Troubleshoot Switch Port and Interface Problems

Contents

Introduction

This document describes how to determine why a port or interface experiences problems.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document applies to Catalyst switches that run on Cisco IOS® System Software.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Conventions

Refer to Cisco Technical Tips Conventions for more information on document conventions.

✎ **Note**: To access tools and websites, you must be a registered Cisco client.

Troubleshoot the Physical Layer

Use the LEDs to Troubleshoot

If you have physical access to the switch, it can save time to look at the port LEDs which give you the link status or can indicate an error condition (if red or orange). The table describes the LED status indicators for Ethernet modules or fixed-configuration switches:

Platform	URL
Catalyst 6000 Series Switches	Ethernet Module LEDs

Catalyst 4000 Series Switches	Ethernet Module LEDs
Catalyst 3750 Series Switches	Front Panel LEDs
Catalyst 3550 Series Switches	Front Panel LEDs
Catalyst 2950/2955 Series Switches	Front Panel LEDs
Catalyst 2900/3500XL Series Switches	Front Panel LEDs
Catalyst 1900 and 2820 Series Switches	Front Panel LEDs

Ensure that both sides have a link. A single broken wire or one shutdown port can cause the problem where one side has a link light, but the other side does not.

A link light does not guarantee that the cable is fully functional. The cable can have encountered physical stress that causes it to be functional at a marginal level. Normally you can identify this situation if the port has many packet errors, or the port constantly flaps (loses and regains link).

Check the Cable and Both Sides of the Connection

If the link light for the port does not come on, you can consider these possibilities:

Possible Cause	**Corrective Action**
No cable connected	Connect cable from switch to a known good device.
Wrong Port	Make sure that both ends of the cable are plugged into the correct ports.
Device has no power	Ensure that both devices have power.
Wrong cable type	Verify the cable selection. Refer to the Catalyst Switch Cable Guide.
Bad cable	Swap suspect cable with known good cable. Look for broken or lost pins on connectors.
Loose connections	Check for loose connections. Sometimes a cable appears to be seated in the jack but is not. Unplug the cable and reinsert it.
Patch Panels	Eliminate faulty patch panel connections. Bypass

	the patch panel if possible to rule it out.
Media Convertors	Eliminate faulty media convertors: fiber-to-copper, and so on. Bypass the media convertor if possible to rule it out.
Bad or wrong Gigabit Interface Convertor (GBIC)	Swap suspect GBIC with known good GBIC. Verify Hw and Sw support for this type of GBIC.
Bad Port or Module Port or Interface or Module not enabled	Move the cable to a known good port to troubleshoot a suspect port or module. Use th **show interface** command for Cisco IOS to look for errdisable, disable or shutdown status. The **show module** command can indicate faulty, which can indicate a hardware problem. See the Common Port and Interface Problems section of this document for more information.

Ethernet Copper and Fiber Cables

Ensure that you have the correct cable for the type of connection you want to make. Category 3 copper cable can be used for 10 Mbps unshielded twisted pair (UTP) connections but must never be used for 10/100 or 10/100/1000Mbps UTP connections. Always use either Category 5, Category 5e, or Category 6 UTP for 10/100 or 10/100/1000Mbps connections.

Warning: Category 5e and Category 6 cables can store high levels of static electricity because of the dielectric properties of the materials used in their construction. Always ground the cables (especially in new cable runs) to a suitable and safe earth ground before you connect them to the module.

For fiber, make sure you have the correct cable for the distances involved and the type of fiber ports that are used. The two options are single mode fiber (SMF) or multimode fiber (MMF). Make sure the ports on the devices that are connected together are both SMF, or both are MMF ports.

Note: For fiber connections, make sure the transmit lead of one port is connected to the receive lead of the other port. Connections for transmit-to-transmit and receive-to-receive do not work.

Ethernet and Fast Ethernet Maximum Transmission Distances

Transceiver Speed	Cable Type	Duplex Mode	Maximum Distance between Station
10 Mbps	Category 3 UTP	Full and half	328 ft (100 m)
10 Mbps	MMF	Full and half	1.2 mi (2 km)

100 Mbps	Category 5 UTP Category 5e UTP	Full and half	328 ft (100 m)
100 Mbps	Category 6 UTP	Full and half	328 ft (100 m)
100 Mbps	MMF	Half	1312 ft (400 m)
		Full	1.2 mi (2 km)
100 Mbps	SMF	Half	1312 ft (400 m)
		Full	6.2 mi (10 km)

For more details on the different types of cables/connectors, cable requirements, optical requirements (distance, type, patch cables, and so on.), how to connect the different cables, and which cables are used by most Cisco switches and modules, refer to Catalyst Switch Cable Guide .

Troubleshoot the Gigabit Ethernet

If you have device A connected to device B over a Gigabit link, and the link does not come up, perform this procedure.

Step-by-Step Procedure

1. Verify device A and B use the same GBIC, short wavelength (SX), long wavelength (LX), long haul (LH), extended wavelength (ZX), or copper UTP (TX). Both devices must use the same type of GBIC to establish link. An SX GBIC needs to connect with an SX GBIC. An SX GBIC does not link with an LX GBIC. Refer to Mode-Conditioning Patch Cord Installation Note for more information.

2. Verify distance and cable used per GBIC as defined in this table.

   1000BASE-T and 1000BASE-X Port Cabling Specifications

GBIC	Wavelength (nm)	Copper/Fiber Type	Core Size[1](Microns)	Modal Bandwidth (MHz / km)	Cable Distance[2]
WS-G54831000Base - T (copper)		Category 5 UTP Category 5e UTP Category 6 UTP			328 ft (100 m)
WS-G54841000BASE-	850	MMF	62.5 62.5 50.0 50.0	160 200 400 500	722 ft (220 m)

SX[3]					902 ft (275 m) 1640 ft (500 m) 1804 ft (550 m)
WS-G54861000BASE-LX/LH	1310	MMF[4]SMF	62.5 50.0 50.0 8.3/9/10	500 400 500 -	1804 ft (550 m) 1804 ft (550 m) 1804 ft (550 m) 6.2 miles (10 km)
WS-G54871000BASE-ZX[5]	1550	MMF SMF[6]	8.3/9/10 8.3/9/10		43.5 miles (70 km)[7]62.1 miles (100 km)

a. The numbers given for multimode fiber-optic cable refer to the core diameter. For single-mode fiber-optic cable, 8.3 microns refers to the core diameter. The 9-micron and 10-micron values refer to the mode-field diameter (MFD), which is the diameter of the portion of the fiber that is light-carrying. This area consists of the fiber core plus a small portion that covers the cladding. The MFD is a function of the core diameter, the wavelength of the laser, and the refractive index difference between the core and the cladding.

b. Distances are based on fiber loss. Multiple splices and substandard fiber-optic cable reduce the cable distances.

c. Use with MMF only.

d. When you use an LX/LH GBIC with 62.5-micron diameter MMF, you must install a mode-conditioning patch cord (CAB-GELX-625 or equivalent) between the GBIC and the MMF cable on both the transmit and receive ends of the link. The mode-conditioning patch cord is required for link distances less than 328 feet (100 m) or greater than 984 feet (300 m). The mode-conditioning patch cord prevents the over use of the receiver for short lengths of MMF and reduces differential mode delay for long lengths of MMF. Refer to Mode-Conditioning Patch Cord Installation Note for more information.

e. Use with SMF only.

f. Dispersion-shifted single-mode fiber-optic cable.

g. The minimum link distance for ZX GBICs is 6.2 miles (10 km) with an 8-dB attenuator installed at each end of the link. Without attenuators, the minimum link distance is 24.9 miles (40 km).

3. If either device has multiple Gigabit ports, connect the ports to each other. This tests each device and verifies that the Gigabit interface functions correctly. For example, you have a switch that has two Gigabit ports. Wire Gigabit port one to Gigabit port two. Does the link come up? If so, the port is good. STP blocks

on the port and prevents any loops (port one receive (RX) goes to port two transmit (TX), and port one TX goes to port two RX).

4. If single connection or Step 3 fails with SC connectors, loop the port back to itself (port one RX goes to port one TX). Does the port come up? If not, contact the TAC, as this can be a faulty port.

5. If steps 3 and 4 are successful, but a connection between device A and B cannot be established, loop ports with the cable that adjoins the two devices. Verify that there is not a faulty cable.

6. Verify that each device supports 802.3z specification for Gigabit auto-negotiation. Gigabit Ethernet has an auto-negotiation procedure that is more extensive than the one used for 10/100 Ethernet (Gigabit auto-negotiation spec: IEEE Std 802.3z-1998). When you enable link negotiation, the system auto-negotiates flow control, duplex mode, and remote fault information. You must either enable or disable link negotiation on both ends of the link. Both ends of the link must be set to the same value or the link cannot connect. Problems have been seen when you connect to devices manufactured before the IEEE 802.3z standard was ratified. If either device does not support Gigabit auto-negotiation, disable the Gigabit auto-negotiation, and it forces the link up. It takes 300msec for the card firmware to notify the software that a 10/100/1000BASE-TX link/port is down. The 300msec default debounce timer comes from the firmware polling timer to the linecards, which occurs every 300 msec. If this link is run in 1G (1000BASE-TX) mode, Gigabit sync, which occurs every 10msec, must be able to detect the link down faster. There is a difference in the link failure detection times when you run GigabitEthenet on copper versus GigabitEthernet over fiber. This difference in detection time is based on the IEEE standards.

Warning: Disable auto-negotiation and this hides link drops or physical layer problems. This is only required if end-devices such as older Gigabit NICs are used which cannot support IEEE 802.3z. Do not disable auto-negotiation between switches unless absolutely required to do so, as physical layer problems can go undetected, which results in STP loops. The alternative is to contact the vendor for software/hardware upgrade for IEEE 802.3z Gigabit auto-negotiation support.

For GigabitEthernet system requirements as well as Gigabit Interface Converters (GBICs), Coarse Wavelength Division Multiplexing (CWDM), and Small Form-Factor Pluggable (SFP) system requirements, refer to these documents:

- System Requirements to Implement Gigabit Ethernet on Catalyst Switches
- Catalyst GigaStack Gigabit Interface Converter Switch Compatibility Matrix
- Cisco Gigabit Ethernet Transceiver Modules Compatibility Matrix
- Cisco 10-Gigabit Ethernet Transceiver Modules Compatibility Matrix

For general configuration information and additional information on how to troubleshoot, refer to Configuring and Troubleshooting Ethernet 10/100/1000 MB Half/Full Duplex Auto-Negotiation .

Connected vs Notconnected

Most Cisco switches have a port in the notconnect state. This means it is currently not connected to anything, but it can connect if it has a good connection to another operational device. If you connect a good cable to two switch ports in the notconnect state, the link light must become green for both ports, and the port status must indicate connected. This means that the port is up as far as Layer 1 (L1) is concerned.

