CHAPTER

16

Troubleshooting

Difficulties exist to be surmounted.
— Ralph Waldo Emerson

We complete this book with a topic that we all hope we never need to worry about, but all need to know. It is nearly impossible to run a LAN full time with nothing ever going wrong. Even the most carefully designed networks experience ''issues.'' As a matter of fact, when designing your network, you decided on your acceptable level of risk for such issues. Now that your network is operational, it's your actions, both proactive and reactive, that are going to get you out of trouble quickly when problems occur within the LAN.

No two LANs are alike. Even if they are alike in design, operationally they are their own entities. There is no one fix-all for any particular issue within the LAN. The complexities of today's networks (high-speed data transfer, complex end-user application, etc.)[1] complicate the troubleshooting process even more. This is why we find ourselves (us writing and you reading) with this chapter. It's not a troubleshooting bible, but it is a guide that you can use to help you get a feel for 1) what is out there, and 2) a little of what you can expect. We hope it gives you the upper hand when you first approach troubleshooting and serves as a useful reference for you in the future.

The quote used for this section, ''Difficulties exist to be surmounted,'' seemed like a perfect thing to remember when taking on the challenge of an issue that has reared its ugly head. The more that we thought about the quote,

[1] According to the NASA Mars Rover update on `http://marsrovers.nasa.gov`, the Mars Rover, *Opportunity*, completed a 7.5 mile journey on the planet Mars that occurred in late September 2008. Think about this. Somewhere in the world (and outside of the world), there was data flowing from here to there that told *Opportunity* exactly where to go and how to get feedback to NASA when it was done.

the more we realized that this really is one of those ideas you should live by. Trouble and difficulty are two of life's guarantees. It is up to us whether we let them overtake the situation. In the network world, we do not have the option of letting the trouble overtake our LAN. That said, let's start surmounting!

16.1 The Little LAN that Cried Wolf

Troubleshooting a LAN can be a seriously tough job to accomplish. Many variables come into play. A specific issue may appear to be one thing and turn out to be something totally different. Sometimes a troubleshooting session can lead to a resolution in a few moments, and other times it may take hours.[2] If the issue is reasonably containable, you can often provide a quick fix without causing too much of an impact for any given set of issues. It's the times when you cannot get a fix in a reasonable amount of time that can cause you to wonder why the heck you even bothered reading this book.[3]

It is never a good thing when a catastrophic event impairs network functionality. Sometimes you may luck out and it will be an issue you have seen or heard of before. Other times, you will end up on a conference call with a team of managers, vendors, and other support staff trying to resolve a LAN problem.

With any luck, you made sure to insist that a robust network management station be installed when you were designing the network. Network management is big business and is well worth the time, money, and effort that are put into deploying one in a LAN. Network management and proactive troubleshooting are optional methods for keeping your LAN up and running. If, for whatever reason, you are unable to do either or both of these, you will eventually come to regret it.

End-user feedback is another way to find out if there are issues in the network. End users can provide invaluable help when you want to isolate an issue and then verify whether a fix worked.[4] Whichever method you use to police your LAN (everything at your disposal, we hope), make sure to establish a specific process to respond to events. You don't want a bunch of unprepared individuals running off to fix an issue with no strategy in mind. If this is the case, a hundred bucks says that you will have more than

[2]Sometimes it can take weeks and possibly even months to reach a full resolution for an issue. These problems are often related to poor design, vendor interoperability issues, or bugs in the code. If you can prove it is an issue with a vendor, you will be able to get them to bend over backwards to get you back up and running. Often there will already be a backup plan, but if not, a temporary resolution will suffice until a permanent solution is developed.

[3]Especially the part that Jim wrote! (Rich just got Jim back for those comments he's been throwing around throughout this book.)

[4]Although you have to fully trust that the end user is giving accurate data.

one individual working on the same issue and possibly heading in opposite directions strategy-wise.[5]

Before we move on in this chapter, let's look at some reasons why a network may begin experiencing problems and at some common feedback you might get when such LAN issues arise.

16.1.1 Feedback

When a problem occurs within your LAN (and remember, it's not *if*, but *when*), you need a notification procedure that alerts you as quickly as possible. After all, you want to narrow down and fix the issue before too many users, services, or applications are affected. As we have mentioned, you may have a network management station that alerts you during these times, but user feedback is also a great way of identifying a problem.

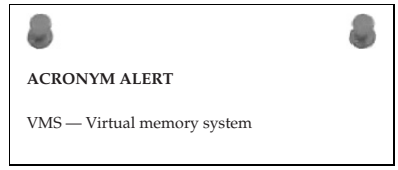

16.1.1.1 End-User Feedback

Often, the issue may be reported from a single individual, and with any luck, will turn out to be something particular to that one person. That person called in, reported what was occurring, and quickly you have them back up and running. Now it is time for some coffee. Have a great day! Issues that may be reported by an end user include the following:

- Unable to print to a network printer
- Unable to send or receive e-mail
- Unable to access a specific application server
- Unable to be authenticated
- Things on the computer seem a lot slower than usual
- Unable to get on the network

End users are not only useful in notifying you of an issue, they are also helpful in describing what is going on as you attempt to resolve an issue. End users can be walked through different tools that are running on

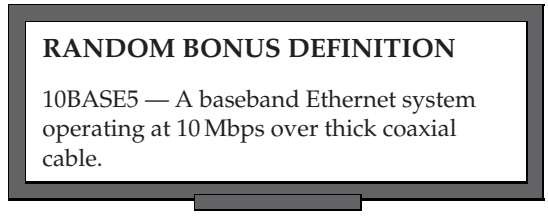

[5]Try to get a couple of bull-headed engineers off their train of thought. Before you know it, they are both making changes and are making matters worse.

their workstation, and you might be surprised how much these tools will tell you.

16.1.1.2 Management Station Feedback

As we discussed in Chapter 15, the network management station is a good way of being alerted to issues in the LAN. When an issue occurs that could potentially affect normal data flow, the network management station gives an alert. The alert can be visual (a red node identifier within a GUI) or audio (a ''ding'' or ''beep''). It can also be an alert that sends a message to a phone number, pager, or even an e-mail account (or set of accounts).

The management station can also provide statistical reports. You can review them for anything that seems abnormal. In other words, if you notice a particular port reporting an excessive number of errors, you may want to investigate to determine whether this indicates a problem in the LAN.

16.1.1.3 Hmm . . .

It may seem that we are making a big to-do about nothing. We have network management. We designed the network, we know it better than the paper route we had in the seventh grade. Therefore, nothing can happen that we can't overcome in mere moments. Seriously, no sweat. We are done! Or are we?

16.1.2 What Could Possibly Go Wrong?

You know for certain that you took every precaution in the design of your network. It should be reliable and fast . . . very fast. As a matter of fact, you made sure to put in a top-of-the-line network management station with all the bells and whistles. So what could possibly go wrong? The answer to that question is, almost anything. And that is what is so challenging to the network professional. Exactly why is the network having issues? Following is a list of possible issues that you might come across when troubleshooting a problem in your LAN:[6]

- Damaged cables
- Dirty fiber
- Excessive signal attenuation
- Insufficient bandwidth

[6]We are going to let you in on a little secret that those who don't read footnotes may never find out. Pay attention to the items on the list because at the end of the chapter, there is a question that tells you to match each item with the OSI layer it applies to.

- Denial-of-service (DoS) attack
- Electrical interference
- Wireless interference
- Damaged nodes
- Damaged interface
- Dirty interface
- Configuration error
- Authentication issues
- Excessive utilization
- Excessive errors
- VLAN configuration error
- Class-of-service issue
- Quality-of-service issue

16.1.3 Food for Thought

So far we have reviewed examples of some of the symptoms that you might come across when monitoring and managing the LAN. Without getting too deep into the task of troubleshooting, let's take a look at some of the instant questions that might come to mind when first learning of a network issue.

If you are notified that something is irregular in the LAN, you should ask the following questions:

- How many users are affected?
- Is only one user experiencing the problem or several?
- Is the problem an expected one? If so, do you have an action plan?
- Have you seen the issue before?
- What is the impact to the LAN? In other words, what nodes are affected?
- How many domains are affected?

Understanding the problem is paramount to effective network troubleshooting (as discussed in depth later in this chapter). Without a good understanding of the problem, you may find yourself chasing the wrong trails instead of working toward resolution.[7]

[7] Although troubleshooting sometimes is a game of cat and mouse.

"OH NO! I BURNED THE DINNER!"

We have provided you with some excellent recipes to enjoy at some point, but now it is time to take a more serious approach to future dinner plans, so you can proactively be prepared when you happen to burn the dinner. Sure, you may laugh at this thought, and we know that burning the dinner is not something you plan on doing anytime in the near future, but there is a real concern here, and we will tell you why. The concern is that by this point in the book, you are so engrossed in reading that time is flying by. Because of this, it is possible that you may be reading this most excellent book and might, in fact, burn your dinner.

We (Rich and Jim) feel that it is important that we give you some helpful tips on things you can do to be prepared for such quick meal emergencies. By all means, feel free to add any of your own favorites to this (or even to your own) list.

Jim's List of Proactive Staples:

Frozen veggies	Bouillon	Refried beans
Cheddar cheese	Frozen pizza	Ground beef
Canned veggies	Peanut butter	Jelly
Crackers	Ramen	Eggs
Tortillas	Salsa	

Rich's List of Proactive Staples:

Peanut butter	Popcorn	Refried beans
Crackers	Provolone cheese	Salted peanuts
Dried figs	Cranraisins	Cheerios
Oatmeal	Toast	Macaroni and cheese
Franks and beans		

You might also want to consider eating a few leftovers. For a quick and easy meal, throw every leftover in a pot with some water and bouillon and make a nice refrigerator stew. Ramen with some lunch meat or tuna, an egg, and some American cheese makes a really yummy meal. If you can handle the taste and smell of kimchi (a Korean pickled cabbage), you can add that to the ramen, too.

Now, if your meal is for a date, break down and take the person out to dinner. Whatever you do, don't call your mom as a backup for the date. (We are sure she would be more than happy to come over and help, though. Probably with your baby pictures in tow.)

16.2 The Proactive Approach Beats the Reactive Approach Hands Down

POP QUIZ

_____ feedback is a great way to be notified of an issue on the LAN.

It is amazing how many network administrators still take an extremely loose approach to network management. Not that we are here to judge anyone, but it really doesn't make sense in most LANs to take a loose approach. In at least one[8] of the authors' opinions, proactively troubleshooting your network is one of the most important things you can do to decrease the time it takes to reach a resolution when an issue occurs on the LAN.

16.2.1 Baseline

To know that something is wrong, you have to know what right is (in this scenario, with regard to your LAN). In other words, what is normal? The baseline of the network is defined by the normal behavior of the network. Traffic patterns and protocol behaviors are only a couple of items that can be baselined and used as a comparison when something just does not look right. Here are a few things that you can use to baseline your network:

- **Traffic analyzer** — The traffic analyzer (also known as a network analyzer, packet analyzer, packet sniffer, or just sniffer) is an application or a specialized node used to capture data transmitted on the network. Each packet is captured and can be manipulated and sorted by RFC, protocol, statistics, or whatever specifications are set by the user. The data can be organized and set to graphs or any other form supported by the analyzer and set by the user.
- **Statistical graphing with the management station** — SNMP management stations can usually record statistical data and output the data in reports and graphs. Maintaining such information can prove useful in determining abnormal conditions within the internetwork.
- **Determine thresholds that need to be maintained** — Know your LAN. Test and analyze traffic patterns, traffic capacity limits, overall throughput operation, routing path costs, Physical layer well-being,

[8] And I know that the other author would agree with this, even if he doesn't say so.

etc. Keep the results on hand for reference when an issue occurs. Not only can this assist in alerting you of abnormal operations, it can offer proof of the abnormal operations should you need to bring in a vendor for assistance. Understand peak traffic periods, protocol usage statistics, and average throughput. Normal traffic patterns are important in understanding problems that occur in the network and resolving them quickly and effectively.[9]

By having a baseline to compare with, when the end user tells you that things are slower than usual, you will be able to compare the baseline data with real-time data to see if, in fact, the issue resides in the LAN. Baselining is an excellent way of reducing the time it takes to find the culprit when there is an issue on the LAN. Keeping a record of captured baseline data is a very important practice to make a habit. You will find that this is helpful not only in reaching a resolution, but also in proving that something is wrong when you start calling vendors for support. Often, you need to have the proof to show to help the vendor help you. In addition, sometimes the vendor won't be able to find an issue if there is no proof the issue exists. By having the baseline data, you have the proof.

