CHAPTER 3

Characterizing the Existing Internetwork

According to Abraham Lincoln:

> If we could first know where we are and whither we are tending, we could better judge what to do and how to do it.

An important step in top-down network design is to examine a customer's existing network to better judge how to meet expectations for network scalability, performance, and availability. Examining the existing network includes learning about the topology and physical structure, and assessing the network's performance.

By developing an understanding of the existing network's structure, uses, and behavior, you can determine whether a customer's design goals are realistic. You can document any bottlenecks or network-performance problems, and identify internetworking devices and links that will need to be replaced because the number of ports or capacity is insufficient for the new design. Identifying performance problems can help you select solutions to solve problems as well as develop a baseline for future measurements of performance.

Most network designers do not design networks from scratch. Instead, they design enhancements to existing networks. Being able to develop a successful network design requires that you develop skills in characterizing an incumbent network to ensure interoperability between the existing and anticipated networks. This chapter describes techniques and tools to help you develop those skills. This chapter concludes with a Network Health Checklist that documents typical thresholds for diagnosing a network as "healthy."

CHARACTERIZING THE NETWORK INFRASTRUCTURE

Characterizing the infrastructure of a network means developing a network map and learning the location of major internetworking devices and network segments. It also includes documenting the names and addresses of major devices and segments, and identifying any standard methods for addressing and naming. Documenting the types and lengths of physical cabling, and investigating architectural and environmental constraints, are also important aspects of characterizing the network infrastructure.

Developing a Network Map

Learning the location of major hosts, interconnection devices, and network segments is a good way to start developing an understanding of traffic flow. Coupled with data on the performance characteristics of network segments, location information gives you insight into where users are concentrated and the level of traffic a network design must support.

At this point in the network design process, your goal is to obtain a map of the already-implemented network. Some design customers may have maps for the new network design as well. If that is the case, then you may be one step ahead, but be careful of any assumptions that are not based on your detailed analysis of business and technical requirements.

Tools for Developing Network Maps

Not all customers can provide a detailed and up-to-date map of the existing network. In many cases, you need to develop the map yourself. Companies that are constantly working in "fire-fighting" mode do not have time to proactively document the existing network.

To develop a network drawing, you should invest in a good network-diagramming tool. Visio Corporation's Visio Professional is one of the premiere tools for diagramming networks. Visio Professional ships with templates for typical LANs and WANs, icons for common network and telecommunications devices, and the ability to draw WANs on top of a geographical map and LANs on top of a building or floor plan.

To create more detailed network diagrams, you can use the Visio Network Equipment product, an add-on library of 10,000 manufacturer-specific shapes with port-level detail. If a customer has equipment documented in a spreadsheet or database, you can

use the Visio Network Diagram Wizard to draw a diagram based on the network-equipment spreadsheet or database.

Some companies offer diagramming and network documentation tools that automatically discover the existing network. Pinpoint Software's ClickNet Professional is one such tool. ClickNet Professional uses various network-management protocols and other mechanisms to automatically learn and document the infrastructure of a customer's network. The tool automatically learns about internetworking devices and workstations, including CPU type, software versions, amount of memory, and the number of ports and network-interface cards. The tool includes the ability to customize a network map with backgrounds, floor plans, icons, and text. It also supports "what-if" analysis to determine the impact of projected network design changes.

NetSuite Development is another company that specializes in network-discovery and design tools. NetSuite Professional Audit is similar to ClickNet in its support for automatic discovery. The information gathered in a discovery session can be linked to the NetSuite Advanced Professional Design application to populate a design schematic. NetSuite Advanced Professional Design helps you design complex multi-layer networks. The application provides access to a library of network devices and includes a validation engine to let you test some aspects of a network design.

What Should a Network Map Include?

Regardless of the tools you use to develop a network map, your goal should be to develop (or obtain from your customer) a map (or set up maps) that includes the following:

- Geographical information, such as countries, states or provinces, cities, and campuses
- WAN connections between countries, states, and cities
- Buildings and floors, and possibly rooms or cubicles
- WAN and LAN connections between buildings and between campuses
- An indication of the data-link layer technology for WANs and LANs (Frame Relay, ISDN, 10-Mbps or 100-Mbps Ethernet, Token Ring, and so on)
- The name of the service provider for WANs
- The location of routers and switches, though not necessarily hubs

- The location and reach of any Virtual Private Networks (VPNs) that connect corporate sites via a service provider's WAN
- The location of major servers or server farms
- The location of mainframes
- The location of major network-management stations
- The location and reach of any virtual LANs (VLANs). (If the drawing is in color, you can draw all devices and segments within a particular VLAN in a specific color.)
- The topology of any firewall security systems
- The location of any dial-in and dial-out systems
- Some indication of where workstations reside, though not necessarily the explicit location of each workstation
- A depiction of the logical topology or architecture of the network

While documenting the network infrastructure, take a step back from the diagrams you develop and try to characterize the logical topology of the network as well as the physical components. The logical topology illustrates the architecture of the network, which can be hierarchical or flat, structured or unstructured, layered or not, and other possibilities. The logical topology also describes methods for connecting devices in a geometric shape, for example, a star, ring, bus, hub and spoke, or mesh.

The logical topology can affect your ability to upgrade a network. For example, a flat topology does not scale as well as a hierarchical topology. A typical hierarchical topology that does scale is a core layer of high-end routers and switches that are optimized for availability and performance, a distribution layer of routers and switches that implement policies, and an access layer that connects users via hubs, switches, and other devices. Logical topologies are discussed in more detail in Chapter 5, "Designing a Network Topology."

Figure 3–1 shows a typical high-level network diagram for an electronics manufacturing company. The drawing shows a physical topology, but it is not hard to step back and visualize that the logical topology is a hub-and-spoke shape with three layers. The core layer of the network is a 16-Mbps Token Ring network. The distribution layer includes routers and bridges, and Frame Relay and T1 links. The access layer comprises 4-Mbps and 16-Mbps Token Ring networks. An Ethernet network hosts the company's World Wide Web server. As you can see from the figure, the network included some rather old equipment. The company required design consultation to select new technologies to eliminate performance problems caused by Token-Ring bridges dropping frames.

