



Checking Port Status and Connectivity

- [Check Connected Modules, on page 1](#)
- [Check Interface Status, on page 2](#)
- [Displaying PORT SET ENABLED LED Status, on page 3](#)
- [Display MAC Addresses, on page 5](#)
- [Using Telnet, on page 6](#)
- [Check Cable Status Using Time Domain Reflectometer, on page 7](#)
- [Change the Logout Timer, on page 8](#)
- [Monitor User Sessions, on page 9](#)
- [Using Ping, on page 9](#)
- [Using IP Traceroute, on page 11](#)
- [Layer 2 Traceroute, on page 12](#)
- [Configure ICMP, on page 13](#)
- [Feature History for Checking Port Status and Connectivity, on page 14](#)

Check Connected Modules

The Catalyst 9400 series switch is a modular system. You can see which modules are installed, and the MAC address ranges and version numbers for each module, by entering the show module command. Use the *mod_num* argument to specify a particular module number and display detailed information on that module.

This example shows how to check the status for all modules on your switch:

```
Device# show module
```

```
Chassis Type: C9410R
```

Mod	Ports	Card Type	Model	Serial No.
1	48	48-Port UPOE w/ 24p mGig 24p RJ-45	C9400-LC-48UX	JAE2229053D
2	48	48-Port 5Gig/mGig 90W BT (RJ-45)	C9400-LC-48HN	JAE24530BF3
3	48	48-Port UPOE w/ 24p mGig 24p RJ-45	C9400-LC-48UX	JAE2128068Z
4	48	48-Port 5Gig/mGig 90W BT (RJ-45)	C9400-LC-48HN	JAE24241WAY
5	11	Supervisor 1 Module	C9400-SUP-1	JAE22280PL8
6	11	Supervisor 1 Module	C9400-SUP-1	JAE22280PHT
7	48	48-Port UPOE w/ 24p mGig 24p RJ-45	C9400-LC-48UX	JAE2229055N
8	48	48-Port UPOE w/ 24p mGig 24p RJ-45	C9400-LC-48UX	JAE22280DBU
9	48	48-Port UPOE w/ 24p mGig 24p RJ-45	C9400-LC-48UX	JAE22080BWS
10	48	48-Port UPOE w/ 24p mGig 24p RJ-45	C9400-LC-48UX	JAE230707YP

Mod	MAC addresses	Hw	Fw	Sw	Status
1	BC26.C7A4.E738 to BC26.C7A4.E767	1.0	17.5.1r	17.05.01	ok
2	ECCE.13E2.B670 to ECCE.13E2.B69F	1.0	17.5.1r	17.05.01	ok
3	E4AA.5D59.A868 to E4AA.5D59.A897	1.0	17.5.1r	17.05.01	ok
4	A0B4.3982.43C0 to A0B4.3982.43EF	1.0	17.5.1r	17.05.01	ok
5	2C5A.0F1C.1EEC to 2C5A.0F1C.1EF6	2.0	17.5.1r	17.05.01	ok
6	2C5A.0F1C.1EF6 to 2C5A.0F1C.1F00	2.0	17.5.1r	17.05.01	ok
7	BC26.C7A4.D820 to BC26.C7A4.D84F	1.0	17.5.1r	17.05.01	ok
8	BC26.C772.E91C to BC26.C772.E94B	1.0	17.5.1r	17.05.01	ok
9	707D.B9C8.B5F8 to 707D.B9C8.B627	2.1	17.5.1r	17.05.01	ok
10	70EA.1ADB.7E74 to 70EA.1ADB.7EA3	3.0	17.5.1r	17.05.01	ok

Mod	Redundancy Role	Operating Mode	Configured Mode	Redundancy Status
5	Active	sso	sso	Active
6	Standby	sso	sso	Standby Hot

Chassis MAC address range: 44 addresses from 2c5a.0f1c.1ec0 to 2c5a.0f1c.1eeb

Check Interface Status

You can view the summary or detailed information on the switch ports using the **show interface status** command. To see the summary information on all ports on the switch, enter the **show interface status** command with no arguments. Specify a particular module number to see information on the ports on that module only. Enter both the module number and the port number to see detailed information about the specified port.

