



Configuring Layer 2 Switching

This chapter describes how to configure Layer 2 switching using Cisco NX-OS.

This chapter includes the following sections:

- [Finding Feature Information, on page 1](#)
- [Information About Layer 2 Switching, on page 1](#)
- [Prerequisites for Configuring MAC Addresses, on page 5](#)
- [Guidelines and Limitations for Configuring MAC Addresses, on page 5](#)
- [Default Settings for Layer 2 Switching, on page 6](#)
- [Configuring Layer 2 Switching, on page 6](#)
- [Verifying the Layer 2 Switching Configuration, on page 13](#)
- [Configuration Example for Layer 2 Switching, on page 14](#)
- [Additional References for Layer 2 Switching, on page 14](#)
- [Feature History for Configuring Layer 2 Switching, on page 15](#)

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

Information About Layer 2 Switching



Note See the *Cisco Nexus 7000 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 7000 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 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. 1/10-Gigabit Ethernet operates in full duplex only.

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. Beginning with Cisco NX-OS Release 5.2(1), multicast MAC addresses can be configured as static MAC addresses. For further information, see the “Configuring IGMP Snooping” of the *Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide*. The static MAC entries are retained across a reboot of the device.

Beginning with Cisco NX-OS Release 4.1(5), 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 7000 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.

See the *Cisco Nexus 7000 Series NX-OS Security Command Reference* for information on MAC port security.

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. Beginning with Cisco NX-OS 4.1(2), when you enter the **show forwarding consistency I2** command, the device displays discrepant, missing, and extra MAC address entries.

Layer 3 Static MAC Addresses

Beginning with Release 4.2, you can configure a static MAC address for all Layer 3 interfaces. The default MAC address for the Layer 3 interfaces is the VDC MAC address.

You can configure a static MAC address for the following Layer 3 interfaces:

- Layer 3 interfaces
- Layer 3 subinterfaces
- Layer 3 port channels
- VLAN network interface



Note You cannot configure static MAC address on tunnel interfaces.

See the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide* for information on configuring Layer 3 interfaces.

High Availability for Switching

You can upgrade or downgrade the software seamlessly, with respect to classical Ethernet switching. Beginning with Release 4.2(1), 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 7000 Series NX-OS High Availability and Redundancy Guide* for complete information on high availability features.

Virtualization Support for Layer 2 Switching

The device supports virtual device contexts (VDCs), and the configuration and operation of the MAC address table are local to the VDC.



Note See the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide* for complete information on VDCs and assigning resources.

MAC Address Movement

Rapid MAC address movement, caused by either Layer 2 loop or other system events (for example: misconfiguration, dual-active server cluster, and so on), if not limited could eventually overload the supervisor and potentially impact other processes. Such situation might lead to an overall instability of the control plane. To avoid this situation rapid MAC move protection has been implemented in the processes that handle MAC addresses learning.

MAC Move Protection

The following methods protect the SUP from excessive mac move:

- Software throttle: Using **mac address loop-detect flow-control-fe** command.
- Hardware throttle: Using **mac address loop-detect disable-learn-vlan** command.

Software throttling is enabled by default and this is the recommended method. You can use only one throttling method at a time. The throttling commands should be executed within a VDC.

Software Throttle

In software throttle, the mac-move notifications are throttled so the rate of mac-move notification is limited from the module to the supervisor.

This throttling is usually done per Forwarding Engine [FE] (per ASIC level) on a specific module. If necessary (for example: during rapid mac move across all modules in the system) global throttling is invoked that would throttle notification from all FEs on all modules in order protect the supervisor.

Hardware Throttle

In hardware throttle, mac-learning is disabled on a particular VLAN (for all FE and all modules) for specific time and then re-enabled. This throttling can be done per VLAN level (per VLAN throttle) or for all VLANs (global throttle).

Increasing the Throttle

In case the software throttle is found to be inadequate, in extreme cases, the mac-move information sent from the line card module is reduced.

This method is not a recommended option and should be exercised with caution.



Note Increasing the threshold could make the system unstable if not set accordingly to the device scale.

