

Troubleshooting Cisco Catalyst Switches to NIC Compatibility Issues

Document ID: 17053

Contents

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

Background Information

- Purpose

Why Do Autonegotiation and Compatibility Issues Exist?

General Troubleshooting for 10/100/1000 Mbps NICs

- Autonegotiation Valid Configuration Table
- EtherChannel and Trunking Between Catalyst Switches and NICs
- Verifying Physical Connection and Link
- Verifying Switch Port Configuration

Maintaining Link (Link Up/Down Situations)

- Performance Notes
- Understanding Data Link Errors
- Sniffer Trace
- Teaming of Network Interface Cards

Additional Troubleshooting for 1000BASE-X NICs

Gigabit Autonegotiation (No Link to Connected Device)

- Verifying GBIC

Cisco Catalyst Switch Compatibility and Operation-Specific Issues

- Catalyst 8510 and 8540 CSR
- Catalyst 6000 and 6500 Switches
- Catalyst 5000 and 5500 Switches
- Catalyst 4000, 2948G, and 2980G Switches
- Catalyst 2950 and 3550 Switches

NIC Compatibility and Operation Issues

Appendix A: Information to Gather Before Creating a Service Request

Appendix B: Understanding How Autonegotiation Works

Related Information

Introduction

The purpose of this document is to cover common issues associated with network interface cards (NICs) that interoperate with Cisco Catalyst switches. Network issues, such as slow performance and connectivity problems, as well as Catalyst switch issues that deal with physical connectivity and data link errors, can be related to NIC issues.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

Conventions

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

Background Information

Purpose

This document discusses how to troubleshoot these issues:

- Autonegotiation
- Physical Connectivity
- Port Errors (Data Link Errors)
- Continuous Link Up/Down Situations
- Gigabit Port Configuration
- Common Catalyst Switch Software Issues
- Common NIC Issues and Resolutions

When you troubleshoot NIC issues with Catalyst switches, the first step is to verify that the issue is not related to a possible configuration issue with the Catalyst switch. For useful information that pertains to common connectivity issues with the configuration of the Catalyst switch, refer to these documents:

- This document addresses initial connectivity delays that occur when workstations connected to Catalyst switches are unable to log in to a network domain (Microsoft Windows NT or Novell), or are unable to obtain a Dynamic Host Configuration Protocol (DHCP) address, due to the Catalyst switch configuration. The first step in order to troubleshoot these scenarios is to confirm that the switch configuration is correct, as shown in [Using PortFast and Other Commands to Fix Workstation Startup Connectivity Delays](#).
- Excessive data link errors cause ports on some Catalyst switches to go into an `errdisabled` state. [Recovering From errDisable Port State on the CatOS Platforms](#) describes what the `errdisable` state is, explains how to recover from it, and provides two examples of recovery from this state.

Why Do Autonegotiation and Compatibility Issues Exist?

Autonegotiation issues can result from nonconforming implementation, hardware incapacities, or software defects. When NICs or vendor switches do not conform exactly to the IEEE specification 802.3u, problems can result. Hardware incompatibility and other issues can also exist as a result of vendor-specific advanced features, such as autopolarity or cable integrity, which are not described in IEEE 802.3u for 10/100 Mbps autonegotiation. Generally, if both the NIC and the switch adhere to IEEE 802.3u autonegotiation specifications and all additional features are disabled, autonegotiation must properly negotiate speed and duplex, and no operational issues exist.

General Troubleshooting for 10/100/1000 Mbps NICs

Autonegotiation Valid Configuration Table

Speed determination issues can result in no connectivity. However, issues with autonegotiation of duplex generally do not result in link establishment issues. Instead, autonegotiation issues mainly result in performance-related issues. The most common problems with NIC issues deal with speed and duplex configuration. Table 1 summarizes all possible settings of speed and duplex for FastEthernet NICs and switch ports.

Note: This section is only applicable for 10/100/1000 Mbps (1000BASE-T) NICs, and not 1000BASE-X NICs.

Table 1 Autonegotiation Valid Configuration

Configuration NIC (Speed/Duplex)	Configuration Switch (Speed/Duplex)	Resulting NIC Speed/Duplex	Resulting Catalyst Speed/Duplex	Comments
AUTO	AUTO	1000 Mbps, Full-duplex	1000 Mbps, Full-duplex	Assuming maximum capability of Catalyst switch, and NIC is 1000 Mbps, full-duplex.
1000 Mbps, Full-duplex	AUTO	1000 Mbps, Full-duplex	1000 Mbps, Full-duplex	Link is established, but the switch does not see any autonegotiation information from NIC. Since Catalyst switches support only full-duplex operation with 1000 Mbps, they default to full-duplex, and this happens only when operating at 1000 Mbps.
AUTO	1000 Mbps, Full-duplex	1000 Mbps, Full-duplex	1000 Mbps, Full-duplex	Assuming maximum capability of NIC is 1000 Mbps, full-duplex.

1000 Mbps, Full-duplex	1000 Mbps, Full-duplex	1000 Mbps, Full-duplex	1000 Mbps, Full-duplex	Correct Manual Configuration
100 Mbps, Full-duplex	1000 Mbps, Full-duplex	No Link	No Link	Neither side establishes link, due to speed mismatch
100 Mbps, Full-duplex	AUTO	100 Mbps, Full-duplex	100 Mbps, Half-duplex	Duplex Mismatch ¹
AUTO	100 Mbps, Full-duplex	100 Mbps, Half-duplex	100 Mbps, Full-duplex	Duplex Mismatch ¹
100 Mbps, Full-duplex	100 Mbps, Full-duplex	100 Mbps, Full-duplex	100 Mbps, Full-duplex	Correct Manual Configuration ²
100 Mbps, Half-duplex	AUTO	100 Mbps, Half-duplex	100 Mbps, Half-duplex	Link is established, but switch does not see any autonegotiation information from NIC and defaults to half-duplex when operating at 10/100 Mbps.
10 Mbps, Half-duplex	AUTO	10 Mbps, Half-duplex	10 Mbps, Half-duplex	Link is established, but switch does not see Fast Link Pulse (FLP) and defaults to 10 Mbps half-duplex.
10 Mbps, Half-duplex	100 Mbps, Half-duplex	No Link	No Link	Neither side establishes link, due to speed mismatch.
AUTO	100 Mbps, Half-duplex	100 Mbps, Half-duplex	100 Mbps, Half-duplex	Link is established, but NIC does not see any autonegotiation information and defaults to 100 Mbps, half-duplex.
AUTO	10 Mbps, Half-duplex	10 Mbps, Half-duplex	10 Mbps, Half-duplex	Link is established, but NIC does not see FLP and

				defaults to 10 Mbps, half-duplex.
--	--	--	--	-----------------------------------

¹ A duplex mismatch can result in performance issues, intermittent connectivity, and loss of communication. When you troubleshoot NIC issues, verify that the NIC and switch use a valid configuration.

