

TROUBLESHOOTING NETWORK PROBLEMS

After reading this chapter and completing the exercises, you will be able to:

- ▶ Describe the elements of an effective troubleshooting methodology
- ▶ Follow a systematic troubleshooting process to solve networking problems
- ▶ Use a variety of software and hardware tools to diagnose problems
- ▶ Discuss practical issues related to troubleshooting



ON THE JOB

Our ISP division hosts Web sites for a number of corporate clients. Each site requires a separate Web server, but multiple Web servers can run on the same machine. Once, at about 3:00 A.M. on a Sunday morning, one of our engineers began upgrading the hardware that supported about 100 of these corporate Web servers. The engineer finished the work on schedule, and everything appeared to be fine. We expected the sites to perform much better with the new hardware installed.

At roughly 6:00 A.M. that morning, I was with a customer working on a network topology conversion project, when the engineer called with bad news. Contrary to our expectations, the Web sites were performing dismally on the new hardware. The exact problem wasn't clear. I made my apologies to the customer (a down condition always takes precedence) and headed back to the office to do some troubleshooting.

Indeed, the performance of the Web sites on the new hardware was awful. I plugged in a Network General Sniffer to our core Ethernet switch and then set the sniffer port to spanning mode so that the sniffer could examine all traffic on the Web server VLAN. I set the sniffer filters such that I was monitoring only packets to the Web server in question. Almost immediately, a problem became apparent. The packets destined for the Web server were plainly seen on the network, but no replies came from the Web server. For some reason, the Web server was not "seeing" the traffic directed to it. A number of causes seemed possible: For example, the wiring to the new Web server might be bad, the Web server might have a defective network interface card (NIC), or, less likely, the switch might have a bad Ethernet port.

I tried the easiest option first, replacing the Category 5 Ethernet cable to the Web server. Sure enough, the problem went away, and the Web pages were quickly served. The sniffer showed normal network protocol behavior.

The company had engaged in discussions about getting a cable tester several times, but never quite got around to making a purchase. After this incident, we ordered a cable tester immediately. We also implemented a policy requiring engineers to test each cable before installing it on the network.

James G. Berbee
Berbee Information Networks, Inc.

By now, you know how networks should work. Like other complex systems, however, they don't always work as planned. Many things can go wrong on a network, just as many things can go wrong with your car, house, or a project at work. In fact, a network professional probably spends more time fixing network problems than designing or upgrading a network. Some breakdowns (such as an overtaxed processor) come with plenty of warning, but others (such as a hard disk controller failure) can strike instantly.

As with your car, the best defense against problems is prevention. Just as you should have your car serviced regularly, so you should monitor the health of your network regularly. Of course, even the most well-monitored network will sometimes experience unexpected problems. For example, a utility company could dig a new hole for its cable and accidentally cut your dedicated line to the Internet. In such a situation, your network can go from perfect to disastrous performance in an instant. In this chapter, you learn how to diagnose and solve network problems in a logical, step-by-step fashion, using a variety of tools.

TROUBLESHOOTING METHODOLOGY

Successful troubleshooters proceed logically and methodically. This section introduces a basic troubleshooting methodology, leading you through a series of general problem-solving steps. Bear in mind that experience in your network environment may prompt you to follow the steps in a different order or to skip certain steps entirely. For example, if you know that one segment of your network is poorly cabled, you may try replacing a section of cable in that area to solve a connectivity problem before attempting to verify the physical and logical integrity of the workstation's NIC. In general, however, it is best to follow each step in the order shown. Such a logical approach can save you from undertaking wasteful, time-consuming efforts such as unnecessary software or hardware replacements.

Steps for troubleshooting network problems are as follows:

1. Identify the symptoms. Carefully document what you learn from people or systems that alerted you to the problem and keep that documentation handy.
2. Identify the scope of the problem. Is it universal—that is, are all users on the network experiencing the problem at all times? Or is the problem limited to a specific geographic area of the network, to a specific demographic group of users, or to a particular period of time? In other words, is the problem subject to geographic, demographic, or chronological constraints?
3. Establish what has changed on the network. Recent hardware or software changes may be causing the symptoms.
4. Determine the most probable cause of the problem. This determination may include the following techniques:
 - a. Verify user competency.
 - b. Re-create the problem, and ensure that you can reproduce it reliably.
 - c. Verify the physical integrity of the network connection (such as cable connections, NIC installations, and power to devices), starting at the affected nodes and moving outward toward the backbone.
 - d. Verify the logical integrity of the network connection (such as addressing, protocol bindings, software installations, and so on).

5. Implement a solution.
6. Test the solution.
7. Recognize the potential effects of the solution. For example, if you have to reassign IP addresses, how will the change of an IP address on a server affect its clients? Or, in another case, if you upgrade the type of client software used on a workstation, how will that affect a user's daily routine?
8. Document the solution. Make sure that both you and your colleagues understand the cause of the problem and how you solved it. This information should be kept in a centrally available repository, such as an online database.

Depending on your findings, you may skip from one step to another step further down in the list, eliminating the need to carry out the intervening steps. For example, if you determine that a NIC has been improperly seated in a workstation's system board, you may skip directly to Step 5 (in this case, reinstall the NIC) without analyzing recent changes to the network. Above all, use common sense in your troubleshooting efforts. As you read through the following sections, you will understand how the suggested troubleshooting steps are interrelated and how answering a question under one step might prompt you to skip to another step.

The flowchart in Figure 12-1 illustrates how these steps are related. Each decision step in the flowchart is discussed in more detail in the following sections, and in some sections the flowchart is expanded to reflect different outcomes based on different findings. The following sections also explain how to narrow down the possible causes of a problem by answering specific questions. In particular, you can question users to get clues about the problem. Finally, the chapter describes ways to test your attempted resolution of a network problem.



In addition to the organized method of troubleshooting described in this section, a good, general rule for troubleshooting can be stated as follows: Pay attention to the obvious! Although some questions may seem too simple to bother asking, don't discount them. You can often save much time by checking cable connections first. Every networking professional can tell a story about spending half a day trying to figure out why a computer wouldn't connect to the network, only to discover that the network cable was not plugged into the wall jack or the device's NIC.

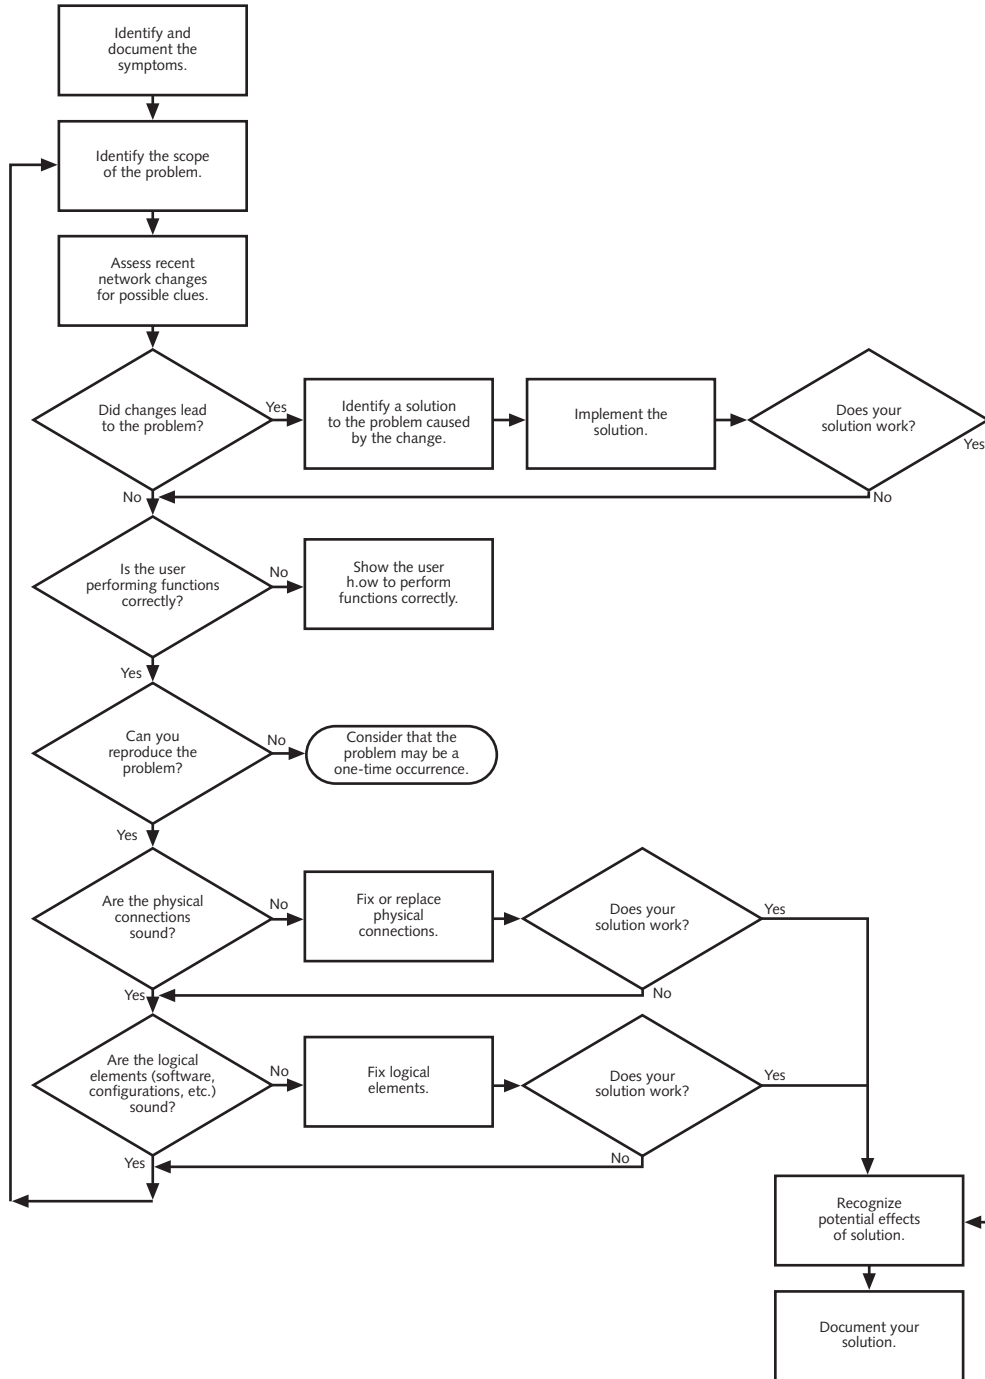


Figure 12-1 A simple flowchart of troubleshooting steps

Identify the Symptoms

When troubleshooting a network problem, act like a doctor diagnosing a patient's illness. Your first step should be to identify the specific symptoms of the problem. In a broad sense, this step brings you closer to pinpointing the cause of the problem. For example, identifying a patient's sore throat and headache as symptoms, rules out carpal tunnel syndrome and a host of other ailments. Nevertheless, the problem may still be anything from mononucleosis to allergies.

In a network, symptoms of a single problem might include a user's inability to access a network drive, send e-mail, or print to a specific printer. The problem may be caused by a number of things, including a faulty NIC, a faulty cable, a faulty hub, a faulty router, an incorrect client software configuration, a server failure, or a user error. On the other hand, you can probably rule out a power failure, a printer failure, an Internet connectivity failure, an e-mail server failure, and a host of other problems.

Answering the following questions may help you identify the symptoms of a network problem:

- Is access to the network affected?
- Is network performance affected?
- Are data or programs affected? Or are both affected?
- Are only certain network services (such as printing) affected?
- If programs are affected, does the problem include one local application, one networked application, or multiple networked applications?
- What specific error messages do users report?
- Is one user or are multiple users affected?
- Do the symptoms manifest themselves consistently?

One danger in troubleshooting technical problems lies in jumping to conclusions about the symptoms. For example, you might field 12 questions from users one morning about a problem printing to the network printer in the Facilities Department. You might have already determined that the problem is an addressing conflict with the printer and be in the last stages of resolving the problem. Minutes later, when a 13th caller says, "I'm having problems printing," you might immediately conclude that she is another Facilities staff member and that her inability to print results from the same printer addressing problem. In fact, this user may be in the Administration Department, and her inability to print could represent a symptom of a larger network problem.

Take time to pay attention to the users, system and network behaviors, and any error messages. Treat each symptom as unique (but potentially related to others). In this way, you will avoid the risk of ignoring problems or—even worse—causing more problems.



Take note of the error messages reported by users. If you aren't near the users, ask them to read the messages to you directly off their screens or, better yet, print the screens that contain the error messages. (On some computers, pressing the Print Screen button—which is sometimes labeled “Print Scrn” or “Prt Sc”—will perform the Print Screen function. On other computers, you can use the Shift-Print Screen or Alt-Print Screen keystroke combinations.) Keep a record of these error messages along with your other troubleshooting notes for that problem.

Identify the Scope of the Problem

After you have identified the problem's symptoms and ruled out user error, you should determine the scope of the problem—whether the problem appears only with a certain group of users, with certain areas of the organization, or at certain times. For example, if a problem affects only users on one network segment, you may deduce that the problem lies with that network segment's cabling, configuration, router port, or gateway. On the other hand, if symptoms are limited to one user, you can typically narrow the cause of the problem down to a single cable, workstation (hardware or software) configuration, or user.

In the doctor/patient analogy, this scope identification process is similar to that of the doctor who asks a patient how long his sore throat has lasted and whether anyone else in his family is affected. If the patient answers that the sore throat started yesterday and his twin toddlers both have colds, the doctor might suspect a cold virus. Conversely, if the patient indicates that no one he knows is ill and that his sore throat has lingered for 10 days, the doctor might suspect something other than a simple cold.

Answering the following questions may help you ascertain the scope of a network problem:

- How many users or network segments are affected?
 - One user or workstation?
 - A workgroup?
 - A department?
 - One location within an organization?
 - An entire organization?
- When did the problem begin?
 - Has the network, server, or workstation ever worked properly?
 - Did the symptoms appear in the last hour or day?
 - Have the symptoms appeared intermittently for a long time?
 - Do the symptoms appear only at certain times of the day, week, month, or year?

Like identifying symptoms, narrowing down a problem's scope can eliminate some causes and point to others. In particular, narrowing down the affected groups of users or areas of your organization can help to distinguish workstation (or user) problems from network

problems. If the problem affects only a department or floor of your organization, for example, you will probably need to examine that network segment, its router interface, its cabling, or a server that provides services to those users. If a problem affects users at a remote location, you should examine the WAN link or its router interfaces. If a problem affects all users in all departments and locations, a catastrophic failure has occurred, and you should assess critical devices such as central switches and backbone connections.



If a problem is universal—that is, if it affects the entire LAN or WAN—you will naturally want to answer these questions very quickly. In the doctor/patient analogy, this situation would be similar to performing triage in an emergency room.

Usually, network problems are not catastrophic, and you can take a little time to troubleshoot them correctly, by asking specific questions designed to identify their scope. For example, suppose a user complains that his mail program isn't picking up e-mail. You should begin by asking when the problem began, whether it affects only that user or everyone in his department, and what error message (or messages) the user receives when he attempts to pick up mail. In answering your questions, he might say, "The problem began about 10 minutes ago. Both my neighbors are having problems with e-mail, too. And as a matter of fact, a network technician was working on my machine this morning and installed a new graphics program."

As you listen to the user's response, you may need to politely filter out information that is unlikely to be related to the problem. In this situation, the user relayed two significant pieces of information: (1) the scope of the problem includes a group of users, and (2) the problem began 10 minutes ago. With this knowledge, you can then delve further in your troubleshooting. In this example, you would proceed by focusing on the network segment rather than on one workstation.

