Configure and Verify Ethernet 10/100/1000Mb Half/Full Duplex Auto-Negotiation

Contents

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

Background Information

When to Use Ethernet 10/100 Mb Auto-Negotiation

When to Use Ethernet 1000 Mb Auto-Negotiation

Auto-Negotiation on Catalyst Switches that Run Cisco IOS Software

Appendix A - Catalyst Switch Modules

Appendix B - Ethernet Crossover Cables

Appendix C - Explanation of Auto-MDIX and Supported Switch Platforms

Appendix D - Explanation of Fields in the show interfaces Command

- Output Fields for the show interface Command

Appendix E - Frequently Asked Questions

Related Information

Introduction

This document describes the guidelines to troubleshoot, as well as isolate and resolve, Ethernet auto-negotiation issues.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- How to troubleshoot issues with 10/100 Network Interface Cards (NICs).
- Gigabit negotiation
- Operational issues on specific Cisco platforms
- Operational issues with specific NICs
- Table that shows all possible settings and results of speed and duplex between a NIC and a switch.
- Discussion of the auto-negotiation protocol itself (includes FLP).

Note: Refer to Troubleshooting Cisco Catalyst Switches to NIC Compatibility Issues for more information on auto-negotiation.

Components Used

The information in this document is based on these software and hardware versions:

- Cisco IOS System Software

This equipment was used to create the examples in this document:

- A terminal
- A console cable suitable for the Supervisor Engine in the switch. Refer to Connecting a Terminal to the Console Port on Catalyst Switches for more information.
- Two Catalyst switches in a lab environment with cleared configurations.
- Two 10/100/1000 Mb TX full-duplex capable interfaces
- An Ethernet crossover cable

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.

Note: The **write erase** command was issued on each switch to ensure that they have default configurations.

Conventions

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

Background Information

This document provides a general description of auto-negotiation and explains the procedure to configure and verify auto-negotiation on Catalyst switches that run the Cisco IOS Software on both the Supervisor Engine and MSFC (Native). This document also shows an example of why the most common duplex-mismatch error occurs and describes how to configure and verify auto-negotiation on Catalyst switches that run Cisco IOS® System Software.

Note: The Catalyst switches/modules, such as the Catalyst 6500/6000, 4500/4000, 3550, and 2950, support 10/100/1000 Mbps negotiated Ethernet interfaces or ports. These ports work on 10 Mbps, 100 Mbps, or 1000 Mbps speed based on their connection to the other end. These 10/100/1000 Mbps ports can be configured for speed and duplex negotiation similar to 10/100 Mbps ports Cisco IOS Software-based switches. Therefore, the configurations described in this document for 10/100 Mbps port negotiation apply to 10/100/1000 Mbps ports as well.

When to Use Ethernet 10/100 Mb Auto-Negotiation

Auto-negotiation is an optional function of the IEEE 802.3u Fast Ethernet standard that enables devices to automatically exchange information over a link about speed and duplex abilities.

Auto-negotiation is targeted at ports. These ports are allocated to areas where transient users or devices connect to a network. For example, many companies provide shared offices or cubes for Account Managers and System Engineers to use when they are in the office. Each office or cube has an Ethernet port permanently connected to the office network. Because it cannot be possible to ensure that every user has either a 10 Mb, a 100 Mb Ethernet, or a 10/100 Mb card in their laptop, the switch ports that handle these connections must be able to negotiate their speed and duplex mode. The alternative is to provide both a 10 Mb and a 100 Mb port in each office, or cube and label them accordingly.

One of the most common causes of performance issues on 10/100 Mb Ethernet links occurs when one port on the link operates at half-duplex while the other port operates at full-duplex. This occurs when one or both ports on a link are reset and the auto-negotiation process does not result in both link partners with the same configuration. It also can occur when users reconfigure one side of a link and forget to reconfigure the other side. Both sides of a link must have auto-negotiation on, or both sides must have it off. Cisco recommends to leave auto-negotiation on for those devices compliant with 802.3u.

Many performance-related support calls are avoided if you correctly configure auto-negotiation. Many Catalyst Ethernet switching modules support 10/100 Mb and half-duplex or full-duplex. Exceptions include the Ethernet Group switch modules. The **show interfaces capabilities** command shows if the interface or module you work on supports 10/100/1000 Mb and half-duplex or full-duplex. This document uses two WS-X5530 Supervisor Engine IIIs, each with two optional uplink 10/100 BaseTX Ethernet ports installed.



Note: When the WS-6748-GE-TX module is connected to a network tap device, automatic negotiation does not work. In order to resolve this issue, you must configure auto-negotiation manually. Go to the interface mode and execute this command:

```
<#root>

Cat6K-IOS(config-if)#

speed auto
```

When to Use Ethernet 1000 Mb Auto-Negotiation

Basically, auto-negotiation in GigabitEthernet covers these items:

- Duplex settings —While Cisco devices only support full-duplex, the IEEE 802.3z standard does have support for half-duplex GigabitEthernet. Because of this, duplex is negotiated between GigabitEthernet devices.
- Flow Control —Because of the amount of traffic that can be generated by GigabitEthernet, there is a PAUSE functionality built into GigabitEthernet. The PAUSE frame is a packet that tells the far-end device to stop the transmission of packets until the sender is able to handle all the traffic and clear its buffers. The PAUSE frame has a timer included which tells the far-end device when to start to send packets again. If that timer expires without another PAUSE frame sent, the far-end device can then send packets again. Flow-Control is an optional item and must be negotiated. Devices can send or receive to a PAUSE frame, and they possibly do not agree to the flow-control request of the far-end neighbor.