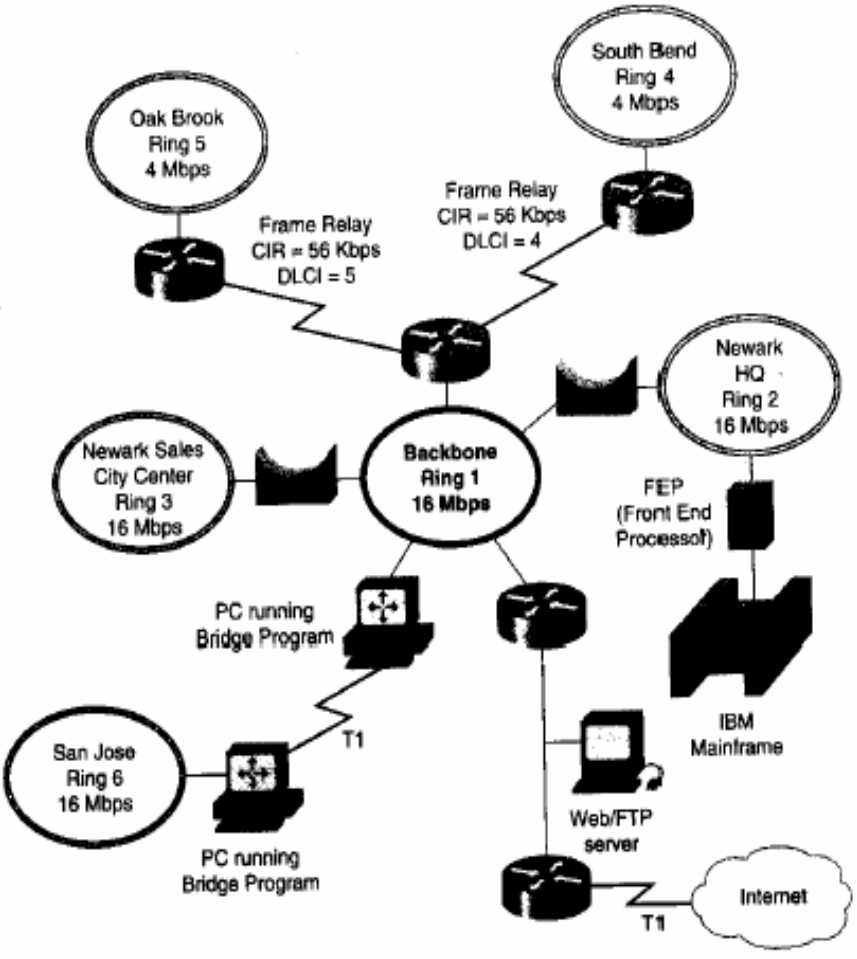

Figure 3–1 *Network diagram for an electronics manufacturing company.*

Characterizing Network Addressing and Naming

Characterizing the logical infrastructure of a network involves documenting any strategies your customer has for network addressing and naming. Addressing and naming are discussed in greater detail in Part II of this book, "Logical Network Design."

When drawing detailed network maps, include the names of major sites, routers, network segments, and servers. Also document any standard strategies your customer uses for naming network elements. For example, some customers name sites using airport codes. (San Francisco = SFO, Oakland = OAK, and so on.) You may find that a customer suffixes names with an alias that describes the type of device, for example, *rtr* for router. Some customers use a standard naming system, such as the Domain Name System (DNS), for IP networks.

You should also investigate the network-layer addresses your customer uses. Your customer's addressing scheme (or lack of any scheme) can influence your ability to adapt the network to new design goals. For example, your customer might use illegal IP addresses that will need to be changed or translated before connecting to the Internet. As another example, current IP subnet masking might limit the number of nodes in a LAN or VLAN.

Your customer might have a goal of using route summarization, which is also called *route aggregation* or *supernetting*. *Route summarization* reduces routes in a routing table, routing-table update traffic, and overall router overhead. Route summarization also improves network stability and availability, because problems in one part of a network are less likely to affect the whole internetwork. Summarization is most effective when address prefixes have been assigned in a consistent and contiguous manner, which is often not the case.

Your customer's existing addressing scheme might affect the routing protocols you can select. Some routing protocols do not support classless addressing, variable-length subnet masking (VLSM), or discontiguous subnets. A *discontiguous subnet* is a subnet that is divided, as shown in Figure 3–2. Subnet 108 of network 10 is divided into two areas that are separated by network 192.168.49.0.

Characterizing Wiring and Media

To help you meet scalability and availability goals for your new network design, it is important to understand the cabling design and wiring of the existing network. Doc-

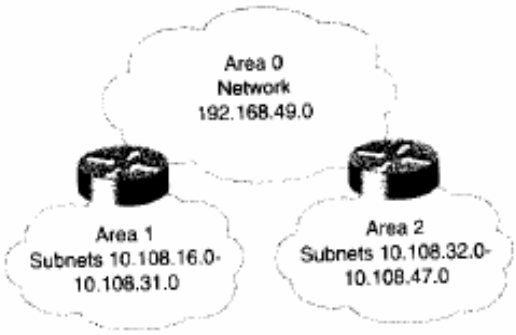

Figure 3-2 *An example of a discontiguous subnet.*

umenting the existing cabling design can help you plan for enhancements and identify any potential problems. If possible, you should document the types of cabling in use as well as cable distances. Distance information is useful when selecting data-link-layer technologies based on distance restrictions.

While exploring the cabling design, assess how well equipment and cables are labeled in the current network. The extent and accuracy of labeling will affect your ability to implement and test enhancements to the network.

Your network diagram should document the connections between buildings. The diagram should include information on the number of pairs of wires and the type of wiring (or wireless technology) in use. The diagram should also indicate how far buildings are from one another. Distance information can help you select new cabling. For example, if you plan to upgrade from copper to fiber cabling, the distance between buildings can be much longer. (Selecting cabling is discussed in more detail in Chapters 9 and 10.)

Probably the wiring (or wireless technology) between buildings is one of the following:

- Single-mode fiber
- Multi-mode fiber
- Shielded twisted pair (STP) copper
- Category-5 unshielded-twisted-pair (UTP) copper
- Coaxial cable

- Microwave
- Laser
- Radio
- Infra-red

Within buildings, try to locate telecommunications wiring closets, cross-connect rooms, and any laboratories or computer rooms. If possible, determine the type of cabling that is installed between telecommunications closets and in work areas. (Some technologies, for example, 100Base-TX Ethernet, require Category-5 cabling.) Gather information about both vertical and horizontal wiring. As shown in Figure 3–3, *vertical wiring* runs between floors. *Horizontal wiring* runs from telecommunications closets to wallplates in cubicles or offices. *Work-area wiring* runs from the wallplate to a workstation in a cubicle or office.