For Cisco IOS, you can use the **show interfaces** command to verify whether the interface is **up, line protocol is up (connected)** . The first **up** refers to the physical layer status of the interface. The **line protocol up** message shows the data link layer status of the interface and says that the interface can send and receive keepalives.

```
<#root>

Router#

show interfaces fastEthernet 6/1

FastEthernet6/1 is down, line protocol is down (notconnect)

!---

Reasons

: In this case,
!--- 1) A cable is not properly connected or not connected at all to this port.
!--- 2) The connected cable is faulty.
!--- 3) Other end of the cable is not connected to an active port or device.

!---

Note

: For gigabit connections, GBICs need to be matched on each
!--- side of the connection.
!--- There are different types of GBICs, depends on the cable and
!--- distances involved: short wavelength (SX),
!--- long-wavelength/long-haul (LX/LH) and extended distance (ZX).
!--- An SX GBIC needs to connect with an SX GBIC;
!--- an SX GBIC does not link with an LX GBIC. Also, some gigabit
!--- connections require conditioning cables,
!--- that depend on the lengths involved.
```

```
<#root>

Router#

show interfaces fastEthernet 6/1

FastEthernet6/1 is

up

, line protocol is

down

 (notconnect)

!--- The interface is up (or not in a shutdown state), but line protocol down.
!--- Reason: In this case, the device on the other side of the wire is a
!--- CatOS switch with its port disabled.
```

```
<#root>

Router#

show interfaces fastEthernet 6/1 status
```

```
Port    Name     Status        Vlan     Duplex    Speed  Type
Fa6/1

notconnect

    1        auto      auto    10/100BaseTX
```

If **show interfaces** shows up/line protocol up (connected) but you see errors increment in the output of either command, refer to the Common Port and Interface Problems section of this document for advice.

Troubleshoot the Most Common Port and Interface Commands for Cisco IOS

This table shows the most common commands used to troubleshoot the port or interface problems on switches that run Cisco IOS System Software on the Supervisor.

Note: The right hand column on the next table gives a brief description of what the command does and lists any exceptions to the use per platform.

If you have the output of the supported commands from your Cisco device, you can use Cisco CLI Analyzer to display potential issues and fixes.

Cisco IOS Commands	Description
show version	This command displays output similar to a Cisco router, like software image name and version information and system memory sizes. Helpful with the search for software/hardware incompatibilities (with the Release Notes or Software Advisor) and bugs (with the Bug Search Tool). **Note**: Only registered Cisco users can access internal Cisco tools and information.

show module	This command displays what cards are present in the switch, the version of software they are that run, and what state the modules are in: ok, faulty, and so on. This is helpful to diagnose a hardware problem on a module or port. For more information about how to troubleshoot hardware problems with the **show module** command, see the Port or Interface Status is disabled or shutdown or the Hardware Problems sections of this document.
show run-config	This command displays the current configuration file of the switch. Changes are saved to the config in Cisco IOS with the **write memory** command. This is helpful to use to determine whether a misconfiguration of the mod/port or interface can cause a problem.
show interfaces	The **show interface** command displays the administrative and operational status of a switch port, input and output packets, buffer failures, errors, and so on.
clear counters	Use the **clear counters** command to zero the traffic and error counters so that you can see if the problem is only temporary, or if the counters continue to increment. **Note**: The Catalyst 6500/6000 series switches do not clear the bit counters of an interface with the clear counters command. The only way to clear the bit counters in these switches is to reload.
show interfaces counters	This is the command to use on the Catalyst 6000,

	4000, 3550, 2950, and 3750 series.
show counters interface show controllers ethernet-controller	The **show counters interface** command was introduced in software version 12.1(13)E for the Catalyst 6000 series only and displays 32-bit and 64-bit error counters. For Cisco IOS on 2900/3500XL, 2950/2955, 3550, 2970 and 3750 series switches, the **show controllers Ethernet-controller** command displays discarded frames, deferred frames, alignment errors, collisions, and so on.
show interfaces counters	This is the command to use on the Catalyst 6000, 4000, 3550, 2950, and and 3750 series.
show diagnostic(s) show post	The command **show diagnostic** was introduced in 12.1(11b)E for the Catalyst 6000 series and **show diagnostics** (with an **s**) was introduced in for Catalyst 4000 Series. On the 2900/3500XL, 2950/2955, 3550, 2970 and 3750 series switches the equivalent command is **show post** which displays the results of the switch POST. For more information on troubleshoot hardware related errors on Catalyst switches, see the Hardware Problems section of this document.

Understand the Specific Port and Interface Counter Output for Cisco IOS

Most switches have some way to track the packets and errors that occur on a port or interface. The common commands used to find this type of information are described in the Most Common Port and Interface Troubleshooting Commands for Cisco IOS section of this document.

Note: There can be differences in the implementation of the counters across various platforms and releases. Although the values of the counters are largely accurate, they are not very precise by design. In order to pull the exact statistics of the traffic, it is suggested that you use a sniffer to monitor the necessary ingress and egress interfaces.

Excessive errors for certain counters usually indicate a problem. When you operate at half-duplex setup, some data link errors increment in Frame Check Sequence (FCS), alignment, runts, and collision counters are normal. Generally, a one percent ratio of errors to total traffic is acceptable for half-duplex connections. If the ratio of errors to input packets is greater than two or three percent, performance degradation can be noticed.

In half-duplex environments, it is possible for both the switch and the connected device to sense the wire and transmit at exactly the same time and result in a collision. Collisions can cause runts, FCS, and alignment errors due to the frame not completely copied to the wire, which results in fragmented frames.

When you operate at full-duplex, errors in FCS, Cyclic Redundancy Checks (CRC), alignment, and runt counters must be minimal. If the link operates at full-duplex, the collision counter is not active. If the FCS, CRC, alignment, or runt counters increment, check for a duplex mismatch. Duplex mismatch is a situation where the switch operates at full-duplex and the connected device operates at half-duplex, or vice versa. The results of a duplex mismatch are extremely slow performance, intermittent connectivity, and loss of connection. Other possible causes of data link errors at full-duplex are bad cables, faulty switch ports, or NIC software/hardware issues. See the Common Port and Interface Problems section of this document for more information.

Show Interfaces for Cisco IOS

The **show interfaces card-type {slot/port}** command is the used command for Cisco IOS on the Supervisor to display error counters and statistics. An alternative to this command (for Catalyst 6000, 4000, 3550, 2970 2950/2955, and 3750 series switches) is the **show interfacescard-type <slot/port> counters errors** command which only displays the interface error counters. Refer to Table 1 for explanations of the error counter output.

Note: For 2900/3500XL Series switches use the **show interfaces card-type {slot/port}** command with the **show controllers Ethernet-controller** command.

<#root>

```
Router#sh

interfaces fastEthernet 6/1

FastEthernet6/1 is

up, line protocol is up (connected)

   Hardware is C6k 100Mb 802.3, address is 0009.11f3.8848 (bia 0009.11f3.8848)


MTU 1500 bytes

, BW 100000 Kbit, DLY 100 usec,
      reliability 255/255, txload 1/255, rxload 1/255
   Encapsulation ARPA, loopback not set


Full-duplex, 100Mb/s

   input flow-control is off, output flow-control is off
   ARP type: ARPA, ARP Timeout 04:00:00
   Last input 00:00:14, output 00:00:36, output hang never
   Last clearing of "show interface" counters never
   Input queue: 0/2000/0/0 (size/max/drops/flushes); Total output drops: 0
   Queueing strategy: fifo
   Output queue :0/40 (size/max)
   5 minute input rate 0 bits/sec, 0 packets/sec
   5 minute output rate 0 bits/sec, 0 packets/sec
```

The **show interfaces** command output up to this point is explained here (in order) :

- **up, line protocol is up (connected)** - The first **up** refers to the physical layer status of the interface. The **line protocol up** message shows the data link layer status of the interface and says that the interface can send and receive keepalives.

- **MTU** - The Maximum Transmission Unit (MTU) is 1500 bytes for Ethernet by default (for the max data portion of the frame).
- **Full-duplex, 100Mb/s** - Full-duplex and 100Mbps is the current speed and duplex setup of the interface. This does not tell you whether autoneg was used to achieve this. Use the **show interfaces fastEthernet 6/1 status** command to display this:

  <#root>

  ```
  Router#

  show interfaces fastEthernet 6/1 status

  Port    Name               Status       Vlan       Duplex  Speed Type
  Fa6/1                      connected    1

  a-full  a-100

   10/100BaseTX

  !--- Autonegotiation was used to achieve full-duplex and 100Mbps.
  ```

- **Last input, output** - The number of hours, minutes, and seconds since the last packet was successfully received or transmitted by the interface. This is useful to know when a dead interface failed.
- **Last clearing of show interface counters** - The last time the **clear counters** command was issued since the last time the switch was rebooted. The **clear counters** command is used to reset interface statistics.

Note: Variables that can affect routing (for example, load and reliability) are not cleared when the counters are cleared.

- **Input queue** - The number of packets in the input queue.**Size/max/drops**= the current number of frames in the queue / the max number of frames the queue can hold before it must start to drop frames / the actual number of frames dropped because the max queue size was exceeded. **Flushes** is used to count Selective Packet Discard (SPD) drops on the Catalyst 6000 Series that run Cisco IOS. (The flushes counter can be used but never increments on the Catalyst 4000 Series that run Cisco IOS.) SPD is a mechanism that quickly drops low priority packets when the CPU is overloaded in order to save some process capacity for high priority packets. The flushes counter in the show interface command output increments as part of selective packet discard (SPD), which implements a selective packet drop policy on the IP process queue of the router. Therefore, it applies to only process switched traffic.

  The purpose of SPD is to ensure that important control packets, such as routing updates and keepalives, are not dropped when the IP input queue is full. When the size of the IP input queue is between the minimum and maximum thresholds, normal IP packets are dropped based on a certain drop probability. These random drops are called SPD flushes.

- **Total output drops** - The number of packets dropped because the output queue is full. A common cause is traffic from a high bandwidth link that is switched to a lower bandwidth link or traffic from multiple inbound links that are switched to a single outbound link. For example, if a large amount of traffic flow comes in on a gigabit interface and is switched out to a 100Mbps interface, this can cause output drops to increment on the 100Mbps interface. This is because the output queue on that interface is overwhelmed by the excess traffic due to the speed mismatch between the inbound and outbound bandwidths.

- **Output queue** - The number of packets in the output queue. Size/max means the current number of frames in the queue/the max number of frames the queue can hold before it is full and must start to drop the frames.
- **5 minute input/output rate** - The average input and output rate seen by the interface in the last five minutes. Specify a shorter period of time to get an accurate read (to better detect traffic bursts for example and issue the **load-interval <seconds>** interface command.

See Table 1 for explanations of the error counter output.

