Cisco Nexus 1000V InterCloud Layer 2 Switching Configuration Guide, Release 5.2(1)IC1(1.1)

First Published: June 28, 2013

Last Modified: October 10, 2013

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Text Part Number: OL-29147-01

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED “AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

CONTENTS

Preface

This preface contains the following sections:

Audience

This publication is for network administrators who configure and maintain Cisco Nexus devices.

This guide is for network and server administrators with the following experience and knowledge:

- An understanding of virtualization
- Using Virtual Machine Manager (VMM) software to create a virtual machine and configure a VMware vSwitch
- Ability to create an account on provider cloud such as Amazon Web Services (AWS).
- Knowledge of VMware vNetwork Distributed Switch is not required.

Document Conventions

Command descriptions use the following conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
Italic	Italic text indicates arguments for which the user supplies the values.

Convention	Description
[x]	Square brackets enclose an optional element (keyword or argument).
[x | y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.
{x | y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.
[x {y | z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.
`variable`	Indicates a variable for which you supply values, in context where italics cannot be used.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Examples use the following conventions:

Convention	Description
`screen font`	Terminal sessions and information the switch displays are in screen font.
`boldface screen font`	Information you must enter is in boldface screen font.
italic screen font	Arguments for which you supply values are in italic screen font.
< >	Nonprinting characters, such as passwords, are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:

Note Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the manual.

Caution Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Related Documentation for Cisco Nexus 1000V InterCloud

This section lists the documents used with the Cisco Nexus 1000V InterCloud and available on Cisco.com at the following URL:

http://www.cisco.com/en/US/partner/products/ps12904/tsd_products_support_series_home.html

General Information

Cisco Nexus 1000V InterCloud Release Notes

Install and Upgrade

Cisco Nexus 1000V InterCloud Installation Guide

Configuration Guides

Cisco Nexus 1000V InterCloud License Configuration Guide

Cisco Nexus 1000V InterCloud High Availability and Redundancy Configuration Guide

Cisco Nexus 1000V InterCloud Interface Configuration Guide

Cisco Nexus 1000V InterCloud Layer 2 Configuration Guide

Cisco Nexus 1000V InterCloud Port Profile Configuration Guide

Cisco Nexus 1000V InterCloud Security Configuration Guide

Cisco Nexus 1000V InterCloud System Management Configuration Guide

Reference Guides

Cisco Nexus 1000V InterCloud Command Reference

Cisco Nexus 1000V InterCloud Verified Scalability Reference

Cisco Nexus 1000V MIB Quick Reference

Troubleshooting and Alerts

Cisco Nexus 1000V Password Recovery Procedure

Cisco Nexus 1000V Documentation

Cisco Nexus 1000V for VMware vSphere Documentation

http://www.cisco.com/en/US/products/ps9902/tsd_products_support_series_home.html

Cisco Prime Network Services Controller Documentation

http://www.cisco.com/en/US/products/ps13213/tsd_products_support_series_home.html

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to nexus1k-docfeedback@cisco.com. We appreciate your feedback.

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html

Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.

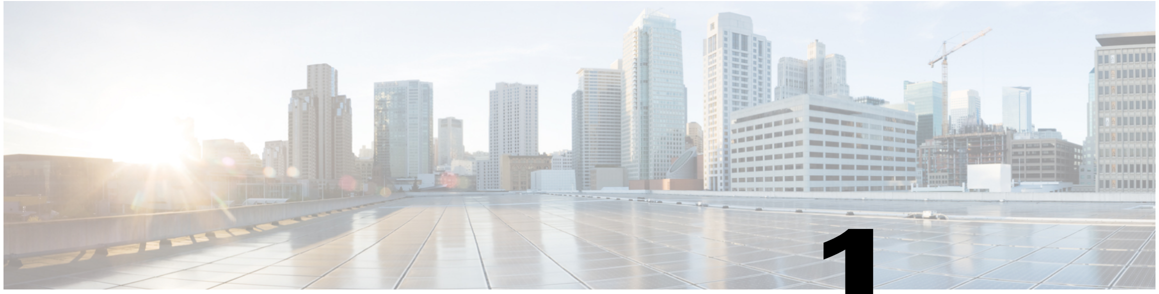

CHAPTER 1

Overview

This chapter contains the following sections:

