



Multicast IGMP Snooping

This chapter describes how to identify and resolve problems that relate to multicast Internet Group Management Protocol (IGMP) snooping.

Information About Multicast IGMP Snooping

IP multicast is a method of forwarding the same set of IP packets to a number of hosts within a network. You can use multicast in both IPv4 and IPv6 networks to provide efficient delivery of data to multiple destinations.

Multicast involves both a method of delivery and discovery of senders and receivers of multicast data, which is transmitted on IP multicast addresses called groups. A multicast address that includes a group and source IP address is often referred to as a channel.

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. The IGMP snooping feature 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.

IGMP snooping works as follows:

- Ethernet switches, such as Catalyst 6500 series switches, parse and intercept all IGMP packets and forward them to a CPU, such as a supervisor module, for protocol processing.
- Router ports are learned by using IGMP queries. The switch returns IGMP queries; it remembers which port the query comes from and marks the port as a router port.
- IGMP membership is learned by using IGMP reports. The switch parses IGMP report packets and updates its multicast forwarding table to keep track of IGMP membership.
- When the switch receives multicast traffic, it checks its multicast table and forwards the traffic only to those ports interested in the traffic.
- IGMP queries are flooded to the whole VLAN.
- IGMP reports are forwarded to the uplink port (the router ports).
- Multicast data traffic is forwarded to uplink ports (the router ports).

The operation of multicast IGMP snooping depends on the correct configuration of the upstream switch. Because the IGMP process needs to know which upstream port connects to the router that supports IGMP routing, you must turn on IP multicast routing on the upstream switch by entering the **ip multicast-routing** command.

This example shows how to turn on global multicast routing, configure an SVI interface, and turn on the PIM routing protocol:

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# ip multicast-routing
switch(config)# end

switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# int vlan159
switch(config-if)# ip pim dense-mode
switch(config-if)# end
```

Troubleshooting Guidelines

Follow these guidelines when troubleshooting multicast IGMP issues:

- Verify that IGMP snooping is enabled by entering the **show ip igmp snooping** command.
- Make sure the upstream switch has IGMP configured.
- Verify that the Cisco Nexus 1000V switch is configured correctly and is ready to forward multicast traffic by entering the **show ip igmp snooping groups** command. In the displayed output of the command, look for the letter R under the port heading. The R indicates that the VSM has learned the uplink router port from the IGMP query that was sent by the upstream switch, and means that the Cisco Nexus 1000V is ready to forward multicast traffic.

Multicast IGMP Snooping Troubleshooting Commands

You can use the commands in this section to troubleshoot problems related to multicast IGMP snooping.

Command	Purpose
show cdp neighbor	Displays Cisco Discovery Protocol (CDP) neighbors. However, if you have disabled CDP on the upstream switch by entering the no cdp enable command, the show cdp neighbor command does not display any information. See Example 11-1 on page 11-2 .
show ip igmp snooping groups	Displays if IGMP snooping is enabled on the VLAN. See Example 11-2 on page 11-3 .
show ip igmp snooping groups	Displays snooping information for the group addresses.
show ip igmp snooping vlan	Generates additional logs to debug IGMP snooping events on all VLANs. See Example 11-3 on page 11-3 .

Example 11-1 show cdp neighbor command

```
n1000V# show cdp neighbor
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute
```

Device ID	Local Infrfce	Hldtme	Capability	Platform	Port ID
n1000V	Eth3/2	179	R S I	WS-C6506-E	Gig5/16
n1000V	Eth3/4	179	R S I	WS-C6506-E	Gig5/23

Example 11-2 show ip igmp snooping vlan command

```
n1000V# show ip igmp snooping vlan 159
IGMP Snooping information for vlan 159
IGMP snooping enabled <-- IGMP SNOOPING is enabled for vlan 159
Optimised Multicast Flood (OMF) enabled
IGMP querier none
Switch-querier disabled
IGMPv3 Explicit tracking enabled (initializing, time-left: 00:03:20)
IGMPv2 Fast leave disabled
IGMPv1/v2 Report suppression enabled
IGMPv3 Report suppression disabled
Link Local Groups suppression enabled
Router port detection using PIM Hellos, IGMP Queries
Number of router-ports: 0
Number of groups: 0
VLAN vPC function disabled
Active ports:
```

Example 11-3 debug ip igmp snooping vlan command

```
n1000V(config)# debug ip igmp snooping vlan
2008 Sep  2 13:29:36.125661 igmp: SNOOP: <vlan 159> Process a valid IGMP packet
2008 Sep  2 13:29:36.126005 igmp: SNOOP: <vlan 159> Received v2 report: group 224.0.0.251
fro 7.159.159.54 on Vethernet3
2008 Sep  2 13:29:36.126086 igmp: SNOOP: <vlan 159> Added oif Vethernet3 for (*,
224.0.0.251) entry
2008 Sep  2 13:29:36.126157 igmp: SNOOP: <vlan 159> Forwarding report for (*, 224.0.0.251)
came on Vethernet3
2008 Sep  2 13:29:36.126225 igmp: SNOOP: <vlan 159> Forwarding the packet to router-ports
2008 Sep  2 13:29:36.126323 igmp: SNOOP: <vlan 159> Forwarding packet to router-port
Ethernet3/6 (iod 42)
```

On the VSM, use the following command:

- **module vem module-number execute vemcmd show vlan**

In [Example 11-4](#), the output shows that LTL 18 corresponds to vmnic3, and LTL 47 corresponds to VM fedora8, interface eth0.

The multicast group table for 224.1.2.3, shows the interfaces the VEM forwards to when it receives multicast traffic for group 224.1.2.3. If fedora8 has multicast group 224.1.2.3 on its eth0 interface, LTL 47 should be in the multicast group table for 224.1.2.3.

LTL 18 is also in multicast group 224.1.2.3, which means that it is a VM and generates multicast traffic to 224.1.2.3. The traffic is forwarded to vmnic3, which is the uplink to the upstream switch.

The multicast group table entry for 0.0.0.0 serves as a default route. If any multicast group traffic does not match any of the multicast groups, the address uses the default route, which means that the traffic is forwarded to an upstream switch through vmnic3.

Example 11-4 module vem module-number execute vemcmd show vlan Command

```
n1000V# module vem 3 execute vemcmd show vlan 159
BD 159, vdc 1, vlan 159, 3 ports
Portlist:
  18  vmnic3
  47  fedora8.eth0
```

```

Multicast Group Table:
Group 224.1.2.3 RID 1 Multicast LTL 4408
    47
    18
Group 0.0.0.0 RID 2 Multicast LTL 4407
    18

```

Problems with Multicast IGMP Snooping

The following are symptoms and solutions for problems with multicast IGMP snooping.

Symptom	Solution
A VM is interested in multicast traffic but is not receiving the multicast traffic	Determine if IGMP snooping is working as expected by entering the debug ip igmp snooping vlan command. Examine the output to see if the port is receiving the IGMP report and if the interface has been added to the multicast traffic interface list for the VM.
	Verify that the multicast distribution table in the VEM has the correct information by entering the module vem module-number execute vemcmd show vlan command.
	View the port table by entering the module vem module-number execute vemcmd show port command. Make sure that the table has the correct information and that the state of the trunk port and the access port is UP/UP.