



Configuring Layer 2 Interfaces

- [Information About Ethernet Interfaces, on page 1](#)
- [Guidelines and Limitations for Layer 2 Interfaces, on page 3](#)
- [Interface Speed, on page 3](#)
- [40-Gigabit Ethernet Interface Speed, on page 4](#)
- [SVI Autostate, on page 4](#)
- [Cisco Discovery Protocol, on page 5](#)
- [Error-Disabled State, on page 5](#)
- [Default Interfaces, on page 6](#)
- [Debounce Timer Parameters, on page 6](#)
- [MTU Configuration, on page 7](#)
- [Default Physical Ethernet Settings , on page 8](#)
- [Displaying Interface Information, on page 9](#)

Information About Ethernet Interfaces

The Ethernet ports can operate as standard Ethernet interfaces connected to servers or to a LAN.

The Ethernet interfaces are enabled by default.

Interface Command

You can enable the various capabilities of the Ethernet interfaces on a per-interface basis using the **interface** command. When you enter the **interface** command, you specify the following information:

The interface numbering convention is extended to support use with a Cisco Nexus Fabric Extender as follows:

switch(config)# **interface ethernet** [*chassis*]/*slot*/*port*

- The chassis ID is an optional entry that you can use to address the ports of a connected Fabric Extender. The chassis ID is configured on a physical Ethernet or EtherChannel interface on the switch to identify the Fabric Extender discovered through the interface. The chassis ID ranges from 100 to 199.

Unidirectional Link Detection Parameter

The Cisco-proprietary Unidirectional Link Detection (UDLD) protocol allows ports that are connected through fiber optics or copper (for example, Category 5 cabling) Ethernet cables to monitor the physical configuration of the cables and detect when a unidirectional link exists. When the switch detects a unidirectional link, UDLD shuts down the affected LAN port and alerts the user. Unidirectional links can cause a variety of problems, including spanning tree topology loops.


UDLD is a Layer 2 protocol that works with the Layer 1 protocols to determine the physical status of a link. At Layer 1, autonegotiation takes care of physical signaling and fault detection. UDLD performs tasks that autonegotiation cannot perform, such as detecting the identities of neighbors and shutting down misconnected LAN ports. When you enable both autonegotiation and UDLD, Layer 1 and Layer 2 detections work together to prevent physical and logical unidirectional connections and the malfunctioning of other protocols.

A unidirectional link occurs whenever traffic transmitted by the local device over a link is received by the neighbor but traffic transmitted from the neighbor is not received by the local device. If one of the fiber strands in a pair is disconnected, and if autonegotiation is active, the link does not stay up. In this case, the logical link is undetermined, and UDLD does not take any action. If both fibers are working normally at Layer 1, then UDLD at Layer 2 determines whether those fibers are connected correctly and whether traffic is flowing bidirectionally between the correct neighbors. This check cannot be performed by autonegotiation, because autonegotiation operates at Layer 1.

A Cisco Nexus device periodically transmits UDLD frames to neighbor devices on LAN ports with UDLD enabled. If the frames are echoed back within a specific time frame and they lack a specific acknowledgment (echo), the link is flagged as unidirectional and the LAN port is shut down. Devices on both ends of the link must support UDLD in order for the protocol to successfully identify and disable unidirectional links.

The following figure shows an example of a unidirectional link condition. Device B successfully receives traffic from Device A on the port. However, Device A does not receive traffic from Device B on the same port. UDLD detects the problem and disables the port.

Figure 1: Unidirectional Link



Default UDLD Configuration

The following table shows the default UDLD configuration.

Table 1: UDLD Default Configuration

Feature	Default Value
UDLD global enable state	Globally disabled
UDLD aggressive mode	Disabled
UDLD per-port enable state for fiber-optic media	Enabled on all Ethernet fiber-optic LAN ports

Feature	Default Value
UDLD per-port enable state for twisted-pair (copper) media	Disabled on all Ethernet 10/100 and 1000BASE-TX LAN ports

UDLD Aggressive and Nonaggressive Modes

UDLD aggressive mode is disabled by default. You can configure UDLD aggressive mode only on point-to-point links between network devices that support UDLD aggressive mode. If UDLD aggressive mode is enabled, when a port on a bidirectional link that has a UDLD neighbor relationship established stops receiving UDLD frames, UDLD tries to reestablish the connection with the neighbor. After eight failed retries, the port is disabled.

