Configuring Layer 2 Switching

Information About Layer 2 Switching

Note See the Cisco Nexus 9000 Series NX-OS Interfaces Configuration GuideCisco Nexus 9000 Series NX-OS Interfaces Configuration Guide*Cisco Nexus 9000 Series NX-OS Interfaces Configuration Guide*, for information on creating interfaces.

You can configure Layer 2 switching ports as access or trunk ports. Trunks carry the traffic of multiple VLANs over a single link and allow you to extend VLANs across an entire network. All Layer 2 switching ports maintain MAC address tables.

Note See the *Cisco Nexus 9000 Series NX-OS High Availability and Redundancy Guide*Cisco Nexus 9000 Series NX-OS High Availability and Redundancy GuideCisco Nexus 9000 Series NX-OS High Availability and Redundancy Guide, for complete information on high-availability features.

Layer 2 Ethernet Switching Overview

The device supports simultaneous, parallel connections between Layer 2 Ethernet segments. Switched connections between Ethernet segments last only for the duration of the packet. New connections can be made between different segments for the next packet.

The device solves congestion problems caused by high-bandwidth devices and a large number of users by assigning each device (for example, a server) to its own collision domain. Because each LAN port connects to a separate Ethernet collision domain, servers in a switched environment achieve full access to the bandwidth.

Because collisions cause significant congestion in Ethernet networks, an effective solution is full-duplex communication. Typically, 10/100-Mbps Ethernet operates in half-duplex mode, which means that stations can either receive or transmit. In full-duplex mode, which is configurable on these interfaces, two stations can transmit and receive at the same time. When packets can flow in both directions simultaneously, the effective Ethernet bandwidth doubles.

Switching Frames Between Segments

Each LAN port on a device can connect to a single workstation, server, or to another device through which workstations or servers connect to the network.

To reduce signal degradation, the device considers each LAN port to be an individual segment. When stations connected to different LAN ports need to communicate, the device forwards frames from one LAN port to the other at wire speed to ensure that each session receives full bandwidth.

To switch frames between LAN ports efficiently, the device maintains an address table. When a frame enters the device, it associates the media access control (MAC) address of the sending network device with the LAN port on which it was received.

Building the Address Table and Address Table Changes

The device dynamically builds the address table by using the MAC source address of the frames received. When the device receives a frame for a MAC destination address not listed in its address table, it floods the frame to all LAN ports of the same VLAN except the port that received the frame. When the destination station replies, the device adds its relevant MAC source address and port ID to the address table. The device then forwards subsequent frames 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 addresses as static MAC addresses. The static MAC entries are retained across a reboot of the device.

You must manually configure identical static MAC addresses on both devices connected by a virtual port channel (vPC) peer link. The MAC address table display is enhanced to display information on MAC addresses when you are using vPCs.

See the Cisco Nexus 9000 Series NX-OS Interfaces Configuration GuideCisco Nexus 9000 Series NX-OS Interfaces Configuration Guide*Cisco Nexus 9000 Series NX-OS Interfaces Configuration Guide* for information about vPCs.

The address table can store a number of MAC address entries depending on the hardware I/O module. The device uses an aging mechanism, defined by a configurable aging timer, so if an address remains inactive for a specified number of seconds, it is removed from the address table.

Consistent MAC Address Tables on the Supervisor and on the Modules

Optimally, all the MAC address tables on each module exactly match the MAC address table on the supervisor. When you enter the **show forwarding consistency l2** command or the **show consistency-checker l2** command, the device displays discrepant, missing, and extra MAC address entries.

High Availability for Switching

You can upgrade or downgrade the software seamlessly, with respect to classical Ethernet switching. If you have configured static MAC addresses on Layer 3 interfaces, you must unconfigure those ports in order to downgrade the software.

Note See the *Cisco Nexus 9000 Series NX-OS High Availability and Redundancy Guide*Cisco Nexus 9000 Series NX-OS High Availability and Redundancy GuideCisco Nexus 9000 Series NX-OS High Availability and Redundancy Guide, for complete information on high availability features.

Prerequisites for Configuring MAC Addresses

MAC addresses have the following prerequisites:

- You must be logged onto the device.
- If necessary, install the Advanced Services license.

Default Settings for Layer 2 Switching

This table lists the default setting for Layer 2 switching parameters.

Table 1: Default Layer 2 Switching Parameters

Parameters	Default
Aging time	1800 seconds

MAC Move Loop Detection

Cisco Nexus 9000 Series switches leverage L2FM for software MAC learning (and, subsequently, loop detection). If a host (MAC address) moves between two interfaces within the same VLAN, it would trigger a MAC move. If there are a large number of such MAC moves in a short duration of time, the control plane of the switch and the CPU performance could get impacted. L2FM protects the switch from such scenarios by disabling MAC learning on the specific VLAN once the number of MAC moves for the corresponding MAC address exceeds a threshold.

