



Configuring IGMP Snooping

This chapter describes how to configure the Internet Group Management Protocol (IGMP) Snooping on Cisco NX-OS devices for IPv4 networks.

- [About IGMP Snooping, on page 1](#)
- [Prerequisites for IGMP Snooping, on page 3](#)
- [Guidelines and Limitations for IGMP Snooping, on page 4](#)
- [Default Settings, on page 4](#)
- [Configuring IGMP Snooping Parameters, on page 5](#)
- [Verifying the IGMP Snooping Configuration, on page 12](#)
- [Displaying IGMP Snooping Statistics, on page 12](#)
- [Clearing IGMP Snooping Statistics, on page 13](#)
- [Configuration Examples for IGMP Snooping, on page 13](#)

About IGMP Snooping



Note We recommend that you do not disable IGMP snooping on the device. If you disable IGMP snooping, you might see reduced multicast performance because of excessive false flooding within the device.


IGMP snooping software examines Layer 2 IP multicast traffic within a VLAN to discover the ports where interested receivers reside. Using the port information, IGMP snooping can reduce bandwidth consumption in a multi-access LAN environment to avoid flooding the entire VLAN. IGMP snooping tracks which ports are attached to multicast-capable routers to help the routers forward IGMP membership reports. The IGMP snooping software responds to topology change notifications. By default, IGMP snooping is enabled on the device.

This figure shows an IGMP snooping switch that sits between the host and the IGMP router. The IGMP snooping switch snoops the IGMP membership reports and Leave messages and forwards them only when necessary to the connected IGMP routers.



Note Owing to {Vlan,Mac} lookup IGMPv2 reports are flooded/forwarded to the receivers already attached, this results in report-suppression. This is specific to Cisco Nexus 3550-T only.

Figure 1: IGMP Snooping Switch



The IGMP snooping software operates upon IGMPv1, IGMPv2, and IGMPv3 control plane packets where Layer 3 control plane packets are intercepted and influence the Layer 2 forwarding behavior.

The Cisco NX-OS IGMP snooping software has the following proprietary features:

- Layer 2 multicast forwarding is only done based on MAC address on the Cisco Nexus 3550-T hardware.
- Optimized Multicast Flooding (OMF) forwards unknown traffic to only routers and performs no data-driven state creation.

For more information about IGMP snooping, see [RFC 4541](#)

IGMPv1 and IGMPv2

Both IGMPv1 and IGMPv2 support membership report suppression, which means that if two hosts on the same subnet want to receive multicast data for the same group, the host that receives a member report from the other host suppresses sending its report. Membership report suppression occurs for hosts that share a port.

If no more than one host is attached to each VLAN switch port, you can configure the fast leave feature in IGMPv2. The fast leave feature does not send last member query messages to hosts. As soon as the software receives an IGMP leave message, the software stops forwarding multicast data to that port.

IGMPv1 does not provide an explicit IGMP leave message, so the software must rely on the membership message timeout to indicate that no hosts remain that want to receive multicast data for a particular group.



Note The software ignores the configuration of the last member query interval when you enable the fast leave feature because it does not check for remaining hosts.

IGMPv3

The IGMPv3 snooping implementation on Cisco NX-OS supports full IGMPv3 snooping, which provides constrained flooding based on the (S, G) information in the IGMPv3 reports. The source based filtering is not supported for L2 multicast on Cisco NX-OS 3550-T series switches, owing to MAC based multicast forwarding.

By default, the software tracks hosts on each VLAN port. The explicit tracking feature provides a fast leave mechanism. Because every IGMPv3 host sends membership reports, report suppression limits the amount of traffic that the device sends to other multicast-capable routers. When report suppression is enabled, and no

IGMPv1 or IGMPv2 hosts requested the same group, the software provides proxy reporting. The proxy feature builds the group state from membership reports from the downstream hosts and generates membership reports in response to queries from upstream queriers.

Even though the IGMPv3 membership reports provide a full accounting of group members on a LAN segment, when the last host leaves, the software sends a membership query. You can configure the parameter last member query interval. If no host responds before the timeout, the software removes the group state.



Note Owing to cut-through forwarding of L2 multicast based on destination MAC, (S,G) information in the IGMPv3 reports is ignored when PIM is not enabled on VLAN. Cisco Nexus 3550-T series switch does not support PIM on trunk ports.

