

Analyzing Technical Goals and Constraints

This chapter provides techniques for analyzing a customer's technical goals for an enterprise network design. Analyzing your customer's technical goals can help you confidently recommend technologies that will perform to your customer's expectations.

Typical technical goals include scalability, availability, performance, security, manageability, usability, adaptability, and affordability. Of course, there are tradeoffs associated with these goals. For example, meeting strict requirements for performance can make it hard to meet a goal of affordability. The section, "Making Network Design Tradeoffs," later in this chapter discusses tradeoffs in more detail.

One of the objectives of this chapter is to give you terminology that will help you discuss technical goals with your customer. Network designers and users have many terms for technical goals, and, unfortunately, many different meanings for the terms. This chapter can help you use the same terms as your customer and mean the same things by the terms.

This chapter concludes with a checklist to help you determine whether you have addressed all of your customer's technical goals and constraints.

SCALABILITY

Scalability refers to how much growth a network design must support. For many enterprise network design customers, scalability is a primary goal. Large companies are adding users, applications, additional sites, and external network connections at a rapid rate. The network design you propose to a customer should be able to adapt to increases in network usage and scope.

Planning for Expansion

Your customer should be able to help you understand how much the network will expand in the next year and in the next two years. (Ask your customer to analyze goals for growth in the next five years also, but be aware that not many companies have a clear five-year vision.) You can use the following list of questions to analyze your customer's short-term goals for expansion:

- How many more sites will be added in the next year? The next two years?
- How extensive will the networks be at each new site?
- How many more users will access the corporate internetwork in the next year? The next two years?
- How many more servers (or hosts) will be added to the internetwork in the next year? The next two years?

Expanding the Data Available to Users

Chapter 1, "Analyzing Business Goals and Constraints," talked about a common business goal of expanding the data available to employees who use enterprise networks. Managers are empowering employees to make strategic decisions that require access to sales, marketing, engineering, and financial data. Traditionally this data was stored on departmental LANs. Today this data is often stored on centralized servers.

For years, networking books and training classes taught the 80/20 rule for capacity planning: 80 percent of traffic stays local in departmental LANs and 20 percent of traffic is destined for other departments or external networks. This rule is no longer universal and is rapidly moving to the other side of the scale. Many companies have centralized servers residing on server farms located on building or campus backbone

networks. In addition, corporations are increasingly implementing intranets that enable employees to access centralized World Wide Web servers using Internet Protocol (IP) technologies.

At some companies, employees can access intranet Web servers to arrange business travel, search online phone directories, order equipment, and attend distance-learning training classes. The Web servers are centrally located, which breaks the classic 80/20 rule.

In the 1990s, there has also been a trend of companies connecting internetworks with other companies to collaborate with partners, resellers, suppliers, and strategic customers. The term *extranet* is gaining popularity to describe an internal internetwork that is accessible by outside parties. If your customer has plans to implement an extranet, you should document this in your list of technical goals so you can design a topology and provision bandwidth appropriately.

In the 1980s, mainframes running Systems Network Architecture (SNA) protocols stored most of a company's financial and sales data. In the 1990s and beyond, the value of making this data available to more than just financial analysts has been recognized. The business goal of making data available to more departments often results in a technical goal of merging an SNA network with an enterprise IP network. Chapter 7, "Selecting Bridging, Switching, and Routing Protocols," provides more detail on how to migrate SNA data to an IP network.

In summary, the business goal of making more data available to users results in the following technical goals for scaling and upgrading corporate enterprise networks:

- Connect separated departmental LANs into the corporate internetwork
- Solve LAN/WAN bottleneck problems caused by large increases in internetwork traffic
- Provide centralized servers that reside on server farms or an intranet
- Merge an independent SNA network with the enterprise IP network
- Add new sites to support field offices and telecommuters
- Add new sites to support communication with customers, suppliers, resellers, and other business partners

Constraints on Scalability

When analyzing a customer's scalability goals, it is important to keep in mind there are impediments to scalability inherent in networking technologies. Selecting technologies that can meet a customer's scalability goals is a complex process with significant ramifications if not done correctly. For example, selecting a flat network topology with Layer 2 switches can cause problems as the number of users scales, especially if the users' applications or network protocols send numerous broadcast frames. (Switches forward broadcast frames to all connected segments.)

Subsequent chapters in this book consider scalability again. Chapter 4, "Characterizing Network Traffic," discusses the fact that network traffic—for example, broadcast traffic—affects the scalability of a network. Part 2, "Logical Network Design," provides details on the scalability of routing and bridging protocols. Part 3, "Physical Network Design," provides information on the scalability of LAN and WAN technologies and internetworking devices. Remember that top-down network design is an iterative process. Scalability goals and solutions are revisited during many phases of the network design process.


AVAILABILITY

Availability refers to the amount of time a network is available to users and is often a critical goal for network design customers. Availability can be expressed as a percentage uptime per year, month, week, day, or hour, compared to the total time in that period. For example, in a network that offers 24-hour, seven-days-a-week service, if the network is up 165 hours in the 168-hour week, availability is 98.21 percent.

Network design customers don't use the word *availability* in everyday English and have a tendency to think it means more than it does. In general, availability means how much time the network is operational. Availability is linked to redundancy, but redundancy is not a network goal. Redundancy is a solution to the goal of availability. Redundancy means adding duplicate links or devices to a network to avoid downtime.

Availability is also linked to reliability, but has a more specific meaning (percent uptime) than reliability. Reliability refers to a variety of issues, including accuracy, error rates, stability, and the amount of time between failures. Some network users use the term *recoverability* to specify how easily and in what timeframe a network can recover from problems. Recoverability is an ingredient of availability.

Availability is also associated with *resiliency*, which is a word that is becoming more popular in networking magazines. Resiliency means how much stress a network can handle and how quickly the network can rebound from problems. A network that has good resiliency usually has good availability.



Sometimes network engineers classify *capacity* as part of availability. The thinking is that even if a network is available at Layer 1 (the physical layer), it is not available from a user's point of view if there is not enough capacity to send the user's traffic.

For example, Asynchronous Transfer Mode (ATM) has a connection admission control (CAC) function that regulates the number of cells allowed into an ATM network. If the capacity and quality of service requested for a connection are not available, cells for the connection are not allowed to enter the network. This problem could be considered an availability issue. However, this book classifies capacity with performance goals. Availability is considered simply a goal for percent uptime.