To apply configuration commands to a particular port, you must specify the appropriate logical module.

This example shows how to display the status of all interfaces on a Catalyst 9400 series switch, including transceivers:

```
Switch# show interface status
```

Port	Name	Status	Vlan	Duplex	Speed	Type
Gi1/0/1		connected	1	a-full	a-1000	10/100/1000BaseTX
Gi1/0/2		connected	1	a-full	a-1000	10/100/1000BaseTX
Gi1/0/3		connected	1	a-full	a-1000	10/100/1000BaseTX
Gi1/0/4		notconnect	1	auto	auto	10/100/1000BaseTX
Gi1/0/5		notconnect	1	auto	auto	10/100/1000BaseTX
Gi1/0/6		notconnect	1	auto	auto	10/100/1000BaseTX
Gi1/0/7		notconnect	1	auto	auto	10/100/1000BaseTX
Gi1/0/8		notconnect	1	auto	auto	10/100/1000BaseTX
Gi1/0/9		notconnect	1	auto	auto	10/100/1000BaseTX
Gi1/0/10		notconnect	1	auto	auto	10/100/1000BaseTX
Gi1/0/11		notconnect	1	auto	auto	10/100/1000BaseTX
Gi1/0/12		notconnect	1	auto	auto	10/100/1000BaseTX
Gi1/0/13		notconnect	1	auto	auto	10/100/1000BaseTX
Gi1/0/14		notconnect	1	auto	auto	10/100/1000BaseTX
Gi1/0/15		notconnect	1	auto	auto	10/100/1000BaseTX
Gi1/0/16		notconnect	1	auto	auto	10/100/1000BaseTX
Gi1/0/17		notconnect	1	auto	auto	10/100/1000BaseTX
Gi1/0/18		notconnect	1	auto	auto	10/100/1000BaseTX
Gi1/0/19		notconnect	1	auto	auto	10/100/1000BaseTX
Gi1/0/20		notconnect	1	auto	auto	10/100/1000BaseTX
Gi1/0/21		notconnect	1	auto	auto	10/100/1000BaseTX
Gi1/0/22		notconnect	1	auto	auto	10/100/1000BaseTX
Gi1/0/23		notconnect	1	auto	auto	10/100/1000BaseTX
Gi1/0/24		notconnect	1	auto	auto	10/100/1000BaseTX

This example shows how to display the status of interfaces in error-disabled state:

```

Device# show interfaces status err-disabled
Port Name Status Reason
Fa9/4 err-disabled link-flap
informational error message when the timer expires on a cause
-----
5d04h:%PM-SP-4-ERR_RECOVER:Attempting to recover from link-flap err-disable state on Fa9/4
Switch#

```

Displaying PORT SET ENABLED LED Status

PORT SET ENABLED LED Status on Supervisor 1

There are four PORT SET ENABLED LEDs on the Supervisor 1 faceplate:

- One for port numbers 1 to 4, termed G1.
- One for port numbers 5 to 8, termed G2
- One for port number 9, termed G3
- One for port number 10, termed G4

Ports 1 to 8 are tengigabit ports and ports 9 and 10 are fortygigabit ports.

Standalone Supervisor 1

With a Standalone Supervisor, a single Supervisor is active and has ten ports. Group G1 and group G3 are mutually exclusive which means that either ports 1 to 4 are active or port 9 is active. Similarly, group G2 and group G4 are mutually exclusive; either ports 5 to 8 are active or port 10 is active. The status of the groups is decided by the configuration of the fortygigabit interfaces.