The reduction in mac-move information sent is done in two ways:

- Reduce number of notifications that can be batched.
- Change/increase the time-period after which this notification batch can be sent from the module to the supervisor module.

Use the **mac address throttle-buffer-intv { max |optimal}** command (to be executed within a VDC) to increase the throttle by tuning the throttle buffer and the scan duration on the line card module.

When the **max** keyword is used, the throttling is maximum. It means information sent from the line card module to the supervisor is reduced and are spaced out more.

When the **optimal** keyword is used, the throttling is medium.

When this command is not used, the throttling is minimum (which is the default).

Prerequisites for Configuring MAC Addresses

MAC addresses have the following prerequisites:

- You must be logged onto the device.

Guidelines and Limitations for Configuring MAC Addresses

MAC addresses have the following configuration guidelines and limitations:

| MAC Address Table | Age Group |
|-----------------------|---------------------------|
| M1 Line Cards | 128,000 entries |
| F1 Line Cards | 16,000 to 256,000 entries |
| F2 and F2e Line Cards | 16,000 to 192,000 entries |



Note

The F2 and F2e modules synchronize the MAC address tables for a VLAN across all Switch on Chips (SoCs) present in a virtual device context (VDC) when a switch virtual interface (SVI) for the VLAN is configured. Synchronizing the MAC address tables can reduce the number of MAC addresses supported in a VDC to 16,000.

Beginning with NX-OS Release 6.0.1, the learning mode feature is supported. Learning mode has the following configuration guidelines and limitations:

| Line Cards | Classic Ethernet (CE) Nonconversational Learning Supported | Classic Ethernet (CE) Conversational Learning Supported | Fabric Path Conversational Learning | Fabric Path Nonconversational Learning |
|------------|--|---|-------------------------------------|---|
| M1 | Yes | NA | NA | NA |
| F1 | Yes | Yes | Yes | No |
| F2 and F2e | Yes | Yes | Yes | Yes, if the switch virtual interface (SVI) is configured. |



Note When you configure a static MAC address on a vPC switch, ensure to configure a corresponding static MAC address on the other vPC switch. If you configure the static MAC address only on one of the vPC switches, the other vPC switch will not learn the MAC address dynamically.

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 |

Beginning with NX-OS Release 6.0.1, the learning mode feature is supported. This table lists the default learning mode parameters.

Table 2: Default Learning Mode Parameters

| Parameters | Default |
|----------------------------|-------------------|
| Classic Ethernet (CE) VLAN | Nonconversational |
| Fabric Path VLANs | Conversational |

Configuring Layer 2 Switching



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 addresses as static MAC addresses. Beginning with Cisco NX-OS Release 5.2(1), multicast MAC addresses can be configured as static MAC addresses. For further information, see the "Configuring IGMP Snooping" of the *Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide*.

Before you begin

Before you configure static MAC addresses, ensure that you are in the correct VDC (or enter the **switchto vdc** command).

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | config t Example: switch# config t switch(config)# | Enters global configuration mode. |
| Step 2 | mac address-table static <i>mac-address</i> vlan <i>vlan-id</i> {[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 global 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 a Static MAC Address on a Layer 3 Interface

Beginning with Release 4.2(1), you can configure static MAC addresses on Layer 3 interfaces. You cannot configure broadcast addresses as static MAC addresses. Beginning with Cisco NX-OS Release 5.2(1), multicast

MAC addresses can be configured as static MAC addresses. For further information, see the "Configuring IGMP Snooping" of the *Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide*.



Note You cannot configure static MAC addresses on tunnel interfaces.

See the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide* for information on configuring Layer 3 interfaces.

Before you begin

Before you configure static MAC addresses, ensure that you are in the correct VDC (or enter the **switchto vdc** command).