² Some third-party NIC cards can fall back to half-duplex operation mode, even though both the switchport and NIC configuration are manually configured for 100 Mbps, full-duplex. This is because NIC autonegotiation link detection still operates when the NIC is manually configured. This causes duplex inconsistency between the switchport and the NIC. Symptoms include poor port performance and frame check sequence (FCS) errors that increment on the switchport. In order to troubleshoot this issue, try to manually configure the switchport to 100 Mbps, half-duplex. If this action resolves the connectivity problems, this NIC issue is the possible cause. Try to update to the latest drivers for your NIC, or contact your NIC card vendor for additional support.

Why Is It That the Speed and Duplex Cannot Be Hardcoded on Only One Link Partner?

As indicated in Table 1, a manual setup of the speed and duplex for full-duplex on one link partner results in a duplex mismatch. This happens when you disable autonegotiation on one link partner while the other link partner defaults to a half-duplex configuration. A duplex mismatch results in slow performance, intermittent connectivity, data link errors, and other issues. If the intent is not to use autonegotiation, both link partners must be manually configured for speed and duplex for full-duplex settings.

Recommended Port Configuration (Autonegotiation or Manual Configuration)

There are many opinions on the subject of autonegotiation. Previously, many engineers advised customers not to use autonegotiation with any switch-connected device. However, improvements in the interoperability of autonegotiation and the maturity of the technology has recently changed the view of autonegotiation and its use. In addition, performance issues due to duplex mismatches, caused by the manual setting of speed and duplex on only one link partner, are more common. Because of these recent issues, the use of autonegotiation is regarded as a valid practice.

EtherChannel and Trunking Between Catalyst Switches and NICs

EtherChannel can be configured dynamically with Port Aggregation Protocol (PAgP), and trunking can also be configured dynamically with Dynamic Trunking Protocol (DTP). Both PAgP and DTP are Cisco proprietary protocols and supported only on Catalyst switches. If you want to configure EtherChannel or trunking between Catalyst switches and NICs, it is recommended that you configure these features statically, as other vendor NICs can potentially not support PAgP and DTP. On Catalyst switches, configure the EtherChannel mode to `on` and trunking mode to `nonegotiate`, which disables the PAgP and DTP protocols. If you configure the switch port with `auto` or `desirable` mode, it is possible you can not be able to form the EtherChannel or trunk with NICs.

Verifying Physical Connection and Link

When you troubleshoot NIC issues, the first step is to verify physical connectivity. Visual inspection of the switch must show a LINK light indicator when connected to a link partner. In addition, the NIC can also have a LINK light indicator. The Command Line Interface (CLI) of the switch must be checked in order to verify physical connectivity. The port in question must show `connected` for Catalyst OS software and `line protocol up` for Cisco IOS® Software on the switch.

and that spanning tree PortFast is enabled.

Refer to Using PortFast and Other Commands to Fix Workstation Startup Connectivity Delays for more documentation with regard to this configuration change.

Maintaining Link (Link Up/Down Situations)

Under certain circumstances, interoperability issues between Cisco switches and various NICs can result in continuous or intermittent link up/down situations. These link up/down situations are usually a result of power management features or jitter tolerance issues associated with the NIC.

- For link up/down situations for CatOS, these messages appear and are normal for link up/down situations:

```
PAGP-5-PORTTOSPT: Port [dec]/[dec] joined bridge port [dec]/[chars]
PAGP-5-PORTFROMSPT: Port [dec]/[dec] left bridge port [dec]/[chars]
```

This is an example:

```
%PAGP-5-PORTFROMSTP:Port 3/3 left bridge port 3/3
%PAGP-5-PORTTOSTP:Port 3/3 joined bridge port 3/3
```

- For Cisco IOS Software–based switches, these messages appear for link up/down situations:

```
%LINK-3-UPDOWN: Interface interface, changed state to up
%LINK-3-UPDOWN: Interface interface, changed state to down
```

This is an example:

```
%LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to up
%LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to down
```

In order to resolve these issues, troubleshoot with these techniques:

- **Disable Windows 2000 and Windows Millennium Edition (ME) power management functions.**

Windows 2000 and Windows ME employ a power management capability that can disable the NIC. When the NIC is disabled for power management, it drops the link to the switch. If there is a concern about the link going up/down on NICs with the Windows 2000 or Windows ME operating systems, disable the power management feature as a first step in order to troubleshoot link up/down situations.

- **Disable the NIC power management functionality. Many NICs support their own power management capability.**

When you troubleshoot link up/down issues, disable this feature. For information on how to disable power management, refer to the NIC documentation.

- **Adjust switch jitter tolerance.**

Jitter tolerance, based on the IEEE 802.3u–1995, clause 25, must not exceed 1.4 nanoseconds. However, there are situations in which NICs that operate out-of-specification with respect to excessive jitter cause link up/down situations on Catalyst 6000 and 6500 10/100 ports. The workaround for this issue is to increase the jitter tolerance on the Catalyst 6000 and 6500 switches for 10/100 ports to 3.1 seconds. The **set port debounce *mod/port* enable** command enables the feature. As an ultimate solution, replace the out-of-specification NICs, instead of using the debounce option. This feature is first integrated into software version 5.3(5)CSX.

For the Catalyst 2900XL and 3500XL, the interface command **carrier-delay *time*** can be adjusted to

four seconds as a possible workaround for this same issue.

Refer to Fast Ethernet Consortium Physical Medium Dependent Test Suite [☐](#) for more information about jitter tolerance.

