

Chapter 4

Network and Security Management

LEARNING OBJECTIVES

By the end of this chapter, you should be able to:

- Describe strategic network planning principles.
- Discuss network quality of service (QoS) and how to specify service level agreement (SLA) guarantees.
- Design a network layout based on required traffic volumes between sites, considering redundancy.
- Be able to describe options for dealing with momentary traffic peaks.
- Describe and apply strategic security planning principles.
- Describe the importance of centralized network and security management and discuss tools for centralizing network and security management. Explain how software-defined networking may revolutionize the way that networks are managed and what benefits SDN may bring.

FAILURES IN THE TARGET BREACH

After every breach, companies should pause to take lessons from the experience. This type of reflection, if it leads to appropriate changes, will reduce the odds of similar breaches in the future.

One lesson from the Target breach is that you cannot trust external businesses you deal with to have good security. In the case of Fazio Mechanical Services, an employee fell for a spear phishing attack. This could happen in any company. However, Fazio made it more likely. It used the free consumer version of an antivirus program, Malwarebytes

Anti-Malware.¹ This free version did not do real-time assessment for arriving e-mail messages and attachments. If Fazio had used a commercial antivirus program for their e-mail, the employee probably would have seen a warning that opening an attachment was a bad idea or even that a specific threat existed in the attachment.

The breach also taught a number of lessons about Target's security. After the attackers gained a foothold on the vendors' server, they were able to move into more sensitive parts of the network in order to download malware onto the POS terminals, compromise a server to create a holding server, and compromise another server to act as an extrusion server. The low-security and highly sensitive parts of the network should have been segregated. They were not, or at least not enough.

Another issue is that Target received explicit warnings when the attackers were setting up the extrusion server. The thieves had to download malware onto the extrusion server in order to take it over and to manage subsequent FTP transmission. Target used the FireEye intrusion detection program. Target's intrusion detection team notified the Minneapolis security staff that a high-priority event had occurred on November 30, 2013.² In addition, the thieves had trouble with the initial malware. They had to make additional updates on December 1 and December 3. These resulted in additional FireEye warnings being sent to Target's Minneapolis security group. Had Target followed up on these warnings, they could have stopped or at least reduced the data extrusion, which began on December 2.³

Target may have been lax in understanding the danger of POS attacks. In April and August, VISA had sent Target and other companies warnings about new dangers regarding POS data theft.⁴ It appears that Target's own security staff expressed concern for the company's exposure to charge card data theft.⁵ If target did not respond to this risk aggressively, this would have been another serious lapse.

Overall, Figure 3-1 showed that the thieves had to succeed at every step in a complex series of actions. Lockheed Martin's Computer Incident Response Team⁶ staff called this a *kill chain*, which is a term borrowed from the military. The kill chain concept was designed to visualize all of the manufacturing, handling, and tactical steps needed for a weapon to destroy its target. Failure in a single step in a kill chain will

¹ Brian Krebs, Email Attack on Vendor Set Up Breach at Target, February 14, 2014. <http://krebsonsecurity.com/2014/02/email-attack-on-vendor-set-up-breach-at-target/>.

² Michael Riley, Ben Elgin, Dune Lawrence, and Carol Matlack, *Missed Alarms and 40 Million Stolen Credit Card Numbers: How Target Blew It*, *Bloomberg Businessweek*, March 13, 2014. <http://www.businessweek.com/articles/2014-03-13/target-missed-alarms-in-epic-hack-of-credit-card-data>.

³ Aviv Raff, *Pos Malware Targeted Target*, *Seculert*, January 16, 2014. <http://www.seculert.com/blog/2014/01/pos-malware-targeted-target.html>.

⁴ Jim Finkle and Mark Hosenball, *Exclusive: More Well-Known U.S. Retailers Victims of Cyber Attacks – Sources*, *Reuters*, January 12, 2014. <http://www.reuters.com/article/2014/01/12/us-target-databreach-retailers-idUSBREA0B01720140112>.

⁵ Danny Yadron, Paul Ziobro, Devlin Barrett, *Target Warned of Vulnerabilities Before Data Breach*, *The Wall Street Journal*, February 14, 2014. <http://online.wsj.com/news/articles/SB10001424052702304703804579381520736715690>.

⁶ Eric M. Hutchins, Michael J. Cloppert, and Rohan M. Amin, *Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains*, Lockheed Martin, 2011. <http://www.lockheedmartin.com/content/dam/lockheed/data/corporate/documents/LM-White-Paper-Intel-Driven-Defense.pdf>.

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) Compare trends in network demand and network budgets. b) What are the implications of these trends?

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) What are QoS metrics? (Do not just spell out the acronym.) b) Why are QoS metrics important?

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.



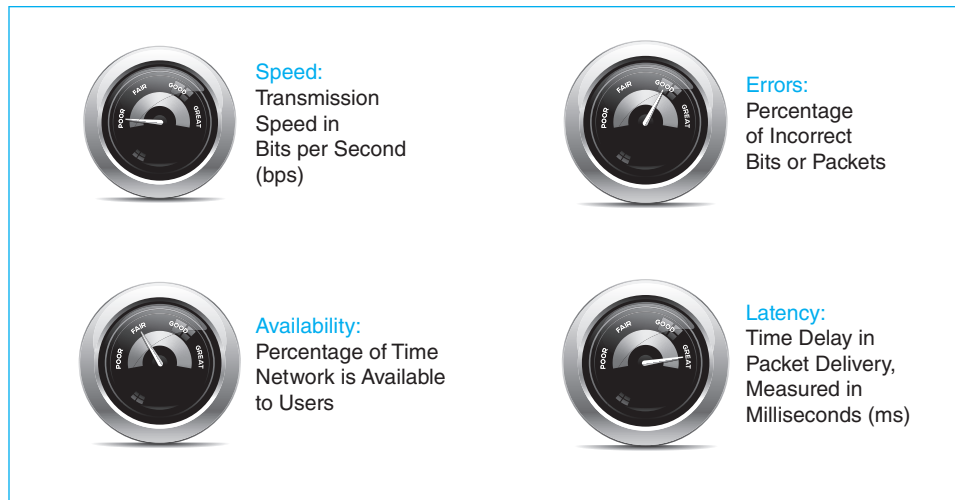


FIGURE 4-2 Quality-of-Service (QoS) Metrics

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

Rated Speed
The speed a system should achieve
According to vendor claims or to the standard that defines the technology
Throughput
The data transmission speed a system <i>actually</i> provides to users
Aggregate versus Rated Throughput on Shared Lines
The aggregate throughput is the total throughput available to all users in part of a network
Individual Throughput
The individual throughput is an individual's share of the aggregate throughput

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.

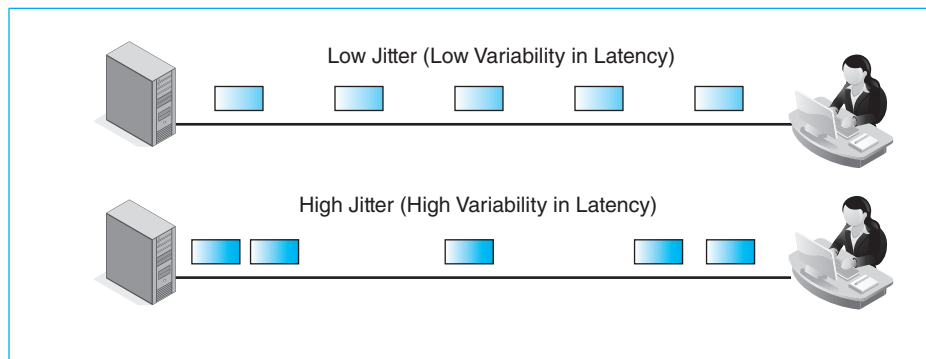


FIGURE 4-4 Jitter

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.

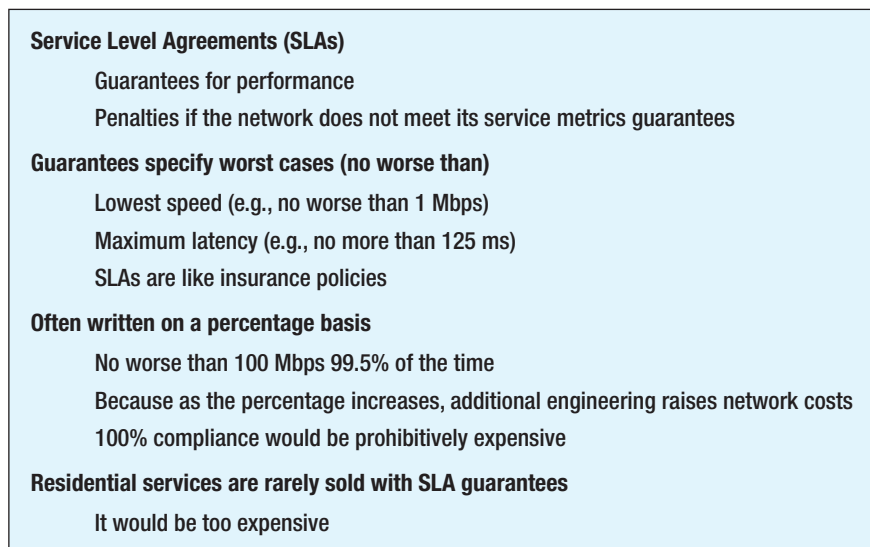


FIGURE 4-5 Service Level Agreements (SLAs) (Study Figure)

You certainly would not want a speed SLA to specify a *maximum* speed. More speed is good. Why would you want to impose penalties on the network provider for exceeding some maximum speed? That would give them a strong incentive not to increase speed! Making things better is not the SLA's job.