IGMP Snooping Querier

When PIM is not enabled on an interface because the multicast traffic does not need to be routed, you must configure an IGMP snooping querier to send membership queries. You define the querier in a VLAN that contains multicast sources and receivers but no other active querier.

The querier can be configured to use any IP address in the VLAN.

As a best practice, a unique IP address, one that is not already used by the switch interface or the Hot Standby Router Protocol (HSRP) virtual IP address, should be configured so as to easily reference the querier.



Note The IP address for the querier should not be a broadcast IP address, multicast IP address, or 0 (0.0.0.0).

When an IGMP snooping querier is enabled, it sends out periodic IGMP queries that trigger IGMP report messages from hosts that want to receive IP multicast traffic. IGMP snooping listens to these IGMP reports to establish appropriate forwarding.

The IGMP snooping querier performs querier election as described in RFC 2236. Querier election occurs in the following configurations:

- When there are multiple switch queriers configured with the same subnet on the same VLAN on different switches.
- When the configured switch querier is in the same subnet as with other Layer 3 SVI queriers.

Prerequisites for IGMP Snooping

IGMP snooping has the following prerequisites:

- You are logged onto the device.
- For global commands, you are in the correct virtual routing and forwarding (VRF) mode. The default configuration mode shown in the examples in this chapter applies to the default VRF.

Guidelines and Limitations for IGMP Snooping

IGMP snooping has the following guidelines and limitations:

- Cisco Nexus 3550-T Release 10.2(3t) does not enable PIM on trunk ports
- If there is a host on a network segment sending periodic reports, hosts on other ports suppress IGMPv2 reports resulting in a timeout. If such hosts are present, use static configuration of receivers to prevent host timeouts.
- Cisco Nexus® 3550-T switches support IGMP snooping for IPv4 but do not support MLD snooping for IPv6.
- Layer 3 IPv6 multicast routing is not supported.
- Layer 2 IPv6 multicast packets are flooded on the incoming VLAN.
- Cisco Nexus® 3550-T switch does not flood/forward unknown L2/L3 multicast packets on incoming VLAN. As a result, multicast packets are not sent to OMF ports (all external multicast router ports, either statically configured or dynamically learned).
- You must enable the **ip igmp snooping group-timeout** command when you use the **ip igmp snooping proxy general-queries** command. We recommend that you set it to "never". Otherwise, you might experience multicast packet loss.
- Cisco Nexus® 3550-T switch forwards known multicast packets based on Multicast DestMAC of packets on incoming ports where PIM is dis-abled to provide lower latency. Hence, Cisco Nexus® 3550-T switch IGMPv1/v2 incoming reports are forwarded to known multicast receivers.
- Enable the IGMP query flood parameter to send queries one port at a time.

Default Settings

| Parameters | Default |
|---|----------|
| IGMP snooping | Enabled |
| Explicit tracking | Enabled |
| Fast leave | Disabled |
| Last member query interval | 1 second |
| Snooping querier | Disabled |
| Report suppression | Enabled |
| Link-local groups suppression | Enabled |
| Optimise-multicast-flood | Enabled |
| IGMPv3 report suppression for the entire device | Disabled |

| Parameters | Default |
|------------------------------------|---------|
| IGMPv3 report suppression per VLAN | Enabled |

Configuring IGMP Snooping Parameters



Note If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.



Note You must enable IGMP snooping globally before any other commands take effect.

Configuring Global IGMP Snooping Parameters

To affect the operation of the IGMP snooping process globally, you can configure various optional IGMP snooping parameters.

Notes for IGMP Snooping Parameters

- IGMP Snooping Proxy parameter

To decrease the burden placed on the snooping switch during each IGMP general query (GQ) interval, the Cisco NX-OS software provides a way to decouple the periodic general query behavior of the IGMP snooping switch from the query interval configured on the multicast routers.