One other aspect of availability is disaster recovery. Most large institutions have a plan for recovering from natural disasters, such as floods, fires, hurricanes, and earthquakes. Also, some large enterprises (especially service providers) must plan how to recover from satellite outages. Satellite outages can be caused by meteorite storms, collisions with space debris, solar flares, or system failures. Unfortunately, some institutions have also found the need to specify a recovery plan for man-made disasters, such as bombs or hostage situations. A disaster recovery plan includes a process for keeping data backed up in a place that is unlikely to be hit by disaster, as well as a process for switching to backup technologies if the main technologies are affected by a disaster. The details of disaster-recovery planning are outside the scope of this book.

Specifying Availability Requirements

You should encourage your customers to specify availability requirements with precision. Consider the difference between an uptime of 99.70 percent and an uptime of 99.95 percent. An uptime of 99.70 percent means the network is down 30 minutes per week, which is not acceptable to many customers. An uptime of 99.95 percent

means the network is down five minutes per week, which probably is acceptable. Availability requirements should be specified with at least two digits following the decimal point.

It is also important to specify a timeframe with percent uptime requirements. Go back to the example of 99.70 percent uptime, which equated to 30 minutes of downtime per week. A downtime of 30 minutes in the middle of a working day is probably not acceptable. But a downtime of 30 minutes every Saturday evening for regularly scheduled maintenance might be fine.

Not only should your customers specify a timeframe with percent uptime requirements, but they should also specify a time unit. Availability requirements should be specified as uptime per year, month, week, day, or hour. Consider an uptime of 99.70 percent again. This uptime means 30 minutes of downtime during a week. The downtime could be all at once, which would be a problem, or it could be spread out over the week. An uptime of 99.70 percent could mean that approximately every hour the network is down for 10.70 seconds. Will users notice a downtime of 10.70 seconds? Perhaps. For many applications, however, a downtime of 10.70 seconds every hour is tolerable.

Try doing the math yourself for a network goal of 99.80 percent uptime. How much downtime is permitted in hours per week? How much downtime is permitted in minutes per day and seconds per hour? Which values are acceptable?

The Cost of Downtime

In general, a customer's goal for availability is to keep mission-critical applications running smoothly, with little or no downtime. A method to help both you and your customer understand availability requirements is to specify a cost of downtime. For each critical application, document how much money the company loses per hour of downtime. (For some applications, such as order processing, specifying money lost per minute might have more impact.) If network operations will be outsourced to a third-party network management firm, explaining the cost of downtime can help the firm understand the criticality of applications to a business's mission.

Specifying the cost of downtime can also help clarify whether in-service upgrades must be supported. In-service upgrades refer to mechanisms for upgrading network equipment and services without disrupting operations. Most internetworking vendors sell high-end internetworking devices that include hot-swappable components for in-service upgrading.

Mean Time Between Failure and Mean Time to Repair

In addition to expressing availability as a percent uptime, you can define availability as a mean time between failure (MTBF) and mean time to repair (MTTR). You can use MTBF and MTTR to calculate availability goals when the customer wants to specify explicit periods of uptime and downtime, rather than a simple percent uptime value.

MTBF is a term that comes from the computer industry and is best suited to specifying how long a computer or computer component will last before it fails. When specifying availability requirements in the networking field, MTBF is sometimes designated with the more cumbersome phrase *mean time between service outage* (MTBSO), to account for the fact that a network is a service, not a component. Similarly, MTTR can be replaced with the phrase *mean time to service repair* (MTTSR). This book uses the simpler and better-known terms MTBF and MTTR.

A typical MTBF goal for a network that is highly relied upon is 4,000 hours. In other words, the network should not fail more often than once every 4,000 hours or 166.67 days. A typical MTTR goal is one hour. In other words, the network failure should be fixed within one hour. In this case, the mean availability goal is

$$4,000 / 4,001 = 99.98 \text{ percent}$$

A goal of 99.98 percent is typical for mission-critical operations.

When specifying availability using MTBF and MTTR, the equation to use is as follows:

$$\text{Availability} = \text{MTBF} / (\text{MTBF} + \text{MTTR})$$

Using this availability equation allows a customer to clearly state the acceptable frequency and length of network outages.

Remember that what is calculated is the mean. The variation in failure and repair times can be high and must be considered as well. It is not enough to just consider mean rates, especially if you depend on external service agents (vendors or contractors) who are not under your tight control. Also, be aware that customers might need

to specify different MTBF and MTTR goals for different parts of a network. For example, the goals for the core of the enterprise network are probably much more stringent than the goals for a switch port that only affects one user.

Although not all customers can specify detailed application requirements, it is a good idea to identify availability goals for specific applications, in addition to the network as a whole. Application availability goals can vary widely depending on the cost of downtime. For each application that has a high cost of downtime, you should document the acceptable MTBF and MTTR.

For MTBF values for specific networking components, you can generally use data supplied by the vendor of the component. Most router, switch, and hub manufacturers can provide MTBF and MTTR figures for their products. You should also investigate other sources of information, such as trade publications, to avoid any credibility problems with figures published by manufacturers. Search for variability figures as well as mean figures. Also, try to get written commitments for MTBF, MTTR, and variability values from the providers of equipment and services.

NETWORK PERFORMANCE

When analyzing technical requirements for a network design, you should isolate your customer's criteria for accepting the performance of a network, including throughput, accuracy, efficiency, delay, and response time.

Many mathematical treatises have been written on network performance. This book approaches network performance in a practical and mostly non-mathematical way, avoiding the daunting equations that appear in mathematical treatments of performance. Although the equations are much simpler than they seem, they are usually not necessary for understanding a customer's goals. The objective of this section is to offer an uncomplicated view of network performance, including real-world conclusions you can draw when there is no time to do a mathematical analysis.

Analyzing a customer's network performance goals is rightly tied to analyzing the existing network, which is covered in Chapter 3. Analyzing the existing network will help you determine what changes need to be made to meet performance goals. Network performance goals are also tightly linked to scalability goals. You should gain an understanding of plans for network growth before analyzing performance goals.