- Negotiation —Usually built-in Gigabit Ethernet ports are capable of negotiation, but in cases like modular SFP or GBIC types, they do not negotiate. Line protocol can be down for a Gigabit Ethernet port when connected to a Fast Ethernet port. This can be verified via the **show interfaces interface capabilities** command:

```
<#root>

Switch#

show interfaces Gig 5/3 capabilities


GigabitEthernet5/3
  Model:                 VS-S720-10G
  Type:                  10/100/1000BaseT

  Speed:                 10,100,1000,auto
  Duplex:                half,full

  Trunk encap. type:     802.1Q,ISL
  Trunk mode:            on,off,desirable,nonegotiate
  Channel:               yes
  Broadcast suppression: percentage(0-100)


Flowcontrol:           rx-(off,on,desired),tx-(off,on,desired)

  Membership:            static
  Fast Start:            yes
  QOS scheduling:        rx-(2q4t), tx-(1p3q4t)
  QOS queueing mode:     rx-(cos), tx-(cos)
  CoS rewrite:           yes
  ToS rewrite:           yes
  Inline power:          no
  SPAN:                  source/destination
  UDLD                   yes
  Link Debounce:         yes
  Link Debounce Time:    no
  Ports-in-ASIC (Sub-port ASIC) : 1-5 (3-4)
  Remote switch uplink:  no
  Port-Security:         yes
  Dot1x:                 yes
```

Assume that there are two devices, A and B. Assume that each device can have auto-negotiation enabled or disabled. The correct behavior of link status with auto-negotiation in accordance to the IEEE Std 802.3z-1998 must be like this:

- If A is enabled and B is enabled, then link status must be reported on both devices as link up.
- If A is disabled and B is enabled, then A must report link up and B must report link down.
- If A is enabled and B is disabled, then A must report link down and B must report link up.

By default, all devices are supposed to perform auto-negotiation. 802.3z does not specifically define a way to turn auto-negotiation off, for both 1GigabitEthernet and 10GigabitEthernet.

Auto-Negotiation on Catalyst Switches that Run Cisco IOS

Software

The commands described in this section apply to different types of Catalyst switch products that runs Cisco IOS System Software such as Catalyst 4500, and the Catalyst 6500. There are some outputs taken from Catalyst 3850 and 9500 platforms as well. Devices in this section were connected with an Ethernet crossover cable. See Appendix B for more information on crossover cables and Auto-MDIX feature.

The switches that run Cisco IOS Software default to auto-negotiation for speed and are set to on for the duplex. Run the **show interface interface status** command to verify these settings.

The first output is taken from a Catalyst 6500/6000 that runs Cisco IOS Software Release 12.1(6)E. It shows a connected port that auto-negotiates a link to 100 Mbps and half-duplex. The configuration that runs for this switch has no duplex or speed commands underneath interface FastEthernet 3/1 because auto-negotiation is the default. Issue the **show interface interface** command (without the **status** keyword) to see the port speed and duplex.

The **a** prefixes on the half and 100 indicate that this port is not hard coded (configured) for a specific duplex mode or speed. Therefore, it auto-negotiates the duplex mode and speed if the device it is connected to also auto-negotiates duplex mode and speed. The status is connected, which means that a link pulse is detected from the other port. The status can be connected even if duplex is incorrectly negotiated or incorrectly configured. Also, notice that there is no speed or duplex commands under the interface configuration, this is because auto-negotiate speed and duplex is the default configuration.

```
<#root>

NativeIOS#

show interfaces fastethernet 3/1 status


Port    Name               Status       Vlan       Duplex Speed Type
Fa3/1                      connected    routed

a

-half

a

-100 10/100BaseTX

NativeIOS#

show run

...
!
interface FastEthernet3/1
 ip address 172.16.84.110 255.255.255.0
!

NativeIOS#

show interfaces fastethernet 3/1

FastEthernet3/1 is up, line protocol is up
 Hardware is C6k 100Mb 802.3, address is 0002.7ef1.36e0 (bia 0002.7ef1.36e0)
 Internet address is 172.16.84.110/24
 MTU 1500 bytes, BW 100000 Kbit, DLY 100 usec,
    reliability 255/255, txload 1/255, rxload 1/255
```

```
 Encapsulation ARPA, loopback not set
 Keepalive set (10 sec)

Half-duplex, 100Mb/s

 ...
```

If you want to hard code the speed and duplex on a switch that runs Cisco IOS Software (turn off auto-negotiation), issue the **speed** and **duplex** commands underneath the specific interface. Duplex is subservient to speed in the sense that if speed is set to auto, then the duplex cannot be manually set. You can see cyclic redundancy check (CRC) error messages when both the speed and duplex settings are hardcoded on the two devices. This can be because any one of the devices runs an earlier version of Cisco IOS. You can upgrade the Cisco IOS or set the speed and duplex to auto on both devices in order to resolve this.