<#root>

```
!--- ...

show interfaces

 command output continues.
     1117058 packets input, 78283238 bytes, 0 no buffer
      Received 1117035 broadcasts, 0 runts, 0 giants, 0 throttles
      0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
      0 watchdog, 0 multicast, 0 pause input
      0 input packets with dribble condition detected
      285811 packets output, 27449284 bytes, 0 underruns
      0 output errors, 0 collisions, 2 interface resets
      0 babbles, 0 late collision, 0 deferred
      0 lost carrier, 0 no carrier
      0 output buffer failures, 0 output buffers swapped out
```

Note: **There is a difference between the counter of show interface command output for a physical interface and a VLAN interface.** The input packet counters increment in the output of **show interface** for a VLAN interface when that packet is Layer 3 (L3) processed by the CPU. Traffic that is Layer 2 (L2) switched never makes it to the CPU and is not counted in the **show interface** counters for the VLAN interface. It would be counted on the **show interface** output for the appropriate physical interface.

The **show interfaces** < **card-type> <slot/port> counters errors** command is used in Cisco IOS to display the output of the interface errors only. See Table 1 for explanations of the error counter output.

<#root>

```
Router#

show interfaces fastEthernet 6/1 counters errors


 Port        Align-Err    FCS-Err   Xmit-Err    Rcv-Err UnderSize OutDiscards
 Fa6/1               0          0          0          0         0           0

 Port      Single-Col Multi-Col  Late-Col Excess-Col Carri-Sen     Runts    Giants
 Fa6/1              0         0         0          0         0         0         0
```

Table 1. Cisco IOS error counter output for **show interfaces** or **show interfaces <** card-type> <x/y> counters errors for the Catalyst 6000 and 4000 Series.

Counters (in alphabetical order)	**Issues and Common Causes that Increase Error Counters**
Align-Err	**Description: show interfaces counters errors** . Alignment errors are a count of the number of frames received that do not end with an even number of octets and have a bad Cyclic Redundancy Check (CRC). **Common Causes:** These are usually the result of a duplex mismatch or a physical problem (such as cabling, a bad port, or a bad NIC). When the cable is first connected to the port, some of these errors can occur. Also, if there is a hub connected to the port, collisions between other devices on the hub can cause these errors. **Platform Exceptions:** Alignment errors are not counted on the Catalyst 4000 Series Supervisor I (WS-X4012) or Supervisor II (WS-X4013).
babbles	**Description: show interfaces** counter indicates that the transmit jabber timer expired. A jabber is a frame longer than 1518 octets (which exclude frame bits, but include FCS octets), which does not end with an even number of octets (alignment error) or has a bad FCS error.
Carri-Sen	**Description: show interfaces counters errors** . The Carri-Sen (carrier sense) counter increments every time an Ethernet controller wants to send data on a half-duplex connection. The controller senses the wire and checks if it is not busy before it transmits. **Common Causes:** This is normal on an half-duplex Ethernet segment.
collisions	**Descriptions: show interfaces** counter. The number of times a collision occurred before the interface transmitted a frame to the media successfully. **Common Causes:** Collisions are normal for interfaces configured as half-duplex but must not be seen on full duplex interfaces. If collisions increase dramatically, this points to a highly utilized link or possibly a duplex mismatch with the attached device.
CRC	**Description: show interfaces** counter. This increments when the CRC generated by the LAN station or far-end device that originates the traffic

	does not match the checksum calculated from the data received. **Common Causes:** This usually indicates noise or transmission problems on the LAN interface or the LAN itself. A high number of CRCs is usually the result of collisions but can also indicate a physical issue (such as cabling, bad interface or NIC) or a duplex mismatch.
deferred	**Description: show interfaces** counter. The number of frames that have been transmitted successfully after they wait because the media was busy. **Common Causes:** This is usually seen in half-duplex environments where the carrier is already in use when it tries to transmit a frame.
pause input	**Description: show interfaces** counter. An increment in pause input counter means that the connected device requests for a traffic pause when its receive buffer is almost full. **Common Causes:** This counter is incremented for informational purposes since the switch accepts the frame. The pause packets stop when the connected device is able to receive the traffic.
input packets with dribble condition	**Description: show interfaces** counter. A dribble bit error indicates that a frame is slightly too long.**Common Causes:**This frame error counter is incremented for informational purposes, since the switch accepts the frame.
Excess-Col	**Description**: **show interfaces counters errors** . A count of frames for which transmission on a particular interface fails due to excessive collisions. An excessive collision happens when a packet has a collision 16 times in a row. The packet is then dropped. **Common Causes** : Excessive collisions are typically an indication that the load on the segment needs to be split across multiple segments but can also point to a duplex mismatch with the attached device. Collisions must not be seen on interfaces configured as full duplex.
FCS-Err	**Description**: **show interfaces counters errors** . The number of valid size frames with Frame Check Sequence (FCS) errors but no frame errors. **Common Causes** : This is typically a physical issue (such as cabling, a bad port, or a bad Network Interface Card (NIC)) but can also indicate a duplex mismatch.

frame	**Description**: **show interfaces counter** . The number of packets received incorrectly that has a CRC error and a non-integer number of octets (alignment error). **Common Causes** : This is usually the result of collisions or a physical problem (such as cabling, bad port or NIC) but can also indicate a duplex mismatch.
Giants	**Description**: **show interfaces** and **show interfaces counters errors**. Frames received that exceed the maximum IEEE 802.3 frame size (1518 bytes for non-jumbo Ethernet) and have a bad Frame Check Sequence (FCS). **Common Causes**: In many cases, this is the result of a bad NIC. Try to find the offending device and remove it from the network. Platform Exceptions: Catalyst Cat4000 Series that run Cisco IOS Previous to software Version 12.1(19)EW, the giants counter incremented for a frame > 1518bytes. After 12.1(19)EW, a giant in show interfaces increments only when a frame is received >1518bytes with a bad FCS.
ignored	**Description**: **show interfaces** counter. The number of received packets ignored by the interface because the interface hardware ran low on internal buffers. **Common Causes** : Broadcast storms and bursts of noise can cause the ignored count to be increased.
Input errors	**Description**: **show interfaces** counter. **Common Causes** : This includes runts, giants, no buffer, CRC, frame, overrun, and ignored counts. Other input-related errors can also cause the input errors count to be increased, and some datagrams can have more than one error. Therefore, this sum cannot balance with the sum of enumerated input error counts. Also refer to the section Input Errors on a Layer 3 Interface Connected to a Layer 2 Switchport.
Late-Col	**Description**: **show interfaces** and **show interfaces counterserrors**. The number of times a collision is detected on a particular interface late in the transmission process. For a 10 Mbit/s port this is later than 512 bit-times into the transmission of a packet. Five hundred and twelve bit-times corresponds to 51.2 microseconds on a 10 Mbit/s system. Common Causes : This error can indicate a duplex mismatch among other things. For the duplex mismatch scenario, the late collision is seen on the half-duplex side. As the half-duplex side transmits, the full duplex side does not wait its turn and transmits simultaneously which causes a late

	collision. Late collisions can also indicate an Ethernet cable or segment that is too long. Collisions must not be seen on interfaces configured as full duplex.
lost carrier	**Description** : **show interfaces** counter. The number of times the carrier was lost in transmission. **Common Causes**: Check for a bad cable. Check the physical connection on both sides.
Multi-Col	**Description** : **show interfaces counters errors**. The number of times multiple collisions occurred before the interface transmitted a frame to the media successfully. **Common Causes**: Collisions are normal for interfaces configured as half-duplex but must not be seen on full duplex interfaces. If collisions increase dramatically, this points to a highly utilized link or possibly a duplex mismatch with the attached device.
no buffer	**Description**: **show interfaces** counter. The number of received packets discarded because there is no buffer space. **Common Causes**: Compare with ignored count. Broadcast storms can often be responsible for these events.
no carrier	**Description**: **show interfaces** counter. The number of times the carrier was not present in the transmission. **Common Causes**: Check for a bad cable. Check the physical connection on both sides.
Out-Discard	**Description**: The number of outbound packets chosen to be discarded even though no errors have been detected. **Common Causes**:One possible reason to discard such a packet can be to free up buffer space.
output buffer failures output buffers swapped out	**Description**: **show interfaces** counter. The number of failed buffers and the number of buffers swapped out. **Common Causes**: A port buffers the packets to the Tx buffer when the rate of traffic switched to the port is high and it cannot handle the amount of traffic. The port starts to drop the packets when the Tx buffer is full and thus increases the underruns and the output buffer failure counters. The increase in the output buffer failure counters can be a sign that the ports are run at an inferior speed and/or duplex, or there is too much traffic that goes through the port. As an example, consider a scenario where a 1gig

	multicast stream is forwarded to 24 100 Mbps ports. If an egress interface is over-subscribed, it is normal to see output buffer failures that increment along with Out-Discards. For troubleshoot information, see the Deferred Frames (Out-Lost or Out-Discard) section of this document.
output errors	**Description**: **show interfaces** counter. The sum of all errors that prevented the final transmission of datagrams out of the interface. **Common Cause**:This issue is due to the low Output Queue size.
overrun	**Description**:The number of times the receiver hardware was unable to hand received data to a hardware buffer. **Common Cause**:The input rate of traffic exceeded the ability of the receiver to handle the data.
packets input/output	**Description**: **show interfaces** counter. The total error free packets received and transmitted on the interface. Monitor these counters for increments as it is useful to determine whether traffic flows properly through the interface. The bytes counter includes both the data and MAC encapsulation in the error free packets received and transmitted by the system.
Rcv-Err	**Description**: For the Catalyst 6000 Series only - **show interfaces counters error** . **Common Causes**: See Platform Exceptions.**Platform Exceptions**: **Catalyst 5000 Series** rcv-err = receive buffer failures. For example, a runt, giant, or an FCS-Err does not increment the rcv-err counter. The rcv-err counter on a 5K only increments as a result of excessive traffic. On **Catalyst 4000 Series** rcv-err = the sum of all receive errors, which means, in contrast to the Catalyst 5000, that the rcv-err counter increments when the interface receives an error like a runt, giant or FCS-Err.
Runts	**Description**: **show interfaces** and **show interfaces counters errors** . The frames received that are smaller than the minimum IEEE 802.3 frame size (64 bytes for Ethernet), and with a bad CRC. **Common Causes**: This can be caused by a duplex mismatch and physical problems, such as a bad cable, port, or NIC on the attached device. **Platform Exceptions**: **Catalyst 4000 Series that run Cisco IOS** Previous to software Version 12.1(19)EW, a runt = undersize. Undersize = frame < 64bytes. The runt counter only

	incremented when a frame less than 64 bytes was received. After 12.1(19EW, a runt = a fragment. A fragment is a frame < 64 bytes but with a bad CRC. The result is the runt counter now increments in **show interfaces** , along with the fragments counter in **show interfaces counters errors** when a frame <64 bytes with a bad CRC is received. **Cisco Catalyst 3750 Series Switches** In releases prior to Cisco IOS 12.1(19)EA1, when dot1q is used on the trunk interface on the Catalyst 3750, runts can be seen on **show interfaces** output because valid dot1q encapsulated packets, which are 61 to 64 bytes and include the q-tag, are counted by the Catalyst 3750 as undersized frames, even though these packets are forwarded correctly. In addition, these packets are not reported in the appropriate category (unicast, multicast, or broadcast) in receive statistics. This issue is resolved in Cisco IOS release 12.1(19)EA1 or 12.2(18)SE or later.
Single-Col	**Description**: **show interfaces counters errors** . The number of times one collision occurred before the interface transmitted a frame to the media successfully. **Common Causes**:Collisions are normal for interfaces configured as half-duplex but must not be seen on full duplex interfaces. If collisions increase dramatically, this points to a highly utilized link or possibly a duplex mismatch with the attached device.
throttles	**Description**:**show interfaces**. The number of times the receiver on the port is disabled, possibly because of buffer or processor overload. If an asterisk (*) appears after the throttles counter value, it means that the interface is throttled at the time the command is run. **Common Causes**: Packets which can increase the processor overload include IP packets with options, expired TTL, non-ARPA encapsulation, fragmentation, tunnels, ICMP packets, packets with MTU checksum failure, RPF failure, IP checksum and length errors.
underruns	**Description**: The number of times that the transmitter has been that run faster than the switch can handle. **Common Causes**: This can occur in a high throughput situation where an interface is hit with a high volume of traffic bursts from many other interfaces all at once. Interface resets can occur along with the underruns.