Performance Notes

Most performance issues are related to switch port configuration, duplex mismatches, link up/down situations, and data link errors. When you troubleshoot performance issues, review all previous sections of this document. After you review these sections, proceed to the next section, Understanding Data Link Errors. The final step in order to resolve any performance issue is to obtain a sniffer trace. A sniffer trace is very conclusive with regard to any specific performance problem because it details packet transfer.

Understanding Data Link Errors

Many performance issues with NICs can be related to data link errors. Excessive errors usually indicate a problem. When operating at a half-duplex setting, some data link errors such as FCS, alignment, runts, and collisions 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, caused when the frame is not completely copied to the wire, which results in fragmented frames.

When operating at full-duplex, FCS, cyclic redundancy checks (CRC), alignment errors, and runt counters are probably 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 in which the switch operates at full-duplex and the connected device operates at half-duplex, or the other way around. The result of a duplex mismatch is extremely slow performance, intermittent connectivity, and loss of connection. Other possible causes of data link errors at full-duplex are bad cables, a faulty switch port, or NIC software or hardware issues.

When you troubleshoot NIC performance issues, view the output of the **show port mod/port** command and the **show mac mod/port** command, and note the counter information.

Table 2 Explanation of CatOS show port Command Counters

Counter	Description
Alignment 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 CRC.
FCS	FCS error count is the number of frames that were transmitted or received with a bad checksum (CRC value) in the Ethernet frame. These frames are dropped and not propagated onto other ports.
Xmit-Err	This is an indication that the internal transmit buffer is full.
Rcv-Err	This is an indication that the receive buffer is full.
UnderSize	These are frames that are smaller than 64 bytes, which includes FCS, and have a good FCS value.

Single Collisions	Single collisions are the number of times the transmitting port had one collision before successfully transmitting the frame to the media.
Multiple Collisions	Multiple collisions are the number of times the transmitting port had more than one collision before successfully transmitting the frame to the media.
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 that 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 the other sends until after it puts the entire packet on the network. Late collisions are detected by the transmitter after the first time slot of the 64-byte transmit time occurs. They are only detected during transmissions of packets longer than 64 bytes. Its detection is exactly the same as it is for a normal collision; it just happens later than it does for a normal collision.
Excessive Collisions	Excessive collisions are the number of frames that are dropped after 16 attempts to send the packet resulted in 16 collisions.
Carrier Sense	Carrier sense occurs every time an Ethernet controller wants to send data and the counter is incremented when there is an error in the process.
Runts	These are frames smaller than 64 bytes with a bad FCS value.
Giants	These are frames that are greater than 1518 bytes and have a bad FCS value.

Table 3 Possible Causes for Incrementing CatOS Counters

Counter	Description
Alignment Errors	These are the result of collisions at half-duplex, duplex mismatch, bad hardware (NIC, cable, or port), or a connected device that generates frames that do not end with on an octet and have a bad FCS.
FCS	These are the result of collisions at half-duplex, duplex mismatch, bad hardware (NIC, cable, or port), or a connected device that generates frames with bad FCS.
Xmit-Err	This is an indication of excessive input rates of traffic. This is also an indication that the transmit buffer is full. The counter must only increment in situations in which the switch is unable to forward

	out the port at a desired rate. Situations such as excessive collisions and 10 Mb ports cause the transmit buffer to become full. If you increase speed and move the link partner to full-duplex, it minimizes this occurrence.
Rcv-Err	This is an indication of excessive output rates of traffic. This is also an indication that the receive buffer is full. This counter must be zero unless there is excessive traffic through the switch. In some switches, the Out-Lost counter has a direct correlation to the Rcv-Err.
UnderSize	This is an indication of a bad frame generated by the connected device.
Single Collisions	This is an indication of a half-duplex configuration.
Multiple Collisions	This is an indication of a half-duplex configuration.
Late Collisions	This is an indication of faulty hardware (NIC, cable, or switch port) or a duplex mismatch.
Excessive Collisions	This is an indication of overutilization of the switch port at half-duplex or duplex mismatch.
Carrier Sense	This is an indication of faulty hardware (NIC, cable, or switch port).
Runts	This is an indication of the result of collisions, duplex mismatch, IEEE 802.1Q (dot1q), or an Inter-Switch Link Protocol (ISL) configuration issue.
Giants	This is an indication of faulty hardware, dot1q, or an ISL configuration issue.

Table 4 Explanation of CatOS show mac Command Counters

Counter	Description
Rcv-Unicast	This is an indication of the number of unicast packets received.
Rcv-Multicast	This is an indication of the number of multicast packets received.
Rcv-Broadcast	This is an indication of the number of broadcast packets received.
Xmit-Unicast	This is an indication of the number of unicast packets transmitted.
Xmit-Multicast	This is an indication of the number of multicast packets transmitted.
Xmit-Broadcast	This is an indication of the number of broadcast packets transmitted.

Delay Exceeded	This is an indication of the number of frames discarded because of excessive delay in the switching process.
MTU–Exced	This is an indication that one of the devices on the port or segment transmits more than the allowed frame size.
In–Discard ²	Count of valid frames received which were discarded or filtered by the forwarding process.
Lrn–Discard ²	Packets which are forwarded and must not be forwarded.
In–Lost	Packets which can not be received because the input buffers are full.
Out–Lost	Packets which can not be transmitted because the output buffers are full.

² In–Discard and Lrn–Discard do not exist on all Catalyst platforms.

Table 5 Possible Causes for Incrementing CatOS Counters

Counter	Possible Cause
Delay Exced	Severe problem with the switch. Create a service request with Cisco Technical Support.
MTU–Exced	Verify ISL and dot1q configurations. Verify another switch or router does not inject frame over maximum transmission unit (MTU) into the switch network.
Lrn–Discard ²	Increments when the switch receives traffic on a trunk for a specific VLAN while the switch does not have any other ports on that VLAN. Counter also increments when the destination address of the packet is learned on the port on which the packet is received.
Lrn–Discard ²	This counter must remain at zero. If counter increments, create a service request with Cisco Technical Support.
In–Lost	Excessive input rate of traffic.
Out–Lost	Excessive output rate of traffic. Increments in this counter are more likely to occur when connected to low–speed devices. The first step to troubleshoot Out–Lost increments is to verify the link partner runs 100 Mbps, full–duplex without any errors.

² In–Discard and Lrn–Discard do not exist on all Catalyst platforms.