Note: If you hard code the speed on a port, it disables all auto-negotiation functionality on the port for speed and duplex.

```
<#root>

NativeIOS#

show run

...
interface FastEthernet3/2
 no ip address
!
NativeIOS#

configure terminal

Enter configuration commands, one per line.  End with CNTL/Z.
NativeIOS(config)#

interface fastethernet3/2

NativeIOS(config-if)#

duplex full

Duplex will not be set until speed is set to non-auto value

!--- Error: On this platform, you must set the speed before the duplex.
!--- Not all switch platforms have this command ordering requirement.

NativeIOS(config-if)#

speed 100


NativeIOS(config-if)#

duplex full


NativeIOS(config-if)#

^Z
```

```
NativeIOS#
show interfaces fastethernet 3/2 status
Port      Name               Status       Vlan       Duplex  Speed Type

Fa3/2     notconnect         routed
full
 100 10/100BaseTX

NativeIOS#
show run

...
interface FastEthernet3/2
no ip address
duplex full

speed 100

!--- Notice that the speed and duplex commands appear in the configuration
!--- now because they have been manually set to a non-default behavior.
```

The next outputs were taken from a 3850 and a 9500 Catalyst switches. In this example, these two switches are directly connected on one side speed, and duplex was hardcoded, and on the other side, auto-negotiation is used. As it is observed, the absence of the **a** prefix in the status fields of the output from the show interface TwentyFiveGigE1/0/2 status command on Switch_1 shows that the duplex mode is configured for full and the speed is configured for 1000.

<#root>

```
Switch_1#
show run interface TwentyFiveGigE1/0/2

Building configuration...

Current configuration : 37 bytes
!
interface TwentyFiveGigE1/0/2
end

Switch_1#
configure terminal

Enter configuration commands, one per line.  End with CNTL/Z.
Switch_1(config)#
interface TwentyFiveGigE1/0/2

Switch_1(config-if)#
```

```
duplex full

Switch_1(config-if)#

speed 1000

Switch_1(config-if)#

end

*Aug  1 19:26:33.957: %LINEPROTO-5-UPDOWN: Line protocol on Interface TwentyFiveGigE1/0/2, changed stat
*Aug  1 19:26:34.913: %SYS-5-CONFIG_I: Configured from console by console
*Aug  1 19:26:34.957: %LINK-3-UPDOWN: Interface TwentyFiveGigE1/0/2, changed state to down

*Aug  1 19:26:38.819: %LINK-3-UPDOWN: Interface TwentyFiveGigE1/0/2, changed state to up
*Aug  1 19:26:39.820: %LINEPROTO-5-UPDOWN: Line protocol on Interface TwentyFiveGigE1/0/2, changed state


Switch_1#

show interface TwentyFiveGigE1/0/2 status


Port         Name               Status       Vlan       Duplex  Speed Type
Twe1/0/2

connected

    1

full   1000

 10/100/1000BaseTX SFP

Switch_1#

show cdp neighbors TwentyFiveGigE1/0/2

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone,
                  D - Remote, C - CVTA, M - Two-port Mac Relay

Device ID        Local Intrfce     Holdtme    Capability  Platform  Port ID
Switch_2         Twe 1/0/2         124              S I   WS-C3850- Gig 1/0/1

Total cdp entries displayed : 1

Switch_2#

show run interface GigabitEthernet1/0/2

Building configuration...

Current configuration : 38 bytes
!
interface GigabitEthernet1/0/2
end

Switch_2#

show interfaces GigabitEthernet1/0/2 status


Port         Name               Status       Vlan       Duplex  Speed Type
Gi1/0/2

connected
```

```
    1
a
-full
a
-1000 10/100/1000BaseTX
```

If you try to configure half duplex on a GigabitEthernet interface, an error message similar to the next output can be seen:

```
<#root>
Switch_1#
 configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch_1(config)#
interface twentyFiveGigE 1/0/2
Switch_1(config-if)#
duplex half

% Duplex cannot be set to half when speed autonegotiation subset contains 1Gbps,2.5Gbps,5Gbps or 10Gbps
```

Only interfaces with a speed of 100 can accept the half duplex configuration:

```
<#root>
Switch_1(config-if)#
speed 100
Switch_1(config-if)#
duplex half

Switch_1(config-if)#
Switch_1(config-if)#
speed 1000
Cannot change speed to 1000Mbps when in half duplex
Switch_1(config-if)#
end
Switch_1#
```

The next message is about a duplex mode mismatch. It is displayed on a switch after it detects that there is a duplex mismatch on the interface. This mismatch can occur due to a misconfiguration on the device connected on interface GigabitEthernet2/0/20:

```
%CDP-4-DUPLEX_MISMATCH: duplex mismatch discovered on GigabitEthernet2/0/20 (not half duplex), with XXXXX GigabitEthernet0 (half
```

It is important to notice that this message is created by the Cisco Discovery Protocol (CDP), not the 802.3 auto-negotiation protocol. CDP can report problems it discovers, but it does not automatically fix them.

A duplex mismatch can or cannot result in an error message. Another indication of a duplex mismatch is the rapid increase of FCS and alignment errors on the half-duplex side, and runs on the full-duplex port.

Appendix A - Catalyst Switch Modules

This document contains information on how to install Catalyst modules and the functionality of each module. It also contains explanations of the LEDs on each module. In general, the LEDs indicate the status of the module as well as which ports are active.