You can configure the device to consume IGMP general queries from the multicast router, rather than flooding the general queries to all the switchports. When the device receives a general query, it produces proxy reports for all currently active groups and distributes the proxy reports over the period specified by the MRT that is specified in the router query. At the same time, independent of the periodic general query activity of the multicast router, the device sends an IGMP general query on each port in the VLAN in a round-robin fashion. It cycles through all the interfaces in the VLAN at the rate given by the following formula.

$$\text{Rate} = \{\text{number of interfaces in VLAN}\} * \{\text{configured MRT}\} * \{\text{number of VLANs}\}$$

When queries are run in this mode, the default MRT value is 5,000 milliseconds (5 seconds). For a device that has 500 switchports in a VLAN, it would take 2,500 seconds (40 minutes) to cycle through all the interfaces in the system. This is also true when the device itself is the querier.

This behavior ensures that only one host responds to a general query at a given time, and it keeps the simultaneous reporting rate below the packet-per-second IGMP capability of the device (approximately 3,000 to 4,000 pps).



Note When you use this option, you must change the `ip igmp snooping group-timeout` parameter to a high value or to never time out.

The **ip igmp snooping proxy general-queries [mrt]** command causes the snooping function to proxy reply to general queries from the multicast router while also sending round-robin general queries on each switchport with the specified MRT value. (The default MRT value is 5 seconds.)

- IGMP Snooping Group-timeout parameter

Configuring the group-timeout parameter disables the behavior of an expiring membership based on three missed general queries. Group membership remains on a given switchport until the device receives an explicit IGMP leave on that port.

The **ip igmp snooping group-timeout {timeout | never}** command modifies or disables the behavior of an expiring IGMP snooping group membership after three missed general queries.

- IGMP query flood

On Cisco Nexus 3550-T platform, forwarding of IPv4 multicast packets is based on DMAC which can result in host report suppression for IGMP snooping, as all learned ports for the group receive the IGMPv1 and IGMPv2 reports. Due to report suppression by IGMPv1 and IGMPv2 hosts, the querier times out the host resulting in a drop in multicast traffic to the hosts. To reduce host timeout probability, this configuration enables you to send queries one port at a time



Note

- The changes are applicable only to a Cisco Nexus 3550-T switch acting as a IGMP snooping and querier device.
- Enabling this is required only when PIM is disabled on VLAN.

Procedure

Step 1 configure terminal

Example:

```
switch# configure terminal
switch(config)#
```

Enters global configuration mode.

Step 2 Use the following commands to configure global IGMP snooping parameters.

| Option | Description |
|---|---|
| ip igmp snooping | Enables IGMP snooping for the device. The default is enabled. |
| <code>switch(config)# ip igmp snooping</code> | |

| Option | Description |
|--|---|
| | <p>Note If the global setting is disabled with the no form of this command, IGMP snooping on all VLANs is disabled, whether IGMP snooping is enabled on a VLAN or not. If you disable IGMP snooping, Cisco Nexus® 3550-Tswitch only disables IGMP snoop packet handling. Hence, even with the no form of this command, IGMP packets including multicast packets are not forwarded in hardware.</p> |
| <p>ip igmp snooping event-history</p> <pre>switch(config)# ip igmp snooping event-history</pre> | <p>Configures the size of the event history buffer. The default is small.</p> |
| <p>ip igmp snooping group-timeout {minutes never}</p> <pre>switch(config)# ip igmp snooping group-timeout never</pre> | <p>Configures the group membership timeout value for all VLANs on the device.</p> |
| <p>ip igmp snooping link-local-groups-suppression</p> <pre>switch(config)# ip igmp snooping link-local-groups-suppression</pre> | <p>Configures link-local groups suppression for the entire device. The default is enabled.</p> |
| <p>ip igmp snooping proxy general-inquiries [mrt seconds]</p> <pre>switch(config)# ip igmp snooping proxy general-inquiries</pre> | <p>Configures the IGMP snooping proxy for the device. The default is 5 seconds.</p> |
| <p>ip igmp snooping v3-report-suppression</p> <pre>switch(config)# ip igmp snooping v3-report-suppression</pre> | <p>Limits the membership report traffic sent to multicast-capable routers. When you disable report suppression, all IGMP reports are sent as-is to multicast-capable routers. The default is enabled.</p> |
| <p>ip igmp snooping report-suppression</p> <pre>switch(config)# ip igmp snooping report-suppression</pre> | <p>Configures IGMPv3 report suppression and proxy reporting. The default is disabled.</p> |

| Option | Description |
|--|--|
| [no] ip igmp snooping query flood switch(config)# ip igmp snooping query flood | Configure globally to enable feature. Use the no form of the command to enable sending queries one port at a time. Note Global setting applies only on VLAN where per VLAN setting is also enabled. |

Step 3 copy running-config startup-config**Example:**

```
switch(config)# copy running-config startup-config
```

(Optional) Copies the running configuration to the startup configuration.