Undersize	**Description**: **show interfaces counters errors** . The frames received that are smaller than the minimum IEEE 802.3 frame size of 64 bytes (which excludes frame bits but includes FCS octets) that are otherwise well formed. **Common Causes**:Check the device that sends out these frames.
Xmit-Err	**Description**: **show interfaces counters errors** . This is an indication that the internal send (Tx) buffer is full. **Common Causes**: A common cause of Xmit-Err can be traffic from a high bandwidth link that is switched to a lower bandwidth link, or traffic from multiple inbound links that are switched to a single outbound link. For example, if a large amount of traffic bursts comes in on a gigabit interface and is switched out to a 100Mbps interface, this can cause Xmit-Err to increment on the 100Mbps interface. This is because the output buffer of the interface is overwhelmed by the excess traffic due to the speed mismatch between the inbound and outbound bandwidths.

Show Interfaces Counters for Cisco IOS

To monitor inbound and outbound traffic on the port as displayed by the next output, for unicast, multicast, and broadcast traffic. The **show interfaces card-type {slot/port} counters** command is used when you run Cisco IOS on the Supervisor.

Note: There is, an Out-Discard counter in the Cisco IOS **show interfaces counters errors** command which is explained inTable 1.

```
<#root>

Router#

show interfaces fas 6/1 counters


  Port            InOctets   InUcastPkts   InMcastPkts   InBcastPkts
  Fa6/1           47856076            23        673028           149

  Port           OutOctets  OutUcastPkts  OutMcastPkts  OutBcastPkts
  Fa6/1           22103793            17        255877          3280
  Router#

!--- Cisco IOS counters used to monitor inbound and outbound unicast, multicast
!--- and broadcast packets on the interface.
```

Show Counters Interface for Cisco IOS

The **show counters interface card-type {slot/port}** command was introduced in Cisco IOS software version 12.1(13)E for the Catalyst 6000 series only, it offers even more detailed statistics for ports and interfaces. This command displays the 32-bit and 64-bit error counters per port or interface.

Show Controller Ethernet-Controller for Cisco IOS

For Catalyst 3750, 3550, 2970, 2950/2955, 2940, and 2900/3500XL switches use the command **show controller ethernet-controller** to display traffic counter and error counter output that is similar to theoutput for Catalyst 6000 series switches.

<#root>

```
3550-1#

show controller ethernet-controller fastEthernet 0/1

  !--- Output from a Catalyst 3550.

    Transmit FastEthernet0/1           Receive
          0 Bytes                        0 Bytes
          0 Unicast frames               0 Unicast frames
          0 Multicast frames             0 Multicast frames
          0 Broadcast frames             0 Broadcast frames
          0 Discarded frames             0 No dest, unicast
          0 Too old frames               0 No dest, multicast
          0 Deferred frames              0 No dest, broadcast
          0  1 collision frames
          0  2 collision frames          0 FCS errors
          0  3 collision frames          0 Oversize frames
          0  4 collision frames          0 Undersize frames
          0  5 collision frames          0 Collision fragments
          0  6 collision frames
          0  7 collision frames          0 Minimum size frames
          0  8 collision frames          0 65 to 127 byte frames
          0  9 collision frames          0 128 to 255 byte frames
          0 10 collision frames          0 256 to 511 byte frames
          0 11 collision frames          0 512 to 1023 byte frames
          0 12 collision frames          0 1024 to 1518 byte frames
          0 13 collision frames
          0 14 collision frames          0 Flooded frames
          0 15 collision frames          0 Overrun frames
          0 Excessive collisions         0 VLAN filtered frames
          0 Late collisions              0 Source routed frames
          0 Good (1 coll) frames         0 Valid oversize frames
          0 Good(>1 coll) frames         0 Pause frames
          0 Pause frames                 0 Symbol error frames
          0 VLAN discard frames          0 Invalid frames, too large
          0 Excess defer frames          0 Valid frames, too large
          0 Too large frames             0 Invalid frames, too small
          0 64 byte frames               0 Valid frames, too small
          0 127 byte frames
          0 255 byte frames
          0 511 byte frames
          0 1023 byte frames
          0 1518 byte frames

 3550-1#

!--- See the next table for additional counter output for 2900/3500XL Series switches.
```

Counter	**Description**	**Possible Causes**
Transmitted Frames		
Discarded frames	The total number of frames whose transmission attempt is abandoned due to insufficient resources. This total includes frames of all destination types.	The traffic load on the interface is excessive and causes the frames to be discarded. Reduce the traffic load on the interface if there are increments in the number of packets in this field.
Too old frames	Number of frames that took longer than two seconds to travel through the switch. For this reason, they were discarded by the switch. This only happens under extreme, high stress conditions.	The traffic load for this switch is excessive and causes the frames to be discarded. Reduce the switch load if the number of packets in this field increase. You can need to modify your network topology to reduce the traffic load for this switch.
Deferred frames	The total number of frames whose first transmission attempt was delayed, due to traffic on the network media. This total includes only those frames that are subsequently transmitted without error and not affected by collisions.	The traffic load destined for this switch is excessive and causes the frames to be discarded. Reduce the switch load if the number of packets in this field increase. You can need to modify your network topology to reduce the traffic load for this switch.
Collision frames	The collision frames counters are the number of times a packet was attempted to be transmitted but was not successful but was successful on its next attempt. This means that if the 2 collision frames counter incremented, the switch attempted to send the packet twice and failed but was successful on its third attempt.	The traffic load on the interface is excessive and causes the frames to be discarded. Reduce the traffic load on the interface if you see the number of packets increase in these fields.
Excessive collisions	The excessive collisions counter increases after 16 consecutive late collisions have occurred in a row. After 16 attempts have been made to send the packet the packet is dropped, and the counter	If this counter increments, it is an indication of a wiring problem, an excessively loaded network, or a duplex mismatch. An excessively loaded network can be caused by too many devices on a shared

	increments.	Ethernet.
Late collisions	A late collision occurs when two devices transmit at the same time, and neither side of the connection detects a collision. The reason for this occurrence is because the time to propagate the signal from one end of the network to another is longer than the time to put the entire packet on the network. The two devices that cause the late collision never see that each sends until after it puts the entire packet on the network. Late collisions are not detected by the transmitter until after the first 64 byte slot time. This is because they are only detected in transmissions of packets longer than 64 bytes.	Late collisions are a result of incorrect cabling or a non-compliant number of hubs in the network. Bad NICs can also cause late collisions.
Good (1 coll) frames	The total number of frames which experience exactly one collision and are then successfully transmitted.	Collisions in a half-duplex environment are normal expected behavior.
Good (>1 coll) frames	The total number of frames which experience between 2 and 15 collisions, inclusive, and are then successfully transmitted.	Collisions in a half-duplex environment are normal expected behavior. Frames that increment at the upper end of this counter can exceed the 15 collisions and can be counted as Excessive collisions.
VLAN discardframes	The number of frames dropped on an interface because the CFI bit is set.	The Canonical Format Indicator (CFI) bit in the TCI of an 802.1q frame is is set to 0 for the ethernet canonical frame format. If the CFI bit is set to 1, this indicates the presence of a RIF (Routing Information Field) or Token Ring noncanonical frame which is discarded.
Received Frames		
No bandwidth frames	*2900/3500XL only.*The number of times that a port received a packet from the network, but the switch	The traffic load on the interface is excessive and causes the frames to be discarded. Reduce the traffic

	did not have the resources to receive it. This only happens under stress conditions but can happen with bursts of traffic on several ports. So, a small number of No bandwidth frames is not a cause for concern. (It still must be far less than one percent of the frames received.)	load on the interface if you see the number of packets increase in these fields.
No buffers frames	*2900/3500XL only.*The number of times that a port received a packet from the network, but the switch did not have the resources to receive it. This only happens under stress conditions but can happen with bursts of traffic on several ports. So, a small number of No buffers frames is not a cause for concern. (It still must be far less than one percent of the frames received.)	The traffic load on the interface is excessive and causes the frames to be discarded. Reduce the traffic load on the interface if you see the number of packets increase in these fields.
No dest, unicast	No destination unicast are the number of unicast packets that the port did not forward to any other ports.	These are brief descriptions of when the No dest, (unicast, multicast, and broadcast) counters can increment: • If a port is an access port, and the port is connected to an Inter-Switch Link Protocol (ISL) trunk port, the No dest counter is very large since all inbound ISL packets are not forwarded. This is an invalid configuration. • If a port is blocked by Spanning Tree Protocol (STP), most packets are not forwarded, which results in No dest packets. If a port just acquired a link, there is a very brief (less than one second) period where inbound packets are not forwarded. • If the port is in a VLAN by itself, and no other ports on the switch belong to that
No dest, multicast	No destination multicast are the number of multicast packets that the port did not forward to any other ports.	
No dest, broadcast	No destination broadcast are the number of broadcast packets that	

	the port did not forward to any other ports.	VLAN, all inbound packets are dropped and the counter increments. • The counter also increments when the destination address of the packet is learned on the port that the packet was received on. If a packet was received on port 0/1, with destination MAC address X, and the switch has already learned that MAC address X resides on port 0/1, it increments the counter and discards the packet. This can happen in these situations: ◦ If a hub is connected to port 0/1, and a workstation connected to the hub transmits a packets to another workstation connected to the hub, port 0/1 does not forward this packet anywhere because the destination MAC resides on the same port. ◦ This can also occur if a switch is connected to port 0/1 and starts to flood packets to all of its ports to learn MAC addresses. • If a static address has been set up on another port in the same VLAN, and no static address was set up for the receiving port, the packet is dropped. For example, if a static map for MAC address X was configured on port 0/2 to forward traffic to port 0/3, the packet must be received on port 0/2 otherwise the packet is dropped. If a packet is sent from any other port, in the same VLAN as port 0/2, the packet is dropped.