SLA guarantees are expressed as worst cases. Service will be no worse than a specific number.

For latency, in turn, an SLA would require that latency will be *no higher* than a certain value. You might specify an SLA guarantee of a maximum of 65 ms (milliseconds). This means that you will not get worse (higher) latency.

Percentage-of-Time Elements In addition, most SLAs have percentage-of-time elements. For instance, an SLA on speed might guarantee a speed of at least 480 Mbps 99.9% of the time. This means that the speed will nearly always be at least 480 Mbps but may fall below that 0.1% of the time without incurring penalties. A smaller exception percentage might be attractive to users, but it would require a more expensive network design. Nothing can be guaranteed to work properly 100% of the time, and beyond some point, cost grows very rapidly with increasing percentage guarantees.

Corporations versus Individuals Companies that use commercial networks expect SLA guarantees in their contracts, despite the fact that engineering networks to meet these guarantees will raise costs and prices. Consumer services, however, rarely have SLAs because consumers are more price sensitive. In particular, residential Internet access service using DSL, cable modem, or cellular providers rarely offer SLAs. This means that residential service from the same ISP may vary widely across a city.

Test Your Understanding

6. a) What are service level agreements? b) Does an SLA measure the best case or the worst case? c) Would an SLA specify a highest speed or a lowest speed? d) Would an SLA specify a highest availability or a lowest availability? e) Would an SLA specify highest latency or lowest latency? f) Would an SLA guarantee specify a highest jitter or a lowest jitter? g) What happens if a carrier does not meet its SLA guarantee? h) If carrier speed falls below its guaranteed speed in an SLA, under what circumstances will the carrier *not* have to pay a penalty to the customers? i) Does residential ISP service usually have SLA guarantees? Why? j) A business has an Internet access line with a maximum speed of 100 Mbps. What two things are wrong with this SLA?

NETWORK DESIGN

Implementing a network project requires a company to go through all phases of the systems development life cycle. In most cases, these stages are similar to those for other IT projects. One special area in the SDLC is the design of a new network or of a modified network.

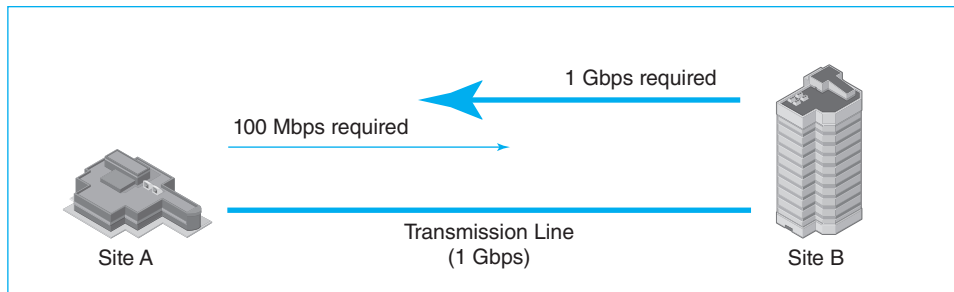


FIGURE 4-6 Two-Site Traffic Analysis

Traffic Analysis

Network design always begins with traffic requirements. **Traffic analysis** asks how much traffic must flow over each of the network's many individual transmission links. Figure 4-6 shows a trivial traffic analysis. A company only has two sites, A and B. A needs to be able to transmit to B at 100 Mbps. B needs to be able to transmit to A at 1 Gbps. Transmission links usually are symmetric, meaning that they have the same speed in both directions. Therefore, the company must install a transmission link that can handle 1 Gbps.

As soon as the number of sites grows beyond two, traffic analysis becomes difficult. Figure 4-7 shows a three-site traffic analysis. For simplicity, we will assume that transmission is symmetric between each pair of sites.

The figure shows that Site Q attaches to Site R, which attaches to Site S. Site Q is west of Site R. Site S is east of Site R. Site Q needs to be able to communicate with Site R at 45 Mbps. Site R needs to be able to communicate with Site S at 2 Gbps. Site Q needs to be able to communicate with Site S at 300 Mbps. There are two links—Link Q-R and Link R-S.

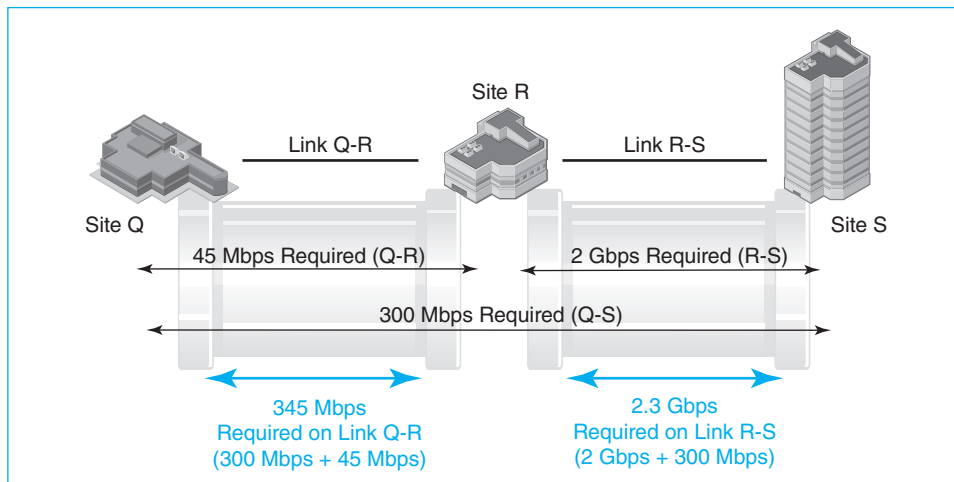


FIGURE 4-7 Three-Site Traffic Analysis

Are you overwhelmed by the last paragraph? Anyone would be! In traffic analysis, it is critical to draw the picture. Figure 4-7 shows how the three sites are laid out and what links connect them. After laying out the sites and links, you draw the three required traffic flows.

Note that the link between Q and R must handle both Q–R traffic (45 Mbps) and the Q–S traffic (300 Mbps). It does not handle any of the traffic between R and S, however. Consequently, Link Q–R must be able to handle 345 Mbps.

Similarly, Link R–S must be able to handle R–S traffic (2 Gbps) and Q–S traffic (300 Mbps). This means that the transmission link between R and S must be able to handle 2.3 Gbps.

If a company has more than two or three sites, doing traffic analysis calculations manually becomes impossible. Companies use simulation programs that try different options for using links to connect its many sites. For each case, traffic analysis is done on each link. However, you need to understand what the program is doing, and the way to do that is to work through a few examples with only a few sites.

Test Your Understanding

7. Do a three-site traffic analysis for the following scenario. Site X attaches to Site Y, which attaches to Site Z. Site X is east of Site Y. Site Z is west of Site Y. Site X needs to be able to communicate with Site Y at 3 Gbps. Site Y needs to be able to communicate with Site Z at 1 Gbps. Site X needs to be able to communicate with Site Z at 700 Mbps. Supply a picture giving the analysis. You may want to do this in Office Visio or a drawing program and then paste it into your homework. a) What traffic capacity will you need on the link between Sites X and Y? (Check Figure: 3.7 Gbps.) b) On the link between Y and Z?

Redundancy

Transmission links sometimes fail. Suppose that the transmission link between R and S in Figure 4-7 failed. Then Q would still be able to communicate with R, but Q and R would not be able to communicate with S. Obviously, this is undesirable.

The solution is to install redundant transmission capacity. **Redundant transmission capacity** is extra transmission capacity on links that is normally not needed but will be needed if another link fails. To illustrate this, Figure 4-8 again shows Sites Q, R, and S. This time, there is a direct link between Q and S. Now, each site can talk to each other site directly.

What happens if the link between Q and R fails? The answer is that Site Q can still talk to Site S through the direct link. In addition, Q can still talk to R by sending its transmissions to S, which will send them on to R.

However, this will only be possible if the remaining links have the redundant capacity to handle the rerouted traffic as well as its normal traffic. For instance, if the link between Q and S is only 300 Mbps, this will be enough if there are no failures. However, if Link Q–R fails, the link will need another 45 Mbps. So it will need to have 345 Mbps of capacity to handle a Link Q–R failure. Link R–S will also need 45 Mbps more capacity. It will need 2.045 Gbps of capacity to handle both R–S traffic and Q–R traffic.

