10

The Small Network Implementation

Terms you’ll need to understand:

✓ Standalone

✓ Branch

Techniques you’ll need to master:

✓ Matching threats to mitigation techniques

✓ Analyzing design alternatives

As you put the pieces of the SAFE Blueprint together, you need to consider the following aspects of each module's design:

- Assets to be protected
- Threats to those assets
- Devices used and their implementation and configuration
- Threats mitigated
- Design alternatives

In the case of the small business network, we're talking about a small edge and a small campus. It is worth remembering that this small network might be an entire business in its own right (standalone) or a branch operation of a larger organization (branch). Either way, it must be secured as an entity.

> EXAM ALERT
>
> The SAFE exam description says that you will be tested on the knowledge and skills to implement the principles and axioms in the SAFE SMR Blueprint. Implementing principles and axioms is less about command syntax than it is about which commands to enter and about which devices to put where *and why*. You should be prepared to answer questions about why a device is used in a certain place in the design instead of another device. In other words, know the alternatives within each module and what each alternative brings to the Blueprint.

For review, Figure 10.1 shows the small business network model in its entirety.

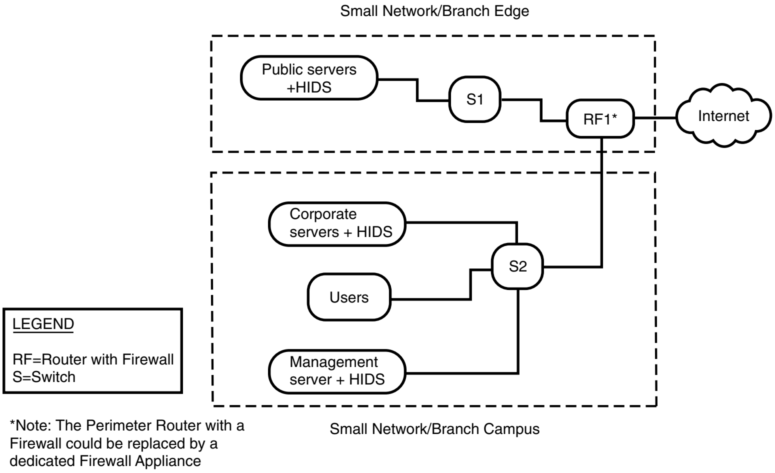
Small Network/Branch Edge
Public servers +HIDS
S1
RF1*
Internet
Corporate servers + HIDS
S2
Users
Management server + HIDS
LEGEND
RF=Router with Firewall
S=Switch
*Note: The Perimeter Router with a Firewall could be replaced by a dedicated Firewall Appliance
Small Network/Branch Campus

Figure 10.1 The small business/branch network.

The Small Network Edge

The small network's edge is composed of a Corporate Internet module that is far from complex: an Internet connection via a router with a firewall (or a dedicated firewall appliance, as noted), with incoming traffic for public servers sent to a switch that further sends it to the correct host. This module terminates incoming VPNs (quite possibly for professionals working from home as well as teleworkers or personnel who are traveling). The public-facing servers can include Web, mail, file transfer, and name servers.

Assets

Assets in this module are the router or firewall itself, the switch, the servers, and their contents. The router with a software firewall or the dedicated firewall appliance (henceforth, just referred to as the router/firewall) is certainly a valuable resource, representing the ingress choke point. The switch segregates the traffic of one server from another, which makes it a target for manipulating traffic inside the module. The servers themselves are targets, of course, for the information they contain and for the possibility that an application vulnerability on one can be capitalized upon and turned into access elsewhere in the network. Concerning the content of the servers and the asset value, you have to think like a bad guy: How can I turn this into money or satisfaction?

If present, the Web server might have Web pages to be defaced (for pure vandalism or competitive sabotage). The file server might have product information files available for download (a colleague of mine once browsed anonymously from home into the FTP server of a hardware contract manufacturer and found the complete production specifications of a Layer 3 switch—not from Cisco—openly available). Even though the mail server should be a relay, not the main mail server, it might have copies of significant business correspondence. The DNS server has zone files, which can be altered to point prospective business contacts to a "404 Not Found" error (causing this business to suffer reputation damage and perhaps lost business) or simply to a competitor's Web site. All these could be done for hire.

