Cisco Nexus 9000 Series NX-OS IP Fabric for Media Solution Guide, Release 7.0(3)I4(2)

First Published: 2016-07-15

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

CONTENTS

Preface

This preface includes the following sections:

Audience

This publication is for network administrators who install, configure, and maintain Cisco Nexus switches.

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.
[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.

Convention	Description
[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.

Related Documentation for Cisco Nexus 9000 Series Switches

The entire Cisco Nexus 9000 Series switch documentation set is available at the following URL:

http://www.cisco.com/en/US/products/ps13386/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 nexus9k-docfeedback@cisco.com. We appreciate your feedback.

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see *What's New in Cisco Product Documentation* at: http://www.cisco.com/c/en/us/td/docs/general/whatsnew/whatsnew.html.

Subscribe to *What's New in Cisco Product Documentation*, which lists all new and revised Cisco technical documentation as an RSS feed and delivers content directly to your desktop using a reader application. The RSS feeds are a free service.

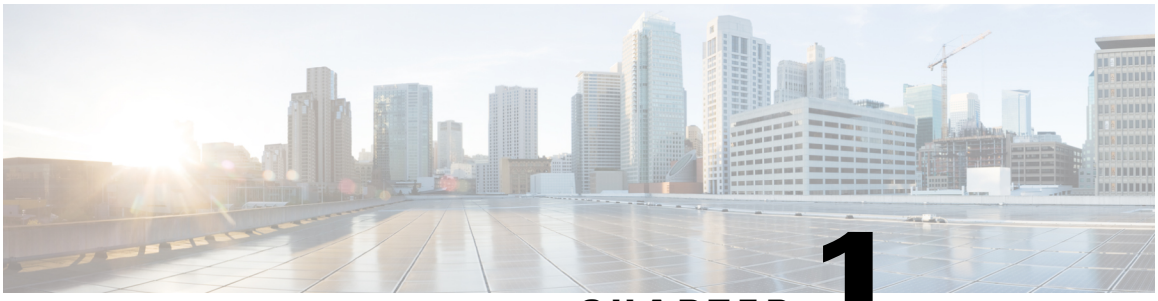

CHAPTER 1

New and Changed Information

This chapter provides release-specific information for each new and changed feature in the *Cisco Nexus 9000 Series NX-OS IP Fabric for Media Solution Guide, Release 7.0(3)I4(2)*.

New and Changed Information

This table summarizes the new and changed features for the *Cisco Nexus 9000 Series NX-OS IP Fabric for Media Solution Guide, Release 7.0(3)I4(2)* and tells you where they are documented.

Table 1: New and Changed Features for Cisco NX-OS Release 7.0(3)I4(2)

Feature	Description	Changed in Release	Where Documented
IP fabric for media	First release. .	7.0(3)I4(2)	

New and Changed Information

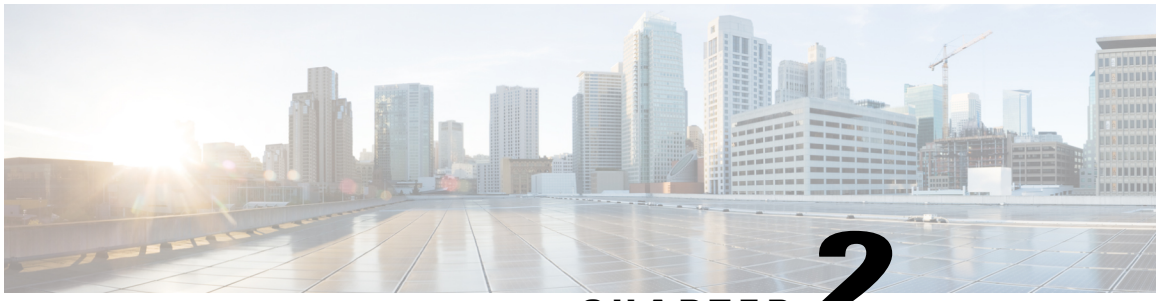

CHAPTER 2

Overview of Cisco's IP Fabric for Media Solution

This chapter contains information about Cisco's IP fabric for media solution.

About the IP Fabric for Media Solution

Today, the broadcast industry uses a serial digital interface (SDI) router and SDI cables to transport video and audio traffic. The SDI cables can carry only a single unidirectional signal. As a result, a large number of cables, frequently stretched over long distances, are required, making it difficult and time-consuming to expand or change an SDI-based infrastructure.