		• If the port is a secure port, packets with disallowed source MAC addresses are not forwarded and increment the counter.
Alignment errors	Alignment errors are the number of frames received that do not end with an even number of octets and have a bad CRC.	Alignment errors are due to the frame that is not completely copied to the wire, which results in fragmented frames. Alignment errors are the result of collisions at half-duplex, a duplex mismatch, bad hardware (NIC, cable, or port), or connected devices that generate frames that do not end with an octet and have a bad FCS.
FCS errors	FCS error count is the number of frames that were received with a bad checksum (CRC value) in the Ethernet frame. These frames are dropped and not propagated onto other ports.	FCS errors are the result of collisions at half-duplex, a duplex mismatch, bad hardware (NIC, cable, or port), or connected devices that generate frames that do not end with an octet and have a bad FCS.
Undersize frames	These are the total number of packets received that were less than 64 octets long (which excludes frame bits but includes FCS) and have a good FCS value.	This is an indication of a bad frame generated by the connected device. Verify that the connected device operates correctly.
Oversize frames	Number of packets received by the port from the network, where the packets were more than 1514 bytes.	This can be an indication of faulty hardware, dot1q or ISL trunking configuration issues.
Collision fragments	The total number of frames whose length is less than 64 octets (which excludes frame bits, but includes FCS) and have a bad FCS value.	If this counter increments, this is an indication that the ports are configured at half-duplex. Set the duplex to full-duplex.
Overrun frames	The number of times the receiver hardware was unable to hand received data to a hardware buffer.	The input rate of traffic exceeded the ability of the receiver to handle the data.
VLAN filtered frames	The total number of frames which are filtered because of the type of	The port can be configured to filter 802.1Q tagged frames. When a

	VLAN information contained in the frame.	frame is received which contains an 802.1Q tag the frame is filtered and this statistic is incremented.
Source routed frames	The total number of receive frames that are discarded due to situation that the source route bit is set in the source address of the native frame.	This kind of source routing is only defined for Token Ring and FDDI. The IEEE ethernet specification forbids this bit to be set in any Ethernet frame. Therefore, the switch discards such frames.
Valid oversize frames	The total number of frames received whose length exceeds the System MTU, yet which have good FCS values.	This statistic counts frames that exceed the configured System MTU, but which can have been increased from 1518 bytes to allow for Q-in-Q or MPLS encapsulations.
Symbol error frames	Gigabit Ethernet (1000 Base-X) uses 8B/10B Encoding to translate 8bit data from the MAC sublayer(layer 2) to a 10bit Symbol to send over the wire. When a port receives a Symbol, it extracts the 8 bit data from the Symbol (10 bits).	A Symbol error means the interface detects an undefined (invalid) Symbol received. Small amounts of symbol errors can be ignored. Large amounts of symbol errors can indicate a bad device, cable, or hardware.
Invalid frames, too large	Giant frames or frames received that exceed the maximum IEEE 802.3 frame size (1518 bytes for non-jumbo Ethernet) and have a bad Frame Check Sequence (FCS).	In many cases, this is the result of a bad NIC. Try to find the offending device and remove it from the network.
Invalid frames, too small	Runt frames or frames received that are less than 64 bytes (which includes the FCS bits and excludes the frame header) and have either an FCS error or an alignment error.	This can be caused by a duplex mismatch and physical problems, such as a bad cable, port, or NIC on the attached device.

Common System Error Messages

For the Cisco IOS system messages format, you can refer to the **Messages and Recovery Procedures Guide** for the release of software you run. For example, you can look at the Messages and Recovery Procedures for Cisco IOS Releases.

%AMDP2_FE-3-UNDERFLO

This error message is caused when a frame is transmitted, and the local buffer of the controller chip local buffer receives insufficient data. The data cannot be transferred to the chip fast enough to keep pace with output rate. Normally, such a condition is temporary, dependent upon transient peak loads within the system. The issue occurs when an excessive amount of traffic is processed by the Fast Ethernet interface. The error message is received when the traffic level reaches about 2.5 Mb. This traffic level constrain is due to hardware limitation. Because of this, a chance exists for the device connected to the catalyst switch to drop packets.

The resolution is that ordinarily the system recovers automatically. No action is required. If the switch overwhelms the Ethernet interface, check the speed and duplex setup. Also, use a sniffer program to analyze packets that come in and out of the router fast Ethernet interface. In order to avoid packet drops on the device connected to the catalyst switch, issue the **ip cef** command on the fast Ethernet interface of the device connected to the switch.

%INTR_MGR-DFC1-3-INTR: Queueing Engine (Blackwater) [1]: FIC Fabric-A Received Unexpected Control Code

The reason for this error message is the receipt of a packet from the switch fabric, where the CRC value in the fabric header on that packet did not match the CRC value calculated by the Fabric Interface Controller (FIC) subblock of the Blackwater ASIC. This indicates that a corruption of the packet occurred within transfer, and Blackwater received the corrupted packet.

Command Rejected: [Interface] not a Switching Port

In switches that support both L3 interfaces and L2 switchport, the message "Command rejected: [interface] not a switching port″ displays when you try to enter a command related to layer 2 on a port that is configured as a layer 3 interface.

In order to convert the interface from layer 3 mode to layer 2 mode, issue the interface configuration command **switchport** . After you issue this command, configure the port for any layer 2 properties.

Common Port and Interface Problems

Port or Interface Status is Disable or Shutdown

An obvious but sometimes overlooked cause of port connectivity failure is an incorrect configuration on the switch. If a port has a solid orange light, this means the software inside the switch shut down the port, either by way of the user interface or by internal processes.

Note: Some port LEDs of the platform work differently in regard to STP. For example, the Catalyst 1900/2820 turns ports orange when they are in STP block mode. In this case, an orange light can indicate the normal functions of the STP. The Catalyst 6000/4000 does not turn the port light orange when it blocks for STP.

Make sure the port or module has not been disabled or powered down for some reason. If a port or module is manually shut down on one side of the link or the other, the link does not come up until you re-enable the port. Check the port status on both sides. Use the **show run interface** command and check to see if the interface is in a **shutdown** state:

```
<#root>

Switch#
```

```
show run interface fastEthernet 4/2

!
interface FastEthernet4/2
 switchport trunk encapsulation dot1q
 switchport mode trunk


shutdown

 duplex full
 speed 100
end

!--- Use the no shut command in config-if mode to re-enable this interface.
```

If the port goes into shutdown mode immediately after a reboot of the switch, the probable cause is the port security setup. If unicast flooding is enabled on that port, it can cause the port to shut down after a reboot. Cisco recommends that you disable the unicast flooding because it also ensure that no flooding occurs on the port once the MAC address limit is reached.

Port or Interface Status is errDisable

By default, software processes inside the switch can shut down a port or interface if certain errors are detected.

When you look at **show interfacecard-type {slot/port} status** command for Cisco IOS:

<#root>

```
Router#

show interface fastethernet 2/4 status


  Port    Name               Status       Vlan       Duplex  Speed Type
  Gi2/4

err-disabled

 1            full   1000 1000BaseSX

  !--- The show interfaces card-type {slot/port} status command for Cisco IOS
  !--- displays a status of errdisabled.
  !--- The show interfaces status errdisabled command shows all the interfaces
  !--- in this status.
```

The **show logging** command for Cisco IOS also display the error messages (exact message format varies) that relate to the errdisable state.

Wheb ports or interlaces are shut down as a result of errdisable are referred to as causes in Cisco IOS. The causes for this range from EtherChannel misconfiguration that causes a PAgP flap, duplex mismatch, BPDU port-guard and portfast configured at the same time, UDLD that detects a one-way link, and so on.

You have to manually enable the port or interface again to take it out the errdisable state unless you configure an errdisable recovery option. In Cisco IOS software you have the ability to automatically re-enable a port after a configurable amount of time spent in the errdisable state. The bottom line is that even if you configure the interface to recover from errdisable the problem reoccurs until the root cause is determined.

Note: Use this Recover Errdisable Port State on Cisco IOS Platforms for more information on errdisable status on switches that run Cisco IOS.

This table shows an example of the commands used to configure verify and troubleshoot the errdisable status on switches. Navigate to the link for more information about the commands Recover Errdisable Port State on Cisco IOS Platforms:

Action	Cisco IOS errdisable Commands
Configure	**errdisable detect cause**
Configure	**errdisable recovery cause**
Configure	**errdisable recovery interval <timer_interval_in_seconds>**
Verify & Troubleshoot	**show errdisable detect**
Verify & Troubleshoot	**show interfaces status err-disabled**

Port or Interface Status is Inactive

One common cause of inactive ports on switches that run Cisco IOS is when the VLAN they belong to disappears. This can occur when interfaces are configured as layer 2 switchports that use the **switchport** command.

Every port in a Layer 2 switch belongs to a VLAN. Every port on a Layer 3 switch configured to be a L2 switchport must also belong to a VLAN. If that VLAN is deleted, then the port or interface becomes inactive.

Note: Some switches show a steady orange (amber) light on each port when this happens.

Use the **show interfaces card-type {slot/port} switchport** command along with **show vlan** to verify.

```
<#root>

Router#

show interfaces fastEthernet 4/47 switchport

  Name: Fa4/47Switchport: Enabled
  Administrative Mode: static access
  Operational Mode: static access
  Administrative Trunking Encapsulation: negotiate
  Operational Trunking Encapsulation: native
  Negotiation of Trunking: Off
  Access Mode VLAN: 11 ((
Inactive

))

  !--- FastEth 4/47 is inactive.

 Router#

show vlan


 VLAN Name                             Status    Ports
 ---- -------------------------------- --------- -------------------------------
 1    default                          active    Gi1/1, Gi2/1, Fa6/6
 10   UplinkToGSR's                    active    Gi1/2, Gi2/2

 !--- VLANs are displayed in order and VLAN 11 is not available.

 30   SDTsw-1ToSDTsw-2Link             active    Fa6/45
```

If the switch that deleted the VLAN is a VTP server for the VTP domain, every server and client switch in the domain has the VLAN removed from their VLAN table as well. When you add the VLAN back into the VLAN table from a VTP server switch, the ports of the switches in the domain that belong to that restored VLAN become active again. A port remembers what VLAN it is assigned to, even if the VLAN itself is deleted. Refer to Understanding and Configuring VLAN Trunk Protocol (VTP) for more information on VTP.

Note: If the output of the **show interface <interface> switchport** command displays the port as a trunk port even after you configure the port as an access port with the **switchport access vlan <vlan>** command, issue the **switchport mode access** command in order to make the port an access port.