Additional counter information can be viewed with the command **show counters mod/port** . The command must be issued for a single port at a time. Refer to this document for information on the counter displayed:

- **show counters** Command Documentation

For additional information about the Cisco IOS Software **show interfaces** command counters, refer to:

- **show interfaces** Command Documentation

Sniffer Trace

Sniffer trace analysis can be very useful when you troubleshoot switch and NIC performance or connectivity issues when the issues persist after all other sections of this document are reviewed. Sniffer trace analysis reveals every packet on the wire and pinpoints the exact problem. It can be important to obtain several sniffer traces from different ports on different switches. Generally, it is very useful to monitor or *span* ports rather than spanning VLANs when you troubleshoot switch and NIC performance and connectivity issues.

Refer to Catalyst Switched Port Analyzer (SPAN) Configuration Example for more information on the use of the Switched Port Analyzer (SPAN) feature required to obtain sniffer traces.

Teaming of Network Interface Cards

Teaming of Network Interface Cards, or NIC Teaming, can cause instability in the networks. Such setups can introduce disruptions to the Spanning tree and can make it undergo frequent recomputations. If intermittent loss of connectivity to NIC teamed servers occurs for devices or hosts in the same VLAN, try to disable NIC teaming. If the connectivity stabilizes, refer to the NIC vendor documentation in order to tune the NIC teaming configuration.

Use one of these methods in order to implement NIC teaming:

- **Server Virtual Address (SVA):** The SVA is used when you want other devices in the network to see the teamed NICs as one physical device with one MAC address. When you use this setup, you must have one of the NICs in a standby state, and the other in active state. Otherwise, you would experience duplicate MAC addresses sent around the network from the SVA.
- **Separate NIC MAC Addresses:** In this setup, you can use both of your NIC cards that run separate MAC addresses. In this mode, both NICs appear from a network perspective to be two separate physical devices. You can configure the Fault Tolerant Mode with Load Balancing option in order to avoid the problem of duplicate MAC addresses on the network.

Additional Troubleshooting for 1000BASE-X NICs

Gigabit Autonegotiation (No Link to Connected Device)

Gigabit Ethernet has an autonegotiation procedure that is more extensive than what is used for 10/100 Mbps Ethernet (Gigabit autonegotiation specification IEEE 802.3z-1998). The Gigabit autonegotiation 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 does not connect.

If either device does not support Gigabit autonegotiation, disable Gigabit autonegotiation in order to force the link up. The default configuration of all Cisco switches is autonegotiation-enabled. If you disable autonegotiation, it hides link drops and other physical layer problems. Only disable autonegotiation to end-devices, such as older Gigabit NICs that do not support Gigabit autonegotiation. Do not disable autonegotiation between switches unless absolutely required, as physical layer problems can go undetected and result in spanning tree loops. Rather than disable autonegotiation, you can contact the vendor for a software or hardware upgrade for IEEE 802.3z Gigabit autonegotiation support.

Table 6 Gigabit Autonegotiation Configuration Table

Autonegotiation Setting	NIC Gigabit Autonegotiation Setting	Switch Port Gigabit Link	Alternate Switch Link/NIC Link
Enabled	Enabled	Up	Up
Disabled	Disabled	Up	Up
Enabled	Disabled	Down	Up
Disabled	Enabled	Up	Down

Issue these commands for Gigabit autonegotiation configuration:

- CatOS command:

```
set port negotiation mod/port enable | disable
```

- Cisco IOS Software command:

```
negotiation auto no negotiation auto
```

Verifying GBIC

When you troubleshoot link issues on Gigabit Ethernet, it is also important to verify the use of the correct Gigabit Interface Converter (GBIC) adapter with the correct cable distance. Refer to Gigabit Interface Converter Installation Note for information on distances and cable specifications required for different versions of GBIC adapters.

Cisco Catalyst Switch Compatibility and Operation–Specific Issues

These sections deal with specific Cisco Catalyst switch issues that can affect performance, compatibility, and interoperation of certain NICs.

Catalyst 8510 and 8540 CSR

In the campus switch router (CSR) Cisco IOS Software Release 12.0(5)W5(13), autonegotiation for speed and duplex are enabled by default. In earlier releases, autonegotiation is not supported by default. As a result, each connected interface must be set to run in full–duplex mode. If you upgrade to Cisco IOS Software Release 12.0(5)W5(13) with a router that runs in manually full–duplex mode, you experience performance issues. The symptoms include high collision rate, reduced throughput, and more packet drops. This is because the Catalyst 8500 waits to autonegotiate with the connected device. Now that the connected device is forced to run in full–duplex mode, it does not participate in autonegotiation. Based on the specification, this causes the Catalyst 8500 interface to settle in half–duplex mode, which causes a mismatch between the device and the Catalyst 8500 at the interface level. Catalyst 8500 interface defaults to half–duplex mode when the peer is incapable of negotiation.

Catalyst 6000 and 6500 Switches

This table describes Cisco bug IDs found on the Catalyst 6000 and 6500 switches.

Table 7

Cisco Bug ID	Resolved In	Description
CSCdm48887 🔗 (registered customers only)	5.2.3, 5.3.1a	When a port goes into the <code>errdisable</code> state on a Catalyst 6000 or 6500, the switch inadvertently learns the MAC addresses of other devices to be learned from the <code>errdisable</code> port. The <code>errdisable</code> state is caused by excessive late collisions on a port. Since all traffic on that VLAN is forwarded incorrectly out of the incorrect port, there is a loss of connectivity. This <code>errdisable</code> state can appear as a result of a duplex mismatch or faulty NIC.
CSCdm80035 🔗 (registered customers only)	5.2.3, 5.3.1a	When a Gigabit connection is reset on a Catalyst 6000 or 6500, the connection can potentially not reconnect. A symptom of this issue can be that a Gigabit NIC does not connect after reset or disconnect.
CSCdm88013 🔗 (registered customers only)	5.2.3, 5.3.1a	Occasionally, host NICs connected to WS-X6248-TEL or WS-X6248-RJ-45 modules can incorrectly revert to half-duplex after an autonegotiation failure.

For detailed information on Cisco bug IDs, refer to the Bug Toolkit [🔗](#) (registered customers only) .

Refer to Cisco Catalyst 6500 Series Switches Release Notes for more Catalyst 6000 and 6500 documented bug fixes.