Cisco's IP fabric for media solution replaces the SDI router in live production studios. In an IP-based infrastructure, a single cable has the capacity to carry multiple bidirectional traffic flows and can support different flow sizes without requiring changes to the physical infrastructure. The Cisco Nexus 9000 Series switches in conjunction with the Cisco non-blocking multicast (NBM) algorithm (an intelligent traffic management algorithm) provide a reliable, scalable IP fabric for the broadcast industry. In addition, this IP fabric solution provides zero-drop multicast transport and support for PTP media profiles.

Cisco Nexus 9200 Series Switches

The following Cisco Nexus 9200 Series switches are used to transport video and audio traffic through the IP fabric:

Cisco Nexus 9200 Series Switch	Number and Size of Ports
Cisco Nexus 9236C switch	36 x 40/100-Gbps ports
Cisco Nexus 9272Q switch	72 x 40-Gbps ports

Cisco Nexus 9200 Series Switch	Number and Size of Ports
Cisco Nexus 92160YC-X switch	48 x 1/10/25-Gbps ports

IP Fabric for Media Topology

Cisco's IP fabric for media solution supports a spine-leaf topology that consists of one spine switch and up to nine leaf switches.

The Cisco Nexus 9236C or 9272Q switch can serve as the spine switch. The Cisco Nexus 9236C, 9272Q, and 92160YC-X switches can act as leaf switches. The topology supports any combination of these leaf switches, including using just one type of leaf switch.

Figure 1: IP Fabric for Media Topology

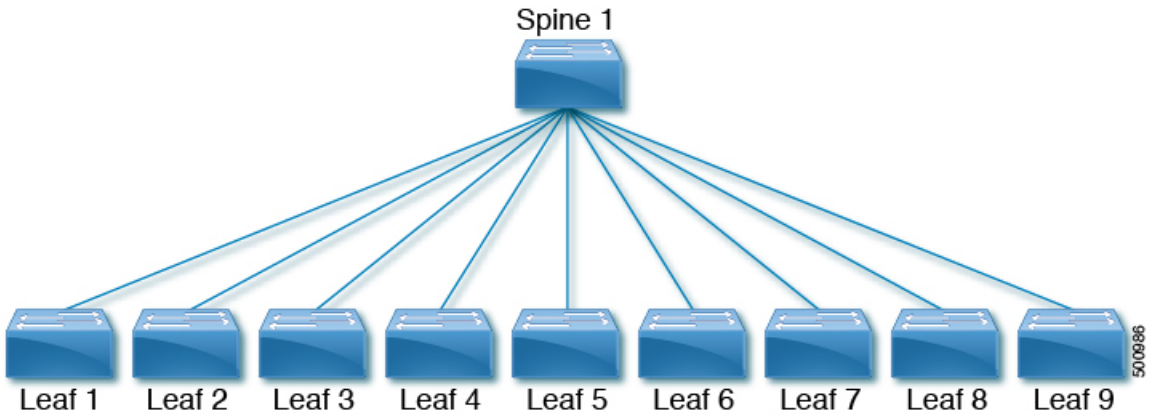

Media sources and receivers connect to the leaf switches, and receivers initiate IGMP join requests to the leaf switches in order to receive the media traffic.

Failure Handling

Cisco's IP fabric for media solution supports deterministic failure handling. When a link fails, the following actions occur:

- When a link fails between a source leaf switch and the spine switch, the spine switch executes the NBM algorithm and requests for flows on other links, provided sufficient bandwidth is available.
- When a link fails between the spine switch and a receiver leaf switch, the leaf switch executes the NBM algorithm and requests for flows on other links, provided sufficient bandwidth is available.
- Flows that are not impacted are not moved to other links.
- When a link comes up, only new flows are forwarded across it, and previous flows are not moved back to the link.

Benefits of the IP Fabric for Media Solution

Cisco's IP fabric for media solution provides the following benefits:

- Replaces specialized hardware (SDI routers) with a general-purpose switching infrastructure.
- Supports various types and sizes of broadcasting equipment endpoints with port speeds up to 100 Gbps.
- Provides up to 3.6 Tb of bandwidth to support the latest video technologies, including 4K and 8K ultra HD.