RANDOM BONUS DEFINITION

campus switch — A switch used within a campus backbone.

16.2.2 Proactive Documentation

Keeping documentation that outlines the physical and logical topology is essential for proactive troubleshooting of the LAN. Even more essential is making sure that this documentation stays current and updated at all times. For instance, it does absolutely no good to maintain network diagrams that do not keep up with the changes that are ever present in the LAN. By documentation, we don't mean you must keep a paper copy of everything you retain; rather, we are suggesting the storage of critical information in either electronic (soft copy) or paper (hard copy) form.

We are offering an overview of some recommended documentation to maintain and have readily available. Like many things in networking, you can look at this information as a reference model. If you don't want to retain any of the information in this section, don't. Feel free to add to this list anything that you feel you need to add. A good saying to live by is, ''There is no such thing as too much documentation.''

[9]This is not meant to infer that you will never reach a resolution without this information. This just makes the job of troubleshooting easier.

Make sure that your baseline documentation is saved and is regularly updated. Make it a habit to change the baseline documentation as changes are introduced into the network. In addition to the baseline documentation, keep good records that pertain to the physical and logical topologies of the network.[10] Keep in mind that even the most impressive network monitoring system is not going to do you any good if you cannot locate a problem node.

At the very least, keep a record of the following:

- A logical topology diagram
- A physical topology diagram
- A spreadsheet or other listing of where user nodes reside
- Documentation that relates to the DMZ
- Documentation that contains information pertaining to nodes outside of the DMZ
- Documentation that outlines the location and node interconnectivity for:
  - Network firewalls
  - Network management stations
  - Mainframe servers
  - Remote access servers
  - Routers
  - Switches
  - Layer 4–7 nodes
  - Wiring closets
  - VLANs
  - VPNs
- Information pertaining to the LAN to WAN connectivity
- Information pertaining to campus-to-campus connectivity
- Documentation listing:
  - Network layer addressing
  - Major node names
  - LAN identification
  - Node serial number
  - Node makes and model names/numbers

[10]The topology documentation is especially important for the portions of the network that are related to major host nodes, server nodes, interconnecting nodes, and individual segments within the network.

- Software and hardware version numbers
- Licensing information

- Circuit identification for WAN connections
- Geographical information for nodes listed above, including address and contact name and number
- Node access login information
- Backup copies of:
  - Node configurations
  - System logs
  - System software
- Name and number for support with the ISP
- Name and number for support for the vendors of nodes and application in use within the LAN

Again, these are all simply recommendations. It is entirely up to you what you choose to keep track of. At the very least, keep a record of addressing schemes.[11] It's nice to have if you need a reference point to start troubleshooting from.

POP QUIZ

Name 10 issues that you might have on the LAN.

16.2.3 There Is No Such Thing as Too Much

Establishing a baseline for your network is not the only thing you can do to simplify the task of troubleshooting. You can take a few other proactive steps to help keep the LAN running and get it back to running when you have an issue:

- **Shared knowledge** — It's always good to share what you know. Have discussions with others when you are not sure, and freely offer up any helpful information you may have with anyone who might have a need to know.
- **Proper tools** — Make sure that you and your staff have access to the correct tools. Examples include backup laptops, serial cables, baseline documentation, etc. The staff can only be as effective as their tools allow them to be.

[11] Although having just the addressing schemes is really inadequate for today's high-speed LANs

- **Knowledge requirements** — Make sure there is no single point of failure in the network. Make sure that any and all tasks are performed by at least two individuals. This way if someone is on vacation there will be at least one person in the know. Also, make sure that you allow only trained personnel to work on an issue. If this is not an option, make sure they have access to someone who is experienced.[12]
- **Spares** — If it was within your budget when you designed the network, you purchased (we hope) and have network spares on hand. These prevent shipping delays when there is a hardware problem. Be careful when you use network spares; make sure not to do a hardware swap if the issue really isn't the hardware.[13]

Your network is only as effective as you are. If you follow all these proactive recommendations, you can help alleviate some of the pain you might experience when troubleshooting without some of this information.[14]

16.3 Troubleshooting Tools

When there is a problem on the network, it is important to have access to some important tools that can assist your troubleshooting. Most network nodes have event logs that assist in diagnosing issues with a specific node. In addition, you can refer to MAC address tables in Layer 2 nodes, IP routing tables in Layer 3 nodes, ARP caches, command-line statistics, etc. Because there are many different types of nodes, and the commands to gather such data are vendor-specific, it is nearly impossible to introduce many of these in this book. These are things you will learn and that will be specific to the vendor your organization uses. Fortunately, there are some tools and utilities that are available for you to use, whether built in, downloadable, or otherwise available for purchase.

16.3.1 Helpful TCP/IP Utilities

If you are using a TCP/IP-compatible node, it will most likely include many of the utilities that we discuss in this section. Examples used in this section are

[12]Your vendor should have a support staff and/or documentation available. Check with your vendor for details.
[13]Sometimes a hardware swap may appear to fix a problem, but in reality the swap ''bounced'' the issue, and it may clear up or return depending on what the problem really is.
[14]Not to mention the potential cost to the organization when an issue occurs.

from a Windows-based PC,[15] so be aware that if you are trying to use them on a non-Windows node, the command may differ from what you see here. The result will still be the same, however.

These helpful utilities provide you with commands to determine whether the node that you are issuing the command from is able to reach a destination, verify the correct path is being taken to a destination, verify name-to-address mapping is correct, verify that MAC-to-IP address mapping is correct, and issue tests to other upper-layer protocols.

16.3.1.1 Ping

Ping is an acronym for *packet Internet groper*. The `ping` utility enables you to test whether a destination node is reachable. The `ping` command is most often issued at the beginning of (and several times during) a troubleshooting session. The ping is an ICMP echo request/reply that determines whether a node is reachable and outputs the round-trip time for the process to complete, any packet-loss percentages, and a statistical summary for a given remote node.

The `ping` command is simple to use. Open your command-line window and issue the command `ping`, followed by a DNS name or an IP address. Following is an example of a successful ping to the DNS name yahoo.com:

```
C:\>ping yahoo.com

Pinging yahoo.com [206.190.60.37] with 32 bytes of data:

Reply from 206.190.60.37: bytes=32 time=23ms TTL=50
Reply from 206.190.60.37: bytes=32 time=21ms TTL=50
Reply from 206.190.60.37: bytes=32 time=22ms TTL=50
Reply from 206.190.60.37: bytes=32 time=22ms TTL=50

Ping statistics for 206.190.60.37:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 21ms, Maximum =  23ms, Average =  22ms
```

Notice that, by default, the `ping` command will send four ICMP requests and will expect four replies. At the end of the session, the statistics are outputted,[16] showing the average success rate, round-trip

RANDOM BONUS DEFINITION

collapsed backbone — A method of interconnecting networks by using a switch or router as a central relay device.

[15] Chances are, this is what you will be using in a networking environment.
[16] Also known as *printed*, or *printed on the screen*.

times, etc. So what happens if we issue a command to a node that ping is unable to reach? Following is such an example:

```
C:\>ping testshow.com

Pinging testshow.com [207.215.79.16] with 32 bytes of data:

Request timed out.
Request timed out.
Request timed out.
Request timed out.

Ping statistics for 207.215.79.16:
    Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum =  0ms, Average =  0ms
```

As you can see, we were not able to get even one reply back from the destination node. The responses (*request timed out*) mean that the reply timer expired before a reply was received. If you receive a *request timed out* response, you cannot automatically assume that something is wrong with the remote node. The problem could be anywhere between your PC and the destination node.

Request timed out is a message that indicates that there is no route to the remote node. One of the first things you can do is issue a ping to the interface of your PC (127.0.0.1) to see whether you get a reply. If you do receive a reply, you know that TCP is running and is working on your PC. The next thing you want to do is ping the next hop router, or your default gateway. Continue pinging until your ping fails; doing so will indicate where the issue is occurring.[17]

You have several options available in most ping utilities. These can assist you in different stages of your troubleshooting session. The options available on a Windows-based PC include the following:

```
Options:
    -t             Ping the specified host until stopped.
                   To see statistics and continue - type Control-Break;
                   To stop - type Control-C.
    -a             Resolve addresses to hostnames.
    -n count       Number of echo requests to send.
    -l size        Send buffer size.
    -f             Set Don't Fragment flag in packet.
    -i TTL         Time To Live.
```

[17] Another option is to use the traceroute utility, which we discuss in Section 16.3.1.2.

```
-v TOS          Type Of Service.
-r count        Record route for count hops.
-s count        Timestamp for count hops.
-j host-list    Loose source route along host-list.
-k host-list    Strict source route along host-list.
-w timeout      Timeout in milliseconds to wait for each reply.
```

As you can see, several options are available to help you narrow down a problem in the LAN. An example of what you can do with the command is to issue a permanent ping to a remote destination. You do this by issuing the `ping` command, the identification ID, followed by `-t`. For example, you want to issue a constant ping to the IP address of 10.10.10.1:

```
C:\>ping 10.10.10.1 -t
```

You can use a constant ping when working on an issue between two endpoints within your network, to see if the node comes up and stays up. Also, some VPN tunnel connections require a constant ping through the tunnel, or the connection between the remote site and the corporate site will be dropped and the tunnel will be brought down.

POP QUIZ

Proactive troubleshooting or reactive troubleshooting — which is better?

16.3.1.2 Traceroute

The `traceroute` utility (`tracert` in Windows) enables you to view the sequence of hops that a packet takes to a destination. Each and every router that the packet passes through is listed in the output of the command. This output will continue until the destination is reached, or when the replies are no longer being sent. Following is an example of the `tracert` command from a PC to the `yahoo.com` domain:

```
C:\>tracert yahoo.com

Tracing route to yahoo.com [206.190.60.37]
over a maximum of 30 hops:

  1     1ms   <10ms   <10ms  192.168.1.1
  2     7ms     7ms     6ms  c-3-0-ubr01.boston.cast.net
                               [43.16.12.1]
  3     9ms     8ms     7ms  ge-1-37-ur01.boston.cast.net
                               [43.16.12.193]
  4     6ms     7ms     6ms  po-20-ur02.boston.cast.net
                               [43.16.12.158]
  5     7ms     7ms     7ms  po-24-ur01.boston.cast.net
                               [43.16.12.161]
```

```
  6     9 ms     9 ms     8 ms  po-21-ar01.needham.cast.net
                                   [43.16.12.157]
  7    14 ms    13 ms    13 ms  te-3-2-ar01.hartford.cast.net
                                   [43.16.12.62]
  8    15 ms    15 ms    15 ms  cr01.newyork.ny.ibone.cast.net
                                   [43.16.12.61]
  9    16 ms    18 ms    15 ms  TenGigabitEthernetNYC1.gblx.net
                                   [43.16.12.217]
 10    16 ms    15 ms    15 ms  te2-4nyc1.gblx.net [43.16.12.237]
 11    24 ms    17 ms    18 ms  NewYork1.Level3.net [44.69.14.13]
 12    26 ms    20 ms    32 ms  Washington1.Level3.net [44.69.12.3]
 13    29 ms    21 ms    31 ms  Washington1.Level3.net [44.69.14.1]]
 14    27 ms    22 ms    20 ms  4.79.228.2
 15    23 ms    21 ms    23 ms  ae2-p140.msr1.re1.yahoo.com
                                   [216.115.108.57]
 16    22 ms    21 ms    21 ms  ge-9-3.bas-a2.re4.yahoo.com
                                   [216.39.49.7]
 17    24 ms    21 ms    22 ms  w2.rc.vip.re4.yahoo.com
                                   [206.190.60.37]

Trace complete.
```

