The SAFE Security Blueprint

Terms you'll need to understand:

- ✓ Defense in depth
- ✓ SMR
- ✓ IDS
- ✓ False positive
- ✓ HIDS
- ✓ NIDS
- ✓ OOB
- ✓ IPT

Techniques you'll need to master:

- ✓ Managing devices in-band
- ✓ Managing device out-of-band
- ✓ Implementing resiliency
- ✓ Analyzing design alternatives

More than anything else, the SAFE Blueprint is a design philosophy presented as a whitepaper.

Whitepapers

Lots of networking documents are presented as "whitepapers," but how many network engineers—and aspiring network engineers—know what a whitepaper is? When in doubt, go to the dictionary. Merriam-Webster OnLine says that a whitepaper is "a detailed or authoritative report" on a topic. So now you know what's implied when someone says he has a whitepaper on a subject you need to know more about.

Naturally, you should be a little skeptical. One vendor I know is fond of publishing whitepapers about its products—but the whitepapers are pure marketing information, although at least they are somewhat technical. In fact, all vendors' whitepapers have at least some marketing bent to them. *Caveat emptor*—let the buyer beware—applies to whitepapers as well as everything else.

Cisco has provided the SAFE Blueprint as a whitepaper on secure network design. However, it is actually more than just that: The goal of the Blueprint is to provide what Cisco calls best-practice information to interested parties on designing and implementing secure networks. When you add implementation as a goal, you automatically add some important characteristics, such as feasibility, manageability, practicality, minimal user inconvenience, and so forth. You require the design to live in the real world, not just a lab.

Philosophy

SAFE is actually not terribly concerned about the physical placement of devices (such as specifying that a firewall must go on this circuit). Instead, its focus is on assessing the threats to information assets and then placing appropriate technologies where they can mitigate those threats.

Mitigate is an important idea: You can alleviate your security problems and make them less painful, but as long as you choose to have a network, you will have security problems. You can keep them from being as bad as they might otherwise be, but you cannot eliminate them.

Because there will always be problems, the SAFE Blueprint uses defense in depth. This concept applies defense like an onion (with many layers that never get easier to deal with) instead of like a candy (with a hard shell outside and a soft, chewy center). By having many layers of defense, the network

is protected from threats that originate inside as well as outside (and remember that insiders are potentially more dangerous). The network is also protected from threats that manage to get through the outer layer of protection, despite your best efforts.

The SAFE Blueprint is designed to introduce and provide an easy guide—demonstrated using Cisco products, of course—to the concept of defense in depth. That concept has become an important security industry standard, implementing security in layers that include firewalls, backup systems, redundant systems, disaster recovery, and incident handling.

In addition, SAFE does not use repeated layers of the same approach, such as three layers of firewalls. Different security technologies require different capabilities to get past. If a threat manages to penetrate the first layer, it's unlikely to get past the second layer of a different security technology, although it could probably penetrate another layer of the same technology. The result is that you make successful penetration of multiple layers unlikely if the layers are different.

Modular Approach

The OSI model takes a modular approach to data communications for two reasons: to keep the problem manageable by breaking it into smaller pieces, and to internally optimize each module (layer) on a schedule that is independent of that of other layers, as long as it meets interface standards with its adjacent layers. For similar reasons, the SAFE Blueprint organizes the enterprise into modules. Each module can be optimized independently of the optimization schedule of the others, and the architecture can focus on the security relationship between any pair of modules.

This modular approach is helpful in studying the SAFE Blueprint as well. We'll break down the problem into pieces and then tackle each piece in some depth. The basic SAFE Blueprint is intended to describe a design for a secure enterprise network, including e-commerce. As such, this is a design that is very busy and very detailed. It includes redundancy and high-availability (HA) features that can both be present in smaller networks. But because it is the biggest picture, it's worth examining first. To be sure that you're clear about which SAFE whitepaper we're discussing, we refer to this one as the Enterprise SAFE Blueprint.

The Enterprise SAFE Blueprint was the first document of the series of SAFE Blueprints. Thus, it is the template for those that followed. It is also the most complex because it provides what Cisco considers the best practices for the most complex network: a full-blown enterprise, including internally managed e-commerce. However, although the CSI exam is about SAFE in general, it focuses on the SMR SAFE Blueprint—the extension of the SAFE design principles to the small and midsize business and remote-access networks.

Therefore, we look over the Enterprise SAFE in this chapter, followed by the other SAFE models (including SMR SAFE) in Chapter 7, "The Extended SAFE Blueprints." After laying the foundations of SAFE in this chapter, you will be able to focus on how the SAFE Blueprint is applied to the SMR environment. Bear in mind as we do this that techniques are not specific to a particular SAFE model; private VLANs, for instance, work the same way and provide the same benefits, regardless of which SAFE model implements them.

Enterprise SAFE Assumptions

Two fundamental assumptions are made in the Enterprise SAFE model (these assumptions also apply to every extension of the SAFE model). First, the SAFE Blueprint specifically assumes that you already have a security policy in place. Unspoken is the corollary to that, in which we assume that the policy is applied or enforced.

Without a security policy in place, you have not defined the following:

- What you need to protect
- What you choose to protect those things from
- What means of protection are appropriate for your organization
- What you will do if and when protection fails

With those things defined, you can design security for your network; without them, you don't have a basis on which to plan. Of course, if the policy is not enforced or actually applied, it might as well not be present; its value as a guide depends on your being able to assume that it will hold (your position will not be undercut by someone with the authority to operate beyond or outside the rules).