Note For example, the bandwidth needed to replace an existing SDI router can be calculated as follows: a 1024 x 1024 SDI router with HD flows would require a bandwidth equivalent to 1024 x 1.5 Gbps = 1.5 Tbps. Using a Cisco Nexus 9236C switch as the spine switch, the topology would support a bandwidth of 3.6 Tbps, or the equivalent of a 2400 x 2400 SDI router.

- Scales horizontally. When you need more capacity, you can add a leaf switch to support more endpoints.
- Provides a deterministic network with zero packet loss, ultra low latency, and minimal jitter.
- Capable of synchronizing all media sources and receivers.
- Provides deterministic failure handling that sends traffic to the receiver when a link fails between a leaf and the spine.
- Supports the coexistence of live and file-based traffic flows for post-production work.
- Offers increased network security.
- Provides a non-blocking network design to prevent the oversubscription of links.
- Requires no changes to the existing operator workflow.

CHAPTER 3

Setting Up the IP Fabric for Media

This chapter describes how to set up an IP fabric for media network.

Determining the Number and Types of Leaf Switches Required in the IP Fabric

The number and types of leaf switches required in your IP fabric depend on the number and types of endpoints in your broadcasting center.

Follow these steps to help determine how many leaf switches you need:

1. Count the number of endpoints (cameras, microphones, etc.) in your broadcasting center (for example, 90 10-Gbps endpoints and 30 40-Gbps endpoints).
2. Determine the type of leaf switches required based on the type of endpoints in your broadcasting center.
   - For 10-Gbps endpoints, you need to use the Cisco Nexus 92160YC-X leaf switches.
   - For 40-Gbps endpoints, you can use the Cisco Nexus 9236C or 9272Q leaf switches.
3. Determine the number of leaf switches required based on the number of endpoints and uplinks that each leaf switch supports.

Table 2: Endpoints and Uplinks Supported per Leaf Switch

Leaf Switch	Endpoint Capacity	Uplink Capacity
Cisco Nexus 9236C switch	25 x 40-Gbps endpoints	10 x 100-Gbps (1000-Gbps) uplinks
Cisco Nexus 9272Q switch	36 x 40-Gbps endpoints	36 x 40-Gbps (1440-Gbps) uplinks

Leaf Switch	Endpoint Capacity	Uplink Capacity
Cisco Nexus 92160YC-X switch	40 x 10-Gbps endpoints	4 x 100-Gbps (400-Gbps) uplinks

For example:

- For 90 10-Gbps endpoints, you need three Cisco Nexus 92160YC-X leaf switches because each switch can support up to 40 10-Gbps endpoints.
- For 30 40-Gbps endpoints, you need two Cisco Nexus 9236C leaf switches because each switch can support up to 25 40-Gbps endpoints or one Cisco Nexus 9272Q leaf switch because each switch can support up to 36 40-Gbps endpoints.

4 Make sure that the uplink bandwidth (toward the spine switch) and the downstream bandwidth (toward the endpoints) are equal.

a Use this equation to determine the uplink bandwidth:

```
Uplink Capacity per Leaf Switch x Number of Leaf Switches = Uplink Bandwidth
```

For example:

400 Gbps (uplink capacity for each Cisco Nexus 92160YC-X switch) x 3 Cisco Nexus 92160YC-X leaf switches = 1200-Gbps uplink bandwidth

1000 Gbps (uplink capacity for each Cisco Nexus 9236C switch) x 2 Cisco Nexus 9236C leaf switches = 2000-Gbps uplink bandwidth

1200-Gbps uplink bandwidth (for 3 Cisco Nexus 92160YC-X leaf switches) + 1200-Gbps uplink bandwidth (for 2 Cisco Nexus 9236C leaf switches) = 2400-Gbps total uplink bandwidth

b Use this equation to determine the downstream bandwidth:

```
Endpoint Capacity per Leaf Switch x Number of Leaf Switches = Downstream Bandwidth
```

For example:

40 x 10-Gbps (400-Gbps endpoint capacity) for each Cisco Nexus 92160YC-X leaf switch x 3 leaf switches = 1200-Gbps downstream bandwidth

25 x 40-Gbps (1000-Gbps endpoint capacity) for each Cisco Nexus 9236C leaf switch x 2 leaf switches = 2000-Gbps downstream bandwidth

1200-Gbps downstream bandwidth (for 3 Cisco Nexus 92160YC-X leaf switches) + 1200-Gbps downstream bandwidth (for 2 Cisco Nexus 9236C leaf switches) = 2400-Gbps total downstream bandwidth

5 If the total uplink bandwidth and downstream bandwidth are equal, your topology is valid. You can now determine the number of achievable flows. If the uplink bandwidth and downstream bandwidth are not equal, you must rework your topology until it is valid.