From the responses received, you can see the amount of time it takes for each particular router to respond. If you notice delays, investigate to determine whether the segment that the delay appears in is having issues. Now, let's issue the `tracert` command to the IP address 207.215.79.16. Notice that this is the same address we used in our failed ping above. You should be able to see this `tracert` fail at some point during the session:

```
C:\>tracert 207.215.79.16

Tracing route to www.testshow.com [207.215.79.16]
over a maximum of 30 hops:

  1     1 ms   <10 ms   <10 ms  192.168.1.1
  2     7 ms     7 ms     6 ms  c-3-0-ubr01.boston.cast.net
                                   [43.16.12.1]
  3     9 ms     8 ms     7 ms  ge-1-37-ur01.boston.cast.net
                                   [43.16.12.193]
  4     6 ms     7 ms     6 ms  po-20-ur02.boston.cast.net
                                   [43.16.12.158]
  5     7 ms     7 ms     7 ms  po-24-ur01.boston.cast.net
                                   [43.16.12.161]
  6     9 ms     9 ms     8 ms  po-21-ar01.needham.cast.net
                                   [43.16.12.157]
  7    14 ms    13 ms    13 ms  te-3-2-ar01.hartford.cast.net
                                   [43.16.12.62]
  8    15 ms    15 ms    15 ms  cr01.newyork.ny.ibone.cast.net
                                   [43.16.12.61]
```

```
 9    16 ms    18 ms    15 ms  TenGigabitEthernetNYC1.gblx.net
                                  [43.16.12.217]
10    16 ms    15 ms    15 ms  te2-4nyc1.gblx.net [43.16.12.237]
11    24 ms    17 ms    18 ms  NewYork1.Level3.net [44.69.14.13]
12    26 ms    20 ms    32 ms  Washington1.Level3.net [44.69.12.3]
13    29 ms    21 ms    31 ms  Washington1.Level3.net [44.69.14.1]]
14    16 ms    16 ms    15 ms  ex1-tg2-0.eqnwnj.sbcglobal.net
                                  [151.164.89.249]
15     *        *        *     Request timed out.
16     *        *        *     Request timed out.
17     *        *        *     Request timed out.
```

As you can see in this example, the last hop that we were able to reach is 151.164.89.249. If you were troubleshooting, you would focus on that router to start with to see why it is not able to reach the destination. It is possible that it can reach the next destination, but the next hop may be filtering ICMP traffic, in which case you would not be able to push the ICMP packet through the node that is blocking that type of traffic.

RANDOM BONUS DEFINITION

mirror port — A port configured on a Layer 2 switch used to copy the traffic appearing on another port on the same switch.

Here is a test for you. You are doing a `tracert` and you notice the following:

```
C:\>tracert 207.215.79.16

Tracing route to www.testshow.com [207.215.79.16]
over a maximum of 30 hops:

  1     1 ms   <10 ms   <10 ms  192.168.1.1
  2     7 ms     7 ms     6 ms  c-3-0-ubr01.boston.cast.net
                                  [43.16.12.1]
  3     9 ms     8 ms     7 ms  ge-1-37-ur01.boston.cast.net
                                  [43.16.12.193]
  4      *        *        *     Request timed out.
  5     7 ms     7 ms     7 ms  po-24-ur01.boston.cast.net
                                  [43.16.12.161]
```

Notice that the fourth hop timed out, but the fifth hop didn't. Why do you think this happened?[18]

As with the `ping` command, you have several options available in most `traceroute` utilities. These can assist you in different stages of your

[18]The answer will be provided at the end of Section 16.3.1.2.

troubleshooting session. The options available on a Windows-based PC include the following:

```
Options:
    -d                 Do not resolve addresses to hostnames.
    -h maximum_hops    Maximum number of hops to search for target.
    -j host-list       Loose source route along host-list.
    -w timeout         Wait timeout milliseconds for each reply.
```

Okay, now getting back to the question at hand: Why did the fourth hop time out, but the fifth hop didn't? This happened because the preferred next hop was not available, so an alternative route was taken.[19]

16.3.1.3 Netstat

By default, the `netstat` utility displays both incoming and outgoing network connections. These commands prove useful in troubleshooting issues on the network, and when gathering traffic statistics to measure network performance. Following is an example of the `netstat` command:

```
C:\>netstat

Active Connections

Proto  Local      Foreign                                    State
       Address    Address
TCP    PC:1693    PC:1694                                ESTABLISHED
TCP    PC:1694    PC:1693                                ESTABLISHED
TCP    PC:1663    es.com:http                            ESTABLISHED
TCP    PC:1671    bs1.ads.vip.ac4.yahoo.com:http           TIME_WAIT
TCP    PC:1674    jkhiuyg.com:http                       ESTABLISHED
TCP    PC:1675    a96-17-73-9.deploy.com:http            ESTABLISHED
TCP    PC:1678    a96-17-72-144.depl.com:http            ESTABLISHED
TCP    PC:1680    bs1.ads.vip.ac4.yahoo.com:http           TIME_WAIT
TCP    PC:1683    209.62.185.43:http                     ESTABLISHED
TCP    PC:1685    bs1.ads.vip.ac4.yahoo.com:http           TIME_WAIT
TCP    PC:1687    bs1.ads.vip.ac4.yahoo.com:http           TIME_WAIT
TCP    PC:1690    a96-17-73-27.depogies.com:http         ESTABLISHED
TCP    PC:1695    a96-17-73-35.deplo.com:http            ESTABLISHED
TCP    PC:1707    bs1b1.ads.vip.re2.yahoo.com:http         TIME_WAIT
TCP    PC:1728    8.12.222.126:http                      ESTABLISHED
TCP    PC:1729    8.12.222.126:http                      ESTABLISHED
TCP    PC:1735    8.12.222.126:http                       CLOSE_WAIT
TCP    PC:1737    8.12.222.126:http                       CLOSE_WAIT
```

[19] Thanks to our friend Mr. Redundancy.

```
TCP    PC:1738    64.236.29.103:http                    CLOSE_WAIT
TCP    PC:1739    64.236.29.103:http                    CLOSE_WAIT
TCP    PC:1743    209.62.185.43:http                    ESTABLISHED
TCP    PC:1746    208.215.179.180:http                  ESTABLISHED
TCP    PC:1749    od-in-f166.google.com:http            ESTABLISHED
TCP    PC:1751    66.235.142.3:http                     ESTABLISHED
TCP    PC:1752    157.166.226.31:http                   CLOSE_WAIT
TCP    PC:1754    157.166.224.32:http                   CLOSE_WAIT
TCP    PC:1756    157.166.226.30:http                   CLOSE_WAIT
```

Take a moment to review the data that you can capture with the default `netstat` command:

- **Protocol** — The first field provided is the Protocol field, which identifies whether the connection is TCP or UDP.
- **Local Address** — This is the name of the local node[20] and the port number that is being used.
- **Foreign Address** — This is the name of the remote node and the port number that the socket is connected to.
- **State** — This field identifies the TCP connection state. Following are the possible TCP states:
  - **LISTEN** — Indicates that the node is waiting for a connection request.
  - **SYN-SENT** — Indicates that a connection request has been sent and TCP is waiting for a matching connection request.
  - **SYN-RECEIVED** — Indicates that TCP is waiting for a confirming connection request acknowledgment after having both sent and received a connection request.
  - **ESTABLISHED** — Indicates that the TCP connection is open. This is the normal state for the data transfer phase of the connection.
  - **FIN-WAIT-1** — Indicates that TCP is waiting for a connection-termination request from the remote TCP. It can also serve as an acknowledgment of a connection termination request.
  - **FIN-WAIT-2** — Indicates that TCP is waiting for a connection termination request from the remote TCP.
  - **CLOSE-WAIT** — Indicates that TCP is waiting for a connection-termination request from the local user.
  - **CLOSING** — Indicates that TCP is waiting for a connection-termination request acknowledgment from the remote TCP.

[20]The local node is always the node that you are close to.

- **LAST-ACK** — Indicates that TCP is waiting for an acknowledgment of a connection-termination request previously sent to the remote TCP.

RANDOM BONUS DEFINITION

copy port — Synonymous with mirror port.

- **TIME-WAIT** — Indicates that TCP is waiting for enough time to pass to be sure that the remote TCP received the acknowledgment of its connection-termination request.

As with the other utilities discussed so far, you have several options available in most `netstat` utilities. These can assist you in different stages of your troubleshooting session. The options available on a Windows-based PC include the following:

```
Options:

-a              Displays all connections and listening ports.
-e              Displays Ethernet statistics. This may be combined
                with the -s option.
-n              Displays addresses and port numbers in numerical
                form.
-p proto        Shows connections for the protocol specified by
                proto; proto may be TCP or UDP. If used with the -s
                option to display per-protocol statistics, proto may
                be TCP, UDP, or IP.
-r              Displays the routing table.
-s              Displays per-protocol statistics. By default,
                statistics are shown for TCP, UDP and IP; the -p
                option may be used to specify a subset of the
                default.
interval        Redisplays selected statistics, pausing interval
                seconds between each display. Press CTRL+C to stop
                redisplaying statistics. If omitted, netstat will
                print the current configuration information once.
```

You can gather so much statistical information from the `netstat` command. Try the command a few times on your PC to get an idea about what you can obtain with this helpful utility.

POP QUIZ

What are two types of topology diagrams that come in handy when troubleshooting an issue within your LAN?

16.3.1.4 Route

You can use the `route` utility to view and manipulate the routing table of a Windows-based node. Note that any Layer 3 node will have this capability.

Following is an example of the `route print` command, which can be used to display the routing table of a PC:

```
C:\>route print
===========================================================

Interface List
0x1 ........................... MS TCP Loopback interface
0x1000003 ...00 10 b5 65 4d 1a ...... NDIS 5.0 driver

===========================================================

===========================================================

Active Routes:
Network          Netmask          Gateway      Interface     Metric
Destination
0.0.0.0          0.0.0.0         192.168.1.1   192.168.1.104    1
127.0.0.0        255.0.0.0       127.0.0.1     127.0.0.1        1
192.168.1.104    255.255.255.255 127.0.0.1     127.0.0.1        1
192.168.1.255    255.255.255.255 192.168.1.104 192.168.1.104    1
255.255.255.255  255.255.255.255 192.168.1.104 192.168.1.104    1
Default Gateway:       192.168.1.1
===========================================================

Persistent Routes:
  None
```

In the routing table, you can see the destination IP addresses, the subnet mask, the gateway, the local interface, and the number of hops to the gateway.

Following are the optional flags and commands that can be used with the `route` command for advanced display and configuration options:

```
  Options:

  -f           Clears the routing tables of all gateway entries. If
               this is used in conjunction with one of the commands,
               the tables are cleared prior to running the command.
  -p           When used with the ADD command, makes a route
               persistent across boots of the system. By default,
               routes are not preserved when the system is
               restarted. Ignored for all other commands, which
               always affect the appropriate persistent routes. This
               option is not supported in Windows 95.
  command      One of these:
                 PRINT     Prints  a route
                 ADD       Adds    a route
                 DELETE    Deletes a route
                 CHANGE    Modifies an existing route
```

```
destination  Specifies the host.
MASK         Specifies that the next parameter is the 'netmask'
             value.
netmask      Specifies a subnet mask value for this route entry.
             If not specified, it defaults to 255.255.255.255.
gateway      Specifies gateway.
interface    The interface number for the specified route.
METRIC       Specifies the metric, ie. cost for the destination.
```

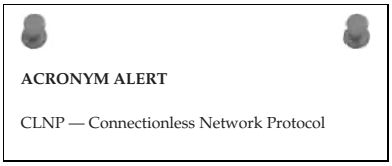

Let's assume that during troubleshooting you discover that you need to add a default static route to your routing table to reach a remote node. The network you want to reach is the 10.10.10.0 subnet. You know that you have a route to that subnet from the gateway node (192.168.1.1), so you want to set up the static route to the 10.10.10.0 network and the 192.168.1.1 as the default gateway. The syntax is

```
C:\>route add <IP address> mask <mask> <gateway IP address>
```

So using the addressing that we discussed previously, the following is the input needed to add the static route:

```
C:\>route add 10.10.10.0 mask 255.255.255.0 192.168.1.1
```

To verify that the route was added, check the routing table:

```
C:\>route print
===========================================================

Interface List
0x1 ........................... MS TCP Loopback interface
0x1000003 ...00 10 b5 65 4d 1a ...... NDIS 5.0 driver

===========================================================

===========================================================
Active Routes:
Network          Netmask          Gateway       Interface     Metric
Destination
0.0.0.0          0.0.0.0         192.168.1.1     192.168.1.104   1
10.10.10.0       255.255.255.0   192.168.1.1     192.168.1.104   1
127.0.0.0        255.0.0.0       127.0.0.1       127.0.0.1       1
192.168.1.104    255.255.255.255 127.0.0.1       127.0.0.1       1
192.168.1.255    255.255.255.255 192.168.1.104   192.168.1.104   1
255.255.255.255  255.255.255.255 192.168.1.104   192.168.1.104   1
Default Gateway:       192.168.1.1
===========================================================

Persistent Routes:
  None
```

RANDOM BONUS DEFINITION

monitored port — A port on a switch that is being mirrored.