For Broadcom ASIC based switches, the MAC move learn disable threshold criteria is when a single MAC address moves 10 or more times in a duration of 1 second within the same VLAN.

For Cisco Nexus 9300-EX/FX/FX2/FX3/GX/GX2/H2R/H1, 9804/9808 switches and Cisco Nexus 9500 switches with 9700-EX/FX/GX line cards, the MAC move learn disable threshold criteria is when a single MAC address moves 10 or more times in 10 seconds within the same VLAN.

Once the threshold limit is hit, all new MAC learning on the corresponding VLAN gets disabled for a period of 120 seconds. After 120 seconds, new MAC learning is re-enabled on that VLAN. There is no impact of this on the rest of the VLANs on the switch.

Generating Syslog Error Messages

To see MAC move notifications in syslogs, follow the below steps:

Procedure

	Command or Action	Purpose

Step 1	**config t**
	Example:
	`switch# config t`
	`switch(config)#`
		Enters configuration mode.

Step 2	**logging level l2fm 5** **Example:** `switch(config)# logging level l2fm 5`	Enables logging of all L2FM events from level 5 up to the highest severity events.
Step 3	(Optional) **mac address-table notification mac-move** **Example:** `switch(config)# mac address-table notification mac-move`	Enables MAC move notification on the switch. **Note** • MAC move notification is enabled by default. • This command ensures that the syslog for L2FM detect displays when there is a MAC address move.

Following are the sample generated syslog messages:

- When MAC move is detected:

  ```
  2023 Nov 29 21:42:04 N-3164Q-40G %L2FM-4-L2FM_MAC_MOVE2: Mac
  0003.0001.005d in vlan 500 has moved from Eth1/24 to Eth1/63
  ```

- When MAC learning on VLAN is disabled:

  ```
  2023 Nov 29 21:23:29 N-3164Q-40G %L2FM-2-L2FM_MAC_FLAP_DISABLE_LEARN:
  Disabling learning in vlan 500 for 120s due to too many mac moves
  ```

- When MAC learning on VLAN is re-enabled:

  ```
  2023 Nov 29 21:23:19 N-3164Q-40G %L2FM-2-L2FM_MAC_FLAP_RE_ENABLE_LEARN:
  Re-enabling learning in vlan 500
  ```

Example

In order to check if the MAC addresses move, enter the command:

```
switch# show mac address-table notification mac-move
MAC Move Notify Triggers: 1206
Number of MAC Addresses added: 944088
Number of MAC Addresses moved: 265
Number of MAC Addresses removed: 943920
```

Note The following are the possible causes for MAC moves:

- MAC addresses move because of server NIC teaming and moving between Active-Active, Active-Standby states, etc.
- MAC addresses move because the source of the data is physically moved across all switches while STP states are converged and in correct states.
- Due to loops in the network.

Configuring Layer 2 Switching by Steps

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.

Configuring a Static MAC Address

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.

SUMMARY STEPS

1. **config t**
2. **mac address-table static** *mac-address* **vlan** *vlan-id* {[**drop** | **interface** {*type slot/port*} | **port-channel** *number*]}
3. **exit**
4. (Optional) **show mac address-table static**
5. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**config t**
	Example:
	`switch# config t`
	`switch(config)#`
		Enters configuration mode.

	Command or Action	Purpose

Step 2	**mac address-table static** *mac-address* **vlan** *vlan-id* {[**drop** | **interface** {*type slot/port*} | **port-channel** *number*]}
	Example:
	`switch(config)# mac address-table static 1.1.1 vlan 2 interface ethernet 1/2`
		Specifies a static MAC address to add to the Layer 2 MAC address table.

Step 3	**exit**
	Example:
	`switch(config)# exit`
	`switch#`
		Exits the configuration mode.

Step 4	(Optional) **show mac address-table static** **Example:** `switch# show mac address-table static`	Displays the static MAC addresses.
Step 5	(Optional) **copy running-config startup-config** **Example:** `switch# copy running-config startup-config`	Copies the running configuration to the startup configuration.

Example

This example shows how to put a static entry in the Layer 2 MAC address table:

```
switch# config t
switch(config)# mac address-table static 1.1.1 vlan 2 interface ethernet 1/2
switch(config)#
```

Configuring the Aging Time for the MAC Table

You can configure the amount of time that a MAC address entry (the packet source MAC address and port on which that packet was learned) remains in the MAC table, which contains the Layer 2 information.

Note MAC addresses are aged out up to two times the configured MAC address table aging timeout.

Note You can also configure the MAC aging time in interface configuration mode or VLAN configuration mode.

