Troubleshooting Tips

Troubleshooting Methodology Facts

Good troubleshooting is a process that combines knowledge, experience, and intuition. As you practice service and support in a work environment, you will add to your experience and develop intuition that will help you to quickly solve a variety of problems.

Regardless of your current troubleshooting abilities, you will benefit from following a systematic approach to problem-solving. The following process has proven effective in a variety of situations:

1. Identify the symptoms and potential causes. Ask the user to describe the problem, check for error messages, and recreate the problem. Resist the urge to start fixing things at this point.
2. Identify the affected area and determine how large the problem is. For example, fixes for one client workstation would likely be very different than fixes for an entire network segment.
3. Establish what has changed. Most often, problems are caused by new hardware or software or changes to the configuration. If necessary, ask questions to discover what might have changed that could have caused the problem.
4. Review the list of potential causes and select the most probable cause. Look for common errors or solutions that can be tried quickly.
5. Escalate the problem if it is beyond your ability to fix or your scope of management. For example, the problem might be in a router configuration that you are not authorized to correct. When forwarding the problem on to someone else, be sure to describe the nature of the problem, the actions you have already taken, and the symptoms that lead you to believe the problem is outside of your area of responsibility.
6. Create an action plan and account for side effects of the proposed plan. Your plan might include purchases for hardware or equipment that need approval before proceeding. In addition, your plan might involve taking some services offline for a period of time. Identifying the affects ahead of time helps you put measures into place to eliminate or reduce any potential negative consequences.
7. When side effects have been weighed against the fix and all concerns have been addressed, fix the problem. If necessary, implement additional steps to correct the problem if your first solution did not work. After you think you have resolved the problem, test the result.
8. Identify the results and effects of the solution. Make sure that the solution has fully fixed the problem and has not caused any other problems.
9. Document the solution and process. In the future, you can check your documentation to see what has changed or to help you remember the solution to common problems.

Remember, however, that troubleshooting is a process of both deduction and induction. Experience will show you when deviating from this process can save both time and effort.

Troubleshooting Utility Facts

The table below describes the tools you can use to troubleshoot network problems.

Task	Tool(s)	Description
View the ARP table	**arp** (Windows)	Shows MAC address-to-IP address mappings including the local MAC and IP addresses.

View IP configuration information	**ipconfig** (Windows 2000 and higher) **ifconfig** (Linux)	Displays IP configuration information for network adapters including: • IP address and mask • Default gateway • DNS servers • WINS servers • DHCP server used for configuration • MAC address
View IP and routing statistics	**netstat** (Windows)	Shows IP-related statistics including: • Current connections • Incoming and outgoing connections • Active sessions, ports, and sockets • The local routing table
View NetBIOS over TCP/IP information	**nbtstat** (Windows)	Displays the NetBIOS name tables for both the local computer and remote computers and the NetBIOS name cache.
Test host-to-host connectivity	**ping**	Sends an ICMP echo request/reply packet to a remote host. A response from the remote host indicates that both hosts are correctly configured and a connection exists between them. Using the **-t** switch with **ping** can be useful in determining whether the network is congested, as such a condition will cause sporadic failures in the ping stream. **Note**: Many firewalls block ICMP messages. To use **ping**, the firewall must allow ICMP messages. Most devices and firewalls allow you to select the specific ICMP messages that are allowed.
Identify the path between two hosts	**tracert** (Windows) **traceroute** (Linux) **mtr** (Linux)	Like **ping**, **traceroute** uses ICMP packets to test connectivity between devices, but as it does so it shows the path between the two devices. Responses from each hop on the route are measured three times to provide an accurate representation of how long the packet takes to reach, and be returned by that host. **Note**: The **mtr** command on Linux is a combination of the **ping** and **traceroute** commands.
Test host-to-host connectivity using ARP	**arping** (Linux)	Sends an ARP request to the specified IP address. **arping** works much like ping in that the host with the specified IP address will respond. Be aware of the following: • **arping** works only on the local subnet (not through routers). • **arping** will often work even if the destination host is blocking ICMP messages.
Test name resolution	**nslookup** (Windows and Linux) **dig** (Linux, this is the preferred	Resolves (looks up) the IP address of a host name. Displays other name resolution-related information such as the DNS server used for the lookup request.

	tool on Linux) **host** (Linux)	
View and modify the routing table	**route**	Use the **route** command to display the contents of the routing table or to add or remove static routes.