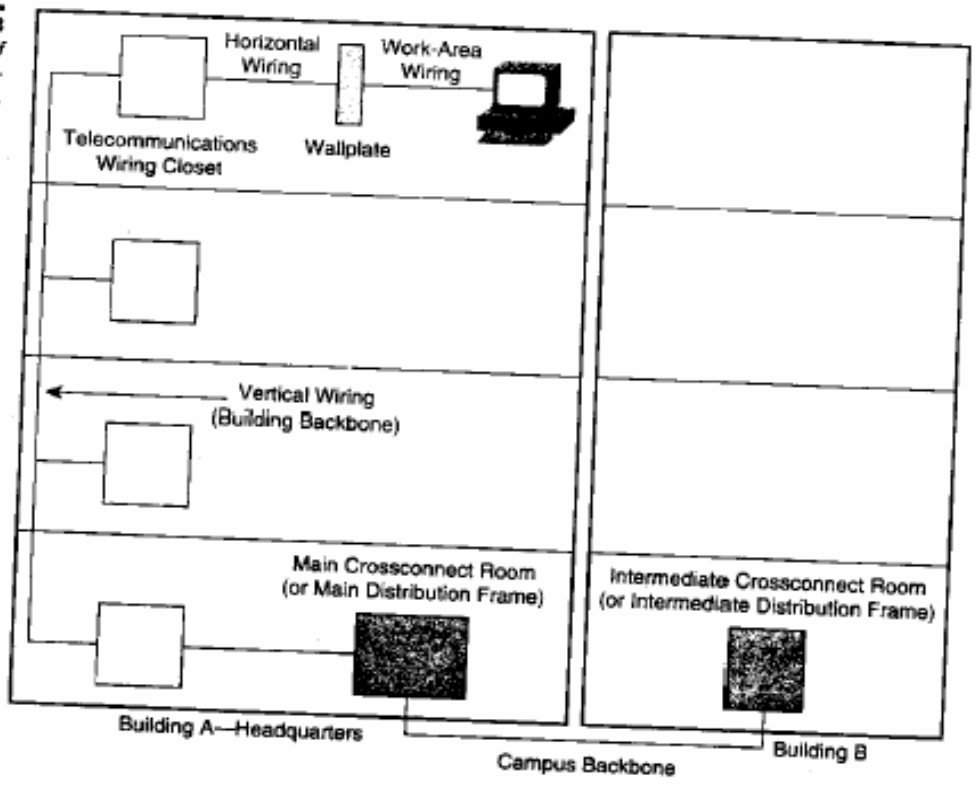

Figure 3–3 *An example of campus network wiring.*

In most buildings, the cabling from a telecommunications closet to a workstation is approximately 100 meters, including the work-area wiring which is usually just a few meters. If you have any indication that the cabling might be longer than 100 meters, you should use a *time-domain reflectometer (TDR)* to verify your suspicions. (TDR functionality is included in most cable testers.) Many network designs are based on the assumption that workstations are no more than 100 meters from the telecommunications closet.

For each building, you can fill out the chart shown in Table 3–1. The data that you fill in depends on how much time you have to gather information and how important you think cabling details will be to your network design. If you do not have a lot of information, than just put an **X** for each type of cabling present and document any assumptions (for example an assumption that workstations are no more than 100 meters from the telecommunications closet). If you have time to gather more details, then include information on the length and number of pairs of cables. If you prefer, you can document building wiring information in a network diagram instead of in a table.

Table 3-1 Building Wiring

Building Name:						
Location of telecommunications closets:						
Location of cross-connect rooms and demarcations to external networks:						
Logical wiring topology (structured, star, bus, ring, centralized, distributed, mesh, tree, or whatever fits):						
Vertical Wiring:						
	Coaxial	**Fiber**	**STP**	**Category 3 UTP**	**Category 5 UTP**	**Other**
Vertical Shaft 1						
Vertical Shaft 2						
Vertical Shaft *n*						

Table 3-1 Building Wiring, Continued

Horizontal Wiring:						
	Coaxial	**Fiber**	**STP**	**Category 3 UTP**	**Category 5 UTP**	**Other**
Floor 1						
Floor 2						
Floor 3						
Floor *n*						
Work-Area Wiring:						
	Coaxial	**Fiber**	**STP**	**Category 3 UTP**	**Category 5 UTP**	**Other**
Floor 1						
Floor 2						
Floor 3						
Floor *n*						

Checking Architectural and Environmental Constraints

When investigating cabling, pay attention to such environmental issues as the possibility that cabling will run near creeks that could flood, railroad tracks or highways where traffic could jostle cables, or construction or manufacturing areas where heavy equipment or digging could break cables.

Be sure to determine if there are any legal right-of-way issues that must be dealt with before cabling can be put into place. For example, will cabling need to cross a public street? Will it be necessary to run cables through property owned by other companies? Finally, for some wireless technologies, such as laser or infra-red, make sure there aren't any obstacles blocking the line of sight.

Within buildings, pay attention to architectural issues that could affect the feasibility of implementing your network design. Make sure the following architectural elements are sufficient to support your design:

- Air conditioning
- Heating
- Ventilation

- Power
- Protection from electromagnetic interference
- Clear paths for wireless transmission and an absence of confusing reflecting surfaces
- Doors that can lock
- Space for:
  - —Cabling (conduits)
  - —Patch panels
  - —Equipment racks
  - —Work areas for technicians installing and troubleshooting equipment

CHECKING THE HEALTH OF THE EXISTING INTERNETWORK

Studying the performance of the existing internetwork gives you a baseline measurement from which to measure new network performance. Armed with measurements of the present internetwork, you can demonstrate to your customer how much better the new internetwork performs once your design is implemented.

Many of the network-performance goals discussed in Chapter 2, "Analyzing Technical Goals and Constraints," are overall goals for an internetwork. Since the performance of existing network segments will affect overall performance, it is important that you study the performance of existing segments to determine how to meet overall network performance goals.

If an internetwork is too large to study all segments, then you should analyze the segments that will interoperate the most with the new network design. Pay particular attention to backbone networks and networks that connect old and new areas.

In some cases, a customer's goals might be at odds with improving network performance. The customer might want to reduce costs, for example, and not worry about performance. In this case, you will be glad that you documented the original perfor-

mance so that you can prove that the network was not optimized to start with and your new design has not made performance worse.

By analyzing existing networks, you can also recognize legacy systems that must be incorporated into the new design. Sometimes customers are not aware that older protocols are still running on their internetworks. By capturing network traffic with a protocol analyzer as part of your baseline analysis, you can identify which protocols are really running on the network and not rely on customers' beliefs.

The Challenges of Developing a Baseline of Network Performance

Developing an accurate baseline of a network's performance is not an easy task. One challenging aspect is selecting a time to do the analysis. It is important that you allocate a lot of time (multiple days) if you want the baseline to be accurate. If measurements are made over too short a timeframe, temporary errors appear more significant than they are.