As you can see, there is now a route to the 10.10.10.0 network. If you were to reboot your PC at this point, the static route will be removed and will have to be re-added. If you want to ensure that the static route remains until it is manually removed, you need to add the `-p` flag:

```
C:\>route add 10.10.10.0 mask 255.255.255.0 192.168.1.1 -p
```

16.3.1.5 Arp

The `arp` utility enables you to view and manipulate the ARP table of a Windows-based node. Note that any Layer 3 node will have this capability. Following is an example of the `arp -a` (to view) command, which can be used to display a PC's routing table:

```
C:\>arp -a

Interface: 19.108.1.14 on Interface 0x1000003
  Internet Address      Physical Address      Type
  19.108.1.1           00-f8-f8-ea-08-bc     dynamic
  19.108.1.102         00-19-ea-f8-36-95     dynamic
```

We want to force our node to learn the route to an IP that is not listed in the ARP table, so we `ping` the remote node's IP address:

```
C:\>ping 19.108.1.106

Pinging 19.108.1.106 with 32 bytes of data:

Reply from 19.108.1.106: bytes=32 time=23ms TTL=50
Reply from 19.108.1.106: bytes=32 time=23ms TTL=50
Reply from 19.108.1.106: bytes=32 time=23ms TTL=50
Reply from 19.108.1.106: bytes=32 time=23ms TTL=50

Ping statistics for 206.190.60.37:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
```

Next, we will view the ARP table again:

```
C:\>arp -a

Interface: 19.108.1.14 on Interface 0x1000003
  Internet Address      Physical Address      Type
  19.108.1.1           00-f8-f8-ea-08-bc     dynamic
  19.108.1.102         00-19-ea-f8-36-95     dynamic
  19.108.1.106         00-95-6f-af-84-6d     dynamic
```

The options available on a Windows-based PC include the following:

```
Options:

-a            Displays current ARP entries by interrogating the
              current protocol data. If inet_addr is specified,
              the IP and Physical addresses for only the specified
              computer are displayed. If more than one network
              interface uses ARP, entries for each ARP table are
              displayed.
-g            Same as -a.
inet_addr     Specifies an internet address.
-N if_addr    Displays the ARP entries for the network interface
              specified by if_addr.
-d            Deletes the host specified by inet_addr. inet_addr
              may be wildcarded with * to delete all hosts.
-s            Adds the host and associates the Internet address
              inet_addr with the Physical address eth_addr. The
              Physical address is given as 6 hexadecimal bytes
              separated by hyphens. The entry is permanent.
eth_addr      Specifies a physical address.
if_addr       If present, this specifies the Internet address of
              the interface whose address translation table should
              be modified. If not present, the first applicable
              interface will be used.
```

16.3.1.6 Ipconfig

The `ipconfig` utility provides the TCP/IP configuration for the PC and enables you to refresh many of the TCP/IP components (such as DNS). `ipconfig` is the Windows version of this tool, although many operating systems have a similar utility. Examples of this include the `ifconfig` command that is used by Unix and some versions of Macintosh operating systems. Following is an example of the `ipconfig` command:


```
C:\>ipconfig

Windows 2000 IP Configuration

Ethernet adapter Local Area Connection:

        Connection-specific DNS Suffix  . : hs.comcast.net.
        IP Address. . . . . . . . . . . . : 192.168.1.1
        Subnet Mask . . . . . . . . . . . : 255.255.255.0
        Default Gateway . . . . . . . . . : 192.168.1.100
```

As you can see in the above example, the `ipconfig` command provides you with the DNS name pertaining to the ISP connection, node IP address, subnet mask, and the default gateway IP address. A lot of helpful options are available with this command. To view the options, use the `/?` option:

```
C:\>ipconfig /?

Windows 2000 IP Configuration

Ethernet adapter Local Area Connection:

USAGE:
   ipconfig [/? | /all | /release [adapter] | /renew [adapter]
            | /flushdns | /registerdns
            | /showclassid adapter
            | /setclassid adapter [classidtoset] ]

   adapter    Full name or pattern with '*' and '?' to 'match',
              * matches any character, ? matches one character.
   Options
       /?           Display this help message.
       /all         Display full configuration information.
       /release     Release the IP address for the specified
                    adapter.
       /renew       Renew the IP address for the specified adapter.
       /flushdns    Purges the DNS Resolver cache.
       /registerdns Refreshes all DHCP leases and re-registers DNS
                    names
       /displaydns  Display the contents of the DNS Resolver Cache.
       /showclassid Displays all the dhcp class IDs allowed for
                    adapter.
       /setclassid  Modifies the dhcp class id.

The default is to display only the IP address, subnet mask and
default gateway for each adapter bound to TCP/IP.

For Release and Renew, if no adapter name is specified, then the IP
address leases for all adapters bound to TCP/IP will be released or
renewed.

For SetClassID, if no class id is specified, then the classid is
removed.

Examples:
    > ipconfig                      ... Show information.
    > ipconfig /all                 ... Show detailed information
    > ipconfig /renew               ... renew all adapaters
    > ipconfig /renew EL*           ... renew adapters named EL....
    > ipconfig /release *ELINK?21*  ... release all matching adapters
```

Many tools are available to help you test and troubleshoot issues in your network.[21] When troubleshooting end-user issues, these tools allow you to relearn IP addressing, DNS information, and DHCP information.

POP QUIZ

Name a proactive step you can take to help keep the network running.

At the beginning of this section, we told you that the `ipconfig` utility will provide you with the TCP/IP settings. To see a detailed list of these settings, just enter the `/all` option:

```
C:\>ipconfig /all

Windows 2000 IP Configuration

        Host Name . . . . . . . . . . . . : Widget Net.
        Primary DNS Suffix  . . . . . . . :
        Node Type . . . . . . . . . . . . : Broadcast
        IP Routing Enabled. . . . . . . . : No
        WINS Proxy Enabled. . . . . . . . : No
        DNS Suffix Search List. . . . . . : hsd1.nh.comcast.net.

Ethernet adapter Local Area Connection:

        Connection-specific DNS Suffix  . : hs.comcast.net.
        Description . . . . . . . . . . . : blahblahblah-based
                                            Ethernet Adapter
        Physical Address. . . . . . . . . : 1A-10-B5-4D-01-56
        DHCP Enabled. . . . . . . . . . . : Yes
        Autoconfiguration Enabled . . . . : Yes
        IP Address. . . . . . . . . . . . : 192.168.1.1
        Subnet Mask . . . . . . . . . . . : 255.255.255.0
        Default Gateway . . . . . . . . . : 192.168.1.100
        DHCP Server . . . . . . . . . . . : 192.168.1.100
        DNS Servers . . . . . . . . . . . : 67.87.71.26
                                            67.87.71.22
                                            67.87.73.16
        Lease Obtained . . . . . . . . . . : Sunday, October 19,
                                            2008 4:00:00 PM
        Lease Expires . . . . . . . . . . : Monday, October 20,
                                            2008 4:00:00 PM
```

A lot of information is output when you use the `ipconfig /all` command. It is a useful command to view and resolve network connection issues.

[21] As well as connection issues with your PC.

RANDOM BONUS DEFINITION

LAN segmentation — The practice of dividing a single LAN into a set of multiple LANs.

This concludes the discussion of TCP/IP utilities in this chapter. We highly recommend that you try as many of these as you can to build a good understanding of how these tools can help you in the future. The next section travels outside of our TCP/IP world and discusses more specialized equipment and processes that assist in network troubleshooting, diagnosis, and resolution.

16.3.2 More Helpful Tools

The TCP/IP utilities that we discussed in the previous section are great tools when troubleshooting Layer 2 and above issues. But what happens when a cable breaks? There are several tools that are handy in troubleshooting Physical layer issues. Following is a list of some of these tools:

- **Volt-ohm meter** — This device is used to measure voltage, current, and resistance.[22] It is used to test electrical continuity through the physical medium. Volt-ohm meters are often called multimeters, which is entirely appropriate. There are two main types of multimeters, analog and digital. Digital multimeters are popular because they are accurate, durable, and may provide additional functionality not supported by an analog multimeter.

  The volt-ohm meter provides basic troubleshooting assistance, but can be helpful in determining whether you have an issue at the Physical layer. As you will find in Section 16.5, you will often start your troubleshooting sessions by looking for Physical layer issues. Having the capability to test the Physical layer at a moment's notice can save valuable time in the long run.

- **Cable tester** — Network cable testers are used to check the integrity of the cabling in the LAN. Cable testers are made for all types of network cabling (STP, UTP, coaxial, and optical fiber), and some cable testers can even choose the type of cabling that will be tested. There are different types of testers out there, and they vary in functionality. Normally, at a minimum, they enable you to test for crosstalk, attenuation, and noise. More advanced testers can test to ensure that the shield used in STP is not damaged, can display MAC address information, and can capture information pertaining to data rates, errors, and collisions.

[22]Some volt-ohm meters provide additional functionality to broaden their testing ability.

- **Breakout box** — This is a device used to troubleshoot issues over a serial port. The breakout box is placed between two nodes (for instance, between a router and a CSU). The breakout box monitors the signals and uses LEDs to display information pertaining to the signals.
- **Bit Error Rate Test (BERT) tester** — This device is used to generate a test pattern that the tester and a remote node can use to test the line for errors. Both the tester and the remote side can be set to the same pattern so that the integrity of the test is maintained.

It's not absolutely necessary for you to have any of these tools, but considering the affordability of many of these, it sure makes sense to have them around. You might never use them,[23] but they can prove extremely helpful when you are troubleshooting.

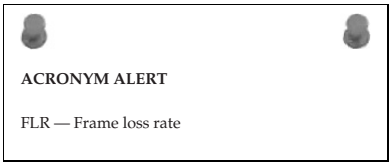

16.3.3 Even More Helpful Tools

These are helpful tools that you will find yourself referring to often when troubleshooting, monitoring, and baselining your network. Of course, some of these tools are optional, and some may not be available on every system or network (for instance, some older LAN bridges did not have the capability of generating an event log). If you do have them available, however, by all means use them:

- **Event logs** — Most network nodes have an event log[24] that reports major changes to the status of any particular function of the node. Event logs can be helpful in troubleshooting an issue. In addition, the event log can be set to report catastrophic events. Often, the node will provide a dump of information that can be analyzed by the vendor.
- **Network analyzer** — The network analyzer is used to capture packets that are being transmitted on the LAN. The analyzer will log the data it receives and can analyze the data based on a specified RFC or standard.

  Network analyzers (*sniffers*) can be software based or can be a node running the analyzer software. Sniffers can be configured to capture all the traffic to and from a specified segment, or can be set to capture specific data only (for instance, only ICMP traffic). Sniffers are

[23]Scratch that — we all know that if you buy one, you will use it, even if you are just fiddling around.

[24]Many upper-layer nodes provide additional logs, such as security log, radius log, authentication log, hardware log, etc. While these are beyond the scope of this book, they should be reviewed when an applicable issue is occurring.

helpful any time you are having an issue with a portion of the network sending or receiving data. They provide a line-by-line analysis of the packets that were captured.

Sniffers are also helpful in troubleshooting protocol specific issues that may be occurring within the LAN. Intrusion attempts by an unwelcome remote party can be detected with the sniffer.