Identifying Communication Problems

The first sign of a communication problem often comes when a user says "The network is down" or "I can't reach the server." As part of the troubleshooting process, you need to identify the scope of the problem so you can take the proper actions to correct the problem.

The following example shows one way to troubleshoot communication problems. In this scenario, workstation A can't communicate with workstation C.

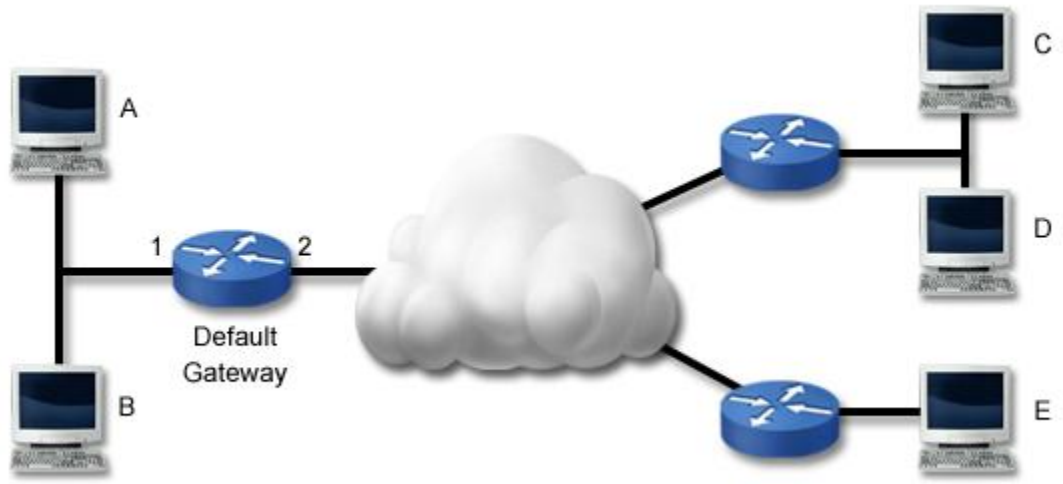

The following table lists several tasks you can perform to troubleshoot connectivity problems. The tasks listed here are listed in order of one way to troubleshoot the reported problem. These steps trace the problem backwards from the remote host to the local host (another way to troubleshoot the issue is to work through these steps in reverse order). Be aware that depending on the situation, you might be able to troubleshoot the problem more efficiently by skipping some tests or changing the order.

Task	Description
Ping host C	Often the best place to start in troubleshooting a problem is to ping the host you are trying to contact. Performing this test *first* verifies the reported problem. If successful, the problem is not related to network connectivity. Check other problems such as name resolution or service access. **Note:** If you have access to another computer, try pinging the destination host from that computer. If successful, then skip the remaining tasks and troubleshoot the local host configuration or physical connection.
Ping host D	If you cannot contact a specific remote host, try pinging another host in the *same* remote network. If successful, then the problem is with the remote host (either a

	misconfiguration, broken link, or unavailable host).
Ping host E	If you cannot contact *any* host in the remote network, try pinging hosts on *other* remote networks (you might try several other networks). If successful, or if you can contact some remote networks and not others, then the problem is with the routing path between your network and the specific remote network. You can then use the **traceroute**/**tracert** commands to check the path to the problem network.
Ping the default gateway	If you cannot contact any remote network, ping the default gateway router. If successful, and you still cannot contact any remote host, have the router administrator verify the router configuration. Check for broken links to the remote network, interfaces that have been shut down, or access control lists or other controls that might be blocking traffic.
Ping host B	If you cannot contact the default gateway router, ping other hosts on the local network. If successful, then check the default gateway router.
Troubleshoot the local host connection or configuration	If you cannot communicate with any host on the local network, then the problem is likely with the local host or its connection to the network. Troubleshoot the following: • Check physical connectivity • Validate the TCP/IP configuration on the local host • Validate IP configuration settings

One special ping test you can perform is to ping the local host. When you ping the local host, you are verifying that TCP/IP is correctly installed and configured on the local host. In essence, you are finding out if the workstation can communicate with itself. To ping the local host, use the following command:

ping 127.0.0.1

If this test fails, check to make sure the TCP/IP is correctly configured on the system. **Note:** This test does not check physical connectivity. The ping can succeed even if the host is disconnected from the network.