Appendix B - Ethernet Crossover Cables

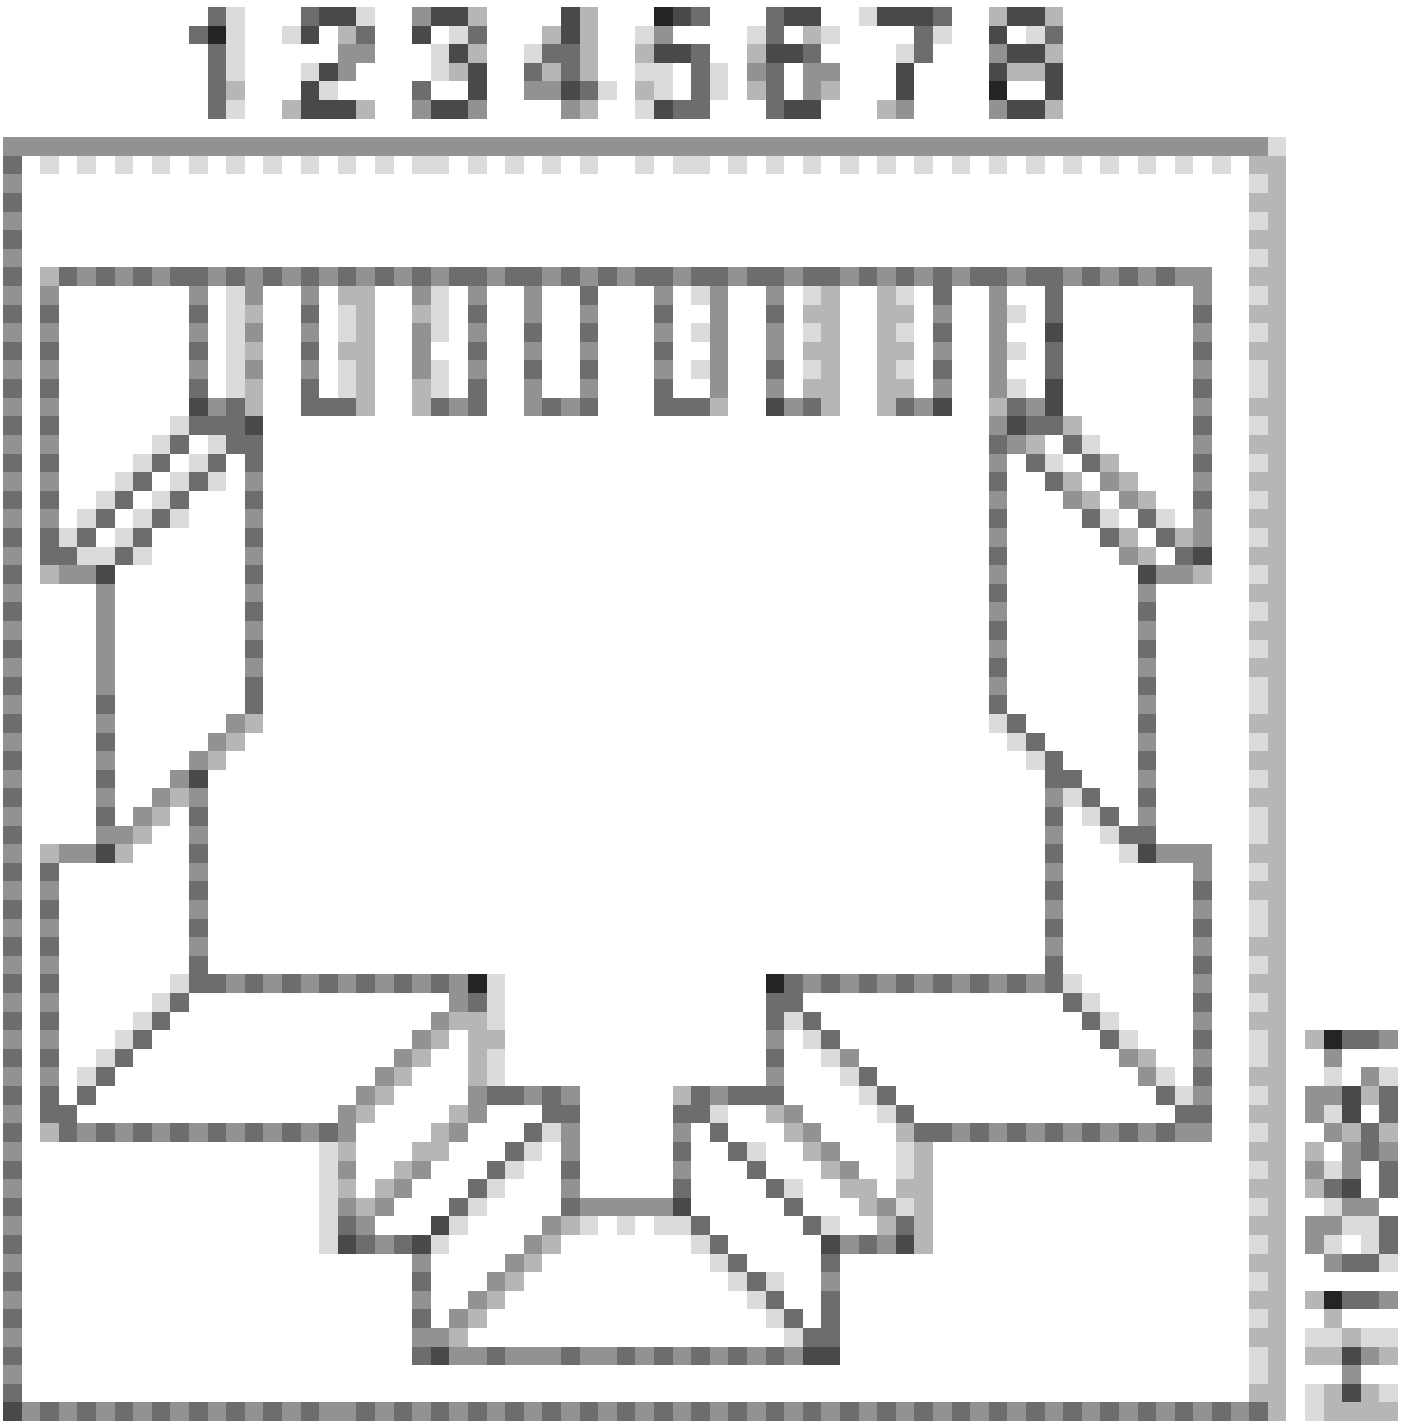
Ethernet ports on Catalyst switches have built-in (on-board) Ethernet transceivers. Devices that connect to Ethernet ports can have on-board Ethernet transceivers or use external transceivers.