- **Documentation captures** — When troubleshooting, make sure you keep a copy of everything you do. If you are running commands on the command line, do them through a terminal emulator, such as Windows HyperTerminal. At the very least, capture the following (from all applicable nodes) before you do anything else:
  - Node configurations
  - Event logs (and any other available logs)
  - Status of the interfaces of the node
  - Memory-usage statistics
  - Any vendor-recommended documentation
  - Screenshots that may provide insight to the issue

As you are troubleshooting, capture anything you believe may be important. Any details that may help lead to resolution are recommended. In addition, if you have a good idea of what the issue is and can capture proof, do so. Document any error messages you encounter. If at some point you need to contact a vendor or another technical support person, you will have a lot of documentation that they can start working with. It is much easier to bring someone up to speed if you have the visual evidence and the troubleshooting performed thus far (and the results).

POP QUIZ

What does `tracert` do?

16.4 A Logical Order

There are various strategies for troubleshooting issues in a LAN. The strategy discussed in this section provides a systematic approach to troubleshooting. There are variations to this troubleshooting model, but they are all geared to work toward the ultimate goal of resolving network issues in a timely manner. Following is an example of a logical model that you can follow to troubleshoot network issues:

1. Define the problem.
2. Consider the possibilities.

3. Determine the issue.
4. Find a possible solution.
5. Test the possible solution.
6. Develop an action plan.
7. Implement the action plan.
8. Monitor the results.

16.4.1 Define the Problem

Before you do anything, you need a good understanding of what the problem is. Sometimes this is a simple step to take, sometimes it isn't. Without fully understanding problem, you will find yourself going around in circles instead of progressing toward resolution.

Narrow down where the problem exists. How many users are affected? What is the overall impact? How many VLANs are down? What range of IPs cannot be reached? These are just a few questions that you should ask when gathering information. It's your opportunity to play detective and investigate. Most organizations have large, high-speed networks that are the core for the daily operations of the business. You need to isolate the problem not only to help you define what it is, but also to keep it from spreading to other segments of the network.

Don't always believe everything you hear. This is not to say that people are lying. A person may make an assumption and provide you more of an opinion than an explanation. You don't have to discount everything you hear. Instead, investigate the issue to confirm whether you can see what users are seeing (or saying that they are seeing). Don't take everything you hear to heart, and make sure you keep control over everything that is going on, especially when troubleshooting in a group setting.[25]

RANDOM BONUS DEFINITION

learning process — The process whereby a bridge builds its filtering database by gleaning address-to-port mappings from received frames.

16.4.2 Consider the Possibilities

Once you have clearly defined the problem as much as you possibly can, the next step is to start looking at all the things that could be contributing to it. This is when you want to get out all of that baseline information you have

[25]This assumes you are the most experienced. If you are not, then have the most experienced person take control of the troubleshooting session, but make sure you keep one person in charge of all the operations. Otherwise, you may find yourself in a position where several members are making changes and not communicating with one another.

been keeping updated and start comparing the baseline data with the current status of the portion of the network you are troubleshooting. What can you determine from this data that is not normal operation? Narrow this down as much as you can.

It's a good possibility that you will never be troubleshooting an issue alone. Complex issues will sometimes require conference calls, and many other technical and nontechnical people will probably join, all hoping to work toward a resolution. Share the known. Share as much as you can. Remember that the more people who have all the facts, the better. When information is withheld (perhaps because someone doesn't want to admit a mistake) is when troubleshooting gets harder.

Once you have determined all the possibilities that may be causing your problem, start narrowing down the list of possibilities by going through them. Take a hard look at each one and see if you can prove whether each is working as intended. If so, you can cross that item off the list. If not, keep it on the list for further evaluation.

16.4.3 Determine the Issue

Now it's time to narrow down your list and go through the remaining items to see whether they can still seriously be considered as culprits of the problem. Continue to gather as much data as you can. Keep a running copy of your command-line sessions. Pull the system logs, get traces, etc. Any data that can be relevant to the issue, grab it. Document, document, document!

Analyze the data. Start with the Physical layer and then move on up the OSI reference model (more on this in Section 16.5). Once you have ruled out the physical medium, continue layer by layer until you determine where the problem lies. Rule out problems with the cables and the interface cards. Next, ensure that data is flowing across the cables. If it is, you will step up to Layer 3: Is information getting passed via the proper route?

The OSI reference model is a great model to follow when troubleshooting. Although some protocols were built without regard to the model, for the most part it is, and probably will forever be, the model to follow.

16.4.4 Find a Possible Solution

By this point, you should have a good idea of the issue. Once you have determined what the issue is, try to develop a solution. Make sure to consider all possibilities. Sometimes a problem may be a symptom of another issue. Sometimes more than one solution may be possible. In such cases, decide

which is less intrusive and try that one first. Keep a list of other possibilities so you can continue looking into the issue should one proposed solution fail to resolve the issue.

16.4.5 Test the Possible Solution

Once you believe you have a possible solution, try to test it in a lab environment before trying it out in your network. If you don't have a lab to try it out in, you will have to do it on the network. Only in dire circumstances should you work on the network during peak periods.[26] Try to schedule downtime so you can try a solution without affecting too many users.

If you were smart enough to have proposed funds and lucky enough to have received them, you may have enough equipment for a lab. If you do, try to replicate and then test an issue in the lab before rolling the proposed solution out on the network.[27] Doing so provides an opportunity to test the solution, to ensure that the solution works, without negatively affecting the production network.

If you do not have the capability to test the solution, find out if one of the vendors can. If a theorized solution does not resolve the issue in a lab environment, you can work on another solution without any impact to the current network status. Of course, if the network is down and you need to get it back up, you may not have the opportunity to test before putting the testing into action. However, because the network is down, you will have an opportunity to try anything that will get it back up.

16.4.6 Develop an Action Plan

An action plan is one of the most helpful tools you can have available to you when you are implementing a proposed solution to a network problem.[28] Granted, sometimes you might not have the time to build an action plan, but even a verbal action plan is better than making changes without a plan in mind. Action plans should be written in an easy-to-understand manner so that anyone who might be joining the troubleshooting session will have a good idea of what is expected. In addition, the action plan is a great way to coordinate efforts.

The type and amount of data that you include in your action plan depends on what exactly you are working on. The action plan can be as simple as instructions for logging into a node, or as complex as who is involved, where

[26] Sometimes you won't have a choice.

[27] The network during troubleshooting sessions is often referred to as the production network. This is to distinguish it from a lab environment or from any question that is not affecting the network.

[28] An action plan is helpful any time you are making a change on the network.

they are involved, steps to be taken, recovery plans, etc. Once an action plan is drawn up, make sure that you have others review it for completeness and sensibility. It wouldn't hurt to have the vendor involved confirm the plan and offer recommendations. Finally, run through the action plan in the lab (if you have one) to see if your action plan resolves the issue.

Again, having an action plan is not a requirement, nor will you always be able to have one. In the long run, though, you will find that the action plan sure makes things smoother.

RANDOM BONUS DEFINITION

local area network — A network that covers a small geographical area.

16.4.7 Implement the Action Plan

It is now time to roll out the action plan. Set up the conference call, notify the appropriate user groups, and implement the action plan. The action plan should be an easy-to-follow document that lays out, in chronological order, the step-by-step troubleshooting procedure, so the implementation phase should run smoothly.

Make sure you have a backup plan so that if anything does go wrong during the change window you have a way to get the network back to where it was before you started the change.

16.4.8 Monitor the Results

Once you have implemented the action plan, it is time to monitor the results. Utilize all tools that you have for the baseline and monitoring so you can gather and compare statistics to ensure the desired effect has been obtained.

Make sure you allow ample time for all the standards and processes running on the node to make the appropriate decisions before panicking that something isn't working. For example, routing updates take time, so you need to allow the appropriate amount of time before you start worrying about something being wrong.

Test to confirm that the problem you are working on is resolved. Also, check other statistics to ensure that no new issues have arisen (in case a change has fixed one problem but introduced another one). Also be aware that sometimes a fix doesn't work in the production environment, and the troubleshooting process will have to begin again.

POP QUIZ

What does the command `ipconfig /all` provide for you?

16.4.9 Another Fantastic Bonus from the Authors

That's right. Because you might need to have a reference at some point in your networking career,[29] we have included a handy-dandy flowchart to follow (see Figure 16-1).

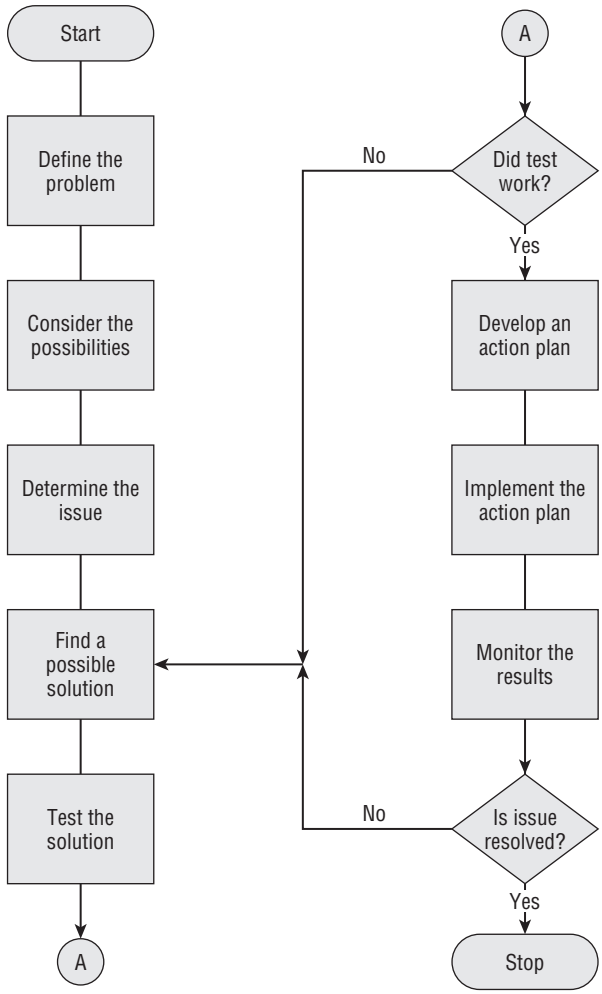

Figure 16-1 The bonus, handy-dandy logical troubleshooting reference flowchart

Continuing on, in the next section we show you how to exploit the OSI reference model when troubleshooting.

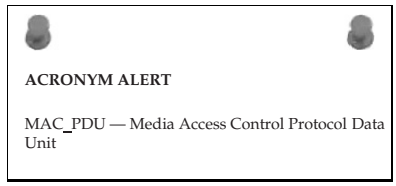

ACRONYM ALERT

MAC_PDU — Media Access Control Protocol Data Unit

[29] We also realized that we had gone too far into the chapter without a picture for you.

TROUBLESHOOTING TIPS

Keep in mind the following tips and follow them to the best of your ability. These tips will help keep the troubleshooting focus headed in a positive direction:

1. **Troubleshoot methodically. Make sure that you stay focused on the task at hand and try not to get sidetracked while you are working on issue resolution. It is easy sometimes to get sidetracked when reviewing and troubleshooting the network. Make sure that you test each step completely before moving to the next step. In other words, test the Layer 1 possibilities fully before moving to Layer 2. It doesn't make sense to make sure that a cable is plugged into every interface but the last hop. Be sure that everything along the way is checked and confirmed.**
2. **Document! Document! Document!**
3. **Approach every possibility with an open mind. Don't discount anything just because you have never heard of it before. Review all tests, and retest if you need convincing, but do not ignore or discount any possibility. Also, don't ever assume. Just because you see a symptom that seems like one you have seen before, you can't be sure it is the same until you confirm that fact.**
4. **Make sure you research error messages fully.**
5. **Remember that there can always be human error. Also, humans are not always aware of, or honest about, something they might have done.**
6. **Whenever possible, test and replicate an issue.**
7. **Keep anyone who might be affected by the trouble aware that the issue is being worked on. Provide a timeline when possible.**

16.5 Layered Strategy

An effective method of troubleshooting is to follow the OSI reference model (see Figure 16-2). Each layer in the OSI reference model communicates with its peer on the remote node (represented by the dotted line in the figure). Each layer passes data down to the next layer[30] for processing, and each layer provides services to the next-higher layer.