Catalyst 5000 and 5500 Switches

This table lists several known issues found on the Catalyst 5000 and 5500 switches.

Table 8

Cisco Bug ID	Resolved In	Description
CSCdt28585 🔗 (registered customers only)	5.5(6)	Directly connected hosts (PCs, routers, and servers) can indicate a <code>connected</code> state in the output of a <code>show port</code> command, but do not forward <code>Xmit-Broadcast</code> frames. This causes connectivity issues that are only fixed when you issue

		the set port disable <i>mod/port</i> and the set port enable <i>mod/port</i> commands.
CSCdr50629 ↗ (registered customers only)	5.5(3)	Ports on WS–X5225R, WS–X5234, and WS–X5201R modules do not transmit unicast frames after a scheduled test of the packet buffers. The workaround is to disable the packet buffer test.
CSCdr03818 ↗ (registered customers only)	4.5(7), 5.4(2)	The WS–X5225R and WS–X5234 modules fail to negotiate the duplex mode correctly after a system reset or power cycle with Sun workstations Ultra 5 .
CSCdm51653 ↗ (registered customers only)	4.5(3), 5.1(2a)	Autonegotiation between Sun 10/100 NICs and certain Catalyst 5000 family modules (such as the WS–X5225R) can result in speed or duplex mismatches under certain conditions. The problem typically occurs after the module is reset or the switch port is disabled and reenabled. The workaround is to disconnect and reconnect the cable that connects the workstation to the switch port.
CSCdk32984 ↗ (registered customers only)	4.2(2)	The 48–port, 10BASE–T Ethernet module (WS–X5012) incorrectly drops valid frames with dribble bits (extra bits added to frames by some end stations and transceivers).
CSCdj82035 ↗ (registered customers only)	3.2(2), 4.1(3)	Under heavy traffic conditions, ports 1 through 24 (or ports 25 through 48) of the 48–port, 10BASE–T Ethernet module (WS–X5012) can stop transmitting frames.

For detailed information on Cisco bug IDs, refer to the Bug Toolkit [↗](#) (registered customers only) .

Refer to Cisco Catalyst 5000 Series Switches Release Notes for more Catalyst 5000 and 5500 documented bug fixes.

Catalyst 4000, 2948G, and 2980G Switches

This table lists several known issues found on the Catalyst 4000, 2948G, and 2980G switches.

Table 9

Bug ID	Resolved In	Description
CSCds38973 ↗ (registered customers only)	4.5(8), 5.2(7), 5.5(2)	The Catalyst 2948G and Catalyst 4000 switches can experience issues with total or intermittent loss of connectivity. The frequency of these issues can vary from once a day to once a month. This issue can reoccur even after the switch is power cycled. The purpose of this Cisco bug ID is to combine several software reworks, fixes that resolve and reduce the loss of connectivity issues, and additional troubleshooting checks in software.
CSCdr37645 ↗ (registered customers only)	4.5(8), 5.5(2), 6.1(1)	An invalid packet with a length of less than 64 bytes received on a 10/100 port causes both the Runts and FCS–Error counters to increment on the port. To determine the actual number of FCS–Errors on valid length packets received on the port, subtract the value of the port Runts counter from the value of the port FCS–Error counter.
CSCdm38405 ↗ (registered customers only)	5.1(1)	Some Sun Gigabit Ethernet NICs do not reliably autonegotiate flow control with certain ports on the Catalyst 4000 series oversubscribed Gigabit Ethernet modules. The 18–port server switching 1000BASE–X (GBIC) Gigabit Ethernet module (WS–X4418) is affected.
CSCdm51653 ↗ (registered customers only)	4.5(3), 5.1(2a)	In some cases, autonegotiation with some Sun NICs can result in a nonoptimal configuration (such as 10 Mbps, half–duplex instead of 100 Mbps, full–duplex).
CSCdt80707 ↗ (registered customers only)	5.5.7, 6.1.3, 6.2.1	On Catalyst 4006 (with a Supervisor Engine II), switch ports in the same VLAN can lose connectivity with one another. The loss of connectivity results in a VLAN that appears to be partitioned into several isolated segments. A host can be able to ping one set of devices in its VLAN, while it cannot ping

		another set of devices in the same VLAN. This loss of connectivity is independent of the slot in which a linecard is installed; that is, the same set of ports on a given linecard are affected regardless of the slot in which the linecard is installed. The workaround is to reset the switch.
CSCds89148 ↗ (registered customers only)	5.5.6, 6.2.1	Xmit-Err counter increments on unconnected ports for unexplained reasons. The bug also resolves an issue of high-CPU utilization conditions that can be caused by unconnected ports that are configured as host ports.

For detailed information on Cisco bug IDs, refer to the Bug Toolkit [↗](#) (registered customers only) .

Refer to Catalyst 4500 Series Release Notes for more Catalyst 2948G, 2980G, and 4000 documented bug fixes.

Catalyst 2950 and 3550 Switches

This table lists several known issues found on the Catalyst 2950 and 3550 switches.

Table 10

Cisco Bug ID	Resolved In	Description
CSCdz44520 ↗ (registered customers only)	12.1(13)EA1	The Catalyst 3550-24PWR inline power interfaces do not link up to certain 10/100/1000 interfaces configured as auto/auto. To connect the Catalyst 3550-24PWR inline power interfaces to a 10/100/1000 interface configured as auto/auto on a Catalyst 3550-12G or 3550-12T does not work.
CSCdz32789 ↗ (registered customers only)	12.1(13)EA1	The link to certain NICs fails to come up when the switch port is hardcoded to 100 Mbps, full-duplex or 100 Mbps, half-duplex.
CSCdy72718 ↗ (registered customers only)	12.1(13)EA1	The switch port does not receive packets if the port is hardcoded to a speed of 100,
CSCea36322	12.1(14)EA1	however it transmits properly. If a 10/100 port on the Catalyst

☞ (registered customers only)		3550–24PWR switch is connected to a Gigabit Ethernet NIC, with the speed/duplex setting set to auto, and the port speed changed from 100 Mbps to 10 Mbps, or from 10 Mbps to 100 Mbps, the link between the port and the NIC can potentially not come up.
---	--	---

For detailed information on Cisco bug IDs, refer to the Bug Toolkit [☞](#) (registered customers only) .