In addition to allocating sufficient time for a baseline analysis, it is also important to find a typical time period to do the analysis. A baseline of normal performance should not include non-typical problems caused by exceptionally large traffic loads. For example, at some companies, end-of-the quarter sales processing puts an abnormal load on the network. In a retail environment, network traffic can increase five times around Christmas time. Network traffic to a Web server can unexpectedly increase as much as 10 times if the Web site gets linked to other popular sites or listed in search engines.

In general, errors, packet/cell loss, and latency increase with load. To get a meaningful measurement of typical accuracy and delay, try to do your baseline analysis during periods of normal traffic load. (On the other hand, if your customer's main goal is to improve performance during peak load, then be sure to study performance during peak load. The decision whether to measure normal performance, performance during peak load, or both, depends on the goals of the network design.)

Some customers do not recognize the value of studying the existing network before designing and implementing enhancements. Your customer's expectations for a speedy design proposal might make it difficult for you to take a step back and insist on time to develop a baseline of performance on the existing network. Also, your other job tasks and goals, especially if you are a sales engineer, might make it impractical to spend days developing a precise baseline.

The work you do before the baseline step in the top-down network design methodology can increase your efficiency in developing a baseline. A good understanding of your customer's technical and business goals can help you decide how thorough to make your study. Your discussions with your customer on business goals can help you identify segments that are important to study because they carry critical and/or backbone traffic. You can also ask your customer to help you identify typical segments from which you can extrapolate conclusions about other segments.

Analyzing Network Availability

To document availability characteristics of the existing network, gather any statistics that the customer has on the mean time between failure (MTBF) and mean time to repair (MTTR) for the internetwork as a whole as well as major network segments. Compare these statistics with information you have gathered on MTBF and MTTR goals, as discussed in Chapter 2, "Analyzing Technical Goals and Constraints." Does the customer expect your new design to increase MTBF and decrease MTTR? Are the customer's goals realistic considering the current state of the network?

Talk to the network engineers and technicians about the causes of the most recent and most disruptive periods of downtime. Acting like a forensic investigator, try to get many sides to the story. Sometimes myths develop about what caused a network outage. (You can usually get a more accurate view of problem causes from engineers and technicians than from users and managers.)

You can use Table 3–2 to document availability characteristics of the current network.

Table 3-2 Availability Characteristics of the Current Network

	MTBF	MTTR	Date and Duration of Last Major Downtime	Cause of Last Major Downtime
Enterprise (as a whole)				
Segment 1				
Segment 2				
Segment 3				
Segment *n*				

Analyzing Network Utilization

Network utilization is a measurement of how much bandwidth is in use during a specific time interval. Utilization is commonly specified as a percentage of capacity. If a network monitoring tool says that network utilization on an FDDI segment is 70 percent, for example, this means that 70 percent of the 100-Mbps capacity is in use, averaged over a specified timeframe or window.

Different tools use different averaging windows for computing network utilization. Some tools let the user change the window. Using a long interval can be useful for reducing the amount of statistical data that must be analyzed, but granularity is sacrificed. As Figure 3–4 shows, it can be informative (though tedious), to look at a chart that shows network utilization averaged every minute.

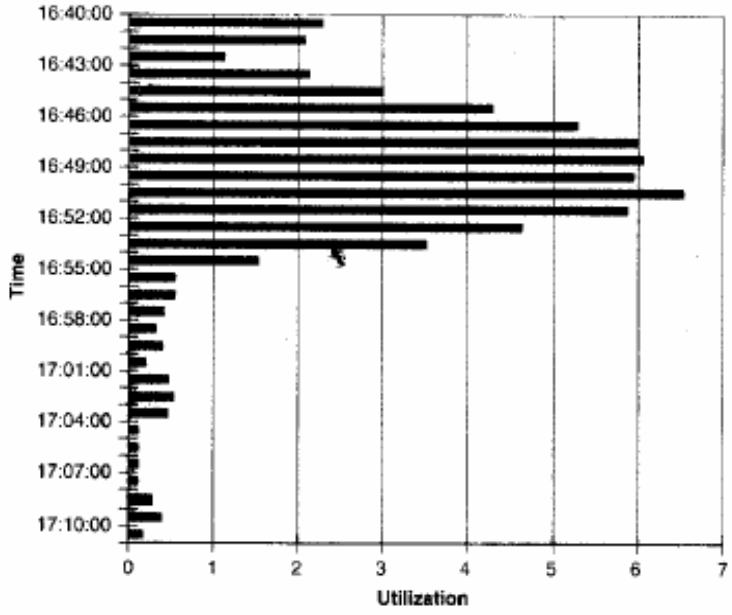

Figure 3–4 *Network utilization in minute intervals.*

Figure 3–5 shows the same data averaged over hour intervals. Note that the network was not very busy so neither chart goes above seven percent utilization. Note also that changing to a long interval can be misleading because peaks in traffic get averaged out

(the detail is lost). In Figure 3–4, you can see that the network was relatively busy around 4:50 PM. You cannot see this in Figure 3–5, when the data was averaged every hour.

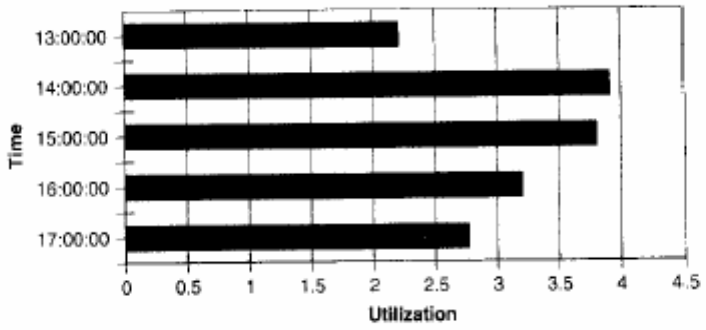

Figure 3–5
Network utilization in hour intervals.

As it turns out, 4:50 PM was the time of day that many of the network users turned their machines off and went home for the day. Each station leaving the Token-Ring network caused an error frame to be sent to the ring error monitor and a configuration frame to be sent to the configuration server, causing the peak in network utilization. So the peak in network traffic wasn't really interesting in this case, unless you were spying on the workers to make sure they didn't go home before 5 PM! Usually peaks in network utilization are something you want to know about when conducting a baseline analysis. In general, you should record network utilization with sufficient granularity in time to see short-term peaks in network traffic so that you can accurately assess the capacity requirements of devices and segments.