Fault Domain Troubleshooting Facts

When troubleshooting physical problems, it helps to identify the *fault domain*. The fault domain is the location of a physical problem and is often manifested by identifying the boundary between communicating devices. For example, if a cable break occurs, a given host might be able to communicate with some devices but not others. When you identify the fault domain, you identify the boundaries of communication and identify the most probable location of the physical problem.

The following table compares how a single break in the network affects device-to-device communication for specific topologies.

Topology	Effect
Bus	A break in the network bus means that the end of the network bus is no longer terminated. For this reason, a break in the bus typically means that no devices can communicate. Identifying

	the location of the break is difficult on a true bus network.
Star	A break in a cable in a star means that the device connected to the central device (hub or switch) through that cable can no longer communicate on the network. All other hosts will be able to communicate with all other devices.
Ring	A break in the ring means that messages can only travel in one direction (downstream) up to the break. Computers can send messages downstream to other devices, but because of the break will not be able to receive any responses.
Mesh	A break in a single link in a mesh topology has no effect on communications. Data can be routed to the destination device by taking a different (sometimes longer) path through the mesh topology.

Link Status Troubleshooting Facts

If a single device is unable to communicate on the network, begin by verifying the physical network connection. Most network cards include link and status lights that can help you verify physical connectivity. The following table describes various light combinations and their meaning in troubleshooting.

Light			Meaning
Link	**Activity**	**Collision**	
Unlit	Unlit	Unlit	The network card does not have a connection to the network. For the link light to be lit, the computer must detect a connection to another device. Possible causes of no link light include: • Bad NIC • Faulty cable • Missing device on the other end (unplugged cable) • Switch or hub port turned off or bad
Red/Amber	Unlit	Unlit	If the link light comes on but is not green, then the NIC has detected a signal but the signal is not what was expected. Possible causes include: • Faulty transceiver on the NIC or on the remote device • Incorrectly configured network cabling • Incompatible networking standards **Note:** On some switches, an amber link light indicates a slower connection (such as 10 Mbps compared to a 100 Mbps connection which might show a green light).
Solid Green	Unlit	Unlit	A solid (normally green) link light indicates a valid network connection. However, an Activity light that *never* lights up means that no data is being received. Check all components and

			connections.
Solid Green	Flashing	Unlit	This is a normal condition that indicates a valid, active connection. The Activity light will periodically flash, even if you are not currently sending data over the link (this is known as a *heartbeat* or *keepalive* signal that lets the NIC know it has an active connection).
Solid Green	Flashing	Flashing/Lit occasionally	This is a normal condition. A small number of collisions are to be expected on an Ethernet network. **Note:** If your network uses full-duplex switches, there should be no collisions on the network.
Solid Green	Flashing constantly	Flashing/Lit occasionally	An Activity light that is constantly flashing indicates constant traffic being sent or received on the link. This could be caused by a device that is very busy (such as a server). In most cases, however, there should be at least some periods of little or no activity. A constantly flashing activity light could be caused by a faulty NIC on the link, constantly sending out data. This condition is known as *chattering* or *jabbering*.
Solid Green	Flashing	Flashing/Lit constantly	If the collision light is constantly flashing, then there are too many collisions on the network. Possible causes include: • Too many devices on the segment. As the number of devices increases, so too will collisions. Reducing the number of devices, or using switches, bridges, or routers to divide the network will reduce the number of collisions. • Faulty cabling, or cable runs that are too long. • A faulty NIC that does not properly sense the medium before transmitting.

Wiring Troubleshooting Facts

The following table describes several conditions that are caused by faulty wiring.