Configuring IGMP Snooping Parameters per VLAN

To affect the operation of the IGMP snooping process per VLAN, you can configure various optional IGMP snooping parameters.



Note You configure the IGMP snooping parameters that you want by using this configuration mode; however, the configurations apply only after you specifically create the specified VLAN. See the *Cisco Nexus 3550-T Series NX-OS Layer 2 Switching Configuration Guide* for information on creating VLANs.

Procedure

Step 1 configure terminal**Example:**

```
switch# configure terminal
switch(config)#
```

Enters global configuration mode.

Step 2 ip igmp snooping**Example:**

```
switch(config)# ip igmp snooping
```

Enables IGMP snooping. The default is enabled.

Note If the global setting is disabled with the **no** form of this command, IGMP snooping on all VLANs is disabled, whether IGMP snooping is enabled on a VLAN or not. If you disable IGMP snooping, Layer 2 multicast frames flood to all modules.

Step 3 **vlan configuration** *vlan-id***Example:**

```
switch(config)# vlan configuration 2
switch(config-vlan-config)#
```

Configures the IGMP snooping parameters you want for the VLAN. These configurations do not apply until you create the specified VLAN.

Step 4 Use the following commands to configure IGMP snooping parameters per VLAN.

| Option | Description |
|--|--|
| <p>ip igmp snooping</p> <pre>switch(config-vlan-config)# ip igmp snooping</pre> | <p>Enables IGMP snooping for the current VLAN. The default is enabled.</p> <p>Note Cisco Nexus 3550-T switches can flood layer 2 multicast packets only to multicast router and IGMP querier ports. This behaviour is not modified with no form of the ip igmp snooping command.</p> |
| <p>ip igmp snooping access-group {prefix-list route-map} <i>policy-name</i> interface <i>interface slot/port</i></p> <pre>switch(config-vlan-config)# ip igmp snooping access-group prefix-list plist interface ethernet 1/2</pre> | <p>Configures a filter for IGMP snooping reports that is based on a prefix-list or route-map policy. The default is disabled.</p> |
| <p>ip igmp snooping explicit-tracking</p> <pre>switch(config-vlan-config)# ip igmp snooping explicit-tracking</pre> | <p>Tracks IGMPv3 membership reports from individual hosts for each port on a per-VLAN basis. The default is enabled on all VLANs.</p> |
| <p>ip igmp snooping fast-leave</p> <pre>switch(config-vlan-config)# ip igmp snooping fast-leave</pre> | <p>Supports IGMPv2 hosts that cannot be explicitly tracked because of the host report suppression mechanism of the IGMPv2 protocol. When you enable fast leave, the IGMP software assumes that no more than one host is present on each VLAN port. The default is disabled for all VLANs.</p> |
| <p>ip igmp snooping group-timeout {<i>minutes</i> never}</p> <pre>switch(config-vlan-config)# ip igmp snooping group-timeout never</pre> | <p>Configures the group membership timeout for the specified VLANs.</p> |
| <p>ip igmp snooping last-member-query-interval <i>seconds</i></p> | <p>Removes the group from the associated VLAN port if no hosts respond to an IGMP query message before the last member query interval expires. Values range from 1 to 25 seconds. The default is 1 second.</p> |