Displaying PORT SET ENABLED LED in a Standalone Supervisor 1 Mode

The following sample configuration enables the fortygigabit port number 10:

```

interface FortyGigabitEthernet4/0/9
end

interface FortyGigabitEthernet4/0/10
  enable
end

```

Following is a sample output of the **show hardware led** command:

```

SUPERVISOR: ACTIVE
PORT STATUS: (10) Te4/0/1:BLACK Te4/0/2:BLACK Te4/0/3:BLACK Te4/0/4:BLACK Te4/0/5:BLACK
Te4/0/6:BLACK Te4/0/7:BLACK Te4/0/8:BLACK Fo4/0/9:BLACK Fo4/0/10:BLACK

BEACON: BLACK

GROUP LED: UPLINK-G1:GREEN UPLINK-G2:BLACK UPLINK-G3:BLACK UPLINK-G4:GREEN

```

In this sample, you can see that group 4 is active (GREEN) and hence group 2 is inactive (BLACK). Since group 3 is not enabled and is inactive (BLACK), group 1 is active (GREEN)

High Availability or Dual Supervisor 1 Mode

In a dual supervisor mode, the Ten-gigabit ports numbered 1 to 4 (G1) and the Forty-gigabit port numbered 9 (G3) are operational on both the supervisors. The other Ten-gigabit ports numbered 5 to 8 (G2) and the Forty-gigabit port numbered 10 (G4) are disabled by default. Of the groups G1 and G3 which are mutually exclusive, either of the groups are active based on the configuration of the Forty-gigabit port number 9.

Displaying PORT SET ENABLED LED in a Dual Supervisor 1 Mode

```
Switch#show run int fo4/0/9
Building configuration...
```

```
Current configuration : 52 bytes
!
interface FortyGigabitEthernet4/0/9
  enable
end
```

```
Switch#
```

```
SUPERVISOR: STANDBY
PORT STATUS: (10) Te3/0/1:BLACK Te3/0/2:BLACK Te3/0/3:BLACK Te3/0/4:BLACK Te3/0/5:BLACK
Te3/0/6:BLACK Te3/0/7:BLACK Te3/0/8:BLACK Fo3/0/9:BLACK Fo3/0/10:BLACK
```

```
BEACON: BLACK
```

```
GROUP LED: UPLINK-G1:GREEN UPLINK-G2:BLACK UPLINK-G3:BLACK UPLINK-G4:BLACK
```

```
SUPERVISOR: ACTIVE
PORT STATUS: (10) Te4/0/1:BLACK Te4/0/2:BLACK Te4/0/3:BLACK Te4/0/4:BLACK Te4/0/5:BLACK
Te4/0/6:BLACK Te4/0/7:BLACK Te4/0/8:BLACK Fo4/0/9:BLACK Fo4/0/10:BLACK
```

```
BEACON: BLACK
```

```
GROUP LED: UPLINK-G1:BLACK UPLINK-G2:BLACK UPLINK-G3:GREEN UPLINK-G4:BLACK
```

PORT SET ENABLED LED Status on Supervisor 2

There are five PORT SET ENABLED LEDs on the Supervisor 2 faceplate:

- One for port numbers 1 to 4, termed G1.
- One for port number 5, termed G2
- One for port number 6, termed G3
- One for port number 7, termed G4
- One for port number 8, termed G5

Standalone Supervisor 2

With a Standalone Supervisor, a single Supervisor is active and has eight ports. Group G1 and group G2 are mutually exclusive which means that either ports 1 to 4 are active or port 5 is active. In a redundant setup, groups G1, G2, and G3 are used, and groups G4 and G5 are inactivate.

Displaying PORT SET ENABLED LED in a Standalone Supervisor 2 Mode

The following sample configuration enables the hundred-gigabit port number 5:

```
interface HundredGigE3/0/5
end

interface HundredGigE3/0/5
  enable
end
```

Following is a sample output of the **show hardware led** command:

```
SUPERVISOR: ACTIVE
PORT STATUS: (10) Twe3/0/1:BLACK Twe3/0/2:ACT_GREEN Twe3/0/3:BLACK Twe3/0/4:BLACK
Hu3/0/5:BLACK Hu3/0/6:ACT_GREEN Hu3/0/7:BLACK Hu3/0/8:BLACK
BEACON: BLACK
STATUS: GREEN
GROUP LED: UPLINK-G1:GREEN UPLINK-G2:BLACK UPLINK-G3:GREEN UPLINK-G4:GREEN UPLINK-G5:GREEN
```

In this sample, you can see that groups 1, 3, 4, and 5 are active (GREEN) and group 2 is inactive (BLACK).

