2

Information Assets

Terms you'll need to understand:

- ✓ Asset
- ✓ Malware
- ✓ AAA
- ✓ NTP
- ✓ SNMP
- ✓ Campus
- ✓ Edge
- ✓ DMZ
- ✓ ISP
- ✓ PSTN
- ✓ WAP

Techniques you'll need to master:

- ✓ Understanding what various devices do on the network
- ✓ Appreciating exposure to threats based on device function or application
- ✓ Recognizing what other devices might be connected to a given device
- ✓ Understanding how those connections can be exploited
- ✓ Appreciating exposure to threats based on location

The CSI Exam tests your knowledge of the SAFE Blueprint, what its architecture can and cannot do for an enterprise of any size, and where you can make compromises or employ design alternatives. You can expect questions from any part of the document called "Extending the Security Blueprint to Small, Midsize, and Remote-User Networks" (sometimes called the SMR Blueprint). However, you will also see exam questions about the details of products employed in the SAFE architecture, and you will need to be able to think through the implications of the SAFE design when questions are not phrased in the way you might expect. That's why we are going to be sure that you have some conceptual foundations in place first. This chapter addresses some ideas that Cisco assumes you already have cemented in place before you read the SMR Blueprint.

What Is the SAFE Blueprint For?

Cisco's SAFE architecture is intended to be a security implementation blueprint for networks in an environment that includes threats outside and valuable components inside—always a dangerous combination. SAFE is really about protecting a network while maintaining its usability. This will always be a balancing act, and different organizations will find different points of balance—SAFE is intended for all nonhome networks, including nonprofit organizations and government agencies, and businesses of all sizes. Many will balance the two about equally, as shown in Figure 2.1.

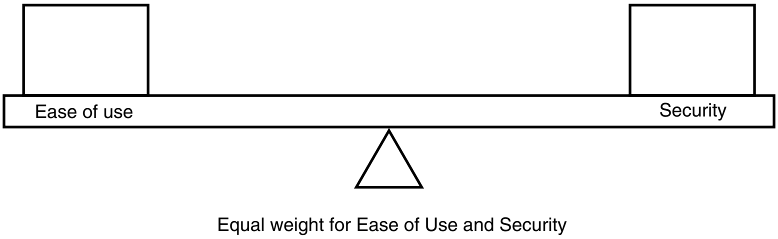

Figure 2.1 Equal weight for ease of use and security.

For some, security is less important than ease of use or simplicity of maintenance. In these situations, the balance point shifts to reflect that, as in Figure 2.2. These organizations implement fewer of the technologies described in the SAFE blueprint, which makes it important to very carefully plan the placement of those technologies.

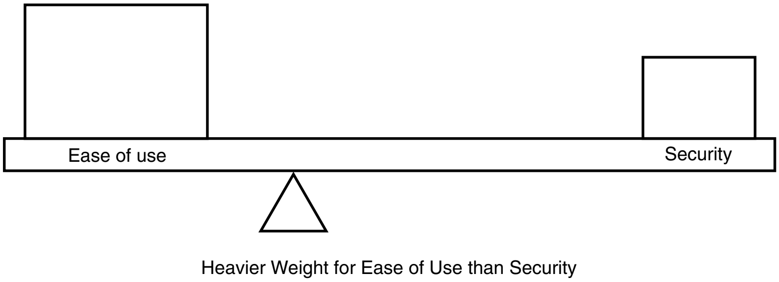

Figure 2.2 Heavier weight for ease of use than security.

For others, the risk of damage from a security breach is far higher, so the balance point shifts to reflect the greater weight that security must bear, as shown in Figure 2.3. These organizations employ more security technologies. Careful thought must go into the monitoring effort because useful information could easily be lost among the "noise" of abundant data.

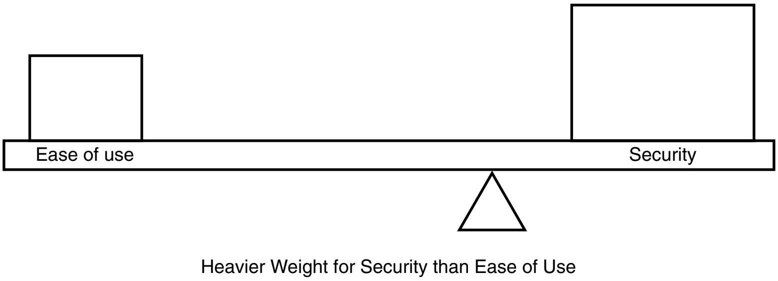

Figure 2.3 Heavier weight for security than ease of use.

All three of these points of balance—and all points along the plane—are equally valid. Which is best for an organization is a decision that already should be made before the SAFE blueprint is applied. That decision should be apparent in the wording of the organization's security policy.

No organization has an unlimited budget for anything, and the budget for security will be less than you would like. That means you need to be able to identify just what your information assets are, what you need to protect them from, and what tools you might have available to do that job—and do it while keeping the network usable by the ordinary user. You wanted a challenge, right?

It is indeed a challenge to protect assets from those who shouldn't have access to them while ensuring that those who should, do. Let's start by looking at how to define the assets that you need to protect.