Use a straight-through patch cable, such as a CAT5/CAT6 10/100/1000BaseT unshielded twisted pair (UTP) patch cable, when you connect a PC, server, printer, or other end-user devices (such as a router) to a switch. Straight-through means that pin 1 on one end of the cable is connected to pin 1 on the other end, pin 2 on one end of the cable is connected to pin 2 on the other end, and so forth.

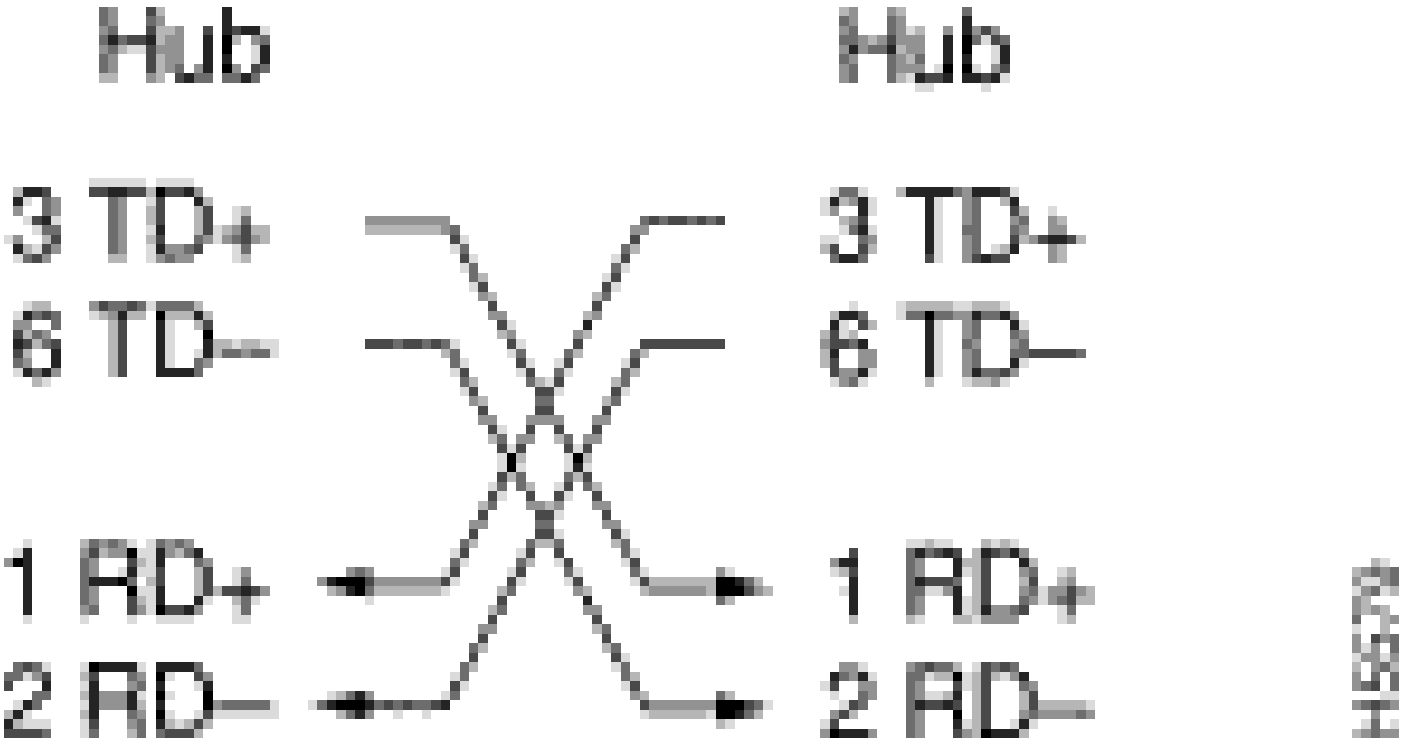
Use a crossover cable, such as a CAT5/CAT6 10/100/1000BaseT UTP crossover patch cable, when you connect another switch port, or other Layer 2 port to an Ethernet port on a switch. In this case, the pins are connected (see Images).

A convenient rule of thumb is to use a crossover cable when the two ports that are connected are in the same layer of the OSI model. If you cross OSI layers, use a straight-through cable. Treat PCs as Layer 3 ports, hubs and most Layer 3 switches as Layer 2 ports. Some devices, especially common on hubs, have a button that can toggle that accepts a straight-through or crossover cable. Therefore, this rule of thumb does not always apply.