SUMMARY STEPS

1. **config t**
2. **mac address-table aging-time** *seconds*
3. **exit**

4. (Optional) **show mac address-table aging-time**
5. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**config t**
	Example:
	`switch# config t`
	`switch(config)#`
		Enters configuration mode.

Step 2	**mac address-table aging-time** *seconds* **Example:** `switch(config)# mac address-table aging-time 600`	Specifies the time before an entry ages out and is discarded from the Layer 2 MAC address table. The range is from 120 to 918000; the default is 1800 seconds. Entering the value 0 disables the MAC aging.

Step 3	**exit**
	Example:
	`switch(config)# exit`
	`switch#`
		Exits the configuration mode.

Step 4	(Optional) **show mac address-table aging-time** **Example:** `switch# show mac address-table aging-time`	Displays the aging time configuration for MAC address retention.
Step 5	(Optional) **copy running-config startup-config** **Example:** `switch# copy running-config startup-config`	Copies the running configuration to the startup configuration.

Example

This example shows how to set the ageout time for entries in the Layer 2 MAC address table to 600 seconds (10 minutes):

```
switch# config t
switch(config)# mac address-table aging-time 600
switch(config)#
```

Checking Consistency of MAC Address Tables

You can check the match between the MAC address table on the supervisor and all the modules.

SUMMARY STEPS

1. **show consistency-checker l2 module** <*slot_number*>

DETAILED STEPS

	Command or Action	Purpose

Step 1	**show consistency-checker l2 module** <*slot_number*>
	Example:
	`switch# show consistency-checker l2 module 7`
	`switch#`
		Displays the discrepant, missing, and extra MAC addresses between the supervisor and the specified module.

Example

This example shows how to display discrepant, missing, and extra entries in the MAC address tables between the supervisor and the specified module:

```
switch# show consistency-checker l2 module 7
switch#
```

Clearing Dynamic Addresses from the MAC Table

You can clear all dynamic Layer 2 entries in the MAC address table. (You can also clear entries by designated interface or VLAN.)

SUMMARY STEPS

1. **clear mac address-table dynamic** {**address** *mac_addr*} {**interface** [**ethernet** *slot/port* | **port-channel** *channel-number*]} {**vlan** *vlan_id*}
2. (Optional) **show mac address-table**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**clear mac address-table dynamic** {**address** *mac_addr*} {**interface** [**ethernet** *slot/port* | **port-channel** *channel-number*]} {**vlan** *vlan_id*}
	Example:
	`switch# clear mac address-table dynamic`
		Clears the dynamic address entries from the MAC address table in Layer 2.

Step 2	(Optional) **show mac address-table** **Example:** `switch# show mac address-table`	Displays the MAC address table.

Example

This example shows how to clear the dynamic entries in the Layer 2 MAC address table:

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

Configuring Dynamic MAC Address Limits Per VLAN

You can set the limit on number of dynamic MAC entries per VLAN to protect the control plane from MAC flood attacks.

Currently this configuration is supported only on Cisco Nexus 9500 switches with 9600-R/RX/R2 line cards. Beginning with Cisco NX-OS Release 10.4(2)F, this configuration support is extended to Cisco Nexus 9300-EX/FX/FX2/FX3/GX/GX2/H2R/H1 platform switches.

Note The configuration is supported only on the default template and is not supported with the L2 heavy template.

Before you begin

This is global level configuration. However, VLAN must be created before specifying the limit on VLAN.

SUMMARY STEPS

1. **config t**
2. **vlan** {*vlan-id* | *vlan-range*}
3. **mac address-table limit vlan** *vlan-id limit -value*
4. **exit**
5. **exit**
6. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**config t**
	Example:
	`switch# config t`
	`switch(config)#`
		Enters configuration mode.

Step 2	**vlan** {*vlan-id* | *vlan-range*}
	Example:
	`switch(config)# vlan 5`
	`switch(config-vlan)#`
		Places you into the VLAN configuration submode. If the VLAN does not exist, the system creates the specified VLAN and then enters the VLAN configuration submode.

Step 3	**mac address-table limit vlan** *vlan-id limit -value* **Example:** `switch(config-vlan)# mac address-table limit vlan 40 108`	Specifies the VLAN to which the MAC address limits should be applied. The permissible value for the limit is 100–196000.

	Command or Action	Purpose
		Note • This command is not supported on EoRs. • This command must not be used with vPC and VXLAN. • Enabling/Disabling the Mac-limit or modification of the mac-limit leads to flushing of all the dynamic MACs learned on that VLAN. However, the static or gateway MACs learning are not impacted. • The user is prompted to confirm before flushing.

Step 4	**exit**
	Example:
	`switch(config-vlan)# exit`
	`switch(config)#`
		Exits the VLAN configuration mode.