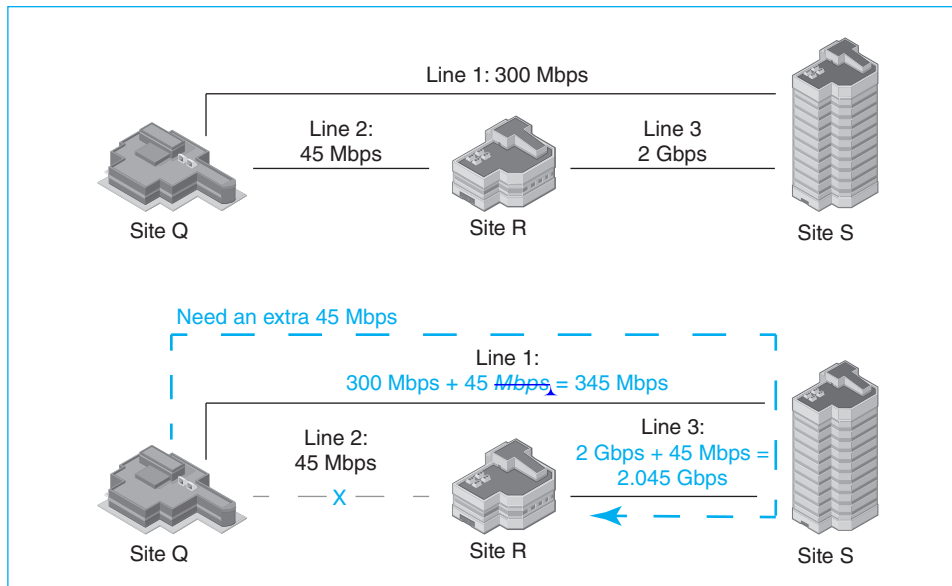


FIGURE 4-8 Three-Site Traffic Analysis with Redundancy

Test Your Understanding

8. a) What is the purpose of redundancy in transmission links? b) If the link between R and S fails in Figure 4-8, how much capacity will the other links need? (Draw the picture.) (Check Figure: Q-R will need to be able to carry 2.045 Gbps.) c) If the link between Q and S fails, how much capacity will the other links need? (Draw the picture.) d) What if both links in the previous two question parts fail? (Draw the picture.)

Momentary Traffic Peaks

Traffic volume varies constantly. Some of this is systematic. In addition, network traffic has a strong random component, and when there is randomness, there will *always* be occasional traffic spikes. These **momentary traffic peaks** typically last only a fraction of a second or a second or two, but they can be disruptive. As Figure 4-9 shows, we are concerned with momentary traffic peaks that exceed the network's transmission link capacities.

Switches and routers have small *buffers* that can hold frames or packets they cannot transmit because of the momentary congestion. They will have to wait to forward frames or packets. This produces latency. Even when buffers are present, they are limited in size. When the buffer size is exceeded, frames or packets will be lost. Applications that use TCP at the transport layer will retransmit lost segments, but retransmission will increase the traffic volume, adding to the overload.

Overprovisioning Figure 4-9 shows three techniques for addressing momentary traffic peaks. The first is overprovisioning, which simply means installing so much more capacity than you normally need that momentary traffic peaks will be so rare and




FIGURE 4-9 Addressing Momentary Traffic Peaks

so brief that they can be ignored. The advantage of overprovisioning is that it places no additional labor burden on the network staff. The disadvantage is that overprovisioning is expensive in terms of transmission links costs. Overprovisioning may make sense in LANs where additional capacity is rather inexpensive. In wide area networks, however, the high cost of transmission capacity makes this impractical.

Priority A second approach is to assign a **priority level** to frames or packets, based on their tolerance for latency and loss. Voice over IP is extremely latency intolerant. Delays in transmission make turn taking in conversations very difficult. When you hear silence, you begin talking, but as soon as you do, you realize that the other person has been talking. In addition, lost frames or packets create silences that VoIP systems must fill with artificial sound. On the other hand, e-mail can easily tolerate

a delay of several seconds. Consequently, VoIP frames and packets get high priority, so that they will get through immediately. E-mail would get low priority because a delay of a few seconds is not a problem in e-mail. All switches and routers from corporations come with the ability to use priority, so priority does not increase capital expense. Priority will bring lower transmission link costs than overprovisioning, but it requires more labor in assigning priority to different traffic flows and configuring devices.

Quality of Service Guarantees A more extreme approach is to give **QoS guarantees** to certain traffic flows such as VoIP. To provide QoS guarantees, the company must allocate **reserved capacity** on each switch and transmission line. This is great for traffic flows with QoS guarantees. However, it means that all other traffic only gets what is left over, even if the reserved capacity is not being used.

Test Your Understanding

9. a) What are momentary traffic peaks? b) How long do they last? c) What two problems do they create? d) What choices do you have for reducing the impact of delays for latency intolerant traffic? e) What is the advantage of each compared to the others? f) What is the disadvantage of each compared to the other? g) Compared to e-mail and voice over IP, what priority would you give to network control messages sent to switches and routers? (The answer is not in the text.) h) Which of the three options would work if you have chronic (frequent) traffic loads that exceed your network's capacity? (The answer is not in the text.)

STRATEGIC SECURITY PLANNING PRINCIPLES

Security Is a Management Issue

People tend to think of security as a technology issue, but security professionals know that security is primarily a management issue. Unless a firm does excellent planning, implementation, and day-to-day execution, the best security technology will be wasted. As noted security expert Bruce Schneier has often said, "Security is a process, not a product."⁹ Unless firms have good security processes in place, the most technologically advanced security products will do little good.

Security is primarily a management issue, not a technology issue.

One thing that sets security management apart from other aspects of network management and IT management in general is that the security team must battle against *intelligent adversaries*, not simply against human mistakes and technical unreliability. Companies today are engaged in an escalating arms race with attackers, and security threats and defenses are evolving at a frightening rate.

⁹ Bruce Schneier, *Crypto-Gram Newsletter*, May 15, 2000. <http://www.schneier.com/crypto-gram-0005.html>.

Test Your Understanding

10. a) Why is security primarily a management issue, not a technology issue? b) What sets security management apart from other aspects of network management and IT management in general?

The Plan–Protect–Respond Cycle

Figure 4-10 shows the overall process that companies follow to deal with threats. On the left is the threat environment, which consists of the attackers and attacks the company faces. We looked at the threat environment in Chapter 3.

The rest of the figure illustrates how companies mount their defenses against the threats they face. The figure shows that companies constantly cycle through three phases of security management. This is the **plan–protect–respond cycle**.

Planning In the **plan phase**, companies assess the threat environment and decide how they will meet these threats. In strategic network management, we talked about the need to close network performance gaps by creating a project portfolio that creates the maximum benefits for the company's limited budget. Companies must do the same for security. In our discussion of the planning stage, we will focus on core principles that companies adopt to make their planning effective.

Protecting In the **protect phase**, companies provide actual protections on a day-to-day basis. We looked at protections such as firewalls in Chapter 3. In Figure 4-10, the protect phase bubble is larger than the other three. This emphasizes the fact that the protect phase is much larger than the other two phases in terms of time

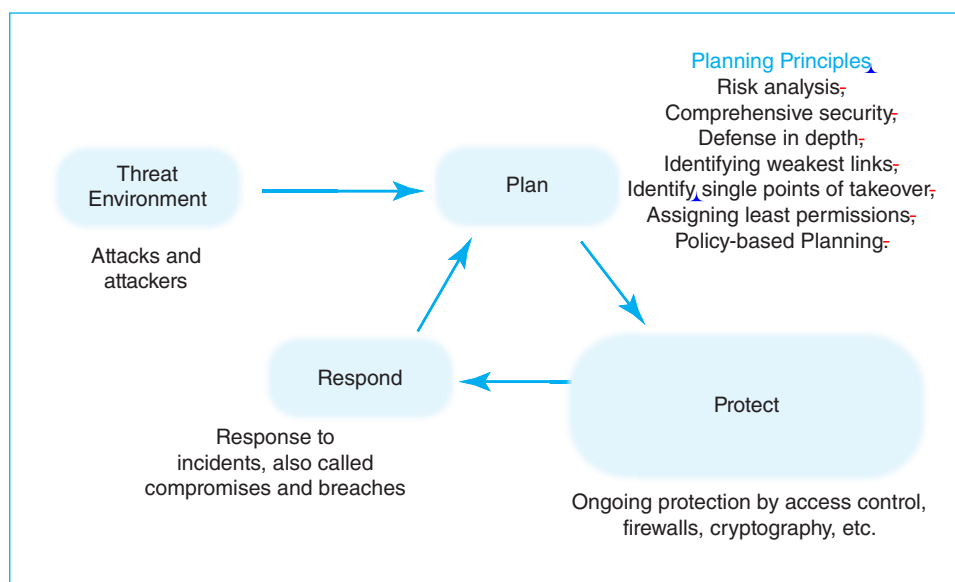


FIGURE 4-10 The Threat Environment and the Plan–Protect–Respond Cycle

spent and resource expenditure. However, without extensive and insightful planning, it is possible to spend a great deal of time and effort mounting protections without being very effective.

Responding In the **response phase**, the company must respond when it suffers a successful security attack. We call successful attacks **compromises, incidents, or breaches**. It would be nice if compromises never occurred. In fact, they will. Like fire departments, security teams must respond immediately and effectively. This requires careful planning and rehearsal because every second counts in reducing the cost of breaches.

Test Your Understanding

11. a) What happens in each stage of the Plan–Protect–Respond cycle? b) Which stage consumes the most time?

Security Planning Principles

Perhaps more than any other aspect of IT, effective security depends on effective planning. Security planning is a complex process that we can discuss only briefly. We will focus on some key planning principles that must be observed in all security thinking.