Changing the interval to a very small amount of time, say a fraction of a second, can be misleading also, however. Be wary of switch vendors who measure utilization on your network in millisecond increments and recommend that you update from a shared medium to a switched technology because short-term network utilization is dangerously high. To understand the concern, consider a very small time interval. In a packet-sized window, at a time when a station is sending traffic, the utilization is 100 percent, which is what is desired.

The size of the averaging window for network utilization measurements depends on your goals. When troubleshooting network problems, keep the interval very small, either minutes or seconds. A small interval helps you recognize peaks caused by problems such as broadcast storms or stations retransmitting very quickly due to a misconfigured timer. For performance analysis and baselining purposes, use an interval of 1 to 5 minutes. For long-term load analysis, to determine peak hours, days, or months, set the interval to 10 minutes.

When developing a baseline, it is usually a good idea to err on the side of gathering too much data. You can always summarize the data later. When characterizing network utilization, use protocol analyzers or other monitoring tools to measure utilization in 1 to 5 minute intervals on each major network segment. If practical, leave the monitoring tools running for at least one or two typical days. If the customer's goals include improving performance during peak times, measure utilization during peak times and typical times. To determine if the measured utilization is healthy, use the Network Health Checklist that appears at the end of this chapter.

Bandwidth Utilization by Protocol

Developing a baseline of network performance should also include measuring utilization from broadcast traffic versus unicast traffic, and by each major protocol. As discussed in Chapter 4, "Characterizing Network Traffic," some protocols send excessive broadcast traffic, which can seriously degrade performance, especially on switched networks.

To measure bandwidth utilization by protocol, place a protocol analyzer on each major network segment and fill out a chart such as the one shown in Table 3–3. If the analyzer supports relative and absolute percentages, specify the bandwidth used by protocols as relative and absolute. *Relative usage* specifies how much bandwidth is used by the protocol in comparison to the total bandwidth currently in use on the segment. *Absolute usage* specifies how much bandwidth is used by the protocol in comparison to the total capacity of the segment (for example, in comparison to 10 Mbps on Ethernet).

Table 3-3 Bandwidth Utilization by Protocol

	Relative Network Utilization	Absolute Network Utilization	Broadcast/Multicast Rate
IP			
IPX			
AppleTalk			
DECnet			
Banyan			
NetBIOS			
SNA			
Other			

Analyzing Network Accuracy

The previous chapter talked about specifying network accuracy as a bit error rate (BER). You can use a BER tester (also called a *BERT*) on serial lines to test the number of damaged bits compared to total bits.

With packet-switched networks, it makes more sense to measure frame (packet) errors because a whole frame is considered bad if a single bit is changed or dropped. In packet-switched networks, a sending station calculates a cyclic redundancy check (CRC) based on the bits in a frame. The sending station places the value of the CRC in the frame. A receiving station determines if a bit has been changed or dropped by calculating the CRC again and comparing the result to the CRC in the frame. A frame with a bad CRC is dropped and must be retransmitted by the sender. Usually an upper-layer protocol has the job of retransmitting frames that do not get acknowledged.

A protocol analyzer can check the CRC on received frames. As part of your baseline analysis, you should track the number of frames received with a bad CRC every hour for one or two days. Because it is normal for errors to increase with utilization, document errors as a function of the number of bytes seen by the monitoring tool. A good rule-of-thumb threshold for considering errors unhealthy is that a network should not have more than one bad frame per megabyte of data. (Calculating errors this way lets you simulate a serial BERT. Simply calculating a percentage of bad frames compared to good frames does not account for the size of frames and hence does not give a good indication of how many bits are actually getting damaged.)

Some network monitors let you print a report of the top 10 stations sending frames with CRC errors. Token-Ring monitors let you print a report of the top 10 stations sending error reports to the ring error monitor. You should correlate the information on stations sending the most errors with information you gathered on network topology to identify any areas of a network that are prone to errors, possibly due to electrical noise or cabling problems.

Network problems are usually not caused by the stations sending bad frames or error reports. The stations reporting problems are usually the "victims" not the "perpetrators." In the case of Token Ring, the problem is usually caused by a station or cabling problem upstream from the station reporting the problem. In the case of Ethernet, it is more difficult to pinpoint the cause of problems. With a thorough investigation, however, you usually can isolate a problematic area of the network where frames are damaged by a bad repeater, electrical problem, cabling fault, or misbehaving network interface card.

In addition to tracking data-link layer errors, such as CRC errors, a baseline analysis should include information on upper-layer problems. A protocol analyzer that includes an expert system, such as Network Associate's Sniffer network analyzer, speeds the identification of upper-layer problems by automatically generating diagnoses and symptoms for network conversations and applications.

Accuracy should also include a measurement of lost packets. You can measure lost packets while measuring response time, which is covered later in this chapter in the "Analyzing Delay and Response Time" section. When sending packets to measure how long it takes to receive a response, document any packets that do not receive a response, presumably because either the request or the response got lost. Correlate the information about lost packets with other performance measurements to determine if the lost packets indicate a need to increase bandwidth, decrease CRC errors, or upgrade internetworking devices. You can also measure lost packets by looking at statistics kept by routers on the number of packets dropped from input or output queues.

Analyzing ATM Errors

The ATM Forum specifies ATM accuracy in terms of a cell error ratio (CER), cell loss ratio (CLR), cell misinsertion rate (CMR), and severely errored cell block ratio (SECBR).

The CER is the number of errored cells divided by the total number of successfully transferred cells plus errored cells. The CLR is the number of lost cells divided by the total number of transmitted cells.

CMR on a connection is caused by an undetected error in the header of a cell being transmitted on a different connection. SECBR occurs when more than a certain number of errored cells, lost cells, or misinserted cells are observed in a received cell block. A *cell block* is a sequence of cells transmitted consecutively on a given connection.

If you do not have tools that can measure cell errors, you can still check the performance of an ATM network by analyzing the level of frame errors and upper-layer problems. With ATM, if a cell is lost or damaged, all cells comprising a complete LAN/WAN frame must be retransmitted. A good protocol analyzer can measure frame errors and upper-layer problems to help you characterize performance on an internetwork that includes ATM segments.

Analyzing Network Efficiency

The previous chapter talked about the importance of using maximum frame sizes to increase network efficiency. Bandwidth utilization is optimized for efficiency when applications and protocols are configured to send large amounts of data per frame, thus minimizing the number of frames and round-trip delays required for a transaction. The number of frames per transaction can also be minimized if the receiver is configured with a large receive window allowing it to accept multiple frames before it must send an acknowledgment. The goal is to maximize the number of data bytes compared to the number of bytes in headers and in acknowledgement packets sent by the other end of a conversation. Changing transmit and receive packet-buffer sizes at clients and servers can result in optimized frame sizes and receive windows.