Step 5	**exit**
	Example:
	`switch(config)# exit`
	`switch#`
		Exits the configuration mode.

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

Configuring L2 Heavy Mode

The purpose of this feature is to increase the current 96k MAC address scale to 200k by carving out a new L2-heavy template, changing FP tile hardware resource allocations, making necessary control plane changes and ISSU restore support to accommodate new scale.

Command	Purpose
sh system routing mode	Shows the configured and applied mode

Command	Purpose
no system routing template-l2-heavy	Enables 200K MAC. 200K MAC is enabled only when this mode is configured and the system is reloaded. Use **no** form of this command to to disable this feature. **Note** Beginning with Cisco NX-OS Release 10.2(2)F, 200K MAC is supported on Cisco N9K-9332D-GX2B platform switches.
sh run | i system	Runs the applied mode

Guidelines & Limitations:

- This feature is supported for Layer 2 unidimensional scale only. SVI, Layer 3 interface, and VXLAN VLANs are not supported.
- Beginning with Cisco NX-OS Release 9.2(3), this feature is supported on the following platforms: N9K-C9264PQ, N9K-C9272Q, N9K-C9236C, N9K-C92300YC, N9K-C92304QC, N9K-C9232C, N9K-C92300YC and 9300-EX
- Beginning with Cisco NX-OS Release 10.2(2)F, the 200K MAC feature is supported on Cisco N9K-9332D-GX2B platform switches.

Following is an example for configuring L2 heavy mode:

```
switch (config)# sh system routing mode
switch# Configured System Routing Mode: L2 Heavy
switch# Applied System Routing Mode: L2 Heavy
switch# switch# switch# sh run | i system
switch# system routing template-l2-heavy
```

Verifying the Layer 2 Switching Configuration

To display Layer 2 switching configuration information, perform one of the following tasks:

Command	Purpose
show mac address-table	Displays information about the MAC address table.
show mac address-table limit	Displays information about the limits set for the MAC address table.

Command	Purpose
show mac address-table aging-time	Displays information about the aging time set for the MAC address entries. **Note** Beginning with Cisco NX-OS Release 10.2(1), Cisco Nexus 9000 and Nexus 3000 switches that use Cloudscale ASICs do not report the MAC age in show mac address outputs. The age column can be ignored as, instead of a fixed value of 0s that was reported in earlier releases, now, NA is reported. This is only a display limitation. MAC aging is still functionally enforced.
show mac address-table static	Displays information about the static entries on the MAC address table.
show mac address-table limit vlan	Displays information about the VLANs configured with MAC learn limit.
show interface [*interface*] **mac-address**	Displays the MAC addresses and the burn-in MAC address for the interfaces.
show forwarding consistency l2 {*module*}	Displays discrepant, missing, and extra MAC addresses between the tables on the module and the supervisor.

Configuration Example for Layer 2 Switching

The following example shows how to add a static MAC address and how to modify the default global aging time for MAC addresses:

```
switch# configure terminal
switch(config)# mac address-table static 0000.0000.1234 vlan 10 interface ethernet 2/15
switch(config)# mac address-table aging-time 120
```

The following example shows how to configure dynamic MAC limit per VLAN:

```
switch(config)# mac address-table limit vlan 251 100
Configuring MAC address limit will result in flushing existing Macs in the specified
VLAN/System. Proceed (yes/no)? [n] yes
Warning : MAC limit per VLAN feature isn't supported along with VPC/VxLAN. Please remove
the config if VPC/VxLAN config is present in this system !!!
switch(config)#
switch(config)# mac address-table limit vlan 252-253 100
Configuring MAC address limit will result in flushing existing Macs in the specified
VLAN/System. Proceed (yes/no)? [n] yes
switch(config)#
switch(config)# mac address-table limit vlan 254 300
Configuring MAC address limit will result in flushing existing Macs in the specified
VLAN/System. Proceed (yes/no)? [n] yes
```

Note When you configure this feature for the first time on the switch, a warning message stating that this feature is not supported with vPC/VXLAN is displayed. This warning message will not be displayed for the subsequent configurations.

To check the configured dynamic MAC limit and current count, use the following show command:

```
switch# show mac address-table limit vlan

Vlan       Conf Limit     Curr Count
----      ------------    ----------
 251         100             100
 252         100             100
 253         100             75
 254         300             60
```

Additional References for Layer 2 Switching -- CLI Version

Related Documents

Related Topic	Document Title
Static MAC addresses	*Cisco Nexus 9000 Series NX-OS Security Configuration Guide*
Interfaces	*Cisco Nexus 9000 Series NX-OS Interfaces Configuration Guide*
High availability	*Cisco Nexus 9000 Series NX-OS High Availability and Redundancy Guide*
System management	*Cisco Nexus 9000 Series NX-OS System Management Configuration Guide*