Using the examples provided in this section, the topology would look like this:

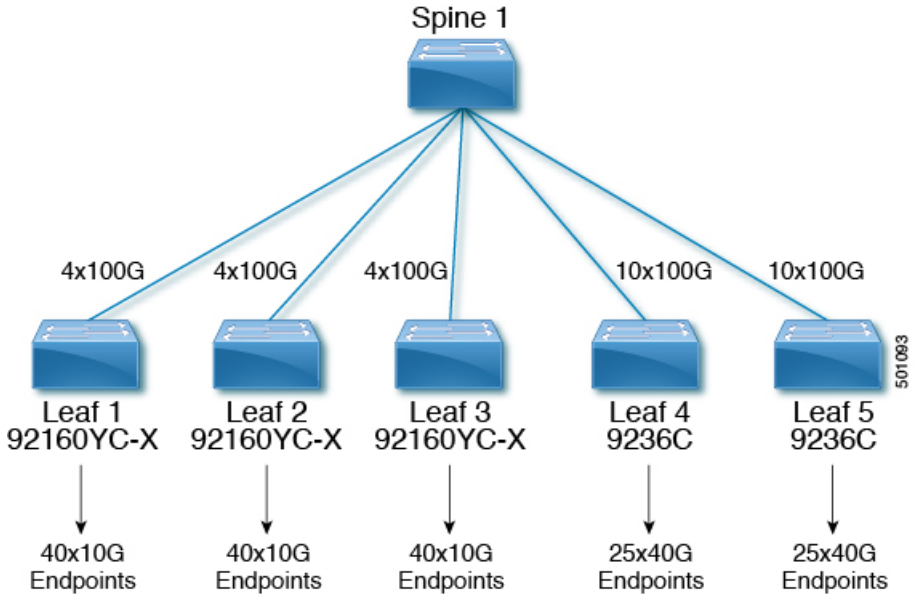

Determining the Number of Achievable Flows in the IP Fabric

Use this equation to determine the number of possible flows in your IP fabric:

```
Total Bandwidth ÷ Flow Size = Number of Achievable Flows
```

The flow size is configurable and is typically based on the type of video technology used in your broadcasting center. For more information, see Configuring NBM.

For example:

3200-Gbps total bandwidth ÷ 1.5-Gbps flow size (for HD video) = 2133 possible flows

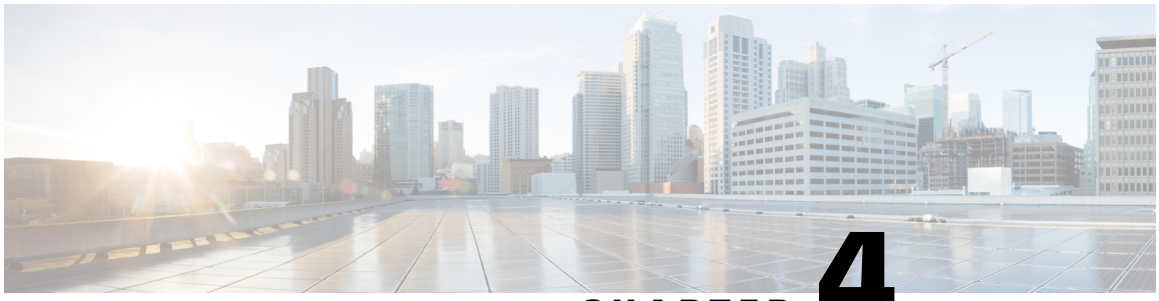

CHAPTER 4

Configuring IP Fabric for Media

This chapter describes how to configure the Cisco Nexus 9200 Series switches for Cisco's IP fabric for media solution.