Risk Analysis Many would say that the goal of security is to stop all threats to the corporation. Surprisingly, that is not true. Stopping all attacks is impossible. Despite strong security efforts, there will still be some risk of a compromise. There has always been crime in society, and there always will be. The same is true of security incidents. No matter how much money a company spends on security, it cannot stop all threats. It could go bankrupt trying. Rather, the goal of security is to reduce the risk of attacks to the extent that is **economically feasible**.

The goal of security is to reduce the risk of attacks to the extent that is economically feasible.

Security Is a Management Issue, Not a Technology Issue

Without good management, technology cannot be effective

A company must have good security processes

Security Planning Principles

Risk analysis

Comprehensive security

Defense in depth

Weakest link analysis

Single points of takeover

Least permissions in access control

FIGURE 4-11 Security Planning Principles


Risk analysis is the process of balancing risks and protection costs. Corporate security planners have to ask whether investing in a countermeasure against a particular threat is economically justified. For example, if the probable annual loss from a threat is \$10,000 but the security measures needed to thwart the threat will cost \$200,000 per year, the firm obviously should *not* spend the money. Instead, it should accept the probable loss.

Risk analysis is the process of balancing risks and protection costs.

Figure 4-12 gives an example of a risk analysis for a hypothetical situation. Without a countermeasure, the damage per successful attack is expected to be \$1,000,000, and the annual probability of a successful attack is 20%. Therefore, the annual probable damage is \$200,000 without a countermeasure. The probable net annual outlay therefore is \$200,000 if no action is taken.

Countermeasure A is designed to cut the damage of a successful attack in half. So the damage per successful attack is expected to be \$500,000 instead of a million dollars. The countermeasure will not reduce the probability of a successful attack, so that continues to be 20%. With Countermeasure A, then, the annual probable damage will be \$100,000. However, the countermeasure is not free. It will cost \$20,000 per year. Therefore, the net annual probable outlay is \$120,000 with the countermeasure.

Countermeasure A, then, will reduce the net annual probable outlay from \$200,000 to \$120,000. The countermeasure has a value of \$80,000 per year. This is positive, so Countermeasure A is justified.

There is also a second candidate countermeasure, Countermeasure B. This countermeasure will reduce the probability of a successful attack by 25%, from 20% to 15%. The loss would not be reduced at all. This countermeasure would cost \$60,000 annually, giving a net annual probable outlay of \$210,000. This exceeds the no-countermeasure's figure of \$200,000. The annual probable outlay is negative \$10,000 if the countermeasure is used. This countermeasure would not make sense even if  was the only candidate countermeasure.

Security professionals may be tempted to think of costs in terms of hardware and software spending. However, most countermeasures require extensive security labor. In fact, labor is often the biggest cost. More broadly, security often increases labor costs

Countermeasure	None	A	B
Damage per successful attack	\$1,000,000	\$500,000	\$1,000,000
Annual probability of a successful attack	20%	20%	15%
Annual probable damage	\$200,000	\$100,000	\$150,000
Annual cost of countermeasure	\$0	\$20,000	\$60,000
Net annual probable outlay	\$200,000	\$120,000	\$210,000
Annual value of countermeasure	\$0	\$80,000	(\$10,000)

FIGURE 4-12 Risk Analysis Calculation

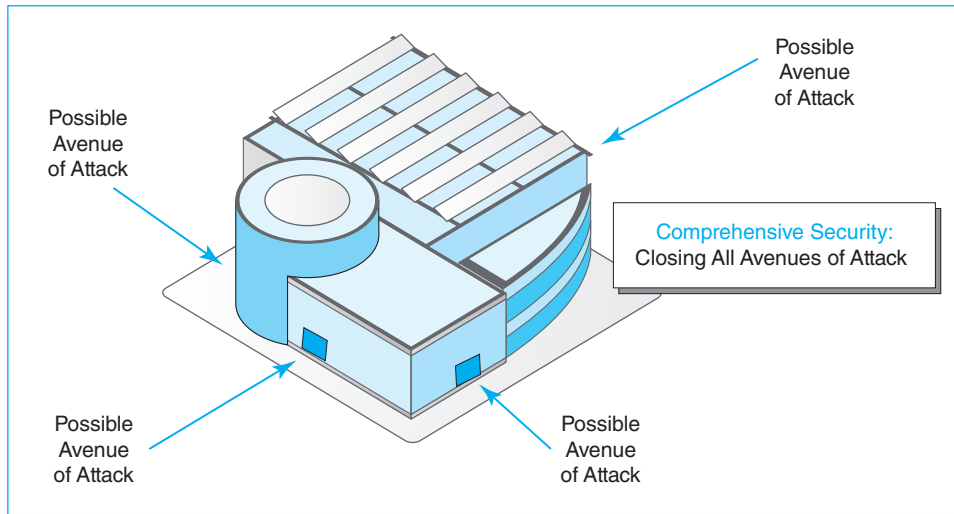


FIGURE 4-13 Comprehensive Security

for users. If users spend even a few extra minutes each time they must use a particular resource, this can lead to substantial cost. It could tip the scales against installing the countermeasure.

Comprehensive Security To be safe from attack, a company must close off *all* avenues of attack. Figure 4-13 illustrates this principle. In contrast, an attacker only needs to find one unprotected avenue to succeed. Although it is difficult to achieve **comprehensive security**, it is essential to come as close as possible.

Comprehensive security is closing off all avenues of attack.

Defense in Depth Another critical planning principle is defense in depth. Every protection will break down occasionally. If attackers have to break through only one line of defense, they will succeed during these vulnerable periods. However, if an attacker has to break through two, three, or more lines of defense, the breakdown of a single defense technology will not be enough to allow the attacker to succeed. Having successive lines of defense that must *all* be breached for an attacker to succeed is called **defense in depth**. Figure 4-14 illustrates the principle.

Having several lines of defense that must all be breached for an attacker to succeed is called defense in depth.

In the figure, there are four protections in **series**. The first is a border firewall at the connection between the company site and the Internet. The second is a host firewall on a particular server. The third is the use of good practice in patching application

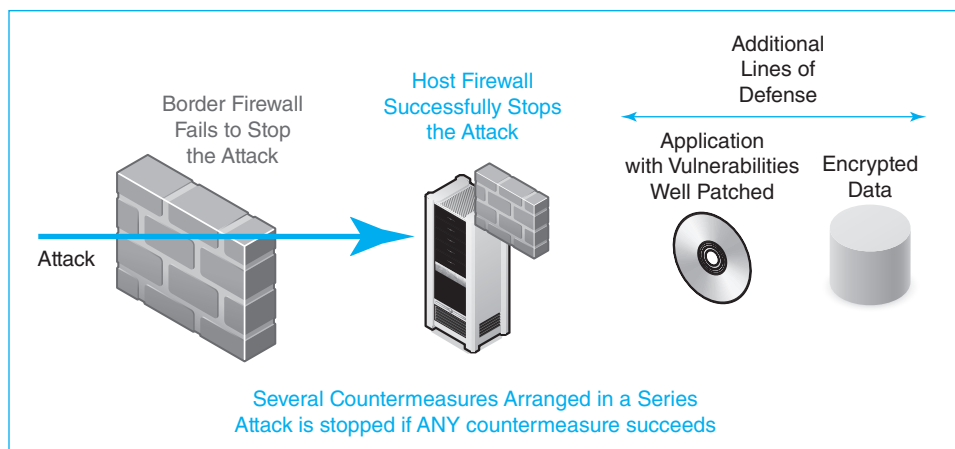


FIGURE 4-14 Defense in Depth

vulnerabilities. The fourth is encrypting all data for confidentiality so that the attacker cannot read sensitive information even **all if** other defenses fail.

The figure shows what happens if the border firewall does not stop an attack. In this case, the host firewall catches the attack and stops it. The company should fix the border firewall quickly, so that it becomes part of the effective defense, but attack packets will not get through to the target data while the border firewall is being fixed.

Identifying Weakest Links Defense in depth is a way to increase security by having a series of protections so a single failure will not compromise security. In contrast, many individual protections consist of a series of internal steps that must *all* work if the protection is to succeed. If one fails, the countermeasure fails. For example, an antivirus program may protect a user by identifying a malicious attachment. However, if the user fails to use good judgment and opens the attachment anyway, there is no protection.

Figure 4-15 shows how weakest links can compromise a countermeasure. Here the countermeasure is a firewall. The firewall has five components, all of which must be effective for the firewall to be effective. These are the firewall hardware, firewall software, a firewall access control list, the firewall log file, and the practice of reading the log file frequently. In the figure, the ACL is defective. Even if all the other elements are fully effective, the firewall will fail to stop an attack. Similarly, if the company fails to read the firewall log file regularly, it will fail to keep the ACL up to date, and this will cause the firewall to fail.

It is easy to confuse defense in depth and weakest link analysis because a series of elements is present in both.

- Typically, defense in depth involves a series of different countermeasures, while weakest link analysis involves a single countermeasure with multiple components.
- In defense in depth, ANY element must be effective to stop an attack. In weakest link analysis, ALL elements must be effective to stop an attack.