Information about Layer 2 Switching

VEM Port Model

The Cisco Nexus 1000V InterCloud differentiates the following Virtual Ethernet Module (VEM) ports:

- VEM Virtual Ports

Virtual Ethernet Ports

A virtual Ethernet port (vEth) represents a port on the Cisco Nexus 1000V Distributed Virtual Switch. The Cisco Nexus 1000V has a flat space of vEth ports, 0...n. These vEth ports are what the virtual cable plugs into and are moved to the host that the VM is running on. Virtual Ethernet ports are assigned to port groups.

InterCloud Extender and InterCloud Switch

Each VEM that is attached to the VSM forwards traffic to and from theInterCloud Extender and InterCloud Switch as an independent and intelligent line card. Each VLAN uses its forwarding table to learn and store MAC addresses for ports that are connected to the VEM.

VSM Port Model

The Cisco Nexus 1000V InterCloud VSM is a Virtual machine vNic on the InterCloud Switch and InterCloud Extender.

MAC Address Tables

To switch frames between LAN ports efficiently, a MAC address table is maintained. The MAC address of the sending network is associated with the LAN port on which it was received.

VLANs

A VLAN is a switched network that is logically segmented by function, project team, or application, without regard to the physical locations of the users. VLANs have the same attributes of physical LANs, but you can group end stations even if they are not physically located on the same LAN segment.

Any switchport can belong to a VLAN, and unicast, broadcast, and multicast packets are forwarded and flooded only to end stations in that VLAN. Each VLAN is considered a logical network, and packets destined for stations that do not belong to the VLAN must be forwarded through a bridge or a router.

All ports, including the management port, are assigned to the default VLAN (VLAN1) when the device first comes up.

A total number of 128 VLANs are supported, and the valid range for VLANs is 1- 4094

These VLANs are organized into several ranges for different uses. Some of these VLANs are reserved for internal use by the device and are not available for configuration.

Note Inter-Switch Link (ISL) trunking is not supported on the Cisco Nexus 1000V.

IGMP Snooping

The Internet Group Management Protocol (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. By default, IGMP snooping is enabled on the device.

CHAPTER 2

Configuring MAC Address Tables

This chapter contains the following sections:

Information About MAC Address Tables

Layer 2 ports correlate the MAC address on a packet with the Layer 2 port information for that packet using the MAC address table. A MAC address table is built using the MAC source addresses of the frames received. When a frame is received for a MAC destination address not listed in the address table, the frame is flooded to all LAN ports of the same VLAN with the exception of the port that received the frame. When the destination station replies, the relevant MAC source addresses and port IDs are added to the address table. Subsequent frames are forwarded to a single LAN port without flooding all LAN ports.

You can configure MAC addresses, which are called static MAC addresses, to statically point to specified interfaces on the device. These static MAC addresses override any dynamically learned MAC addresses on those interfaces. You cannot configure broadcast or multicast addresses as static MAC addresses. The static MAC entries are retained across reboots.

The address table per VEM can store up to 32,000 MAC entries. An aging timer triggers removal of addresses from the table when they remain inactive for the default time of 300 seconds. The aging timer can be configured on a global basis but not per VLAN.

You can configure the length of time an entry remains in the MAC address table, clear the table, and so forth.

Guidelines and Limitations

- The forwarding table for each VLAN in a VEM can store up to 4094 MAC addresses.