To determine if your customer's goals for network efficiency are realistic, you should use a protocol analyzer to examine the current frame sizes on the network. Many protocol analyzers let you output a chart such as the one in Figure 3–6 that documents how many frames fall into standard categories for frame sizes.

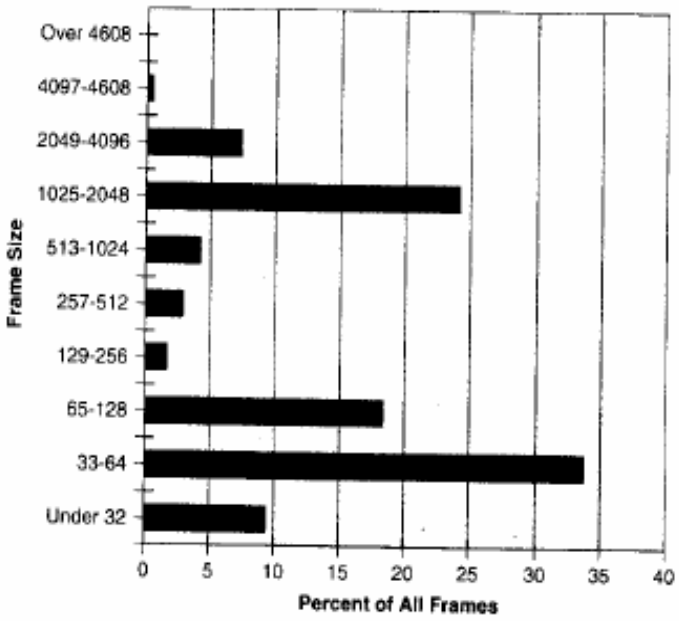

Figure 3–6 *Bar graph of frame sizes on a Token Ring network.*

A simple way to determine an *average* frame size is to divide the total number of megabytes seen on a segment by the total number of frames in a specified timeframe. Unfortunately, this is a case where a simple statistical technique does not result in useful data. The average frame size is not a very meaningful piece of information. On most networks, there are many small frames, many large frames, but very few average-sized frames. Small frames consist of acknowledgments and control information. Data frames fall into the large frame-size categories (if the network has been optimized). A line graph of frame sizes, such as the graph in Figure 3–7, helps demonstrate this point.

Frame sizes typically fall into what is called a *bimodal distribution*, also known as a *camel-back distribution*. There is a "hump" on either side of the average but not many values near the average.

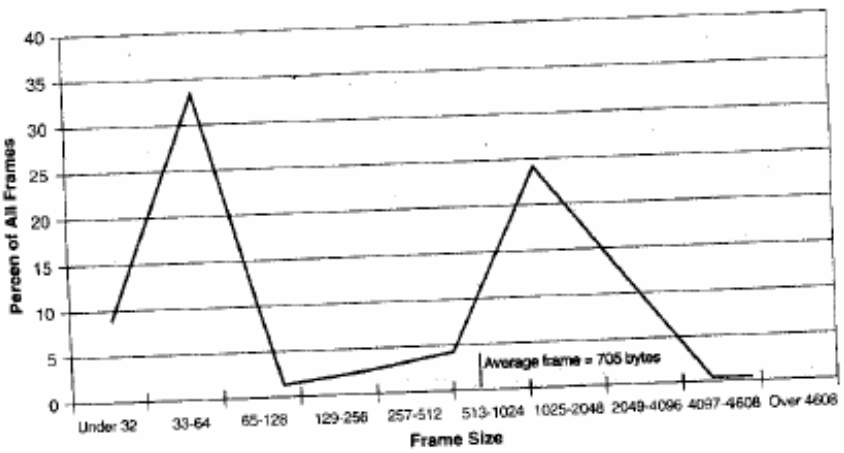

Figure 3–7 *Line graph of frame sizes on a Token Ring network.*

Network performance data is /often bimodal, multi-modal, or skewed from the mean. (*Mean* is another word for average.) Frame size is usually bimodal. Response time from a server can also be bimodal, if sometimes the data is quickly available from random-access memory (RAM) cache and sometimes the data is retrieved from a slow mechanical disk drive.

When network-performance data is bimodal, multi-modal, or skewed from the mean, you should document a standard deviation with any measurements of the mean. *Standard deviation* is a measurement of how widely data disperses from the mean. If you do not have time to calculate standard deviation, a graph of the data can illustrate the deviation, as shown in Figure 3–7. Figure 3–7 shows that very few data points fall at the mean and many data points fall in two "humps" away from the mean.

Analyzing frame sizes can help you understand the health of a network, not just the efficiency. For example, an excessive number of Ethernet runt frames (less than 64 bytes) can indicate too many collisions. It is normal for collisions to increase with utilization that results from access contention. If collisions increase even when utilization

does not increase or even when only a few nodes are transmitting, there could be a component problem, such as a bad repeater or network interface card.

On Token-Ring networks, frames that are less than 32 bytes are probably media-access control (MAC) frames. Network stations use MAC frames to identify themselves during the ring-poll process and to report changes and errors. Too many MAC frames can indicate a problem. (Knowing how many is "too many" requires a previous baseline.) If the number of MAC frames seems suspicious, view the contents of the frames on a protocol analyzer. Check for excessive error reports and beacon frames. *Beacon frames* indicate a serious problem on both Token Ring and FDDI networks.

Analyzing Delay and Response Time

To verify that performance of a new network design meets a customer's requirements, it is important to measure response time between significant network devices before and after a new network design is implemented. Response time can be measured many ways. Using a protocol analyzer, you can look at the amount of time between frames and get a rough estimate of response time at the data-link layer, transport layer, and application layer. (This is a rough estimate because packet arrival times on an analyzer can only approximate packet arrival times on end stations.)

A more common way to measure response time is to send ping packets and measure the round trip time (RTT) to send a request and receive a response. While measuring RTT, you can also measure an RTT variance. *Variance measurements* are important for applications that cannot tolerate much jitter, for example, voice and video applications. You can also document any loss of packets.

In an IP environment, a *ping packet* is an Internet Control Message Protocol (ICMP) echo packet. To measure response time on AppleTalk networks, use the AppleTalk Echo Protocol (AEP). For Novell NetWare networks, you can use the Internetwork Packet Exchange (IPX) ping packet. When testing with an IPX ping, be careful to use the right ping version. There is a Cisco Systems, Inc., proprietary IPX ping to which only Cisco routers respond, and a different IPX ping packet specified by Novell. Novell servers and Cisco routers respond to the Novell IPX ping (as long as the Cisco routers are running a recent version of the Cisco Internetwork Operating System [IOS] software).