The second assumption in all the SAFE Blueprints is that, although security is designed in, the network remains fully usable for those who need to use it. This does not mean that security will be transparent to the users; they will need to make some accommodations, such as living with stronger security on VPNs. Nonetheless, the idea behind SAFE is to implement network security in a manner complementary to the network and its users, not to make design choices or network use difficult.

Enterprise SAFE Design Objectives

The following are listed as design objectives for the SAFE Blueprint (listed in priority order):

- Security and attack mitigation based on policy
- Security implementation through the infrastructure (not just on specialized security devices)
- Cost-effective deployment
- Secure management and reporting
- Authentication and authorization of users and administrators to critical network resources
- Intrusion detection for critical resources and subnets

Not listed as an objective, per se, but also to be accomplished, is ensuring that the network is resilient and scalable. Resiliency requires that there be no single point of failure; this makes the network more complex because of the additional devices required, as well as the more involved configurations required on all devices. Scalability implies a hierarchical structure, with patterns that can be replicated to yield a larger structure that can be managed in units instead of one device at a time.

Trade-offs will be needed when applying the Blueprint, and some trade-offs will be between the cost savings inherent in adding another function to an existing device (or acquiring one multifunction device) and the performance capabilities of using dedicated devices. The SAFE Blueprint recommends that the decision be made based on performance, not cost—the driver should be the capabilities of the dedicated device compared to the advantages gained from integrating that function with another device. That does not mean that cost will not be a factor because, of course, it will. However, all costs must be considered, including the direct and indirect costs of a security incident (indirect costs include loss of reputation, work not accomplished by those repairing damage done, and so on). In that light, you must have a certain level of performance to avoid those costs, so performance should be your first criterion. (If you have multiple possible solutions—both dedicated appliances and integrated devices will perform as you need—cost could be a deciding factor among the acceptable choices.)

Enterprise SAFE Axioms

Axioms are general truths accepted on their intrinsic merit. They are used as the underlying principles to prove other, less obvious truths. The SAFE Blueprint includes a number of axioms, and you should know what they are.

Routers Are Targets

This is the first axiom. Routers are the traffic cops of the network, determining what traffic is allowed and what is denied, and what those on the other side of a connection know about what is on this side (advertisement). Another view of routers is that they are the gatekeepers of a network. Whichever view you prefer, their role makes them a high-priority target for anyone trying to enter or misuse any information asset on the network. Control of a router makes traffic redirection or illegal entry possible. Routers should be protected by the following:

- Locking down Telnet access
- Locking down SNMP access
- Applying access control through TACACS+
- Disabling all unneeded services
- Requiring authentication of routing updates

These methods are discussed in more detail in Chapter 9, "Products in the Edge." Remember, of course, that the router you are protecting is not an isolated element, but is one element of a larger network design. That design can help protect the router; for instance, an external firewall can help limit the potential threats that can reach the router. Likewise, you can always take measures that are not listed in this part of the SAFE Blueprint or in the IOS Security Configuration Guide, such as logging activities that might have security impacts—and reviewing those logs.

Switches Are Targets

Perversely, there is less readily available information on securing switches than there is on securing routers, yet switches touch more resources directly—especially the high-value resources, such as servers—than routers do. The recommendations for routers certainly apply (with the possible exception of the last item, depending on whether the switch also operates at Layer 3). However, because they access so many devices via so many ports, and because

they operate primarily at Layer 2, where filtering and traffic control is more difficult, they require extra precautions:

- Disable all unused ports.
- Enable trunking (including automatic trunking) only on ports that actually need it.
- Trunking ports require extra precautions:
  - Use a dedicated VLAN ID.
  - VLAN 1 might have special meaning for some vendors.
  - Eliminate native VLANs from any 802.1q trunks.
- If possible, limit the MAC addresses associated with a given port to two or three. (Remember that a switch essentially directs traffic according to the MAC addresses from which traffic enters a given port. If you have many addresses at one port—such as a many-device hubbed segment attached to the port—controlling traffic to and from the port is more complex than if you have implemented a switched network with very few hosts per port.)
- Manage change control (especially on switches that can be modified by multiple departments).
- Use private VLANs to limit traffic between hosts in the same VLAN (force the traffic to Layer 3 for filtering).

These methods are discussed in more detail in Chapter 8, "Products in the Campus."

Private VLANs?

If you are not familiar with private VLANs, you should be: They are used in almost every module of the SAFE Blueprint. In an ordinary VLAN, members of the VLAN exchange traffic via Layer 2; traffic is never decapsulated as far as Layer 3, so security features made possible by filtering on IP or upper-layer headers are not available. Even if traffic is on the same physical segment, traffic belonging to different VLANs is not seen by hosts except those that belong to the appropriate VLAN (hosts that belong to VLAN 3 do not see traffic that belongs to VLAN 6, for instance). To see such traffic, it must be handled at Layer 3 by inter-VLAN routing, sometimes called a "router on a stick," or by another Layer 3 device.

Private VLANs mitigate the capability of VLAN members to see the traffic of other members of their VLAN. Ports of a private VLAN fall into one of three categories: isolated, community, or promiscuous. Isolated ports can communicate only at Layer 2 with promiscuous ports of the same VLAN. Community ports can communicate at Layer 2 with other members of the same

community or with promiscuous ports (within that VLAN). Obviously, promiscuous ports can communicate with any other port in the same VLAN at Layer 2.

By isolating ports within a VLAN, it is more difficult for malware to use port redirection to migrate problems from one host to another in that VLAN. Because you will often find servers (high-value targets) grouped for ease of management on one VLAN, this is a good way to protect them.