- You can configure only 1024 static MAC addresses on a single interface.
- Cisco Nexus 1000V InterCloud supports a maximum of 2000 user configured static MAC addresses on a VSM

Default Settings

Table 1: Default MAC Address Aging Time

Parameters	Default
Aging time	300 seconds

Configuring the MAC Address Table

Configuring a Static MAC Address

Use this procedure to configure a MAC address to statically point to a specific interface.

Before You Begin

- You are logged in to the CLI in EXEC mode.
- You cannot configure broadcast or multicast addresses as static MAC addresses.
- Static MAC addresses override dynamically-learned MAC addresses on an interface.

Note Be aware that the Cisco NX-OS commands may differ from those used in Cisco IOS.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **mac address-table static** *mac_address* **vlan** *vlan-id* {[**drop** | **interface** { **type** *if_id* }]}	Adds a static MAC address in the Layer 2 MAC address table and saves it in the running configuration. Interface that can be specified is veth *number*
Step 3	switch(config)# **show mac address static interface** [**type** *if_id*]	(Optional) Displays static MAC addresses.

	Command or Action	Purpose
Step 4	switch(config)# **copy running-config startup-config**	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

```
switch# configure terminal
switch(config)# mac address static
switch(config)# show mac address static
switch(config)#
```

Configuring the Aging Time

Use this procedure to configure the amount of time that packet source MAC addresses, and the ports on which they are learned, remain in the MAC table containing the Layer 2 information.

Note The aging time is a global setting that cannot be configured per VLAN. Although it is a global setting, you can also configure MAC aging time in interface configuration mode or VLAN configuration mode.

Before You Begin

You are logged in to the CLI in EXEC mode.

Note Be aware that the Cisco NX-OS commands may differ from those used in Cisco IOS.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch# **mac address-table aging-time** *seconds*	Specifies and saves in the running configuration the amount of time that will elapse before an entry in the Layer 2 MAC address table is discarded. Allowable entries include: • 120 to 918000 seconds (default is 300) • If you specify zero (0), MAC aging is disabled.

```
switch# configure terminal
switch(config)# mac address-table aging-time 600
switch(config)# show mac address-table aging-time
Vlan  Aging Time
----- ----------
```

```
101   600
100   600
1     600
switch#
```

Clearing Dynamic Addresses from the MAC Address Table

Before You Begin

You are logged in to the CLI in EXEC mode.

Note Be aware that the Cisco NX-OS commands may differ from those used in Cisco IOS.

Procedure

	Command or Action	Purpose
Step 1	switch# **clear mac address-table dynamic** [**vlan** *vlan_id*]	Clears the dynamic address entries from the Layer 2 MAC address table.
Step 2	switch# **show mac address-table**	(Optional) Displays the MAC address table.

The following example clears the entire MAC address table of all dynamic entries:

```
switch# clear mac address-table dynamic
switch#
```

The following example clears the MAC address table of only those dynamic MAC addresses learned on VLAN 5:

```
switch# clear mac address-table dynamic vlan 5
switch#
```

Verifying the MAC Address Table Configuration

Use one of the following commands to verify the configuration:

Command	Purpose
show mac address-table	Displays the MAC address table.
show mac address-table static	Displays information about the MAC address table static entries.
show mac address-table static | inc veth	Displays the static MAC address of vEthernet interfaces in case a VEM physical port learns a dynamic MAC and the packet source is in another VEM on the same VSM.

Command	Purpose
show mac address static interface [**type** *if_id*]	Displays all static MAC addresses.
show mac address-table aging-time	Displays the aging time in the MAC address table.
show mac address-table count	Displays a count of MAC address entries.
show interface *interface_id* **mac**	Displays the MAC addresses and the burn-in MAC address for an interface.

