GLOBAL EDITION

Business Data Networks and Security

TENTH EDITION

Raymond R. Panko • Julia A. Panko

ALWAYS LEARNING

PEARSON

Business Data Networks and Security

create overall failure. Lockheed has suggested that companies should actively consider security kill chains and look for evidence that one of the steps is occurring. Success in identifying an operating kill chain may allow the company to stop it or at least disrupt or degrade it. The warnings when malware was put on the extrusion server could have done exactly that.

Until one understands likely kill chains in depth, however, it is impossible to understand that events are part of each kill chain. Conversely, understanding the kill chain can allow the company to act before a kill chain fitting that pattern begins. For example, even cursory thinking about charge card data theft would lead the company to realize that thieves would probably use FTP transfers to unusual servers, that command communication would probably use certain ports in firewalls, and so forth.

Even well-defended companies suffer security compromises. However, when strategic planning is not done, if protections are not put into place, or if the security staff is not aggressive in doing the work required for the protections to work, the risk of compromises becomes a near certainty. Security expert Ben Schneier said "Security is a process, not a product."[7] Boxes and software are not magic talismans.

Test Your Understanding

1. a) What security mistake did Fazio Mechanical Services make? b) Why do you think it did this? (This requires you to give an opinion.) c) How might segregation of the network have stopped the breach? d) Why do you think the Minneapolis security staff not heed the FireEye warning? (This requires you to give an opinion.) e) What warnings had Target not responded to adequately? f) What happens in a kill chain if a single action fails anywhere in the chain? g) How can kill chain analysis allow companies to identify security actions it should take? h) Explain why security is a process, not a product."

INTRODUCTION

In the first three chapters, we looked at general network concepts and security. However, technology means nothing unless a company manages it well. In this chapter, we will look at network and security planning. Although the concepts are broad, they apply to everything networking professionals do at every level.

Management is critical. Today, we can build much larger networks than we can manage easily. For example, even a mid-size bank is likely to have 500 Ethernet switches and a similar number of routers. Furthermore, network devices and their users are often scattered over large regions—sometimes internationally. While network technology is exciting to talk about and concrete conceptually, it is chaos without good management.

A pervasive issue in network management is cost. In networking, you never say, "Cost doesn't matter." Figure 4-1 illustrates that network demand is likely to grow rapidly

[7] Ben Schneier, "Computer Security: Will We Ever Learn?" *Crypto-Gram Newsletter*, May 15, 2000. https://www.schneier.com/crypto-gram-0005.html.

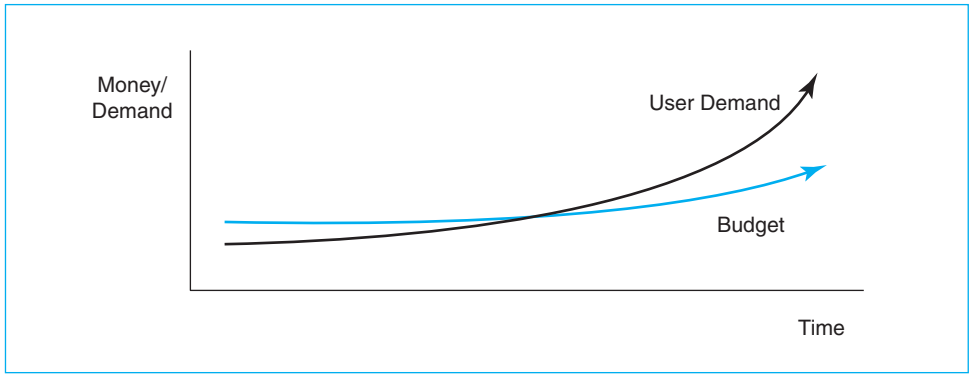

FIGURE 4-1 Network Demand and Budgets

in the future, just as it has in the past. The figure also illustrates that network budgets are growing slowly if they are growing at all.[8]

Taken together, these curves mean that network budgets are always stretched thin. If the network staff spends too much money on one project, it will not have enough left to do another important project. Although there are many concerns beyond costs, cost underlies everything in network management.