High Availability or Dual Supervisor 2 Mode

In a dual supervisor mode, the TwentyFiveGigabit ports numbered 1 to 4 (G1) and the HundredGigabit port numbered 6 (G3) are operational on both the supervisors. The other HundredGigabit ports numbered 7 and 8 (G4 and G5) are disabled by default.

Displaying PORT SET ENABLED LED in a Dual Supervisor 2 Mode

```
Switch#show run interface HundredGigE3/0/5
Building configuration...

Current configuration : 43 bytes
!
interface HundredGigE3/0/5
  enable
end

Switch#

SUPERVISOR: ACTIVE
PORT STATUS: (10) Twe3/0/1:BLACK Twe3/0/2:ACT_GREEN Twe3/0/3:BLACK Twe3/0/4:BLACK
Hu3/0/5:BLACK Hu3/0/6:ACT_GREEN Hu3/0/7:BLACK Hu3/0/8:BLACK
BEACON: BLACK
STATUS: GREEN
GROUP LED: UPLINK-G1:GREEN UPLINK-G2:BLACK UPLINK-G3:GREEN UPLINK-G4:BLACK UPLINK-G5:BLACK

SUPERVISOR: STANDBY
PORT STATUS: (10) Twe4/0/1:ACT_GREEN Twe4/0/2:ACT_GREEN Twe4/0/3:BLACK Twe4/0/4:BLACK
Hu4/0/5:BLACK Hu4/0/6:ACT_GREEN Hu4/0/7:BLACK Hu4/0/8:BLACK
BEACON: BLACK
STATUS: GREEN
GROUP LED: UPLINK-G1:GREEN UPLINK-G2:BLACK UPLINK-G3:GREEN UPLINK-G4:BLACK UPLINK-G5:BLACK
```

Display MAC Addresses

In addition to displaying the MAC address range for a module using the **show module** command, you can display the MAC address table information of a specific MAC address or a specific interface in the switch using the **show mac address-table address** and **show mac address-table interface** commands.

This example shows how to display MAC address table information for all MAC addresses:

```
Switch# show mac address-table
      Mac Address Table
-----
Vlan    Mac Address      Type      Ports
----    -
All     0100.0ccc.cccc   STATIC    CPU
All     0100.0ccc.cccd   STATIC    CPU
All     0180.c200.0000   STATIC    CPU
All     0180.c200.0001   STATIC    CPU
All     0180.c200.0002   STATIC    CPU
All     0180.c200.0003   STATIC    CPU
All     0180.c200.0004   STATIC    CPU
All     0180.c200.0005   STATIC    CPU
All     0180.c200.0006   STATIC    CPU
All     0180.c200.0007   STATIC    CPU
All     0180.c200.0008   STATIC    CPU
All     0180.c200.0009   STATIC    CPU
All     0180.c200.000a   STATIC    CPU
All     0180.c200.000b   STATIC    CPU
All     0180.c200.000c   STATIC    CPU
All     0180.c200.000d   STATIC    CPU
All     0180.c200.000e   STATIC    CPU
All     0180.c200.000f   STATIC    CPU
All     0180.c200.0010   STATIC    CPU
All     0180.c200.0021   STATIC    CPU
All     ffff.ffff.ffff   STATIC    CPU
      1     188b.45eb.cc01   DYNAMIC   Gi1/0/1
Total Mac Addresses for this criterion: 22
Switch#
```

This example shows how to display MAC address table information for a specific interface:

```
Switch# show mac address-table interface Gi1/0/1
      Mac Address Table
-----
Vlan    Mac Address      Type      Ports
----    -
      1     188b.45eb.cc01   DYNAMIC   Gi1/0/1
Total Mac Addresses for this criterion: 1
Switch#
```

Using Telnet

You can access the switch command-line interface (CLI) using Telnet. In addition, Telnet allows you to access other devices in the network. You can have up to eight simultaneous Telnet sessions.

Before you can open a Telnet session to the switch, you must first set the IP address (and in some cases the default gateway) for the switch. For information about setting the IP address and default gateway, see the section on *Configuring the Switch for the First Time*.