Issue	Description
Interference	*Interference* is an electrical signal on a wire that is not part of the original signal sent on the wire. *Electromagnetic interference* (EMI) is interference that comes from an external source. Common sources of EMI include nearby generators, motors (such as elevator motors), radio transmitters, welders, transformers, and fluorescent lighting. To protect against EMI, use one of the following: • Use fiber optic instead of copper cables. Fiber optic cables are immune from EMI. • Use shielded twisted pair cables. Shielded cables have a metal foil that encloses all of the wires. Some cables might also include a drain wire, which is a bare wire in the middle of the wire bundle that helps to reduce EMI. More expensive cable might also use a metal foil around each pair of wires.

Crosstalk	*Crosstalk* is interference that is caused by signals within the twisted pairs of wires. For example, current flow on one wire causes a current flow on a different wire. • The twisting of wires into pairs helps reduce crosstalk between the two wires *within* the pair. • Each pair of wires is twisted at a different rate to reduce crosstalk *between* pairs. • Crosstalk is often introduced within connectors, where the twists are removed to add the connector. Crosstalk can also occur where wires are crushed or where the plastic coating is worn. There are several forms of crosstalk: • *Near end crosstalk* (NEXT) is crosstalk measured on the same end as the transmitter. For example, when a signal is sent on one wire, near end crosstalk measures the interference on another wire at the same connector near the source of the original signal. • *Far end crosstalk* (FEXT) is crosstalk measured on the opposite end from the transmitter. For example, when a signal is sent on one wire, far end crosstalk measures the interference on another wire at the opposite end from the source signal. • *Alien crosstalk* is crosstalk introduced from adjacent, parallel cables. For example, a signal sent on one wire causes interference on a wire that is within a separate twisted pair cable bundle.
Attenuation	*Attenuation* is the loss of signal strength from one end of a cable to the other. • The longer the cable, the more attenuation. For this reason, it is important to not exceed the maximum cable length defined by the networking architecture. • Cables at a higher temperature experience more attenuation than cables at a lower temperature. • A repeater regenerates the signal and removes the unwanted effects caused by attenuation.
Open impedance mismatch (echo)	*Impedance* is the measure of resistance within the transmission medium. • Impedance is measured in ohms. • All cables must have the same impedance rating. The impedance rating for the cable must match the impedance of the transmitting device. • Impedance is mostly a factor in coaxial cables used for networking. Be sure to choose cable with the correct rating (50 or 75 ohm) based on the network type, and do not mix cables with different ratings. • When signals move from one cable to another with a different impedance rating, some of the signal is reflected back to the transmitter, distorting the signal. For video, impedance mismatch is manifested by ghosting of the image. • Cable distance does not affect the impedance of the cable.
Shorts	An electrical *short* occurs when electrical signals take a different path other than the intended path. In the case of twisted pair wiring, a short means that a signal sent on one wire might arrive on a different wire. Shorts are caused by worn wire jackets or crushed wires so that two wires touch, or by something metal piercing the wire and causing an alternate path.

Open circuit	An *open circuit* is a cut in the wire that prevents the original signal from reaching the end of the wire. With a short, the signal travels a different path, but with an open circuit, the signal does not travel anywhere (electricity cannot flow because the path is disconnected).

Miswired	A *miswired* cable is one in which the individual wires are not in the correct positions within both connectors. Several wiring problems might exist:
	• A *reverse connection* is a cable wired using one standard on one end and a different standard on another end, creating a crossover cable. While this condition might be intentional, it can cause problems when a crossover cable is used instead of a straight-through cable.
	• *Wiremapping* refers to the matching of a wire with a pin in a connector on one end with the same pin on the other end. For example, an error in the wiremapping results when the wire at pin 1 connects to pin 4 on the other connector.
	• A *split pair* condition is where a single wire in two different pairs is reversed at both ends. For example, instead of matching the green and green/white wires in pins 1 and 2, you swap the solid green wire with the solid brown wire. With a split pair configuration, the cable might still work (especially if it is short), but could introduce crosstalk.
	  ○ When you follow the 568A/B standards for making drop cables, one pair is split to meet the standards. In this case, a common split pair error is to not split the pair according to the standard, but to simply place all wire pairs in order in the connector.
	  ○ When connecting cables using a punchdown block, you do not split pairs.

Troubleshooting Tool Facts

The following troubleshooting tools can be useful in troubleshooting physical connectivity problems.