NIC Compatibility and Operation Issues

Disclaimer: Use this table as a guide in order to troubleshoot NIC issues. Consult the NIC vendor for verification and proper resolution of the issue.

Table 11

NIC Model/Manufacture	Symptom	Description	Resolution
Apple Macintosh G3	Intermittently loses network services when the built-in Ethernet interface is used.	Driver versions earlier than 2.04 can experience this issue. Contact vendor technical support for more information.	Upgrade to driver version 2.04 or later.
Apple Macintosh, Power Macintosh G3, and Powerbook G3	Unable to set speed and duplex of built-in Ethernet interface manually.	Apple Speed/Duplex tool is required to manually set speed/duplex of Ethernet interface.	Download Apple Speed/Duplex tool from Apple support website.
Apple Macintosh OS with Open Transport 2.5.1 and 2.5.2	Unable to obtain DHCP address from DHCP server.	Upon bootup, Macintosh can fail to obtain IP address from DHCP server.	See Apple technical information library article 25049.
Apple Macintosh Built-In Ethernet	Unable to determine hardware MAC address.	In order to troubleshoot network connectivity issues, host MAC address can be required.	Contact vendor technical support.
Apple Macintosh Performance Issues and NuBus	Built-in Ethernet interfaces outperform NuBus	Concern of maximum data transfer rate possible with the built-in Ethernet.	See Apple technical information library article 12000.
Apple Powerbook G3/G4 with Internal NIC	Ethernet cards. Slow performance when large file transfers are performed.	Some NICs can operate out-of-specification, as published in IEEE 802.3. Some	Use of external or PC card. Contact Apple technical support.

		Catalysts are more tolerant of out-of-spec NICs and do not notice any performance degradation.	
Various Apple G3/G4 Laptops and Workstations with Internal NICs	Slow performance.	Notable slow performance.	Upgrade to latest NIC driver and load Duplexer utility. Verify Set network card settings. TID1084 at Asanté support website.
AsantéFast 10/100 PCI Adapter	Slow login or fails to log in to server.		See technical document TID1084 at Asanté support website.
AsantéFast 10/100 PCI Adapter	Numerous CRC and FCS errors reported on switch when connected to Power Macintosh 9500.		See technical document TID1109 at Asanté support website.
AsantéFast 10/100 PCI Adapter	Slow network throughput after Macintosh OS 8.5 or 8.6 upgrade.		See technical document TID1976 at Asanté support website.
Asanté GigaNIC 1064SX PCI Card-Macintosh	Network performance fluctuates.	When the energy saver mode is active under OS 8.6, the network speed becomes quite slow as soon as the monitor goes dim.	Turn off the energy saver mode in the control panel. The network speed remains constant. See technical document TID2095 at Asanté support website.
Asanté GigaNIC 1064SX PCI Card-Macintosh	Slow performance with AppleShare IP servers and PCI Ethernet cards.	Customers report AppleShare IP servers slow down and eventually crash over time. This occurs with both built-in Ethernet and various PCI cards.	See technical document TID2227 at Asanté support website.
3Com 3C574/575 PCMCIA 10/100	Extremely slow performance when operating at 10 MB.	The 3C574/3C575 experiences slow performance when connecting at 10 MB with Catalyst 2948G, 2980G, 4000, 5000, and 6000 switches. This issue is caused by the NIC doing autopolarity upon linkup.	Upgrade to the latest NIC card driver and disable autopolarity.
3Com 3C595	FCS or alignment errors recorded on	FCS or alignment errors when the	Upgrade to the latest NIC driver and disable Bus

	the switch. Noted slower performance. When the 3C595 adapter is used in 100 MB, half-duplex. This problem usually only represents one percent to two percent of the total traffic.	3C595 adapter is used in 100 MB, half-duplex. This problem usually represents one percent to two percent of the total traffic.	Master. These steps reduce FCS and alignment errors.
3Com 3C905/3C905B	Intermittent DHCP issues.	Despite the proper configuration of Catalyst switch ports, workstations still experience some intermittent DHCP issues.	Upgrade to driver version 4.01b or later, which resolves DHCP issues.
3Com 3C905/3C905B	Unable to log in to Novell Internetwork Packet Exchange (IPX) network.	Despite the proper configuration of Catalyst switch ports, workstations still experience intermittent Novell IPX login issues.	Upgrade to driver 4.01b or later, which resolves IPX autoframe-type issues. Alternatively, manually configure workstations for IPX
3Com 3C905B	Slow performance when large files are received.	Notable slow performance when large files are received. Problem only occurs with standard Microsoft NT 4.0, regardless of service pack.	frame type. Download latest driver from 3Com technical support.
3Com 3C905C	Layer 2 (L2) errors reported on switch port (FCS, alignment, CRC, and runts), and slow performance on high-speed workstations.	Under normal conditions, a Catalyst that reports numerous L2 (physical) errors on ports connected to 3C905C NIC adapters.	Load latest driver and diagnostic tools available from 3Com. Test performance back-to-back between two PCs, and note errors on diagnostic tools. Errors reported, such as transmit under-run and receive over-run, cause physical layers to be reported by the switch and minor performance issues. For more details, refer to Cisco bug ID CSCdt68700 🔗 (registered customers only) .

3C905CX-TX-M	The link fails to come up when the switch port is hardcoded to 100 Mbps, full-duplex or 100 Mbps, half-duplex, and the NIC is set to autonegotiate.	Refer to Cisco bug ID CSCdz32789 ↗ (registered customers only) .	Upgrade to driver version 5.4, and in the advanced NIC properties, set LnkChk to enable. Refer to 3Com technical support reference 1.0.33921641.2241835.
3Com 3C980	Data corruption with Novell.		
3Com	3C985/3C985B	Novell 5.0 issues	Refer to 3Com technical support reference 1.0.16744826.2027011.
3Com 3C985/3C985B	Clients are unable to log in or browse server, but pings work correctly.		Refer to 3Com technical support reference 2.0.4428387.2305072.
3Com 3C985/3C985B	Packets larger than Ethernet MTU (1518 bytes) are generated. These packets are noted as giants on Catalyst switches.		Contact 3Com technical support.
3Com 3C905C or 3C920 integrated NIC on Dell Dimension XPS	Network connectivity is dropped every 2> minutes or the network card must be reinitialized several times to gain network connectivity.	A 3C905C or 3C920 integrated NIC on Dell Dimension XPS can experience network connectivity issues when Windows 2000 is run due to a power management issue.	Disable all power management functions. Contact Dell for details on how to disable power management or for more details about this issue. For more documentation, refer to 3Com technical support reference 2.0.47464140.2853794.
Compaq Netflex-3 Model NIC Adapters	Slow performance.	Autonegotiation can fail to Catalyst 5000 and 5500 switches.	This problem is resolved in software release 4.5(1) and later for the Catalyst 5000 and 5500 switches. For more details, refer to Cisco bug ID CSCdk87853 ↗ (registered customers only) .
Dell Optiplex GX200	The link flaps when it connects to a Dell Optiplex GX200 PC (Intel Pro 10/100). The NIC works	For more details, refer to Cisco bug ID CSCdz60677 ↗ (registered customers only) .	Upgrade to the latest drivers from Dell.