In addition to business disruption, whether vandalism or competitive, the servers could be "borrowed" or hijacked to support activities that are illegal in that area, from storing and serving pornography (as defined locally) to gambling, serving as a mail relay for fraudulent solicitations (email scams), and so forth. Even if not locally illegal, the organization's hosting of this material would probably be extremely damaging to its reputation—and such damage leaves essentially permanent effects because the taint lingers in people’s minds. “Data kidnappings” have actually occurred—sensitive corporate data was encrypted in place, and the key to decrypt it was obtained for a ransom payment. In addition, other forms of extortion (protection payments to not be hacked again, for instance) are believed to be widely underreported to law enforcement authorities.

The point is that the business can suffer damage—real, substantial financial damage—if its public servers are hacked. The way into those servers is via the router/firewall and the switch, which makes these networking devices targets to be neutralized (if not taken over) in the approach to the servers.

Threats

We have just covered some of the threats possible to these assets, but those were threats of what could happen to them. We need to concern ourselves here with the threat technologies and, to a lesser extent, the people who might wield them. The reason we focus less on the people is that they are harder to predict.

Threats that can be encountered in this module are grouped into one of two broad categories: low-profile and forceful. The low-profile threats include the following:

- IP spoofing
- Network reconnaissance
- Packet sniffers
- Trust exploitation
- Unauthorized access

Other threats attempt to apply force (digital force, if you will) to exploit the network including these:

- Application-layer attacks
- Denial of service (DoS)
- Password attacks
- Port redirection
- Virus and trojan horse attacks

These are the same attacks as listed in the SMR SAFE Blueprint, but I have rearranged them into the two qualitative groupings based on shared characteristic (attack style). Smaller groups also are simply easier to remember and associate.

You might remember that one of the axioms of the SAFE Blueprint involved reading your logs, not just keeping them. Unfortunately, smaller businesses rarely have a multiperson IT staff; they might have no one and outsource their network management altogether. Such networks are readily susceptible to low-profile attacks, which, if done with any reasonable level of skill, can go on for some time before being noticed in a relatively unsecured or "unattended to" network. Trust exploitation is especially hard to detect because (technically) nothing has gone wrong. A device that was trusted has accessed a device that did the trusting, and that relationship is intentional (the FTP server trusts the Web server, for instance, because the request came from an internal address, so the FTP server automatically honors any request from it).

Digital force attacks, however, are usually not subtle. Even password attacks can involve repeated login attempts, such as battering at a door until the lock gives way. Although a login function doesn't "give way" in the physical, materiel-fatigue sense, a dictionary attack on a password function is likely to succeed in a shockingly short period of time.

NOTE

If a password is six characters, at 8 bits per ASCII character (any character, letter, number, or special character), that is 48 bits, with 2^{48} possible resulting binary strings. After 2^{24} (16,777,216) tries, however, even a brute force "try every combination" attack has a 50-50 chance of finding the correct string. An "average" PC can attempt approximately one million strings per second (and a customized PC can do an order of magnitude more), so you can see how little time it would take to match the hash recovered from a sniffer or copied from a password archive. Make the password eight characters, and it takes 2^{32} (4,294,967,296) tries for the 50-50 chance. That's 256 times as many tries.

Of course, that assumes truly random searches, but people almost never create random string passwords because those are too hard to remember. Instead, people use words and names, which come from a much smaller universe (the dictionary, and typically a small subset of that); they are much more easily cracked. Applying precalculated hashes from a dictionary-like database, a program that is readily available on the Internet can crack the typical NT-kernel password (Windows NT 4.0, Windows 2000, Windows XP) in seconds. A link to statistics on password cracks using this dictionary approach is given at the end of the chapter.

Password attacks are only as difficult as your password policy makes them.

Port redirection, application-layer attacks, and virus and trojan horse attacks all exploit technological weaknesses present in software. Port redirection is especially difficult to counter because the traffic enters on a legitimate port—port 80, HTTP, for instance—and a malicious software application redirects that traffic to another port used by a different process. This can be used, for instance, to open an FTP connection to upload more malicious software, enabling the hijacking of a server, or to facilitate the progressive compromise of a host.

These more subtle attacks can proceed slowly if they are skillfully crafted, but more often they run amok, presenting you with a cascade of errors and problems that seem critical. Of course, DoS attacks simply consume resources so that others cannot use them—but the consumption means that real work is not getting done.