Discovering the time or frequency with which a problem occurs can reveal more subtle network problems. For example, if multiple users throughout the organization cannot log onto the server at 8:05 A.M., you may deduce that the server needs additional resources to handle the processing burden of accepting so many logins. If a network fails at noon every Tuesday, you may be able to correlate this problem with a test of your building's power system, which causes a power dip that affects the servers, routers, hubs, and other devices.

Identifying the scope of the problem will lead you to your next troubleshooting steps. The path may not always be clear-cut, but as the flowcharts in Figures 12-2 and 12-3 illustrate, some direction can be gained from narrowing both the demographic (or geographic) and chronological scope of a problem. Notice that these flowcharts end with the process of further troubleshooting. In the following sections, you will learn more about these subsequent troubleshooting steps.

The processes of identifying a problem's scope by demographics and by chronology are not mutually exclusive, but rather can be followed simultaneously. For example, you might quickly determine that users in the Software Department experience frequent network

disconnections, but only during the hours between midnight and 2:00 A.M. Knowing that the only staff members working at that time are software engineers, you might choose not to continue through the process of narrowing the problem's demographic scope. Instead, you would want to focus on the network activity during those two hours.

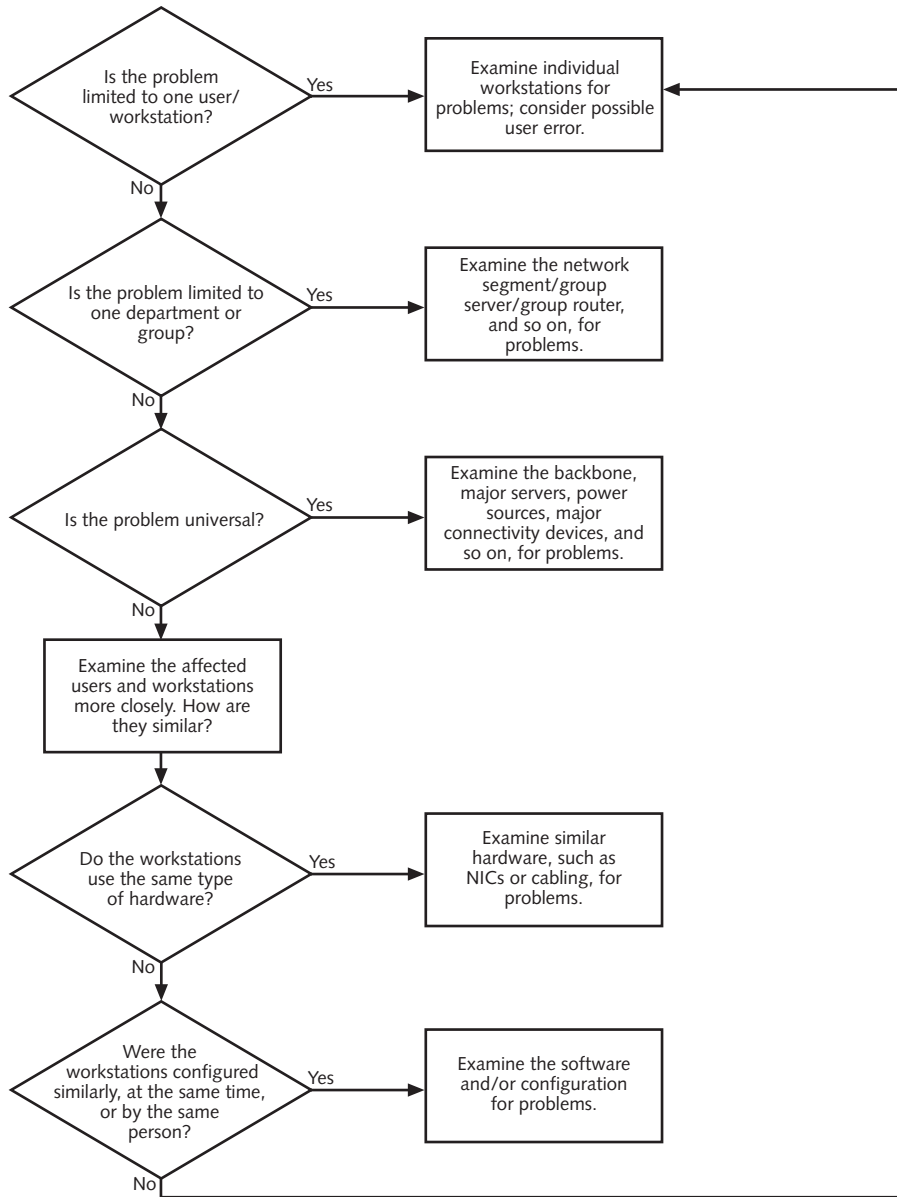


Figure 12-2 Troubleshooting while identifying the demographic scope of a problem

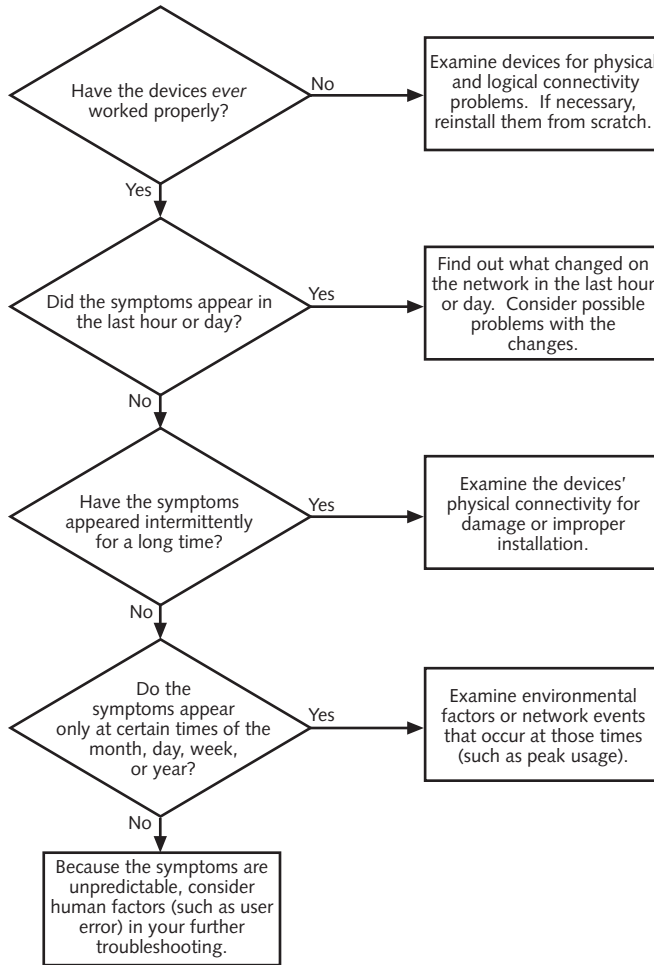


Figure 12-3 Troubleshooting while identifying the chronological scope of a problem



One fascinating example of scope-based (or chronological) troubleshooting was experienced by a wireless networking engineer working on a small metropolitan area network. His spread-spectrum RF network links, which connected businesses to a carrier's POP via a transmitter and receiver on a hospital's roof, worked perfectly all day, but failed when the sun went down each day. When the sun came up the next morning, the wireless links worked again. The engineer confirmed that the equipment was fully operational (as he suspected), then talked with the hospital personnel. The hospital's director informed him that the hospital had installed security cameras on the outside of the building. The cameras used the same RF frequency as the network's wireless links. When the security cameras were activated at sunset, their signals interfered with the wireless network's signals, preventing data from reaching their destination.

Establish What Has Changed

One could argue that considering recent network changes is not a separate step, but rather a continual and integral part of the troubleshooting process. As you begin troubleshooting, you should be aware of any recent changes to your network. These changes may include—among other things—the introduction of new equipment (cabling, connectivity devices, servers, and so on); repair of existing equipment; removal of equipment; installation of new components on existing equipment; installation of new services or applications on the network; equipment moves; addressing or protocol changes; software configuration changes on servers, connectivity devices, or workstations; and modifications to rights, groups, or users. As you can imagine, such changes can create problems if not planned and implemented carefully.

To determine what has changed on a network, you and your colleagues in the IT department should keep complete network change records. You will learn more about maintaining change records in Chapter 13. The more precisely you describe a change, its purpose, and the time and date when it occurred, in your records, the easier your troubleshooting will be if the change subsequently causes problems.

In addition to keeping thorough records, you must make them available to staff members who might need to reference them. For example, you might want to keep a record of changes in a spreadsheet file on a file server, and then use a Web-based form to retrieve and submit information from and to the spreadsheet. That way, no matter where a network technician was working in the organization, she could retrieve the information from any Web-enabled workstation. A simpler alternative is to keep a clipboard in the computer room with notes about changes.

Often, network changes cause unforeseen problems. For example, if you have narrowed a connectivity problem to a group of six users in the Marketing Department, you might refer to your network's change log and find that a hub in the Marketing Department's telecommunications closet was recently moved from one end of the closet to another. Reviewing the record of this change can help you more quickly pinpoint the hub as a possible cause of the problem. Perhaps the hub was incorrectly reconnected to the backbone after the move, or perhaps it became damaged in the move or lost its configuration.

The following questions may help you pinpoint a problem that results from a network change:

- Did the operating system or configuration on a server, workstation, or connectivity device change?
- Were new components added to a server, workstation, or connectivity device?
- Were old components removed from a server, workstation, or connectivity device?
- Was a server, workstation, or connectivity device moved from its previous location to a new location?

- Was a server, workstation, or connectivity device replaced?
- Was new software installed on a server, workstation, or connectivity device?
- Was old software removed from a server, workstation, or connectivity device?

If you suspect that a network change has generated a problem, you can react in two ways: you can attempt to correct the problem that resulted from the change, or you can attempt to reverse the change and restore the hardware or software to its previous state. Both options come with hazards. Of the two, reverting to a previous state is probably less risky and less time-consuming.

However, correcting the problem is sometimes the best solution. For example, if you immediately suspect that a change-related problem can be fixed easily, try correcting the problem first. If it is impossible to restore a software or hardware configuration to its previous state, you must solve the problem with the change in place. You will learn more about modifying a network and then reversing the change in Chapter 13.



Before changing a network device or configuration, develop a plan and gather the proper resources for reversing the change in case things go wrong. For example, if you replace the memory module in a server, you should keep the old memory module handy in case the new one has flaws. In another situation, you might keep a backup of device or application configurations—perhaps by making a copy of the directory that stores the target configuration.

Select the Most Probable Cause

Once you have identified the scope of the problem and analyzed recent changes to the network, you are close to determining the problem’s cause. The following sections provide techniques on how to zero in on the most likely cause among several plausible scenarios.

Verify User Competency

You have probably experienced a moment in your dealings with computers in which you were certain you were doing everything correctly, but still couldn’t access the network, save a file, or pick up your e-mail. For example, you may have typed your case-sensitive network password without realizing that the Caps Lock function was turned on. Even though you were certain that you typed the right password, you received a “password incorrect” error message each time you tried to enter your password. All users experience such problems from time to time.

It’s natural for human beings to make mistakes. Thus, as a troubleshooter, one of your first steps should be to ensure that human error is not the source of the problem. This approach will save you time and worry. In fact, a problem caused by human error is usually simple to solve. It’s much quicker and easier to assist a user in remapping a network drive, for example, than to perform diagnostics on the file server.

Often, an inability to log onto the network results from a user error. Users become so accustomed to typing their passwords every morning and logging onto the network that, if something changes in the logon process, they don't know what to do. In fact, some users might never log out, so they don't know how to log on properly. Although these kinds of problems may seem simple to solve, unless a user receives training in the proper procedures and understands what might go wrong, he or she will never know how to solve a logon problem without assistance. Even if the user took a computer class that covered logging on, he or she may not remember what to do in unfamiliar situations.

When diagnosing user errors, your most powerful tool may be patience. The best way to verify that a user is performing network tasks correctly is by watching the user. If this tactic isn't practical, the next best way is to talk with the user by phone while he or she tries to replicate the error. At every step, calmly ask the user to explain what appears on the screen and what, exactly, he or she is doing. After every keystroke or command, ask the user again what appears on the screen. With this methodical approach, you will be certain to catch any user-generated mistakes. At the same time, if the problem does not result from human error, you will gain important clues for further troubleshooting.

Re-create the Problem

An excellent way to learn more about the causes of a problem is to try to re-create the symptoms yourself. If you cannot reproduce the symptoms, you may suspect that a problem was a one-time occurrence or that a user performed an operation incorrectly.

You should try to reproduce symptoms both while logged on as the user who reported the problem and while logged on under a privileged account (such as an administrator-equivalent ID). If the symptoms appear only when you're logged on under the user's ID, you may suspect that the problem relates to the user's limited rights on the network. For example, a user may complain that he was able to edit a particular spreadsheet in the Accounting directory on the file server on Friday, but was unable to open the file on Monday. When you visit his workstation, you can verify this sequence of events while logged on with his user name. When you then log on as Administrator, however, you may be able to open and edit the file. The difference in your experiences points to a user rights problem. At that point, you should check the user's privileges—especially whether they have changed since he could last retrieve the file. Perhaps someone removed him from a group that had Read and Modify rights to the Accounting directory.

Answering the following questions may help you determine whether a problem's symptoms are truly reproducible and, if so, to what extent:

- Can you make the symptoms recur every time?
- Can you make the symptoms recur some of the time?

- Do the symptoms happen only under certain circumstances? For instance, if you log on under a different ID or try the operation from a different machine, do the symptoms still appear?
- Do the symptoms ever happen when you try to repeat them?

When attempting to reproduce the symptoms of a problem, you should follow the same steps that the person reporting the symptoms followed. As you know, many computer functions can be achieved through different means. For example, in a word-processing program, you might save a file by using the menu bar, using a keystroke combination, or clicking a button on a toolbar. All three methods result in the same outcome. Similarly, you might log onto the network from a command prompt, from a predefined script inside a batch file, or from a window presented by the client software. If you attempt to reproduce a problem by performing different functions than those employed by the user, you may not be able to reproduce a legitimate problem and thus might assume that the symptoms resulted from user error. In fact, you may be missing a crucial clue to solving the problem.

To reproduce a symptom reliably, ask the user precisely what she did before the error appeared. For example, if a user complains that her network connection mysteriously drops when she's in the middle of surfing the Web, you should try to replicate the problem at her workstation; also, find out what else was running on the user's workstation or what kind of Web sites she was surfing.



Use good judgment when attempting to reproduce problems. In some cases, reproducing a problem could wreak havoc on the network, its data, and its devices; you should not attempt to reproduce such a problem. An obvious example involves a power outage in which your backup power source failed to supply power. After your network equipment comes back online, you would not want to try cutting the power again simply to verify that the problem derived from a faulty backup power source.

Verify Physical Connectivity

After you have reproduced the problem's symptoms, you should examine the most straightforward potential flaw in network communications—the physical connectivity. Physical connectivity may include the cabling from workstation or server to data jack, from data jack to punch-down block, from punch-down block to patch panel, or from patch panel to hub or switch. It may also include the proper physical installation of devices such as NICs, hubs, routers, servers, and switches. As noted earlier, you can save much time by checking the obvious first. Physical connectivity problems can be easy to spot and easy to fix.

