

IGMP Snooping on EX-series Switches Overview

Internet Group Management Protocol (IGMP) snooping regulates multicast traffic in a switched network. With IGMP snooping enabled, a LAN switch monitors the IGMP transmissions between a host (a network device) and a multicast router, keeping track of the multicast groups and associated member ports. The switch uses that information to make intelligent multicast-forwarding decisions and forward traffic to the intended destination interfaces. EX-series switches support IGMPv1 and IGMPv2.

For details on IGMPv1 and IGMPv2, see the following standards:

- For IGMPv1, see RFC 1112, *Host extensions for IP multicasting* at <http://www.faqs.org/rfcs/rfc1112.html>
- For IGMPv2, see RFC 2236, *Internet Group Management Protocol, Version 2* at <http://www.faqs.org/rfcs/rfc2236.html>

This IGMP snooping topic covers:

- How IGMP Snooping Works on page 1
- How IGMP Snooping Works with Routed VLAN Interfaces on page 2
- How Hosts Join and Leave Multicast Groups on page 4

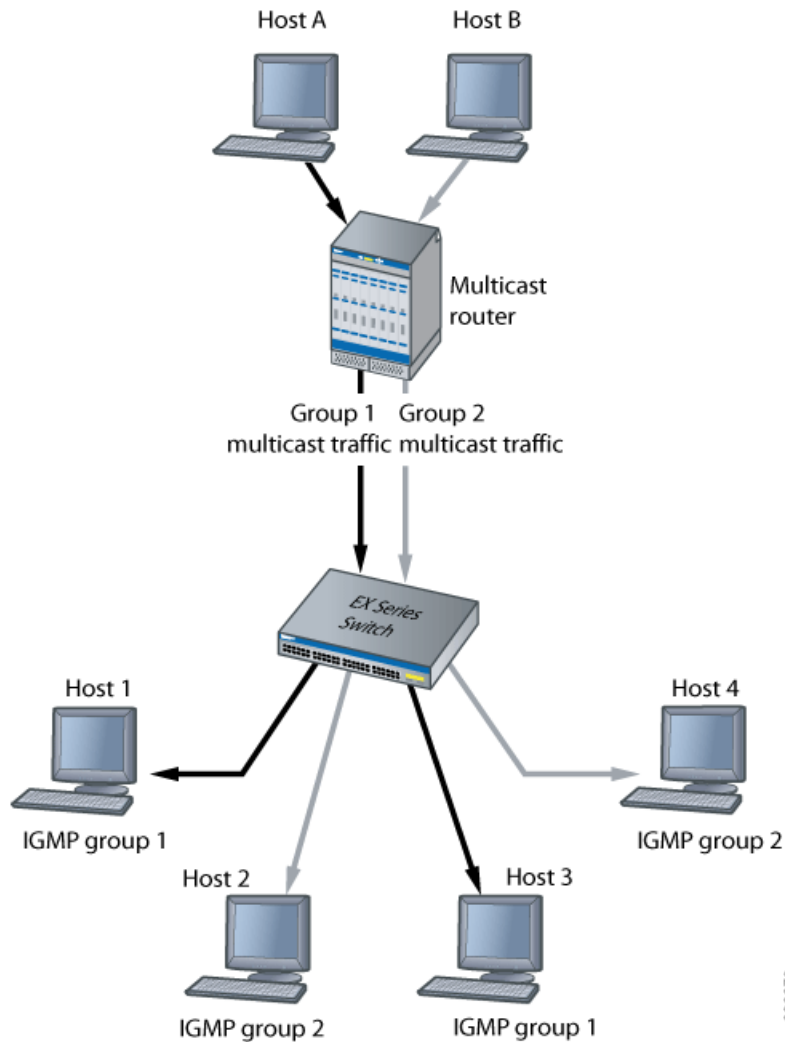
How IGMP Snooping Works

An EX-series switch usually learns *unicast* MAC addresses by checking the source address field of the frames it receives. However, a *multicast* MAC address can never be the source address for a packet. As a result, the switch floods multicast traffic on the VLAN, consuming significant amounts of bandwidth.

IGMP snooping regulates multicast traffic on a VLAN to avoid flooding. When IGMP snooping is enabled, the switch intercepts IGMP packets and uses the content of the packets to build a multicast cache table. The cache table is a database of multicast groups and their corresponding member ports. The cache table is then used to regulate multicast traffic on the VLAN.