Tool/Method	Description
Loopback plug	A *loopback plug* reflects a signal from the transmit port on a device to the receive port on the same device. Use the loopback plug to verify that a device can both send and receive signals. • A failure in the loopback test indicates a faulty network card. • A successful loopback test means the problem is in the network cabling or other connectivity devices. You can purchase loopback plugs already made, or make an inexpensive one by cutting the end off a cable and manually connecting the transmit wires to the receive wires (connect the wire from pin 1 with the wire at pin 3, and the wire at pin 2 with the wire at pin 6).
Smart jack	A *smart jack* is a special loopback plug installed at the demarcation point for a WAN service. Technicians at the central office can send diagnostic commands to the smart plug to test connectivity between the central office and the demarc. When you contact a WAN service provider for assistance, they might execute a test using the smart jack. A successful test indicates that the problem is within your customer premises equipment

	(CPE).
Known good spares	One valuable troubleshooting method is to keep a set of components that you know are in proper functioning order. If you suspect a problem in a component, swap it with the known good component. If the problem is not resolved, troubleshoot other components. Examples of using this strategy are: • Changing the drop cable connecting a computer to the network. • Replacing a NIC with one you know to be working. • Moving a device from one switch port to another.
Cable tester	A *cable tester* verifies that the cable can carry a signal from one end to the other, and that all wires are in their correct positions. • A good cable tester can check for various miswire conditions (wiremapping, reversals, split pairs, shorts, or open circuits). • You can use a cable tester to quickly tell the difference between a crossover and a straight-through cable. • Most testers have a single unit that tests both ends of the cable at once. Many testers come with a second unit that you can plug into one end of a long cable run to test the entire cable.
Time Domain Reflector (TDR)	A TDR is a special device that sends electrical pulses on a wire in order to discover information about the cable. The TDR measures impedance discontinuities, or in other words the echo received on the same wire in response to a signal on the wire. The results of this test can be used to: • Estimate the length of a wire. • Measure the cable impedance. • Identify locations of splices and connectors on the wire. • Identify shorts and open circuits, and the location of the fault. An Optical Time Domain Reflector (OTDR) performs the same function for fiber optic cables using light waves.
Certifier	A cable *certifier* is a multi-function tool that verifies or validates that a cable or an installation meets the requirements for a specific architecture implementation. For example, you would use a certifier to verify that a specific drop cable meets the specifications for 1000BaseT networking. • A certifier is very important for Cat 6 cable used with bandwidths at or above 1000 Mbps. Slight errors in adding connectors to cables might mean that the network functions at 100 Mbps instead of the desired 1000 Mbps or 10 Gbps. • Certifiers can also validate the bandwidth capabilities of network interface cards and switches. Many can detect the duplex settings of network devices. • Most certifiers include features of a toner probe, TDR, and cable tester. • Certifiers are very expensive, and are typically used by organizations that specialize in wiring installations.
Toner probe	A *toner probe* is two devices used together to trace the end of a wire from a known endpoint into the termination point in the wiring closet. • Connect the tone generator to one end of the wire. The tone generator sends a

	signal on the wire. • In the wiring closet, touch the probe to wires or place the probe close to wires. A sound at the probe indicates that the generated tone has been detected. The wire that you are touching is the termination point for the wire you want to trace.
Butt set	A *butt set* (also called a *lineman's handset*) is a device used for testing analog telephone installations. The butt set includes an earpiece, mouthpiece, dialing interface, and various connectors. When troubleshooting a telephone line, you connect the butt set to a set of wires. You can use the butt set to monitor a call, make an outgoing call, or answer an incoming call. The butt set helps you identify portions of the telephone circuit that are working, so you can better identify where a problem in the telephone wiring is at. You can make a simple butt set (called a *beige box*) by cutting the wires on a regular telephone and connecting alligator clips that can be used for connecting to the bare wires within the telephone wiring distribution.
Multimeter	A *multimeter* is a device for testing various electrical properties. For example, most multimeters can measure: • AC and DC voltage • Current (amps) • Resistance (ohms) • Capacitance • Frequency
Voltage event recorder	A *voltage event recorder* keeps track of voltage conditions on a power line. Basic recorders simply keep track of the occurrence of under or over voltage conditions, while more advanced devices track conditions over time and create a graph, saving data from a program running on a computer. Some UPS systems include a simple voltage event recorder. Use a voltage event recorder to be notified or identify periods of low or high voltage that can adversely affect network components.
Temperature monitor	A temperature monitor does what its name implies, it records temperature over time. • Some monitors are internal to a computer, and monitor case and CPU temperature. • Other monitors keep track of temperatures within server rooms, and can monitor other conditions such as humidity, water, smoke, motion, and air flow. Many monitors include the ability to notify you when an excessive temperature threshold is reached.