Uplink Port or Interface Status is Inactive

On a Catalyst 4510R series switch, in order to enable both the 10-Gigabit Ethernet and the Gigabit Ethernet SFP uplink ports, there is an optional configuration. In order to enable the simultaneous use of 10-Gigabit Ethernet and the Gigabit Ethernet SFP interfaces, issue the **hw-module uplink select all** command. After you issue the command, re-boot the switch or else the output of the **show interface status module <module number>** command shows the uplink port as inactive.

Cisco IOS Software Release 12.2(25)SG supports the simultaneous use of 10-Gigabit Ethernet and the Gigabit Ethernet SFP interfaces on Catalyst 4500 switches.

Note: On the Catalyst 4503, 4506, and 4507R series switches, this capability is automatically enabled.

Deferred Counter on the Catalyst Switch Interface Increments

The issue is because the traffic load destined for the switch is excessive and causes the frames to be discarded. Normally the deferred frames are the number of frames that have been transmitted successfully after waiting for the media, because the media was busy. This is usually seen in half-duplex environments where the carrier is already in use when it tries to transmit a frame. But in full duplex environments the issue occurs when the excessive load is destined for the switch.

This is the workaround:

- Hardcode both ends of the link to full duplex so that the negotiation mismatch can be avoided.
- Change the cable and patch panel cord to ensure that the cable and patch cords are not defective.

Note: If the Deferred Counter error increments on a GigabitEthernet of a Supervisor 720, turn on speed negotiation on the interface as a workaround.

Intermittent Failure to set timer [value] from vlan [vlan no]

The issue occurs when Encoded Address Recognition Logic (EARL) is unable to set the CAM aging time for the VLAN to the required number of seconds. Here, the VLAN aging time is already set to fast aging.

When the VLAN is already in fast aging, EARL cannot set the VLAN to fast aging, and aging timer set process is blocked. The default CAM aging time is five minutes, which means that the switch flushes the table of learned MAC addresses every five minutes. This ensures that the MAC address table (the CAM table) contains the most recent entries.

Fast aging temporarily sets the CAM aging time to the number of seconds that the user specifies, and is used in conjunction with the Topology Change Notification (TCN) process. The idea is that when a topology change occurs, this value is necessary to flush the CAM table faster, to compensate for the topology change.

Issue the **show cam aging** command to check the CAM aging time on the switch. TCNs and fast aging are fairly rare. As a result, the message has a severity level of 3. If the VLANs are frequently in fast aging, check the reason for fast aging.

The most common reason for TCNs is client PCs connected directly to a switch. When you power up or down the PC, the switch port changes state, and the switch starts the TCN process. This is because the switch does not know that the connected device is a PC; the switch only knows that the port has changed the state.

In order to resolve this issue, Cisco has developed the PortFast feature for host ports. An advantage of PortFast is that this feature suppresses TCNs for a host port.

Note: PortFast also bypasses spanning-tree calculations on the port, and is therefore only suitable for use on a host port.

Trunking Mode Mismatch

Check the trunking mode on each side of the link. Make sure both sides are in the same mode (both trunking with the same method: ISL or 802.1q, or both not trunking). If you turn the trunking mode to on (as opposed to auto or desirable) for one port and the other port has the trunking mode set to off, they are not able to communicate. Trunking changes the formatting of the packet. The ports need to be in agreement as to what format they use on the link, or they do not understand each other.

For Cisco IOS, use the **show interfaces card-type {mod/port}trunk** command to verify the trunking configuration and Native VLAN.

```
<#root>

Router#

show interfaces fastEthernet 6/1 trunk
```

```
  Port      Mode         Encapsulation  Status        Native vlan
  Fa6/1     desirable    802.1q

trunking      1


  Port      Vlans allowed on trunk
  Fa6/1     1-4094
!--- Output truncated.
```

Refer to these documents for more information on the different trunking modes, guidelines, and restrictions:

- System Requirements to Implement Trunking
- Trunking Technology Support Page

Jumbos, Giants, and Baby Giants

The Maximum Transmission Unit (MTU) of the data portion of an ethernet frame is 1500 bytes by default. If the transmitted traffic MTU exceeds the supported MTU the switch does not forward the packet. Also, dependent upon the hardware and software, some switch platforms increment port and interface error counters as a result.

- Jumbo frames are not defined as part of the IEEE Ethernet standard and are vendor-dependent. They can be defined as any frame bigger than the standard ethernet frame of 1518 bytes (which includes the L2 header and Cyclic Redundancy Check (CRC)). Jumbos have larger frame sizes, typically > 9000 bytes.
- Giant frames are defined as any frame over the maximum size of an ethernet frame (larger than 1518 bytes) that has a bad FCS.
- Baby Giant frames are just slightly larger than the maximum size of an ethernet frame. Typically this means frames up to 1600 bytes in size.

Support for jumbo and baby giants on Catalyst switches varies by switch platform, sometimes even by modules within the switch. The software version is also a factor.

Refer to Configuring Jumbo/Giant Frame Support on Catalyst Switches for more information on system requirements, configure and troubleshoot for jumbo and baby giant issues.

Cannot Ping End Device

Check the end device with a ping sent from the directly connected switch first, then work your way back port by port, interface by interface, trunk by trunk until you find the source of the connectivity issue. Make sure each switch can see the end device MAC address in its Content-Addressable Memory (CAM) table.

Use the **show mac address-table dynamic** command or substitute the **interface** keyword.

```
<#root>

Router#

show mac-address-table interface fastEthernet 6/3
```

```
Codes: * - primary entry

  vlan   mac address     type    learn qos            ports
------+----------------+--------+-----+---+--------------------------
*    2

 0040.ca14.0ab1

  dynamic  No    --  Fa6/3

!--- A workstation on VLAN 2 with MAC address 0040.ca14.0ab1 is directly connected
!--- to interface fastEthernet 6/3 on a switch running Cisco IOS.
```

Once you know the switch actually has the MAC address of the device in the CAM table, determine whether this device is on the same or different VLAN from where you try to ping.

If the end device is on a different VLAN from where you try to ping, a L3 switch or router must be configured to allow the devices to communicate. Make sure your L3 addressing on the end device and on the router/ L3 switch is correctly configured. Check the IP address, subnet mask, default gateway, dynamic routing protocol configuration, static routes, and so on.

Use of Switchport Host to Fix Startup Delays

If stations are not able to talk to their primary servers when they connect through the switch, the problem can involve delays on the switch port when it tries to become active after the physical layer link comes up. In some cases, these delays can be up to 50 seconds. Some workstations simply cannot wait this long to find their server and then they give up. These delays are caused by STP, trunking negotiations (DTP), and EtherChannel negotiations (PAgP). All of these protocols can be disabled for access ports where they are not needed, so the switch port or interface starts forwarding packets a few seconds after it establishes a link with its neighbor device.

In Cisco IOS, you can use the **switchport host** command to disable channeling and to enable spanning-tree portfast and the **switchport nonegotiate** command to turn off DTP negotiation packets. Use the **interface-range** command to do this on multiple interfaces at once.

```
<#root>

Router6k-1(config)#

interface range fastEthernet 6/13 - 18

Router6k-1(config-if-range)#

switchport

Router6k-1(config-if-range)#

switchport host

switchport mode can be set to access
spanning-tree portfast can be enabled
channel group can be disabled
!--- Etherchannel is disabled and portfast is enabled on interfaces 6/13 - 6/18.
Router6k-1(config-if-range)#

switchport nonegotiate

!--- Trunking negotiation is disabled on interfaces 6/13 - 6/18.
```

```
Router6k-1(config-if-range)#end
Router6k-1#
```

Cisco IOS has the option to use the **global spanning-tree portfast default** command to automatically apply portfast to any interface configured as a layer 2 access switchport. Check the Command Reference for your release of software to verify the availability of this command. You can also use the **spanning-tree portfast** command per interface, but this requires that you turn off trunking and etherchannel separately to help fix workstation startup delays.

Note: Refer to Using Portfast and Other Commands to Fix Workstation Startup Connectivity Delays for more information how to fix startup delays.

Speed/Duplex, auto-negotiation, or NIC Issues

If you have a large amount of alignment errors, FCS errors, or late collisions, this can indicate one of these:

- Duplex Mismatch
- Bad or Damaged Cable
- NIC Card Issues

Duplex Mismatch

A common issue with speed/duplex is when the duplex setup are mismatched between two switches, between a switch and a router or between the switch and a workstation or server. This can occur when you manually hardcode the speed and duplex or from auto-negotiation issues between the two devices.

If the mismatch occurs between two Cisco devices with the Cisco Discovery Protocol (CDP) enabled, you see the CDP error messages on the console or in the logging buffer of both devices. CDP is useful to detect errors, as well as port and system statistics on nearby Cisco devices. CDP is Cisco proprietary and works when you send packets to a well-known MAC address 01-00-0C-CC-CC-CC.

The example shows the log messages that result from a duplex mismatch between two Catalyst 6000 series switches that runs Cisco IOS. These messages generally tell you what the mismatch is and where it occurs.

```
Jun  2 11:16:45 %CDP-4-DUPLEX_MISMATCH: duplex mismatch discovered on FastEthernet6/2
(not half duplex), with TBA04251336 3/2 (half duplex).
```

Use the **show cdp neighbors card-type <slot/port> detail** command to display CDP information for Cisco neighbor devices.

```
<#root>

Router#

show cdp neighbors fastEthernet 6/1 detail

-------------------------
Device ID: TBA04251336
```

```
Entry address(es):
  IP address: 10.1.1.1
Platform: WS-C6006,  Capabilities: Trans-Bridge Switch IGMP
Interface:

FastEthernet6/1,  Port ID (outgoing port): 3/1

Holdtime : 152 sec
Version :
WS-C6006 Software, Version McpSW: 6.3(3) NmpSW: 6.3(3)
Copyright (c) 1995-2001 by Cisco Systems
!--- Neighbor device to FastEth 6/1 is a Cisco Catalyst 6000 Switch
!--- on port 3/1 running CatOS.
advertisement version: 2
VTP Management Domain: 'test1'
Native VLAN: 1

Duplex: full

!--- Duplex is full.
Router#
```

setup auto speed/duplex on one side and 100/Full-duplex on the other side is also a misconfiguration and can result in a duplex mismatch. If the switch port receives a lot of late collisions, this usually indicates a duplex mismatch problem and can place the port in an errdisable status in a result. The half-duplex side only expects packets at certain times, not at any time, and therefore counts packets received at the wrong time as collisions. There are other causes for late collisions besides duplex mismatch, but this is one of the most common reasons. Always set both sides of the connection to auto-negotiate speed/duplex or set the speed/duplex manually on both sides.

Use the **show interfaces <card-type> <slot/port> status** command to display speed and duplex setup as well as other information. Use the **speed** and **duplex** commands from interface configuration mode to hardcode both sides to 10 or 100 and half or full as necessary.

<#root>

```
Router#

show interfaces fastEthernet 6/1 status

Port    Name                Status      Vlan       Duplex  Speed Type
Fa6/1                       connected   1

a-full  a-100 10

/100BaseTX
```

If you use the**show interfaces**command without the **status** option, you see a setup for speed and duplex, but you do not know whether this speed and duplex was achieved through auto-negotiation or not.