Feature History for the MAC Address Table

Feature Name	Feature Name	Releases
MAC Address Tables	Release 5.2(1)IC1(1.1)	This feature was introduced

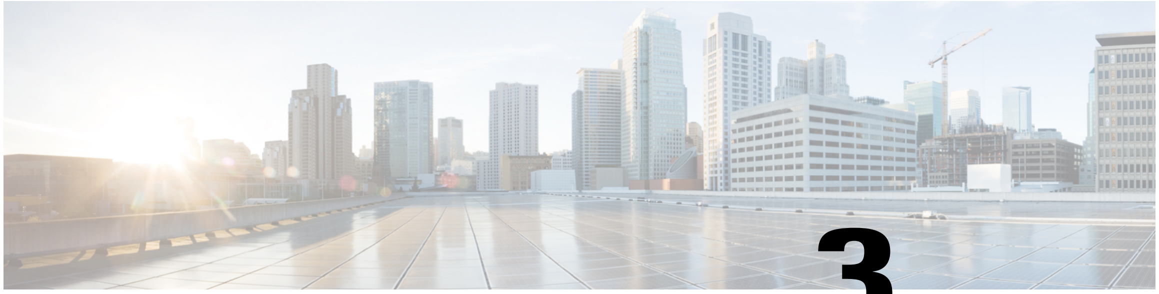

CHAPTER 3

Configuring VLANs

This chapter contains the following sections:

Information About VLANs

vEthernet interfaces that are assigned to specific VLANs are tagged with the VLAN when transmitted. A vEthernet interface that is not assigned to a specific VLAN, or assigned to VLAN 0, is transmitted as untagged on the physical NIC interfaces. When the VLAN is not specified, it is assumed to be 1.

The following table summarizes the actions taken on packets that are received by the Virtual Ethernet Module (VEM) based on VLAN tagging.

Table 2: VEM Action on VLAN Tagging

Port Type	Packet received	Action
Access	Tagged	The packet is dropped.
Access	Untagged	The VEM adds access VLAN to the packet.
Trunk	Tagged	No action is taken on the packet.
Trunk	Untagged	The VEM adds native VLAN tag to packet.

Guidelines and Limitations

In accordance with the IEEE 802.1Q standard, up to 128 VLANs are supported in Cisco Nexus 1000V and the valid range is 1-4094, and are organized in the following table:

Table 3: Cisco Nexus 1000V VLAN Numbering

VLANs Numbers	Range	Usage
1	Normal	Cisco Nexus 1000V default. You can use this VLAN, but you cannot modify or delete it.
2–1005	Normal	You can create, use, modify, and delete these VLANs.
1006-4094	Extended	You can create, name, and use these VLANs. You cannot change the following parameters: • State is always active. • VLAN is always enabled. You cannot shut down these VLANs. The extended system ID is always automatically enabled.
3968-4047 and 4094	Internally allocated	You cannot use, create, delete, or modify these VLANs. You can display these VLANs. Cisco Nexus 1000V allocates these 80 VLANs, plus VLAN 4094, for features, like diagnostics, that use internal VLANs for their operation.

Note For information about diagnostics, see the document, .

Default Settings

Table 4: Default VLAN Settings

Parameters	Default
VLAN assignment for all interfaces and all ports configured as switchports	VLAN 1
VLAN name	VLANxxxx where xxxx represent four numeric digits (including leading zeroes) equal to the VLAN ID number
Shut state	No shutdown
Operational state	Active
External switch tagging (EST)	Enabled
IGMP snooping	Enabled

Configuring a VLAN

Creating a VLAN

Use this procedure to do one of the following:

- Create a single VLAN that does not already exist.
- Create a range of VLANs that do not already exist.
- Delete an existing VLAN.

Note All interfaces and all ports configured as switchports are in VLAN 1 by default.

Before You Begin

- You are logged in to the CLI in EXEC mode.
- VLAN characteristics are configured in the VLAN configuration mode.
- You are familiar with the VLAN numbering.
- Newly-created VLANs remain unused until Layer 2 ports are assigned to them.