Procedure

| | Command or Action | Purpose |
|---------------|--|---|
| Step 1 | config t Example: switch# config t switch(config)# | Enters global configuration mode. |
| Step 2 | interface [ethernet slot/port ethernet slot/port.number port-channel number vlan vlan-id] Example: switch(config)# interface ethernet 7/3 | Specifies the Layer 3 interface and enters interface configuration mode. Note You must create the Layer 3 interface before you can assign the static MAC address. |
| Step 3 | mac-address mac-address Example: switch(config-if)# mac-address 22ab.47dd.ff89 switch(config-if)# | Specifies a static MAC address to add to the Layer 3 interface. |
| Step 4 | exit Example: switch(config-if)# exit switch(config)# | Exits interface configuration mode. |
| Step 5 | (Optional) show interface [ethernet slot/port ethernet slot/port.number port-channel number vlan vlan-id] Example: switch# show interface ethernet 7/3 | Displays information about the Layer 3 interface. |
| Step 6 | (Optional) copy running-config startup-config Example: | Copies the running configuration to the startup configuration. |

| | Command or Action | Purpose |
|--|--|---------|
| | switch# copy running-config startup-config | |

Example

This example shows how to configure the Layer 3 interface on slot 7, port 3 with a static MAC address:

```
switch# config t
switch(config)# interface ethernet 7/3
switch(config-if)# mac-address 22ab.47dd.ff89
switch(config-if)#
```

Configuring the Aging Time for the MAC Address 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 address table, which contains the Layer 2 information.



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

Before you begin

Before you configure the aging time for the MAC address table, ensure that you are in the correct VDC (or enter the **switchto vdc** command).

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | config t Example: switch# config t switch(config)# | Enters global configuration mode. |
| Step 2 | mac address-table aging-time seconds [vlan vlan_id] 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 global configuration mode. |

| | Command or Action | Purpose |
|---------------|---|--|
| 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)#
```

Configuring Learning Mode for VLANs

Beginning with NX-OS Release 6.0.1, configuring the learning mode for VLANs is supported. Based on the learning mode configured, the Cisco NX-OS software can install MAC addresses in hardware either conversationally or nonconversationally.

Before you begin

Before you configure the learning mode for VLANs, ensure that you are in the correct VDC (or enter the **switchto vdc** command).

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | config t Example: switch# config t switch(config)# | Enters global configuration mode. |
| Step 2 | mac address-table learning-mode conversational <i>vlan-range of CE-vlans</i> Example: switch(config)# mac address-table learning-mode conversational vlan1 | Specifies the learning mode for the Layer 2 MAC address table. The options are conversational learning and nonconversational learning. |

| | Command or Action | Purpose |
|---------------|---|----------------------------------|
| Step 3 | exit Example: <pre>switch(config)# exit switch#</pre> | Exits global configuration mode. |

Example

This example shows how to set the learning mode to conversational for the VLANs:

```
switch# config t
switch(config)# mac address-table learning-mode conversational vlan1
switch(config)# end
switch(config)# show mac address-table learning-mode
```

Enabling MAC Move Protection

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | config t Example: <pre>switch# config t switch(config)#</pre> | Enters global configuration mode. |
| Step 2 | mac address loop-detect flow-control-fe global-thresh-time <i>threshold-time</i> global-thresh-count <i>threshold-count</i> Example: <pre>switch(config)# mac address loop-detect flow-control-fe global-thresh-time 5 global-thresh-count 500</pre> | Enables FE-based flow control to turn on the software throttle for mac-move protection for all FEs on all line cards. |
| Step 3 | mac address loop-detect flow-control-fe threshold-time <i>threshold-time</i> threshold-count <i>threshold-count</i> Example: <pre>switch(config)# mac address-table loop-detect flow-control-fe threshold-time 5 threshold-count 500</pre> | Enables FE-based flow control to turn on the software throttle for mac-move protection for a specific FE (per ASIC level). |
| Step 4 | mac address loop-detect disable-learn-vlan global-thresh-time <i>threshold-time</i> global-thresh-count <i>threshold-count</i> Example: | Disables the mac-learning for all VLANs (global throttle). |

| | Command or Action | Purpose |
|---------------|--|---|
| | <pre>switch(config)# mac address loop-detect disable-learn-vlan global-thresh-time 5 global-thresh-count 500</pre> | |
| Step 5 | <p>mac address loop-detect disable-learn-vlan threshold-time <i>threshold-time</i> threshold-count <i>threshold-count</i></p> <p>Example:</p> <pre>switch(config)# mac address loop-detect disable-learn-vlan-thresh-time 5 thresh-count 500</pre> | Disables the mac-learning per VLAN (per VLAN throttle). |
| Step 6 | <p>mac address throttle-buffer-intv max</p> <p>Example:</p> <pre>switch(config)# mac address throttle-buffer-intv max</pre> | Uses maximum scan interval and buffer size to increase/decrease the throttle; and to effect maximum throttle. |
| Step 7 | <p>mac address throttle-buffer-intv optimal</p> <p>Example:</p> <pre>switch(config)# mac address throttle-buffer-intv optimal</pre> | Uses optimal scan interval and buffer size to throttle at a medium level.. |
| Step 8 | <p>exit</p> <p>Example:</p> <pre>switch(config)# exit switch#</pre> | Exits global configuration mode. |

Checking the Consistency of MAC Address Tables

Beginning with Release 4.1(2), you can check the match between the MAC address table on the supervisor and all the modules.

Procedure

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | <p>show forwarding consistency l2 <i>{module_number}</i></p> <p>Example:</p> <pre>switch# show forwarding consistency l2 7 switch#</pre> | 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 forwarding consistency 12 7
switch#
```