What Are Information Assets?

Information assets can be physical devices, the data stored within devices, or the data itself as it moves over the network and through its devices. In a broad sense, then, the entire IT system is an information asset, each physical device within it is an information asset, and all the data that resides in or is transported over the system is an information asset.

> This exam requires a different mind-set from that of most Cisco certification exams, which usually only address the networking hardware and software. For the CSI exam, you have to think about the servers and the workstations, as well as how they must be protected—including from each other—when necessary. You also must think about protecting the data when it is on the wire (or in the air, with wireless).

You should be aware that not all information assets are created equal. Some are more valuable to the operation (whether business or other organization) than others. Naturally, that means you must think about applying different degrees of protection to them. Let's take a look at the kind of assets you might find in a network.

Classifying Assets

The method of classifying assets that follows is not one that you will actually see used on the exam or in the SAFE Blueprint. However, Cisco has been making its exams a little more rigorous (or a lot more rigorous, depending on whom you ask), so you can expect to find questions that you didn't see on a practice test or questions that cover something you didn't think a great deal about during your preparation. It will help if you can break things down into their component parts and think through a problem.

For instance, if you face a question concerning routing traffic, you might realize that "routing" refers to Layer 3: The answer must relate to IP or IPX, not Ethernet, Frame Relay, or ATM (if those were your other choices). In the same way, you need to be able to think about a network and its major constituent parts for the CSI exam. The breakout that follows is good practice for thinking outside just the networking components themselves. Cisco groups all elements of the network—networking devices, servers, workstations, IP phones, and so on—into modules generally described as either *edge* or *campus*. We will make that grouping clear shortly.

Hardware Assets

The short description of an asset is "any physical device connected to the network, directly or indirectly." That doesn't help you prepare for the exam very much, so let's start by breaking things out a little. Hardware assets can be end-user devices, devices intended to support many users (we can call them user-support devices), and networking devices.

End-User Devices

These are the devices that your users actually operate: desktops, laptops, PDAs, and so on. These devices might reside within your network fairly permanently (desktops don't wander around often), or they might be exposed to the outside world some of the time. Depending on the users and how much they actually work inside your offices, some end-user devices might spend most of their time outside your network; when they do access information assets on the inside, they make that connection from somewhere outside.

NOTE

Because of the exposure of end-user devices to networks outside your control (including their exposure to the Internet), you will have to pay extra attention to them when you get to the remote-user model in Chapter 12, "The Remote-User Design." Remember, even if you create secure access into your network for these devices, they sometimes access other networks and can be exposed to malware (malicious software) that they then can ferry into your network.

Other end-user devices might not be quite so obvious. IP phones also are end-user devices that have (or can access) information over the network. Bear in mind, too, that some end-user devices might not be your assets (and you might not know they are there). For example, an individual might supply his own PDA, which is synchronized with his desktop calendar, or he might have a copy of information, such as a spreadsheet.

In addition, many networks handling industrial processes or physical production or product distribution include appliances such as data-collection devices (bar-code scanners that update inventory data) and automated control devices that can open or close valves and relays. Also in this group are industrial robots, which are the end users of their control networks.

A problem also could enter your network when a mobile end-user device, such as a PDA, is reconnected after being used elsewhere. If it does, you likely will learn this after the fact; the problem will appear without warning. The SAFE architecture includes multiple layers of protection largely because, despite your best efforts to secure the network, some risks will remain. The multilayer approach minimizes the damage that those remaining risks can cause.

User-Support Devices

User-support devices are devices that individuals do not operate directly; instead, the users access such devices from time to time as they work. Many users might simultaneously access a single user-support device. Alternatively, a device might remain idle for some time because no user is accessing it.

Obvious choices for user-support devices are file and print servers, along with the printers controlled by the print servers. But other user-support devices should be considered. Some of these perform network-support services, without which users couldn't get much done. Examples of these devices are DNS and DHCP servers, proxy servers, Web and mail servers, and (for those IP phones) call managers.

Speaking of management, some servers are on a support network managing the network; these include Authentication, Authorization, and Accounting (AAA) servers; Network Time Protocol (NTP) servers; Simple Network Management Protocol (SNMP) servers; and log servers. These servers are not necessary for data to flow in the same way that routers and switches are; they support managing the network without actually being a part of the data flow. This might seem like a fine distinction, but you will see in the SAFE models that some devices have interfaces that connect to more than one module. Cisco places these devices in the module where they serve a primary function, and that is the kind of distinction we are making here: These servers support networking but are not networking devices (they support the networking devices, their users, in the same way that a printer supports a human user).

Unlike some of the end-user devices, user-support devices should always be under your complete control. However, because these user-support devices interact with many other hosts on the network, you often will be protecting them more than the end-user devices.

NOTE

This idea might seem a little backward, but much of the SAFE architecture is intended to help you contain a problem that gets into your network, however it managed to do that. (You'll hear that idea repeated a few more times: It's important.) Think about the interactions a given device *should* have and those that it *could* have, based on how the device connects in the network. You should understand why user-support devices get more extensive security attention than end-user devices.