- When you delete a specified VLAN, the ports associated to that VLAN are shut down and no traffic flows. When you delete a specified VLAN from a trunk port, only that VLAN is shut down and traffic continues to flow on all the other VLANs through the trunk port. However, the system retains all the VLAN-to-port mapping for that VLAN, and when you reenable, or re-create, that specified VLAN, the system automatically reinstates all the original ports to that VLAN. Note that the static MAC addresses and aging time for that VLAN are not restored when the VLAN is reenabled.

Note Be aware that the Cisco NX-OS commands may differ from those used in Cisco IOS.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **show vlan**	Displays the VLANs that already exist.
Step 3	switch(config)# { **no** } **vlan** { *vlan-id* | *vlan-range* }	Creates or deletes, and saves in the running configuration, a VLAN or a range of VLANs. **Note** If you enter a VLAN ID that is assigned to an internally allocated VLAN, the system returns an error message. From the VLAN configuration mode, you can also create and delete VLANs.
Step 4	switch(config-vlan)# **show vlan id** *vlan-id*	(Optional) Displays the VLAN configuration.
Step 5	switch(config-vlan)# **copy running-config startup-config**	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

In the following example VLAN 5 is created and you are automatically placed into the VLAN configuration mode for VLAN 5:

```
switch# configure terminal
switch(config)# vlan 5
switch(config-vlan)#
```

The following example shows the range, VLAN 15-20, being created. The VLANs in the range are activated, and you are automatically placed into VLAN configuration mode for VLANs 15-20.

Note If you create a range of VLANs that includes an unusable VLAN, all VLANs in the range are created except those that are unusable; and Cisco Nexus 1000V returns a message listing the failed VLANs.

```
switch# configure terminal
switch(config)# vlan 15-20
switch(config-vlan)#
```

The following example shows VLAN 3967 being deleted, using the no form of the command:

```
switch# configure terminal
switch(config)# no vlan 3967
switch(config)#
```

Configuring VLAN Characteristics

Use this procedure to configure the following for a VLAN that has already been created:

Note Commands entered in the VLAN configuration mode are immediately saved to the running configuration.

- Name the VLAN.
- The operational state (active, suspend) of the VLAN.
- The VLAN media type .
- Shut down switching on the VLAN.

Before You Begin

You are logged in to the CLI in EXEC mode.

Note Some characteristics cannot be modified on some VLANs. For more information, see the VLAN numbering described in the Guidelines and Limitations, on page 10section.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vlan** { *vlan-id* | *vlan-range* }	Enters VLAN configuration mode for the specified VLAN. **Note** If the VLAN does not already exist, the system creates it and then enters the VLAN configuration mode for that VLAN.
Step 3	switch(config-vlan)# **name** *vlan-name*	Adds a name to the VLAN of up to 32 alphanumeric characters. • You cannot change the name of VLAN1 nor the VLANs reserved for internal use. • The default name is VLANxxxx where xxxx represent four numeric digits (including leading zeroes) equal to the VLAN ID number.
Step 4	switch(config-vlan)# **state** { **active** | **suspend** }	Changes the operational state of the VLAN and saves it in the running configuration. Allowable entries are:

	Command or Action	Purpose
		• Active (default) • Suspend While the VLAN state is suspended, the ports associated with this VLAN are shut down, and that VLAN does not pass any traffic. **Note** You cannot suspend the state for the default VLAN or VLANs 1006 to 4094.
Step 5	switch(config-vlan)# **no shutdown**	Enables VLAN switching in the running configuration. Allowable entries are: • no shutdown (default) • shutdown **Note** You cannot shut down the default VLAN, VLAN1, or VLANs 1006 to 4094.
Step 6	switch(config-vlan)# **show vlan** [**id** *vlan-id*]	(Optional) Displays the VLAN configuration.
Step 7	switch(config-vlan)# **copy running-config startup-config**	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

```
n1000v# configure terminal
n1000v(config)# vlan 5
n1000v(config-vlan)# name accounting
n1000v(config-vlan)# state active
n1000v(config-vlan)# no shutdown
n1000v(config-vlan)# show vlan brief
```

Verifying the Configuration

Use one of the following commands to verify the configuration:

Command	Purpose
show running-config vlan *vlan-id*	Displays VLAN information in the running configuration.
show vlan [**all-ports** | **brief** | **id** *vlan-id* | **name** *name* | **dot1q tag native**]	Displays the specified VLAN information.
show vlan summary	Displays a summary of VLAN information.