It's important that you understand the functions of each layer[31] and some typical issues that you will see within each layer. As you are trying to diagnose the cause of an issue, you should follow the OSI reference model. In this

[30] The exception to this is the Physical layer because that layer doesn't have a lower layer to pass data to.

[31] And by now you should.

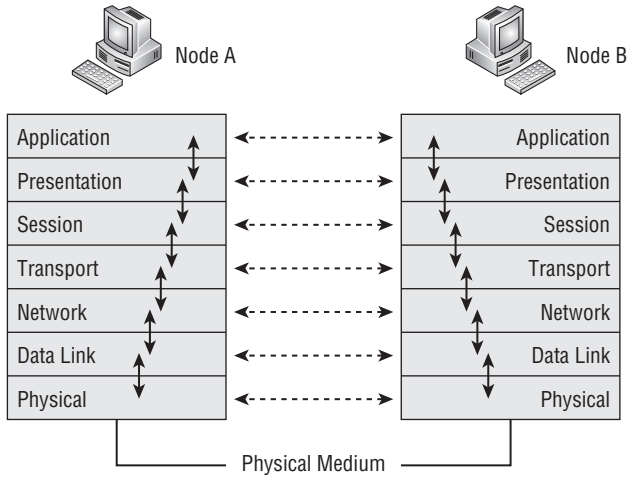

Figure 16-2 The OSI reference model

chapter, we discuss starting at the Physical layer and working your way up,[32] although it really doesn't matter which direction you head. As a matter of fact, you can start at the layer that you suspect is causing the issue and can continue from there. Here is a quick refresher on the roles of the OSI layers:

- **Application layer** — Provides services used by applications (the transfer of files and messages, authentication, name lookups, etc.) within the network.
- **Presentation layer** — Ensures that information received by one host can be read by that host.
- **Session layer** — Sets up, manages, and ends application sessions.
- **Transport layer** — Ensures the transmission of data from one endpoint to another.
- **Network layer** — Provides a path from endpoint to endpoint.
- **Data Link layer** — Provides a way to transport data over a physical link.
- **Physical layer** — Determines the specifications for the operations of the physical links between network devices.

If working from the bottom up, the first layer you check is the Physical layer. Check the cables and the node interfaces to confirm that everything is working correctly. Once you have ruled out the cables and interfaces, check for media access and data encapsulation configurations to ensure that everything

[32]This is because the book is geared toward networking newcomers, and following the OSI reference model from the bottom up makes it easier for a newcomer to work toward a resolution.

is working at the Data Link layer. After the integrity of Layer 2 has been confirmed, check for Network layer issues. Are you having routing problems or problems with addressing in the network? Finally, move on to the upper layers where you verify if the problem is related to memory or buffer issues, authentication, encryption, data compression, etc., all dependent on the layer you are troubleshooting.

As we have stated before, not all standards and protocols follow strict adherence to the OSI reference model. This fact does not change the basic troubleshooting strategy of using the OSI reference model. Regardless of the network issue, the OSI reference model is a good overall base model to follow while troubleshooting.

16.5.1 Common Lower-Layer Issues

Sometimes issues can occur at more than one layer. (For example, you can have a configuration issue in both the Data Link and the Network layer.) Other issues can occur at only one layer of the OSI model (for example, a dirty fiber cable). At the upper layers, many of the processes can overlap layers, which can make the job of troubleshooting a challenge.

RANDOM BONUS DEFINITION

network — A set of nodes that connect to one another over a shared communication link.

In this section, we discuss common issues that relate to the lower three layers of the OSI reference model.

16.5.1.1 Layer 1

Often a network communication error can be caused by something as simple as a cable not being plugged in. This is why we recommend that one of the first things you do is to check the items that function at the Physical layer. Make sure that your cable is plugged in and that there are no errors (for instance, CRC errors, input/output errors, buffer failures, excessive collisions) accumulating on the interface. Most of these errors can be detected with some simple commands that are run within the command-line interface of a given network node. In addition, most network nodes have LEDs that indicate the status of a given interface. Check with the vendor for the definitions of the LEDs for any particular node.

Common Layer 1 issues that can occur in a network include the following:

- Damaged cables
- Dirty fiber

- Excessive signal attenuation
- Insufficient bandwidth
- Electrical interference
- Wireless interference
- Damaged interface
- Dirty interface

You can find yourself chasing the wrong issue if the Physical layer functions are not checked and verified operational. Just think about it. It makes no sense trying to figure out why you are not able to route to a specific subnet if the problem is as simple as a cable that someone has kicked out.

16.5.1.2 Layer 2

In a given LAN, there are plenty of Layer 2 nodes. Configured correctly, their operation should be transparent to end users, meeting and often exceeding expectations as they perform the function of getting frames from point to point. When a connection issue is occurring and the Physical layer has been ruled out, take a look at the processes operating at Layer 2 to see whether the issue originates or resides at this layer.

So, what will you want to look at? One thing to note is whether you are having an issue with a particular VLAN or are seeing a loop in the network. Then you need to focus on what is occurring at Layer 2 before continuing with your troubleshooting.

Common issues that occur at Layer 2 include the following:

- Configuration error
- Denial-of-service (DoS) attack
- VLAN configuration error
- Class-of-service issue
- Excessive utilization
- Excessive errors
- VLAN issue
- Spanning tree issue
- MAC address table issue
- Hardware compression issue
- Software compression issue
- VRRP issue

In networks where data is transported at both Layer 2 and Layer 3, it is important to ensure that the Layer 2 services are operational before assuming that a routing issue is due to a Layer 3 issue. Make sure that all Layer 2 functions are verified for any IP connectivity issue.

16.5.1.3 Layer 3

You have determined that the issue does not reside at Layer 1 and Layer 2, so now you want to focus on Layer 3. The simplest way to tell if Layer 3 routing is working is to issue a ping to a remote portion of the LAN. If you are having problems communicating with a particular node or subnet, issue the ping to an IP address of the affected area to see what the results are.

Most routers will provide a command-line interface in which you can issue commands to view the status of your IP interfaces. These `show` commands will be valuable in reporting the current status of particular interfaces and will assist in narrowing down the source of an issue. Compare the interface information with the information on the remote side of the connection to make sure they match. For the IP interface to be up and operational, the following holds true:

- The interfaces must be on the same IP network.
- The interfaces must have the same IP subnet masking scheme.
- Filters should be checked to make sure that no rules are configured for the interface. If there are rules, they need to be checked to ensure they are correct.
- Make sure that the interface configuration parameters are correct.

If all these items are checked and verified, the issue may reside in other layers, processes, or external to the node. Following are other potential issues that may be occurring at this layer that could be the cause of an IP issue in the network:

- ARP issue
- IP addressing issue
- Configuration error
- Authentication issue
- VLAN configuration error

Keep in mind that for one LAN to communicate with another, there has to be a route to the destination. The route to the destination is either static

or dynamic, depending on the environment in which the router is placed. A helpful tool in troubleshooting is the configuration of a static route. Because of this manual configuration capability, network connectivity can often be restored with a static route, at least until the dynamic routing issue is resolved.

Layer 3 nodes (or nodes that provide Layer 3 services) have TCP/IP utilities that enable users to view the routing table, purge the routing table, etc. The routing table, ARP table, and the many `show` commands offered by the nodes are helpful in troubleshooting Layer 3 issues.

POP QUIZ

What are the eight logical steps that we provided in this chapter that you can use to troubleshoot an issue on the LAN?

16.5.2 Thoughts Pertaining to the Upper Layers

If you are able to successfully ping between nodes in different LANs, it is safe to assume that the connectivity issue is not a problem with Layers 1 through 3. If you can get a ping across, but some other things are not working, the problem is in the upper layers. Following are some of the things that you can look at in each layer:

- **Layer 4** — At this layer, focus on whether TCP or UDP is operating as intended. Take a sniffer trace and determine whether acknowledgments are being sent in response to requests. Also, check whether fragmentation is working as intended. Finally, check to see if any filters or QoS parameters may be affecting the flow of data.
- **Layer 5** — Things to look for at this layer include whether the Session layer protocols are receiving errors while trying to communicate. A sniffer trace can be used to determine if the protocols are behaving correctly.
- **Layer 6** — Encryption, formatting, and compression of data occur in this layer. Is data being encrypted/decrypted appropriately? Are encryption configuration settings correct? Are the correct data formats in use?[33] Another concern at this layer is whether a VPN tunnel is operating as it is configured to do.
- **Layer 7** — Are applications working correctly? Sometimes a version of a standard may support new features, and older clients may no longer interoperate with the new versions. Also, end users may still try to connect to a server with the incorrect client, and, of course, the user will not be able to connect.[34]

[33]In other words, is the end user viewing the data in the same format as it was sent? Make sure that the end user is using the correct program to view the data.
[34]You just cannot Telnet to a destination when the destination requires SSH.

RANDOM BONUS DEFINITION

overprovisioning — A technique of providing more capacity than is actually needed for a given application.

As you can probably see, there is a method to this madness. Some issues require a bit of thought and investigation, whereas others can become second nature. Following the OSI reference model may or may not be for you, but it is a tried and true method . . . and methods really do help.

AN UNRELATED MOMENT OF PAUSE

Jim once worked with another engineer who received a phone call from a customer one day. The customer was insistent on getting a field engineer out to his site because the network was having a service-affecting issue. When the engineer arrived, the customer announced that the issue was no longer happening. The field engineer offered to take a look to ensure that the problem was resolved, but the customer kept turning him down.

After a few moments of discussion, the customer reached up on a shelf and produced a router that was no longer serviced by our company. The customer handed the router to the engineer and said, "Well, since you are here, you can help us configure this router." Although we were never able to prove it, we suspect that the customer told a little story to get an engineer on site and then to bully him into configuring the router. If that was the motive, it worked.

After a few years, this customer upgraded their entire LAN. New equipment, with all the latest and greatest (at the time) bells and whistles. With the upgrades, our company recommended that the customer begin training their support engineers on the new gear that would now be supported. The customer didn't think that this training was all that important, so they decided that the engineers would be trained on the job, rather than in a structured training process.

To make matters worse, multiple nontechnical managers took over the implementation phase. The implementation of the network wasn't the smoothest that ever occurred. The problem that ended up happening was that many of the engineers would make changes that would cause an issue, and no one would admit that they had made the change. The fear of being fired in the high-stress environment was very real, so many would not say a thing.

The moral of the story is to listen to others, but make sure you have proof.

16.6 Troubleshooting Examples

There is no way that we could list every possible scenario that you might come across when troubleshooting within your LAN. There are so many different vendors, standards, and configuration possibilities that the task of

troubleshooting this week will be far different from what it was five years ago or five years from now. We could go into deep discussion about troubleshooting IPX and you might not ever come across a network running that protocol.

We do think it might be helpful if we run through a few possible situations you might come across in your LAN, and some recommended steps you can take to reach a resolution. We recommend that you follow the command examples we have listed throughout this section.

16.6.1 Example 1: PC Can't Connect

A user on your LAN is not able to connect to the network. The computer comes up, but the user reports that no network connection can be obtained. This is considered a critical issue from the user's standpoint, because she cannot access servers that run applications she needs to do her job.

So far, you know that there is at least one user who is unable to connect to the network. Ask the user if others are affected; if not, you can assume the issue is local to the one user.

Have the user verify that there are LEDs on the network adapter. If there are, have her describe them. Are they blinking or solid? What color are they? Generally, a solid green light means an interface is up and there is a medium (wireless/physical medium) present. A blinking green light normally indicates that there are datagrams being passed through the interface. Any other color or activity indicates abnormal behavior. If the LEDs are not working correctly, have the user ensure that the network cable is properly connected.

If you have checked the interface and the connection and they appear to be intact, but there is still no activity, you can assume that the cable is bad, that the card is bad or is not configured correctly, or that the interface is not receiving a signal from the network for some other reason. Check the network card. Make sure it is inserted into the node correctly and is snug. Check the configuration of the interface and the card and confirm that they are configured correctly and that they match the next hop's interface settings.