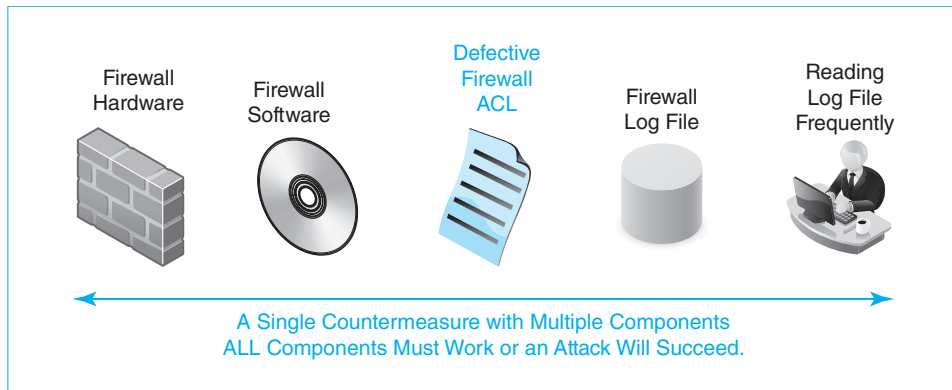


FIGURE 4-15 Weakest Link Analysis

Potential Single Points of Takeover Another principle is to focus on **potential single points of takeover**. Later in this chapter, we will see that companies often control many individual firewalls through a single firewall policy server (see Figure 4-16). If an attacker takes over the firewall policy server, there is no end to the damage that he or she can do. The central firewall policy server is a **potential single point of takeover**, which means that if an attacker can take it over, they gain control over a significant portion of your network.

Companies usually cannot and do not want to eliminate potential single points of failure. Having a central firewall policy server greatly improves a company's control over its firewalls, eliminating inconsistencies and reducing management costs. Eliminating this single point of failure by going back to configuring firewalls individually is not an answer to the threat. Rather, it is critical for companies to identify all single points of takeover and harden them very well against attacks.

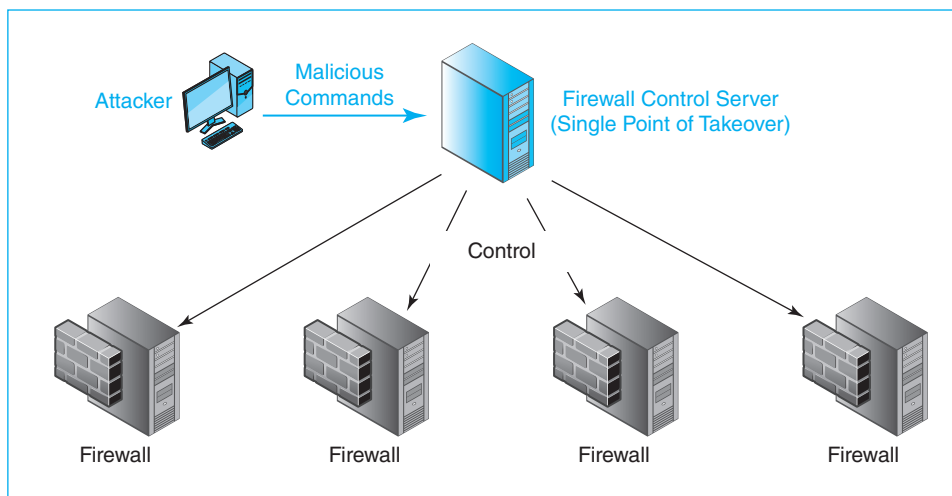


FIGURE 4-16 Potential Single Point of Takeover

Assigning Least Permissions in Access Control Security planners constantly worry about access to resources. People who get access to resources can do damage to those resources. Not surprisingly, companies work very hard to control access to their resources. **Access control** is limiting who may have access to each resource and limiting his or her permissions when using the resource.

Access control is limiting who may have access to each resource and limiting his or her permissions when using the resource.

One aspect of access control that we saw in the previous chapter is authentication, which is requiring users requesting access to prove their identities. However, just because you know who someone is does not mean that he or she should have unfettered access to your resources. (There undoubtedly are several people you know whom you would not let drive your car.)

Authorizations or **Permissions** are the actions that an authenticated user is allowed to take on the resource. For example, although everyone is permitted to view the U.S. Declaration of Independence, no one is allowed to add his or her own signature at the bottom.

Authorizations or Permissions are the actions that an authenticated user is allowed to take on the resource.

An important principle in assigning permissions is to give each person **least permissions**—the minimum permissions that the user needs to accomplish his or her job. In the case of access to team documents, for example, most team members may be given read-only access, in which the user can read team documents but not change them. It is far less work to give the user extensive or full permissions so that he or she does not have to be given additional permissions later. However, it is a terrible security practice. If even one unnecessary permission is assigned to a person, this may be a security risk.

Least permissions are the minimum permissions that the user needs to accomplish his or her job.

Figure 4-17 shows some examples of limited permissions for particular resources. These resources include files, folders, servers, and network elements. To know what resources ~~must be specified and limited~~, you must understand ~~the resource thoroughly~~.

Test Your Understanding

12. a) Comment on the statement, “The goal of security is to eliminate risk.” b) What is risk analysis? c) Repeat the risk analysis in Figure 4-12, this time with Countermeasure C reducing damage severity by a quarter and the likelihood of an attack by 75%. The annual cost of Countermeasure C is \$175,000. Show the full table. What do you conclude? Justify your answer.

Access Control

If attackers cannot get access to a resource, they cannot exploit it

Access control is limiting who may have access to each resource

And limiting his or her permissions when using the resource

Authentication versus Authorizations (Permissions)

Authentication: Proof of identity

Authorizations: Permissions a particular authorized user is given with a resource

Just because a user is authenticated does not mean that he or she will be permitted to do everything

Principle of Least Permissions

Give each authenticated user only the minimum permissions he or she needs to do his or her job

Cannot do unauthorized things that will compromise security

Examples of Limited Permissions

Create files but not delete files

Cannot access files above a specified level of sensitivity

Read files but not write (edit) them

See files in own folders but not all folders

Connect to the person's department server but not to the Finance server

Do certain things but cannot give others permission to do them

FIGURE 4-17 Least Permissions in Access Control

13. a) What is comprehensive security? b) Why is it important? c) What is defense in depth? d) Why is defense in depth necessary? e) Distinguish between defense in depth and weakest link analysis. f) What must companies do about potential single points of takeover? g) Distinguish between authentication and authorizations. h) What is another term for authorizations? i) What is the principle of least permissions? j) Why is it important?

Policy-Based Security

We have discussed the importance of security planning and major security principles. Now we will look at how plans are implemented in well-run organizations.

Policies The heart of security management is the creation and implementation of security policies. Figure 4-18 illustrates how policies should be used. **Policies** are broad statements that specify *what should be accomplished*. For example, a policy might be, “All information on USB RAM sticks must be encrypted.” Policymakers have the overview knowledge that operational people do not have. For instance, policymakers may know that new laws create serious liabilities unless USB RAM sticks are encrypted. Operation-level people may not realize this.

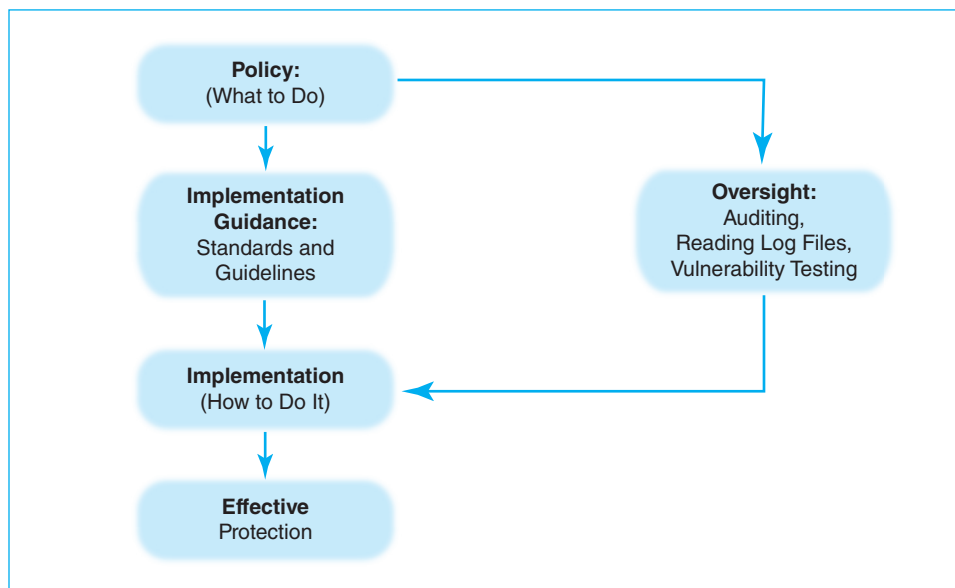


FIGURE 4-18 Policy-Based Security

Policy versus Implementation Note that the policy does not specify which encryption technology should be used or other implementation details. Put another way, policies describe *what* should be done, not *how* to do it.

This separation of policy from implementation permits the implementers to carry out the policy in the best way possible for particular situations. Policymakers should have superior overview knowledge. However, implementers know about specific technologies and the local situation. They have the specific knowledge that policymakers do not, including technical knowledge. Unless they are given the latitude to use this knowledge, weak implementation is likely to doom the policy's effectiveness. Separating policies from implementation prevents senior security professionals from micromanaging operating-level people inappropriately. The separation of policies from implementation uses the strengths of both policy makers and implementers.