Network Performance Definitions

Many network design customers cannot quantify their performance goals beyond, “It has to work with no complaints from users.” If this is the case, you can make assumptions regarding throughput, response time, and so on. On the other hand, some customers have specific performance requirements, based on a service level that has been agreed upon with network users. The following list provides definitions for network performance goals that you can use when analyzing precise requirements:

- **Capacity (bandwidth).** The data-carrying capability of a circuit or network, usually measured in bits per second (bps)
- **Utilization.** The percent of total available capacity in use
- **Optimum utilization.** Maximum average utilization before the network is considered saturated
- **Throughput.** Quantity of error-free data successfully transferred between nodes per unit of time, usually seconds
- **Offered load.** Sum of all the data all network nodes have ready to send at a particular time
- **Accuracy.** The amount of useful traffic that is correctly transmitted, relative to total traffic
- **Efficiency.** A measurement of how much effort is required to produce a certain amount of data throughput
- **Delay (latency).** Time between a frame being ready for transmission from a node and delivery of the frame elsewhere in the network
- **Delay variation.** The amount of time average delay varies
- **Response time.** The amount of time between a request for some network service and a response to the request

Optimum Network Utilization

Network utilization is a measurement of how much bandwidth is used during a specific time period. Utilization is commonly specified as a percentage of capacity. For example, a network-monitoring tool might state that network utilization on an Ethernet segment is 30 percent, meaning that 30 percent of the capacity is in use.

Network analysis tools use varying methods for measuring bandwidth usage and averaging the usage over elapsed time. Usage can be averaged every millisecond, every second, every minute, every hour, and so on. Some tools use a weighted average whereby more recent values are weighted more prominently than older values. Chapter 3, "Characterizing the Existing Internetwork," discusses measuring network utilization in more depth.

Your customer might have a network design goal for the maximum average network utilization allowed on shared segments. Actually, this is a design constraint more than a design goal. The design constraint states that if utilization on a segment is more than a pre-defined threshold, then that segment must be divided into multiple shared or switched segments.

A typical "rule" for shared Ethernet is that average utilization should not exceed 37 percent, because beyond this limit, the collision rate allegedly becomes excessive. This is not a hard-and-fast rule. The 37 percent limit comes from studies done by the Institute of Electrical and Electronics Engineers (IEEE) comparing carrier sense multiple access collision detection (CSMA/CD) to token passing.

Token passing makes a node wait for a token before sending. At modest loads, this wait means that token passing causes more delay (latency) than Ethernet. If more stations are added to a token ring, then the latency is even worse because the token must pass through each station.

However, at around 37 percent utilization on a medium shared by 50 stations, Ethernet frames experience more delay than token ring frames, because the rate of Ethernet collisions becomes significant. (The study used 128-byte frames and compared 10-Mbps Ethernet to 10-Mbps token passing. The results are only slightly different if 4-Mbps or 16-Mbps token ring is used.)

The key point of the IEEE study was that token passing extracts a higher toll for each station added. For 100 stations, Ethernet frames start experiencing more delay than token ring frames at 49 percent load, instead of the 37 percent load for 50 stations. Armed with this knowledge about the IEEE study, you can help your customer understand which, if any, maximum network-utilization goals are appropriate.

Consider the case of an Ethernet segment that is shared by only two stations: a client that sends requests and a server that responds after receiving requests. In this case, is it a problem if network utilization exceeds 37 percent? There are no collisions because the server and client never try to send at the same time, so the 37 percent rule, which is concerned with collisions, does not apply. The load should be almost 100 percent unless the client or server are slow.

In the case of token passing technologies, such as Token Ring and Fiber Distributed Data Interface (FDDI), a typical goal for optimum average network utilization is 70 percent. Collisions are not a factor for token ring networks, but a goal of 100 percent average utilization is unrealistic. A 70 percent threshold for average utilization means that peaks in network traffic can probably be handled without obvious performance degradation. (The variability of typical loads would need to be studied to determine the exact behavior of peaks.)

If customers have a goal of reducing network utilization to free up capacity for new applications, you can work with them to select technologies that curtail bandwidth usage. For example, in a LAN environment, Novell NetWare users should upgrade to client software that takes advantage of Novell's burst-mode and Sequenced Packet Exchange version II (SPX II) technologies.

For wide area networks (WANs), optimum average network utilization is about the same as for token ring networks—70 percent. A network utilization of 70 percent should support peaks in network traffic caused by unexpected downloads from remote sites. Most WANs have less capacity than LANs, so more care is needed in selecting bandwidth that can cover actual, reasonable variations. Customers have many options for technologies that can reduce bandwidth utilization on WANs, including advanced routing-protocol features, compression, repetitive pattern suppression (RPS), and voice activity detection (VAD). Optimizing bandwidth utilization is covered in more detail in Chapter 12, "Optimizing Your Network Design."



RPS, also called *data frame multiplexing (DFM)* or *run-length encoding*, is an option for WAN data circuits that replaces repeating strings of data by a single occurrence of the string and a code that indicates to the far end how many repetitions of the string were in the original data.

VAD is a technology that compresses voice traffic by not sending packets in the absence of speech. Other types of traffic can use the extra bandwidth saved.


Throughput

Throughput is defined as the quantity of error-free data that is transmitted per unit of time. Throughput is often defined for a specific connection or session, but in some cases the total throughput of a network is specified. Ideally, throughput should be the same as capacity. However, this is not the case on real networks.

Capacity depends on the physical-layer technologies in use. The capacity of a network should be adequate to handle the offered load, even when there are peaks in network traffic. (Offered load is the data that all nodes have to send at a particular moment in time.) Theoretically, throughput should increase as offered load increases, up to a maximum of the full capacity of the network. However, network throughput depends on the access method (for example, token passing or collision detection), the load on the network, and the error rate.

Figure 2-1 shows the ideal situation, where throughput increases linearly with the offered load, and the real world, where throughput tapers off as the offered load reaches a certain maximum.

Figure 2-1
Offered
load and
throughput.



Throughput of Internetworking Devices

Some customers specify throughput goals in terms of the number of packets per second (PPS) an internetworking device must process. (In the case of an ATM device, the goal is cells per second, or CPS.) The throughput for an internetworking device is the maximum rate at which the device can forward packets without dropping any packets.

Most internetworking vendors publish PPS ratings for their products, based on their own tests and independent tests. To test an internetworking device, engineers place the device between traffic generators and a traffic checker. The traffic generators send packets ranging in size from 64 bytes to 1,518 bytes for Ethernet. By running multiple generators, the investigation can test devices with multiple ports.

The generators send bursts of traffic through the device at an initial rate that is half of what is theoretically possible for test conditions. If all packets are received, the rate is increased. If all packets are not received, the rate is decreased. This process is repeated until the highest rate at which packets can be forwarded without loss is determined. PPS values for small frames are much higher than PPS values for large frames, so be sure you understand which value you are looking at when reading vendor test results for an internetworking device.