Answering the following questions may help you identify a problem pertaining to physical connectivity:

- Is the device turned on?
- Is the NIC properly inserted?
- Is a device's network cable properly (that is, not loosely) connected to both its NIC and the wall jack?
- Do patch cables properly connect punch-down blocks to patch panels and patch panels to hubs or switches?
- Is the hub, router, or switch properly connected to the backbone?
- Are all cables in good condition (without signs of wear or damage)?
- Are all connectors (for example, RJ-45) in good condition and properly seated?
- Do network (maximum and segment) lengths conform to the IEEE 802 specifications?



A first step in verifying the physical integrity of a connection is to follow that connection from one endpoint on the network to the other. For example, if a workstation user cannot log onto the network, and you have verified that he is typing his password correctly, check the physical connectivity from his workstation's NIC and patch cable. Follow his connection all the way through the network to the server that he cannot reach.

Often, physical connectivity problems will manifest as a continuous or intermittent inability to connect to the network and perform network-related functions. Physical connectivity problems do not typically (but occasionally can) result in application anomalies, the inability to use a single application, poor network performance, protocol errors, software licensing errors, or software usage errors. Some software errors, however, can point to a physical connectivity problem. For example, a user might be able to log onto his file server without problems. When he chooses to run a query on a database, however, his report software might produce an error message indicating that the database is unavailable or not found. If the database resides on a separate server, this symptom could point to a physical connectivity problem with the database server.

In addition to verifying the connections between devices, you must verify the soundness of the hardware used in those connections. A sound connection means that cables are inserted firmly in ports, NICs, and wall jacks; NICs are seated firmly in the system board; connectors are not broken; and cables are not damaged. Damaged or improperly inserted connectivity elements may result in only occasional (and therefore difficult-to-troubleshoot) errors.

For example, you might receive a call from a user who cannot log onto the network in two out of every five attempts. The user might say that she could previously log onto the network without errors and that she thinks the errors have recently become more frequent. Because the error doesn't occur every time, it is probably caused by damaged or improperly installed connectivity hardware or by a segment length that exceeds IEEE 802 specifications. Because the errors are increasing in frequency, they are probably caused by hardware that is sustaining progressively more damage and will eventually fail. Assuming that no one else in this user's department is receiving similar errors, you might examine the cable connecting the user's workstation to the wall jack. Quite possibly, a chair rolling over it could damage this cable.

Even if a cable does not show obvious physical damage, it may still have flaws. For example, it might have been poorly manufactured or damaged internally from age or misuse. If you suspect a flawed cable, the quickest way to test your theory may be to replace the cable and note whether the errors disappear. Alternately, you could use a cable tester to verify the quality of a cable. You will learn more about cable testers later in this chapter.

Other physical components (such as NICs, hubs, or ports on any device) may also have flaws. Often, you can perform diagnostics on the device to determine whether it works correctly. For example, in Chapter 6, you learned that most NIC manufacturers ship a diagnostics program on a floppy disk with the NIC. In some cases, you may need to replace (or "swap out") a part. Later in this chapter, you will learn about the techniques and potential hazards of swapping equipment.

Finally, if symptoms seem to point to a physical connectivity problem, but you cannot find any loose or missing connections or flawed cables, the problem may relate to a network segment whose length exceeds IEEE 802 standards. Recall from Chapter 4 that the different types of networks must adhere to maximum segment lengths. For example, a 10BaseT network segment (the total amount of cabling between the connectivity device and a node) cannot exceed 100 meters. If your segment spans a greater distance, the devices at the end of the segment will experience intermittent connectivity errors or excessive transmission delays. If you have exceeded the maximum segment length, you must rearrange that segment to bring devices closer to the connectivity equipment.

The flowchart in Figure 12-4 illustrates how a logical approach to checking physical connectivity can help you solve a network problem. The steps in this flowchart apply to a typical problem: a user's inability to log onto the network. They assume that you have already ruled out user error and that you have successfully reproduced the problem under both your and the user's login IDs.

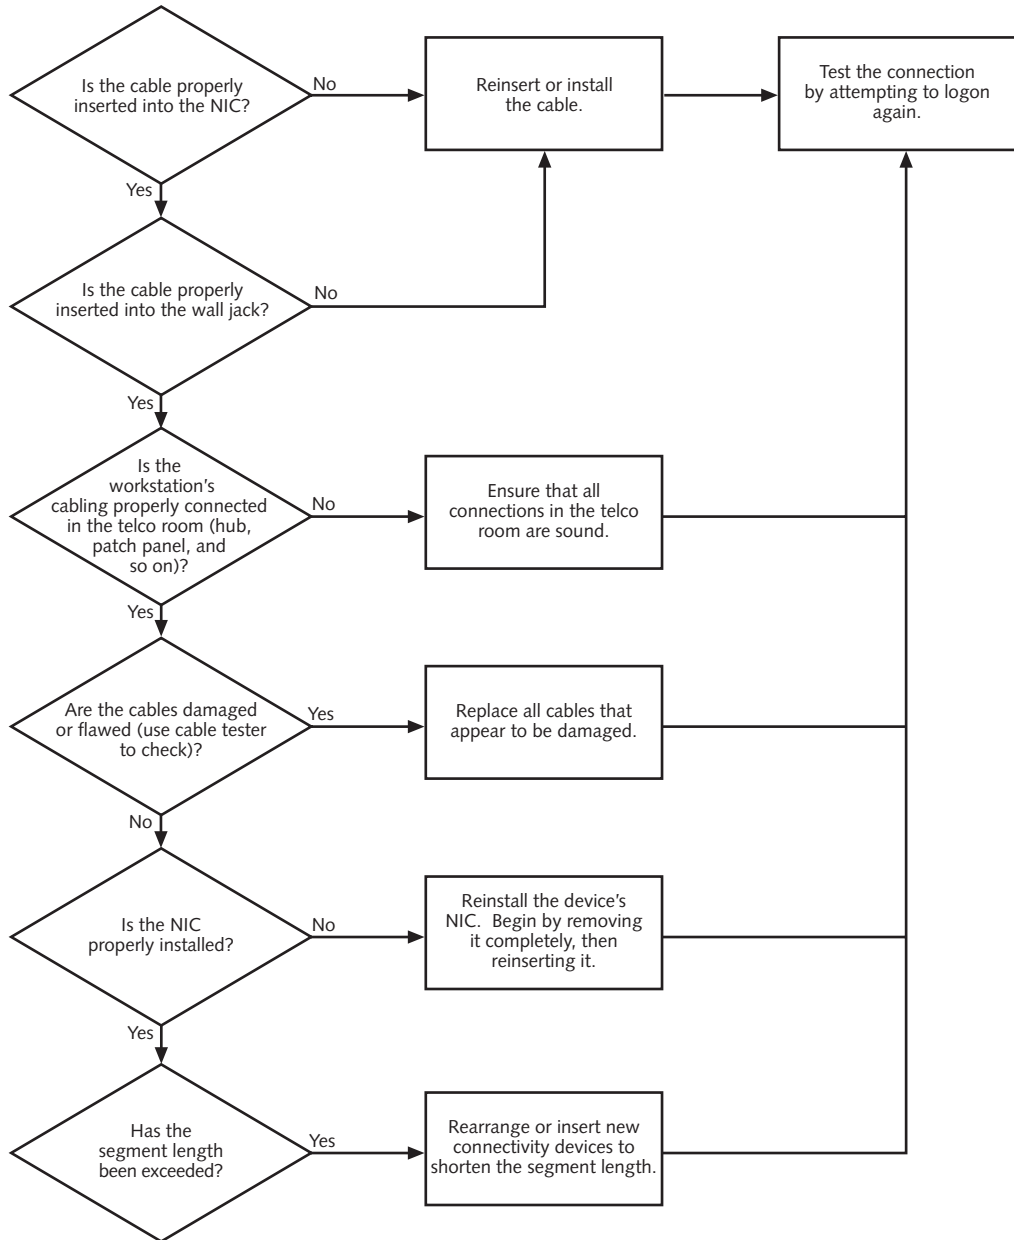


Figure 12-4 Troubleshooting while verifying physical connectivity



As noted in Figure 12-4, physical connectivity errors can frequently be traced to recent changes in the network, such as a replaced hub or a moved server. If you suspect a physical connectivity problem, you should find out whether anything on the network has changed recently. The potential effect of changes on network integrity is covered in detail later in this section. Most modern NICs have at least one LED that flashes green or amber, indicating the NIC's status. Although the meaning and number of these lights may vary according to the NIC model, typically a steady green light indicates that the NIC has successfully connected to the network. The LED will usually blink as the NIC searches for and finds a network connection. A steady blinking amber light generally means that the NIC can't make a network connection. For specific information on your NIC's LEDs, read the NIC's user manual.

Verify Logical Connectivity

Once you have verified the physical connections, you must examine the firmware and software configurations, settings, installations, and privileges. Depending on the type of symptoms, you may need to investigate networked applications, the network operating system, or hardware configurations, such as NIC IRQ settings. All of these elements belong in the category of “logical connectivity.”

Answering the following questions may help you identify a problem with logical connectivity:

- Do error messages reference damaged or missing files or device drivers?
- Do error messages reference malfunctioning or insufficient resources (such as memory)?
- Has an operating system, configuration, or application been recently changed, introduced, or deleted?
- Does the problem occur with only one application or a few, similar applications?
- Does the problem happen consistently?
- Does the problem affect a single user or one group of users?

Logical connectivity problems often prove more difficult to isolate and resolve than physical connectivity problems because they can be more complex. For example, a user might complain that she has been unable to connect to the network for the last two hours. After you go to her workstation and find that you can reproduce the symptoms both under her login ID and your own ID, you check the physical connections. Everything seems to be in order. Next, you may ask the user whether anything changed on her machine approximately two hours ago. She tells you that she didn't do a thing to the machine—it just stopped working.

At this point, you may investigate the workstation's logical connectivity. Some possible software-based causes for a failure to connect to the network include (but are not limited to) the following: resource conflicts with the NIC's configuration, an improperly configured NIC (for example, it may be set to the wrong data rate), improperly installed or configured client software, and improperly installed or configured network protocols or services. In this example, you may take another look at the client login screen and notice that the wrong server is selected as the default. Once you change the default server setting in the user's client software, she will likely be able to log onto the network.

Like many physical connectivity problems, many logical connectivity problems are created by changes to network elements. In the next section, you will learn how to trace the symptoms of a problem to a recent change in the network.

Implement a Solution

At last, after you have found the problem, you can implement a solution. This step may be very brief (such as correcting the default server designation in a user's client login screen) or it may take a long time (such as replacing the hard disk of a server). In either event, record your solution in a central location, such as a call-tracking database. You will learn more about documenting problems and solutions later in this chapter.

Implementing a solution requires foresight and patience, whether it consists of talking a user through changing a setting in his e-mail program or reconfiguring a router. As with finding the problem, the more methodically and logically you can approach the solution, the more efficient the correction process will be. If a problem is causing catastrophic outages, however, you should solve the problem as quickly as possible.

The following steps will help you implement a safe and reliable solution:

1. Collect all the documentation you have about a problem's symptoms from your investigation and keep it handy while solving the problem.
2. If you are reinstalling software on a device, make a backup of the device's existing software installation. If you are changing hardware on a device, keep the old parts handy in case the solution doesn't work. If you are changing the configuration of a program or device, take the time to print out the program or device's current configuration. Even if the change seems minor, jot down notes about the original state. For example, if you intend to add a user to a privileged group to allow her to access the Accounting spreadsheets, first write down the groups to which she currently belongs.
3. Perform the change, replacement, move, or addition that you believe will solve the problem. Record your actions in detail so that you can later enter the information into a database.
4. Test your solution (see the following section).

5. Before leaving the area in which you were working, clean it up. For instance, if you created a new patch cable for a telco room, remove the debris from splicing the cable.
6. If the solution fixes the problem, record the details you have collected about the symptoms, the problem, and the solution in your organization's call tracking database.
7. If your solution involved a significant change or addressed a significant problem (one that affected more than a few users), revisit the solution a day or two later to verify that the problem has, indeed, been solved and that it hasn't created additional problems.

Test the Solution

After implementing your solution, you must test it to verify that it works properly. Obviously, the type of testing you perform depends on your solution. For example, if you replaced a patch cable between a hub port and a patch panel, a quick test of your solution would be to determine whether you could connect to the network from the device that relies on that patch cable. If the device does not successfully connect to the network, you may have to try another cable or reconsider whether the problem stems from physical or logical connectivity or some other cause.

Suppose you replaced a switch that served four different departments in an organization. To test your solution, you might not only test connectivity from each department's workstations, but also use a network analysis tool (such as those discussed later in this chapter) to verify that the switch is handling data correctly.

It's often a good idea to enlist the user who reported the problem in testing your solution, too. That strategy ensures that you will get an objective assessment of the results. You may have been working on the solution so long that you've forgotten the original problem. You might also have enough technical knowledge to circumvent small problems that might flummox the average user. In addition, having the user test your solution will prevent you from leaving a device in a state that is familiar to you, but unfamiliar to the user.

For example, in the process of diagnosing a problem with a user's access to a mail directory, you may have reconfigured his mail settings to log on with your own ID and rule out the possibility of a physical connectivity error. After discovering that the problem was actually due to an IP addressing conflict, you may fix the IP addressing problem but forget that you changed the user's e-mail configuration. Having the user test your solution would reveal this oversight—and prevent you from having to return to the workstation to solve another problem.

You may not be able to test your solution immediately after implementing it. In some cases, you may have to wait days or weeks before you know for certain whether it worked. For example, you may have discovered that a server was sometimes running out of processor capacity when handling clients' database queries, causing users to experience unacceptably slow response times. To solve this problem, you might add two processors and reconfigure the server to use symmetric multiprocessing. The timing of the database usage may be unpredictable, however. As a result, you may not find out whether the added processors eliminated the problem until a certain number of users attempt the operations that will push the server to its peak processor usage.



A copy of all questions included in the preceding sections appears on a form in Appendix D, "Examples of Standard Networking Forms." You might want to create your own form based on these questions but tailored to your particular networking environment. Take your form along whenever you set out on a troubleshooting mission. It will help remind you of possibilities that you might otherwise forget to investigate.

Recognize the Potential Effects of a Solution

Even before fixing a problem on your network, you should consider how the change might affect users and network functionality. Consider the scope, tradeoffs, security, scalability, and cost when implementing a solution. These factors are discussed further in the following section.

One of the most important aspects to consider is the breadth, or scope, of your change. For example, replacing a cable that connects a workstation to a hub may affect only one user, but replacing a cable that connects a server to a hub will affect all users who access that server. Assess the scope of your solution—whether it is a single workstation, a workgroup, a location, or the entire network—before implementing that solution. If the problem does not pose an emergency, wait until no one is on the network before implementing solutions that will affect many users. That way, you will have time to assess the solution's effects systematically and fix any new problems that might arise.

Along with the scope, another factor to consider is the tradeoff your solution might impose. In other words, your solution may restore functionality for one group of users, but remove it for others. For example, let's say you are a network technician at a stationery company that uses specialized software to program custom logos and control its embossing machines. When you add a group of new Windows 2000 workstations to your network, you discover that the embossing control software doesn't work properly with them. The software vendor tells you that to be compatible with Windows 2000, you must install a new version of the software on your file server. You may be thrilled to hear of such a simple solution and begin to install the software immediately. In the next half hour, you receive numerous phone calls from employees using Windows 98 workstations

who cannot properly use the embossing control software. Now you have solved one problem, but created another. In this situation, it would have been wise to ask the software vendor about their upgrade's compatibility with all the other operating systems your company uses. If the vendor told you about a problem with Windows 98 workstations, you could have kept the old installation on the server for these users, then installed the new version of the software in another directory for use by Windows 2000 users.