To prevent spanning tree loops, nonaggressive UDLD with the default interval of 15 seconds is fast enough to shut down a unidirectional link before a blocking port transitions to the forwarding state (with default spanning tree parameters).

When you enable the UDLD aggressive mode, the following occurs:

- One side of a link has a port stuck (both transmission and receive)
- One side of a link remains up while the other side of the link is down

In these cases, the UDLD aggressive mode disables one of the ports on the link, which prevents traffic from being discarded.

Guidelines and Limitations for Layer 2 Interfaces

Layer 2 interfaces have the following configuration guidelines and limitations:

- Auto-negotiation is not supported.
- 1G autonegotiation not supported on N3K-C36180YC-R and N9K-X96136YC-R switches. To work around this issue, you must manually set speed to 1000. If autonegotiation is enabled on the neighbors, you must disable autonegotiation on those neighbors.
- On Cisco Nexus N3K-C3636C-R and N3K-C36180YC-R switches, auto-negotiation may not work on ports 49-64 when bringing up 100G links using QSFP-100G-CR4 cable. To work around this issue, you must hard-code the speed on ports 49-64 and disable auto-negotiation

Interface Speed

Cisco Nexus 36180YC-R switches have 48 small form-factor pluggable (SFP) ports with a default speed of 10 G and 6 quad small form-factor pluggable (QSFP) ports with a default speed of 100 G. 48 SFP interface ports can support 25 G, 10 G, 1 G speeds. 6 QSFP interface ports can support 100 G and 40 G speeds.

In the first 48 ports, each 4 ports in the port group must have the same speed configured. You cannot configure one port at a time which might result in an error. For more information, see [CSCve80686](#).

Table 2: Breakout Modes Support Matrix

Switches	4x10G	4x25G	2x50G
N3K-C3636C-R	Yes	Yes	Yes
N3K-C36180YC-R	Yes	Yes	Yes

40-Gigabit Ethernet Interface Speed

You can operate QSFP ports as either 40-Gigabit Ethernet or 4x10-Gigabit Ethernet modes on Cisco Nexus 3600 platform switches. By default, there are 6 QSFP interface ports numbered 49 to 54 which can operated in 40-Gigabit Ethernet mode. These 40-Gigabit Ethernet ports are numbered in a 2-tuple naming convention. For example, the second 40-Gigabit Ethernet port is numbered as 1/50. The process of changing the configuration from 40-Gigabit Ethernet to 10-Gigabit Ethernet is called breakout and the process of changing the configuration from 10-Gigabit Ethernet to Gigabit Ethernet is called breakin. When you break out a 40-Gigabit Ethernet port into 10-Gigabit Ethernet ports, the resulting ports are numbered using a 3-tuple naming convention. For example, the break-out ports of the second 40-Gigabit Ethernet port are numbered as 1/49/1, 1/49/2, 1/49/3, 1/49/4.



Note The breakout ports are in administratively enabled state after the breakout of the 40G ports into 4x10G mode or the breaking of the 100G ports into 4x25G mode. On upgrade from the earlier releases, the configuration restored takes care of restoring the appropriate administrative state of the ports.



Note When you break out from 40-Gigabit Ethernet to 10-Gigabit Ethernet, or break in from 10-Gigabit Ethernet to 40-Gigabit Ethernet, all interface configurations are reset, and the affected ports are administratively unavailable. To make these ports available, use the **no shut** command.



Note A new QSFP+ 40-Gb transceiver is supported on the Cisco Nexus 3600 platform switches. The new QSFP+ (40-Gb) transceiver has a cable that splits into four 10Gb SFP-10G-LR transceivers. To use it, you need the port to be in 4x10G mode. If you are using the breakout cable, you need to run that 40G port in 4x10G mode.

The ability to break out a 40-Gigabit Ethernet port into four 10-Gigabit Ethernet ports and break in four 10-Gigabit Ethernet ports into a 40-Gigabit Ethernet port dynamically allows you to use any of the breakout-capable ports to work in the 40-Gigabit Ethernet or 10-Gigabit Ethernet modes without permanently defining them.