The previous information on routers and switches is no surprise if you already have a CCNP. Recall that the CCSP was originally a specialization for those with a CCNP. The guidance in the SAFE Blueprints assumes, at least implicitly, an intermediate-level knowledge of routing and switching. However, the only prerequisite to earning a CCSP now is the CCNA, which does not cover this material as deeply. If you do not have your CCNP yet, you will need to study and think about traffic flows and which devices do what tasks at which layers of the OSI model. By the time one has earned a CCNP, this knowledge is virtually automatic. You can pass this exam without a CCNP; you just have to work on routing and switching concepts more.

Hosts Are Targets

Hosts are the most difficult asset to protect because the network administrator cannot entirely prevent users from modifying configurations or adding unauthorized software and services. Plus, of course, there are so many of them: There are generally far more hosts than there are any other class of asset. Hosts are often a mix of hardware and software sets, reflecting generations of the organization's acquisitions. Adding to the difficulty, hosts are more than workstations; hosts include servers, including the public-facing servers, which are frequent attack targets. Securing hosts is a matter of staying on top of the security status of every component—hardware as well as software—for patches and upgrades. When it comes to this work, smart administrators prioritize their efforts, protecting the most valuable and/or the most vulnerable first. However, as you have seen with recent worm attacks, all hosts are targets and must be secured—one bad apple can indeed spoil the barrel, not to mention your entire day.

Networks Are Targets

Up to now, we've looked at targeting specific devices, but it pays to remember that the entire system can be a target as well. Network attacks often take advantage of weaknesses in the networking protocols themselves. Network attacks can use Layer 2 or Layer 3 protocols (ARP- and MAC-based attacks, IP spoofing, and so on). *The worst attack, though, is the one you cannot stop.* If you can't stop it, how can you begin to recover? These are often DoS or DDoS attacks, and the only way to counter them is with help from

upstream—your carrier or service provider must rate limit your incoming traffic. But to get effective help, you must be able to characterize the traffic to be throttled, which means that you will be examining your logs for sources, protocols, ports, and so on to specify as tightly as possible the traffic to deny so that legitimate traffic (hopefully) still gets through.

Some techniques to help you minimize the likelihood of such attacks are described in Chapter 9, when we look at the Edge module.

Applications Are Targets

We all know that some applications are easier to attack than others. Those are (hopefully) the applications that you patch the most; of course, they could also be applications for which patches have not been prepared. Applications, like operating systems, have become bloated as vendors have offered more options and features. With software code running to hundreds of thousands and (too often) millions of lines, human error alone will introduce vulnerabilities. If the software also suffers from design flaws, the vulnerabilities could be more pervasive or more severe. However, that leads us to the next axiom….

Intrusion-Detection Systems Help

Intrusion-detection systems (IDS) can run on individual hosts (host IDS, or HIDS) or on a networking device (network IDS, or NIDS), monitoring a given traffic flow. The goal is to protect the applications that are being targeted for exploitation. IDSs operate on the same principle as antivirus packages: They have a data set of known characteristics of attacks. When a traffic flow occurs matching the (known) characteristics, an alarm can be generated, the traffic can be dropped, or a combination of actions can be taken.

HIDSs monitor specific traffic types—to a mail server, for instance—for which compromise is a known high-risk event. Therefore, HIDSs are often set to automatically drop suspicious traffic (a better-safe-than-sorry approach). However, HIDS sees only a subset of the traffic, so NIDS is needed as well. Because NIDS sees so much more traffic, it is often set only to alarm rather than to presumptively drop traffic. This is one of those tradeoffs to be taken with careful thought, depending on what traffic should reasonably flow on the segment being monitored, as well as the nature and value of the assets on the segment being protected. As you will see in Chapters 8 and 9, in high-value segments of the network, you might need both HIDS and NIDS deployed to maximize your ability to find problems before they become too big and much too expensive.

Having made IDSs sound like your best friend, some caution is in order: You must expect to take some time to "settle in" your IDS implementation. There will be false positives, instances in which IDS falsely categorizes traffic as an attack. You will find these when legitimate traffic does not get through and people complain. Unfortunately, false negatives are a problem, too, but you learn of them when the IDS did not catch a real problem (its negative reading was false) and the problem becomes your network problem to solve. Likewise, the IDS can detect only what it knows about: You must depend on the vendor to update the attack signature profile promptly. If you are the lucky first network to suffer from a new attack, the IDS cannot help you until your vendor can characterize the attack and add the relevant data to the profile.

When IDS blocks traffic, it can use what is called "shunning" the traffic: dynamically adding entries to access lists and filters to drop the offending traffic. Typically, the traffic is blocked for a short period of time, long enough for the network administrator to determine whether there is a real problem (a true or false positive). Alternatively, if the traffic uses TCP, the IDS can send a TCP Reset (the RST flag in the TCP header is set). A reset prevents whatever action is ongoing between the attacking host and its target on your network from continuing; as with shunning, however, it is likely to continue only long enough to give network monitors—humans watching the system—time to do something in response to the attack.

IDSs will help, but they are not a miracle cure.

Secure Management and Reporting

The last axiom in the Enterprise SAFE Blueprint is that you must read your logs for them to do you any good. This is hardly a surprise; in fact, you already know that there are many logs, and they each contain many messages, only some of which are truly important. But those two characteristics of system logs too often lead to spot checks (at best). Even if you do read the logs religiously, you have to wonder how confident you can be of their content—after all, hackers know that logs can reveal what they did, so they target logging. That is the point of this axiom: You must secure your network-management and reporting system for it to be reliable. Then you can confidently use one of the many utilities to parse and analyze those logs, creating reports that direct your attention to suspected problems.