Networking Devices

Now we're finally talking about the devices you're probably most familiar with, at least as you prepare for this exam. These are the devices that you learned about to pass your CCNA exam (routers and switches) and those that you studied in depth for the more advanced exams, such as CSPFA (PIX firewalls), CSIDS (intrusion detection), CSVPN (VPN concentrators and clients), and the Network Access Server (NAS) and AAA that you worried about for the SECUR exam.

For those exams, you focused on how to configure those devices. For this exam, you need to think more about how these devices connect with each other, what happens when they connect with the outside world (where malware seemingly lurks in every connectivity cloud), and how they can be configured to do the following:

1. Protect themselves
2. Protect each other
3. Protect the rest of the network devices (end-user and user-support devices)

If that order seems a bit odd, remember that if the devices don't protect themselves and then each other, they can't protect the rest of the network.

> This is probably the biggest difference between the CSI exam and other Cisco certification exams, except possibly the CCIE. In this exam, you aren't demonstrating how to make the networking devices exchange more information as much as you are demonstrating how to be sure that they exchange only what's necessary and block anything else. Other exams want you to *facilitate* information exchange; this exam wants you to *control* it.

Software Assets

Of course, hardware without software is pretty much useless. But is there really any point to considering the software assets separately? Yes—but, of course, I would say that.

A group calendar is a kind of software that, if compromised, might reveal something as trivial as vacation schedules or something as important as project schedules and timelines. Database software is often the key to accessing very valuable information—information that unscrupulous parties can turn into money, often with little effort. That simple example shows that software assets can vary in their importance, just as hardware assets can.

Often considered in the same general grouping as software assets are the data sets that the software manipulates. These often contain the highest-valued information the organization has: In the event of a disaster, hardware and the operating and application software that it runs can all be replaced, funded at least partly by insurance payments. But the information unique to the organization—the data files—cannot be bought from any vendor. If the data files have been corrupted, restoring from a regular, sufficient backup program might not help, depending on how long the corruption has existed.

So we've established three broad categories of devices to protect and realized that they have some different hardware and software characteristics to consider when protecting them. But there's another big factor in how you protect them, and it's the old rule of real estate: location, location, location.

Location Matters

Although the SAFE architecture designates several categories of location (discussed in more detail when we address the different SAFE models), for this asset-value discussion, there are really two broad categories: internal-only assets and external-facing assets.

Internal-Only Assets

These are the assets completely inside your network space. There is no access to them except through paths under your control. These assets should be easier to protect because you have some sort of buffering device between them and the outside world. You no doubt noticed the words *should be* instead of *are* used when it comes to the ease of protecting them. Much depends on what kind of security policy is in place and how well it is adhered to. Remember, the SAFE Blueprint is based on the assumption that a security policy is in place and that it is supported (we'll discuss that more fully in Chapter 4, "The Security Policy"). Internal assets are found in any of Cisco's standard three-layer network design model: Access, Distribution, and Core.

Another Network Model?

Some people get confused when they look at a picture of the SAFE Blueprint because they see various modules with unfamiliar names. However, those modules contain the network components of the "standard" Access, Distribution, and Core layers. In fact, Cisco designed the routing and switching structure of the SAFE network design based on the Access, Distribution, and Core model; SAFE did not replace it.

It might help you to understand that the "standard" model is based on a *transportation* functionality approach, while the SAFE architecture is based on a *security* functionality approach. Both are valid ways to describe the parts of the network, and you should use the approach that is best suited to the problem you are trying to solve.

In fact, the Distribution and Core layers should be entirely internal-only assets. Access to them should always be filtered through other devices; reasons for that are covered when we discuss Cisco's SAFE Axioms in Chapters 6, "The SAFE Security Blueprint," and 7, "The Extended SAFE Blueprints." Much of the Access layer is composed of internal-only assets (which sometimes interact only with other internal assets); however, some elements of the Access layer face the dangerous outside world (they are accessible to outside users entering from the Internet or other networks). Cisco often refers to the internal-only part of the network, regardless of its layer, as the *campus module*.

External-Facing Assets

External-facing assets are those that you control but that connect directly to devices that you do not control. These are your edge or perimeter routers, NASs, and firewalls—the guardians of your gates. Cisco refers to this part of the network as the *edge* in the SAFE architecture. Although many of us learned about demilitarized zones (DMZs) as areas to host public-facing servers, the edge is much more than just the DMZ. In fact, the edge contains all the devices that connect to your Internet service provider (ISP), the public switched telephone network (PSTN), your wireless access points (WAPs), and so on. The edge often has more than one device, to ensure that the incoming traffic is acceptable, has been properly filtered, and is then distributed only to places where it has legitimate business going. SAFE is about traffic control, and the edge is the entire zone where incoming traffic meets that control.

Because incoming traffic can be of any type, from anyone, and from anywhere, much of the hardest work in the SAFE model goes into securing the edge. Cisco recommends tighter monitoring and surprisingly tight controls even inside the campus. To understand why, take a look at Chapter 3, "Threats," which discusses the threats your network faces.