Issue a ping to the interface itself. The well-known IP address for a local host interface is always 127.0.0.1. For example:

```
C:\>ping 127.0.0.1

Pinging 127.0.0.1 with 32 bytes of data:

Reply from 127.0.0.1: bytes=32 time<10ms TTL=128
Reply from 127.0.0.1: bytes=32 time<10ms TTL=128
Reply from 127.0.0.1: bytes=32 time<10ms TTL=128
Reply from 127.0.0.1: bytes=32 time<10ms TTL=128
```

```
Ping statistics for 127.0.0.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum =  0ms, Average =  0ms
```

The results tell you that your interface is up and your NIC is working as it should. If there had not been four successful replies in the ping test, you could assume that either the NIC is not inserted correctly or that the NIC is failing. Next, check the user's PC to make sure that the correct protocols are set up and are running and that no firewall policies are preventing network connectivity.

POP QUIZ

What are three common issues you can have at the Physical layer?

If there is nothing local to the user's PC that is preventing her from connecting to the network, it is time to migrate to the broadcast domain to see if issues at that level are causing the problem. Broadcast domain issues can be tricky to diagnose. Thank goodness that VLANs help keep the size of the domain down. Your network topology diagram will be needed when you are having an issue within a VLAN. The topology diagram ensures that you are looking at all possibilities within the VLAN. Otherwise, you will have to rely on memory and will have to explain the topology (possibly multiple times) to others who are assisting in the troubleshooting. Always remember that the logical topology is tougher to troubleshoot than the physical topology.

When an issue arises within the VLAN, the first step is to check all the members in the VLAN to see if there is one in particular that exhibits abnormal behavior. Begin troubleshooting bridges that are in the middle and then move outward. Following logical steps, try to pinpoint the issue:

1. Is there any maintenance going on?
2. Have there been any recent changes that may have created the issue?
3. Verify the configurations of the bridges.
4. Check statistics. Review logs and traces. Utilize `show` and `debug` commands that are available and are applicable.
5. Is this a new configuration? Make sure that all necessary configuration parameters are set correctly.
6. Are the VLANs configured correctly? Make sure that all VLAN rules were followed. Have there been any changes made? Did someone delete a VLAN, or are there other issues going on?

7. Are tagging rules applied correctly?
8. Is a routing issue preventing devices within the VLAN from communicating with other devices?

RANDOM BONUS DEFINITION

port number — A locally assigned, bridge-unique number that identifies each port on the bridge.

Regardless of the variables that apply to your LAN, this section gave you a good baseline strategy to work from when finding the cause of a user connection issue in any given LAN. However, there may be other steps that come into play, depending on the installed protocols, so make sure that you cover anything that may be applicable to your LAN.

16.6.2 Example 2: Reading a Sniffer Trace

Having the ability to read a sniffer trace is a must in any networking career. It is one of the most effective ways to prove an issue exists and what might be happening with a particular process you are analyzing. We use Wireshark version 1.0.4 for the examples in this section. If you have not done so already, we recommend that you download the latest stable version and follow along. (You need to learn how to use a packet analyzer at some point. If for some reason you don't want to use Wireshark, there are a lot of other free options. Wireshark is available at `www.wireshark.org/download.html`.)[35]

Once you have Wireshark (or some other traffic analyzer) installed on your PC, open it and start a capture.[36] Capture packets on the interface that you are using on your PC. The PC that was used in capturing the examples is a Microsoft Windows platform. The interface that is used on this PC is a NDIS 5.0 NIC (Network Driver Interface Specification, version 5.0), which is the interface architecture included in many Microsoft Windows packages. Make sure the traffic analyzer is set to capture and display the packets in real time so you can watch the packets as they are hitting the interface.

Once the traffic analyzer is running and capturing packets on the interface, we issue a `ping` command to Rich's website (richardbramante.com):

```
C:\>ping richardbramante.com

Pinging richardbramante.com [68.180.151.74] with 32 bytes of data:

Reply from 68.180.151.74: bytes=32 time=100ms TTL=49
```

[35]While you are at it, download the users guide for the software version that you will be using. The users guide is very informative and will explain the features and how to use them.

[36]Because all the basic functions of Wireshark are beyond the scope of this book, we are not going into the details of how to start the capture. The users guide will help you through this. If you are generally quick to learn menu navigation, the process is fairly straightforward.

```
Reply from 68.180.151.74: bytes=32 time=99ms TTL=49
Reply from 68.180.151.74: bytes=32 time=102ms TTL=49
Reply from 68.180.151.74: bytes=32 time=100ms TTL=49

Ping statistics for 68.180.151.74:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 99ms, Maximum =  102ms, Average =  100ms
```

Now take a look at the example in Figure 16-3. This is a picture of the trace that we captured when issuing the above command. As you can see in the figure, there is a DNS query for richardbramante.com. A DNS response was sent back with the IP address associated with the domain name richardbramante.com. Once the IP address is located, you can see the four successful ICMP echo requests and the corresponding replies to each of this.

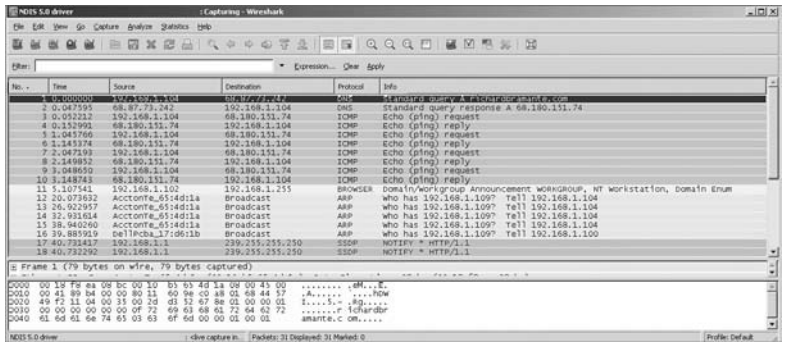

Figure 16-3 Viewing the sniffer trace

Wireshark (and most other analyzers) provides three views to the user. In this description, we used the term *window*, which means a section of the application. You are sure to see what we mean as we go along. The top window of the application shows the list of packets that have been captured, the middle window provides a tree of information for a selected packet, and the bottom window lists the byte information for the packet. Take a look at Figure 16-4 for an example of these windows.

In the example, we highlighted the fourth packet in the trace that we took. This was an ICMP reply packet from source node 68.180.151.74 to the destination node 192.168.1.104. Click on the packet in the upper area of the application to display a tree listing of the details for the fields of the packet. If you can remember the frame format of an ICMP reply, that'll help because this is the data that is provided in each of those fields. The bottom area is the byte information breakdown.

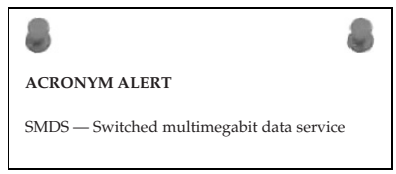

ACRONYM ALERT

SMDS — Switched multimegabit data service

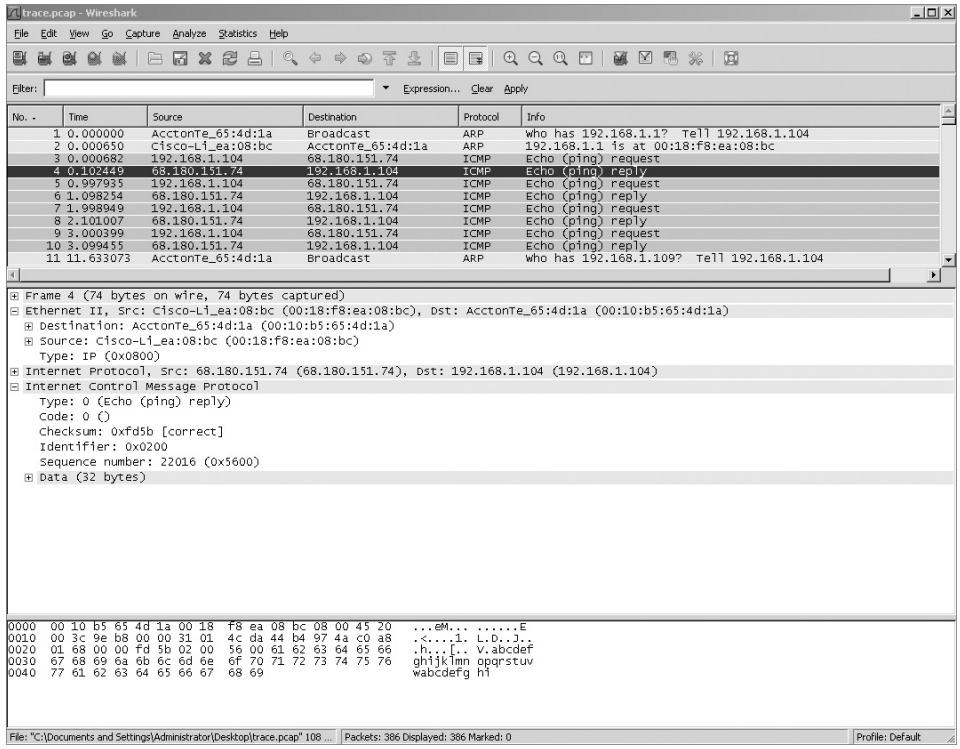

Figure 16-4 Viewing the sniffer trace details

Now that you have had an opportunity to see how a trace is taken and how to break down the information in the packet when you are analyzing/troubleshooting your LAN. Now, we highly recommend that you start another new trace (save your old one for future reference), and this time start multiple processes. Send an e-mail, open an HTTP session, do an FTP, and ping a remote destination (why not richardbramante.com?). While these processes are running, watch the packets processing through the traffic analyzer. This is a great way to better understand the processes of many of the applications used by the users of the LAN.

Another thing you can do is download example traces, which are available in many locations online. These are a good way to learn how to recognize some issues that you may come across. To locate these, just do an Internet search for ''packet capture example.''

16.6.3 Example 3: Identifying a Broadcast Storm

Network broadcasts and multicasts are a necessary operation for any LAN. If you have baselined your network, you should have a good idea of what is a normal amount of broadcasts or multicasts for your LAN. Therefore,

you should notice when you are having more broadcasts than usual. (This condition could be catastrophic if not treated.)

If you are working on a LAN that is starting to run slower than normal,[37] it is a good possibility that you are experiencing a broadcast or multicast storm. Other symptoms of a storm include

- Network operations are timing out.
- Users are not able to connect to the network.
- Users are not able to log on.
- Application access is slow or is not available.
- The network is down.

A network could be experiencing a broadcast storm for several reasons. Hardware may have failed, there may be a configuration error on an interface or a node, spanning tree may have failed, a new application may be causing the issue, a virus or a worm may be attacking the LAN, and many others.

Broadcast storms can be identified in a sniffer trace and via network management systems. Finding the originator can be tricky at times. If the storm is propagating and is affecting multiple areas, you might need to have some coffee and donuts on hand; it's going to be a long night. Once you find the offending interface, you can disable it, and that should calm the storm. A rogue node on the network can be removed after it is located. A worm can be filtered and eventually eradicated with software upgrades and installations. The list goes on and on.

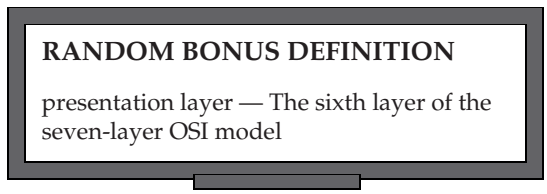

In the case of disabling an interface, you need to track down where the node is located. You can do so with information contained in the trace that you have taken. Some network management applications will also do the tracking for you. Once you have located the offending interface, disable it. You can do that through the command line or you can pull out the cable. This can always be reversed, once you are able to clear whatever is causing or contributing to the storm.