Test Your Understanding

2. a) Why is cost a crucial factor in network management? b) Identify other factors involved in network management.

NETWORK QUALITY OF SERVICE (QOS)

In the early days of the Internet, networked applications amazed new users. However, new users soon added, "Too bad it doesn't work better." Today, networks are mission-critical for corporations. If the network breaks down, much of the organization comes to a grinding and expensive halt. Networks must not only work. They must work *well*. Companies are increasingly concerned with network **quality-of-service (QoS) metrics**, that is, quantitative measures of network performance. Figure 4-2 shows that companies use a number of QoS metrics. Collectively, these metrics track the service quality that users receive.

Test Your Understanding

3. a) Why is it important for a firm's network to perform well? b) How do firms ensure this?

[8] In fact, costs for equipment and transmission lines are falling. This is especially true in cellular transmission. The chief technology officer of Ericsson has said that network efficiencies reduced the price per bit transmitted 50 percent per year from 2008 to 2013. During this time, the cost per megabit fell from 46 cents to 1 to 3 cents. However, transmission volume has doubled each year, so customer bills have not gone down. Stephen Lawson, "5G Will Have to Do More than Just Speed Up Your Phone, Ericsson Says," *PC World*, October 17, 2013. http://www.pcworld.com/article/2055880/5g-will-have-to-do-more-than-send-speed-up-your-phone-ericsson-says.html?tk=rel_news.

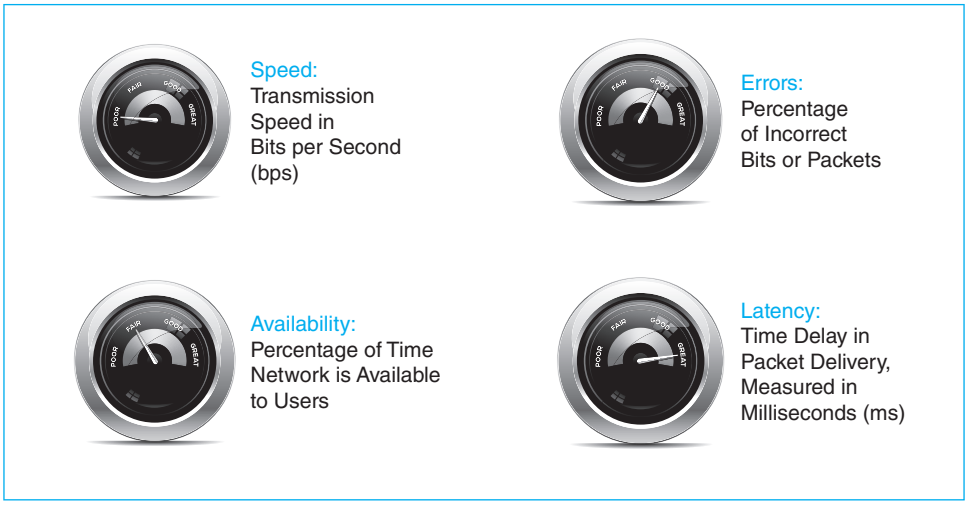

FIGURE 4-2 Quality-of-Service (QoS) Metrics

Transmission Speed

There are many ways to measure how well a network is working. The most fundamental metric, as we saw in Chapter 1, is speed. While low speeds are fine for text messages, the need for speed becomes very high as large volumes of data must be delivered, and video transmission requires increasingly higher transmission speeds.

Rated Speed versus Throughput and Aggregate Throughput

NOTE: Some students find the distinction between rated speed and throughput difficult to learn. However, we must use this distinction throughout this book, so be sure to take the time to understand it.

Rated Speed versus Throughput The term *transmission speed* is somewhat ambiguous. A transmission link's **rated speed** is the speed it *should* provide based on vendor claims or on the standard that defines the technology. For a number of reasons, transmission links almost always fail to deliver data at their full rated speeds. In contrast to rated speed, a network's **throughput** is the data transmission speed the network *actually* provides to users.

