12

The Remote-User Design

Terms you'll need to understand:

✓ Software access option

✓ Remote site firewall option

✓ Hardware VPN client option

✓ Remote site broadband router option

✓ CRL

Techniques you'll need to master:

✓ Performing threat analysis against VPN services

✓ Evaluating remote-user connectivity needs

Chapter 10, “The Small Network Implementation,” and Chapter 11, “The Medium Network Implementation,” dealt with fixed locations, where everything is (nominally, at least) under the control of the IT organization, even if that is outsourced. When we look at securing remote users, however, we face a different kind of problem: Sometimes although the principal information asset is the organization’s property, it is not under the organization’s actual control. When that is the case, problems can develop in the remote user’s host and can travel from that host into the headend network, where they can cause even greater problems.

As an example of what could happen, several teleworkers were disabled by the recent MSBlast worm because they were outfitted with hardware and software, but were not allowed to update their configurations (they were not allowed to have Administrator rights). Their IT organization never got around to patching the remote users. The teleworkers became infected when they connected to the Internet just to log in to work. This particular worm had such a rapid cycle (rebooting its host every 60 seconds) that they weren’t able to infect their headend networks—but that was not as a result of anything the IT did properly. Remote users can be a great benefit to an organization, but they must be kept as current in all respects as though their hosts were permanently in the LAN—or more so because they are more likely to be exposed to trouble.

The Remote-User Problem

Securing remote users is clearly both a priority and a problem whose solution depends on people who often do not understand IT procedures or reasons for doing things a certain way (a justification for no Administrator rights). The SAFE SMR Blueprint offers the same four solutions (called options) as the SAFE VPN Blueprint. Because they are alternative solutions to one overall problem, we look at the problem from the same four angles used in the previous two chapters. Then we look at how each of the options (the design alternatives) solves the problem, as much as it can be “solved.” The four angles are as follows:

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

> The remote-user network assumes a need for communication between physically separated entities, the headend and the remote user. *Their communications must be as much as possible like the communications that would occur if they were located on the same campus and were working in the same LAN.* Subject to bandwidth limitations, of course, that is your goal, provided that the communication is secure. Security comes first; after that, make it as much like it would be inside the LAN as you can.

The remote user (or users—we just use the generic term *user* even when referring to a site-to-site connection) typically connects to the headend to access corporate resources—email, databases, documents, spreadsheets, and so forth. These files travel across a public infrastructure, reside locally, are manipulated, and return (when being saved) to the headend via the same public infrastructure.

At the same time, because people multitask, the user might need or want to browse the Web, exchange personal email, engage in IM chats (which might or might not be work related), check on newsgroups, and do many other typical online activities. Engaging in these activities while the connection to the corporate LAN is open provides a vector for malicious software or an unknown connection to piggyback its way in.

The remote-user model therefore has to concern itself with two separate but related problems:

- The security of the connection between the two endpoints
- The security of the remote hosts (the uncontrolled endpoint), lest those hosts provide an ingress path for trouble

Protecting both of these is necessary.

Assets

As mentioned, the solo remote user might have locally stored copies of corporate information as well as temporary files, which could be recovered with a little patience and access (not necessarily local access) to the hard drive. In a branch office connected via a site-to-site LAN, there will be hosts and a network connectivity device (of some sort), and there might be local servers with email, financial, and other business files stored on them. In other words, the assets to be secured are hosts, possibly servers, and networking devices. That is essentially the same situation in the small business network, except that the numbers might be smaller.

Threats

The remote-user network faces a somewhat different set of threats than the small and midsize models:

- IP spoofing
- Man-in-the-middle attacks
- Network reconnaissance
- Virus and trojan horse attacks
- Unauthorized access

You could argue that denial of service (DoS) also belongs in this list, but there is little that anyone at the headend or the remote location can do about it. Only action by the ISP upstream of either can have any effect on DoS blocking the VPN. The SAFE Blueprints are practical; DoS mitigation in this instance is not practical, so it is not included.