Policymakers have the overview knowledge that operational people do not have. Implementers know about specific technologies and the local situation that policymakers do not. Separating policies from implementation uses the strengths of both.

The separation of policy from implementation certainly does not mean that policy is irrelevant to implementation. It is easy to get lost in implementation details. Having a clear policy permits everybody involved in implementation to stay synchronized by checking whether what he or she is doing will lead to the successful implementation of the policy.

Implementation Guidance In many cases, the policymaker will only specify the policy. However, in some cases, the policymakers will also create some implementation guidance. **Implementation guidance** consists of instructions that are more specific than policies but less specific than implementation.

Implementation guidance consists of instructions that are more specific than policies but less specific than implementation.

For example, after establishing a policy that USB RAM sticks must be encrypted, implementation guidance might be added in the form of a directive that the encryption must use keys at least 128 bits long. This ensures that implementers will not have the latitude to choose weak encryption that can be defeated by an attacker.

There are two general forms of implementation guidance: standards and guidelines. **Standards** are *mandatory* directives that *must* be followed. Requiring 128-bit encryption is a standard. It is mandatory for implementers to follow the directive.

Standards are mandatory directives that must be followed.

In contrast, **guidelines** are directives that *should* be followed. This gives the implementer not only guidance but also some leeway in deciding whether to follow the guidance. This does not mean that implementers can ignore guidelines. They *must* consider them carefully. However, for good reason, they can elect not to follow them.¹⁰ For example, a guideline that security staff members should have three years of security work experience indicates that someone hiring a security staff member must consider that having at least three years of experience is a reasonable expectation. If the person doing the hiring selects someone with only two years of work experience, he or she should have a very good reason for doing so. Following guidelines is optional, but seriously considering guidelines is mandatory.

Guidelines are implementation guidance directives that should be followed but that need not be followed, depending on the context.

When do firms use guidelines instead of standards for implementation guidance? The answer is that they use guidelines for situations that are not amenable to black-and-white rules. Encryption strength is relatively easy to specify. The quality of work experience requires human judgment.

¹⁰ In the *Pirates of the Caribbean* movies, there was a running joke that the Pirate Code is “more of a guideline, really.”

Oversight Figure 4-18 also shows that policymakers cannot merely toss policies and implementation guidance out and ignore how implementation is done. It is essential for management to exercise **oversight**, which refers to a collection of methods for ensuring that policies have been implemented appropriately in a particular implementation.

Oversight is a collection of methods for ensuring that policies have been implemented appropriately in a particular implementation.

One form of oversight is an audit. An **audit** samples actions taken within the firm to ensure that policies are being implemented properly. Note that an audit only *samples* actions. It does not look at everything, which would be impossible to do. However, if the sampling is done well, the auditor can issue an opinion on whether a policy is being carried out appropriately based on well-considered data.

An audit samples actions taken within the firm to ensure that policies are being implemented properly.

Another form of oversight is reading **log files**. Whenever users take actions, their actions should be recorded in log files. Reading log files can reveal improper behavior. Of course, if these log files are not read, they are useless. Log files should be read daily or even several times each day. Few people enjoy reading log files to look for problems, so enforcement must be carefully tracked.

Reading log files can reveal improper behavior.

Yet another important oversight mechanism is vulnerability testing. Simply put, **vulnerability testing** is attacking your own systems before attackers do, so that you can identify weaknesses and fix them before they are exploited by attackers. It is important to set up vulnerability tests cautiously, however. Before doing a vulnerability test, the tester must have explicit written permissions for each test based on a detailed description of what will be done and what damage might be done accidentally. Vulnerability testers who do not take these precautions have been accused of making malicious attacks. This has resulted in firings and even jail terms.

Vulnerability testing is attacking your own systems before attackers do, so that you can identify weaknesses and fix them before they are exploited by attackers.

Note that the policy drives both implementation and oversight. Implementers who attempt to implement the policy must interpret the policy. Auditors and other

oversight professionals must also interpret the policy. If the implementers are lax, the auditors should be able to identify this. However, if oversight practitioners and implementers disagree, this may simply mean that they are interpreting the policy differently. Policymakers may find that one or the other has made a poor choice in interpreting the policy. They may also find that the policy itself is ambiguous or simply wrong. The important thing is to identify problems and then resolve them.

Policies drive both implementation and oversight.

Effective Protection Policies certainly do not give protection by themselves. Neither may unexamined implementations. Protection is most likely to be effective when excellent implementation is subject to strong oversight.

Test Your Understanding

14. a) What is a policy? b) Distinguish between policy and implementation. c) What is the benefit of separating policies from implementation? d) Why is oversight important? e) Compare the specificity of policies, implementation guidance, and implementation. f) Distinguish between standards and guidelines. g) Which **must** be followed? h) Must guidelines be considered? i) List the three types of oversight listed in the text. j) What is vulnerability testing, and why is it done? k) Why is it important for policy to drive both implementation and oversight?

CENTRALIZED NETWORK MANAGEMENT

Given the complexity of networks, network managers need to turn to network management software to support much of their work. Many of these are network visibility tools, which help managers comprehend what is going on in their networks.

Ping

The oldest network visibility tool is the basic ping command available in all operating systems. If a network is having problems, a network administrator can simply **ping** a wide range of IP addresses in the company. When a host receives a ping, it should send back a reply. If it replies, it is obviously in operation. If it does not, it may not be. In Figure 4-19, host 10.1.2.5 does not respond to its ping. This signals a potential problem.

By analyzing which hosts and routers respond or do not respond, then drawing the unreachable devices on a map, the administrator is likely to be able to see a pattern that indicates the root cause of the problem. Of course, manually pinging a wide range of IP addresses could take a prohibitive amount of time. Fortunately, there are many programs that ping a range of IP addresses and portray the results.

Test Your Understanding

15. If you ping a host and it does not respond, what can you conclude?

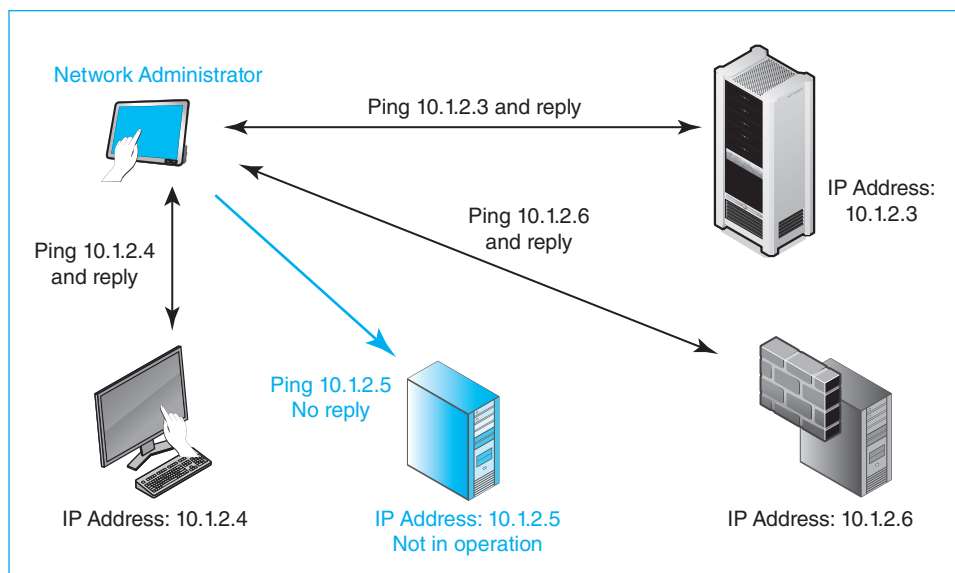


FIGURE 4-19 Ping

The Simple Network Management Protocol (SNMP)

Ping can tell you if a host is available. It can also tell you the latency in reaching that host. For remote device management, most network operation centers use more powerful network visualization products based on the **simple network management protocol (SNMP)**, which is illustrated in Figure 4-20. In the network operations center (NOC), there is a computer that runs a program called the **manager**. This manager manages a large number of **managed devices**, such as switches, routers, servers, and PCs.

Agents Actually, the manager does not talk directly with the managed devices. Rather, each managed device has an **agent**, which is hardware, software, or both. The manager talks to the agent, which in response talks to the managed device. To give an analogy, recording stars have agents who negotiate contracts with studios and performance events. Agents provide a similar service for devices.

Get Commands and the Management Information Base The network operations center constantly collects data from the managed devices using SNMP **Get** commands. It places this data in a **management information base (MIB)**. Data in the MIB allows the NOC managers to understand the traffic flowing through the network. This can include failure points, links that are approaching their capacity, or unusual traffic patterns that may indicate attacks on the network.

Set In addition, the manager can send **Set** commands to the switches and other devices within the network. Set commands can reroute traffic around failed equipment or transmission links, reroute traffic around points of congestion, or turn off expensive transmission links during periods when less expensive links can carry the traffic adequately.