<#root>

```
Router#

show interface fas 6/1

FastEthernet6/1 is up, line protocol is up (connected)
```

```
  Hardware is C6k 100Mb 802.3, address is 0009.11f3.8848 (bia 0009.11f3.8848)
  MTU 1500 bytes, BW 100000 Kbit, DLY 100 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set

  Full-duplex, 100Mb/s

!--- Full-duplex and 100Mbps does not tell you whether autoneg was used to achieve this.
!--- Use the sh interfaces fas 6/1 status command to display this.
```

Bad or damaged cable

Always check the cable for marginal damage or failure. A cable can be just good enough to connect at the physical layer, but it corrupts packets as a result of subtle damage to the wiring or connectors. Check or swap the copper or fiber cable. Swap the GBIC (if removable) for fiber connections. Rule out any bad patch panel connections or media convertors between source and destination. Try the cable in another port or interface if one is available and see if the problem continues.

Auto negotiation and NIC Card Issues

Problems sometimes occur between Cisco switches and certain third-party NIC cards. By default, Catalyst switch ports and interfaces are set to autonegotiate. It is common for devices like laptops or other devices to be set to autonegotiate as well, yet sometimes autonegotation issues occur.

In order to troubleshoot auto-negotiation problems it is often recommended to try and hardcode both sides. If neither auto-negotiation or hardcode setup seem to work, there can be a problem with the firmware or software on your NIC card. Upgrade the NIC card driver to the latest version available on the web site of the manufacture to resolve this.

Refer to Configuring and Troubleshooting Ethernet 10/100/1000 MB Half/Full Duplex Auto-Negotiation for details on how to resolve speed/duplex and auto-negotiation issues.

Refer to Troubleshooting Cisco Catalyst Switches to NIC Compatibility Issues for details on how to resolve third-party NIC issues.

Spanning Tree Loops

Spanning Tree Protocol (STP) loops can cause serious performance issues that masquerade as port or interface problems. In this situation, your bandwidth is used by the same frames over and over again, which leaves little room for legitimate traffic.

The STP loop guard feature provides additional protection against Layer 2 forwarding loops (STP loops). An STP loop is created when an STP block port in a redundant topology erroneously transitions to the forwarding state. This usually happens because one of the ports of a physically redundant topology (not necessarily the STP block port) no longer receives STP BPDUs. In its operation, STP relies on continuous reception or transmission of BPDUs based on the port role. The designated port transmits BPDUs, and the non-designated port receives BPDUs.

When one of the ports in a physically redundant topology no longer receives BPDUs, the STP conceives that the topology is loop free. Eventually, the block port from the alternate or backup port becomes designated and moves to a forwarding state. This situation creates a loop.

The loop guard feature makes additional checks. If BPDUs are not received on a non-designated port, and loop guard is enabled, that port is moved into the STP loop-inconsistent block state, instead of the listening / learning / forwarding state. Without the loop guard feature, the port assumes the designated port role. The port moves to the STP forwarding state and creates a loop. Refer to Configure STP with Loop Guard and BPDU Skew Detection for more information on the loop guard feature.

This document covers reasons that STP can fail, what information to look for to identify the source of the problem, and what kind of design minimizes STP risks.

Loops can also be caused by a uni-directional link. For more information, refer to the UDLD: One-Way link problems section of this document.

UDLD: One-Way Link

A unidirectional link is a link where traffic goes out one way, but no traffic is received in the ingress direction. The switch does not know that the link ingress direction is bad (the port thinks that the link is up and works).

A broken fiber cable or other cabling/port issues can cause this one-way only communication. These partially functional links can cause problems such as STP loops when the switches involved do not know that the link is partially broken. UDLD can put a port in errdisable state when it detects a unidirectional link. The command **udld aggressive-mode** can be configured on switches that run Cisco IOS (check release notes for command availability) for point-to-point connections between switches where unidirectional links cannot be tolerated. The use of this feature can help you identify difficult to find unidirectional link problems

Refer to Configure the UDLD Protocol Feature for configuration information on UDLD.

Deferred Frames (Out-Lost or Out-Discard)

If you have a large number of deferred frames, or Out-Discard (also referred to as Out-Lost on some platforms), it means that the switch output buffers have filled up and the switch had to drop these packets. This can be a sign that this segment is run at an inferior speed and/or duplex, or there is too much traffic that goes through this port.

Use the **show interfaces counters error** command to look at OutDiscards.

```
<#root>

Router#

show interfaces counters error

Port        Align-Err    FCS-Err   Xmit-Err    Rcv-Err UnderSize OutDiscards
Fa7/47              0          0          0          0         0           0
Fa7/48              0          0          0          0         0

2871800

Fa8/1               0          0          0          0         0

2874203

Fa8/2             103          0          0        103         0

2878032

Fa8/3             147          0          0        185         0           0
Fa8/4             100          0          0        141         0

2876405
```

```
Fa8/5               0          0          0          0         0

2873671

Fa8/6               0          0          0          0         0           2
Fa8/7               0          0          0          0         0           0

!--- The show interfaces counters errors command shows certain interfaces
!--- that increment in large amounts OutDiscards while others run clean.
```

Investigate these common causes of output buffer failures:

Inferior Speed/Duplex for the Amount of Traffic

Your network can send too many packets through this port for the port to handle at its current speed/duplex setup. This can happen where you have multiple high-speed ports flowing to a single (usually slower) port. You can move the device that hangs off this port to faster media. For example, if the port is 10 Mbps, move this device to a 100 Mbps or Gigabit port. You can change the topology to route frames differently.

Congestion Issues: Segment Too Busy

If the segment is shared, other devices on this segment can transmit so much that the switch has no opportunity to transmit. Avoid daisy-chained hubs whenever possible. Congestion can lead to packet loss. Packet loss causes retransmissions at the transport layer which in turn causes users to experience latency at the application level. You can upgrade 10Mbps links to 100Mbps or Gigabit Ethernet links when possible. You can remove some devices from crowded segments to other less populated segments. Make congestion avoidance a priority on your network.

Applications

At times the traffic transmission characteristics of the applications used can lead to output buffer problems. NFS file transfers that come from a Gigabit attached server that uses user datagram protocol (UDP) with a 32K window size is one example of an application setup that can bring out this type of problem. If you have checked or tried the other suggestions in this document (checked speed/duplex, no physical errors on the link, all the traffic is normal valid traffic, and so on), then reduce the unit size that is sent by the application which can help to alleviate this problem.

Software Problems

If you see behavior that can only be considered strange, you can isolate the behavior to a specific box, and you have looked at everything suggested so far, this can indicate software or hardware problems. It is usually easier to upgrade the software than it is to upgrade hardware. Change the software first.

Use the **show version** command to verify the current software version along with the **dir flash** : or **dir bootflash** : (dependent upon the platform) command to verify the available flash memory for the upgrade:

```
<#root>

Router#

show version

Cisco Internetwork Operating System Software
Cisco IOS (tm) Catalyst 4000 L3 Switch Software (cat4000-IS-M), Version 12.1(13)EW, EA
RLY DEPLOYMENT RELEASE SOFTWARE (fc1)
```

```
TAC Support: http://www.cisco.com/tac
Copyright (c) 1986-2002 by cisco Systems, Inc.
Compiled Fri 20-Dec-02 13:52 by eaarmas
Image text-base: 0x00000000, data-base: 0x00E638AC
ROM: 12.1(12r)EW
Dagobah Revision 71, Swamp Revision 24
trunk-4500 uptime is 2 weeks, 2 days, 6 hours, 27 minutes
System returned to ROM by redundancy reset
System image file is "bootflash:cat4000-is-mz.121-13.EW.bin"

!--- Typical Cisco IOS show version output.

Router#

dir bootflash

:
Directory of bootflash:/
1  -rw-     8620144   Mar 22 2002 08:26:21  cat4000-is-mz.121-13.EW.bin
61341696 bytes total (

52721424 bytes free

)

!--- Verify available flash memory on switch running Cisco IOS.
```

How to Upgrade Software

For information on how to upgrade software for your Cisco Switches , navigate to link, choose your platform and look at the Software Configuration section.

Hardware Software Incompatibility

There can be a situation where the software is not compatible with the hardware. This happens when new hardware comes out and requires special support from the software. For more information on software compatibility, use the Software Advisor tool.

Software Bugs

The operating system can have a bug. If you load a newer software version, it can often fix this. You can search known software bugs with the Software Bug Toolkit.

Corrupt Images

An image can have become corrupted. For information in regard to the recovery from corrupted images, choose your platform Switch and look at the Troubleshoot section.

Hardware Problems

Check the results of **show module** for Catalyst 6000 and 4000 series switches that run Cisco IOS.

Check the results of the POST results from the switch to see if there were any failures indicated for any part of the switch. Failures of any test of a module or port show an 'F' in the test results.

For Cisco IOS, on modular switches like the Cat6000, use the command **show diagnostics** . In order to see POST results per module, use the **show diagnostics module < module>** command.

<#root>

```
ecsj-6506-d2#

sh diagnostic module 3

  Current Online Diagnostic Level =

Minimal

  !--- The diagnostic level is set to

minimal

 which is a shorter,
  !--- but also less thorough test result.
  !--- You may wish to configure

diagnostic level complete

 to get more test results.
  Online Diagnostic Result for Module 3 : MINOR ERROR
  Online Diagnostic Level when Line Card came up = Minimal
  Test Results: (. = Pass, F = Fail, U = Unknown)
  1 . TestLoopback :
  Port  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24
  ----------------------------------------------------------------------------
        .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  F  F  F  F F  F

  !--- Notice the

MINOR ERROR

test result and failed loopback test which means
  !--- these ports are currently unusable.
  !--- Use the

hw-module

{mod}

reset

 command or, if necessary, physically reseat the
  !--- module to try and fix this problem.

  !--- If these steps fail, open a case with Cisco Technical Support.
```

Note: For Catalyst 3750, 3550, 2970 , 2950/2955, and 2900/3500XL Series switches use the **show post** command, which indicates a simple pass or fail for the hw status. Use the LEDs on these switches to help you understand the POST results.

For further information on how to troubleshoot hardware problems on Catalyst switches that run Cisco IOS, navigate to the Cisco Switches support pages, choose your platform and look at the Troubleshooting > Hardware section. For possible issues related to Field Notices, refer to Field Notices for LAN and ATM Switches.

Input Errors on a Layer 3 Interface Connected to a Layer 2 Switchport

By default, all layer 2 ports are in `dynamic desirable` mode, so the layer 2 port tries to form a trunk link and sends out DTP packets to the remote device. When a layer 3 interface is connected to a layer 2 switchport, it is not able to interpret these frames, which results in Input errors, WrongEncap errors, and Input queue drops.