SVI Autostate

The Switch Virtual Interface (SVI) represents a logical interface between the bridging function and the routing function of a VLAN in the device. By default, when a VLAN interface has multiple ports in the VLAN, the SVI goes to the down state when all the ports in the VLAN go down.

Autostate behavior is the operational state of an interface that is governed by the state of the various ports in its corresponding VLAN. An SVI interface on a VLAN comes up when there is at least one port in that vlan that is in STP forwarding state. Similarly, this interface goes down when the last STP forwarding port goes down or goes to another STP state.

By default, Autostate calculation is enabled. You can disable Autostate calculation for an SVI interface and change the default value.

Cisco Discovery Protocol

The Cisco Discovery Protocol (CDP) is a device discovery protocol that runs over Layer 2 (the data link layer) on all Cisco-manufactured devices (routers, bridges, access servers, and switches) and allows network management applications to discover Cisco devices that are neighbors of already known devices. With CDP, network management applications can learn the device type and the Simple Network Management Protocol (SNMP) agent address of neighboring devices that are running lower-layer, transparent protocols. This feature enables applications to send SNMP queries to neighboring devices.

CDP runs on all media that support Subnetwork Access Protocol (SNAP). Because CDP runs over the data-link layer only, two systems that support different network-layer protocols can learn about each other.

Each CDP-configured device sends periodic messages to a multicast address, advertising at least one address at which it can receive SNMP messages. The advertisements also contain time-to-live, or holdtime information, which is the length of time a receiving device holds CDP information before discarding it. Each device also listens to the messages sent by other devices to learn about neighboring devices.

The switch supports both CDP Version 1 and Version 2.

Default CDP Configuration

The following table shows the default CDP configuration.

Table 3: Default CDP Configuration

Feature	Default Setting
CDP interface state	Enabled
CDP timer (packet update frequency)	60 seconds
CDP holdtime (before discarding)	180 seconds
CDP Version-2 advertisements	Enabled

Error-Disabled State

An interface is in the error-disabled (err-disabled) state when the interface is enabled administratively (using the **no shutdown** command) but disabled at runtime by any process. For example, if UDLD detects a unidirectional link, the interface is shut down at runtime. However, because the interface is administratively enabled, the interface status displays as err-disabled. Once an interface goes into the err-disabled state, you

must manually reenable it or you can configure an automatic timeout recovery value. The err-disabled detection is enabled by default for all causes. The automatic recovery is not configured by default.

When an interface is in the err-disabled state, use the **errdisable detect cause** command to find information about the error.

You can configure the automatic err-disabled recovery timeout for a particular err-disabled cause by changing the time variable.

The **errdisable recovery cause** command provides automatic recovery after 300 seconds. To change the recovery period, use the **errdisable recovery interval** command to specify the timeout period. You can specify 30 to 65535 seconds.

To disable recovery of an interface from the err-disabled state, use the **no errdisable recovery cause** command.

The various options for the **errdisable recover cause** command are as follows:

- **all**—Enables a timer to recover from all causes.
- **bpduguard**—Enables a timer to recover from the bridge protocol data unit (BPDU) Guard error-disabled state.
- **failed-port-state**—Enables a timer to recover from a Spanning Tree Protocol (STP) set port state failure.
- **link-flap**—Enables a timer to recover from linkstate flapping.
- **pause-rate-limit**—Enables a timer to recover from the pause rate limit error-disabled state.
- **udld**—Enables a timer to recover from the Unidirectional Link Detection (UDLD) error-disabled state.
- **loopback**—Enables a timer to recover from the loopback error-disabled state.

If you do not enable the err-disabled recovery for the cause, the interface stays in the err-disabled state until you enter the **shutdown** and **no shutdown** commands. If the recovery is enabled for a cause, the interface is brought out of the err-disabled state and allowed to retry operation once all the causes have timed out. Use the **show interface status err-disabled** command to display the reason behind the error.

Default Interfaces

You can use the default interface feature to clear the configured parameters for both physical and logical interfaces such as the Ethernet, loopback, management, VLAN, and the port-channel interface.