Many internetworking devices can forward packets at the theoretical maximum, which is also called wire speed. The theoretical maximum is calculated by dividing bandwidth by packet size, including any headers, preambles, and inter-frame gaps. Table 2-1 shows the theoretical maximum PPS for one Ethernet stream, based on frame size.

Table 2-1 Maximum Packets Per Second (PPS)

Frame Size (in bytes)	10-Mbps Ethernet Maximum PPS
64	14,880
128	8,445
256	4,528
512	2,349
768	1,586
1,024	1,197
1,280	961
1,518	812

To understand the PPS value for a multiport device, testers send multiple streams of data and compare the results to the theoretical maximum. For example, a Cisco Catalyst 5000 switch can forward to 30 ports the theoretical maximum throughput of 30 Ethernet streams of 64-byte packets, which is

$$14,880 \times 30 = 446,400 \text{ PPS}$$

It can also forward to 30 ports the theoretical maximum for 1,518-byte packets which is

$$812 \times 30 = 24,360 \text{ PPS}$$

Application-Layer Throughput

Most end users are concerned about the throughput rate for applications. Marketing materials from some networking vendors refer to application-layer throughput as *goodput*. Calling it goodput sheds light on the fact that it is a measurement of good and relevant application-layer data transmitted per unit of time.

It is possible to improve throughput such that more data per second is transmitted, but not increase goodput, because the extra data transmitted is overhead or retransmissions. It is also possible to increase throughput by not using compression. More data is transmitted per unit of time, but the user sees worse performance.

A simple goal for throughput based on data-per-second rates between stations does not identify the requirements for specific applications. When specifying throughput goals for applications, make it clear that the goal specifies good (error-free) application-layer data per unit of time. Application-layer throughput is usually measured in kilobytes or megabytes per second.

Work with your customer to identify throughput requirements for all applications that can benefit from maximized application-layer throughput, such as file transfer and database applications. (Throughput is not important for all applications, for example, some interactive character-based applications that don't need large screen updates.) Explain to your customer the factors that constrain application-layer throughput, which include the following:

- End-to-end error rates
- Protocol functions, such as handshaking, windows, and acknowledgments
- Protocol parameters, such as frame size and retransmission timers
- The PPS or CPS rate of internetworking devices
- Lost packets or cells at internetworking devices

- Workstation and server performance factors:
 - Disk-access speed
 - Disk-caching size
 - Device driver performance
 - Computer bus performance (capacity and arbitration methods)
 - Processor (CPU) performance
 - Memory performance (access time for real and virtual memory)
 - Operating-system inefficiencies
 - Application inefficiencies or bugs

If necessary, work with your customer to identify application throughput problems caused by errors or inefficiencies in protocols, operating systems, and applications. Protocol analyzers are important tools for this. Chapter 3, “Characterizing the Existing Internetwork,” discusses isolating performance problems in more detail.

Accuracy

The overall goal for accuracy is that the data received at the destination must be the same as the data sent by the source. Typical causes of data errors include power surges or spikes, impedance mismatch problems, poor physical connections, failing devices, and noise caused by electrical machinery. Sometimes software bugs can cause data errors also, though software problems are a less common cause of errors than physical-layer problems. Frames that have an error must be retransmitted, which has a negative effect on throughput. In the case of IP networks, the Transmission Control Protocol (TCP) provides retransmission of data.

For WAN links, accuracy goals can be specified as a bit error rate (BER) threshold. If the error rate goes above the specified BER, then the accuracy is considered unacceptable. Analog links have a typical BER threshold of about 1 in 10^5 . Digital circuits have a much lower error rate than analog circuits, especially if fiber-optic cable is used. Fiber-optic links have an error rate of about 1 in 10^{11} . Copper links have an error rate of about 1 in 10^6 .

For LANs, a BER is not usually specified, mainly because measuring tools, such as protocol analyzers, focus on frames, not bits. But you can approximate a BER by comparing the number of frames with errors in them to the total number of bytes seen by the measuring tool. A good threshold to use is that there should not be more than one bad frame per 10^6 bytes of data.

On shared Ethernet, errors are often the result of collisions. Two stations try to send a frame at the same time and the resulting collision damages the frames, causing Cyclic Redundancy Check (CRC) errors. Depending on the size of the Ethernet network, many of these collisions happen in the 8-byte preamble of the frames and are not registered by troubleshooting tools. If the collision happens past the preamble and somewhere in the first 64 bytes of the data frame, then this is registered as a legal collision and the frame is called a runt frame. A general goal for Ethernet collisions is that less than 0.1 percent of the frames should be affected by a legal collision (not counting the collisions that happen in the preamble).

A collision that happens beyond the first 64 bytes of a frame is a late collision. Late collisions are illegal and should never happen. Ethernet networks that are too large experience late collisions because stations sending minimum-sized frames cannot hear other stations within the allowed timeframe. The extra propagation delay caused by the excessive size of the network causes late collisions between the most widely-separated nodes. Faulty repeaters and network interface cards can also cause late collisions.

In the case of Token Ring networks, accuracy goals sometimes include goals for minimizing media-access control (MAC) error reports. Token Ring includes a rich MAC-layer protocol for reporting problems. MAC error-reporting frames are the result of physical-layer problems or the result of the normal insertion and de-insertion of ring stations. Any goals for reducing MAC error-reporting frames must take into account that one or two errors are normal when a station enters or leaves the ring.

It is also normal to see a ring purge frame from the Token Ring active monitor when a station enters or leaves the ring. A *ring purge frame* reinitializes the network. The active monitor sends a ring purge frame when the token gets lost. If the ring purge

frame returns to the active monitor, then the active monitor knows that the ring is operational. The active monitor then initiates a ring poll and, if the poll succeeds, releases a token that can be captured for data transmission. Two seconds after sending the ring purge frame, the active monitor sends an error report, notifying the ring error monitor that the token was lost.

Efficiency


Efficiency is a term borrowed from engineering and scientific fields. It is a measurement of how effective an operation is in comparison to the cost in effort, energy, time, or money. Efficiency specifies how much overhead is required to produce a required outcome. For example, you could measure the efficiency of a method for boiling water. Does most of the energy go to actually boiling the water or does a lot of the energy get wasted heating the electrical wiring, the pot the water is in, and the air around it? How much overhead is required to produce the desired outcome?