A transmission link's rated speed is the speed it should *provide based on vendor claims or on the standard that defines the technology.*

Throughput is the transmission speed a network actually *provides to users.*

Aggregate versus Individual Throughput Sometimes transmission links are shared. For example, if you are using a Wi-Fi computer in a classroom, you share the wireless access point's throughput with other users of that access point. In shared

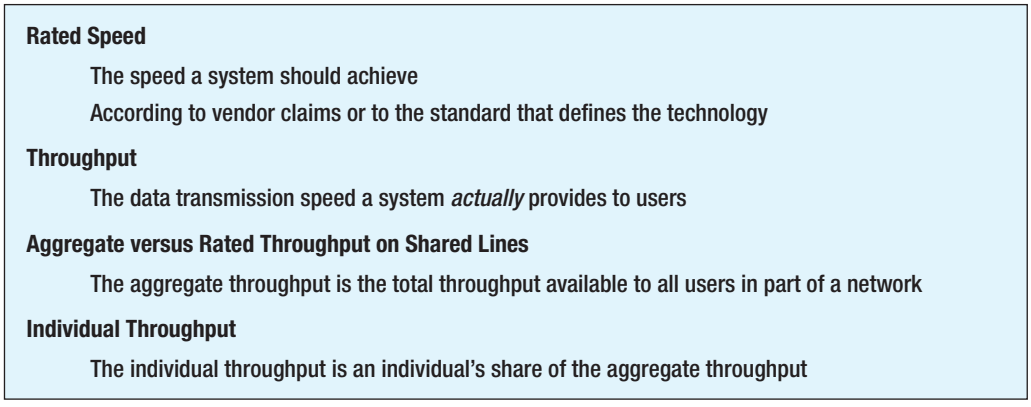

FIGURE 4-3 Rated Speed, Throughput, Aggregate Throughput, and Individual Throughput (Study Figure)

situations, it is important to distinguish between a link's **aggregate throughput**, which is the total it provides to all users who share it in a part of a network, and the link's **individual throughput** that single users receive as their shares of the aggregate throughput. Individual throughput is always lower than aggregate throughput. As you learned as a child, despite what your mother said, sharing is bad.

Test Your Understanding

4. a) Distinguish between rated speed and throughput. b) Distinguish between individual and aggregate throughput. c) You are working at an access point with 20 other people. Three are doing a download at the same time you are. The rest are looking at their screens or sipping coffee. The access point channel you share has a rated speed of 150 Mbps and a throughput of 100 Mbps. How much speed can you expect for your download? (Check figure: 33 Mbps). d) In a coffee shop, there are 10 people sharing an access point with a rated speed of 20 Mbps. The throughput is half the rated speed. Several people are downloading. Each is getting five Mbps. How many people are using the Internet at that moment?

Other Quality-of-Service Metrics

Although network speed is important, it is not enough to provide good quality of service. Figure 4-2 showed that there are other QoS categories. We will look briefly at three of them.

Availability One is **availability**, which is the percentage of time that the network is available for use. Ideally, networks would be available 100% of the time, but that is impossible in reality. On the Public Switched Telephone Network, the availability target usually is 99.999%. Availability on data networks is usually lower, although by carefully adding redundancy, Netflix and some other companies can reach telephone availability levels.

Error Rates Ideally, all packets would arrive intact, but a small fraction do not. The **error rate** is the percentage of bits or packets that are lost or damaged during delivery. (At the physical layer, it is common to measure bit error rates. At the internet layer, it is common to measure packet error rates.)

When the network is overloaded, error rates can soar because the network has to drop the packets it cannot handle. Consequently, companies must measure error rates when traffic levels are high in order to have a good understanding of error rate risks.

The impact of even small error rates can be surprisingly large. TCP tries to avoid network congestion by sending TCP segments slowly at the beginning of a connection. If these segments get through without errors, TCP sends the following segments more quickly. However, if there is a single error, the TCP process assumes that the network is overloaded. It falls back to its initial slow start rate for sending TCP segments. This can produce a major drop in throughput for applications.