Feature History for VLANs

Feature Name	Feature Name	Releases
VLANs	Release 5.2(1)IC1(1.1)	This feature was introduced

CHAPTER 4

Configuring IGMP Snooping

This chapter contains the following sections:

Information about IGMP Snooping

Introduction

The Internet Group Management Protocol (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. By default, IGMP snooping is enabled on the device.

The following 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.

Figure 1: IGMP Snooping Switch

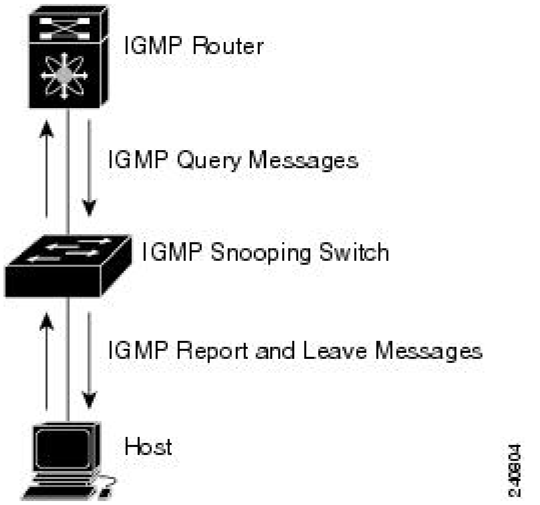

The IGMP snooping software operates upon IGMPv1, IGMPv2, and IGMPv3 control plane packets where Layer 3 control plane packets are intercepted and influence the Layer 2 forwarding behavior.

The Cisco Nexus 1000V IGMP snooping implementation has the following proprietary features:

- Multicast forwarding based on an IP address rather than a MAC address.
- Optimized multicast flooding (OMF) that forwards unknown traffic to routers only and performs no data driven state creation.

For more information about IGMP snooping, see RFC 4541.

IGMPv1 and IGMPv2

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.

Report suppression is not supported and is disabled by default.

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

IGMPv3 snooping provides constrained flooding based on the group IP information in the IGMPv3 reports.

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 switch sends to other multicast capable routers.

Even though the IGMPv3 membership reports provide a full accounting of group members on a LAN segment, when the last host leaves, the querier sends a membership query. You can configure the parameter last member query interval. If no host responds before the time-out, the software removes the group state. If the querier specifies a mean-response-time (MRT) value in the queries, it overrides the last member query interval configuration.

Prerequisites for IGMP Snooping

IGMP snooping has the following prerequisites:

- You are logged in to the switch.
- A querier must be running on the uplink switches on the VLANs that contain multicast sources and receivers.

When the multicast traffic does not need to be routed, you must configure an external switch to query membership. On the external switch, define the query feature in a VLAN that contains multicast sources and receivers but no other active query feature. In Cisco Nexus 1000V, report suppression is not supported and is disabled by default.

When an IGMP snooping query feature is enabled, it sends out periodic IGMP queries that trigger IGMP report messages from hosts wanting to receive IP multicast traffic. IGMP snooping listens to these IGMP reports to identify accurate forwarding.