This discussion revolves around one direction of traffic: severs and networking devices reporting to you. However, the other direction is just as important. You must have a secure means of sending your management traffic to the important devices. As a solution to both sides of this problem, the SAFE

Blueprint recommends that you establish a separate network for your network-management needs, preferably with direct connections to the networking devices. This network means that you will exercise management out-of-band (OOB) with respect to the regular traffic. In fact, this separate network carries no production traffic; however, the syslog hosts are on this network (not the production network). Establishing a separate network requires a few more ports and a separate logical network. This level of structure usually is justifiable only with a large, enterprise-level network or one on which network control and management have exceptional value (perhaps protecting particularly high-value content). If OOB is not practical or perhaps not desirable, there are alternatives.

One of the strengths of the SAFE Blueprint is that not everyone is forced to fit in the same box; there are always design alternatives. Of course, remembering the alternatives is a little more work, but you'll find it useful during the exam—you can expect to see questions about the alternatives available to a certain module's design.

If you are managing your devices in-band, you want to be sure that only you—definitely not anyone else—make configuration changes (and especially any change to logging—with that changed, you might not know that something else was changed). You have two acceptable choices for secure in-band device management:

- If the device has IPSec, perform device management via an IPSec tunnel terminating on the device.
- If IPSec is not available, use SSH instead of Telnet (and disable Telnet so that no one else can use it, either).

Unfortunately, there is no secure substitute for TFTP (or FTP), but you can sharply limit the addresses allowed to run either or both of them with an access list. Likewise, you can limit the acceptable addresses for SNMP, but consider whether you really must use that for management or whether you can work with it only for reporting (if so, use a read-only community). You should also consider using only SNMPv3 and taking advantage of its improved security features.

Finally, the SAFE Blueprint reminds you to manage configurations carefully and archive them to a safe place via TFTP or FTP—you might need to recover a copy, and copying in an archive is better than rebuilding the configuration.

Modular Approach

The Enterprise SAFE model divides the network into two broad categories, the edge and the campus, further subdividing each of these into modules. The enterprise edge is composed of the following modules:

- The E-Commerce module
- The Corporate Internet module
- ISP modules (one per ISP—two are assumed, for reliability, named ISP A and ISP B)
- The VPN/Remote Access module
- The PSTN module, which connects via the VPN/RA module
- The WAN module
- The Frame/ATM module, which connects via the WAN module

We will dig into the components of each module shortly, but first it's useful to look at their relationships, or how they are interconnected. An overview of the modules and their connections is shown in Figure 6.1; the modules are presented "in silhouette," to keep you from being distracted by the wealth of internal details.

A number of factors are worth noting here. First, the E-Commerce module is separated from the Corporate Internet module. Although both access the Internet via redundant links (via redundant ISPs), the E-Commerce module must allow strangers at least some access, although you do not want strangers entering your internal network via your users' Internet access. Notice that there is no direct connectivity between these two modules; their only interconnection is via the links passing through the ISPs. The E-Commerce module is a DMZ, well isolated from the production network's Internet access. However, the E-Commerce module has access to the heart of the production network via the Edge Distribution module; later examination will show that linkage to be heavily protected.

A second point to notice is that the VPN/Remote Access module connects to the outside world via other modules: via the Corporate Internet module for Internet access, and via the PSTN module for dialup service. Finally, the WAN module interconnects the Edge Distribution module to the Frame/ATM Network. Because these are dedicated circuits, you will find this interconnection less heavily protected.

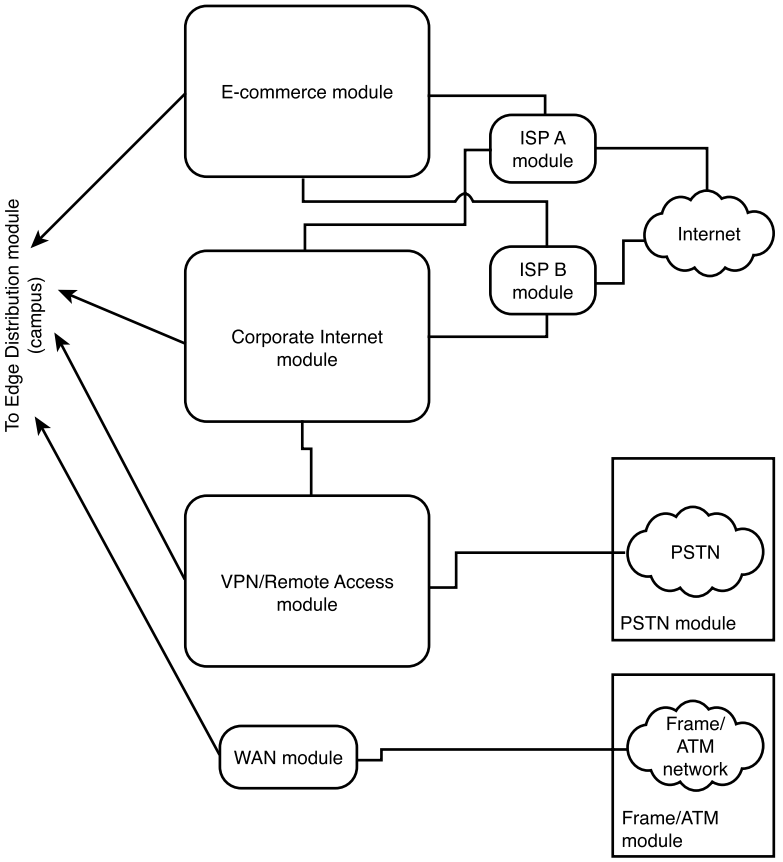

Figure 6.1 Edge module relationships.