Note To establish a Telnet connection to a host by using the hostname, configure and enable DNS.

To establish a Telnet connection to another device on the network from the switch, enter this command:

```
Switch# telnet host [port]
```

This example shows how to establish a Telnet connection from the switch to the remote host named labsparc:

```
Switch# telnet labsparc
Trying 172.16.10.3...
Connected to labsparc.
Escape character is '^]'.
UNIX(r) System V Release 4.0 (labsparc)
login:
```

Check Cable Status Using Time Domain Reflectometer

The Time Domain Reflectometer (TDR) feature allows you to determine if a cable is OPEN or SHORT when it is at fault.

With TDR, you can check the status of copper cables on the 48-port 10/100/1000 BASE-T modules for the Catalyst 9400 Series Switch. TDR detects a cable fault by sending a signal through the cable and reading the signal that is reflected back. All or part of the signal can be reflected back due to defects in the cable.



Note Category 5 cable has four pairs. Each pair can assume one of the following states: open (not connected), broken, shorted, or terminated. The TDR test detects all four states and displays the first three as “Fault” conditions, and displays the fourth as “Terminated”. Although the CLI output is shown, the cable length is displayed only if the state is “Faulty.”

TDR feature is supported on the following modules:

- C9400-LC-48U
- C9400-LC-48T
- C9400-LC-48P

TDR detects a cable fault by sending a signal along its wires. Depending on the reflected signal, it can determine roughly where the cable fault has occurred. The variations on how TDR signal is reflected back determine the results on TDR. On Catalyst 9400 Series Switch, only two types of cable fault types are detected - OPEN or SHORT. We do display Terminated status in case cable is properly terminated and this is done for illustrative purpose.

Running the TDR Test

To start the TDR test, perform this task:

Procedure

	Command or Action	Purpose
Step 1	<code>test cable-diagnostics tdr {interface { interface-number}}</code>	Starts the TDR test.
Step 2	<code>show cable-diagnostics tdr {interface interface-number}</code>	Displays the TDR test counter information.

TDR Guidelines

The following guidelines apply to the use of TDR:

- Do not change the port configuration while the TDR test is running.
- If you connect a port undergoing a TDR test to an Auto-MDIX enabled port, the TDR result might be invalid. In those instances, the port on the device should be administratively down before the start of the TDR test.
- If you connect a port undergoing a TDR test to a 100BASE-T port such as that on the device, the unused pairs (4-5 and 7-8) are reported as faulty because the remote end does not terminate these pairs.
- Due to cable characteristics, you should run the TDR test multiple times to get accurate results.
- Do not change port status (for example, remove the cable at the near or far end) because the results might be inaccurate.
- TDR works best if the test cable is disconnected from the remote port. Otherwise, it might be difficult for you to interpret results correctly.
- TDR operates across four wires. Depending on the cable conditions, the status might show that one pair is OPEN or SHORT while all other wire pairs display as faulty. This operation is acceptable because you should declare a cable faulty provided one pair of wires is either OPEN or SHORT.
- TDR intent is to determine how poorly a cable is functioning rather than to locate a faulty cable.
- When TDR locates a faulty cable, you should still use an offline cable diagnosis tool to better diagnose the problem.
- TDR results might differ between runs on different Catalyst 9400 modules because of the resolution difference of TDR implementations. When this occurs, you should refer to an offline cable diagnosis tool.

Change the Logout Timer

The logout timer automatically disconnects a user from the switch when the user is idle for longer than the specified time. To set the logout timer, enter this command:

```
Switch(config-line)# exec-timeout minutes seconds
```

This command changes the logout timer value (a timeout value of 0 prevents idle sessions from being disconnected automatically).

Use the **no** keyword to return to the default value.

To set the logout for 10 minutes and 10 seconds, enter the following command:

```
Switch(config)# line console 0  
Switch(config-line)# exec-timeout 10 10
```

To set no logout timer for console session:

```
Switch(config)# line console 0  
Switch(config-line)# exec-timeout 0 0
```


Monitor User Sessions

You can display the currently active user sessions on the switch using the **show users** command. The command output lists all active console port and Telnet sessions on the switch.