Debounce Timer Parameters

The debounce timer delays notification of a link change, which can decrease traffic loss due to network reconfiguration. You can configure the debounce timer separately for each Ethernet port and specify the delay time in milliseconds. The delay time can range from 0 milliseconds to 5000 milliseconds. By default, this parameter is set for 100 milliseconds, which results in the debounce timer not running. When this parameter is set to 0 milliseconds, the debounce timer is disabled.

**Caution**

Enabling the debounce timer causes the link-down detections to be delayed, which results in a loss of traffic during the debounce period. This situation might affect the convergence and reconvergence of some Layer 2 and Layer 3 protocols.

MTU Configuration

The switch does not fragment frames. As a result, the switch cannot have two ports in the same Layer 2 domain with different maximum transmission units (MTUs). A per-physical Ethernet interface MTU is not supported. Instead, the MTU is set according to the QoS classes. You modify the MTU by setting class and policy maps.

**Note**

When you show the interface settings, a default MTU of 1500 is displayed for physical Ethernet interfaces.

Counter Values

See the following information on the configuration, packet size, incremented counter values, and traffic.

Configuration	Packet Size	Incremented Counters	Traffic
L2 port – without any MTU configuration	6400 and 10000	Jumbo, giant, and input error	Dropped
L2 port – with jumbo MTU 9216 in network-qos configuration	6400	Jumbo	Forwarded
L2 port – with jumbo MTU 9216 in network-qos configuration	10000	Jumbo, giant, and input error	Dropped
Layer 3 port with default Layer 3 MTU and jumbo MTU 9216 in network-qos configuration	6400	Jumbo	Packets are punted to the CPU (subjected to CoPP configs), get fragmented, and then they are forwarded by the software.
Layer 3 port with default Layer 3 MTU and jumbo MTU 9216 in network-qos configuration	6400	Jumbo	Packets are punted to the CPU (subjected to CoPP configs), get fragmented, and then they are forwarded by the software.

Configuration	Packet Size	Incremented Counters	Traffic
Layer 3 port with default Layer 3 MTU and jumbo MTU 9216 in network-qos configuration	10000	Jumbo, giant, and input error	Dropped
Layer 3 port with jumbo Layer 3 MTU and jumbo MTU 9216 in network-qos configuration	6400	Jumbo	Forwarded without any fragmentation.
Layer 3 port with jumbo Layer 3 MTU and jumbo MTU 9216 in network-qos configuration	10000	Jumbo, giant, and input error	Dropped
Layer 3 port with jumbo Layer 3 MTU and default L2 MTU configuration	6400 and 10000	Jumbo, giant, and input error	Dropped

**Note**

- Under 64 bytes packet with good CRC—The short frame counter increments.
- Under 64 bytes packet with bad CRC—The runts counter increments.
- Greater than 64 bytes packet with bad CRC—The CRC counter increments.

Downlink Delay

You can operationally enable uplink SFP+ ports before downlink RJ-45 ports after a reload on a Cisco Nexus 3048 switch. You must delay enabling the RJ-45 ports in the hardware until the SFP+ ports are enabled.

You can configure a timer that during reload enables the downlink RJ-45 ports in hardware only after the specified timeout. This process allows the uplink SFP+ ports to be operational first. The timer is enabled in the hardware for only those ports that are admin-enable.

Downlink delay is disabled by default and must be explicitly enabled. When enabled, if the delay timer is not specified, it is set for a default delay of 20 seconds.

Default Physical Ethernet Settings

The following table lists the default settings for all physical Ethernet interfaces:

Parameter	Default Setting
Duplex	Auto (full-duplex)
Encapsulation	ARPA

Parameter	Default Setting
MTU ¹	1500 bytes
Port Mode	Access
Speed	Auto (10000)

¹ MTU cannot be changed per-physical Ethernet interface. You modify MTU by selecting maps of QoS classes.