The enterprise campus is actually more complex in its internal structure, but, paradoxically, its modular relationships are simpler. Again, using only the silhouettes to capture the relationships, the architecture of the campus is as shown in Figure 6.2.

If there is only one building, of course, this is even simpler. Note that there is one modification in Figure 6.2 compared to the diagrams in the Enterprise SAFE Blueprint. I added the dashed line for connectivity between the Management module and the Building Distribution module. The SAFE Blueprint itself does not address how the Management module connects to the networking devices in the other modules, although there must be some connections for management to happen. In fact, the Blueprint specifies only that a terminal server can be used to connect directly to devices, a router (with IOS firewall) can be used to connect with encrypted in-band management, and switches can be used for out-of-band management. Because these

connections must get from the devices in the Management module to the other devices wherever they might be throughout the network, I have connected the Management module to the rest of the network via the Building Distribution module, which, of course, connects to everywhere else.

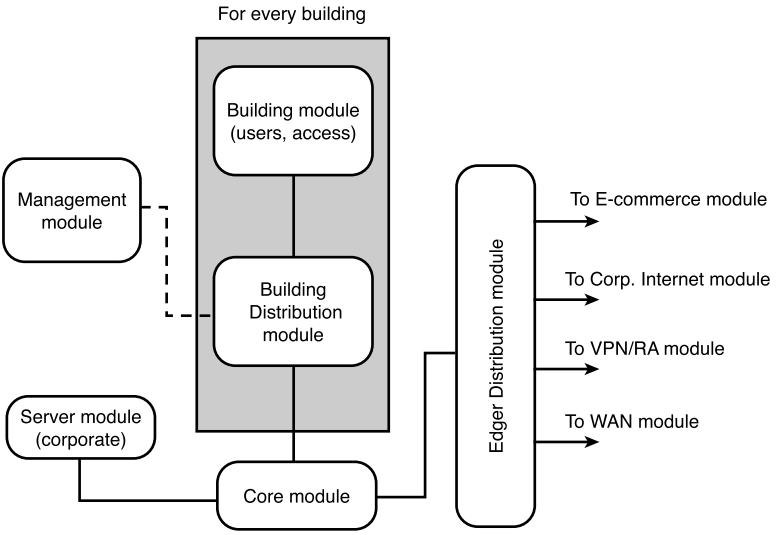

Figure 6.2 Campus module relationships.

Now let's take each module in turn and look at its high points. Much of what we note about these modules will apply to the modules in every version of the SAFE Blueprint, including the SMR SAFE Blueprint.

Edge Modules

The edge, of course, consists of the portions of your network that interact with the outside world, whether that outside is an ISP, the telephone system, or a leased line. It can also include customers, people who actually want to send you money for your goods and services. Of course, when money is involved, the security requirements—and the reliability requirements—become very important. Therefore, it helps to separate commercial activity from the rest of your interactions with the outside world. Even among the others, different kinds of security problems arise with different kinds of connectivity, so the Enterprise SAFE Blueprint addresses each type in a separate module.

E-Commerce Module

The E-Commerce module is fully redundant. It is shown in Figure 6.3; the various server sets are, indeed, multiple servers, each of whom is connected to both switches of the pair serving them:

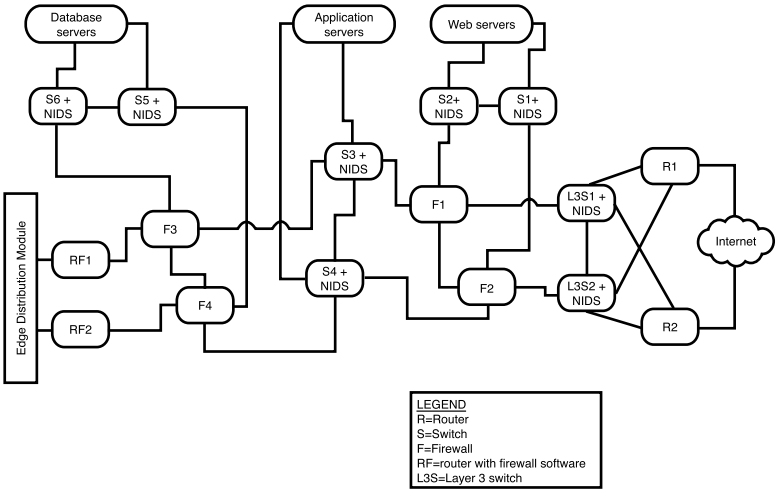

Figure 6.3 E-Commerce module.

Ingress for customer traffic from the Internet can come over either of two routes, from ISP A or ISP B, as represented by the perimeter routers R1 and R2.

Three types of servers sets exist:

- Web servers (the SAFE diagrams show three, but there is no "magic number" beyond the two required for redundancy—there should be as many as are needed to handle the traffic load)
- Application servers (for applications that support the Web servers)
- Database servers (supporting both the application servers and the Web servers)

The server types are isolated from one another, not only by switches, but by firewalls as well. A compromised Web server will not easily infect an application server or a database server.