When the switch receives multicast packets, it uses the cache table to selectively forward the packets only to the ports that are members of the destination multicast group. Figure 1 shows an example of IGMP traffic flow with IGMP snooping enabled.

Figure 1: IGMP Traffic Flow with IGMP Snooping Enabled



How IGMP Snooping Works with Routed VLAN Interfaces

Switches send traffic to hosts that are part of the same broadcast domain, but routers are needed to route traffic from one broadcast domain to another. EX-series switches use a routed VLAN interface (RVI) to perform these routing functions. IGMP snooping works with Layer 2 interfaces and RVIs to regulate multicast traffic in a switched network.

When an EX-series switch receives a multicast packet, the Packet Forwarding Engines in the switch perform an IP multicast lookup on the multicast packet to determine how to forward the packet to its local ports. From the results of the IP multicast lookup, each Packet Forwarding Engine extracts a list of Layer 3 interfaces (which can include VLAN interfaces) that have ports local to the Packet Forwarding Engine. If an RVI is part of this list, the switch provides a bridge multicast group ID for each RVI to the Packet Forwarding Engine.


A bridge multicast ID is assigned to direct Layer 3 interfaces and to RVIs. For VLANs that include multicast receivers, the bridge multicast ID includes a sub-next-hop ID. The sub-next-hop ID identifies the multicast Layer 2 interfaces in that VLAN that are interested in receiving the multicast stream. The switch ultimately assigns a next-hop after it does a route lookup. The next-hop includes all direct Layer 3 interfaces and RVIs. The Packet Forwarding Engine then forwards multicast traffic to the bridge multicast ID that includes all Layer 3 interfaces and RVIs that are multicast receivers for a given multicast group.

Figure 2 shows how multicast traffic is forwarded on a multilayer switch. In this illustration, multicast traffic is coming in through the xe-0/1/0.0 interface. A multicast group has been formed by the Layer 3 interface ge-0/0/2.0, vlan.0 and vlan.1. The ge-2/0/0.0 interface is a common trunk interface that belongs to both vlan.0 and vlan.1. The letter “R” next to an interface name in the illustration indicates that a multicast receiver host is associated with that interface.



NOTE: Traffic sent to an access interface is untagged; traffic sent to a trunk interface is tagged. For more information on VLAN tagging, see Understanding Bridging and VLANs on EX Series Switches.

Figure 2: IGMP Traffic Flow with Routed VLAN Interfaces



The following table shows the bridge multicast IDs and next-hops that are created. The term subnh refers to a sub-next-hop. The Packet Forwarding Engine will forward multicast traffic to bridge multicast ID9.

ID Number	Type of Next-Hop	Next Hop	Tag Information
ID1	RHN_UNICAST	ge-0/0/0.0	tag = off

ID Number	Type of Next-Hop	Next Hop	Tag Information
ID2	RHN_UNICAST	ge-2/0/0.0	tag = on
ID3	RHN_FLOOD	[ID1, ID2]	
ID4	RHN_UNICAST	ge-0/0/1.0	tag = off
ID5	RHN_FLOOD	[ID4, ID2]	
ID6	RHN_UNICAST	vlan.0	subnh = ID3
ID7	RHN_UNICAST	VLAN.1	subnh = ID5
ID8	RHN_UNICAST	ge-0/0/2.0	
ID9	RHN_FLOOD	[ID6, ID7, ID8]	

How Hosts Join and Leave Multicast Groups

Hosts can join multicast groups in either of two ways:

- By sending an unsolicited IGMP join message to a multicast router that specifies the IP multicast that the host is attempting to join.
- By sending an IGMP join message in response to a general query from a multicast router.

A multicast router continues to forward multicast traffic to a VLAN provided that at least one host on that VLAN responds to the periodic general IGMP queries. For a host to remain a member of a multicast group, therefore, it must continue to respond to the periodic general IGMP queries.

To leave a multicast group, a host can either not respond to the periodic general IGMP queries, which results in a “silent leave” (the only leave option for hosts connected to switches running IGMPv1), or send a group-specific IGMPv2 leave message.

- Related Topics**
- Example: Configuring IGMP Snooping on EX Series Switches
 - Configuring IGMP Snooping (CLI Procedure)
 - RFC 3171, *IANA Guidelines for IPv4 Multicast Address Assignments* at <http://tools.ietf.org/html/rfc3171>

Published: 2009-09-14