Note: Use a crossover cable when you connect two ports in the same layer of the OSI model, such as router to router (Layer 3) or switch to switch (Layer 2). Use a straight-through cable if the two ports are in different layers, such as router to switch (Layer 3 to 2) or PC to switch (Layer 3 to 2). For this rule, treat a PC as a Layer 3 device.



Ethernet Crossover Cables

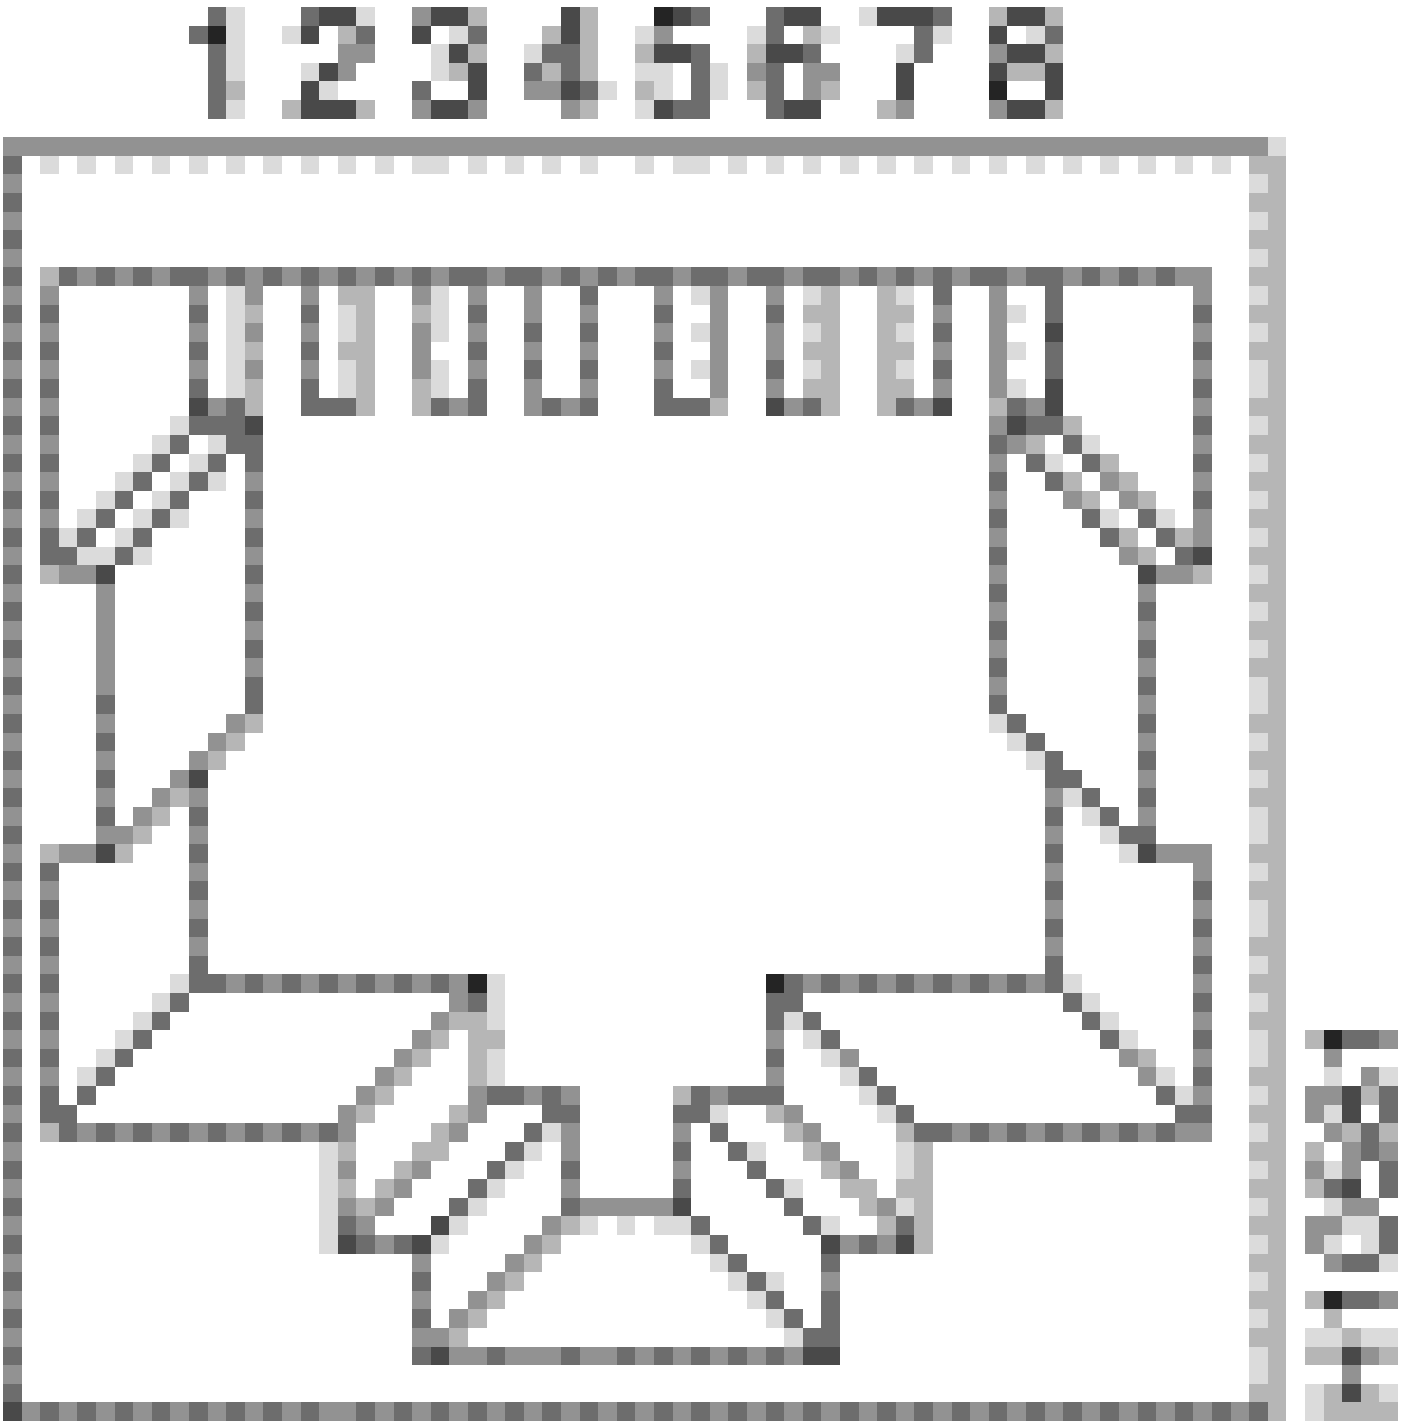


Ethernet Crossover Cables Available at Stores

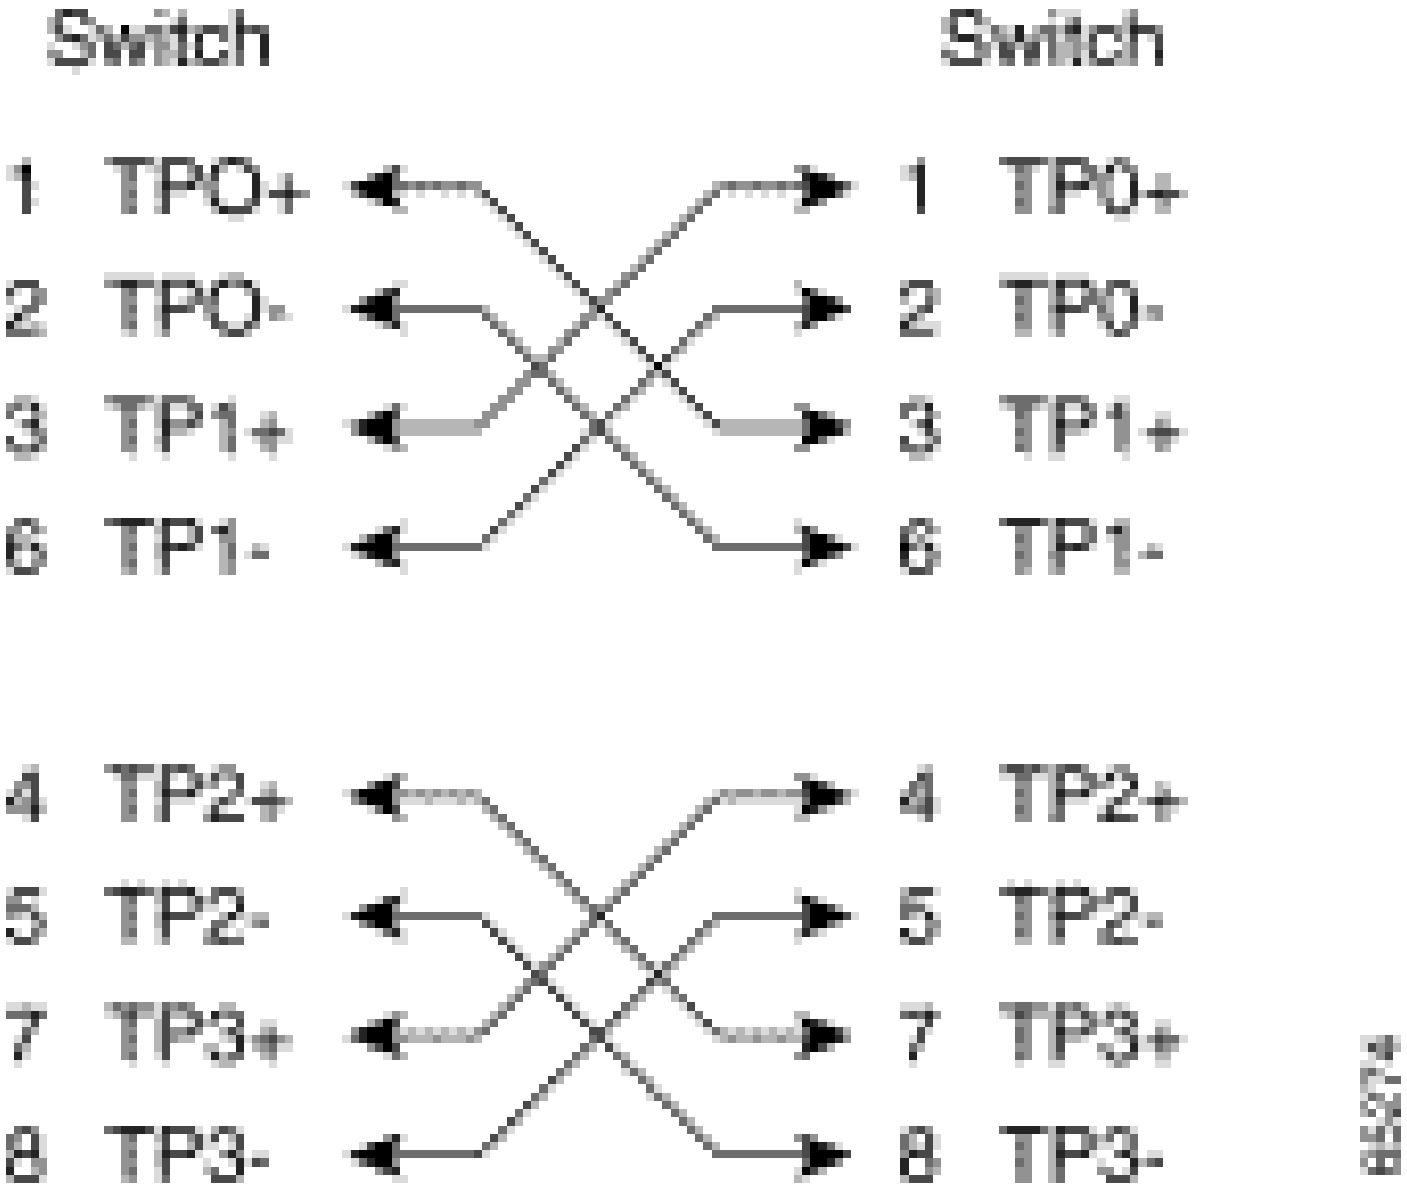
CAT5/CAT6 10/100/1000BaseT UTP crossover patch cables are available from most computer stores.