As for the threats that were included, they principally reflect the public infrastructure character of the communications. At any point between the headend and the remote user, a hacker can interpose himself. A hacker can gain access to either end of the connection via IP spoofing—presenting an address that will be accepted because "of course it's ours." Man-in-the-middle attacks require the capability to interpose a system into the communications path; although this is difficult, with access to a compromised ISP (for instance), it can be done. Network reconnaissance can start with observing the existence of the VPN; if there is encrypted traffic between these two endpoints (a sniffer somewhere between them will facilitate that discovery), there must be something going on worth probing. The reconnaissance now has the two endpoints known and can attempt to probe whichever seems weaker (almost certainly the remote-user end instead of the headend).

Virus and trojan horse attacks are always a threat. As I write this, on my home network, my firewall logs show an increasing number of probes on port 17300, which is the port used by a particular trojan (sometimes known as Milkit) that exploits the back doors left by two other well-known trojans. My home network is essentially stealthed: It responds to external probes on no ports (any that I can't specifically disable are redirected to a bit bucket). Yet my public IP is routinely probed on known trojan ports; your remote-user network will be probed as well. Viruses arrive via email, of course, but also (increasingly) via IRC and IM message exchanges. It is not unreasonable to expect them to arrive via any kind of communication (one of the concerns with converged networks, if you recall, was the cross-connection vector path—viruses and other malware intended for one network type can enter via the other).

Finally, unauthorized access remains a problem wherever a host of any sort is located, physically or logically. Because physical control of a remote host is more easily lost than that of a host inside the headend, it is important to have the capability to readily disable any remote host's access to the headend at any time. You can lose control of the remote host in any number of ways, from compromise via its interconnection to another network (a teleworker who also has a home network), to laptop theft. You must not allow that loss to create an opening for anyone to access the headend.

Devices and Implementation

In the headend of whatever size, the following functionalities are present in some form:

- A firewall and/or router for traffic filtering
- A switch to direct traffic to the correct host
- A means of authenticating traffic flows
- Isolation of public-facing hosts from internal traffic

You need to replicate those functionalities in your remote-user network to secure the hosts and the communications path. Therefore, you need to provide firewalling/traffic filtering, traffic direction (to the appropriate host), authentication, and segregation of Internet or other public-network access from the path between the remote user and the headend. Depending on the situation at the remote end, several technology combinations can accomplish this. We address them in each of the options.

Threats Mitigated

As in previous chapters, we have already listed the threats. Table 12.1 lists the threats and the technologies used to mitigate them, to help you pair them up. A smaller network is used in this model, so we need only match the threats and their mitigations (we need not note to which part of the network a threat applies).

Table 12.1 Remote-User Network Threats and Their Mitigation

Threat	Mitigated By
IP spoofing	RFC 2827 and RFC 1918 filtering at headend ingress and at remote site ingress
Man-in-the-middle attacks	Traffic encryption and content integrity validation
Network reconnaissance	Protocol filtering at the remote site
Virus and trojan horse attacks	Antivirus on every host
Unauthorized access	Filtering at the ingress router/firewall or personal firewall software on the host

Mitigation in the remote-user network is a little different from that of the headend, especially when it comes to protecting against unauthorized access. AAA at the headend protects hosts and services there from unauthorized access via interposition in the connection (a man-in-the-middle attack), but it does not necessarily protect the network if the host at the remote end of the tunnel is compromised. (Remember, most user logons offer the username; the person attempting to log on need only enter the associated password.) That's why the remote-user model must focus so strongly on the remote hosts and the device that filters traffic at their end.

The Four Options

The four options listed in the SMR Blueprint for this model are as follows:

- A software access option
- A remote site firewall option
- A hardware VPN client option
- A remote site broadband router option

The options should look familiar because we discussed them briefly earlier. Here they are again, in Figure 12.1.