To display the active user sessions on the switch, enter this command:

```
Switch# show users [all]
```

To disconnect an active user session on the switch, enter the following command:

```
Switch# disconnect { console | ip_address }
```

Example

This example shows the output of the show users command when local authentication is enabled for console and Telnet sessions (the asterisk [*] indicates the current session)

```
Switch# show users
Line User Host(s) Idle Location
* 0 con 0 idle 00:00:00
Interface User Mode Idle Peer Address

Switch# show users all
Line User Host(s) Idle Location
* 0 con 0 idle 00:00:00
1 vty 0 00:00:00
2 vty 1 00:00:00
3 vty 2 00:00:00
4 vty 3 00:00:00
5 vty 4 00:00:00
Interface User Mode Idle Peer Address
Switch#
```

This example shows how to disconnect an active console port session and an active Telnet session:

```
Switch> disconnect console
Console session disconnected.
Console> (enable) disconnect tim-nt.bigcorp.com
Telnet session from tim-nt.bigcorp.com disconnected. (1)
Switch# show users
Session User Location
-----
telnet jake jake-mac.bigcorp.com
* telnet suzy suzy-pc.bigcorp.com
Switch#
```

Using Ping

These sections describe how to use IP ping:

How Ping Works

The ping command allows you to verify connectivity to remote hosts. If you attempt to ping a host in a different IP subnet, you must define a static route to the network or configure a router to route between those subnets.

The ping command is configurable from normal executive and privileged EXEC mode. A ping returns one of the following responses:

- Normal response—The normal response (hostname is alive) occurs in 1 to 10 seconds, depending on the network traffic.
- Destination does not respond—If the host does not respond, a No Answer message is returned.
- Unknown host—If the host does not exist, an Unknown Host message is returned.
- Destination unreachable—If the default gateway cannot reach the specified network, a Destination Unreachable message is returned.
- Network or host unreachable—If there is no entry in the route table for the host or network, a Network or Host Unreachable message is returned.

To stop a ping in progress, press Ctrl-C.

Run Ping Command

To ping another device on the network from the switch, enter this command in normal executive and privileged EXEC mode:

```
Switch# ping host
```

Checks connectivity to a remote host.

This example shows how to ping a remote host from normal executive mode:

```
Switch# ping labsparc
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Switch#
```

```
Switch# ping 72.16.10.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 72.16.10.3, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Switch#
```

This example shows how to use a ping command in privileged EXEC mode to specify the number of packets, the packet size, and the timeout period:

```
Switch# ping
Protocol [ip]: ip
Target IP address: 1.1.1.1
Repeat count [5]: 10
Datagram size [100]: 100
Timeout in seconds [2]: 10
Extended commands [n]: n
Sweep range of sizes [n]: n
Type escape sequence to abort.
Sending 10, 100-byte ICMP Echos to 1.1.1.1, timeout is 10 seconds:
!!!!!!!!!!!!
Success rate is 100 percent (10/10), round-trip min/avg/max = 1/1/1 ms
Switch#
```

Using IP Traceroute

How IP Traceroute Works

IP traceroute allows you to identify the path that packets take through the network on a hop-by-hop basis. The command output displays all network layer (Layer 3) devices, such as routers, that the traffic passes through on the way to the destination.

Layer 2 switches can participate as the source or destination of the trace command but does not appear as a hop in the trace command output.

The trace command uses the time to live (TTL) field in the IP header to cause routers and servers to generate specific return messages. Traceroute starts by sending a User Datagram Protocol (UDP) datagram to the destination host with the TTL field set to 1. If a router finds a TTL value of 1 or 0, it drops the datagram and sends back an Internet Control Message Protocol (ICMP) Time-Exceeded message to the sender. Traceroute determines the address of the first hop by examining the source address field of the ICMP Time-Exceeded message.

To identify the next hop, traceroute sends a UDP packet with a TTL value of 2. The first router decrements the TTL field by 1 and sends the datagram to the next router. The second router sees a TTL value of 1, discards the datagram, and returns the Time-Exceeded message to the source. This process continues until the TTL is incremented to a value large enough for the datagram to reach the destination host or until the maximum TTL is reached.