Prerequisites

Cisco's IP fabric for media solution has the following prerequisites:

- The Cisco Nexus 9200 Series switches must be running Cisco NX-OS Release 7.0(3)I4(2) or a later release.
- TCAM carving must be configured on the spine switch and each leaf switch using the **hardware access-list tcam region ing-ifacl 2048** command. The following example shows how you might provision the TCAM. For more information on TCAM carving, see the Cisco Nexus 9000 Series NX-OS Security Configuration Guide.

```
hardware access-list tcam region ing-l2-qos 0
hardware access-list tcam region ing-l3-vlan-qos 256
hardware access-list tcam region ing-racl 256
hardware access-list tcam region ing-ifacl 2048
```

Guidelines and Limitations

Cisco's IP fabric for media solution has the following guidelines and limitations:

- Only the Cisco Nexus 9236C and 9272Q can be spine switches.
- Only the Cisco Nexus 9236C, 9272Q, and 92160YC-X can be leaf switches.

- The number of leaf switches depends on the number of uplinks used and the number of ports available on the spine switch.
- The uplink bandwidth from each leaf switch must be equal to the bandwidth provided to the endpoints.
- If possible, Cisco recommends overprovisioning uplinks to account for failures.
- As a best practice, use Layer 3 ports that go to the endpoints with a /30 mask. Assign one IP address to the endpoint and another to the switch interface.
- Cisco recommends that you choose the maximum bandwidth per flow. A constant or almost constant bandwidth rate is assumed. If you expect bursts of traffic in your IP fabric, consider reserving this value at a peak rate to prevent congestion.
- No two sources can transmit to the same multicast group at the same time. Also, when a given source stops transmitting traffic, the multicast group must time out before a new source can start transmitting to the same group.

Configuring NBM

Licensing Requirements

Product	License Requirement
Cisco NX-OS	The IP fabric for media solution requires an Enterprise Services license and a Network Services license. For a complete explanation of the Cisco NX-OS licensing scheme and how to obtain and apply licenses, see the *Cisco NX-OS Licensing Guide*.

Configuring NBM

After you have set up the IP fabric, you must enable the NBM feature and set the flow bandwidth on the spine switch and each leaf switch. The NBM feature ensures that the bandwidth that is coming into the fabric is exactly the same as the bandwidth that is going out.

Before You Begin

Enable the SPT threshold infinity using the **ip pim spt-threshold infinity** command.

Procedure

	Command or Action	Purpose

Step 1	**configure terminal**
	
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

	Command or Action	Purpose
Step 2	[**no**] **feature nbm** **Example:** `switch(config)# feature nbm`	Enables the NBM feature. The **no** form of this command disables this feature.
Step 3	[**no**] **ip multicast multipath nbm** **Example:** `switch(config)# ip multicast multipath nbm`	Configures multicast multipath for NBM. This command is required only if ECMP links are available for multicast traffic in the IP fabric. The **no** form of this command disables this feature.

Step 4	[**no**] **nbm flow bandwidth** *flow-bandwidth*
	
	Example:
	`switch(config)# nbm flow bandwidth 1500`
	`Setting NBM Per Flow Bandwidth as 1500.`
	`Existing NBM Per Flow Bandwidth Value (1000) will continue to be used.`
	`Changes shall take effect after reload`
		Configures the NBM flow size for this device. You should configure a flow size based on the type of video technology used in your broadcasting center. The range is from 1 to 100,000 Mbps.
		Table 3: Flow Sizes per Video Technology
		Technology / **Flow Size**
		HD video / 1.5 Gbps (1500 Mbps)
		3G HD video / 3 Gbps (3000 Mbps)
		4K ultra HD video / 12 Gbps (12,000 Mbps)
		8K ultra HD video / 48 Gbps (48,000 Mbps)
		The **no** form of this command disables this feature.
		Note The NBM flow size should be configured when you are initially setting up the IP fabric. It should not be changed at random. Also, any change to the existing flow bandwidth value takes effect only after the device is rebooted.

Step 5	**show nbm flows bandwidth**
	
	Example:
	`switch(config)# show nbm flows bandwidth`
	`Applied NBM Per Flow Bandwidth Value: 1000 Mbps`
	`Configured NBM Per Flow Bandwidth Value: 1500 Mbps`
		(Optional)
		Displays the configured and applied NBM bandwidth per flow.

Step 6	**copy running-config startup-config** **Example:** `switch(config)# copy running-config startup-config`	(Optional) Copies the running configuration to the startup configuration.

Verifying the NBM Configuration

To display the NBM configuration information, perform one of the following tasks:

Command	Purpose
show ip mroute	Displays the uplink interfaces for each NBM (S,G) or (*,G) flow.
show nbm flows bandwidth	Displays the configured and applied NBM bandwidth per flow.
show running-config | grep nbm	Displays the NBM running-configuration information.