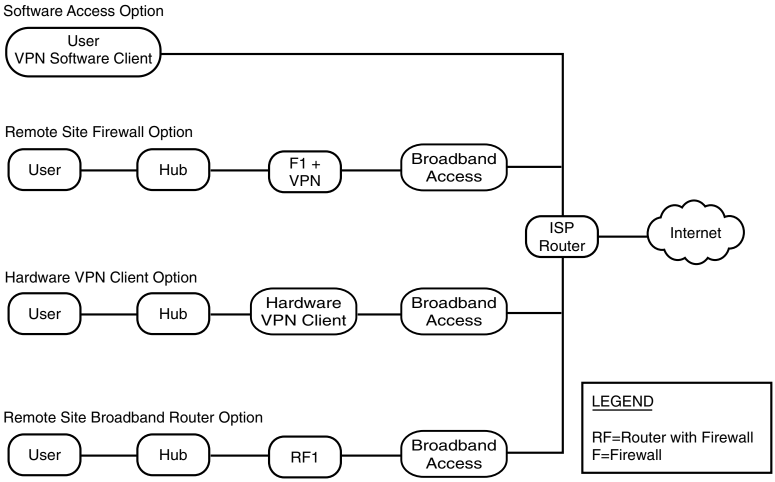

Figure 12.1 The remote-user network options.

The Software Access Option

This is simultaneously the simplest option (topologically) and the most dangerous to the organization. There is no additional hardware between the user's host and the Internet access. This access can be dialup or broadband (while traveling, for instance). All safety precautions for the host must be present on the host. These include antivirus software, regular software (OS and application) maintenance, and a personal (software) firewall. The software firewall must perform all filtering (for IP spoofing and unauthorized access). The software VPN client authenticates the host to the headend and provides IPSec termination.

After authentication and the pushdown of policy settings (such as access rights), the remote host can receive a virtual IP address from the headend's block and the addresses of DNS and WINS servers. In other words, it operates as though it were part of the LAN. The headend can enable or disable split tunneling during this configuration pushdown; the SAFE SMR recommends disabling it. This makes sense when you consider the rather direct path from the Internet through the host to the corporate network when it is enabled (if the personal firewall settings were incorrect, for instance).

The Remote Site Firewall Option

With this option, the incoming broadband enters through its access device (often called a DSL or cable modem) and then immediately passes through a stateful firewall. From there, the data flow passes through a hub (or switch) to the actual host. This configuration is suitable for a teleworker or a small branch office with few hosts.

Threat mitigation on the host need be only the (usual) antivirus software and maintenance of the OS and applications (patching). The bulk of the threat mitigation occurs in the stateful firewall, which provides both the filtering functions and IPSec tunnel termination. Address filtering should mitigate any attempted IP spoofing inbound to the remote user, and the protocol filtering limits any network reconnaissance. The two types of filtering together mitigate unauthorized access. If desired, split tunneling is more reasonable in this configuration because the stateful firewall segregates the connections that it passes. IP addresses used behind the firewalls at different remote sites should not overlap unless NAT is being used on their firewalls.

The second principal function of the firewall, IPSec tunnel termination, provides secure, encrypted communications to the headend, mitigating man-in-the-middle attacks (and ordinary snooping).

> Because IPSec includes per-packet authentication as well as encryption, any interposition by a hacker requires a much more sophisticated handling of the communications stream. No security professional will say that this is impossible, but one will remind you that it is beyond the capabilities of anyone but the most sophisticated hackers and government agencies. It is safe to say that using IPSec ("encrypting the traffic" is often given as a shorthand for this) will protect against a man-in-the-middle attack.

Configuration of the tunnel provides access control and authorization into the corporate network for the users behind the firewall. The remote firewall and tunnel configuration parameters can be managed via a tunnel from the headend, again alleviating end users from needing to perform any configuration tasks (and making their authentication and authorization subject to better control from the headend). Note that the firewall provides device-level authentication to the headend but does not itself provide any user-level authentication.