Clearing Dynamic Addresses from the MAC Address Table

You can clear all dynamic Layer 2 entries in the MAC address table.

Before you begin

Before you clear the dynamic MAC address table, ensure that you are in the correct VDC (or enter the **switchto vdc** command).

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | clear mac address-table dynamic {address <i>mac_addr</i> } {interface [ethernet <i>slot/port</i> loopback <i>number</i> port-channel <i>channel-number</i>]} {vlan <i>vlan_id</i> } 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#
```

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 aging-time | Displays information about the aging time set for the MAC address entries. |
| show mac address-table static | Displays information about the static entries on the MAC address table. |

| Command | Purpose |
|---|--|
| show interface [<i>interface</i>] mac-address | Displays the MAC addresses and the burned in MAC addresses for the interfaces. |
| show forwarding consistency l2 { <i>module</i> } | Displays discrepant, missing, and extra MAC addresses between the tables on the module and the supervisor. |

For information on the output of these commands, see the *Cisco Nexus 7000 Series NX-OS Layer 2 Switching Command Reference*.

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
```

Additional References for Layer 2 Switching

Related Documents

| Related Topic | Document Title |
|-------------------------------------|---|
| Port security, static MAC addresses | <i>Cisco Nexus 7000 Series NX-OS Security Configuration Guide</i> |
| Interfaces | <i>Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide</i> |
| Command reference | <i>Cisco Nexus 7000 Series NX-OS Layer 2 Switching Command Reference</i> |
| High availability | <i>Cisco Nexus 7000 Series NX-OS High Availability and Redundancy Guide</i> |
| VDCs | <i>Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide</i> |
| System management | <i>Cisco Nexus 7000 Series NX-OS System Management Configuration Guide</i> |
| Licensing | <i>Cisco NX-OS Licensing Guide</i> |
| Release Notes | <i>Cisco Nexus 7000 Series NX-OS Release Notes</i> |

Standards

| Standards | Title |
|---|-------|
| No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature. | — |

Feature History for Configuring Layer 2 Switching

This table lists the release history for this feature.

Table 3: Feature History for Configuring Layer 2 Switching

| Feature Name | Releases | Feature Information |
|--|----------|---|
| MAC move protection | 8.2(3) | MAC move protection using software throttle and hardware throttle is supported. |
| Learning mode for VLANs | 6.0(1) | You can configure conversational or nonconversational learning mode for VLANs. |
| Layer 3 interface static MAC addresses | 4.2(1) | You can configure a Layer 3 interface with a static MAC address. |
| show mac address-table | 4.1(2) | This display provides additional information when vPC is enabled and running. |
| Layer 2 consistency | 4.1(2) | The show forwarding consistency l2 command displays inconsistent entries on the MAC address table between the modules. |