Be aware of the security implications of your solution, because it may inadvertently result in the addition or removal of network privileges for a user or group of users. The consequence may be simply that a user can no longer access a data file or application he is used to accessing. But a worse consequence is that you could create a security opening that allows unauthorized people to access your network.

You should also consider the scalability of the solution you intend to implement. Does it position the network for additions and enhancements later on, or is it merely a temporary fix that the organization will outgrow in a year? Ideally, your solution would be perfectly suited to your network and allow for future growth. But a temporary fix is not necessarily wrong, depending on the scenario. For example, you might walk into the office one day to find that none of your users can access the network. You may track down the problem as an internal hardware problem with your IP gateway. Since the gateway is under warranty, you quickly call the manufacturer to either get the gateway replaced or fixed immediately. The manufacturer may tell you that while they don't have the identical gateway available in their local office, they can substitute a different, smaller model to get your users reconnected today and meanwhile order the identical gateway that you can install when you have more time. In this situation, it is preferable to take the temporary gateway and restore functionality than to wait for the ideal solution.

Another factor to consider when implementing your solution is cost. Obviously, replacing one patch cable or faulty network adapter is a fairly inexpensive proposition, and you don't need to analyze cost in these cases. But if the solution you have proposed requires significant dollars for either software or hardware, you should spend time carefully considering your options. For example, you may discover a problem with performance on your network. After some investigation you may determine that the best solution is to replace all of your 400 workstations' network adapters with newer, faster network adapters. If you purchase quality NICs, this solution could cost over \$10,000 for the hardware alone, not to mention the time it will take technicians to replace the devices, which may cost more. Also you should consider when these workstations will be replaced and if you will have to either discard or remove the network adapters you just installed. It may be more prudent to identify where the network's performance is poor and address those areas separately—for example, by adding a switch to a busy segment or adding a more powerful server for a heavily used application.

Last, if you are uncertain about whether your proposed solution is the *best* solution, even after your thorough diagnosis and research, you should consult with others, either within or outside of your organization. Colleagues or consultants may share an experience that leads you to prefer one solution to another.

After your solution is in place, communicate your solution to your colleagues, thus adding to the store of knowledge about your network. Next you will learn about how best to document your troubleshooting efforts.

Document Problems and Solutions

Whether you are a one-person network support team or one of a hundred network technicians at your organization, you should always write down the symptoms of a problem and your solution for it. Given the volume of problems you and other analysts will troubleshoot, it will be impossible to remember the circumstances of each incident. In addition, networking personnel frequently change jobs, and everyone will appreciate clear, thorough documentation. An effective way to document problems and solutions is in a centrally located database to which all networking personnel have online access.

Some organizations use a software program for documenting problems, known as a **call tracking system** (also informally known as help desk software). Examples of popular call tracking systems include Clientele, Expert Advisor, ServiceIT, and Track-It! These programs provide user-friendly graphical interfaces that prompt the user for every piece of information associated with the problem. They assign unique identifying numbers to each problem, in addition to identifying the caller, the nature of the problem, the time necessary to resolve it, and the nature of the resolution.

Most call tracking systems are highly customizable, so you can tailor the form fields to your particular computing environment. For example, if you work for an oil refinery, you might add fields for identifying problems with the plant's flow-control software. In addition, most call tracking systems allow you to enter free-form text explanations of problems and solutions. Some also offer Web-based interfaces.

If your organization does not have a call tracking system, you should at least keep records in a simple electronic form. You can find an example of a network problem record in Appendix D, "Examples of Standard Networking Forms." A typical problem record form should include at least the following fields:

- The name, department, and phone number of the problem originator (the person who first noticed the problem)
- Information regarding whether the problem is software- or hardware-related
- If the problem is software-related, the package to which it pertains; if the problem is hardware-related, the device or component to which it pertains

- Symptoms of the problem, including when it was first noticed
- The name and telephone number of the network support contact
- The amount of time spent troubleshooting the problem
- The resolution of the problem

As discussed earlier in this chapter, many organizations operate a help desk staffed with personnel who have only basic troubleshooting expertise and who record problems called in by users. To effectively field network questions, an organization's help desk staff must maintain current and accurate records for network support personnel. Your department should take responsibility for managing a supported services list that help desk personnel can use as a reference. A **supported services list** is a document (preferably online) that lists every service and software package supported within an organization, plus the names of first- and second-level support contacts for those services or software packages. Anything else you or your department can do to increase communication and availability of support information will expedite troubleshooting.

In addition to communicating problems and solutions to your peers whenever you work on a network problem, you should follow up with the user who reported the problem. Make sure that the client understands how or why the problem occurred, what you did to resolve the problem, and who to contact should the problem recur. This type of education will not only help your clients make better decisions about the type of support or training they need, but will also improve their understanding of and respect for your department.

TROUBLESHOOTING TOOLS

So far, this chapter has focused on using a systematic method of trial and error to diagnose network problems. In the real world, however, this technique may lead nowhere or take too much time. In some cases, the most efficient approach is to use a tool specifically designed to analyze and isolate network problems. Several tools are available, ranging from simple cable testers that indicate whether a cable is faulty, to sophisticated protocol analyzers that capture and interpret all types of data traveling over the network. The tool you choose will depend on the particular problem you need to investigate and the characteristics of your network.

The following sections describe a variety of network troubleshooting tools, their functions, and their relative costs. In the Hands-on Projects at the end of this chapter, you will have the opportunity to try some of these network troubleshooting tools.

Hardware Troubleshooting Tools

This section describes tools that can assist you in identifying a problem with a cable, connector, or network adapter.

Crossover Cable

As you learned in Chapter 4, a crossover cable is one in which the transmit and receive wire pairs in one of the connectors are reversed. This reversal enables you to use a crossover cable to directly interconnect two nodes without using an intervening connectivity device such as a hub. A crossover cable is useful in troubleshooting to quickly and easily verify that a node's network adapter is transmitting and receiving signals properly. For example, suppose you are a network technician on your way to fix urgent network problems. A user flags you down and says that over the last week he has occasionally had problems connecting to the network and as of this morning, he hasn't been able to connect at all. He's very frustrated, so you kindly say that if you can help him in 10 minutes, you will; otherwise, he'll have to call the help desk. You follow him to his workstation and, by asking around, you determine that he is the only one suffering this problem. Thus, you can probably narrow the problem down to his workstation (either hardware or software) or his cabling (or less likely, his port on the hub in the telecommunications closet). Because you have your laptop and troubleshooting gear in your bag, you quickly connect one plug of the crossover cable to his workstation's network adapter and the other plug to your laptop's network adapter. You then try logging onto your laptop from his workstation. Because this process is successful, you suggest that the problem lies with his network cable, and not with his workstation's software or hardware. As you rush off, you hand him a new patch cable to replace his old one.

Tone Generator and Tone Locator

Ideally, you and your networking colleagues would label each port and wire termination in a telecommunications closet so that problems and changes can be easily managed. However, because of personnel changes and time constraints, a telecommunications closet often winds up being disorganized and poorly documented. If this is the case where you work, you may need a tone generator and a tone locator to determine where one pair of wires (out of possibly hundreds) terminates.

A **tone generator** is a small electronic device that issues a signal on a wire pair. A **tone locator** is a device that emits a tone when it detects electrical activity on a wire pair. By placing the tone generator at one end of a wire and attaching a tone locator to the other end, you can verify the location of the wire's termination. Figure 12-5 depicts the use of a tone generator and a tone locator. Of course, you must work by trial and error, guessing which termination corresponds to the wire over which you've generated a signal until the tone locator indicates the correct choice. This combination of devices is also known as a **fox and hound**, because the locator (the hound) chases the generator (the fox).




Figure 12-5 Use of a tone generator and tone locator

Tone generators and tone locators cannot be used to determine any characteristics about a cable, such as whether it has defects or whether its length exceeds IEEE standards for a certain type of network. They are only used to determine where a wire pair terminates. In fact, because of their limited functionality, tone generators and tone locators are rarely used on modern networks. (However, they are still widely used by telephone technicians.)



A tone generator should never be used on a wire that may connect to a device's port or network adapter. Because a tone generator transmits electricity over the wire, it may damage the device or network adapter.

Multimeter

Cable testing tools are essential for both cable installers and network troubleshooters, as cables are often at fault when a network problem arises. Symptoms of cabling problems can be as elusive as occasional lost packets or as obvious as a break in network connectivity. You can easily test cables for faults with specialized tools. In this section and in the ones following, you will learn about different tools that can help isolate problems with network cables. The first device you will learn about is a **multimeter**, a simple instrument that can measure many characteristics of an electric circuit, including its resistance and voltage.

If you have taken any introductory electronics classes, you are probably familiar with a **voltmeter**, the instrument that measures the pressure, or voltage, of an electric current. Recall that voltage is used to create signals over a network wire. Thus, every time data travel over a wire, the wire carries a small voltage. In addition, each wire has a certain amount of **resistance**, or opposition to electric current. Resistance is a fundamental

property of wires that depends on the wire's molecular structure and size. Every type of wire has different resistance characteristics (for example, each type of coaxial cable listed in Table 4-2 has a different amount of resistance). Resistance is measured in ohms, and the device used to measure resistance is called an **ohmmeter**.

Although electricians and network professionals could use separate instruments for measuring resistance and voltage on a wire, it is more convenient to have one instrument that accomplishes both of these functions. The multimeter is such an instrument. Figure 12-6 shows a multimeter.




Figure 12-6 A multimeter

As a network professional, you might use a multimeter to:

- Verify that a cable is properly conducting electricity—that is, whether its signal can travel unimpeded from one node on the network to another
- Check for the presence of noise on a wire (by detecting extraneous voltage)
- Verify that the amount of resistance generated by terminators on coaxial cable networks (such as 10Base5 Ethernet) is appropriate or whether terminators are actually present and functional
- Test for short or open circuits in the wire (by detecting unexpected resistance or loss of voltage)

You should be aware that multimeters are at the low end of the cable testing tool spectrum because of their limited capabilities. More sophisticated tools, such as cable testers, can perform the same tests that multimeters perform, in addition to other, more network-specific, functions.

Cable Checkers

Basic **cable checkers** simply determine whether your cabling can provide connectivity. To accomplish this task, they apply a small voltage to each conductor at one end of the cable, and then check whether that voltage is detectable at the other end. They may also check whether voltage cannot be detected on other conductors in the cable. Figure 12-7 depicts a typical simple cable checker.



Figure 12-7 A basic cable checker

Most cable checkers provide a series of lights that signal pass/fail. Some also indicate a cable pass/fail with an audible tone. A pass/fail test provides a simple indicator of whether a component can perform its stated function.

In addition to checking cable continuity, a good cable checker will verify that the wires are paired correctly and that they are not shorted, exposed, or crossed. Recall from Chapter 4 that different network models use specific wire pairings and follow cabling standards set forth in EIA/TIA 568. Make sure that the cable checker you purchase can test the type of network you use—for example, 10BaseT Ethernet, 100BaseTX Ethernet, or Token Ring.

When you make your own cables, be sure to verify their integrity with at least a cable checker (better yet, a cable tester). Even if you purchase cabling from a reputable vendor, you should make sure that it meets your network's required standards. Just because a cable is labeled "CAT5" does not necessarily mean that it will live up to that standard. Testing cabling before installing it may save many hours of troubleshooting after the network is in place.

Cable checkers cannot test the continuity of fiber-optic cabling, because fiber cable uses light rather than voltage to transmit data. To test fiber-optic cabling, you need a specialized fiber cable tester.



Do not use a cable checker on a live network cable. Disconnect the cable from the network, and then test its continuity.

For convenience, most cable checkers are portable and lightweight and typically use one 9-volt battery. A basic cable checker costs between \$100 and \$300, but it may save many hours of work. Popular cable checker manufacturers include Belkin, Fluke, Microtest, and Paladin.

Cable Testers

The difference between cable checkers and cable testers lies in their sophistication and price. A **cable tester** performs the same continuity and fault tests as a cable checker, but also provides the following functions:

- Ensures that the cable is not too long
- Measures the distance to a cable fault
- Measures attenuation along a cable
- Measures near-end crosstalk between wires
- Measures termination resistance and impedance for Thinnet cabling
- Issues pass/fail ratings for CAT3, CAT5, CAT6, or even CAT7 standards
- Stores and prints cable testing results

Some cable testers may provide even more features—for example, a graphical output depicting a cable's attenuation and crosstalk characteristics over the length of the cable. Because of their sophistication, cable testers cost significantly more than cable checkers. A high-end unit may cost from \$5000 to \$8000, and a low-end unit may cost between \$1000 and \$4000. Popular cable tester manufacturers include Fluke and Microtest. Figure 12-8 shows an example of a high-end cable tester.

When choosing a cable tester for twisted-pair networks, make sure to purchase one that performs attenuation and crosstalk testing for the frequency range used by your network. For example, if you want to test a 100BaseT Ethernet network, purchase a cable tester capable of testing up to 100 MHz.



Figure 12-8 A high-end cable tester

To better appreciate how many problems a good cable tester can diagnose, recall from Chapter 4 that network segments must adhere to strict length limits to ensure that data reach their destinations on time and error-free. If one room of workstations continually experiences intermittent problems logging onto the network or very slow connections, you could use a cable tester to discover whether those workstations are situated beyond their maximum distance from the network hub. If another group of workstations frequently experiences slow responses from the network, a cable tester might reveal the presence of too many stations between the sending and receiving nodes, which causes excessive signal attenuation.

Another significant factor in wire-based data transmission is crosstalk. Recall from Chapter 4 that crosstalk occurs when the signals on one wire interfere with signals on an adjacent wire. The result is interference, much in the same way that the voices from two conversations in a loud room interfere with each other and prevent listeners from understanding the words. Crosstalk often arises when wires are crushed or crossed at the connector end of a cable. For this reason, you can accurately test for crosstalk only after installation of a cable, and you should perform the test at both ends of the wire.

In addition to cable testers for coaxial and twisted-pair networks, you can also find cable testers for fiber-optic networks. Rather than issue an electrical signal over the cable as twisted-pair cable testers do, a fiber-optic cable tester transmits light-based signals of different wavelengths over the fiber. These tests can indicate the amount of attenuation on the cable, and the continuity and the length of the cable. Note that since crosstalk does not apply to light-based signals, a fiber tester cannot (and need not) test for crosstalk. Because of the relatively high cost of installing fiber-optic cable, you should use a fiber tester on your cable before you install it, as well as after you install it.

Time Domain Reflectors (TDRs)

A **time domain reflector (TDR)** is a high-end instrument for testing the qualities of a cable. It works by issuing a signal on a cable and measuring the way the signal bounces back (or reflects) to the TDR. Connectors, crimps, bends, short circuits, cable mismatches, or other defects modify the signal's amplitude before it returns to the TDR, thus changing the way it reflects. The TDR then accepts and analyzes the return signal, and based on its condition and the amount of time the signal took to return, determines cable imperfections. In the case of a coaxial cable network, a TDR can indicate whether terminators are properly installed and functional. A TDR can also indicate the distance between nodes and segments.

As with cable testers, time domain reflectors are also made for fiber-optic networks. Such instruments are called **optical time domain reflectors (OTDRs)**. Rather than issuing an electrical signal, OTDRs issue a light-based signal over the fiber. Based on the type of return light signal, the OTDR can accurately measure the length of the fiber, determine the location of faulty splices, breaks, connectors, or bends, and measure attenuation over the cable.