The Hardware VPN Client Option

The third option looks much like the second, except that the firewall has been replaced by a hardware VPN client. Although this device is optimized for setting up and managing the VPN connection, it lacks some of the other (quite useful) features of a true firewall. It does not provide stateful connection management, so a personal firewall on the host (on every connecting host, if there is more than one) is needed along with the antivirus software (and software maintenance). Disabling split tunneling would somewhat alleviate the need for a personal firewall by forcing all Internet connectivity to travel via the tunnel. Unfortunately, using the tunnel does not prevent foolish downloads from being acquired and placed on the hosts; the software firewall is recommended even if split tunneling is disabled.

Again, all connection management can be done from the headend, although the hardware client uses an SSL tunnel for this purpose (remember, its management is GUI-based, using a Web browser). Access and authorization are centrally controlled at the headend, but, as with the firewall option, the device is authenticated, not the user who happens to be accessing the headend.

The Remote Site Broadband Router Option

This option is very like the firewall option. The differences lie in the differing capabilities of a router versus a firewall. With a software firewall as a part of the router's software, stateful connection management is available, but you also gain the possibilities of the other router software features. This includes the capability to handle other protocols and QoS.

The router also might be the broadband access device (these are sometimes sold or provided by ISPs that are differentiating their service). In this case, you will not likely be able to manage the router from your headend. In such a case, you will need some other termination device (including possibly the software client) for your tunnel.

User Authentication

You might have noticed that these configuration options all lead to the host device (computer, firewall, hardware client, or router) being authenticated at the headend (most likely a VPN concentrator). The tunnel endpoints, not the person who is using the tunnel, are authenticated via this process. User authentication remains the function of the same means by which it is normally done inside the main (headend) network: most often with a Windows domain logon, which can be the next step after the tunnel is authenticated.

In this respect, we come back to the point made earlier: The communicating experience for the users must be as much like the experience they would have locally as possible, but security comes first. In a sense, the tunnel is replacing the bootup and initial network connection to the LAN. Of course, because this occurs over an insecure public infrastructure, a secured tunnel is necessary to protect both the remote user and the headend from malicious parties anywhere in the communications path between them.

Centralized Management

Another item that you might have noticed was the emphasis on centralized (headend) management of the remote devices and the tunnel configurations. An IPSec tunnel depends on matching authentication during the IKE phases (to include using the same preshared key or validating certificates) and then being able to match IPSec configuration criteria, such as the encryption and authentication algorithms, choice of the same traffic to encrypt, and so on.

Putting things bluntly, if IT finds it difficult to trust users to handle the Windows Update function, it is overwhelmingly likely that it does not want to handle a help-desk call trying to troubleshoot what changed in the tunnel parameters. ("I didn't change anything! I swear, it just stopped working!") Whether the tunnel termination is the VPN software client, the VPN hardware client, the firewall, or the router, whenever possible, the SAFE Blueprints assume centralized management of all tunnel-configuration parameters.

Besides reducing the headache level, another reason for centralized management has to do with handling things when a tunnel authentication must be revoked. If a laptop is lost or stolen, if an employee is terminated for copying the firewall parameters (more on how that was addressed and corrected comes later), or for any other reason, when a previously valid tunnel can no longer be accepted, it is easier to manage that from the headend than to try to talk a remote user through the changes to make in the local configuration.

If certificates are being used, it is as simple as adding the certificate to the Certificate Revocation List (CRL), which should be checked before any certificate is accepted as valid. If a preshared key has been compromised, it is faster—and more secure—to immediately establish a tunnel to the remote device and modify the configuration, clear the SAs, and re-establish a new connection under the new key.