Note: Some Ethernet network devices (10/100BaseT hubs) have what is referred to as a media dependent interface (MDI) port. Activate an internal crossover function, and this type of port allows the device to connect to an Ethernet port on a switch that uses a straight-through patch cable. Turn the MDI switch on to perform this. When the MDI switch is in the out position, the port expects to be connected to an end-user device.

Four Twisted-Pair Crossover Cable Schematics for 10/100/1000 and 1000BASE-T GBIC Module Ports



Four Twisted-Pair Crossover Cable Schematics 1



Four Twisted-Pair Crossover Cable Schematics 2

CAT 5, 5e or 6 UTP crossover patch cables are available from most computer stores.

Fiber Cable Connection Guidelines

If you use an Ethernet port on the switch with a fiber interface to connect to another switch port, a router port, or other Layer 2 device, you need to reverse the connection on one of the devices. Rotate the connector one half turn or cross over the individual fiber connectors to reverse the connection. Think about each fiber as either fiber A or fiber B. If a straight-through connection is A-to-A and B-to-B, a crossover connection is A-to-B and B-to-A.

Appendix C - Explanation of Auto-MDIX and Supported Switch Platforms

Automatic medium-dependent interface crossover (Auto-MDIX) is a feature that allows the switch interface to detect the required cable connection type (straight-through or crossover) and automatically configure the connection appropriately. With Auto-MDIX enabled, you can use either a straight-through or crossover type cable to connect to the other device, and the interface automatically corrects for any incorrect cabling.

Appendix D - Explanation of Fields in the show interfaces Command

Output Fields for the show interface Command

Counters (in alphabetical order)	Issues and Common Causes that Increase Error Counters
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.
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.
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: how interfaces show interfaces counters errors. 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 receive 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.

Appendix E - Frequently Asked Questions

1. When must you use auto-negotiation?

   Cisco recommends that auto-negotiation be used when the devices involved are compliant with the 802.3u standard. Refer to Troubleshooting Cisco Catalyst Switches to NIC Compatibility Issues for more information on specific products. Auto-negotiation is very useful for ports where devices with different capabilities are connected and disconnected on a regular basis. An example is when an employee visits the office and brings their own laptop.

2. How can you configure an interface for auto-negotiation?

   Remove the hardcoded speed and duplex settings from the interface configuration. This resets both the speed and duplex mode to auto-negotiate. Or run the interface command **speed auto**.

3. How can you tell how your port is configured?

   Run the **show interface <interface >** status command. Look for the **a** prefix in the status fields. This indicates the port is configured for auto-negotiation. Examples are a-full and a-100. If the **a** prefix is not present, the port is manually configured for the parameters shown. Examples are full and 100. Run the **show run interface <interface>** command to view the configuration of the switch.

4. How can you tell what your interface is capable of?

   Run the **show interface capabilities** command or you can also run the **show interfaces <interface> status** command to view the speed/duplex settings.

5. Why does a port not detect the correct duplex mode when the link partner is not configured for auto-negotiation?

   The port does not detect it because there is no method available to perform this.

6. Why is it possible to have link show connected when the two ports have different duplex modes configured?

   It is possible because the electrical signals the ports use to determine if they are connected do not track the status of the duplex modes.

7. Does the**a**prefix on the duplex and speed status fields always mean the port has auto-negotiated behavior?

   No, it means that the port can perform auto-negotiation.

8. What does the`%CDP-4-DUPLEX_MISMATCH: duplex mismatch discovered`message mean?

   This means the CDP determines, via a configuration comparison dialogue, that a mismatch exists.

CDP does not attempt to resolve the mismatch.

Related Information

- Troubleshooting Cisco Catalyst Switches to NIC Compatibility Issues
- LAN Switching Technology Support
- Technical Support & Documentation - Cisco Systems