Because some loss of a signal is expected with the addition of nodes and connectors, TDRs are a good way of taking a baseline measurement for your network cabling. A **baseline** is a record of how well the network operates under normal conditions (including its performance, collision rate, utilization rate, and so on). Baselines are used for comparison when conditions change. A TDR can provide a baseline for the characteristics and performance of a network's cable infrastructure. Then later, if you suspect cabling problems, you can use the TDR and compare your new results with your baseline measurement to ascertain whether signaling characteristics have changed.

Software Troubleshooting Tools

As noted earlier, once you have ruled out user error and physical connectivity problems (including faulty cabling) in your troubleshooting, a more in-depth analysis of the network may be necessary. Software-based tools that enable you to analyze network traffic include NOS log files, network monitors, and network analyzers. While log files can reveal what has happened on a server, network monitors and analyzers can capture and interpret data traveling across the network.

Network Monitors

A **network monitor** is usually a software-based tool that continually monitors traffic on the network from a server or workstation attached to the network. Network monitors typically can interpret up to Layer 3 of the OSI Model. They can determine the protocols passed by each packet, but can't interpret the data inside the packet. By capturing data they can provide either a snapshot of network activity at one point in time or a historical record of network activity over a period of time.

Network monitoring tools are generally less expensive than network analyzers (discussed next) and may be included in your network operating system software. In the following sections, you will learn about two tools that can be part of your network operating system: Microsoft's Network Monitor (which ships with Windows NT Server version 4.0 or Windows 2000) and Novell's LANalyzer agent (which is bundled with Novell's ManageWise software package). These packages actually blur the distinction between network monitors and network analyzers, because they provide some of the same functionality as high-end protocol analyzers. In addition, you will learn about network analyzers, such as Network Associates' Sniffer Portable software, and sniffer hardware. Once you have worked with one network monitoring or analyzing tool, you will find that other products work in much the same way. Most even use very similar graphical interfaces.



To take advantage of software-based network monitoring and analyzing tools, the network adapter installed in your machine must support promiscuous mode. In **promiscuous mode**, a device driver directs the network adapter card to pick up all frames that pass over the network—not just those destined for the node served by the card. You can determine whether your network adapter supports promiscuous mode by reading its manual or checking with the manufacturer. Some network monitoring software vendors may even suggest which network adapters to use with their software.

Before adopting a network monitor or analyzer, you should be familiar with some of the data errors that these tools can distinguish. The following list defines some commonly used terms for abnormal data patterns and packets, along with their characteristics:

- **Local collisions**—Collisions that occur when two or more stations are transmitting simultaneously. A small number of collisions are normal on an Ethernet network. Excessively high collision rates within the network usually result from cable or routing problems.

- **Late collisions**—Collisions that take place outside the window of time in which they would normally be detected by the network and redressed. Late collisions are usually caused by one of two problems: (1) a defective station (for example, a card or transceiver) that is transmitting without first verifying line status, or (2) failure to observe the configuration guidelines for cable length, which results in collisions being recognized too late.
- **Runts**—Packets that are smaller than the medium's minimum packet size. For instance, any Ethernet packet that is smaller than 64 bytes is considered a runt. Runts are often the result of collisions.
- **Giants**—Packets that exceed the medium's maximum packet size. For example, any Ethernet packet that is larger than 1518 bytes is considered a giant.
- **Jabber**—A device that handles electrical signals improperly, usually affecting the rest of the network. A network analyzer will detect a jabber as a device that is always retransmitting, effectively bringing the network to a halt. A jabber usually results from a bad NIC. Occasionally, it can be caused by outside electrical interference.
- **Negative frame sequence checks**—The result of the cyclic redundancy checksum (CRC) generated by the originating node not matching the checksum calculated from the data received. It usually indicates noise or transmission problems on the LAN interface or cabling. A high number of negative CRCs usually result from excessive collisions or a station transmitting bad data.
- **Ghosts**—Frames that are not actually data frames, but aberrations caused by a repeater misinterpreting stray voltage on the wire. Unlike true data frames, ghosts have no starting delimiter.

Microsoft's Network Monitor (NetMon) Microsoft's **Network Monitor (NetMon)** is a software-based network monitoring tool that comes with Windows NT Server 4.0 and Windows 2000. It offers the following capabilities:

- Capturing network data traveling from one or many segments
- Capturing frames sent by or to a specified node
- Reproducing network conditions by transmitting a selected amount and type of data
- Detecting any other running copies of NetMon on the network (depending on the placement and configuration of routers)
- Generating statistics about network activity

Probably NetMon's most useful capability is capturing data as it travels across the network. As with hardware-based network analyzers, you can instruct NetMon to pay attention to the network for a period of time and to capture all data that travel across the particular segment. (Because NetMon takes advantage of promiscuous mode, it captures all data—not just data to or from the NetMon console.)



If you completed the Hands-on Projects in Chapter 5, you had an opportunity to experiment with Network Monitor. See Figure 5-31 for a view of Network Monitor's interface as it captures network traffic.

How can capturing data help you solve a problem? Imagine that traffic on a segment of the network you administer suddenly grinds to a halt one morning at about 8:00. You no sooner step in the door than everyone from the help desk calls to tell you how slowly the network is running. Nothing has changed on the network since last night, when it ran normally, so you can think of no obvious reasons for problems. You suspect a faulty NIC on one workstation is using network bandwidth by continually transmitting bad packets.

At the workstation where you have previously installed NetMon, you capture all data transmissions for approximately five minutes. You can then sort out the erroneous frames in NetMon, arranging the nodes in order based on how many bad packets each has generated. If your suspicion is correct, the workstation at the top of the list will be the culprit, generating significantly more bad data transmissions than any other node.

Novell's LANalyzer Novell provides a network monitoring tool that is similar to Microsoft's Network Monitor, called the **LANalyzer** agent. It can act as a standalone program on a Windows 9x or 2000 workstation or as part of the ManageWise suite of network management tools on a NetWare server. LANalyzer performs the following functions:

- Initially discovering all network nodes on a segment
- Continuously monitoring network traffic
- Tripping alarms when traffic conditions meet preconfigured thresholds (for example, if usage exceeds 50% of capacity)
- Capturing traffic to and from all or selected nodes

Like Network Monitor, LANalyzer enables you to capture traffic, identify data errors by node, and generate traffic statistics by segment. In addition, as part of the ManageWise suite, the LANalyzer agent can poll the network to find all nodes on a particular segment. It can use this data to build a network management system that can gather more than simple traffic information—for example, discovering how many times a user has logged on at a certain workstation or noting what kind of programs a workstation typically requests from the server.

LANalyzer can also provide real-time network statistics and send alert messages and/or sound alarms when network thresholds are reached. For example, to make sure that average network traffic never exceeds 50% of your network's capacity, you could configure LANalyzer to warn you when the average reaches 49%. If this warning occurs frequently on one segment of your network, you can take steps to redistribute the traffic or reinforce your network's capacity. Note that an average utilization means that LANalyzer would have to measure a 49% reading more than a single time; a single reading represents a **spike**. You can also customize the sensitivity of the triggers.

Network Analyzers

A **network analyzer** (also known as a **protocol analyzer**) is a tool that can capture traffic and analyze packets, typically all the way to Layer 7 of the OSI Model. For example, it can identify that a packet uses TCP/IP and, more specifically, that it is an ARP request from one particular workstation to a server. Analyzers can also interpret the payload portion of packets, translating from binary or hexadecimal code to human-readable form. As a result, network analyzers can capture passwords going over the network, if their transmission is not encrypted. Some network analyzer software packages can run on a standard PC, but others require PCs equipped with special network adapters and operating system software.

In addition to using the software that comes with the network operating system, you can purchase network analyzing software from vendors that specialize in products for network management. One popular example is Network Associates' **Sniffer Portable**, network analyzer software that provides data capture and analysis, node discovery, traffic trending, history, alarm tripping, and utilization prediction. Essentially, Sniffer Portable has the same features as Network Monitor and LANalyzer, plus a few extras. It can also generate traffic in an attempt to reproduce a network problem and monitor multiple network segments simultaneously. Its graphical interface makes this product very easy to use, readily revealing the traffic flow across the network. In addition, Sniffer Portable supports a multitude of protocols and network topologies.

One advantage to using a network monitor or analyzer that is not part of the network operating system relates to mobility. With Sniffer Portable software installed on your laptop, for instance, you can roam from one network segment to another, analyzing traffic without having to install multiple network monitoring consoles. Hardware-based network analyzers, such as the sniffers discussed below, also offer the advantage of mobility.

Network Associates has also led the way in developing hardware-based network analyzers, known as sniffers. **Sniffers** usually resemble regular laptops, but are equipped with a special network adapter and network analysis software. The sole job of a sniffer is to analyze network problems. Unlike laptops that have a network monitoring tool installed,

sniffers typically cannot be used for other purposes, because they don't depend on a familiar desktop operating system such as Windows. They have their own, proprietary operating system (developed by Network Associates, for example). Because they do not rely on a desktop operating system such as Windows, hardware-based network analyzers have an advantage over network monitoring software. Because they do not rely on Windows device drivers (for the NIC), for example, they can capture information that the NIC would automatically discard, such as runt packets.

Sniffers offer a great deal of versatility in the type and depth of information they can reveal. The danger in using this type of tool is that it may collect more information than you or the machine can reasonably process, thus rendering your exercise futile. To avoid this problem, you should set filters on the data gathered. For example, if you suspect that a certain workstation is causing a traffic problem, you should filter the data collection to accept only packets to or from that workstation's MAC address. If you suspect that you have a gateway-related TCP/IP problem, you would set a filter to capture only TCP/IP packets and to ignore other protocols from the gateway's MAC address.

Sniffers are tailored to a particular type of network. For example, one sniffer may be able to analyze both Ethernet and Token Ring networks, but another sniffer may be necessary to analyze fiber or ATM networks. A sniffer represents a significant investment, with costs ranging from \$10,000 to \$30,000.



Recall from Chapter 6 that using a switch logically separates a network into several segments. If a network is fully switched (that is, if every node is connected to its own switch port), your network analyzer can capture only broadcast packets and packets destined for the node on which you're running the software, because those packets are the only ones that will travel through a switched environment. The increasing use of switches has made network monitoring more difficult, but not impossible. One solution to this problem is to reconfigure the switch to reroute the traffic so that your network analyzer can pick up all traffic. Obviously, you would want to weigh the disruptive effects of this reconfiguration against the potential benefits from being able to analyze the network traffic and solve a problem.

PRACTICAL TROUBLESHOOTING

You have learned about following a troubleshooting methodology and using specialized tools to diagnose network problems. You will acquire much of your troubleshooting expertise through experience. But if you don't yet have experience, you can get a head start by learning some practical tips and strategies for troubleshooting based on the experience of others. The following sections provide real-world techniques for network troubleshooting that do not neatly fit into a troubleshooting methodology.

Physical Layer Problems and Symptoms

By now you have probably realized that one symptom, such as a user not being able to log onto the network, could result from a number of possible problems. In addition to systematically following a troubleshooting methodology, you may discover a symptom's cause by identifying the OSI Model layer where it is occurring. That way, you can analyze connections, settings, or traffic within that layer and move closer to a solution. Figure 12-9 summarizes the services and devices that you have already learned about according to their OSI Model layer.

| | |
|---------------------|--|
| | <u>User</u> Programs |
| Application layer: | Program-to-(N)OS interaction |
| Presentation layer: | Text formatting, encryption, code conversion |
| Session layer: | Establishing, maintaining, coordinating connections |
| Transport layer: | Flow control, sequencing, acknowledgment |
| Network layer: | Logical addressing, routing (routers, layer 3 switches) |
| Data Link layer: | Framing and physical addressing (bridges, switches) |
| Physical layer: | Voltage detection, signaling (NICs, hubs, repeaters, cabling) |
| | Topology |

Figure 12-9 Services and devices in the OSI Model

By some estimates, more than half of all network problems occur at the Physical layer of the OSI Model, which includes cabling, network adapters, repeaters, and hubs. The Physical layer also controls signaling and the voltage levels used in signaling. Thus, RFI and EMI noise can cause network problems at the Physical layer. Because Physical layer faults are so common (and often easily fixed), you should be thoroughly familiar with the symptoms of such problems. The following list details some common Physical layer problems and their symptoms on an Ethernet network. It also suggests troubleshooting steps that you can follow to verify whether the problem indeed exists. Finally, it provides some possible solutions to consider once you have verified the problem.

- Segment Problems

Problem: Segment or network lengths that exceed the IEEE maximum standards (for example, an Ethernet 100BaseT segment that exceeds 100 meters)

Symptoms: An excessive number of late collisions. Users recognize this problem as intermittent difficulty connecting to the network or exchanging data over the network.

Troubleshooting hints: A scope limited to a geographical area or workgroup within the LAN could point to this problem. A protocol analyzer can help determine specifically which segments or nodes are experiencing late collisions. Observation (or relying on network documentation) can help determine which network lengths may exceed IEEE standard maximums.

Solution: Reconfigure the topology of the network to avoid excessive segment or network lengths.

- Signal noise

Problem: Noise affecting a signal (from EMI or RFI sources, improper grounding, or crosstalk)

Symptoms: Excessive number of packet errors such as runts, giants (in the case of improper grounding), and damaged frame check sequence fields, but no evidence of excessive collisions. Users recognize this problem as intermittent difficulty connecting to the network or exchanging data over the network.

Troubleshooting hints: The nodes or segments affected by the errors can be identified with the help of a protocol analyzer. A simple AM radio may be used to detect EMI near cables. Examine the cables' environment for noise sources. See if network problems disappear when those sources (for example, fluorescent lights or microwaves) are turned off.

Solution: Remove sources of EMI or RFI from cabling areas, encase cables in conduit, or reroute cabling. If this is not possible, consider changing the cable type to one with better resistance to noise. Ensure proper grounding on coaxial cable networks. Reduce crosstalk on twisted-pair networks by using wires with a higher twist ratio and making sure cables are not bundled too tightly.

- Cable damage

Problem: Damaged cables (for example, crimped, bent, nicked, or partially severed)

Symptoms: Excessive number of normal collisions or packet errors (such as giants and runts), but few late collisions. Users recognize this problem as frequent difficulty connecting to or exchanging data with the network, very poor network response time, or a complete inability to connect to the network (depending on the severity of the cable damage).

Troubleshooting hints: The scope of this problem may be a single user (in the case of a workstation patch cable) or a whole segment or network of users. Once you have identified a suspicious cable, a cable tester or cable checker can help determine the integrity and reliability of that cable. A protocol analyzer can indicate which nodes are experiencing excessive numbers of packet errors. A network monitor can indicate where traffic bottlenecks are occurring, in the case of a severely damaged cable.

Solution: Replace the faulty cable with a good cable.

- Connector flaws

Problem: Improper terminations, faulty connectors, loose connectors, or poorly crimped connections

Symptoms: Excessive number of normal collisions and packet errors (such as giants and runts), but few late collisions. Users will recognize this as frequent problems connecting to or exchanging data with the network, very poor network response time, or a complete inability to connect to the network (depending on the severity of the connector fault).

Troubleshooting hints: As with a faulty cable, the scope of this problem may be a single user or a whole segment or network of users. A simple PING test may help determine the location of the fault. A protocol analyzer can indicate which nodes are experiencing excessive numbers of packet errors. A network monitor can indicate where traffic bottlenecks are occurring, in the case of a very loose or completely faulty connector.

Solution: Replace the connector with a good connector, reseal the loose connector, or correct the termination error.