Summary

The remote-user network offers an opportunity for an organization to expand its operational reach and enable its people to work from where they choose. At the same time, it presents a potentially easy route for a hacker to enter the main network if a remote host can be compromised. Add to that the fact that the communications between the remote host and the headend travel over a public infrastructure and must be secured, and you begin to realize that supporting remote users is what one wag called "an insurmountable opportunity."

The remote-user model of the SAFE SMR Blueprint offers four options—which are really four design alternatives, by another name—to satisfy the need to emulate the inside-the-LAN experience for users as much as possible while securing their hosts and the communications path between them and the headend. These four options provide filtering, traffic direction, authentication, and traffic isolation in different ways. Because no one design ever fits all situations, you should know how each option does this so that you can make a reasonable design choice when faced with the need to support remote users.

Pulling It All Together

We have now covered all the elements you need to understand to pass the Cisco SAFE Implementation (CSI) exam, Exam 642-541. This is an exam that is usually the last in a five-exam series to earn the Cisco Certified Security Professional (CCSP) designation. To earn this, Cisco requires you to hold a valid CCNA and have passed these other exams:

- SECUR (Securing Cisco IOS Networks) or its predecessor, MCNS (Managing Cisco Network Security)
- CSPFA (Cisco Secure PIX Firewall Advanced)
- CSIDS (Cisco Secure Intrusion Detection System)
- CSVPN (Cisco Secure Virtual Private Networks)

The CSI exam pulls all the pieces together into a single unit covering how to design and implement a secured network using Cisco security technologies and products. It assumes that you already have a substantial level of networking and security knowledge.

However, there is no requirement to take Cisco's exams in any particular order; you might not yet have taken any of the exams listed previously. To maximize your chances, and to provide a review for those who have already covered that material (though perhaps not recently), we began with a three-chapter review of how to identify your information assets, what the threats to them are, and what a security policy does to help organize your thinking about both of these (remember, the SAFE Blueprints all assume that a security policy is already in place and either is being or will be adhered to).

We followed that with a chapter on network-management protocols and how they work, with a particular eye to their security strengths and weaknesses. When you consider that much of the security in a secured network results from conscious management of what is allowed to happen and what traffic is allowed to flow, the significance of the management protocols (and especially secured versions of them) becomes apparent.

We spent two chapters after that taking a high-level view of the various SAFE Blueprints. The first one, and the one that set the pattern for all which followed, is the Enterprise SAFE, suited to a large organization or one that engages in e-commerce. Another one that has great bearing on this exam is the SAFE VPN Blueprint, which goes into considerable detail on how to secure communications between separated locations. The IP Telephony and Wireless SAFE Blueprints are newer and not substantially relevant to the exam, so we did not spend too much time on them. We did pull together the fundamental concepts of the Enterprise SAFE—minus the e-commerce and resiliency aspects—and the VPN SAFE to find that these concepts are the basis of the SMR Blueprint, which is the focus of the CSI exam. The assumptions, axioms, design fundamentals, and design alternatives for these three key Blueprints (Enterprise, VPN, and especially SMR) are things that you need to know to pass the exam.

The exam is also about implementing the SAFE Blueprint with Cisco products, so we spent the next two chapters looking at the major products employed in a secured network: switches, routers, IDS, AAA, firewalls, the VPN concentrator, and the VPN client (both hardware and software). You should be able to configure each of these in a simulation using the preferred method (CLI for a router, switch, PIX firewall, and IDS, and the GUI for AAA and the VPN devices).

At that point, after covering the background, the overall design ideas, and the products to make the designs work, we spent a chapter each on the details of the small, midsize, and remote-user networks—the SMR Model. None of these designs is cast in stone; all have alternatives that you should know.

You are almost ready to take the exam, but before you do anything for real, it's always a good idea to practice it a few times. The next four chapters give you a chance to do just that: They contain a pair of practice tests and a pair of answer sets with explanations. They are as much like the actual exam as I can make them without violating the nondisclosure agreement. Review this chapter and then go through those practice tests to see how well you know the whole picture.