Efficiency also provides a useful way to talk about network performance. For example, shared Ethernet, as we have discussed, is inefficient when the collision rate is high. (The amount of effort to successfully send a frame becomes considerable because so many frames experience collisions.) Network efficiency specifies how much overhead is required to send traffic, whether that overhead is caused by collisions, token-passing, error-reporting, rerouting, acknowledgments, large frame headers, and so on.

Large frame headers are an obvious cause for inefficiency. A good network performance goal is that applications should minimize the amount of bandwidth used by headers by using the largest possible frame the MAC layer allows. Using a large frame maximizes the amount of useful application data compared to header data, and improves application-layer throughput.

Figure 2-2 shows a bandwidth pipe used by small frames and the same pipe used by large frames. The header of each frame is shaded. Note that there is an inter-frame gap between each frame in addition to the headers. From the graphic, you can see that large frames use bandwidth more efficiently than small frames.

Figure 2-2
Bandwidth utilization efficiency for small versus large frames.



The maximum frame size is a tradeoff with the BER discussed in the previous section. Bigger frames have more bits and hence are more likely to be hit by an error. If there were no errors, an infinitely big frame would be the most efficient (although not the most fair to other senders!). If a frame is hit by an error, then it must be retransmitted, which wastes time and effort and reduces efficiency. The bigger the frame, the more bandwidth is wasted retransmitting. So, because networks experience errors, frame sizes are limited to maximize efficiency (and provide fairness), as shown in Table 2-2.

Table 2-2 Maximum Frame Sizes

Technology	Maximum Valid Frame
10- and 100-Mbps Ethernet	1,518 bytes (including the header and CRC)
4-Mbps Token Ring	4,500 bytes
16-Mbps Token Ring	18,000 bytes
FDDI	4,500 bytes
ATM with ATM Adaptation Layer 5 (AAL5)	65,535 bytes (AAL5 payload size)
ISDN Basic Rate Interface (BRI) and Primary Rate Interface (PRI) using the Point-to-Point Protocol (PPP)	1,500 bytes
T1	Not specified but 4,500 bytes generally used

Delay and Delay Variation

Users of interactive applications expect minimal delay in receiving feedback from the network. In addition, users of multimedia applications require a minimal variation in

the amount of delay that packets experience. Delay must be constant for voice and video applications. Variations in delay, called *jitter*, cause disruptions in voice quality and jumpiness in video streams.

Older applications, such as SNA-based applications, are also sensitive to delay. In traditional SNA environments, delay could be carefully planned and measured. When SNA is migrated to a multiprotocol network, however, it is harder to predict delay. Unexpected delays can occur because of bursty LAN traffic and routers taking a long time to converge after a link outage. Long delays cause SNA sessions to timeout.

In a multiprotocol internetwork, SNA runs above the connection-oriented (Type 2) Logical Link Control protocol, also known as LLC2. With LLC2, when a station sends a frame, it expects a response within a timeframe defined by the T1 timer. If the station does not receive a response, it resends the frame. The number of times to resend is set by the retries counter. The default settings are two seconds for T1 and six for the number of retries.

In complex multiprotocol networks, you might have to increase the T1 timer or retries counter to guarantee that delay does not exceed

$$T1 + (T1 \times \text{retries})$$

You can also avoid timeouts by configuring a local router to respond to LLC2 frames. (Chapter 7, "Selecting Bridging, Switching, and Routing Protocols," covers these issues in more detail.)

Applications that use the Telnet protocol are also sensitive to delay because the user expects quick feedback when typing characters. With the Telnet `remote echo` option, the character typed by a user doesn't appear on the screen until it has been acknowledged and echoed by the far end, and the near end has sent an acknowledgment for

the echo. You should determine if your customer plans to run any applications based on delay-sensitive protocols, such as Telnet or SNA. Digital Equipment Corporation's Local Area Transport (LAT) protocol is also sensitive to delay and supports few adjustments for improving its behavior on networks with high delay.

Causes of Delay

Any goals regarding delay must take into account fundamental physics. Despite science fiction, any signal experiences a propagation delay resulting from the finite speed of light, which is about 300,000 kilometers per second (or 186,000 miles per second for metric-challenged readers in the United States). Network designers can also remember 1 nanosecond per foot. These values are for light traveling in a vacuum. A signal in a cable or optical fiber travels approximately 2/3 the speed of light in a vacuum.

Delay is relevant for all data transmission technologies, but especially for satellite links and long terrestrial cables. Geostationary satellites are in orbit above the earth at a height of about 36,000 kilometers, or 24,000 miles. This long distance leads to a delay of about 270 milliseconds (ms) for an intercontinental satellite hop. In the case of terrestrial cable connections, delay is about 1 ms for every 200 kilometers (120 miles).

Another fundamental cause for delay is the time to put digital data onto a transmission line, which depends on the data volume and the speed of the line. For example, to transmit a 1,024-byte packet on a 1.544-Mbps T1 line takes about 5 ms.

An additional fundamental delay is packet-switching delay. *Packet-switching delay* refers to the latency accrued when bridges, switches, and routers forward data. The latency depends on the speed of the internal circuitry and CPU, and the switching architecture of the internetworking device. The delay can be quite small. Scott Bradner of the Network Device Testing Laboratory at Harvard University periodically conducts latency tests on submitted network equipment. Bradner has tested Layer-2 and Layer-3 switches with latencies in the 10 to 50 microsecond range for 64-byte Ethernet IP packets. Routers have higher latencies than switches, but router vendors continually make progress on reducing latency.

When a packet comes into a router, the router checks its routing table, decides which interface should send the packet, and encapsulates the packet with the correct header and trailer. Routing vendors, such as Cisco Systems, have advanced caching mechanisms so that a frame destined for a known destination can receive its new encapsu-

lation very quickly without requiring the CPU to do any table lookup or other processing. These mechanisms minimize packet-switching delay.

Packet-switching delay can also include *queuing delay*. The number of packets in a queue on a packet-switching device increases exponentially as utilization increases, as you can see from Figure 2-3. If utilization is 50 percent, the average queue depth is 1 packet. If utilization is 90 percent, the average queue depth is 9 packets. Without going into mathematical queuing theory, the general rule of thumb for queue depth is:

$$\text{queue depth} = \text{utilization} / (1 - \text{utilization})$$

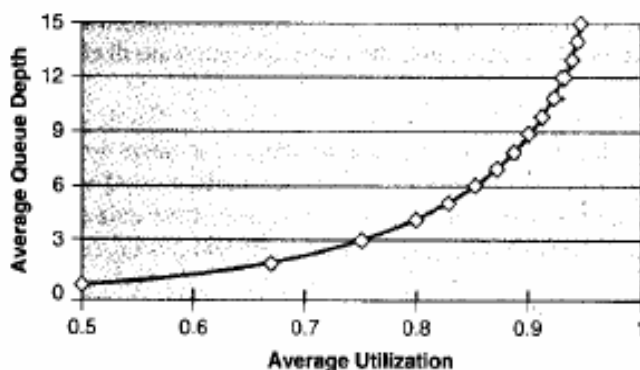


Figure 2-3
Queue
depth and
bandwidth
utilization.