Displaying Interface Information

To view configuration information about the defined interfaces, perform one of these tasks:

Command	Purpose
switch# show interface <i>type slot/port</i>	Displays the detailed configuration of the specified interface.
switch# show interface <i>type slot/port capabilities</i>	Displays detailed information about the capabilities of the specified interface. This option is available only for physical interfaces.
switch# show interface <i>type slot/port transceiver</i>	Displays detailed information about the transceiver connected to the specified interface. This option is available only for physical interfaces.
switch# show interface brief	Displays the status of all interfaces.
switch# show interface flowcontrol	Displays the detailed listing of the flow control settings on all interfaces.

The **show interface** command is invoked from EXEC mode and displays the interface configurations. Without any arguments, this command displays the information for all the configured interfaces in the switch.

This example shows how to display the physical Ethernet interface:

```
switch# show interface ethernet 1/1
Ethernet1/1 is up
Hardware is 1000/10000 Ethernet, address is 000d.eca3.5f08 (bia 000d.eca3.5f08)
MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,
    reliability 255/255, txload 190/255, rxload 192/255
Encapsulation ARPA
Port mode is trunk
full-duplex, 10 Gb/s, media type is 1/10g
Input flow-control is off, output flow-control is off
Auto-mdix is turned on
Rate mode is dedicated
Switchport monitor is off
Last clearing of "show interface" counters never
5 minute input rate 942201806 bytes/sec, 14721892 packets/sec
5 minute output rate 935840313 bytes/sec, 14622492 packets/sec
Rx
  129141483840 input packets 0 unicast packets 129141483847 multicast packets
  0 broadcast packets 0 jumbo packets 0 storm suppression packets
```

```

8265054965824 bytes
0 No buffer 0 runt 0 Overrun
0 crc 0 Ignored 0 Bad etype drop
0 Bad proto drop
Tx
119038487241 output packets 119038487245 multicast packets
0 broadcast packets 0 jumbo packets
7618463256471 bytes
0 output CRC 0 ecc
0 underrun 0 if down drop      0 output error 0 collision 0 deferred
0 late collision 0 lost carrier 0 no carrier
0 babble
0 Rx pause 8031547972 Tx pause 0 reset

```

This example shows how to display the physical Ethernet capabilities:

```

switch# show interface ethernet 1/1 capabilities
Ethernet1/1
Model:                734510033
Type:                 10Gbase- (unknown)
Speed:                1000,10000
Duplex:               full
Trunk encap. type:    802.1Q
Channel:              yes
Broadcast suppression: percentage(0-100)
Flowcontrol:          rx- (off/on), tx- (off/on)
Rate mode:            none
QOS scheduling:        rx- (6qlt), tx- (lp6q0t)
CoS rewrite:          no
ToS rewrite:          no
SPAN:                 yes
UDLD:                 yes
MDIX:                 no
FEX Fabric:           yes

```

This example shows how to display the physical Ethernet transceiver:

```

switch# show interface ethernet 1/1 transceiver
Ethernet1/1
sfp is present
name is CISCO-EXCELIGHT
part number is SPP5101SR-C1
revision is A
serial number is ECL120901AV
nominal bitrate is 10300 Mbits/sec
Link length supported for 50/125mm fiber is 82 m(s)
Link length supported for 62.5/125mm fiber is 26 m(s)
cisco id is --
cisco extended id number is 4

```

This example shows how to display a brief interface status (some of the output has been removed for brevity):

```
switch# show interface brief
```

Ethernet Interface	VLAN	Type	Mode	Status	Reason	Speed	Port Ch #
Eth1/1	200	eth	trunk	up	none	10G(D)	--
Eth1/2	1	eth	trunk	up	none	10G(D)	--
Eth1/3	300	eth	access	down	SFP not inserted	10G(D)	--
Eth1/4	300	eth	access	down	SFP not inserted	10G(D)	--
Eth1/5	300	eth	access	down	Link not connected	1000(D)	--

```

Eth1/6      20      eth  access down   Link not connected   10G(D)  --
Eth1/7      300     eth  access down   SFP not inserted    10G(D)  --
...

```

This example shows how to display the CDP neighbors:

```

switch# show cdp neighbors
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute

Device ID         Local Intrfce  Hldtme  Capability  Platform  Port ID
dl13-dist-1       mgmt0         148     S I         WS-C2960-24TC  Fas0/9
n5k (FLC12080012) Eth1/5        8       S I s       N5K-C5020P-BA  Eth1/5

```