| Option | Description |
|---|---|
| <pre>switch(config-vlan-config)# ip igmp snooping last-member-query-interval 3</pre> | |
| <p>ip igmp snooping proxy general-queries [mrt <i>seconds</i>]</p> <pre>switch(config-vlan-config)# ip igmp snooping proxy general-queries</pre> | Configures an IGMP snooping proxy for specified VLANs. The default is 5 seconds. |
| <p>ip igmp snooping querier <i>ip-address</i></p> <pre>switch(config-vlan-config)# ip igmp snooping querier 172.20.52.106</pre> | Configures a snooping querier when you do not enable PIM because multicast traffic does not need to be routed. The IP address is used as the source in messages. |
| <p>ip igmp snooping querier-timeout <i>seconds</i></p> <pre>switch(config-vlan-config)# ip igmp snooping querier-timeout 300</pre> | Configures a snooping querier timeout value for IGMPv2 when you do not enable PIM because multicast traffic does not need to be routed. The default is 255 seconds. |
| <p>ip igmp snooping query-interval <i>seconds</i></p> <pre>switch(config-vlan-config)# ip igmp snooping query-interval 120</pre> | Configures a snooping query interval when you do not enable PIM because multicast traffic does not need to be routed. The default value is 125 seconds. |
| <p>ip igmp snooping query-max-response-time <i>seconds</i></p> <pre>switch(config-vlan-config)# ip igmp snooping query-max-response-time 12</pre> | Configures a snooping MRT for query messages when you do not enable PIM because multicast traffic does not need to be routed. The default value is 10 seconds. |
| <p>ip igmp snooping report-policy {prefix-list route-map} <i>policy-name</i> interface <i>interface slot/port</i></p> <pre>switch(config-vlan-config)# ip igmp snooping report-policy route-map rmap interface ethernet 1/4</pre> | Configures a filter for IGMP snooping reports that is based on a prefix-list or route-map policy. The default is disabled. |
| <p>ip igmp snooping startup-query-count <i>value</i></p> <pre>switch(config-vlan-config)# ip igmp snooping startup-query-count 5</pre> | Configures snooping for a number of queries sent at startup when you do not enable PIM because multicast traffic does not need to be routed. |

| Option | Description |
|--|---|
| <p>ip igmp snooping startup-query-interval <i>seconds</i></p> <pre>switch(config-vlan-config)# ip igmp snooping startup-query-interval 15000</pre> | Configures a snooping query interval at startup when you do not enable PIM because multicast traffic does not need to be routed. |
| <p>ip igmp snooping robustness-variable <i>value</i></p> <pre>switch(config-vlan-config)# ip igmp snooping robustness-variable 5</pre> | Configures the robustness value for the specified VLANs. The default value is 2. |
| <p>ip igmp snooping report-suppression</p> <pre>switch(config-vlan-config)# ip igmp snooping report-suppression</pre> | Limits the membership report traffic sent to multicast-capable routers. When you disable report suppression, all IGMP reports are sent as-is to multicast-capable routers. The default is enabled. |
| <p>ip igmp snooping mrouter interface <i>interface</i></p> <pre>switch(config-vlan-config)# ip igmp snooping mrouter interface ethernet 1/1</pre> | Configures a static connection to a multicast router. The interface to the router must be in the selected VLAN. You can specify the interface by the type and the number, such as ethernet slot/port . |
| <p>ip igmp snooping static-group <i>group-ip-addr [source source-ip-addr]</i> interface <i>interface</i></p> <pre>switch(config-vlan-config)# ip igmp snooping static-group 230.0.0.1 interface ethernet 1/1</pre> | Configures the Layer 2 port of a VLAN as a static member of a multicast group. You can specify the interface by the type and the number, such as ethernet slot/port . |
| <p>ip igmp snooping link-local-groups-suppression</p> <pre>switch(config-vlan-config)# ip igmp snooping link-local-groups-suppression</pre> | Configures link-local groups suppression for the specified VLANs. The default is enabled. |
| <p>ip igmp snooping v3-report-suppression</p> <pre>switch(config-vlan-config)# ip igmp snooping v3-report-suppression</pre> | Configures IGMPv3 report suppression and proxy reporting for the specified VLANs. The default is enabled per VLAN. |
| <p>ip igmp snooping version <i>value</i></p> | Configures the IGMP version number for the specified VLANs. |

| Option | Description |
|---|--|
| <pre>switch(config-vlan-config)# ip igmp snooping version 2</pre> | |
| <pre>[no] ip igmp snooping query flood switch(config-vlan-config)# ip igmp snooping query flood</pre> | <p>Configures per port query feature, default behaviour is to flood the queries on all ports together. Use the no form of the command to enable sending queries one port at a time.</p> <p>Note This setting is applied only when the feature is enabled globally.</p> |

Step 5 copy running-config startup-config

Example:

```
switch(config)# copy running-config startup-config
```

(Optional) Copies the running configuration to the startup configuration.