- Adapter flaws

Problem: Faulty network adapter

Symptoms: Some types of NIC faults result in an excessive number of packet errors (for example, giants, runts, or damaged frame check sequence fields), but no apparent increase in collisions; other types of NIC faults result in an excessive number of late collisions (when the NIC's carrier sense mechanism is not operating properly). Users will recognize either of these situations as intermittent problems connecting to the network or exchanging data over the network.

Troubleshooting hints: The scope of this problem is limited to the nodes that rely on the network adapter (for example, if it is a workstation NIC, only the workstation user should notice the problem; if it is a switch NIC, all switched connections will share the problem). A protocol analyzer can indicate which nodes are experiencing excessive numbers of late collisions. A network monitor can indicate which node is issuing bad packets (and therefore, which NIC is to blame).

Solution: Replace the faulty network adapter with a good network adapter (making sure they are identical or compatible models).

Staff Involved in Troubleshooting

Many staff members may contribute to troubleshooting a network problem. Often the division of duties is formalized, with a help desk acting as the first, single point of contact for users to call in regarding errors. A help desk is typically staffed with help desk analysts—people proficient in basic (but not usually advanced) workstation and network troubleshooting. Larger organizations may group their help desk analysts into teams based on their expertise. For example, a company that provides users with word-processing, spreadsheet, project planning, scheduling, and graphics software might assign different technical support personnel at the help desk to answer questions pertaining to each application.

The help desk analysts are often considered first-level support, because they provide the first level of troubleshooting. When a user calls with a problem, a help desk analyst typically creates a record for the incident and attempts to diagnose the problem. The help desk analyst may be able to solve a common problem over the phone within minutes by explaining something to the user. On other occasions, the problem may be rare or complex. In such cases, the first-level support analyst will refer the problem to a second-level support analyst. A second-level support analyst is someone who has specialized knowledge in one or more aspects of a network. For example, if a user complains that she can't connect to a server, and the first-level support person narrows the problem down to a failed file server, that first-level support analyst would then refer the problem to the second-level support person. Typically, first-level support analysts stay at the help desk while second-level support analysts are mobile.

In addition to having first- and second-level support analysts, most help desks include a help desk coordinator. The help desk coordinator ensures that analysts are divided into the correct teams, schedules shifts at the help desk, and maintains the infrastructure to enable analysts to better perform their jobs.

Most organizations also have an operations manager, who supervises the help desk coordinator. This person knows less about the day-to-day activities of the help desk, but works with the help desk coordinator to determine how to improve customer service and supply analysts with the needed infrastructure. For example, the operations manager may control the budget that provides help desk analysts with office space, call tracking software, a call distribution system, and any additional resources necessary to perform their jobs.

Examples of How to Investigate Problems

The following scenarios illustrate how to narrow down the cause of a network problem. Notice that all questions do not apply in all situations. You should use common sense to decide which questions apply to a particular situation and to interpret the answers you receive. In addition to reviewing the scenarios given here, you will have more opportunities to exercise your investigative and troubleshooting skills in the Case Projects at the end of this chapter.

Scenario 1: Unable to Access the Network

Perhaps one of the most common problems you'll address as a network troubleshooter is an inability to access the network. This problem can be caused by a variety of failures (either hardware or software) and situations (for example, user error or changes in the network infrastructure). If you receive notice of the problem from a user, rather than from your automated network monitoring system or a fellow computer professional, the initial information you receive may not be very helpful. Your conversation with the user might go something like this:

USER: I can't log onto the network.

YOU: When did the problem begin?

USER: Just this morning. I came into work and I couldn't log on. I really have to get my invoices done now, because they're due to my boss by 10:00 A.M.

YOU: As far as you know, are you the only person in your area who's having this problem?

USER: I think so.

YOU: And what kind of error message do you receive when you try to log on?

USER: It says something about the network being unavailable.

YOU: Let's check to make sure your network cable hasn't accidentally been pulled out or loosened.

USER: I checked it already, and I'm sure it's all right.

YOU: Well, humor me a little. I just want to rule out any possibility of a connection problem. Sometimes the janitors accidentally jar a connection loose when they clean the floors.

USER: OK. (Checks the connections according to your guidance.) Nope, they seem to be plugged in just fine.

YOU: All right, thanks for checking. Has anything changed on your computer in the last day? For example, did you have to add any programs, or did a PC technician work on your machine?

USER: Yeah, someone was in here last night trying to get my sound card working.

YOU: Let's take a look at the configuration for the sound card . . .

By following this set of questions, you have narrowed the scope of the problem to only one workstation, verified that the physical connections work correctly, and discovered that a configuration change on the workstation might have caused the problem. At this point, you can probably assume that whoever worked on the user's sound card created a resource conflict between the sound card and the NIC, preventing the NIC from making a connection to the network. If you feel comfortable talking the user through checking the device settings, you could proceed with that approach. If not, you could visit the workstation yourself and fix the problem.

Scenario 2: A Misbehaving Network Printer

Network printers cause as many problems as network workstations (although they are usually less critical than servers). Typically, a malfunctioning network printer affects everyone who tries to use it. Although user input may prove helpful in solving network printer problems, you will probably get more information faster by checking the printer yourself. Following are some logical steps you might take to assess a network printer problem:

1. Try to narrow the scope of the problem by determining whether everyone or only a few of those who normally use the printer are having problems printing.
2. Try to replicate the error yourself. First, try to print to the printer from your machine (which is properly connected to the network and has the printer device drivers properly installed) to discover whether the problem might derive from workstation configurations. If you receive an error, note the exact wording of the error. If you do not receive an error, the symptoms are not network-wide, and the problem may be caused by either a user error or an incorrect printer device configuration on one workstation.
3. If you cannot replicate the problem from your computer, go to the workstations that have problems and try to replicate the error from them.

4. If the error occurs on only one workstation, the problem may be caused by physical connectivity problems or logical connectivity problems with that workstation. Check that workstation's network cable and NIC, and then check its printer device drivers and settings. Reinstall the device drivers if necessary.
5. If the error occurs on multiple workstations, the problem probably has to do with the printer itself. Visit the printer and verify its physical and logical connectivity. Make sure that the printer is turned on. Verify that the printer is properly connected to the network. Also, verify that the printer is ready to print—that is, it is online and has no internal errors.
6. If the printer is connected and ready to print, print a test page to view the printer's configuration. From this test page, you can determine whether the printer is connecting to the correct server, is receiving protocols correctly, and has a properly setup network configuration (for example, if it's on an Ethernet network and using IPX/SPX, make sure it has the correct frame type setting).

This logical sequence of steps allows you to zero in on the possible causes of the problem. Once you have determined that multiple workstations are experiencing the problem, that the problem is repeatable, that the device drivers are installed correctly on every workstation, and that the printer is properly physically connected to the network, you can turn your attention to the printer's network configuration. By process of elimination, this configuration is probably the source of the fault.

Scenario 3: Unable to Connect to the Internet

If your organization depends on e-mail and other Internet-related services, such as Web databases or e-commerce, an inability to connect to the Internet can quickly hamper productivity and perhaps affect the organization's profitability. At the least, being disconnected from the Internet is an inconvenience. An inability to connect to the Internet, like many network problems, may be caused by errors at a number of different points in the system. In the following scenario, a large group of users is affected by an Internet-related problem. The following steps suggest a way to troubleshoot the problem:

1. A user calls and complains that he can't pick up his e-mail. At the same time, the other two network administrators in your department are fielding similar calls. When you finish your phone calls and compare notes, you realize that the users who called are all located in your company's Finance Department.
2. You call your company's help desk and tell the first-level support analysts that the Finance Department has lost Internet access. You ask the analysts to let you know whether any other departments report the same problem.
3. You attempt to reproduce the problem by trying to access the Internet from your workstation.
4. If you fail to connect to the Internet, you would use the PING utility to see whether you could contact your TCP/IP gateway.

5. In this example, let's assume that you can reach the Internet. You, therefore, know that the problem must be isolated to other areas of the company, which include the Finance Department.
6. Still at your desk, you try pinging the Finance department's default gateway address. A positive response indicates that physical connectivity to that gateway is sound. A negative response tells you that the gateway may be physically disconnected or otherwise incapacitated.
7. If you receive a positive response from the default gateway PING, your next step is to go to a Finance Department workstation and attempt to ping a host on another subnet (perhaps your own workstation, as you know that its TCP/IP resources are functional). A positive response from this test indicates that the workstation can communicate with and through the Finance Department's gateway. Thus the Finance gateway may not be incapacitated, but rather something else in the network (such as cabling from the router to the backbone) may not be working.
8. In this example, let's assume that you receive a negative response from the default gateway PING, which suggests either workstation or subnet connectivity problems from that node. Your next step is to try pinging the loopback address. A positive response to the loopback PING indicates that the workstation's TCP/IP services are installed and operating properly. Thus, you have narrowed the problem down to the subnet that includes the Finance Department.
9. A help desk analyst pages you with a message that the Accounting and Human Resources Departments are experiencing the same problems. You know that these departments are on the same subnet as Finance.

With the information you have gathered, you can conclude that the TCP/IP connectivity fault lies somewhere on the subnet that serves those three departments. You leave the Finance department and begin analyzing the network to find out whether the problem lies in the subnet's router or cabling.

Swapping Equipment

If you suspect a problem lies with a network component, one of the easiest ways to test your theory is to exchange that component for a functional one. In many cases, such a swap will resolve the problem very quickly, so you should consider trying this tactic early in your troubleshooting process. It won't always work, of course, but with experience you will learn what types of problems are most likely due to component failure.

For example, if a user cannot connect to the network, as in Scenario 1 in the “Examples of How to Investigate Problems” section, and even after entering the correct user ID and password still can’t log on, you might consider swapping the user’s network cable with a functional one. As you learned in Chapter 4, network cables must meet specific standards to operate properly. If one becomes damaged (for example, by a chair repeatedly rolling over it), it will prevent a user from connecting to the network. Swapping an old network cable with a new one is a quick test that may save you further troubleshooting.

In addition to swapping network cables, you might need to change a patch cable from one port in a hub or switch to another, or from one data jack to another. Ports and data jacks can be operational one day and faulty the next. You might also swap a network adapter from one machine to another or try installing a new network adapter, making sure it’s precisely the same make and model as the original. It’s more difficult to swap a switch or router because of the number of nodes serviced by these components and the potentially significant configuration they require; if network connectivity has failed, however, this approach may provide a quicker answer than attempting to troubleshoot the faulty device.



A better alternative to swapping parts is to have redundancy built into your network. For example, you might have a server that contains two network adapters, allowing one network adapter to take over for the other if one adapter should fail. If properly installed and configured, this arrangement results in no downtime; in contrast, swapping parts requires at least a few minutes of service disruption. In the case of swapping a router, the downtime might last for several hours.



Before swapping any network component, make sure that the replacement has exactly the same specifications as the original part. By installing a component that doesn’t match the original device, you risk thwarting your troubleshooting efforts, because the new component might not work in the environment. In the worst case, you may damage existing equipment by installing a component that isn’t rated for it.

Using Vendor Information

Some networking professionals pride themselves in being able to install, configure, and troubleshoot devices without reading the instructions—or at least exhausting all possibilities before they submit to reading a manual. Although some manufacturers clearly write better documentation than others, you have nothing to lose by referring to the manual, except a little time. Chances are you will find exactly what you need—jumper settings for a NIC, configuration commands and their arguments for a router, and troubleshooting tips for a network operating system function.

In addition to the booklets that ship with the networking component (which are often lost in a network manager's pile of documentation and miscellaneous equipment), most network software and hardware vendors provide free online troubleshooting information. For example, both Microsoft and Novell offer searchable databases in which you can type your error message or a description of your problem and receive lists of possible solutions. Reputable equipment manufacturers, such as 3Com, Cisco, IBM, Intel, and Hewlett-Packard also offer sophisticated Web interfaces for troubleshooting their equipment. If you cannot find the documentation for a networking component, you should try looking for information on the Web.

Bear in mind that some vendors require you to register for online support, and occasionally you may have to pay for this service. Nevertheless, most vendors provide a significant amount of information (including entire manuals) free of charge from their Web sites. Table 12-1 lists links to technical support Web sites for popular networking vendors. (Note that these URLs were verified at the time of this writing, but may change without notice.)

Table 12-1 Links for troubleshooting resources on the Web

| Vendor | Technical Support Web Site Address |
|-----------------|--|
| 3Com | www.3com.com/support/en_US/index3.html |
| Cisco | www.cisco.com/univercd/home/home.htm |
| Compaq | www.compaq.com/support/ |
| Dell | www.dell.com/support/index.htm |
| Hewlett Packard | welcome.hp.com/country/us/eng/support.htm |
| IBM | www-1.ibm.com/support/ |
| Intel | www-cs.intel.com/ |
| Lucent | www.lucent.com/support/ |
| Microsoft | support.microsoft.com/ |
| Nortel/Bay | www12.nortelnetworks.com/cgi-bin/cnss/cs/main.jsp |
| Novell | support.novell.com/ |
| Oracle | www.oracle.com/support/ |
| SMC | www.smc.com/smc/pages_html/support.html |
| Sun | www.sun.com/service/online/ |

Call the vendor's technical support phone number only after you have read the manual and searched the vendor's Web page. In some cases, you may wait a long time before getting an answer when you call. With some manufacturers, you can talk to a technical support agent only if you have established and paid for a support agreement. With others, you must pay per phone call. Each vendor has a different pricing structure for technical support, so before you agree to pay for technical support, you should find out whether the vendor charges on a per hour or per problem basis.



Keep a list handy (preferably online, either on a Web page or in a shared file on the network) of the hardware and software vendors for your networking equipment; the list should include not only the company's name, but also its technical support phone number, a contact name (if available), its technical support Web site address, policies for technical support, and the type of agreement you currently have with the vendor. You can find an example of such a form in Appendix D, "Examples of Standard Networking Forms." Make sure the list is updated regularly and available to all Information Services personnel who might need it.

Notify Others of Changes

After solving a particularly thorny network problem, you should not only record its resolution in your call tracking system, but also notify others of your solution and what, if anything, you needed to change to fix the problem. This communication serves two purposes: (1) it alerts others about the problem and its solution, and (2) it notifies others of network changes you made, in case they affect other services.

The importance of recording changes cannot be overemphasized. Imagine that you are the network manager for a group of five network technicians who support a WAN consisting of three different offices and 150 users. One day the company's CEO travels from headquarters to a branch office for a meeting with an important client. At the branch office, she needs to print out a financial statement, but encounters a printing problem. Your network technician discovers that her login ID does not have rights to that office's printer, because users on your WAN do not have rights to printers outside the office to which they belong. The network technician quickly takes care of the problem by granting all users rights to all printers across the WAN. What are the implications of this change? If your technician tells no one about this change, at best users may incorrectly print to a printer in Duluth from the St. Paul office. In a worst-case scenario, a "guest" user account may gain rights to a networked printer, potentially creating a security hole in your network.

Large organizations often implement change management systems to methodically track changes on the network. A **change management system** is a process or program that provides support personnel with a centralized means of documenting changes to the network. In smaller organizations, a change management system may be as simple as one document on the network to which networking personnel continually add entries to mark their changes. In larger organizations, it may consist of a database package complete with graphical interfaces and customizable fields tailored to the computing environment. Whatever form your change management system takes, the most important element is participation. If networking personnel do not record their changes, even the most sophisticated software is useless.