Displaying NBM Flows and Flow Statistics

To display the NBM flows and flow statistics, perform one of the following tasks:

Command	Purpose
show nbm flows [**all**]	Displays the current NBM flows in the IP fabric. The flows are ordered by their startup time, with the most recent first. The **all** option shows the following information for expired flows: end time, total number of packets or bytes, and average flow rate.
show nbm flows statistics	Displays the NBM flow statistics.

The following example shows sample output for the **show nbm flows** command:

```
switch# show nbm flows

NBM Active flow(s)
Start-Time         Src-Port  Mcast-Group     Src-IP          L4-S  L4-D
06/17 01:27:11. 53 Eth1/2    225.0.0.9       192.168.102.2   1024  1024
06/17 01:27:11. 52 Eth1/2    225.0.0.8       192.168.102.2   1024  1024
06/17 01:27:11. 51 Eth1/2    225.0.0.7       192.168.102.2   1024  1024
06/17 01:27:11. 50 Eth1/2    225.0.0.6       192.168.102.2   1024  1024
06/17 01:27:11. 50 Eth1/2    225.0.0.4       192.168.102.2   1024  1024
06/17 01:27:11. 50 Eth1/2    225.0.0.3       192.168.102.2   1024  1024
06/17 01:27:11. 49 Eth1/2    225.0.0.2       192.168.102.2   1024  1024
06/17 01:27:11. 49 Eth1/2    225.0.0.1       192.168.102.2   1024  1024
06/17 01:27:11. 49 Eth1/2    225.0.0.0       192.168.102.2   1024  1024
```

The following example shows sample output for the **show nbm flows statistics** command:

```
switch# show nbm flows statistics

NBM Flow Statistics
```

```
Start-Time          Src-Port  Mcast-Group     Packets        Bytes
06/17 01:27:11. 53 Eth1/2    225.0.0.9       65163          268992864
06/17 01:27:11. 52 Eth1/2    225.0.0.8       65163          268992864
06/17 01:27:11. 51 Eth1/2    225.0.0.7       65196          269129088
06/17 01:27:11. 50 Eth1/2    225.0.0.6       65196          269129088
06/17 01:27:11. 50 Eth1/2    225.0.0.4       65198          269137344
06/17 01:27:11. 50 Eth1/2    225.0.0.3       65201          269149728
06/17 01:27:11. 49 Eth1/2    225.0.0.2       65207          269174496
06/17 01:27:11. 49 Eth1/2    225.0.0.1       65208          269178624
06/17 01:27:11. 49 Eth1/2    225.0.0.0       65207          269174496
```

Configuring PTP for Media

Cisco's IP fabric for media solution supports the following IEEE 1588 PTP profiles:

- Audio Engineering Society 67 profile (AES67) - For high-performance streaming audio over IP
- Professional Broadcast Environment profile (SMPTE-2059-2) - For high-performance streaming video over IP

The solution also introduces mixed mode PTP support with multicast sync and announce messages as well as unicast delay request and response messages.

To configure PTP for media, you should use one of these profiles.

Note The PTP configuration for media is different than the PTP configuration for a non-media network. However, you can refer to the *Cisco Nexus 9000 Series NX-OS System Management Configuration Guide* for more information on PTP.

Procedure

	Command or Action	Purpose

Step 1	**configure terminal**
	
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	[**no**] **feature ptp** **Example:** `switch(config)# feature ptp`	Enables or disables PTP on the device. **Note** Enabling PTP on the switch does not enable PTP on each interface.
Step 3	[**no**] **ptp source** *ip-address* [**vrf** *vrf*] **Example:** `switch(config)# ptp source 10.10.10.1`	Configures the source IPv4 address for all PTP packets.

	Command or Action	Purpose

Step 4	**interface ethernet** *slot/port*
	
	Example:
	`switch(config)# interface ethernet 2/1`
	`switch(config-if)#`
		Specifies the interface on which you are enabling PTP and enters the interface configuration mode.