The people perpetrating an attack might be quickly recognizable, especially if they are relatively new and clumsy. More often than not, however, you cannot identify the perpetrator quickly, if at all. To make things more interesting, you might not ever be able to tell with certainty whether the attack originated externally or internally. Remember the outsourced IT: If a contractor present on the site two months ago placed a trojan horse in the system, what appears to be an external attack today actually originated internally two months ago.

This uncertainty is one reason the SAFE Blueprints do not spend significant attention on forensics (aside from the fact that it is a large and highly technical topic in its own right). Another reason is that the SAFE Blueprints are practical: They focus on what you should do to prevent problems and mitigate their effects if they do occur. Your hands will be quite full with that.

Devices and Implementation

The servers are the principal target of most hackers. A strong antivirus package should be installed and kept current on every host in the entire network, including the servers in the edge. Every host should also be periodically scanned for viruses and other malware (trojan horses, keystroke loggers, and so on). Every application and OS on the servers should be kept current on patches, and no unnecessary processes should be allowed to run. This is often referred to as "locking down" the OS or the server (depending on the immediate subject).

Wherever possible, all applications should be configured to accept updates only from specified internal addresses. An example of this is the DNS server: It is the source of name-to-address mappings for the public. You might choose to make this DNS server in the edge the network's authoritative server, in which case another server inside your network is the secondary server. The DNS server in the DMZ should accept a zone transfer only from your legitimate secondary server at a specified inside address. If this is a secondary server, it should accept a zone transfer only from the authoritative server inside the campus. Either way, it should never accept a zone transfer from any host but the specified server—specified by IP address because name mappings can be compromised (if nothing else, through the addition of a hosts file).

> EXAM ALERT
>
> The previous paragraph is an example of thinking about more than networking devices. Remember, we said back in Chapter 2, "Information Assets," that such thinking is one of the differences between this exam and other Cisco exams: You must be prepared to mitigate threats on the hosts—servers and workstations—as well as on the routers, switches, VPN devices, and firewalls.

A host IDS should be one of the processes running on every server in the edge. A number of vendors make HIDS; although the SAFE Blueprint was validated in the lab using HIDS from Entercept, Cisco now offers such a product with its acquisition of Okena (the new product is known as the Cisco Security Agent). Because the HIDS is being applied to a specific host serving a specific function, it can be set to be aggressive in its protection: Rather than merely sending an alarm, it should be configured to send an alarm and to drop offending traffic (sending a reset might depend on whether you want the attacker to realize that he got as far as the server—it can be debated whether this is really wise).

The switch in this module isolates the traffic for each server to a separate wire (in the nature of switches). Using private VLANs, this separation can be made stronger. Servers should be on private ports rather than community ports and certainly not on promiscuous ports. Any unused ports should be disabled rather than left available for device connectivity. Unless there is an overwhelming need, there should be no trunking on any port facing the servers in this DMZ: Traffic between VLANs should have to pass through Layer 3 inspection and filtering—at the router/firewall—before being permitted to pass.

The router or firewall is where the security heavy lifting is done in this small edge. For ingressing traffic, it filters both according to RFC 2827 and RFC 1918, and it filters out fragmented packets (which are often used to consume router and server resources, especially if the fragmentation is "artful"—that is, deliberately misconfigured in the header, for instance). Adding limits on the number of half-open connections protects against SYN floods (a DoS attack). Further rate limiting can be configurable, if needed, or the organization might need the cooperation of its upstream to prevent the limited incoming bandwidth from being consumed by DoS traffic (that cooperation will likely be greater if you can describe the offending traffic carefully based on your logs).

Another external connection managed by the router/firewall might be VPNs for remote users. Preshared keys are the only practical means of tunnel authentication unless this small network is a branch of a very large organization that operates a CA. User authentication for the VPNs is made via an authentication server inside the Campus module (often the management server).

The router/firewall filters traffic ingressing on its DMZ and inside ports as well. No server in this DMZ should need to initiate an outgoing session, nor, in all likelihood, should it need to initiate a session with any other server. If this organization operated an e-commerce module, there might be such a need, but it does not. Filtering such server requests prevents one compromised server from compromising another, just as private VLANs prevent traffic from traveling without Layer 3 filtering.