Interpreting ipconfig

You can use **ipconfig /all** to troubleshoot IP configuration problems. The following table describes how the output for this command changes based on how IP settings are configured and for specific problem situations.

Condition	ipconfig /all Output
Static IP Configuration	If the workstation is configured with static IP information, the following conditions will exist: • The **DHCP Enabled** line will show **No** • The **DHCP Server** line will not be shown
DHCP Configuration	If the workstation has received configuration information from a DHCP server, the following conditions will exist: • The **DHCP Enabled** line will show **Yes** • The **DHCP Server** line will show the IP address of the DHCP server from which configuration information was received
Rogue DHCP Server	A *rogue* DHCP server is an unauthorized DHCP server on the network. Symptoms of a rogue DHCP server include: • Conflicting IP addresses on the network • Incorrect IP configuration information on some hosts To identify a rogue DHCP server using **ipconfig**, verify the DHCP server address. If this address is not the address of your DHCP server, you have a rogue DHCP server. **Note:** When you have a rogue DHCP server on the network, some hosts will likely receive configuration information from the correct DHCP server and some from the rogue DHCP server.
Incorrectly Configured DHCP Server	Your DHCP server can send out various IP configuration values in addition to the IP address and mask. If network hosts are configured with incorrect IP values (such as incorrect default gateway or DNS server addresses), first verify that the workstations are contacting the correct DHCP server. If the correct server is being used, go to the DHCP server to verify that it is sending out correct configuration information.
APIPA Configuration	If the workstation has used APIPA to set configuration information, the following conditions will exist: • The **DHCP Enabled** line will show **Yes** • The **DHCP Server** line will not be shown • The IP address will be in the range of 169.254.0.1 to 169.254.255.254 with a mask of 255.255.0.0 • The **Default Gateway** line will be blank • The **DNS Servers** line will not include any IPv4 addresses **Note:** When APIPA is used, the workstation sets its own IP address and mask. It does not automatically configure default gateway or DNS server values. When APIPA is being used: • Communication is restricted to hosts within the same subnet (there is no default gateway set). • Hosts can communicate with other hosts that have used APIPA. If some hosts are still using an address assigned by the DHCP server (even if the DHCP server is down), these hosts will not be able to communicate with the APIPA

	hosts. • Name resolution will not be performed (there are no DNS server addresses configured).
Alternate Configuration	If the workstation has been configured using an alternate configuration, the following conditions will exist: • The **DHCP Enabled** line will show **Yes** • The **DHCP Server** line will not be shown • The IP address and subnet mask values will be a value other than the APIPA values • Default gateway and DNS server addresses will be configured using the alternate configuration values

If the workstation has received configuration information from the wrong DHCP server or configured itself using APIPA, you might need to retry to contact the DHCP server once DHCP problems have been resolved. Use the following commands:

- Use **ipconfig /release** to stop using the current dynamic IP configuration parameters.
- Use **ipconfig /renew** to retry the DHCP server request process to obtain IP configuration parameters.

Note: To display the TCP/IP configuration on a Linux computer, use the **ifconfig** command.

arp, netstat, and nbtstat Facts

The following table lists several commands on a Windows system that you can use to gather information about network connections.

Tool	Option(s)
arp	**arp -a** shows the IP address-to-MAC address mapping table (the address cache)
netstat	**netstat** shows the active connections
	netstat -a shows detailed information for active connections
	netstat -r or **route print** shows the routing table of the local host
	netstat -s shows TCP/IP statistics
nbtstat	**nbtstat -c** shows the IP address-to-NetBIOS name mapping table (the name cache)

Note: Local computers have a cache of recently-used IP addresses and their corresponding MAC addresses. When a computer needs to contact another computer on its own subnet, it first checks its cache for an entry of the IP address. If found, the corresponding MAC address is used to communicate with the destination computer. The cache can cause problems if the MAC address for a computer has recently changed, such as if the network interface card has been replaced. To correct the problem, use the **netsh** command to clear the ARP cache.