In order to resolve this, change the mode of the switch port to `static access` or `trunk` as per your requirement.

```
<#root>

Switch2(config)#

interface fastEthernet1/0/12

Switch2(config-if)#

switchport mode access
```

Or

```
<#root>

Switch2(config)#

interface fastEthernet1/0/12

Switch2(config-if)#

switchport trunk encapsulation dot1q

Switch2(config-if)#

switchport mode trunk
```

Rapidly Increment Rx-No-Pkt-Buff Counter and Input Errors

The Rx-No-Pkt-Buff counter can increase on ports when it has blades, such as WS-X4448-GB-RJ45, WS-X4548-GB-RJ45, and WS-X4548-GB-RJ45V. Also, some packet drop incrementation is normal and is the result of traffic bursts traffic.

These types of errors increase rapidly, especially when the traffic that passes through that link is high or when it has devices such as servers connected to that interface. This high load of traffic oversubscribes the ports, which exhausts the input buffers and causes the Rx-No-Pkt-Buff counter and input errors to increase rapidly.

If a packet cannot be completely received because the switch is out of packet buffers, this counter is incremented once for every dropped packet. This counter indicates the internal state of the Switching ASICs on the Supervisor and does not necessarily indicate an error condition.

Pause Frames

When the receive part (Rx) of the port has its Rx FIFO queue filled and reaches the high water mark, the transmit part (Tx) of the port starts to generate pause frames with an interval value mentioned in it. The remote device is expected to stop / reduce the transmission of packets for the interval time mentioned in the pause frame.

If the Rx is able to clear the Rx queue or reach low water mark within this interval, Tx sends out a special pause frame that mentions the interval as zero (0x0). This enables the remote device to start to transmit packets.

If the Rx still works on the queue, once the interval time expires, the Tx sends a new pause frame again with a new interval value.

If Rx-No-Pkt-Buff is zero or does not increment and the TxPauseFrames counter increments, it indicates that our switch generates pause frames and the remote end obeys, hence Rx FIFO queue depletes.

If Rx-No-Pkt-Buff increments and TxPauseFrames also increments, it means that the remote end disregards the pause frames (does not support flow control) and continues to send traffic despite the pause frames. In order to overcome this situation, manually configure the speed and duplex, as well as disable the flow control, if required.

These types of errors on the interface are related to a traffic problem with the ports oversubscribed. The WS-X4448-GB-RJ45, WS-X4548-GB-RJ45, and WS-X4548-GB-RJ45V switching modules have 48 oversubscribed ports in six groups of eight ports each:

- Ports 1, 2, 3, 4, 5, 6, 7, 8
- Ports 9, 10, 11, 12, 13, 14, 15, 16
- Ports 17, 18, 19, 20, 21, 22, 23, 24
- Ports 25, 26, 27, 28, 29, 30, 31, 32
- Ports 33, 34, 35, 36, 37, 38, 39, 40
- Ports 41, 42, 43, 44, 45, 46, 47, 48

The eight ports within each group use common circuitry that effectively multiplexes the group into a single, non-block, full-duplex Gigabit Ethernet connection to the internal switch fabric. For each group of eight ports, the frames that are received are buffered and sent to the common Gigabit Ethernet link to the internal switch fabric. If the amount of data received for a port begins to exceed buffer capacity, flow control sends pause frames to the remote port to temporarily stop traffic and prevent frame loss.

If the frames received on any group exceeds the bandwidth of 1 Gbps, the device starts to drop the frames. These drops are not obvious as they are dropped at the internal ASIC rather than the actual interfaces. This can lead to slow throughput of packets across the device.

The Rx-No-Pkt-Buff does not depend on the total traffic rate. It depends on the amount of the packets that are stored in the Rx FIFO buffer of the module ASIC. The size of this buffer is only 16 KB. It is counted with short traffic bursts flow when some packets fill this buffer. Thus, Rx-No-Pkt-Buff on each port can be counted when the total traffic rate of this ASIC port group exceeds 1 Gbps, since WS-X4548-GB-RJ45 is 8:1 oversubscribed module.

When you have devices that need to carry a large amount of traffic through that interface, consider the use of one port of each group so that the common circuitry that shares a single group is not affected by this amount of traffic. When the Gigabit Ethernet switching module is not fully utilized, you can balancee the port connections across port groupings to maximize available bandwidth. For example, with the WS-X4448-GB-RJ45 10/100/1000 switching module, you can connect ports from different groups, such as ports 4, 12, 20, or 30 (in any order), before you connect ports from the same group, such as ports 1, 2, 3, 4, 5, 6, 7, and 8. If this does not solve the issue, you need to consider a module without any oversubscription of ports.

Understand Unknown Protocol Drops

Unknown protocol drops is a counter on the interface. It is caused by protocols that are not understood by the router/switch. This example of the **show run interface** command shows the unknown protocol drops on the GigabitEthernet 0/1 interface.

<#root>

```
Switch#

show run interface GigabitEthernet0/1

GigabitEthernet0/1 is up, line protocol is up
  Hardware is BCM1125 Internal MAC, address is 0000.0000.0000 (via 0000.0000)
  MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation 802.1Q Virtual LAN, Vlan ID  1., loopback not set
  Keepalive set (10 sec)
  Full-duplex, 1000Mb/s, media type is RJ45
  output flow-control is XON, input flow-control is XON
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:05, output 00:00:03, output hang never
  Last clearing of "show interface" counters 16:47:42
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     3031 packets input, 488320 bytes, 0 no buffer
     Received 3023 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 watchdog, 63107 multicast, 0 pause input
     0 input packets with dribble condition detected
     7062 packets output, 756368 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets

     2015 unknown protocol drops



    4762 unknown protocol drops

     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier, 0 pause output
     0 output buffer failures, 0 output buffers swapped out
```

Unknown protocol drops are normally dropped because the interface where these packets are received is not configured for this type of protocol, or it can be any protocol that the router does not recognize. For example, if you have two routers connected and you disable CDP on one router interface, this results in unknown protocol drops on that interface. The CDP packets are no longer recognized, and they are dropped.

Trunking between a Switch and a Router

Trunk links between a switch and a router can make the switchport go down. Trunk can come up after you disable and enable the switchport, but eventually the switchport can go down again.

In order to resolve this issue, complete these steps:

1. Make sure Cisco Discovery Protocol (CDP) runs between the switch and router and both can see each other.
2. Disable the **Keepalives** on the interface of the router.
3. Reconfigure the trunk encapsulation on both devices.

When the keepalives are disabled, the CDP enables link to operate normally.

Connectivity Issues due to Oversubscription

When you use either the WS-X6548-GE-TX or WS-X6148-GE-TX modules, there is a possibility that individual port utilization can lead to connectivity problems or packet loss on the surrounding interfaces. Refer to Interface/Module Connectivity Problems for more information on oversubscription.

Sub Interfaces in SPA Modules

In SPA modules, after you create a sub interface with 802.1Q, the same VLAN is not usable on the switch. Once you have encapsulation dot1q on a subinterface, you can no longer use that VLAN in the system because the 6500 or 7600 internally allocates the VLAN and makes that sub interface its only member. In order to resolve this issue, create trunk ports instead of sub interfaces. That way, the VLAN can be seen in all interfaces.

Troubleshoot Output Drops

Typically, the output drops can occur if QoS is configured and does not provide enough bandwidth to certain class of packets. It also occurs when the hardware hits an oversubscription.

For example, here you see a high amount of output drops on the interface GigabitEthernet 8/9 on a Catalyst 6500 Series Switch:

```
<#root>

Switch#

show interface GigabitEthernet8/9

GigabitEthernet8/9 is up, line protocol is up (connected)
  Hardware is C6k 1000Mb 802.3, address is 0013.8051.5950 (bia 0013.8051.5950)
  Description: Connection To Bedok_Core_R1 Ge0/1
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 18/255, rxload 23/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 1000Mb/s, media type is SX
  input flow-control is off, output flow-control is off
  Clock mode is auto
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:28, output 00:00:10, output hang never
  Last clearing of "show interface" counters never
Input queue: 0/2000/3/0 (size/max/drops/flushes);

Total output drops: 95523364

  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 94024000 bits/sec, 25386 packets/sec
  5 minute output rate 71532000 bits/sec, 24672 packets/sec
     781388046974 packets input, 406568909591669 bytes, 0 no buffer
     Received 274483017 broadcasts (257355557 multicasts)
     0 runts, 0 giants, 0 throttles
     3 input errors, 2 CRC, 0 frame, 0 overrun, 0 ignored
     0 watchdog, 0 multicast, 0 pause input
     0 input packets with dribble condition detected
     749074165531 packets output, 324748855514195 bytes, 0 underruns
     0 output errors, 0 collisions, 3 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier, 0 PAUSE output
     0 output buffer failures, 0 output buffers swapped out
```

In order to analyze the problem, collect the output of these commands:

- **show fabric utilization detail**
- **show fabric errors**
- **show platform hardware capacity**
- **show catalyst6000 traffic-meter**
- **show platform hardware capacity rewrite-engine drop**

Last Input Never from the Output of Show interface Command

This example of the show interface command shows the **Last input never** on the TenGigabitEthernet1/15 interface.

```
<#root>

Switch#

show interface TenGigabitEthernet1/15

TenGigabitEthernet1/15 is up, line protocol is up (connected)
  Hardware is C6k 10000Mb 802.3, address is 0025.84f0.ab16 (bia 0025.84f0.ab16)
  Description: lsnbuprod1 solaris
  MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 10Gb/s
  input flow-control is off, output flow-control is off
  ARP type: ARPA, ARP Timeout 04:00:00


Last input never

, output 00:00:17, output hang never
  Last clearing of "show interface" counters 2d22h
  Input queue: 0/2000/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 46000 bits/sec, 32 packets/sec
     52499121 packets input, 3402971275 bytes, 0 no buffer
```

```
          Received 919 broadcasts (0 multicasts)
          0 runts, 0 giants, 0 throttles
          0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
          0 watchdog, 0 multicast, 0 pause input
          0 input packets with dribble condition detected
          118762062 packets output, 172364893339 bytes, 0 underruns
          0 output errors, 0 collisions, 3 interface resets
          0 babbles, 0 late collision, 0 deferred
          0 lost carrier, 0 no carrier, 0 PAUSE output
          0 output buffer failures, 0 output buffers swapped out
```

This shows the number of hours, minutes, and seconds since the last packet was successfully received by an interface and processed locally on the router. This is useful to know when a dead interface has failed. This counter is updated only when packets are process switched, not when packets are fast switched. **Last input never** means there was no successful interface packet transfer to other end point or terminal. Usually this means there was no packet transfer relative to that entity.

Related Information

- **Troubleshooting Cisco Catalyst Switches to NIC Compatibility Issues**
- **Using PortFast and Other Commands to Fix Workstation Startup Connectivity Delays**
- **Configuring and Troubleshooting Ethernet 10/100/1000Mb Half/Full Duplex Auto-Negotiation**
- **Upgrade Software Images and Working with Configuration Files on Catalyst Switches**
- **Technical Support & Documentation - Cisco Systems**