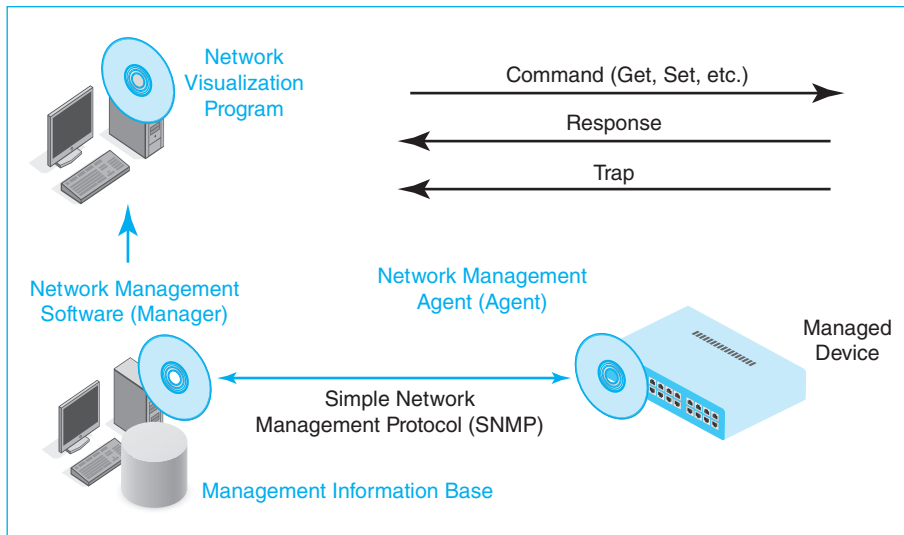


FIGURE 4-20 Simple Network Management Protocol (SNMP)

Trap Normally, the manager sends a command and the agent responds. However, if the agent senses a problem, it can send a **Trap** command on its own initiative. The trap command gives details of the problem.

Network Visualization Program There is one more program in the figure—a **network visualization program**. This program takes results from the MIB and interprets the data to display results in maps, find root causes for problems, and do other tasks. Note that this functionality is *not* included in the simple network management protocol. SNMP simply collects the data in a way that network visualization programs can use. This lack of specification allows network visualization program vendors to innovate without being constrained by standards. ~~What do network visualization programs do?~~

Automation Many other network management chores can be automated to reduce the amount of work that network managers need to spend on minutia. For example, many routers are given a standard corporate configuration when they are installed. Doing this manually can take an hour or more per router. However, it may be possible to create a standard configuration, store it, and simply download it onto new routers. In addition, if corporate standard configurations change or a patch must be installed on all routers, it may be possible simply to “push out” these changes to all routers.

Test Your Understanding

16. a) List the main elements in SNMP. b) Does the manager communicate directly with the managed device? Explain. c) Distinguish between Get and Set commands. d) Where does the manager store the information it receives from Get commands? e) What kinds of messages can agents initiate? f) Why is network automation important?

Software-Defined Networking (SDN)

In Chapter 10, we will look at a radical and extremely promising approach to managing networks. It is called **software-defined networking (SDN)**. Figure 4-21 shows how switches, routers, and wireless access points operate in traditional networking. Each device has a *forwarding function* and a *control function*. The forwarding function actually forwards individual frames or packets. The control function consists of rules that tell the forwarding function how to forward individual frames or packets.

As we saw in Chapter 1, switches operate individually. They do not know a frame's entire path through the network. In Chapter 8, we will see that routers also operate independently. Routers do exchange information to identify possible routes for packets, but there is no simple way to determine flows across an internet. Individual operation simplifies network management, but it does not allow precise control over flows at the data link layer or the internet layer.

One problem with independent operation is that each device's control function has to be configured individually. As the number of devices grows, manual configuration cost and complexity grow rapidly. In networking terms, manual configuration does not scale to very large numbers of devices. There are certainly ways to reduce this problem. For instance, standard configurations for routers can be stored and downloaded rapidly to new routers. However, the need for some manual configuration still remains on a per-router basis.

Although this approach has worked for many years, its limitations are rapidly growing more serious. These were first noticed in cloud server farms. In a multitenant server farm, it is important to ensure that virtual machines for different customers cannot talk to each other through the site's local area network. This means that forwarding rules may have to change each time the cloud service provider spawns a new VM. Given the rate of VM spawning in cloud server farms, manual reconfiguration of the server farm network has become extremely difficult, especially when many routers and switches have to be reconfigured after each VM spawn.

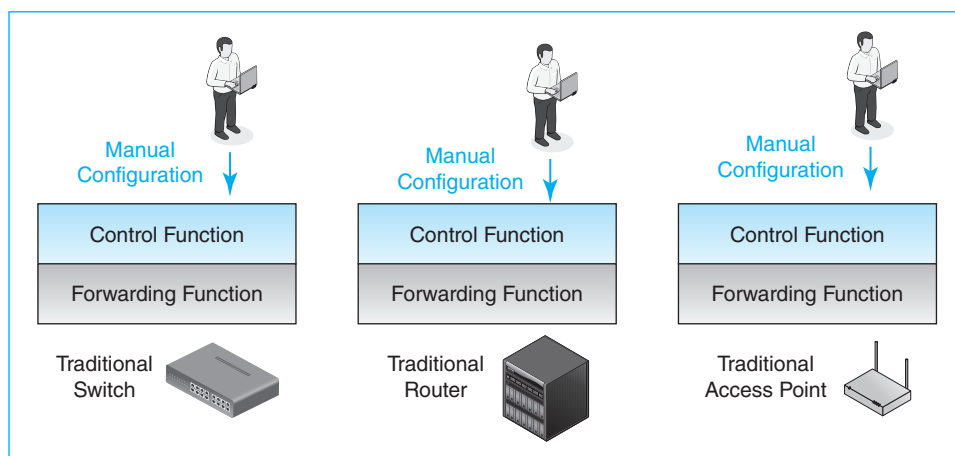


FIGURE 4-21 Traditional Device Control in Networking

In 2008, researchers created a new approach, which they called software-defined networking (SDN). As Figure 4-22 shows, the forwarding function of each switch or router operates as it always has, to forward frames or packets. However, the control function does not stay the same. Instead, the control function accepts commands from an **SDN controller**. Formally, **software-defined networking (SDN)** separates the network forwarding function from the control function and places the control function in an SDN controller.

Software-defined networking (SDN) separates the network device forwarding function from the control function and places the control function in an SDN controller.

SDN controllers permit centralized control of all network devices, including switches, routers, wireless access points, and other possible forwarding devices. The researchers created software-defined networking to study experimental forwarding algorithms in working network environments. In CSP server farms, however, operators soon discovered that this approach was exactly what they needed to manage their own networks as agilely as they managed their VM instances. As we will see in Chapter 10, software-defined networking is beginning to spread not only within data centers but also across local and wide area networks. If SDN does become widespread, it will completely change the way firms manage their networks, and it will do so in a very positive way. Nothing in networking is so potentially disruptive for the field.

Figure 4-22 shows the SDN controller communicating with switches and routers through the OpenFlow Protocol. **OpenFlow** is an open (nonproprietary) protocol for communication between an SDN controllers and SDN-compatible switches, routers, and access points. Strictly speaking, using OpenFlow is not necessary for SDN.

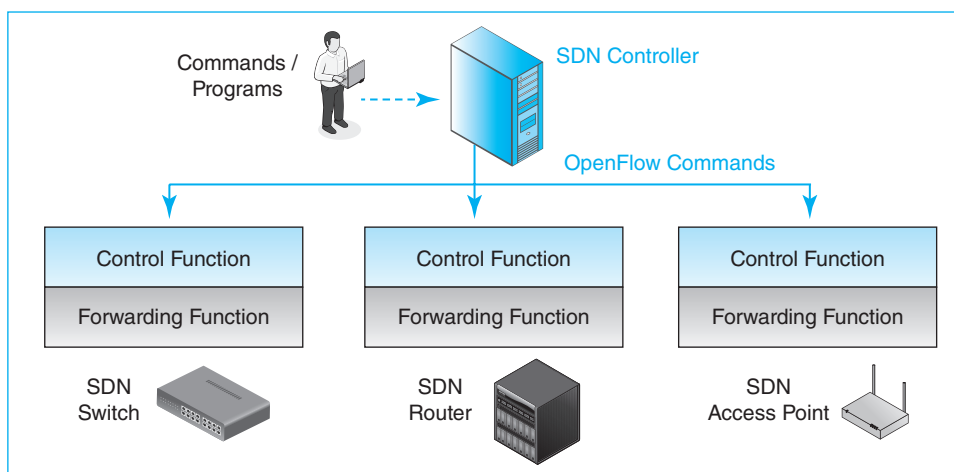


FIGURE 4-22 Software-Defined Networking (SDN) Control in Networking

However, if OpenFlow becomes dominant, today's difficult problem of controlling networks made from the products of multiple vendors will be solved.

OpenFlow is an open (nonproprietary) protocol for communication between an SDN controllers and SDN-compatible switches, routers, and access points.

Test Your Understanding

17. a) What are the two functions in network forwarding devices? b) Traditionally, how is the control function in each device managed? c) What problems does this create? d) In a server farm, what ~~networking control~~ may have to be changed? e) How are device control functions managed in software-defined networking? f) What device manages the control functions on forwarding devices? g) What are benefits of SDN?