Verifying the IGMP Snooping Configuration

| Command | Description |
|---|---|
| show ip igmp snooping [vlan <i>vlan-id</i>] | Displays the IGMP snooping configuration by VLAN. |
| show ip igmp snooping groups [<i>source [group] group [source]</i>] [vlan <i>vlan-id</i>] [detail] | Displays IGMP snooping information about groups by VLAN. |
| show ip igmp snooping querier [vlan <i>vlan-id</i>] | Displays IGMP snooping queriers by VLAN. |
| show ip igmp snooping mroute [vlan <i>vlan-id</i>] | Displays multicast router ports by VLAN. |
| show ip igmp snooping explicit-tracking [vlan <i>vlan-id</i>] [detail] | Displays IGMP snooping explicit tracking information by VLAN. |

Displaying IGMP Snooping Statistics

You can display the IGMP snooping statistics using these commands.

| Command | Description |
|--|--|
| show ip igmp snooping statistics vlan | Displays IGMP snooping statistics. You can see the virtual port channel (vPC) statistics in this output. |

| Command | Description |
|--|--|
| <code>show ip igmp snooping {report-policy access-group} statistics [vlan vlan]</code> | Displays detailed statistics per VLAN when IGMP snooping filters are configured. |

Clearing IGMP Snooping Statistics

You can clear the IGMP snooping statistics using these commands.

| Command | Description |
|---|---|
| <code>clear ip igmp snooping statistics vlan</code> | Clears the IGMP snooping statistics. |
| <code>clear ip igmp snooping {report-policy access-group} statistics [vlan vlan]</code> | Clears the IGMP snooping filter statistics. |

Configuration Examples for IGMP Snooping



Note The configurations in this section apply only after you create the specified VLAN. See the *Cisco Nexus 3550-T Layer 2 Switching Configuration* section for information on creating VLANs.

The following example shows how to configure the IGMP snooping parameters:

```

config t
 ip igmp snooping
 vlan configuration 2
 ip igmp snooping
 ip igmp snooping explicit-tracking
 ip igmp snooping fast-leave
 ip igmp snooping last-member-query-interval 3
 ip igmp snooping querier 172.20.52.106
 ip igmp snooping report-suppression
 ip igmp snooping mrouter interface ethernet 1/1
 ip igmp snooping static-group 230.0.0.1 interface ethernet 1/1
 ip igmp snooping link-local-groups-suppression
 ip igmp snooping v3-report-suppression

```

The following example shows how to configure prefix lists and use them to filter IGMP snooping reports:

```

ip prefix-list plist seq 5 permit 224.1.1.1/32
ip prefix-list plist seq 10 permit 224.1.1.2/32
ip prefix-list plist seq 15 deny 224.1.1.3/32
ip prefix-list plist seq 20 deny 225.0.0.0/8 eq 32

vlan configuration 2
 ip igmp snooping report-policy prefix-list plist interface Ethernet 1/2
 ip igmp snooping report-policy prefix-list plist interface Ethernet 1/3

```

In the above example, the prefix-list permits 224.1.1.1 and 224.1.1.2 but rejects 224.1.1.3 and all the groups in the 225.0.0.0/8 range. The prefix-list is an implicit "deny" if there is no match. If you wish to permit everything else, add **ip prefix-list plist seq 30 permit 224.0.0.0/4 eq 32**.

The following example shows how to configure route maps and use them to filter IGMP snooping reports:

```
route-map rmap permit 10
  match ip multicast group 224.1.1.1/32
route-map rmap permit 20
  match ip multicast group 224.1.1.2/32
route-map rmap deny 30
  match ip multicast group 224.1.1.3/32
route-map rmap deny 40
  match ip multicast group 225.0.0.0/8

vlan configuration 2
  ip igmp snooping report-policy route-map rmap interface Ethernet 1/4
  ip igmp snooping report-policy route-map rmap interface Ethernet 1/5
```

In the above example, the route-map permits 224.1.1.1 and 224.1.1.2 but rejects 224.1.1.3 and all the groups in the 225.0.0.0/8 range. The route-map is an implicit "deny" if there is no match. If you wish to permit everything else, add **route-map rmap permit 50 match ip multicast group 224.0.0.0/4**.