Troubleshooting Name Resolution Facts

Name resolution problems typically have the following symptoms:

- You can ping a destination host using its IP address.
- Methods that use the logical host name to communicate with the host fails. This might include things such as:
    - Typing a URL into the browser.
    - Pinging the host using the host name.
    - Searching for the host by its name.

To troubleshoot DNS name resolution, use one of the following tools:

- **nslookup** for Windows or Linux systems
- **dig** or **host** for Linux systems (**dig** is replacing **nslookup** on Linux systems)

The following table lists several ways to use these commands.

Use...	To...	Example
nslookup *host*	Query the IP address of a host.	`nslookup www.mit.edu`
nslookup	Start **nslookup** in interactive mode. The default interactive mode query is for A records, but you can use the **set type=** command to change the query type.	`nslookup` `set type=ns`
dig *hostname* **host** *hostname*	Query a host. The default query is for A records. You can change the default search by appending one of the record types you see below to the end of the command. • a--address records • any--any type of record • mx--mail exchange records • ns--name server records • soa--sort of authority records • hinfo--host info records • axfr--all records in the zone • txt--text records	`dig www.vulture.com ns` `host www.vulture.com -t ns`
dig *@IP address or host name domain*	Query the root server at the IP address or host name for A records for the domain. You can change the default query type by appending a different record type to the back of the command.	`dig @192.168.1.1 vulture.com ns`

dig -x *IP address* **host** *IP address*	Find the host name for the queried IP address.	`dig -x 62.34.4.72` `host 62.34.4.72`

Note: Local computers have a cache of recently-resolved DNS names. The cache holds the DNS name and the IP address for that name. When you use a DNS name, the computer first checks its cache. If the name is in the cache, the corresponding IP address will be used. This can sometimes cause problems if the IP address of a host has changed. Old values in the cache might continue to be used for a time, making communication using the DNS name impossible. To correct this problem on a Windows computer, run **ipconfig /flushdns** to delete the local DNS name cache.

Switch Troubleshooting Facts

The following table lists several problems that you might encounter when managing switches on your network.

Issue	Description
Collisions	A *collision* occurs when two devices that share the same media segment transmit at the same time. In a switched network, collisions should only occur on ports that have multiple devices attached. • To eliminate collisions, connect only a single device to each switch port. • If collisions are still detected, verify that the port is not being shared by multiple devices, then troubleshoot cable and NIC issues.
Duplex mismatch	A *duplex mismatch* occurs when two devices are using different duplex settings. In this case, one device will try to transmit using full duplex, while the other will expect half duplex communications. By default, devices are configured to use autonegotiation to detect the correct duplex setting to use. If a duplex method cannot be agreed upon, devices should default to using half duplex. A duplex mismatch can occur in the following cases: • Both devices are configured to use different duplex settings. • Autonegotiation does not work correctly on one device. • One device is configured for autonegotiation and the other device is manually configured for full duplex. Symptoms of a duplex mismatch include very slow network communications. Ping tests might appear to complete correctly, but normal communications work well below the expected speeds, even for half duplex communications.
Slow link speed	Most network components are capable of supporting multiple network specifications. For example, a NIC might support both 10BaseT and 100BaseTX. By default, such devices will detect other devices and use the maximum speed supported by all devices. If you find that the speed in use on a segment is lower than expected (for example 10 Mbps instead of 100 Mbps, or 100 Mbps instead of 1000 Mbps):