Consider the following example. A packet switch has five users, each offering packets at a rate of 10 packets per second. The average length of the packets is 1,024 bits. The packet switch needs to transmit this data over a 56-Kbps WAN circuit.

$$\text{Load} = 5 \times 10 \times 1,024 = 51,200 \text{ bps.}$$

$$\text{Utilization} = 51,200 / 56,000 = 91.4 \text{ percent}$$

$$\text{Average number of packets in the queue} = (0.914) / (1 - 0.914) = 10.63 \text{ packets}$$

By increasing bandwidth on a WAN circuit you can decrease queue depth and hence decrease delay. Alternately, to improve performance, you can use an advanced queuing algorithm that outputs certain types of packets first—for example, voice or video packets. Advanced router queuing techniques are discussed in more detail in Chapter 12, “Optimizing Your Network Design.”

Delay Variation

As customers implement new digital voice and video applications, they are becoming concerned about delay and delay variation. Additionally, customers are becoming more aware of the issues associated with supporting bursty LAN traffic on the same network that carries delay-sensitive traffic. If bursts in LAN traffic cause jitter, audio and video streams experience problems that disrupt communications.

Desktop audio/video applications can minimize jitter by providing a buffer that the network puts data into. Display software or hardware pulls data from the buffer. The insulating buffer reduces the effect of jitter because variations on the input side are smaller than the total buffer size and therefore not obvious on the output side. The data is smoothed on the output, and the user experiences no ill effects from the input jitter.

If possible, you should gather exact requirements for delay variation from a customer. For customers who cannot provide exact goals, a good rule of thumb is that the variation should be less than one or two percent of the delay. For example, for a goal of an average delay of 40 ms, the variation should not be more than 400 or 800 microseconds.

Short fixed-length cells, for example ATM 53-byte cells, are inherently better than frames for meeting delay and delay-variance goals. To help understand this concept, consider the analogy of people trying to get onto an escalator. The escalator is like a bandwidth pipe. At first, each person gets onto the escalator in an orderly fashion and the delay is predictable. Then a school class arrives and the children are all holding hands, expecting to get onto the escalator all at once! What happens to your delay if you happen to be behind the children?

A gaggle of school children holding hands is analogous to a large frame causing extra delay for small frames. Consider the case of a user starting a file transfer using 1,518-byte frames. This user's data affects bandwidth usage and queuing mechanisms at internetworking devices, causing unexpected delay for other traffic. Good throughput for one application causes delay problems for another application.

Cell-relay technologies—for example, ATM—were designed to support traffic that is sensitive to delay and jitter. Depending on the class of service, ATM lets a session specify a maximum cell transfer delay (MCTD) and cell delay variation (MCDV). Chapter 4, "Characterizing Network Traffic," describes ATM service classes in more detail.

Response Time

Response time is the network performance goal that users care about most. Users don't know about propagation delay and jitter. They don't understand throughput in packets per second or in megabytes per second. They aren't concerned about bit-error rates, although perhaps they should be! Users recognize the amount of time to receive a response from the network system. They also recognize small changes in the expected response time and become frustrated when the response time is long.

Users begin to get frustrated when response time is more than about 100 ms or 1/10th of a second. Beyond 100 ms, users notice they are waiting for the network to display a Web page, echo a typed character, start downloading e-mail, and so on. If the response happens within 100 ms, most users do not notice any delay.

The 100-ms threshold is often used as a timer value for protocols that offer reliable transport of data. For example, many TCP implementations retransmit unacknowledged data after 100 ms by default. (Good TCP implementations also adjust the retransmit timer based on network conditions. TCP should keep track of the average amount of time to receive a response and dynamically adjust the retransmit timer based on the expected delay.)

The 100-ms response time threshold applies to interactive applications. For bulk applications, such as transferring large files or graphical Web pages, users are willing to wait at least 10 to 20 seconds. Technically savvy users expect to wait even longer if they know the file is large and the transmission medium is slow. If your network users are not technically savvy, you should provide some guidelines on how long to wait, depending on the size of files and the technologies in use (modems, high-speed digital networks, geostationary satellites, and so on).

SECURITY

Security design is one of the most important aspects of enterprise network design, especially as more companies add Internet and extranet connections to their internet networks. An overall goal that most companies have is that security problems should not disrupt the company's ability to conduct business. Network design customers need assurances that a design offers some protection against business data and other resources getting lost or damaged. Every company has trade secrets, business operations, and equipment to protect.

The first task in security design is planning. Planning involves analyzing risks and developing requirements. This chapter briefly discusses security planning. Chapter 8, "Developing Network Security and Network Management Strategies," covers planning for secure networks in more detail.

As is the case with most technical design requirements, achieving security goals means making tradeoffs. Security implementations can add to the cost of deploying and operating a network. Strict security policies can also affect the productivity of users, especially if some ease-of-use must be sacrificed to protect resources and data. Poor security implementations can annoy users, causing them to think of ways to get around security policies. Security can also affect the redundancy of a network design if all traffic must pass through encryption devices.

Security Risks

Ask your customer to help you understand the risks associated with not implementing a secure network. How sensitive is the customer's data? What would be the financial cost of someone accessing the data and stealing trade secrets? What would be the financial cost of someone changing the data?

As companies attach to the Internet they need to consider the additional risks of outsiders getting into the corporate network and doing damage. Customers who access remote sites across a Virtual Private Network (VPN) need to analyze the security features offered by the VPN service provider.

Some customers worry about hackers putting protocol analyzers on the Internet or VPN and sniffing packets to see passwords, credit-cards numbers, or other private data. This is not as big a risk as it appears. Credit-card numbers are almost always sent encrypted, using technologies such as the Secure Sockets Layer (SSL) protocol. Even when passwords or credit cards are not encrypted, it is extremely difficult to find these minute pieces of data in the midst of millions of packets.