You can use Table 3–4 to document response time measurements. The table uses the term *node* to mean *router, server, client*, or *mainframe*.

Table 3-4 Response-Time Measurements

	Node A	Node B	Node C	Node D
Node A	X			
Node B		X		
Node C			X	
Node D				X

Depending on the amount of time you have for your analysis and depending on your customer's network design goals, you should also measure response time from a user's point of view. On a typical workstation, run some representative applications and measure how long it takes to get a response for typical operations, such as checking e-mail, sending a file to a server, downloading a Web page, updating a sales order, printing a report, and so on. Measure how much time a workstation takes to boot.

Sometimes applications or protocol implementations are notoriously slow or poorly written. Some peripherals are known to cause extra delay because of incompatibilities with operating systems or hardware. By joining mailing lists and newsgroups and reading information in journals and on the World Wide Web, you can learn about causes of response-time problems. Be sure to do some testing on your own also, though, since every environment is different.

In addition to testing user applications, test the response time for system protocols, for example Domain Name System (DNS) queries, Dynamic Host Configuration Protocol (DHCP) requests for an IP address, requests for a list of zones on an AppleTalk network, and so on. Chapter 4, "Characterizing Network Traffic," covers protocol issues in more detail.

Although your customer might not give you permission to simulate network problems, it also makes sense to do some testing of response times when the network is experiencing problems or change. For example, if possible, measure response times while routing protocols are converging after a link has gone down. Measure response time during convergence again, after your new design is implemented, to see if the results have improved. As covered in Chapter 11, "Testing Your Network Design," you can test network problems on a pilot implementation.

Checking the Status of Major Routers on the Internetwork

The final step in characterizing the existing internetwork is to check the behavior of the major routers on the internetwork. This includes routers that connect layers of a hierarchical topology, backbone routers, and routers that will have the most significant roles in your new network design.

Checking the behavior and health of a router includes determining how busy the router is (CPU utilization), how many packets the router has processed, how many packets the router has dropped, and the status of buffers and queues. Your method for assessing the health of a router depends on the router vendor and architecture. In the case of Cisco routers, you can use the following Cisco IOS commands:

- **`show interfaces`.** Displays statistics for network interface cards, including the input and output rate of packets, a count of packets dropped from input and output queues, the size and usage of queues, a count of packets ignored due to lack of I/O buffer space on a card, and how often interfaces have restarted.

- **`show processes`.** Displays CPU utilization for the last five seconds, one minute, and five minutes, and the percentage of CPU used by various processes, including routing protocols, buffer management, and user-interface processes.

- **`show buffers`.** Displays information on buffer sizes, buffer creation and deletion, buffer usage, and a count of successful and unsuccessful attempts to get buffers when needed.

You can also use the Simple Network Management Protocol (SNMP) to check the health of a router. Following is a list of useful router performance variables in Cisco's private extension to the Internet standard Management Information Base II (MIB II):

- **`BusyPer`.** CPU busy percentage in the last five-second period.

- **`AvgBusy1`.** One-minute exponentially-decayed moving average of the CPU busy percentage.

- **`AvgBusy5`.** Five-minute exponentially-decayed moving average of the CPU busy percentage.

- **`LocIfInputQueueDrops`.** The number of packets dropped because the input queue was full.

- **LocIfOutputQueueDrops.** The number of packets dropped because the output queue was full.
- **LocIfInIgnored.** The number of input packets ignored by the interface.
- **BufferElMiss.** The number of buffer-element misses. (You can also check misses for small, medium, big, large, and huge buffer pools.)
- **BufferFail.** The number of buffer allocation failures.

To analyze router health, you need to check the variables listed above on a regular basis over a few days. To get a precise and complete portrayal of router performance, a long-term study (lasting a few weeks or months) should be done, using some of the tools mentioned in the next section and in Chapter 8, "Developing Network Security and Network Management Strategies." (As a network designer it is probably not your job to do the long-term study. You should encourage your customer to assign network engineers or consultants to the job of proactively studying long-term router performance.)

TOOLS FOR CHARACTERIZING THE EXISTING INTERNETWORK

This chapter has already mentioned some tools for characterizing an existing network, including network-discovery tools, protocol analyzers, SNMP tools, and Cisco IOS commands. To help you select tools, this section provides more information on tools.

Protocol Analyzers

A *protocol analyzer* is a fault-and-performance-management tool that captures network traffic, decodes the protocols in the captured packets, and provides statistics to characterize load, errors, and response time. Some analyzers include an expert system that automatically identifies network problems.

One of the best known protocol analyzers is the Sniffer Network Analyzer from Network Associates, Inc. (Network Associates purchased Network General, the original manufacturer of the Sniffer Network Analyzer, in 1997). The Sniffer network analyzer decodes hundreds of protocols and applies expert analysis to diagnose problems and recommend corrective action. Because the Sniffer network analyzer has been on

the market longer than most other analyzers, it has the most sophisticated protocol decoding and expert system.

Another noteworthy protocol analyzer is EtherPeek from the AG Group. The AG Group has versions of EtherPeek for the Macintosh operating system, Windows 95, and Windows NT. Because the AG Group developed EtherPeek for the Macintosh first, it is very easy to use and install. EtherPeek decodes all major protocols and includes a nice feature for displaying in real time a tree structure of protocols within protocols. EtherPeek includes plug-in modules for expert analysis.

Remote Monitoring Tools

The Internet Engineering Task Force (IETF) developed the Remote Monitoring (RMON) MIB in the early 1990s to address shortcomings in the standard SNMP MIBs for gathering statistics on data-link and physical-layer parameters. The IETF developed the RMON MIB to enable network managers to collect traffic statistics, analyze Ethernet problems, plan network upgrades, and tune network performance. In 1994, Token-Ring statistics were added. Other types of statistics, for example, application-layer and WAN statistics, are under development.

RMON facilitates gathering statistics on the following data-link-layer performance factors:

- CRC errors
- Ethernet collisions
- Token-Ring soft errors
- Frame sizes
- The number of packets in and out of a device
- The rate of broadcast packets

The RMON MIB alarm group lets a network manager set thresholds for network parameters and automatically deliver alerts to management consoles. RMON also supports capturing packets and sending the captured packets to a network-management station for protocol decoding. RMON is discussed in more detail in Chapter 8, "Developing Network Security and Network Management Strategies."

Cisco Tools for Characterizing an Existing Internetwork

Cisco has a complete range of tools for characterizing an existing internetwork, ranging from the Cisco Discovery Protocol to sophisticated Netsys tools.