To determine when a datagram reaches its destination, traceroute sets the UDP destination port in the datagram to a large value that the destination host is unlikely to be using. When a host receives a datagram with an unrecognized port number, it sends an ICMP Port Unreachable error message to the source. The Port Unreachable error message indicates to traceroute that the destination has been reached.

Perform IP Traceroute

To trace the path that packets take through the network, enter this command in EXEC or privileged EXEC mode:

Procedure

	Command or Action	Purpose
Step 1	<code>traceroute [protocol] [destination]</code>	Runs IP traceroute to trace the path that packets take through the network.

Example

This example shows how to use the traceroute command to display the route that a packet takes through the network to reach its destination:

```
Switch# traceroute ip ABA.NYC.mil
Type escape sequence to abort.
Tracing the route to ABA.NYC.mil (26.0.0.73)
 1 DEBRIS.CISCO.COM (192.180.1.6) 1000 msec 8 msec 4 msec
 2 BARRNET-GW.CISCO.COM (192.180.16.2) 8 msec 8 msec 8 msec
```

```

3 EXTERNAL-A-GATEWAY.STANFORD.EDU (192.42.110.225) 8 msec 4 msec 4 msec
4 BB2.SU.BARRNET.NET (192.200.254.6) 8 msec 8 msec 8 msec
5 SU.ARC.BARRNET.NET (192.200.3.8) 12 msec 12 msec 8 msec
6 MOFFETT-FLD-MB.in.MIL (192.52.195.1) 216 msec 120 msec 132 msec
7 ABA.NYC.mil (26.0.0.73) 412 msec 628 msec 664 msec
Switch#

```

Layer 2 Traceroute

The Layer 2 traceroute feature allows the switch to identify the physical path that a packet takes from a source device to a destination device. Layer 2 traceroute supports only unicast source and destination MAC addresses. It determines the path by using the MAC address tables of the switches in the path. When the switch detects a device in the path that does not support Layer 2 traceroute, the switch continues to send Layer 2 trace queries and lets them time out.

If you want the switch to trace the path from a host on a source device to a host on a destination device, the switch can identify only the path from the source device to the destination device. It cannot identify the path that a packet takes from the source host to the source device or from the destination device to the destination host.

Layer 2 Traceroute Usage Guidelines

These are the Layer 2 traceroute usage guidelines:

- CDP must be enabled on all the devices in the network. For Layer 2 traceroute to function properly, do not disable CDP.
 - If any devices in the physical path are transparent to CDP, the switch cannot identify the path through these devices.
- All switches in the physical path must have IP connectivity. When a switch is reachable from another switch, you can test connectivity by using the ping command in privileged EXEC mode.
- The maximum number of hops identified in the path is ten.
- You can enter the **traceroute mac** or the **traceroute mac ip command** in privileged EXEC mode on a switch that is not in the physical path from the source device to the destination device. All switches in the path must be reachable from this switch.
- The **traceroute mac** command output shows the Layer 2 path only when the specified source and destination MAC addresses belong to the same VLAN. If you specify source and destination MAC addresses that belong to different VLANs, the Layer 2 path is not identified, and an error message appears.
- If you specify a multicast source or destination MAC address, the path is not identified, and an error message appears.
- If the source or destination MAC address belongs to multiple VLANs, you must specify the VLAN to which both the source and destination MAC addresses belong. If the VLAN is not specified, the path is not identified, and an error message appears.
- The **traceroute mac ip** command output shows the Layer 2 path when the specified source and destination IP addresses belong to the same subnet. When you specify the IP addresses, the switch uses Address Resolution Protocol (ARP) to associate the IP address with the corresponding MAC address and the VLAN ID.

- If an ARP entry exists for the specified IP address, the switch uses the associated MAC address and identifies the physical path.
- If an ARP entry does not exist, the switch sends an ARP query and tries to resolve the IP address. If the IP address is not resolved, the path is not identified, and an error message appears.
- When multiple devices are attached to one port through hubs (for example, multiple CDP neighbors are detected on a port), the Layer 2 traceroute feature is not supported. When more than one CDP neighbor is detected on a port, the Layer 2 path is not identified, and an error message appears.
- This feature is not supported in Token Ring VLANs.