On the other hand, hackers do have the ability to access and change sensitive data on enterprise networks. Consider the possibility of a hacker damaging an enterprise's image by changing the enterprise's public Web pages. You may have read about some of the cases of hackers changing U.S. government Web pages. These security breaches affected the government's image in two ways: the changed Web pages had silly graphics and text, and the government lost credibility because it appeared that it was easy to hack into government networks.

In general, hackers have the ability to attack computer networks in the following ways:

- Use resources they are not authorized to use
- Keep authorized users from accessing resources (also called denial-of-service attacks)
- Change, steal, or damage resources
- Take advantage of well-known security holes in operating systems and application software
- Take advantage of holes created while systems, configurations, and software releases are being upgraded

In addition to considering outside hackers as a security risk, companies should heed problems caused by inept or malicious internal network users. According to security surveys that Ernst and Young conducts every year, the biggest security problem facing companies today is software viruses that spread when users download software from untrusted sites. Companies reported that following viruses, the next most significant cause of problems was inadvertent user errors. This was followed by malicious acts by internal users. These problems were more common than malicious acts from the outside or industrial espionage.

Security Requirements

A customer's primary security requirement is to protect resources from being incapacitated, stolen, altered, or harmed. Resources can include network hosts, servers, user systems, internetworking devices, system and application data, and a company's image.

Other more specific requirements could include one or more of the following goals:

- Let outsiders (customers, vendors, suppliers) access data on public Web or FTP servers but not access internal data
- Authorize and authenticate branch-office users, mobile users, and telecommuters

- Detect intruders and isolate the amount of damage they do
- Authenticate routing-table updates received from internal or external routers
- Protect data transmitted to remote sites across a VPN
- Physically secure hosts and internetworking devices (for example, keep devices in a locked room)
- Logically secure hosts and internetworking devices with user accounts and access rights for directories and files
- Protect applications and data from software viruses
- Train network users and network managers on security risks and how to avoid security problems
- Implement copyright or other legal methods of protecting products and intellectual property

MANAGEABILITY

Every customer has different objectives regarding the manageability of a network. Some customers have precise goals, such as a plan to use the Simple Network Management Protocol (SNMP) to record the number of bytes each router receives and sends. Other clients have less specific goals. If your client has definite plans, be sure to document them, because you will need to refer to the plans when selecting equipment. In some cases, equipment has to be ruled out because it does not support the management functions a customer requires.

To help customers who don't have specific goals, you can use International Organization for Standardization (ISO) terminology to define network management functions:

- **Performance management.** Analyzing traffic and application behavior to optimize a network, meet service-level agreements, and plan for expansion
- **Fault management.** Detecting, isolating, and correcting problems; reporting problems to end users and managers; tracking trends related to problems
- **Configuration management.** Controlling, operating, identifying, and collecting data from managed devices

- **Security management.** Monitoring and testing security and protection policies, maintaining and distributing passwords and other authentication and authorization information, managing encryption keys, auditing adherence to security policies
- **Accounting management.** Accounting of network usage to allocate costs to network users and/or plan for changes in capacity requirements

Almost all customers have a need for fault and configuration management. Many customers also need performance and security management. Some customers need accounting management. Network management is discussed in more detail in Chapter 8, “Developing Network Security and Network Management Strategies.”

USABILITY

A goal that is related to manageability, but is not exactly the same as manageability, is usability. Usability refers to the ease-of-use with which network users can access the network and services. Whereas manageability focuses on making network managers’ jobs easier, usability focuses on making network users’ jobs easier.

It is important to gain an understanding of how important usability is to your network design customer, because some network design components can have a negative affect on usability. For example, strict security policies can have a negative affect on usability (which is a tradeoff that most customers are willing to make, but not all customers). You can plan to maximize usability by deploying user-friendly host-naming schemes and easy-to-use configuration methods that make use of dynamic protocols, such as the Dynamic Host Configuration Protocol (DHCP).

ADAPTABILITY

When designing a network, you should try to avoid incorporating any elements that would make it hard to implement new technologies in the future. A good network design can adapt to new technologies and changes. Changes can come in the form of new protocols, new business practices, new fiscal goals, new legislation, and a myriad of other possibilities. For example, some states have enacted environmental laws that require a reduction in the number of employees driving to work. To meet the legal requirement to reduce automobile emissions, companies need their remote-access designs to be flexible enough to adapt to increasing numbers of employees working at home.

The adaptability of a network affects its availability. For example, consider the need for a network to adapt to environmental changes. Some networks must operate in environments that change drastically from day to night or from winter to summer. Extreme changes in temperature can affect the behavior of electronic components of a network. A network that cannot adapt cannot offer good availability.

A flexible network design is also able to adapt to changing traffic patterns and quality of service (QoS) requirements. For some customers, the selected WAN or LAN technology must adapt to new users randomly joining the network to use applications that require a constant-bit-rate service. Chapter 4, "Characterizing Network Traffic," discusses QoS requirements in more detail.

One other aspect of adaptability is how quickly internetworking devices must adapt to problems and to upgrades. For example, how quickly do switches and bridges adapt to another switch failing, causing a change in the spanning-tree topology? How quickly do routers adapt to new networks joining the topology? How quickly do routing protocols adapt to link failures? These issues are discussed in more detail in Chapter 7, "Selecting Bridging, Switching, and Routing Protocols."

AFFORDABILITY

The final technical goal this chapter covers is *affordability*. Affordability is sometimes called *cost-effectiveness*. Most customers have a goal for affordability, though sometimes other goals such as performance and availability are more important. Affordability is partly a business goal, and, in fact, was discussed in Chapter 1, "Analyzing Business Goals and Constraints." It is covered again in this chapter because of the technical issues.

The primary goal of affordability is to carry the maximum amount of traffic for a given financial cost. Financial costs include non-recurring equipment costs and recurring network operation costs.

In campus networks, low cost is often a primary goal. Customers expect to be able to purchase affordable switches that have numerous ports and a low cost per port. They expect cabling costs to be minimal and service-provider charges to be minimal or non-existent. They also expect network interface cards (NICs) for end systems and servers to be inexpensive. Depending on the applications running on end systems, low cost is often more important than availability and performance in campus network designs.