Speaking of switches and firewalls, these are redundant as well (with a pair serving each distribution point), along with the perimeter routers at the two

module edges (the Internet and the Edge Distribution module). Firewalls connect to single switches instead of to both members of a redundant set (to simplify firewall data-passage rules). Each switch has an associated NIDS device, while the Layer 3 switches at the Internet-facing edge have NIDS installed (Layer 3 switches are used in addition to the routers for traffic load handling in hardware, leaving the routers to do what they do best: route traffic).

Each of the Layer 2 switches segregating server sets can implement private VLANs to force traffic between devices to pass through Layer 3 for inspection and filtering. Finally, before any traffic is sent to the Edge Distribution module, it must also pass through a router with a software firewall set installed.

This might seem like a "belt-and-suspenders" approach—in fact, the entire system might seem overdone—but this module is most likely to be attacked, both early and often. At the same time, performance matters to customers, so functions are highly segmented: Each device has only a limited number of functions to perform, enabling you to optimize it to perform them securely at speed (they should not cause congestion or be a choke point). Different security techniques make this frequently attacked and potentially compromised module unlikely to provide an attack path into the corporate network. Both redundancy and resiliency are built in.

Design Alternatives

The simplest design alternative to this module is to let someone else handle it: Offload the e-commerce infrastructure to a service provider. If this is done, the connection to manage the e-commerce resources will take place over the Internet connection, requiring the capability to secure that (such as encrypted tunnels or private lines to the e-commerce management access).

In addition, depending on the sensitivity of the commerce conducted (high monetary value, for instance), it might be desirable to use additional firewalls to isolate devices further. Security generally is enhanced if different firewall types are used, limiting the utility of a single exploit.

Corporate Internet Module

Unlike the E-Commerce module, the Corporate Internet Module is about internal access out (while e-commerce is about external access in, into a very specially protected space called the DMZ). Of course, replies to internal-out traffic are permitted, but unsolicited external-origin traffic is extremely limited, if it is permitted at all.

The internal structure of the Corporate Internet module is shown in Figure 6.4.

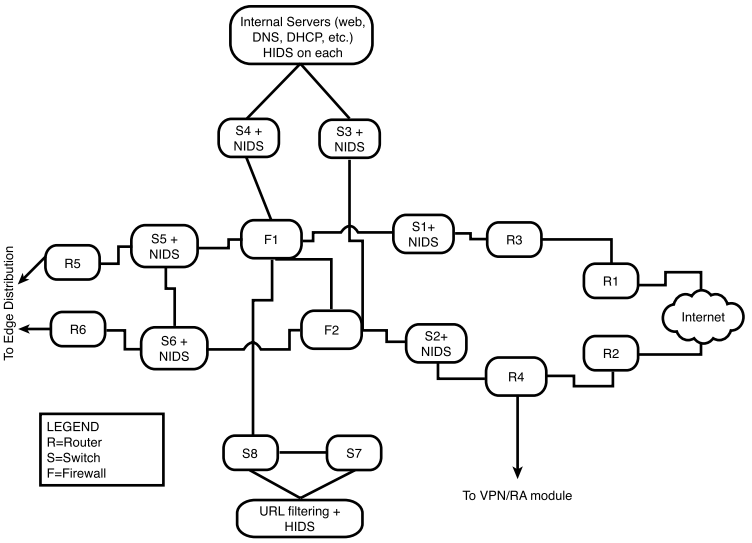

Figure 6.4 Corporate Internet module.

This module is simpler, needing less firewalling (although all traffic to or from the Internet must pass through a firewall). As with the E-Commerce module diagram, Internet access via R1 and R2 passes through the ISP modules (which are not shown, to keep the diagram at a manageable size). In addition to NIDS with most of the switches, the servers themselves have HIDS. The "doubling" of routers at the Internet edge might seem superfluous, but the goal again is efficiency: The perimeter routers (R1 and R2) perform fundamental filtering and rate limiting, while R3 and R4 perform traffic distribution, including to the VPN/Remote Access module. The switches again use private VLANs to protect individual devices. The NIDS are usually appliances attached to the relevant switch; if throughput is an issue, a Layer 3 switch can use a dedicated blade for a NIDS and take advantage of the greater throughput of the switch's backplane.

Design Alternatives

The NIDSs on the path from the Internet to the firewalls record potential attacks that the firewalls might silently discard. This is useful information. However, if no basic filtering is being done at Internet access (on R3 and R4), these NIDSs might be overwhelmed.

One other design alternative is to eliminate the "extra" routers at the Internet edge—R3 and R4—and collapse their functions into R1 and R2. This depends on how comfortable you are with the performance of your devices at R1 and R2, as well as how much work they have to do (how busy

your Internet connections are and how much filtering is actually done at your ISPs before the traffic gets to your edge).

VPN/Remote Access Module

The VPN/Remote Access module needs to accommodate incoming traffic from three different sources: remote sites (via the Corporate Internet module, passing into a router/firewall), remote users (also via the Corporate Internet module, but passing into a VPN concentrator), and dialup users (via the PSTN). Tunnels can be GRE, IPSec, L2TP, or PPTP. The router/firewall and VPN concentrator are both capable of handling tunnel termination for many tunnels; the router handling termination of dialup access need not handle so many. By segregating the ingress of the three types, the load on each ingress device is manageable; the structure is shown in Figure 6.5.