The router/firewall also filters traffic from this network to the outside according to RFC 2827 and RFC 1918. Apart from being good Internet citizenship, this filtering demonstrates intent to essentially quarantine problems that do develop. This makes any negligence claim by an aggrieved party elsewhere harder to prove.

Threats Mitigated

We have already listed the threats. Table 10.1 presents the threats with the technology used to mitigate each of them, to help you make the connections.

Table 10.1 Small Network Edge Threats and Their Mitigation

Threat	Mitigated By
Packet sniffers	Switched network, HIDS
Network reconnaissance	HIDS, protocol filtering
IP spoofing	RFC 2827 and RFC 1918 filtering at ingress (and by ISP)
Trust exploitation	Restrictive trust model, private VLANs
Unauthorized access	Filtering at the router/firewall, AAA
Password attacks	Strong password policy, OS lockdown, IDS detection
Port redirection	Router/firewall filtering, HIDS
Application-layer attacks	HIDS, OS and applications locked down
Virus and trojan horse attacks	Antivirus on every host
Denial of service (DoS)	Rate limiting at edge (preferably by the ISP), TCP setup limits

As you can see from Table 10.1, not all mitigation techniques involve Cisco equipment or software features. Some, such as locking down the OS (three failed login attempts lead to a locked account for 20 minutes is one aspect of a lockdown) and a strong password policy are system administration and management functions. However, they all contribute to the mosaic of a secured network, even for a small business.

Design Alternatives

One alternative already is built into this design: the choice between a router with an integrated (software) firewall and a dedicated firewall appliance. A firewall has fewer external connection options (usually requiring Ethernet input), but this could work well in a small business with DSL or local broadband connectivity.

The design alternative discussed in the SAFE SMR Blueprint is to add a VPN concentrator if many VPNs must terminate here. This could occur in a medical or legal practice, in which several professionals connect from home or when traveling. When the numbers grow large enough to justify a concentrator, however, the size of the overall network begins to look more like that of the midsize network (covered in Chapter 11, “The Medium Network Implementation”).

The Small Network Campus

The small network’s Campus module has three types of hosts—corporate servers, a management server, and users. Traffic is distributed among them by a switch.

Assets

The lowest-value assets in the campus are probably the users’ systems (if the organization practices centralized file management—if not, security is only one of its problems). Theoretically, few corporate information assets should primarily reside on any user’s system. What might be found there, however, are working copies and temporary files that contain sensitive information, and those should be protected. That is made difficult by user behavior, which frequently weakens the security posture of the systems. Users download or copy in files, executables, and assorted utilities, not all of which are benign; some are actual malware, and none is likely to be known about by the person responsible for security.

The primary copies of the important data should reside in the corporate servers. These are not just document files, of course, but also databases, the email server with its archived copies of correspondence, the business's financial records, and so on. If the company provides services, detailed data concerning its customers and the contracts with them could be here. If the company produces tangible goods, product specifications, testing results, safety data, and so on could be here. Depending on the product or service, legal requirements might need to be met, such as those of HIPAA, GLBA, and/or the Sarbanes-Oxley Act in the United States.

All the information stored in these servers must be considered among the most valuable information assets the organization possesses. If this is a branch office of a larger organization, these servers probably have trust relationships with their corollaries at the corporate headquarters, increasing their value to a hacker.

The management server in this module provides AAA, logging, the IDS Director function, and general configuration management. This is a prime target for the savvy hacker because, if it can be compromised, a user account can be created to allow entry, with high (even root/Administrator) privileges on every device in the network. Likewise, logs stored on this server can be "edited" to remove signs that the hacker was ever there. If the hacker can recognize the IDS, he can remove the signature file for a particular attack before it is begun—resulting in no alarm from the IDS. This server protects other resources, so it must be protected at least as well as the highest level of data that it protects.

A switch is the usual campus distribution device; most small networks do not require the network segmentation provided by a router, at least for traffic-management purposes. Again, be sure that you have a switch capable of supporting private VLANs.

Threats

As when discussing the edge, we have alluded to some of the campus's threats in the previous discussion. However, it's important to take a little different perspective here: Most of the threats to the edge are likely to originate somewhere outside the organization's network. This certainly reinforces the need for a firewall in the edge (whether software or a dedicated device) that is capable of handling the load. You cannot afford to reduce its security profile.