CENTRALIZED SECURITY MANAGEMENT

Centralized management is also important in security management. For example, a company may have dozens or even hundreds of firewalls on its network. It would be easy to accidentally misconfigure a few of these firewalls to ignore individual access control rules. Figure 4-23 shows a central firewall management system designed to prevent such oversights. The firewall administrator creates high-level access policies for firewalls to implement. In the figure, the **firewall policy** being implemented is that any IP address not in the accounting department may access any external webserver

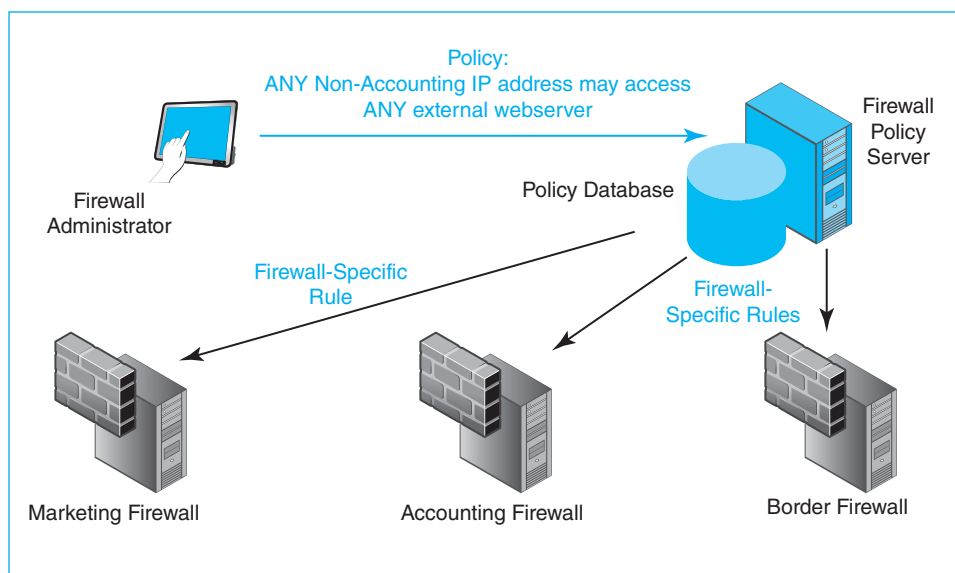


FIGURE 4-23 Centralized Firewall Management

webserver. Many companies allow ~~open~~ access to external webserver ~~for most~~ IP addresses. This security administrator has decided that open access is good but that it should not extend to hosts within the accounting department. The firewall administrator sends this policy to the **firewall policy server**, which places the policy in its policy database.

The firewall policy server then creates detailed **firewall rules** to implement this new policy (such as ACL rules in stateful inspection firewalls). It then pushes these rules out to all of the company's many firewalls. These rules are firewall-specific. The Accounting firewall, for example, may get a different firewall rule to implement this policy than the main border firewall or the Marketing firewall.

Separating firewall policies from firewall rules (which implement firewall policies) is an example of policy-based security. The firewall administrator sets high-level policies. The firewall policy server does the work needed to convert this policy into individual firewall rules. The firewall policy server will not make human mistakes like forgetting to configure a particular firewall. If there is a question about a particular firewall rule on a particular firewall, furthermore, the firewall administrator can ask what policy it implements. Policies are usually easier to understand than specific firewall rules.

Test Your Understanding

18. a) Distinguish between firewall policies and firewall rules. b) When a firewall administrator sends a policy to the policy server, what does the policy server do? c) Which is easier to understand—a firewall policy or a firewall rule?

CONCLUSION

Synopsis

This is the last of four introductory chapters. In this chapter, we looked at network and security management. Technology is never enough. How well people manage the firm's networks and security makes all the difference in the world.

Networks today must work well. Networks must meet goals for quality-of-service (QoS) metrics. We looked at speed, availability, error rates, latency, and jitter. After discussing individual QoS metrics, we looked at service level agreements (SLAs), which guarantee performance levels for certain QoS metrics, usually for a certain percentage of time. Many find it confusing that QoS metrics specify that service will be *no worse* than certain values. For example, SLAs will specify a minimum speed, not a maximum speed.

Designing networks is a complex process. We looked at basic principles of traffic analysis, which identifies the traffic that various transmission links must sustain, including redundancy in case of link failures. Traffic analysis forms the core of network design. We also looked at ways to manage momentary traffic peaks, including overprovisioning, priority, and QoS guarantees with reserved capacity.

Security management follows the plan-protect-respond cycle. Planning prepares the company for day-to-day protection both now and in the future. Response happens when protections break down, as they inevitably do. Of course, experience in managing protections and responses feeds back into the planning process.

Strategic security planning uses the following six planning principles that must be considered in every project plan:

- *Risk analysis.* Many people believe that the purpose of security is to eliminate risk. Unfortunately, all countermeasures cost money. It is always critical to balance potential damage against the cost to prevent it. The purpose of security is to reduce risk to a degree that is economically justified.
- *Comprehensive security.* Attackers only have to find one avenue of attack into a firm or into a specific asset. In contrast, the security staff must discover and close off all avenues of attack. This comprehensive security cannot always be achieved, but firms must constantly look for unprotected avenues of approach.
- *Defense in depth.* Every security protection sometimes breaks down. Defense in depth means creating a *series of protections* that the attacker must break through to succeed. Unless *all* fail, the attacker will not succeed.
- *Avoiding weakest links.* Some *individual protections* have multiple parts that *must all work effectively* if the protection is to work. If one component is quite likely to be defective, then the protection's value will be suspect. For example, if someone is likely to fall for spear phishing and opens an attachment with a virus, the best technical protections will break down.
- *Avoiding potential single points of failure.* The centralization of functions can provide efficiency and control. However, when functions are centralized, this often creates a potential single point of takeover that could lead to serious problems if an attacker were to take control. Potential single points of failure usually cannot be eliminated. Instead, they must be identified and given intense attention.
- *Assigning least permissions.* Even if you authenticate someone, this does not mean that you will let him or her do anything to your resources. It is important to give authenticated users the minimum permissions (authorizations) they need to do their work.

We looked at policy-based security in which a high-level policy group creates security policies and lower-level staff members implement the policy. Policies specify *what is to be done*. Implementation focuses on *how to do it*. This division of labor works because high-level policy people have a broad understanding of security risks and can create policies that will give comprehensive security. Implementation is done by lower-level staff members who know the technology and local situation in far greater detail. They are best suited for deciding how to do implementations. Sometimes, the policy group creates intermediate implementation guidance consisting of standards (which must be followed) and guidelines that must be considered, although they do not have to be followed if there is good reason not to. A separate oversight process ensures that implementations are faithful to appropriate policies.

We closed with a discussion of centralized management for networking and security. By simplifying and automating many actions, centralized management prevents labor costs from increasing as rapidly as networking and security device numbers. We began with ping, which is in the toolbox of every network administrators. Network management depends heavily on the Simple Network Management Protocol (SNMP). We will look at SNMP again in more detail in Chapter 9. In network management, software-defined networking (SDN) may revolutionize the way we manage networks by allowing us to control all network forwarding devices from an SDN server.

For centralized security management, we saw how firewall policy servers accept firewall policies and then send customized firewall rules to individual firewalls to implement the policy.

END-OF-CHAPTER QUESTIONS

Thought Questions

- ~~4-1. a) How can jitter be reduced on a user's PC if there is jitter in incoming packets? b) What does this do to latency?~~
- 4-2. Your home is connected to the Internet. You get to create SLAs that the ISP must follow. Being reasonable, write SLAs you would like to have for the following things: a) Write an SLA for speed. b) Write an SLA for availability. c) Write an SLA for latency. d) Write an SLA for jitter. Do not just say what each SLA should include; actually write the SLAs as the ISP would write them.
- 4-3. A company has offices in Honolulu, Seattle, Ogden, and Dublin, Ireland. There are transmission links between Honolulu and Seattle, Seattle and Ogden, and Ogden and Dublin. Seattle needs to communicate at 1 Gbps with each other site. Seattle and Dublin only need to communicate with each other at 1 Mbps. Ogden and Dublin need to communicate at 2 Gbps, and Ogden and Seattle need to communicate with each other at 10 Mbps. How much traffic will each transmission link have to carry? Show your work. (Check Figure: Honolulu-Seattle needs 1.011 Gbps.)
- 4-4. a) Suppose that an attack would do \$100,000 in damage and has a 15% annual probability of success. Spending \$9,000 per year on "Measure A" would cut the annual probability of success by 75%. Do a risk analysis comparing benefits and costs. Show your work clearly. b) Should the company spend the money? Explain. c) Do another risk analysis if Measure A costs \$20,000 per year. Again, show your work. d) Should the company spend the money? Explain.
- 4-5. An executive opened an e-mail attachment because the content looked like it came from a subordinate. In addition, the executive knew that the company did antivirus filtering. Actually, this was a spear phishing attempt, and the attachment contained malware. What security planning principle does this breakdown represent?
- 4-6. Edward Snowden, a server administrator, was able to copy many CIA secret and top secret files to a USB RAM stick. What security planning principle breakdown allowed this?
- 4-7. Why do you think companies create guidelines for some things instead of creating standards for them?
- 4-8. If oversight practitioners and implementers disagree on whether an implementation is correct, what might be wrong?

Perspective Questions

- 4-9. What was the most surprising thing you learned in this chapter?
- 4-10. What was the most difficult part of this chapter for you?