	• Check individual devices to verify that all support the higher speed. • Check individual devices to see if any have been manually configured to use the lower speed. • Use a cable certifier to verify that the cables meet the rated speeds. Bad cables are often the cause of 1000BaseT networks operating only at 100BaseTX speeds.
Switching loop	A *switching loop* occurs when there are multiple active paths between two switches. Switching loops lead to incorrect entries in a MAC address table, making a device appear to be connected to the wrong port, and causing unicast traffic being circulated in a loop between switches. The spanning tree protocol runs on switches to prevent switching loops by making only a single path between switches active at a single time.
Broadcast storm	A *broadcast storm* is excessive broadcast traffic that renders normal network communications impossible. Broadcast storms can be caused by: • Switching loops that cause broadcast traffic to be circulated endlessly between switches. • Denial of Service (DoS) attacks. To reduce broadcast storms: • Run the spanning tree protocol to prevent switching loops. • Implement switches with built-in broadcast storm detection which limits the bandwidth that broadcast traffic can use. • Use VLANs to create separate broadcast domains on switches.
Incorrect VLAN membership	VLANs create logical groupings of computers based on switch port. Because devices on one VLAN cannot communicate with devices in different VLANs, incorrectly assigning a port to a VLAN might prevent a device from communicating through the switch. **Note**: VLAN membership is defined by switch port, not by MAC address. Connecting a device to a different switch port might change the VLAN membership of the device. On the switch, verify that ports are assigned to the correct VLANs, and that any unused VLANs are removed from the switch.
Frame errors	The switch examines incoming frames and will only forward frames that are complete and correctly formed, while invalid frames are simply dropped. Most switches include logging capabilities to track the number of corrupt or malformed frames. • Frames that are too long are typically caused by a faulty network card that jabbers (constantly sends garbage data). • Frames that are too short are typically caused by collisions. • CRC errors indicate that a frame has been corrupted in transit. • All types of frame errors can be caused by faulty cables or physical layer devices.

Troubleshooting Routing Facts

The general symptoms that indicate a routing problem is the inability to access hosts on a specific network or the inability to access any remote network. The following table lists various problems that are typically caused by routing issues.

Problem	Description
Can't access hosts outside the local subnet	If one or more hosts can only communicate with hosts on the local subnet, the problem is likely with the default gateway configuration. • If a single host is having problems, verify the default gateway setting on that host. • If multiple hosts are having problems, verify the default gateway setting, and verify that the DHCP server is configured to deliver the correct default gateway address. • If all hosts have the same problem, and if the default gateway setting is correct, verify that the default gateway server is up and configured for routing.

Can't communicate with any host on a specific network	If hosts are unable to contact hosts on a specific subnet, but communication with other subnets is working, check the following:
	
	• Verify that the router connected to the subnet is up.
	• Use the **route** command on the default gateway of the local subnet and verify that the router has a route to the remote subnet. If necessary, configure a static route or a routing protocol so that the route can be learned automatically.
	• Use **traceroute** to view the route taken to the destination network. Identify the last router in the path, then troubleshoot routing at that point.
	• Check for *routing loops* in the path to the destination network. A routing loop is caused by a misconfiguration in the routers along the path such that data is sent back along the same path rather than being forwarded to the destination. Routing loops are indicated by:
	  ○ Routing table entries that appear and then disappear (called *route flapping*), often at regular intervals (such as every minute).
	  ○ Routing table entries where the next hop router address oscillates (changes) between two or more different routers.
	  ○ **Traceroute** output that displays the same sequence of routers being repeated.

Can't access the Internet	If hosts are able to reach all internal networks but can't access the Internet, check for the following: • Verify that the Internet connection is up. • Check for a default route on the router connected to the Internet. A default route is indicated by a network address of 0.0.0.0 with a mask of 0.0.0.0. The default route is used for all packets that do not match another entry in the routing table. **Note**: Most routers connecting private networks to the Internet do not know about specific networks and routes on the Internet. In addition, most routers do not share routes for private subnets with Internet routers. Instead, the router is configured with a single default route that is used for all Internet traffic, and a router at the ISP is responsible for sharing a single route into your private network with other Internet

	routers.
Remote clients can't access network resources	If you have remote access clients who can establish a connection to the remote access server, but can't connect to other resources on the private network, check the following: • If remote clients are being assigned an IP address on the same subnet as the private network, make sure that proxy ARP is enabled on the LAN interface of the remote access server. Proxy ARP is required to make it appear as if the remote clients are connected to the same network segment. • If remote clients are being assigned an IP address on a different subnet than the private network, make sure the remote access server has routing configured to route packets between the remote clients and the private network.