## Example

```
mDNS profile configuration
Profile Name: Students
VLANs       : 1-3,25

Rules:
ID  Instance          Service          Action
---  -
```

1	ANY	AppleTV	Deny
2	MyComputer	ANY	Permit

## Show mDNS

```
mDNS enable
mDNS gateway vlan 1-2
mDNS profile "abcd"
  rule 1 instance Host1 action permit
  rule 2 service AppleTv action deny
  vlan 1-2
  exit

vlan 1
  name "DEFAULT_VLAN"
  untagged 1-24
  ip address dhcp-bootp
  exit

vlan 2
  name "VLAN2"
  untagged 2
  ip address 10.1.1.1 255.255.255.0
  exit
```

## Debug mDNS

### Syntax

```
debug mdns
```

### Description

Enable or disable mDNS debug logging.

### Usage

```
debug mdns
no debug mdns
```

## Validation rules

Rule	Error/Warning/Prompt
Profile name exceeds max length	The profile name exceeds the maximum length of %d.
Profile name already exist. It should be unique.	The profile name already exists.
Profile name contains invalid characters	The profile name contains invalid characters.
Trying to delete mDNS profile which does not exist.	The profile is not found.
Trying to add Profile beyond the max limit.	Cannot add the profile. It reached the maximum limit.

*Table Continued*

Rule	Error/Warning/Prompt
Instance name exceeds	The instance name exceeds the maximum length %d.
Instance name contains invalid characters	The instance name contains invalid characters.
Service name exceeds max length	The service name exceeds the maximum length of %d.
Service name contains invalid characters	The rules for Service Names [RFC6335] state that they may be no more than 15 characters long, consisting of only letters, digits, and hyphens, must begin and end with a letter or digit, must not contain consecutive hyphens, and must contain at least one letter.
Trying to add rule beyond the max limit.	Cannot add rule. It reached the maximum limit.
Trying to add gateway vlan beyond the limit.	Maximum number of mDNS gateway VLANs is %s.
Trying to add profile vlan beyond the limit.	Maximum number of mDNS profile VLANs is %s.
Trying to add rule which is already present.	The rule is already configured with this ID.
Trying to delete rule which is not found.	Rule ID %s is not found.
Trying to show mDNS profile which does not exist.	The profile is not found.
Gateway vlan cannot be configured as secondary vlan	mDNS gateway VLAN cannot be configured on secondary VLAN. It should be configured on the primary VLAN
Profile vlan cannot be configured as secondary vlan	mDNS profile VLAN cannot be configured on secondary VLAN. It should be configured on the primary VLAN.
Secondary vlan cannot be configured as gateway vlan	Secondary VLAN cannot be configured on mDNS gateway VLAN.
Secondary vlan cannot be configured as gateway vlan	Secondary VLAN cannot be configured on mDNS profile VLAN.

## RMON table

RMON event	Details
RMON_mDNS_ENABLED	Proposed Display: I 05/22/13 20:39:20 04633 mDNS: mDNS is enabled.
RMON_mDNS_DISABLED	Proposed Display: I 05/22/13 20:39:20 04633 mDNS: mDNS is disabled.
RMON_mDNS_PKT_MAX_LIMIT	Proposed Display: W 05/22/13 20:49:12 04635 mDNS: mDNS packets are dropped. It has exceeded the maximum limit of %d packets per second.