Perform Layer 2 Traceroute

To display the physical path that a packet takes from a source device to a destination device, enter either one of these commands:

```
Switch# traceroute mac source-mac-address destination-mac-address
```

OR

```
Switch# traceroute mac ip source-ip destination-ip
```

The following examples show how to use the **traceroute mac** and **traceroute mac ip** commands to display the physical path that a packet takes through the network to reach its destination:

```
Switch# traceroute mac cc16.7eaa.7203 188b.45eb.cc64
Source cc16.7eaa.7203 found on Switch
1 Switch (1.1.1.1) : V11 => Gi1/0/1
Destination 188b.45eb.cc64 found on Switch
Layer 2 trace completed.
Switch#
```

```
Switch# traceroute mac ip 1.1.1.1 1.1.1.2 detail
Translating IP to mac .....
1.1.1.1 => cc16.7eaa.7203
1.1.1.2 => 188b.45eb.cc64
```

```
Source cc16.7eaa.7203 found on Switch[C9410R] (1.1.1.1)
1 Switch / C9410R / 1.1.1.1 :Gi1/0/1 [auto, auto]
Destination 188b.45eb.cc64 found on Switch[C9410R] (1.1.1.1)
Layer 2 trace completed.
Switch#
```

Configure ICMP

Internet Control Message Protocol (ICMP) provides many services that control and manage IP connections. ICMP messages are sent by routers or access servers to hosts or other routers when a problem is discovered with the Internet header. For detailed information on ICMP, refer RFC 792.

Enable ICMP Protocol Unreachable Messages

If the Cisco IOS software receives a nonbroadcast packet that uses an unknown protocol, it sends an ICMP Protocol Unreachable message back to the source.

Similarly, if the software receives a packet that it is unable to deliver to the ultimate destination because it knows of no route to the destination address, it sends an ICMP Host Unreachable message to the source. This feature is enabled by default.

To enable the generation of ICMP Protocol Unreachable and Host Unreachable messages, enter the following command in interface configuration mode:

```
Switch (config-if)# [no] ip unreachable
```

Use the **no** keyword to disable the ICMP destination unreachable messages.



Note If you enter the **no ip unreachable** command, you will break the path MTU discovery functionality. Routers in the middle of the network might be forced to fragment packets.

To limit the rate that Internet Control Message Protocol (ICMP) destination unreachable messages are generated, enter the following command:

```
Switch (config)# [no] ip icmp rate-limit unreachable [df] milliseconds
```

Use the **no** keyword to remove the rate limit and reduce the CPU usage.

Enable ICMP Mask Reply Messages

Occasionally, network devices must know the subnet mask for a particular subnetwork in the internetwork. To obtain this information, devices can send ICMP Mask Request messages. These messages are responded to by ICMP Mask Reply messages from devices that have the requested information. The Cisco IOS software can respond to ICMP Mask Request messages if the ICMP Mask. Reply function is enabled.

To have the Cisco IOS software respond to ICMP mask requests by sending ICMP Mask Reply messages, enter the following command:

```
Switch (config-if)# [no] ip mask-reply
```

Use the **no** keyword to disable this functionality.

Feature History for Checking Port Status and Connectivity

This table provides release and related information for features explained in this module.

These features are available on all releases subsequent to the one they were introduced in, unless noted otherwise.

Release	Feature	Feature Information
Cisco IOS XE Everest 16.6.1	Port Status and Connectivity Check	This feature includes the steps to check the status of modules, and interfaces; and also how to verify connectivity between devices within the network.

Release	Feature	Feature Information
Cisco IOS XE Fuji 16.8.1a	Command to display LED status	The show hardware led command was introduced to display the LED status.
Cisco IOS XE Cupertino 17.7.1	Port Status and Connectivity Check	This feature was implemented on supervisor modules C9400X-SUP-2 and C9400X-SUP-2XL, which were introduced in this release.

Use Cisco Feature Navigator to find information about platform and software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>.