Cisco Discovery Protocol

The Cisco Discovery Protocol (CDP) specifies a method for Cisco routers and switches to send configuration information to each other on a regular basis. Analyzing CDP data can help you characterize the topology of an existing network (although you should use more sophisticated tools for large networks). If you enable CDP on a router and neighboring routers, you can use the `show cdp neighbors detail` command to display the following information about neighboring routers:

- Which protocols are enabled
- Network addresses for enabled protocols
- The number and types of interfaces
- The type of platform and its capabilities
- The version of Cisco IOS software

Enterprise Accounting for NetFlow

Cisco Enterprise Accounting for NetFlow can help you understand bandwidth usage and allocation, quality of service (QoS) levels, router usage, and router port usage. NetFlow accounting recognizes network flows and characterizes network and router usage by user (IP address), application, and department.

Netsys Service-Level Management Suite

The Cisco Netsys Service-Level Management Suite enables defining, monitoring, and assessing network connectivity, security, and performance. The Cisco Netsys Performance Service Manager is particularly useful for characterizing an existing network as part of a network design proposal.

CiscoWorks

CiscoWorks is a series of SNMP-based internetwork management software applications to allow device monitoring, configuration maintenance, and troubleshooting of Cisco devices. Health Monitor is a CiscoWorks application that lets you view information about the status of a device, including buffer usage, CPU load, available memory, and protocols and interfaces being used. Threshold Manager allows you to set RMON alarm thresholds and retrieve RMON event information. You can set thresholds for network devices using Cisco-provided default or customized policies.

CiscoWorks Blue Internetwork Performance Monitor (IPM) provides mechanisms to isolate performance problems, diagnose latency, perform trend analysis, and determine the possible paths between two devices and display the performance characteristics of each path. Performance measurement capability is supported for both IP and Systems Network Architecture (SNA) session paths.

Other Tools for Characterizing an Existing Internetwork

You can search the Web to learn more about the following tools that have become industry standards for monitoring network and router performance.

The Proactive Management family of products from Network Associates consists of the RouterPM tool for monitoring Cisco routers, RouterPM Blue for SNA networks, SwitchPM for Cisco switches, and FrameRelayPM for multi-vendor frame relay devices. The Proactive Management tools let you use a Web browser to view statistics and reports about router and network usage, and identify and resolve enterprise-wide performance and capacity-planning problems.

In April 1996, the Internet community awarded Merit Network, Inc., a grant to develop a freeware, turn-key network statistics package for managers of Internet sites. Merit Network, Inc., gained a reputation for expertise in network management because of its successful management of the NSFNET backbone of the Internet. The Network Statistics Collection and Reporting Facility (NetSCARF) team at Merit developed the *Scion software package*, which collects network-management information from network routers and makes the information available in HTML (Web) format.

The Multi Router Traffic Grapher (MRTG) is a tool for monitoring network traffic load and other performance characteristics of a routed network. MRTG generates HTML pages containing GIF images that provide a live (real-time) graphical representation of network traffic. MRTG is based on the Perl scripting language and C pro-

gramming language and runs on the UNIX and Windows NT operating systems. Many sites around the world use MRTG to proactively manage network traffic, monitor QoS commitments, and bill customers based on network usage. MRTG is available under a GNU public license. The only thing that author Tobias Oetiker asks users to do is send him a picture postcard. (The author lives in Zurich, Switzerland. A search on the Web can find his most recent address.)

The Cooperative Association for Internet Data Analysis (CAIDA) maintains a taxonomy of network measurement tools. The group's focus is research and network design tools, rather than operational network management, so the taxonomy is quite relevant to the discussions in this chapter. For more information on the CAIDA taxonomy, go to the Web site `www.caida.org/Tools/taxonomy.html`.

Also, several mailing lists discuss traffic measurement and analysis tools, including the Internet Statistics Measurement and Analysis (ISMA) mailing list and the IETF Internet Protocol Performance Metrics (IPPM) working group mailing list. To subscribe to the ISMA mailing list, send mail to `isma-request@nlanr.net`. For more information on the IPPM working group, see the group's Web site at: `io.advanced.org/IPPM/`.

NETWORK HEALTH CHECKLIST

You can use the following Network Health Checklist to assist you in verifying the health of an existing internetwork. The network health checklist is generic in nature and documents a best-case scenario. The thresholds might not apply to all networks.

- ☐ The network topology and physical infrastructure are well documented.
- ☐ Network addresses and names are assigned in a structured manner and are well documented.
- ☐ Network wiring is installed in a structured manner and is well labeled.
- ☐ Network wiring between telecommunications closets and end stations is generally no more than 100 meters.
- ☐ Network availability meets current customer goals.
- ☐ Network security meets current customer goals.

- ☐ No shared Ethernet segments are becoming saturated. (50 percent average network utilization in a 10-minute window.)
- ☐ No shared Token Ring segments are becoming saturated. (70 percent average network utilization in a 10-minute window.)
- ☐ No shared FDDI segments are becoming saturated. (70 percent average network utilization in a 10-minute window.)
- ☐ No WAN links are becoming saturated. (70 percent average network utilization in a 10-minute window.)
- ☐ No segments have more than one CRC error per million bytes of data.
- ☐ On Ethernet segments, less than 0.1 percent of packets are collisions. There are no late collisions.
- ☐ On Token Ring segments, less than 0.1 percent of packets are soft errors not related to ring insertion. There are no beacon frames.
- ☐ Broadcast traffic is less than 20 percent of all traffic on each network segment. (Some networks are more sensitive to broadcast traffic and should use a 10 percent threshold.)
- ☐ Wherever possible, frame sizes have been optimized to be as large as possible for the data-link layer in use.
- ☐ No routers are overutilized. (Five-minute CPU utilization is under 75 percent.)
- ☐ On an average, routers are not dropping more than 1 percent of packets. (For networks that are intentionally oversubscribed to keep costs low, a higher threshold can be used.)
- ☐ The response time between clients and hosts is generally less than 100 milliseconds (1/10 of a second).

SUMMARY

This chapter covered techniques and tools for characterizing a network before designing enhancements to the network. Characterizing an existing network is an important step in top-down network design because it helps you verify that a customer's technical-design goals are realistic. It also helps you understand the current topology and locate existing network segments and equipment, which will be useful information when the time comes to install new equipment. As part of the task of characterizing the existing network, you should develop a baseline of current performance. Baseline performance measurements can be compared to new measurements once your design is implemented to demonstrate to your customer that your new design (hopefully) improves performance.