16.6.4 Example 4: VPN Client Can't Connect to VPN Server

A VPN client reports that they are unable to connect to the VPN server. They have repeatedly tried to connect and have also tried to reboot their PC, but

[37] Jim likes the term *sluggish*.

they are not able to connect. The first thing the user needs to check is whether the VPN client is giving him an error message. If there is one, often it can be looked up at the vendor's website to see what the issue might be.[38] Finally, have the user make sure he has an active connection to the Internet. Has he rebooted his router/modem? Once you have verified that everything is operating correctly on the client side, whether there have been any changes on that end? These changes can be PC modifications, change of service provider, connection setting error, etc. Client issues can be tricky, because the user of the client software often has the capability to manipulate files on the PC. By making changes to their configuration, a conflict may arise that causes issues.

Everything on the client side has been checked and verified that the settings are correct and everything is operating as it should. If the client application allows for local logging, ask the user to set this up so he can log a connection attempt on his end. Also, the VPN server log should be checked for messages pertaining to the user that are on the VPN client side.

Client connection issues can range from a configuration error to a failed authentication attempt. Most of the time, an error will be displayed and should identify the problem that is occurring. If there is no error, check the logs. If you can run a sniffer trace on the client side (and possibly even the public interface of the VPN server), you may be able to see more indications of what the issue may be.

If all else fails, the VPN client software may have to be reloaded on the user's PC. This would have to be accomplished by following the rules of the LAN.[39]

POP QUIZ

What are some Network layer issues that you may come across in a LAN?

16.6.5 Example 5: Two Common LAN Issues

In Section 16.6.1, we introduced a single user's connectivity issue from the application that was running on the user's PC to the VLAN and routing domain that the user belonged to. VLAN issues in a bridged LAN are not the only common issue you might come across. *Duplex mismatching* and *spanning tree loops* are issues that you will come across at some point in your networking career.

In this section, we provide some information on these two issues and some ways to diagnose and resolve the issue. If you have access to a lab, you can

[38]Often the error is self-explanatory. For instance, invalid password means that the password is not right.

[39]Hopefully an automatic process or a script has been developed so the installation process is automatic and as transparent to the end user as possible.

easily replicate these issues and take some traces to view the processes and learn how to recognize them.

16.6.5.1 Duplex Mismatch

One of the most common issues that arise in an Ethernet LAN is a mismatch in duplex settings. If you have one end set to full-duplex while the other end is set to half-duplex, you will see latency when passing data on that link. It's always a good rule of thumb to run *autonegotiation* on your switch ports that connect to end users. Regardless of what PC they may be using and how it is configured, the switch will be able to recognize and mirror the settings of the other end. When autonegotiation was first introduced as a standard, there were a lot of issues with it working with some bridges that were already installed in networks at the time. Because of this, many network administrators hard-set the duplex settings on nodes in the internetwork. When possible, the settings should be set to 100 MB/full-duplex to ensure maximum performance from the link.

Symptoms that you may see that indicate a duplex mismatch include the following:

- FCS/CRC errors
- Runt datagrams
- Late collisions

The hard-coding of the duplex settings requires that the opposite end is set to the same duplex and speed settings. Although simple in concept, if the network grows, maintaining these settings at both ends might prove a real headache. In addition, if a device is set to autonegotiation, it will not communicate with any device that is not set to autonegotiate. It will recognize the speed of the link, but not the duplex settings.

Because the packets can be sent when you have a duplex mismatch, you cannot rely on a ping test to determine whether a problem exists. When you have a duplex mismatch, data can still be passed over the link, but will have failures when data is being forwarded from both ends of the link. Remember, in TCP many datagrams will send an acknowledgment to the sender, causing traffic to be sent in both directions on the link. The full-duplex side will receive the acknowledgment datagrams, but the half-duplex side will not and will recognize the link as having a series of collisions. This will cause most of the datagrams that were sent by the host on the full-duplex side to be lost. Another symptom you will

see when there is a duplex mismatch is that the transfer of data is slower than expected when passing large amounts of data (in either direction).

The simplest way to check for duplex mismatches is to look at the configuration settings on both ends of the link.

16.6.5.2 Spanning Tree

A failure in spanning tree usually creates a loop within the area that the spanning tree group covers. The loop is often called a spanning tree loop, although this really does not make sense because spanning tree is there to prevent, not to cause loops. Therefore we should all start referring to it as a *lack of spanning tree loop*. So what does this loop mean? It means there is probably a port within the spanning tree that is forwarding traffic when it shouldn't be. Remember that a port can change from blocking to forwarding when it does not receive a better advertised BPDU from the designated bridge and therefore elects itself as the designated bridge.

Another problem that might be occurring is a duplex mismatch. If there is a duplex mismatch within the spanning tree area, the resulting collisions on the half-duplex side can cause BPDUs to not reach the other bridges and will therefore trigger a loop in your spanning tree. In addition, if the physical medium is reporting errors, this will cause packet corruption, which could cause excessive BPDU loss. Finally, a bridge could get overloaded and run out of resources, preventing it from sending out BPDUs, although this is not as common. Usually, forming and sending the BPDUs is not too resource intensive and will receive a priority over processes that consume resources.

RANDOM BONUS DEFINITION

Rapid Spanning Tree Protocol (RSTP) — An enhancement to the original Spanning Tree Protocol; provides for a faster convergence when there is a change in the topology of the spanning tree group.

If you experience a loop in your bridged network, you will most likely discover this through system statistics, with a sniffer trace, or because everyone is reporting issues. If you have effectively baselined your network, the network topology diagrams can prove very helpful when troubleshooting the loop. In addition, it helps to know where the root bridge is located in relation to the diagram, and which ports should be blocking, which ones are redundant links, and which are not utilized on the spanning tree. As mentioned previously in this chapter, this information will help you when you are reviewing the traces and the statistics while trying to trace down the problem. This will also assist anyone who may not be familiar with your network.

To get out of the loop as soon as possible, disable redundant links (beginning where you suspect most of the looped traffic is occurring). Although this can be time consuming, it will tell you where your problem lies. To disable the link, you can administratively do so or just pull the link and see if the loop dissipates. If the loop does not dissipate, you can put that link back in and move to the next one.

You also want to rely on your system event logs, system statistics, and any traces that you have captured. Some switches also allow options to verbosely log specific events, so you can capture more specific information in your event logs. Another thing to keep in mind when troubleshooting issues in multivendor environments is that some proprietary standards may cause problems with data passing between different vendor bridges. If your environment is such, check with the vendor for any compatibility issues that may exist.

WASN'T THIS BOOK A SNAP?

We know (well, are pretty sure) that you have enjoyed this book, but not as much as you are going to enjoy these cookies if you make them. We recommend that you bake up a couple of batches to munch on while you go over the 200+ bonus questions that follow in the appendixes. This recipe makes about 48 cookies. It was added to this chapter mainly because Jim and Rich have a diet competition going on, and Jim is trying to tempt Rich into baking and eating these cookies. It's also a bonus for you, the reader. Thank you for reading this book. We hope it really was a snap!

Ingredients:

- **1 cup brown sugar, packed**
- **1/4 cup molasses**
- **3/4 cup vegetable oil**
- **1 egg**
- **2 cups flour**
- **2 tsp. baking soda**
- **1/4 tsp. salt**
- **1/2 tsp. ground cloves**
- **1 tsp. ground ginger**
- **1 tsp. ground cinnamon**
- **1/3 cup sugar (for decoration)**

(continued)

WASN'T THIS BOOK A SNAP? *(continued)*

Preparation steps:

1. **Preheat oven to 375 degrees.**
2. **In a large bowl, mix the brown sugar, oil, molasses, and egg.**
3. **In a separate bowl, combine the flour, baking soda, salt, cloves, cinnamon, and ginger; stir into the molasses mixture.**
4. **Roll the dough into 1-1/4 inch balls.**
5. **Roll each ball in the white sugar and place them on an ungreased cookie sheet, 2 inches apart. Do *not* press the cookies down.**
6. **Bake for 8 to 10 minutes in the preheated oven. The cookie is done when the center becomes firm.**
7. **Cool the cookies on wire racks.**

16.7 Chapter Exercises

1. For each item in the following list, identify the layer of the OSI reference model that item applies to:
   - Damaged cables
   - Dirty fiber
   - Excessive signal attenuation
   - Insufficient bandwidth
   - Denial-of-service (DoS) attack
   - Electrical interference
   - Wireless interference
   - Damaged interface
   - Dirty interface
   - Configuration errors
   - Authentication issues
   - Excessive utilization
   - Excessive errors
   - VLAN configuration errors
   - Class-of-service issues

2. In the following example, explain why there is a missing hop.

```
C:\>tracert 207.215.79.16

Tracing route to www.testshow.com [207.215.79.16]
over a maximum of 30 hops:

 1   1ms  <10ms  <10ms 192.168.1.1
 2   7ms   7ms   6ms c-3-0-ubr01.boston.cast.net
                 [43.16.12.1]
 3   9ms   8ms   7ms ge-1-37-ur01.boston.cast.net
                 [43.16.12.193]
 4    *     *     *   Request timed out.
 5   7ms   7ms   7ms po-24-ur01.boston.cast.net
                 [43.16.12.161]
```

3. True or false: The UDP connection state is one of the fields displayed with the `netstat` utility.
4. Network _____ testers are devices used to check the integrity of LAN cabling.
5. List three options that can be used with the `arp` command in Windows, and what each one does.
6. The network analyzer is also known as a _____.
7. What are three ways you can baseline your network? Describe each method.
8. What command in a Windows environment enables you to retrieve detailed information about the current TCP/IP settings of your PC?
9. From your PC, open a traffic analyzer and point it to the interface of your PC. Issue a constant ping to richardbramante.com (`ping richardbramante.com -t`). Next, disable your network connection. What do you notice in your trace?
10. This chapter listed eight quick checks that you can do when you are having issues within a VLAN. What are they?

16.8 Pop Quiz Answers

1. *User* feedback is a great way to be notified of an issue on the LAN.
2. Name 10 issues that you might have on the LAN.
   - Damaged cables
   - Dirty fiber
   - Excessive signal attenuation
   - Insufficient bandwidth

- Denial-of-service (DoS) attack
- Electrical interference
- Wireless interference
- Damaged nodes.
- Damaged interface
- Dirty interface
- Configuration error
- Authentication issues
- Excessive utilization
- Excessive errors
- VLAN configuration error
- Class-of-service issue
- Quality-of-service issue

3. Proactive troubleshooting or reactive troubleshooting — which is better?

   The proactive approach beats the reactive approach hands down!

4. What are two types of topology diagrams that come in handy when troubleshooting an issue within your LAN?

   Logical diagram and physical diagram

5. Name a proactive step you can take to help keep the network running.

   - **Shared knowledge** — It's always good to share what you know. Have discussions with others when you are not sure, and freely offer any helpful information you may have with anyone who might have a need to know.
   - **Proper tools** — Make sure that you and your staff have access to the correct tools. Examples include backup laptops, serial cables, baseline documentation, etc. The staff can only be as effective as their tools allow them to be.
   - **Knowledge requirements** — Make sure there is no single point of failure in the network. Make sure that any and all tasks are performed by at least two individuals. This way if someone is on vacation there will be at least one person in the know. Also, make sure that you allow only trained personnel to work on an issue. If this is not an option, make sure they have access to someone who is experienced.
   - **Spares** — If it was within your budget when you designed the network, you purchased (we hope) and have network spares on

hand. These prevent shipping delays when there is a hardware problem. Be careful when you do use network spares; make sure not to do a hardware swap if the issue really isn't the hardware.

6. What does `tracert` do?

   Provides you with a visual view of the sequence of hops that a packet takes to a destination

7. What does the command `ipconfig /all` provide for you?

   A detailed listing of the TCP/IP settings on your PC

8. What are the eight logical steps that we provided in this chapter that you can use to troubleshoot an issue on the LAN?

   1. Define the problem.
   2. Consider the possibilities.
   3. Determine the issue.
   4. Find a possible solution.
   5. Test the possible solution.
   6. Develop an action plan.
   7. Implement the action plan.
   8. Monitor the results.

9. What are three common issues you can have at the Physical layer?

   - Damaged cables
   - Dirty fiber
   - Excessive signal attenuation
   - Insufficient bandwidth
   - Electrical interference
   - Wireless interference
   - Damaged interface
   - Dirty interface

10. What are some Network layer issues that you may come across in a LAN?

    - ARP issues
    - IP addressing issues
    - Configuration errors
    - Authentication issues
    - VLAN configuration errors