For enterprise networks, availability is usually more important than low cost. Nonetheless, customers are looking for ways to contain costs for enterprise networks. Recurring monthly charges for WAN circuits are the most expensive aspect of running a large network. To reduce the cost of operating a WAN, customers often have one or more of the following technical goals for affordability:

- Use a routing protocol that minimizes WAN traffic
- Use a routing protocol that selects minimum-tariff routes
- Consolidate parallel leased lines carrying voice and data into fewer WAN trunks
- Select technologies that dynamically allocate WAN bandwidth, for example, ATM rather than time-division multiplexing (TDM)
- Improve efficiency on WAN circuits by using such features as compression, voice activity detection (VAD), and repetitive pattern suppression (RPS)
- Eliminate underutilized trunks from the internetwork and save money by eliminating both circuit costs and trunk hardware
- Use technologies that support oversubscription

With old-style TDM networks, the core backbone capacity had to be at least the sum of the speeds of the incoming access networks. With cell and frame switching, oversubscription is common. Because of the bursty nature of frame-based traffic, access-port speeds can add up to more than the speed of a backbone network, within reason. Enterprise network managers who have a goal of reducing operational costs are especially interested in solutions that will let them oversubscribe their trunks, while still maintaining service guarantees they have offered their users.

The second most expensive aspect of running a network, following the cost of WAN circuits, is the cost of hiring, training, and maintaining personnel to operate and manage the network. To reduce this aspect of operational costs, customers have the following goals:

- Select internetworking equipment that is easy to configure, operate, maintain, and manage

- Select a network design that is easy to understand and troubleshoot
- Maintain good network documentation to reduce troubleshooting time
- Select network applications and protocols that are easy to use so that users can support themselves to some extent

Making Network Design Tradeoffs

When analyzing a customer's goals for affordability, it is important to gain an understanding of how important affordability is compared to other goals. Despite what politicians tell us about federal budgets during an election year, in the real world, meeting goals requires making tradeoffs. This section describes some typical network design tradeoffs.

To meet high expectations for availability, redundant components are often necessary, which raises the cost of a network implementation. To meet rigorous performance requirements, high-cost circuits and equipment are required. To enforce strict security policies, expensive monitoring might be required and users must forgo some ease-of-use. To implement a scalable network, availability might suffer, because a scalable network is always in flux as new users and sites are added. To implement good throughput for one application might cause delay problems for another application.

To implement affordability might mean availability must suffer. One cause of network problems can be inadequate staffing and reduced training due to overzealous cost-cutting. These mistakes have hurt many sizeable organizations and are hard to recover from. Once the network staff is gone, outsourcing becomes a necessity, which *may end up being more costly*.

To help you analyze tradeoffs, ask your customer to identify a single driving network design goal. This goal can be the same overall business goal for the network design project that was identified in Chapter 1, or it can be a rephrasing of that goal to include technical issues. In addition, ask your customer to prioritize the rest of the goals. Prioritizing will help the customer get through the process of making tradeoffs.

One analogy that helps with prioritizing goals is the “kid in the candy store with a dollar bill” analogy. Using the dollar-bill analogy, explain to the customer that he or she is like a child in a candy store that has exactly one dollar to spend. The dollar can be spent on different types of candy: chocolates, licorice, jelly beans, and so on. But each time more money is spent on one type of candy, less money is available to spend on other types. Ask customers to add up how much they want to spend on scalability, availability, network performance, security, manageability, usability, adaptability, and affordability. For example, a customer could make the following selections:

Scalability	20
Availability	30
Network performance	15
Security	5
Manageability	5
Usability	5
Adaptability	5
Affordability	15
Total (must add up to 100)	100

Keep in mind that sometimes making tradeoffs is more complex than what has been described because goals can differ for various parts of an internetwork. One group of users might value availability more than affordability. Another group might deploy state-of-the-art applications and value performance more than availability. In addition, sometimes a particular group’s goals are different than the overall goals for the internetwork as a whole. If this is the case, document individual group goals as well as goals for the network as a whole. Later when selecting network technologies, you might see some opportunities to meet both types of goals—for example choosing LAN technologies that meet individual group goals, and WAN technologies that meet overall goals.

TECHNICAL GOALS CHECKLIST

You can use the following checklist to determine if you have addressed all your client's technical objectives and concerns:

- I have documented the customer's plans for expanding the number of sites, users, and servers/hosts for the next year and next two years.
- The customer has told me about any plans to migrate departmental servers to server farms or intranets.
- The customer has told me about any plans to migrate an SNA network to the multiprotocol internetwork.
- The customer has told me about any plans to implement an extranet to communicate with partners or other companies.
- I have documented a goal for network availability in percent uptime and/or MTBF and MTTR.
- I have documented any goals for maximum average network utilization on shared segments.
- I have documented goals for network throughput.

Table 2-3 Network Applications Technical Requirements

Name of Application	Type of Application	New Application? (Yes or No)	Criticality	Cost of Downtime	Acceptable MTBF

- I have documented goals for PPS throughput of internetworking devices.
- I have documented goals for accuracy and acceptable BERs.
- I have discussed with the customer the importance of using large frame sizes to maximize efficiency.
- I have identified any applications that have a more restrictive response-time requirement than the industry standard of less than 100 ms.
- I have discussed network security risks and requirements with the customer.
- I have gathered manageability requirements, including goals for performance, fault, configuration, security, and accounting management.
- Working with my customer, I have developed a list of network design goals, including both business and technical goals. The list starts with one overall goal and includes the rest of the goals in priority order. Critical goals are marked as such.
- I have updated the Network Applications chart to include the technical application goals shown in Table 2-3.

Chapter 1, “Analyzing Business Goals and Constraints,” provided a Network Applications chart. At this point in the design process, you can expand the chart to include technical application requirements, such as MTBF, MTTR, and throughput and delay goals, as shown in Table 2-3.

Acceptable MTTR	Throughput Goal	Delay must be less than:	Delay variation must be less than:	Comments

SUMMARY

This chapter covered technical requirements for a network design, including scalability, availability, network performance, security, manageability, usability, adaptability, and affordability. It also covered typical tradeoffs that must be made to meet these goals.

Analyzing your customer's technical and business goals prepares you to carry out the next steps in the top-down network design process, including making decisions regarding network technologies to recommend to a customer. Researchers who study decision models say that one of the most important aspects of making a sound decision is having a good list of goals. At this point in the network design process, you have gathered both business and technical goals. You should make a list of your customer's most important technical goals and merge this list with the list of business goals you made in Chapter 1.

You should put the goals in the list in priority order, starting with the overall most important business and technical goal, and following with critical goals and then less critical goals. Later, you can make a list of options and correlate options with goals. Any options that do not meet critical goals can be eliminated. Other options can be ranked by how well they meet a goal. This process can help you select network components that meet a customer's requirements.