	properly when the PC is turned off, but when it is turned back on, the flap occurs.		
Dell Precision 420/530/620	When connecting to the Catalyst 2950 switch with an autonegotiation link, it continuously flaps and the autonegotiation link fails. The motherboards were manufactured with an Intel Chipset that are incompatible with Cisco switches and Netgear hubs.	Manufactured between May 21 and August 1, 2001.	Contact Dell technical support and Cisco Technical Support for more details.
Broadcom NetXtreme 57xx Gigabit Integrated Controller	Link comes up only with autonegotiation for speed and duplex.	NIC card drivers comes bundled with some administrative software, which affects the NIC card when the speed/duplex values are hard-coded. Release Date: 6/17/2005 Version: v7.1.0 , A04 Download Type: Application	Uninstall the administrative programs that were originally installed along with driver files.
IBM 10/100 EtherJet CardBus Adapter	Extremely slow performance when operating at 10 Mbps.	Certain 10/100 switches implement automatic correction for polarity-reversed cables that are not completely compatible with the same correction provided by the IBM 10/100 EtherJet CardBus Adapter. If the network speed is forced to 10 Mbps, severe throughput problems can be experienced.	In order to resolve this problem, a new Auto Polarity keyword is added in the adapter advanced properties. If needed, the default setting of ON, which means that the card compensates for reversed
IBM ThinClient Workstations	Link flaps continuously after extended operation.	Workstations earlier than Service Pack 3.0 bounce the link on the switch after	cables, can be set to OFF. Upgrade IBM ThinClient in order to disable to Service Pack 3.0 polarity correction. This restores normal throughput.

		continuous use when attached to Catalyst 2948G or 4000 switches on software version 6.x and later.	
Intel Pro/100	Consistent link up/down connections to Catalyst switches.	May be caused by power management. Contact Intel technical support for further information.	<ol style="list-style-type: none"> 1. Choose Control Panel > System > Hardware > Device Manager. 2. Choose Network Adapters > Intel Pro 100 +. 3. At the Power Management tab, uncheck Allow the computer to turn off this device.
Intel Pro/1000 T Gigabit Copper NIC	When an Intel Pro/1000 T NIC is connected to a Catalyst switch, you can see poor network connections or excessive numbers of dropped packets. The interoperability issue arises when a module with a ten bit interface (TBI) transmits an odd byte packet to a receiver with a Gigabit media-independent interface (GMII).	The interoperability issue arises from the implementation of Carrier Extension. Carrier Extension is detailed in subsection 35.2.3.5 in the IEEE 802.3 specification. Carrier Extension can	Contact Intel technical support for the latest driver.
Sun Microsystems QFE Card	Unable to manually set speed and duplex correctly.	Manually setting the last byte of a packet, so the packet is aligned on an even-numbered boundary.	Contact vendor technical support to obtain the latest driver to resolve the issue.
Sun Microsystems v1.1 Gigabit Cards	Unable to establish link.	V1.1 can potentially not establish link to switch.	Contact vendor technical support or v2.0 Gigabit Card.
Xircom CreditCard Ethernet 10/100 CE3B-100	Not negotiating or operating correctly at 100 Mbps, full-duplex.	Full-duplex operation is only supported at 10 Mbps. Full-duplex is not supported at	Do not operate this NIC at 100 Mbps, full-duplex.

		<p>100 Mbps. The LineMode keyword has no effect on performance at 100 Mbps. If the LineSpeed keyword is set to 100 Mbps and the LineMode keyword is set to full-duplex, the LineMode keyword is ignored. Full-duplex at 10 Mbps is only available when the adapter is connected to a full-duplex-capable switch or hub.</p>	
<p>Xircom CreditCard Ethernet 10/100 CE3B-100</p>	<p>Not negotiating 10 Mbps full-duplex.</p>	<p>The CE3 and, in some cases, the CE3B are not capable of negotiating to 10 Mbps, full-duplex mode.</p>	<p>On these adapters, in order to operate in full-duplex mode, the LineSpeed keyword must be set to 10 Mbps, and the LineMode keyword must be set to full-duplex. The cable type keyword can be set to Auto Detect or 10BASE-T/100BaseTX. The related port on the attached hub or switch must also be set to 10Mbps, full-duplex.</p>
<p>Xircom RealPort2 CardBus Ethernet 10/100 Adapter (R2BE/RBE/CBE) Models</p>	<p>Extremely slow performance when operating at 10 Mbps.</p>	<p>Certain 10/100 switches implement automatic correction for polarity-reversed cables that are not completely compatible with the same correction provided by the CBE/RBE. If the network speed is forced to 10 Mbps, severe throughput problems can be experienced.</p>	<p>In order to resolve this problem, a new Auto Polarity keyword is added in the adapter advanced properties in driver version 3.01. If needed, the default setting of ON, which means that the card compensates for reversed cables, can be set to OFF to disable polarity</p>
<p>Xircom RealPort2 CardBus Ethernet 10/100 Adapter (R2BE/RBE/CBE)</p>	<p>Initial network connections can fail. DHCP can obtain an IP</p>	<p>Initialization delay. Certain switches and routers are unable to immediately forward</p>	<p>correction. This restores A new keyword, Initialization Delay, is added to the adapter advanced properties that</p>