Latency When packets move through a network, they will encounter some delays. The amount of delay is called **latency**. Latency is measured in **milliseconds (ms)**. A millisecond is a thousandth of a second. When latency reaches about 125 milliseconds, turn taking in telephone conversations becomes difficult. You think the other person has finished speaking, so you begin to speak—only to realize that the other party is still speaking.

Jitter A related concept is **jitter**, which Figure 4-4 illustrates. Jitter occurs when the latency between successive packets varies. Some packets will come farther apart in time, others closer in time. While jitter does not bother most applications, VoIP and streaming media are highly sensitive to jitter. If the sound is played back without adjustment, it will speed up and slow down. These variations often occur over millisecond times. As the name suggests, variable latency tends to make voice sound jittery.

Jitter is the average variability in arrival times (latency) divided by the average latency.

Engineering for Latency and Jitter Most networks were engineered to carry traditional data such as e-mail and database transmissions. In traditional applications, latency was only slightly important, and jitter was not important at all.

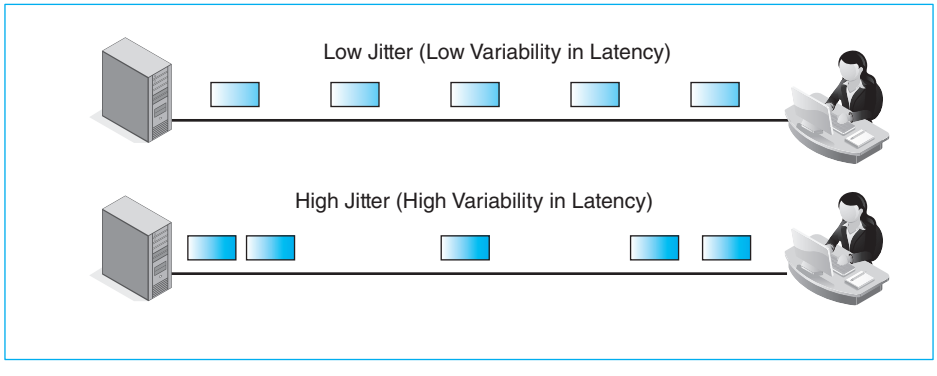

FIGURE 4-4 Jitter

However, as voice over IP (VoIP), video, and interactive applications have grown in importance, companies have begun to worry more about latency and jitter. They are finding that extensive network redesign may be needed to give good control over latency and jitter. This may include forklift upgrades for many of its switches and routers.

Test Your Understanding

5. a) What is availability? b) How does network availability usually compare to availability on the telephone network? c) When should you measure error rates? Why? d) When an application uses TCP at the transport layer, why is error rate a problem for throughput? e) What is latency? f) Give an example not listed in the text of an application for which latency is bad. g) What is jitter? h) Name an application not listed in the text for which is jitter a problem. i) Why may adding applications that cannot tolerate latency and jitter be expensive?

Service Level Agreements (SLAs)

When you buy some products, you receive a guarantee that promises that they will work and specifying what the company will do if they do not work as promised. In networks, service providers often provide **service level agreements (SLAs)**, which are contracts that guarantee levels of performance for various metrics such as speed and availability. If a service does not meet its SLA guarantees, the service provider must pay a penalty to its customers.

Worst-Case Specification SLA guarantees are expressed as **worst cases**. For example, an SLA for speed would guarantee that speed will be *no lower* than a certain amount. If you are downloading webpages, you want at least a certain level of speed.

Service Level Agreements (SLAs)

- Guarantees for performance
- Penalties if the network does not meet its service metrics guarantees

Guarantees specify worst cases (no worse than)

- Lowest speed (e.g., no worse than 1 Mbps)
- Maximum latency (e.g., no more than 125 ms)
- SLAs are like insurance policies

Often written on a percentage basis

- No worse than 100 Mbps 99.5% of the time
- Because as the percentage increases, additional engineering raises network costs
- 100% compliance would be prohibitively expensive

Residential services are rarely sold with SLA guarantees

- It would be too expensive

FIGURE 4-5 Service Level Agreements (SLAs) (Study Figure)