Step 5	[**no**] **ptp** **Example:** `switch(config-if)# ptp`	Enables or disables PTP on an interface.
Step 6	[**no**] **ptp announce interval** [**aes67** | **smpte-2059**] *log-seconds* **Example:** `switch(config-if)# ptp announce interval aes67 3`	(Optional) Configures the interval between PTP announce messages on an interface. ***Table 4: PTP Announcement Interval Range and Default Values*** **Option** / **Range** / **Default Value** **aes67** / 0 to 4 log seconds / 1 log second **smpte-2059** / –3 to 1 log seconds / 1 log second Without the **aes67** or **smpte-2059** option / 0 to 4 log seconds / 1 log second
Step 7	[**no**] **ptp announce timeout** [**aes67** | **smpte-2059**] *count* **Example:** `switch(config-if)# ptp announce timeout aes67 2`	(Optional) Configures the number of PTP intervals before a timeout occurs on an interface. ***Table 5: PTP Announcement Timeout Range and Default Values*** **Option** / **Range** / **Default Value** **aes67** / 2 to 10 intervals / 3 intervals **smpte-2059** / 2 to 10 intervals / 3 intervals Without the **aes67** or **smpte-2059** option / 2 to 4 intervals / 3 intervals

	Command or Action	Purpose
Step 8	[**no**] **ptp delay-request minimum interval** [**aes67** | **smpte-2059**] *log-seconds* **Example:** `switch(config-if)# ptp delay-request minimum interval aes67 -1`	(Optional) Configures the minimum interval allowed between PTP delay messages when the port is in the master state. ***Table 6: PTP Delay-Request Minimum Interval Range and Default Values*** **Option** / **Range** / **Default Value** **aes67** / –4 to 5 log seconds / 0 log seconds **smpte-2059** / –4 to 5 log seconds / 0 log seconds Without the **aes67** or **smpte-2059** option / –1 to 6 log seconds (where –1 = 1 frame per second) / 0 log seconds
Step 9	[**no**] **ptp sync interval** [**aes67** | **smpte-2059**] *log-seconds* **Example:** `switch(config-if)# ptp sync interval aes67 1`	(Optional) Configures the interval between PTP synchronization messages on an interface. ***Table 7: PTP Synchronization Interval Range and Default Values*** **Option** / **Range** / **Default Value** **aes67** / –4 to 1 log seconds / –2 log seconds **smpte-2059** / –4 to –1 log seconds / –2 log seconds Without the **aes67** or **smpte-2059** option / –3 to 1 log seconds / –2 log seconds

	Command or Action	Purpose
Step 10	[**no**] **ptp vlan** *vlan-id* **Example:** `switch(config-if)# ptp vlan 1`	(Optional) Specifies the VLAN for the interface where PTP is being enabled. You can enable PTP only on one VLAN on an interface. The range is from 1 to 4094.
Step 11	**show ptp brief** **Example:** `switch(config-if)# show ptp brief`	(Optional) Displays the PTP status.
Step 12	**show ptp port interface** *interface slot/port* **Example:** `switch(config-if)# show ptp port interface ethernet 2/1`	(Optional) Displays the status of the PTP port.
Step 13	**copy running-config startup-config** **Example:** `switch(config-if)# copy running-config startup-config`	(Optional) Copies the running configuration to the startup configuration.

Configuration Example

The following example shows how to configure an IP network that supports 3G HD video broadcast traffic:

```
switch# configure terminal
switch(config)# hardware access-list tcam region ing-ifacl 2048
switch(config)# feature nbm
switch(config)# ip multicast multipath nbm
switch(config)# nbm flow bandwidth 3000
Setting NBM Per Flow Bandwidth as 3000.
Existing NBM Per Flow Bandwidth Value (1000) will continue to be used.
Changes shall take effect after reload
switch(config)# feature ptp
switch(config)# ptp source 10.10.10.1
switch(config)# interface ethernet 1/1
switch(config-if)# ptp
switch(config-if)# ptp announce interval smpte-2059 1
switch(config-if)# ptp announce timeout smpte-2059 5
switch(config-if)# ptp delay-request minimum interval smpte-2059 -1
switch(config-if)# ptp sync interval smpte-2059 -1
switch(config-if)# ptp vlan 1
```

Related Documentation

Related Topic	Document Title
IP fabric for media	Cisco IP Fabric for Media Solution video

Related Topic	Document Title
Cisco NX-OS release information	Cisco Nexus 9000 Series NX-OS IP Fabric for Media Release Notes
Cisco NX-OS software upgrades	Cisco Nexus 9000 Series NX-OS Software Upgrade and Downgrade Guide
PTP	Cisco Nexus 9000 Series NX-OS System Management Configuration Guide
TCAM carving	Cisco Nexus 9000 Series NX-OS Security Configuration Guide

INDEX