The types of changes that network personnel should record in a change management system include the following:

- Adding or upgrading software on network servers or other devices
- Adding or upgrading hardware components on network servers or other devices
- Adding new hardware on the network (for example, a new server)
- Changing the network properties of a network device (for example, changing the IP address or NetBIOS name of a server)
- Increasing or decreasing rights for a group of users
- Physically moving networked devices
- Moving user IDs and their files/directories from one server to another
- Making changes in processes (for example, a new backup schedule or a new contact for DNS support)
- Making changes in vendor policies or relationships (for example, a new hard disk supplier)

It is not necessary to record minor modifications, such as changing a user's password, creating a new group for users, creating new directories, or changing a network drive mapping for a user. Each organization will have unique requirements for its change management system, and analysts who record change information should clearly understand these requirements.

Preventing Future Problems

If you review the list of questions and the troubleshooting scenarios given at the beginning of this chapter, you can predict how some network problems can be averted by network maintenance, documentation, security, or upgrades. Although not all network problems are preventable, many can be avoided. Just as with your body's health, the best prescription for network health is prevention.

For example, to avoid problems with users' access levels for network resources, you can comprehensively assess users' needs, set policies for groups, use a variety of groups, and communicate to others who support the network why those groups exist. To prevent overusing network segments, you should perform regular network health checks—perhaps even continual network monitoring—and ensure that you have the means to either redesign the network to distribute traffic or purchase additional bandwidth well before utilization reaches critical levels. With experience, you will be able to add more suggestions for network problem prevention. When planning or upgrading a network, you should consciously think about how good network designs and policies can prevent later problems—not to mention, make your job easier and more fun.

CHAPTER SUMMARY

- Before you can resolve a network problem, you need to determine its cause. The key to solving network problems is to approach them methodically and logically, using your experience to inform your decisions, and knowing when to ask for someone else's help.
- When assessing a network problem, act like a doctor diagnosing a patient. First, ask the user a series of standard questions in a logical order to learn about the problem's symptoms. Never ignore the obvious! Although some questions may sound too simple to bother asking, don't discount them.
- Next identify the scope of the problem. In general, a network problem may be limited by the number of users, departments, or areas it affects or by what times of day or week it occurs.
- At each point in the troubleshooting process, stop to consider what kind of changes have occurred on the network that might have created a problem. Changes pertaining to hardware may include the addition of a new device, the removal of an old device, a component upgrade, a cabling upgrade, or an equipment move. Changes pertaining to software may include an operating system upgrade, device driver upgrade, a new application, or a changed configuration.
- Early in the troubleshooting process, you should ensure that the user is performing all functions correctly. It's easy for a user to make mistakes and assume that something is wrong with the network.
- Attempt to reproduce the problem's symptoms. If possible, go to the location where the problem is occurring and try to repeat the steps precisely. Note also whether a problem is repeatable only under specific circumstances.
- Check whether the affected device (or devices) have sound connections to the network, from workstation to backbone. Physical connectivity may be impaired by poorly or incorrectly installed cabling, NICs, or connectivity devices; flawed or damaged components; or excessive segment length.
- If you find no physical connectivity problems, determine whether the affected device(s) have properly configured software, including applications, hardware configurations, operating system software, and client software.
- After implementing your solution, you must test it to ensure that it works correctly. The type of testing you perform will depend on your solution. Enlist the help of users to test the solution. If the solution required significant network changes, revisit the solution a day or two after you implement it to verify that it has truly worked and not caused additional problems.

- A tone generator and tone locator are used to identify the terminating location of a wire pair. Telephone technicians use these tools more often than network technicians, and this combination of devices may also be known as a fox and hound.
- A multimeter is a simple device that can measure the voltage, resistance, and other characteristics of an electrical circuit.
- Basic cable checkers determine whether your cabling can provide connectivity. To accomplish this task, they apply a small voltage to each conductor at one end of the cable, and then check whether that voltage is detectable at the other end. They may also verify that voltage cannot be detected on other conductors in the cable. A good cable checker will also verify that the wires are paired correctly and that they are not shorted, exposed, or crossed.
- A cable tester performs the same continuity and fault tests as a cable checker, but also ensures that the cable length is not too long, measures the distance to a cable fault, measures attenuation along a cable, measures near-end crosstalk between wires, measures termination resistance and impedance for Thinnet cabling, issues pass/fail ratings for CAT3, CAT5, CAT6, or even CAT7 standards, and stores and prints cable testing results.
- Because of their sophistication, cable testers cost significantly more than cable checkers.
- A network monitor is usually a software-based tool that continually monitors traffic on the network from a server or workstation attached to the network. Network monitors typically can interpret up to Layer 3 of the OSI Model. They can determine the protocols passed by each packet, but can't interpret the data inside the packet.
- Network analyzers can typically interpret data up to Layer 7 of the OSI Model. They can also interpret the payload portion of packets, translating from binary or hexadecimal code to human-readable form.
- Before adopting a network monitor or analyzer, you should be familiar with some of the data errors that these tools can distinguish, such as runts, late collisions, jabber, and negative frame sequence checks.
- To take advantage of software-based network monitoring and analyzing tools, the network adapter installed in your machine must support promiscuous mode. Promiscuous mode means that a device driver directs the network adapter card to pick up all frames that pass over the network—not just those destined for the node served by the card.
- Microsoft's Network Monitor (NetMon) is a software-based network monitoring tool that comes with Windows NT Server 4.0 and Windows 2000.

- Novell provides a network monitoring tool called the LANalyzer agent. It can act as a standalone program on a Windows 9x or 2000 workstation or as part of the ManageWise suite of network management tools on a NetWare server. Like Network Monitor, LANalyzer can capture traffic, identify data errors by node, and generate traffic statistics by segment.
- You may choose to purchase network analyzing software from vendors that specialize in products for network management. One popular example is Network Associates' Sniffer Portable, network analyzer software that provides data capture and analysis, node discovery, traffic trending, history, alarm tripping, and utilization prediction.
- Network Associates has also led the way in hardware-based network analyzers, known as sniffers. Sniffers are usually regular laptops equipped with a special network adapter and software dedicated to network analysis.
- Sniffers are tailored to a particular type of network. For example, one sniffer may be able to analyze both Ethernet and Token Ring networks, but another sniffer may be necessary to analyze fiber or ATM networks. The cost of sniffers can range from \$10,000 to \$30,000.
- Most organizations operate a help desk staffed with first-level support personnel who field user questions, perform initial problem diagnosis, and record problems in a call tracking database. Help desks also use second-level support personnel, who are experts in some aspect in a specific area of computing. In addition, help desk coordinators maintain help desk schedules and ensure that help desk staff members have the resources necessary to perform their jobs. An operations manager typically supervises the help desk coordinator and approves the help desk's budget.
- If you suspect that a problem lies with a network component, one of the easiest ways to test your theory is to exchange that component for a functional one. In many cases, this tactic will resolve the problem very quickly, so you should consider trying it early in your troubleshooting process.
- Although some manufacturers clearly write better documentation than others, you have nothing to lose by referring to a product's manual. Most network software and hardware vendors also provide free online troubleshooting information.
- Keep a list of the hardware and software vendors for your networking equipment. This list should include not only the company's name, but also its technical support phone number, a contact name, technical support Web site address, policies for technical support, and the type of agreement that you currently have with the vendor.
- Some organizations use a software program for documenting problems, known as a call tracking system (or help desk software). These programs provide a user-friendly graphical interface that prompts the user for every piece of information associated with the problem.

- Whether you use a formal call tracking system or a simple form, you should record the following details about a problem: the originator's name, department, and phone number; whether the problem is software- or hardware-related; if the problem is software-related, the package to which it pertains; if the problem is hardware-related, the device or component to which it pertains; the symptoms of the problem, including when it was first noticed; the name and telephone number of the network support contact; the amount of time spent troubleshooting the problem; and the resolution of the problem.
- In addition to communicating problems and solutions to your peers whenever you work on a network problem, you should follow up with the person who reported the problem. Make sure that the client understands how or why the problem occurred, what you did to resolve the problem, and who to contact should it recur.
- Organizations often implement change management systems to methodically track changes on the network. A change management system is a process or program that provides support personnel with a centralized means of documenting changes to the network.
- Network personnel should record the following types of changes in a change management system: adding or upgrading software, adding or upgrading hardware, changing the network properties of a network device, increasing or decreasing rights for a group of users, physically moving networked devices, moving user IDs and their files/directories from one server to another, making changes in processes, and making changes in vendor policies or relationships.

KEY TERMS

baseline — A record of how well the network operates under normal conditions (including its performance, collision rate, utilization rate, and so on). Baselines are used for comparison when conditions change.

cable checker — A simple handheld device that determines whether cabling can provide connectivity. To accomplish this task, a cable checker applies a small voltage to each conductor at one end of the cable, then checks whether that voltage is detectable at the other end. It may also verify that voltage cannot be detected on other conductors in the cable.

cable tester — A handheld device that not only checks for cable continuity, but also ensures that the cable length is not excessive, measures the distance to a cable fault, measures attenuation along a cable, measures near-end crosstalk between wires, measures termination resistance and impedance for Thinnet cabling, issues pass/fail ratings for wiring standards, and stores and prints cable testing results.

call tracking system — A software program used to document problems (also known as help desk software). Examples of popular call tracking systems include Clientele, Expert Advisor, Professional Help Desk, Remedy, and Vantive.

change management system — A process or program that provides support personnel with a centralized means of documenting changes made to the network. In smaller organizations, a change management system may be as simple as one document on the network to which networking personnel continually add entries to mark their changes. In larger organizations, it may consist of a database package complete with graphical interfaces and customizable fields tailored to the particular computing environment.

fox and hound — Another term for the combination of devices known as a tone generator and a tone locator. The tone locator is considered the hound because it follows the tone generator (the fox).

ghosts — Frames that are not actually data frames, but rather aberrations caused by a repeater misinterpreting stray voltage on the wire. Unlike true data frames, ghosts have no starting delimiter.

giants — Packets that exceed the medium's maximum packet size. For example, any Ethernet packet that is larger than 1518 bytes is considered a giant.

jabber — A device that handles electrical signals improperly, usually affecting the rest of the network. A network analyzer will detect a jabber as a device that is always retransmitting, effectively bringing the network to a halt. A jabber usually results from a bad NIC. Occasionally, it can be caused by outside electrical interference.

LANalyzer — Novell's network monitoring software package. LANalyzer can act as a standalone program on a Windows 9x or 2000 workstation or as part of the ManageWise suite of network management tools on a NetWare server. LANalyzer offers the following capabilities: discovery of all network nodes on a segment, continuous monitoring of network traffic, alarms that are tripped when traffic conditions meet preconfigured thresholds (for example, if usage exceeds 70%), and the capturing of traffic to and from all or selected nodes.

late collisions — Collisions that take place outside the normal window in which collisions are detected and redressed. Late collisions are usually caused by a defective station (such as a card, or transceiver) that is transmitting without first verifying line status or by failure to observe the configuration guidelines for cable length, which results in collisions being recognized too late.

local collisions — Collisions that occur when two or more stations are transmitting simultaneously. Excessively high collision rates within the network can usually be traced to cable or routing problems.

multimeter — A simple instrument that can measure multiple characteristics of an electric circuit, including its resistance and voltage.

negative frame sequence checks — The result of the cyclic redundancy checksum (CRC) generated by the originating node not matching the checksum calculated from the data received. It usually indicates noise or transmission problems on the LAN interface or cabling. A high number of (nonmatching) CRCs usually results from excessive collisions or a station transmitting bad data.

- network analyzer** — A portable, hardware-based tool that a network manager connects to the network expressly to determine the nature of network problems. Network analyzers can typically interpret data up to Layer 7 of the OSI Model.
- network monitor** — A software-based tool that continually monitors traffic on the network from a server or workstation attached to the network. Network monitors typically can interpret up to Layer 3 of the OSI Model.
- Network Monitor (NetMon)** — A software-based network monitoring tool that comes with Windows NT Server 4.0 or Windows 2000. Its capabilities include capturing network data traveling from one or many segments, capturing frames sent by or to a specified node, reproducing network conditions by transmitting a selected amount and type of data, detecting any other running copies of NetMon, and generating statistics about network activity.
- ohmmeter** — A device used to measure resistance in an electrical circuit.
- optical time domain reflector (OTDR)** — A time domain reflector specifically made for use with fiber-optic networks. It works by issuing a light-based signal on a fiber-optic cable and measuring the way in which the signal bounces back (or reflects) to the OTDR.
- promiscuous mode** — The feature of a network adapter card that allows a device driver to direct it to pick up all frames that pass over the network — not just those destined for the node served by the card.
- protocol analyzer** — See *network analyzer*.
- resistance** — The opposition to an electric current. Resistance of a wire is a factor of its size and molecular structure.
- runts** — Packets that are smaller than the medium's minimum packet size. For instance, any Ethernet packet that is smaller than 64 bytes is considered a runt.
- sniffer** — A laptop equipped with a special network adapter and software that performs network analysis. Unlike laptops that may have a network monitoring tool installed, sniffers typically cannot be used for other purposes, because they don't depend on a desktop operating system such as Windows.
- Sniffer Portable** — Network analyzer software from Network Associates that provides data capture and analysis, node discovery, traffic trending, history, alarm tripping, and utilization prediction.
- spike** — A single (or short-lived) jump in a measure of network performance, such as utilization.
- supported services list** — A document (preferably online) that lists every service and software package supported within an organization, plus the names of first- and second-level support contacts for those services or software packages.
- time domain reflector (TDR)** — A high-end instrument for testing the qualities of a cable. It works by issuing a signal on a cable and measuring the way in which the signal bounces back (or reflects) to the TDR.
- tone generator** — A small electronic device that issues a signal on a wire pair. When used in conjunction with a tone locator, it can help locate the termination of a wire pair.

tone locator — A small electronic device that emits a tone when it detects electrical activity on a wire pair. When used in conjunction with a tone generator, it can help locate the termination of a wire pair.

voltmeter — A device used to measure voltage (or electrical pressure) on an electrical circuit.

REVIEW QUESTIONS

1. If, after several tries, you cannot reproduce symptoms of a problem, what might you suspect as the cause of the problem?
 - a. user error
 - b. faulty cabling
 - c. incorrect software configuration
 - d. incompatible protocols
 - e. an improperly installed NIC
2. Which of the following symptoms probably points to a physical connectivity problem?
 - a. a group of users consistently experiences delays on the network
 - b. a user always loses his drive mappings to file server directories
 - c. a group of users complain that they cannot log onto the network
 - d. a user can send e-mail but can't pick it up
 - e. a user is receiving instant network messages intended for someone else
3. Which part of the network should you examine if a network problem affects a single workstation?
 - a. the segment's router interface
 - b. the cabling between the switch and the backbone
 - c. the workgroup's hub
 - d. the entrance facility connections
 - e. the workstation's NIC and cabling
4. You are troubleshooting a problem in which a dial-in remote user claims he cannot make a connection to your organization's access server. Of the following steps, which should you take first and second as you diagnose this problem?
 - a. Ask the user his password so you can replicate the problem from a workstation at your desk.
 - b. Ask the user how long the problem has been occurring.
 - c. Ask the user to try pinging the organization's Web server and read the results to you.
 - d. Ask the user what type of error message he sees when he tries to connect.
 - e. Ask the user whether he has changed any of his software configurations lately.