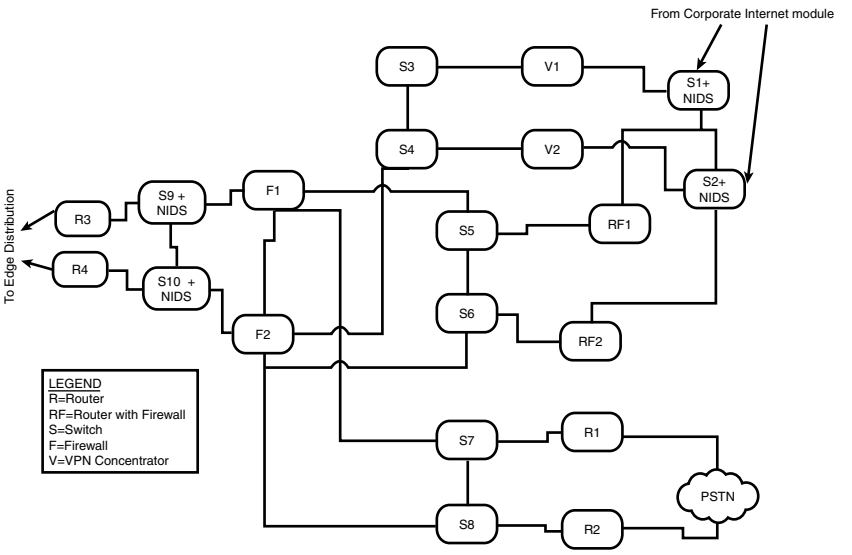

Figure 6.5 VPN/Remote Access module.

Layer 2 switches distribute the traffic, but paths through the firewalls are kept discrete as much as possible while maintaining redundancy. The device roles are quite similar to what you have seen in previous modules.

One item to remember is that VPNs using IPSec might need IKE (UDP port 500) and ESP (protocol 50), plus UDP port 10000 if the ESP traffic must be tunneled inside UDP because of firewalling or NAT traffic management between the two endpoints. Segregating the traffic limits how many openings you need on the ingress devices.

As a further note, remember that the dialup users should always be required to authenticate with CHAP rather than PAP.

Design Alternatives

Design alternatives for this module are discussed in a separate whitepaper, the VPN SAFE Blueprint, which we cover (though not deeply) in Chapter 7.

WAN Module

The WAN module is very simple, consisting of ingress from the service providers' networks and filtering on the ingress routers. It is shown in Figure 6.6.

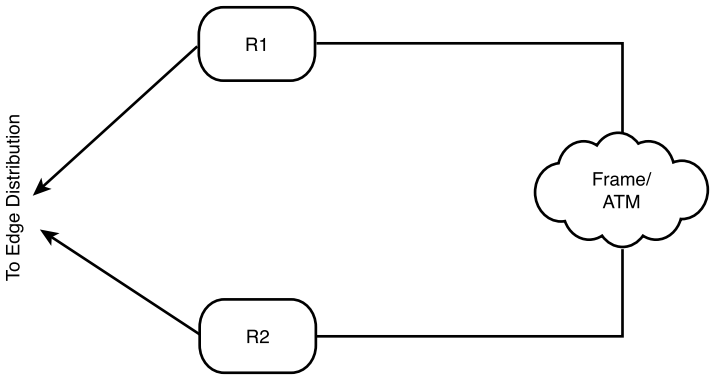

Figure 6.6 WAN module.

Design Alternatives

The simpler the design, the fewer the points from which you can diverge to an alternative. In the case of such a simple design as the WAN module, the only alternative is to protect even further the data traveling over private circuits through encryption (IPSec).

Edge Summary

That is the enterprise edge, a group of modules that are designed to filter and analyze traffic, segregating it into portions of like type that can be processed by devices optimized for just that task. Internally, these modules use segregation again to mitigate the effects of any attack that penetrates—techniques such as a switched architecture, to minimize the knowledge that

can be gained with packet sniffers; private VLANs, to minimize the opportunity to leverage a compromise of one host into the compromise of others; and stateful firewalls.

That leaves the heart of the network, the enterprise campus.

Campus Modules

The campus (or enterprise campus) modules are not as difficult to deal with (the overview of the campus was shown in Figure 6.2). In fact, the basic Access-Distribution-Core design for a given building involves nothing special in security terms. All hosts should have antivirus software, should be kept current, and should scan all files, of course. In addition, Layer 3 filtering (RFC 2827 filtering, discussed in more detail in Chapter 9), should be present. Likewise, the Edge Distribution module serves as a connector between the campus and the edge; protection is placed nearer the sources of traffic that must cross it than in this module.

The modules where you need to make modifications to implement a secure network design are the Management module and, to a lesser extent, the Server module (the group of internal use–only application servers). Because that is the lightweight problem, we deal with it first.

Server Module

The Server module inside the enterprise campus is the home of servers that support getting real work done: the enterprise's applications. Recall that network support servers (Web, DNS, DHCP, and so on) were actually located in the enterprise edge, as part of the Corporate Internet module. The servers we are concerned with here are those run by departments to deliver their portion of the enterprise's business activity.

In terms of security features, a look at Figure 6.7 reveals that you don't need to do too much.

Although these servers are accessible only via switches with NIDS and their exposure is extremely limited, it is advisable to place HIDS on them as well. These servers are the home of sensitive, internal corporate information, whose integrity (not to mention its presence) cannot be questionable. The Layer 3 switches should have sufficient NIDS blades to handle to traffic load (to be able to monitor all of it), and the ports to the servers should be set up as private VLANs.