The more likely threat in the campus is an internal threat, someone authorized (at least, authorized at some time) to use a system on the network. This person starts from a position that the hacker must work to achieve—having a trusted account on the system. From there, the hacker can work at escalating his privileges. In addition, the internal threat starts with some knowledge of the system that the hacker must acquire through reconnaissance.

Even given this stronger human security threat, the technical threats in the campus are somewhat fewer:

- Application-layer attacks
- Packet sniffers
- Port redirection
- Trust exploitation
- Unauthorized access
- Virus and trojan horse attacks

Not present in this list are password attacks (because new user accounts likely will be created or existing account privileges will be escalated), network reconnaissance (not needed), IP spoofing (not needed), and DoS attacks (possible, but unlikely). Note that the ones missing are also ones that might be easier to spot if you are reading logs and have filtering in place.

Devices and Implementation

That naturally leads us to look at how to protect the devices in the campus. All hosts, both users and servers, should have an antivirus software installed, should be updated regularly, and should have the systems scanned frequently. Likewise, every host should be fully maintained (fully patched), in both its OS and its applications. Configuration control must be practiced to ensure that users do not install their own additional applications and utilities—sometimes in the interest of productivity and more often in the interest of convenience.

In addition, both corporate and management servers should have HIDS installed and tuned to operate very aggressively: Although false positives in this environment are indisputably annoying, they are better than false negatives. This is likely to be a more aggressive HIDS posture than you used in the edge, where the public, rather than your internal users, would be inconvenienced.

The switch should be configured to have all servers on private VLANs. In addition, certain user workstations, such as administrative workstations, might benefit from this. As in the edge, all unused ports should be disabled to prevent unknown and unauthorized devices from being able to connect to the network. In addition, ports that need to trunk should be specified, and all others should have their trunking autonegotiate set to Off.

As you might have noticed, most of the security work to be done here is system administration rather than networking. However, this area should not be penetrated by outsiders if the edge is well secured (although there can never be any guarantees). Managing insiders largely is a system-administration function.

Threats Mitigated

Having described the threats and the security measures to be taken, it's time to summarize them in Table 10.2.

Table 10.2 Small Network Campus Threats and Their Mitigation

Threat	Mitigated By
Packet sniffers	Switched network, HIDS on servers
Trust exploitation	Private VLANs, restrictive trust model (where appropriate)
Unauthorized access	HIDS, strict authorization control on applications
Port redirection	HIDS
Application-layer attacks	OS and applications locked down, HIDS
Virus and trojan horse attacks	Antivirus on every host

Even though there are many similarities between this table and Table 10.1, there is more of a system administration orientation in Table 10.2, just as there was with the configuration section. Because this is the heart of the operation, much data exchange must be permitted to get the work done. A restrictive trust model, for instance, can be more aggressively applied in the edge than in the campus because of the differing requirements for data exchange. Note that you should still use the restrictive trust model in the campus as much as possible. Strict authorization controls and the aggressive use of HIDS are required because there is no Layer 3 device in this model to filter traffic.

Design Alternatives

The design alternative in the campus is to add Layer 3 filtering (without this, you must work harder in other areas). This could be done with a router or a small firewall (placed to isolate high-value assets), which might require a second switch for traffic distribution on the isolated segment.

Branch Versus Standalone

In the first paragraph of this chapter, I mentioned that the small network could be a standalone organization, complete unto itself, or a branch of a larger organization. If it is complete, everything that we have described could well apply. However, if it is operating as a branch under a headquarters at the end of a WAN link, a couple of things might change. In the branch's edge, there would probably be no need for incoming VPN termination, although there might well be a VPN connection (site to site or LAN to LAN) to the headquarters. In the branch's campus, the management hosts (both server and any workstations) would probably not exist: They would be provided at the headquarters, and management of the branch's devices would be performed over a VPN tunnel.

Summary

That is the small network: an uncomplicated edge that must nonetheless be thoroughly secured, and a simple campus, which might not even have Layer 3 traffic management. Though small in scale, it is enough to give us a handle on the process of putting all the pieces together into a coherent secured network design secured according to the organization's policy and technology choices. In Chapter 11, we expand our view by going through the same exercise with the midsize network, which follows the same principles but differs in the details.