Default Settings

Table 5: Default IGMP Snooping Settings

Parameters	Default
IGMP snooping	Enabled
IGMPv3 Explicit tracking	Enabled
IGMPv2 Fast leave	Enabled
Last member query interval	1 second
Link-local groups suppression	Enabled
Snooping querier	Disabled

Parameters	Default
IGMPv1/v2 Report suppression	Disabled
IGMPv3 Report suppression	Disabled

Configuring IGMP Snooping

Enabling or Disabling IGMP Snooping Globally for the VSM

Use this procedure to enable or disable IGMP snooping globally for the VSM. IGMP snooping is enabled globally on the VSM (the default). If enabled globally, you can turn it on or off per VLAN.

Before You Begin

You are logged in to the CLI in EXEC mode.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# [**no**] **ip igmp snooping**	Enables or disables IGMP snooping in the running configuration for all VLANs. The default is enabled. If you have previously disabled the feature then you can enable it with this command.
Step 3	switch(config)# **show ip igmp snooping** [**vlan** *vlan-id*]	(Optional) Displays the configuration for verification. **Note** If disabled, then IGMP snooping on all VLANs is disabled.
Step 4	switch(config)# **copy running-config startup-config**	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Configuring IGMP Snooping on a VLAN

Use this procedure to configure IGMP snooping on a VLAN. IGMP snooping is enabled by default for all VLANs in the VSM.

Before You Begin

You are logged in to the CLI in EXEC mode.

Note If IGMP snooping is disabled globally, it takes precedence over the VLAN state.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vlan configuration** *vlan-id*	Enters configuration mode for the specified VLAN.
Step 3	switch(config-vlan-config)# [**no**] **ip igmp snooping**	Enables or disables IGMP snooping in the running configuration for the specific VLAN. If IGMP snooping is enabled for the VSM, then IGMP snooping is enabled for the VLAN by default. **Note** IGMP snooping must be enabled globally (the default) in order to toggle it on or off per VLAN. If IGMP snooping is disabled globally, then it cannot be enabled per VLAN.
Step 4	switch(config-vlan-config)# [**no**] **ip igmp snooping mrouter interface type** *if_id*	(Optional) Configures a static connection for the VLAN to a multicast router in the running configuration. The interface to the router must be in the specified VLAN. You can specify the interface by the type and the number.
Step 5	switch(config-vlan-config)# [**no**] **ip igmp snooping static-group** *group-ip-addr* **interface type** *if_id*	(Optional) Configures a VLAN Layer 2 port as a static member of a multicast group in the running configuration. You can specify the interface by the type and the number.
Step 6	switch(config-vlan-config)# [**no**] **ip igmp snooping link-local-groups-suppression**	(Optional) Configures link-local groups suppression. The default is enabled. **Note** You can apply link-local groups suppression to all interfaces in the VSM by entering this command in global configuration mode.
Step 7	switch(config-vlan-config)# **show ip igmp snooping** [**vlan** *vlan-id*]	(Optional) Displays the configuration for verification.
Step 8	switch(config-vlan-config)# **copy running-config startup-config**	(Optional) (Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

Verifying the IGMP Snooping Configuration

Use the following commands to verify the IGMP snooping configuration information.

Command	Purpose
show ip igmp snooping [**vlan** *vlan-id*]	Displays IGMP snooping configuration by VLAN.
show ip igmp snooping groups [**vlan** *vlan-id*] [**detail**]	Displays IGMP snooping information about groups by VLAN.
show ip igmp snooping querier [**vlan** *vlan-id*]	Displays IGMP snooping queriers by VLAN.
show ip igmp snooping mroute [**vlan** *vlan-id*]	Displays multicast router ports by VLAN.
show ip igmp snooping explicit-tracking [**vlan** *vlan-id*]	Displays IGMP snooping explicit tracking information by VLAN.

Feature History for IGMP Snooping

Feature Name	Releases	Description
IGMP Snooping	Release 5.2(1)IC1(1.1)	This feature was introduced.

CHAPTER 5

Layer 2 Switching Configuration Limits

This chapter contains the following sections:

Layer 2 Switching Configuration Limits

The configuration limits are documented in the *Cisco Nexus 1000V InterCloud Verified Scalability Reference*.