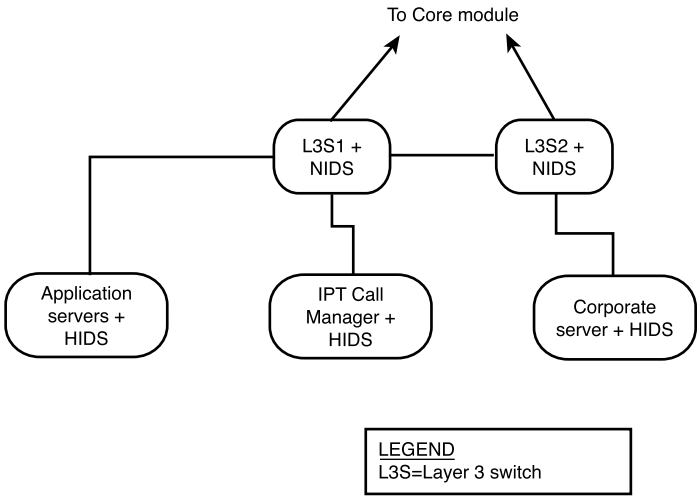

Figure 6.7 Server module.

Design Alternatives

Again, as a simple module, there are only a few design alternatives. If resources are limited, it is possible to collapse this module into the Enterprise Core module, although there are obvious traffic-management questions to be handled in that event.

If IP Telephony (IPT) or a particular server (such as one containing financial data) are considered special risks, you can add a stateful firewall between these sensitive devices and the rest of the network.

Management Module

The Management module is more complex because it must handle many different tasks. Remember, in Figure 6.2, the Management module was connected to the Building Distribution module, but that is conceptual, not from Cisco. The Management module's component structure is shown in Figure 6.8.

The Management module separates network management into two broad zones: outside the firewall, where network connections to devices (both in-band and out-of-band) exist, and inside the firewall, where the management hosts and the connection to console ports exist. Of course, the management network uses a different address block than the rest of the network (and, if using one of the private RFC 1918 address spaces, often a block from a different private address space). Routing protocols on the terminal server routers and the router with the firewall do not advertise routes to the rest of the network.

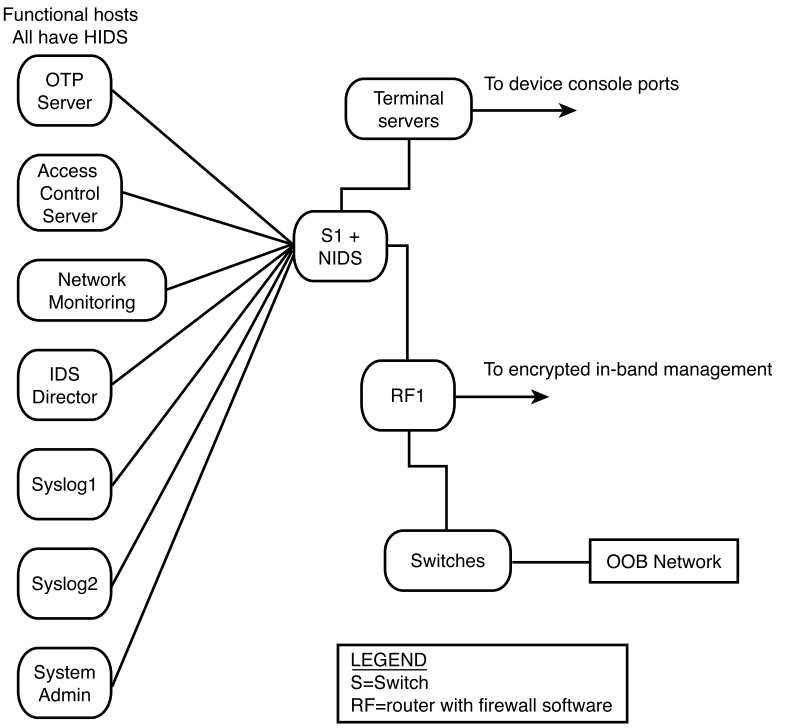

Figure 6.8 Management module.

This module is a priority target for any attacker, first because of its access to the devices that control the network, and, second, for its record-keeping functions (and the opportunity to "sanitize" those records). Naturally, security is many-layered here: the dedicated firewall on the access router, the separate addressing scheme, private VLANS, NIDS, and HIDS.

Design Alternatives

Alternative approaches to the design of this module are somewhat limited, not because of inherent simplicity, but rather because of the critical functions here that must be protected. If resource limitations require the use of in-band management, especially exclusively in-band management, great care must be taken to ensure that only the proper users have access to the network devices (as strong as possible authentication) and that any file manipulation (configurations and images) is done securely (SSH versus Telnet, IPSec tunnels, very restrictive access lists, and so on).

If throughput between the management stations and the network devices is an issue, a dedicated hardware firewall can be added, offloading the firewall function from the router's software.

Summary

This has been a busy chapter, but it has provided you with an important baseline. With its many security technologies and many layers of implementation of them, the SAFE Blueprint mitigates a number of network threats. Here are a few that we have developed just from technology placement:

- *Packet sniffing*—Mitigated through extensive use of switched networks, limiting the amount of traffic to be discovered
- *Port redirection*—Mitigated through the use of private VLANs, to limit Layer 2 traffic flows
- *Unauthorized access*—Mitigated through AAA, especially on critical information assets
- *Network intruders*—Mitigated by extensive use of IDS, both NIDS and HIDS, which recognize known attack traffic

As you look at the products in Chapters 8 and 9, you will see more ways in which the SAFE design mitigates threats to your network. Before you can do that, however, you need to look at the other versions of SAFE—the extensions of the blueprint to VPNs, IP telephony, wireless, and (for our purposes, the most important) the small and midsize business networks and remote access (SMR). Those are the subject of Chapter 7.