5. You have recently resolved a problem in which a user could not print to a particular shared printer, by upgrading her workstation's client software. Which of the following might be an unintended consequence of your solution?
 - a. The user is no longer able to use her e-mail application from her hard disk.
 - b. The user complains that her login screen looks different.
 - c. The shared printer no longer allows users to print double-sided documents.
 - d. The shared printer no longer responds to form feed commands from the print server.
 - e. The workgroup to which the user belongs cannot see the printer on the network.
6. Answering which two of the following questions may help you identify the demographic scope of a problem?
 - a. When did the problem first occur?
 - b. How frequently does the problem occur?
 - c. How many users have similar symptoms?
 - d. Do the symptoms appear on all workstations in one department?
 - e. Are the cables properly inserted into the hub, wall jack, and device NIC?
7. Which of the following is a characteristic symptom of a gateway failure?
 - a. All workstations on a segment are unable to perform networked functions at all times.
 - b. All workstations on a segment are intermittently prevented from connecting to the network.
 - c. All workstations on a segment lose their IP addresses.
 - d. Only one workstation is unable to log onto the network.
 - e. Some workstations on a segment cannot run the same application from the server.
8. Under what circumstances should you try swapping equipment?
9. You have just discovered that your backup device is not properly writing files to your backup media. Which of the following would be the *last* two steps you take in troubleshooting this problem?
 - a. Determine when the last good backup was made.
 - b. Document your solution and share your notes with colleagues.
 - c. Call the backup software vendor's technical support line.
 - d. Check the backup software log for errors.
 - e. Upgrade the backup software according to the vendor's recommendation.

10. Which of the following is an example of a network change that could cause a group of workstations to lose connectivity to one local file server?
 - a. The server is renamed.
 - b. The dedicated line to the Internet fails.
 - c. One of the server's two NICs fails.
 - d. The server's backup device fails.
 - e. The server's external storage device fails.
11. Which of the following tools could you use to determine whether a user's workstation is transmitting packets in the proper Ethernet frame type for your network?
 - a. multimeter
 - b. cable checker
 - c. time domain reflector
 - d. cable tester
 - e. network analyzer
12. Which of the following symptoms would definitely be present if your Ethernet network length exceeds the maximum specified by IEEE standards?
 - a. excessive normal collisions
 - b. excessive late collisions
 - c. giants
 - d. ghosts
 - e. crosstalk
13. Which member of the IT department staff is usually the first to receive notice of a network problem?
 - a. help desk analyst
 - b. IT director
 - c. network administrator
 - d. help desk supervisor
 - e. chief information officer
14. If you don't have the manual for your 3Com NIC, how can you find out whether it supports promiscuous mode?
 - a. Read its label.
 - b. Look up the information on 3Com's Web site.
 - c. Attach it to a network protocol analyzer.
 - d. Attempt to flood it with traffic and gauge its response.
 - e. Read the manual of another type of 3Com NIC.

15. What kind of tool would you use to verify that your new cable meets CAT5 standards?
 - a. cable tester
 - b. cable checker
 - c. cable monitor
 - d. tone generator and tone locator
 - e. multimeter
16. Which TCP/IP command can you use to find out whether a workstation's TCP/IP stack is operating properly?
 - a. netstat
 - b. nbtstat
 - c. ftp
 - d. ping
 - e. nslookup
17. Where is crosstalk most likely to occur?
18. Which two of the following tools can help you determine whether your Thinnet connection has the proper amount of impedance at each end?
 - a. protocol analyzer
 - b. cable tester
 - c. cable gauge
 - d. time domain reflector
 - e. multimeter
19. Which of the following frequently results in negative frame sequence checks?
 - a. improper flow control
 - b. excessive nodes on a segment
 - c. excessive segment length
 - d. incorrect protocol configuration
 - e. noise
20. Which of the following frequently causes a jabber?
 - a. near-end crosstalk
 - b. faulty cabling
 - c. faulty NIC
 - d. excessive segment length
 - e. noise

21. With what operating system does NetMon work?
 - a. UNIX
 - b. NetWare
 - c. Linux
 - d. Windows 98
 - e. Windows 2000
22. The LANalyzer agent can help you determine when network traffic exceeds 50%. True or False?
23. If you wanted to determine the average daily traffic on your network's backbone, what type of tool would you use?
 - a. network analyzer
 - b. cable tester
 - c. time domain reflector
 - d. network monitor
 - e. multimeter
24. Name two advantages of using a sniffer over using NetMon or LANalyzer.
25. Which two of the following functions can both network monitors and network analyzers perform?
 - a. capture and analyze data traveling from one node to another
 - b. identify a faulty cable
 - c. provide trend information on data traffic from a switch port
 - d. capture and interpret unencrypted passwords on the network
 - e. discover nodes on the network
26. How do switches affect network analyzers?
 - a. They prevent network analyzers from working.
 - b. They limit the amount of the traffic that a network analyzer can capture.
 - c. They cause interference that can skew the data captured by a network analyzer.
 - d. They generate excessive numbers of bad packets, thereby flooding the network analyzer with data.
 - e. They initiate frequent broadcasts that require filtering before an analyzer can capture useful data.
27. You can typically use the same sniffer for your Token Ring and ATM networks. True or False?

28. You have just purchased a new network adapter to replace the faulty network adapter in your file server. The adapter is so new that your Windows 2000 Server software does not provide a device driver for it. As you install the network adapter, where should you obtain the device driver from?
 - a. the Windows 2000 Server technical support Web site
 - b. the network adapter manufacturer's Web site
 - c. the floppy disk that came with the network adapter
 - d. a server directory containing device drivers for other network adapters used on your network
 - e. a client on the network that uses the same network adapter
29. You work in a small office with only six employees and a small, peer-to-peer network that uses a single hub to connect all the workstations. One day you glance at the hub and notice that one of the port's LEDs has gone from blinking green to blinking amber. What can you conclude about the workstation connected to that port?
 - a. It has been shut down.
 - b. Its NIC has a problem.
 - c. Its NIC has switched speeds from 10 Mbps to 100 Mbps.
 - d. Its file-sharing capability has been turned off.
 - e. Its TCP/IP settings have been changed to use DHCP rather than static addressing.
30. Which of the following is a network change that does not need to be recorded in the change management system?
 - a. adding a new disk drive to a server
 - b. moving a hub from one closet to another
 - c. replacing the NIC in a router
 - d. changing a user's password
 - e. upgrading the network operating system

HANDS-ON PROJECTS

Until you use a network troubleshooting tool, such as Network Monitor, it's difficult to understand how these programs work. The following Hands-on Projects offer you a chance to try out cable testers and network monitors. In a real networking environment, you will probably use a number of different tools, depending on your network environment. They are similar enough, however, so that if you master one you can easily master another.

For the following exercises, you will need the cable you created during the Hands-on Projects in Chapter 4, a cable tester (such as the Fluke DSP-4000 CableAnalyzer), a penknife (or scissors), and a Windows 2000 server with several clients connected to it.



Project 12-1

In this project, you will find out how a cable tester detects and reports a damaged cable.

1. In the Hands-on Projects in Chapter 4, you created a CAT5 cable with two RJ-45 connectors. Retrieve that cable (or make a new one), and use the cable tester to find out whether it meets CAT5 standards.
2. If your cable does not meet CAT5 standards, cut off both connectors and recrimp it according to the standard. Test it again.
3. If your cable does meet CAT5 standards, use a penknife to slice about one-fourth of the way through the cable, making sure to pass the housing and at least nick one of the twisted pairs.
4. Try testing the cable again with your cable tester. What kind of message (or messages) do you receive?



Project 12-2

In this exercise, you will use Network Monitor from a Windows 2000 server to capture data on the network.

1. With at least five clients connected to your Windows 2000 server, open Network Monitor as follows: Click **Start**, point to **Programs**, point to **Administrative Tools**, and then click **Network Monitor**.
2. Maximize one or both Network Monitor screens, if necessary.
3. Click **Capture** on the menu bar, and then click **Start**. Network Monitor begins capturing frames.
4. Go to (or have one of your classmates go to) one of the clients connected to the Windows 2000 server and start an application from the server. Exit to a DOS prompt, and then ping the server's IP address. Have someone log onto the server from a different client on the network.
5. After you have generated a few minutes of network traffic, click **Capture** on the menu bar, and then click **Stop**.
6. To view more detail on the captured data, click **Capture** on the menu bar, and then click **Display Captured Data**. Close the Capture Summary after viewing it.
7. Use the scroll bars on each pane within the Microsoft Network Monitor window to view the captured data, including network utilization, statistics, and address information.
8. In the bottom window, click a network address to view its data frames in more detail. What kind of protocols do the frames use? If the network is based on Ethernet, what version are you using? What is your server's MAC address?
9. Find the frames that pertain to the logon process mentioned in Step 4. Can you read the person's password in ASCII form?
10. Close the Network Monitor program without saving the data you have captured.



Project 12-3

In this project, you will use the troubleshooting methodology discussed in this chapter to solve a network problem of your own creation. (If you are in a classroom setting and can work in pairs, it may be more fun to have a partner create a connectivity problem with your workstation, and then troubleshoot the problem.) For this exercise, you will use one of the clients from Hands-on Project 12-2, in addition to the Windows 2000 server.

1. Turn off your workstation and remove the cover, as you learned to do when installing network adapters in Chapter 6.
2. Find the network adapter and loosen it from its slot until approximately half of the pins are above the slot connector. (Depending on how far you remove the NIC, you may experience different types of symptoms.)
3. Close your workstation's cover and turn on the workstation, making sure that the network cable is properly connected to the network adapter and data jack.
4. Follow the steps in the troubleshooting methodology described at the beginning of this chapter, answering all questions under each step. Keep your answers on a separate sheet of paper.
5. Once you have followed the troubleshooting steps, summarize how the problem manifested itself. At Step 3 of the troubleshooting process, to how many different types of problems could your symptoms have applied? How many at Step 5?
6. Resolve the problem.
7. After you have resolved the problem, create a method for testing the problem to verify that your solution worked. Did it work? How can you be sure?



Project 12-4

In this project you will have the opportunity to act as if you are either experiencing a network problem or troubleshooting a network problem.

1. First, pair up with another student.
2. Designate one person in your pair as the user and the other person as the troubleshooter.
3. The user should pick one of the network problems listed below and take a few moments to consider the likely symptoms of those problems. The user should also anticipate which questions the troubleshooter will ask and prepare answers to those questions (which the user will deliver acting as if he or she does not know the cause of the problem). The user should not reveal which problem that has been selected to the troubleshooter.
 - Interference from nearby machinery is influencing a group of users' workstations.
 - A mouse has chewed through the cable that connects a print server to the network backbone.

- A network manager has used your workstation to log on as administrator and left the client software configured with his settings.
 - The infrared port on your laptop is covered with grime and preventing the laptop's wireless network adapter from working.
 - The RJ-45 connector for your workgroup's hub has been knocked out of its switch port.
 - The same address assigned to an organization's Web server has been assigned to your workstation.
 - You have inadvertently uninstalled the client software for your workstation.
 - The carrier that supplies your organization's Internet connection has suffered a construction accident that severed its fiber-optic cables.
 - A technician mistakenly replaced your workstation's patch cable with a crossover cable.
 - You are typing in the wrong logon password.
 - On a network you use IPX/SPX for local file and print services, but TCP/IP for Internet connectivity, and you have inadvertently deleted your DNS settings from the TCP/IP configuration.
 - A smoldering fire has broken out in the plenum above the server room where the network's backbone cables lie.
 - Someone has replaced your CAT5 connection between the patch panel and the hub with a CAT3 cable.
 - You are new to the organization, and your workstation has been added to the end of a 100BaseTX segment whose length is 180 meters.
4. While the user is thinking about how to characterize the problem, the troubleshooter should write down four questions he or she will ask sometime early in their conversation.
 5. Next, the user should initiate the conversation with a vague complaint that pertains to the problem. The troubleshooter should ask as many of his or her four questions as are applicable and write down the answers. If the troubleshooter can guess which problem the user has, that's great. If not, he or she should write down four more questions that will lead to the answer.
 6. Now that the troubleshooter knows which problem was selected, the user and the troubleshooter should discuss a possible solution and agree on the best course of action.
 7. After user and troubleshooter have determined a good solution, reverse roles and begin the project again at Step 3.

CASE PROJECTS



1. You are a network support technician for a college with 4,000 users scattered over five locations. A group of users from the downtown location has called your help desk, complaining that they cannot send or receive messages from the Internet, although they can receive messages on the college's internal GroupWise system. List the steps you will take to troubleshoot this problem and describe why each step is necessary.
2. While you're downtown fixing the first problem, a fellow network technician asks you to look at the library's server. She informs you that it's "flaky." Sometimes it doesn't allow users to log on; other times, it works perfectly. Sometimes it responds so slowly to requests for programs or files that users think it's frozen, but after several minutes it does finally respond. How would you troubleshoot this problem in the most efficient manner? Explain why you chose the steps you propose and how each might save you time.
3. You're in high demand because the word has gotten around the college that you can fix problems quickly. A small satellite campus requests that you visit it and examine a group of workstations in a computer lab that often—but not always—has problems connecting to a server. Your contact is a new instructor who teaches Interior Design in the lab. The workstations worked perfectly until the beginning of the semester, and no hardware or software changes have been made to the machines. Explain how you would troubleshoot this problem and why you chose the steps you propose.
4. Suggest ways that the problems in Case Projects 12-1, 12-2, and 12-3 might have been prevented.
5. Your friend Joseph, who works as a network technician for a global long-distance firm with 300 networked locations along the Eastern seaboard, calls you for help. Usually, five other technicians are on duty to help him handle technical problems. Today, two of his co-workers are out sick, one is away on jury duty, and another has not shown up for work yet. That leaves Joseph and one other technician to solve all of the problems that have occurred on this particular morning, including the following:
 - A WAN link is down between the Washington and New York locations, causing traffic to be rerouted from Washington to Boston, then to New York. As a result, customers are complaining about slow performance.
 - The Albany, New York, location's network appears to have suffered a catastrophic failure. This failure has caused outages for thousands of customers in the upstate New York region.
 - Three executive users at Joseph's corporate headquarters in Baltimore cannot pick up their e-mail, and they are calling every five minutes to ask when the problem will be fixed.

- A networked printer that provides services to the Accounting group at the Baltimore headquarters is not accepting any print jobs. The users have asked Joseph to troubleshoot the printer. They need to send invoices out to customers by noon.
 - Half of the workstations in the Advertising Department seem to be infected with a virus, and Joseph is worried that these users will copy the virus to the network, thus risking widespread data damage.
 - Joseph asks for advice about the order in which he and his other colleague should address the problems (or which ones to address simultaneously). What do you tell him, and why would you place them in that order?
6. Joseph is very grateful for your assistance and calls you at the end of the day to tell you how things turned out. One problem was particularly difficult to diagnose, because he didn't get all of the details until well into the troubleshooting process. As it turned out, the three executives—Sal, Martha, and Gabe—who couldn't pick up their e-mail messages were all sitting in a conference room with another two executives, Barb and Darrel. Barb and Darrel are vice presidents in the Operations group and had scheduled the meeting in a conference room down the hall from their offices. Sal, Martha, and Gabe, on the other hand, are vice presidents of Marketing, Engineering, and Research. They had to travel from other buildings on the headquarters' grounds to reach the conference room. Although Barb and Darrel could pick up their e-mail before the meeting started, the other three executives couldn't. He asks you to guess what the problem was. What do you tell him?
7. Joseph tells you that he first received the call for help from Sal at 7:54 A.M. and finally solved the executives' problem by 10:00 A.M. Write a sample tracking record for the incident described in Case Project 12-6. Include all pertinent details that will help future troubleshooters more quickly diagnose the same kind of problem and that will enable you to give the executives thorough, clear answers in case they call to ask why the problem took so long to fix.