Models	address, and Windows NT login and Novell IPX can fail.	network traffic when a network adapter first establishes link to one of its ports because of initialization delays. This problem is most commonly seen when the network adapter is connected directly to ports on the switch. The adapter, by default, when used under some operating systems, has almost no delay between link and the initial network request.	prevents forwarding of network requests for a user-selectable period of time. Delays can be added ranging from 1 0 seconds. In most cases, if you add a delay in the 1> second range, it resolves the problem.
Xircom RealPort2 CardBus Ethernet 10/100 Adapter (R2BE/RBE/CBE) Models	Not able to connect to the network or get an IP address from the DHCP server when connected to a port replicator or docking station.	Possible Basic Input/Output System (BIOS) or driver update required. Contact vendor support for more information.	If you attempt to use a CBE/CBE2/RBE in a port replicator or docking station, with Windows 95, and have problems, confirm that your laptop has the latest BIOS and that the latest manufacturer patches and utility software are installed.
Xircom XE2000 PCMCIA NIC	Does not autonegotiate to 100 Mbps, full-duplex.	NIC only autonegotiates to 100 Mbps, half-duplex.	Known limitation of XE2000 NIC. Refer to the XE2000 release notes.
PROXIM TSUNAMI 5054-R Wireless Bridge	Not negotiating properly with Cisco catalyst 4510R-E.	Ports of Catalyst 4510R-E and PROXIM TSUNAMI 5054-R Wireless Bridge negotiation fails and the success rate is intermittent.	PROXIM TSUNAMI 5054-R is incompatible with Catalyst 4510R-E.

Appendix A: Information to Gather Before Creating a Service Request

If the troubleshooting procedure outlined in this document does not resolve your issue, you need to create a service request with Cisco Technical Support. Before you create a service request, gather this information:

1. Identify the specific problem with NIC-to-switch interoperability.

For example, is the problem only with DHCP, Novell IPX, login, or performance?

2. Issue the **show tech-support** command from all affected Cisco devices, if applicable; or, issue the **show module** , **show config** , **show version** , or the **show port** commands.
3. Know the make and model of the NIC.
4. Know the operating system and the NIC driver version.
5. Verify the consistency of the problem.

For example, does the problem occur across multiple Catalyst switches?

Appendix B: Understanding How Autonegotiation Works

Autonegotiation uses a modified version of the link integrity test that is used for 10BASE-T devices to negotiate speed and exchange other autonegotiation parameters. The original 10BASE-T link integrity test is referred to as Normal Link Pulse (NLP). The modified version of the link integrity test for 10/100 Mbps autonegotiation is referred to as FLP. The 10BASE-T devices expect a burst pulse every 16 (+/- 8) milliseconds (msec) as part of the link integrity test. FLP for 10/100 Mbps autonegotiation sends these bursts every 16 (+/- 8) msec with the additional pulses every 62.5 (+/- 7) microseconds. The pulses within the burst sequence generate code words that are used for compatibility exchanges between link partners. This process of FLP used in autonegotiation maintains backward compatibility with existing 10BASE-T connections, with the pulse burst every 16 (+/- 8) msec to comply with the link integrity test for normal 10BASE-T hardware. If a device sends FLP and only receives NLP, the hardware immediately ceases transmission of the FLP and enables the standard 10BASE-T hardware to continue 10BASE-T operation.

This table describes the possible programmable options of the control register for a FastEthernet interface. These options determine how the FastEthernet interface functions when connected to a link partner. The 0 in the Bits column refers to the programmable register address, and the decimal number after the 0 refers to the bit placement within the 16-bit register.

Table 12 Physical Interface (PHY) Control Register Programmable Options

Bits	Name	Description
0.15	Reset	1 = PHY reset 0 = normal mode
0.14	Loopback	1 = loopback mode switched on 0 = loopback mode switched off
0.13	Rate Selection (least-significant bit [LSB])	0.6 0.13 1 1 reserved 1 0 1000 Mbps 0 1 100 Mbps 0 0 10 Mbps
0.12	Autonegotiation Enable	1 = autonegotiation enabled 0 = autonegotiation disabled
0.11	Power Down	1 = power down 0 = normal down
0.10	Isolated	1 = PHY electrically isolated from media-independent interface (MII) 0 = normal mode
0.9	Restart Autonegotiation	1 = restart the autonegotiation process 0 = normal mode
0.8	Duplex Mode	1 = full-duplex 0 = half-duplex
0.7	Collision Test	1 = collision (COL) signal test active 0 = COL signal test switched off

0.6	Rate Selection (most-significant bit [MSB])	See bit 0.13
-----	---	--------------

The register bits relevant to this document include 0.13, 0.12, 0.8, and 0.6. The other register bits are documented in the IEEE 802.3u specification. Based on IEEE 802.3u, in order to manually set the rate (speed), the autonegotiation bit, 0.12, must be set to a value of 0. As a result, autonegotiation must be disabled in order to manually set the speed and duplex. If the autonegotiation bit 0.12 is set to a value of 1, bits 0.13 and 0.8 have no significance, and the link uses autonegotiation to determine the speed and duplex. When autonegotiation is disabled, the default value for duplex is half-duplex, unless the 0.8 is programmed to 1, which represents full-duplex.

Based on IEEE 802.3u, it is not possible to manually configure one link partner for 100 Mbps, full-duplex and still autonegotiate to full-duplex with the other link partner. If you attempt to configure one link partner for 100 Mbps, full-duplex and the other link partner for autonegotiation, it results in a duplex mismatch. This is because one link partner autonegotiates and does not see any autonegotiation parameters from the other link partner and defaults to half-duplex.

As described in Appendix B: Understanding How Autonegotiation Works, pulses within the FLP are used to derive code words that exchange link partner capabilities. The first code word exchanged is referred to as the base page. It informs each link partner of the message type, IEEE 802.3 or IEEE 802.9a, and a technology ability field. This technology ability field is encoded to exchange the maximum operational speed and duplex of each link partner.

Related Information

- [Configuring and Troubleshooting Ethernet 10/100/1000Mb Half/Full Duplex Auto-Negotiation](#)
- [LAN Switches Product Support](#)
- [LAN Switching Technology Support](#)
- [Technical Support & Documentation – Cisco Systems](#)

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2014 – 2015 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Oct 27, 2009

Document ID: 17053
