The Extended SAFE Blueprints

Terms you'll need to understand:

- ✓ CA
- ✓ PPVPN
- ✓ Split tunneling
- ✓ Headend
- ✓ XAUTH
- ✓ PFS
- ✓ SA
- ✓ PMTUD
- ✓ MODCFG
- ✓ DPD
- ✓ Stale SAs

Techniques you'll need to master:

- ✓ Implementing a SAFE small network design
- ✓ Implementing a SAFE medium network design
- ✓ Implementing a SAFE remote-user network design

The Enterprise SAFE Blueprint is written for large enterprises and those that have an e-commerce operation. The Extended SAFE Blueprints cover smaller organizations (the SMR SAFE Blueprint) and specialized operations (VPNs, IP telephony, and wireless). The CSI exam description states

> The Cisco SAFE Implementation exam tests the knowledge and skills needed to use and implement the principles and axioms presented in the SAFE Small, Midsize, and Remote User (SMR) White Paper. Candidates are tested on knowledge of how the following devices can be used to create a complete end-to-end solution: IOS routers, PIX Firewalls, VPN Concentrators, Cisco IDS Sensors, Cisco Host IDS, and the Cisco VPN Client.

The Enterprise SAFE Blueprint provides the foundation for the design principles used in the SMR and the other SAFE Blueprints. But the SAFE SMR adds a twist that is not in the Enterprise SAFE: remote-user networks. In fact, the SMR is a cross between the Enterprise SAFE (without resiliency and e-commerce) and the SAFE VPN Blueprints. Therefore, we need to look into the VPN Blueprint before we put those pieces together in the SMR Blueprint. Because Cisco might modify the exam between the time this book is published and the time you actually take the exam, we look (lightly) at the specialty SAFE Blueprints, the IP Telephony, and Wireless Blueprints.

To be sure you have it fresh in your mind when looking at products in the modules (Chapters 8, "Products in the Campus," and 9, "Products in the Edge"), we take things in this order: VPN (to cover the second necessary major topic), IPT, and wireless (just in case questions are added concerning them). We close by pulling together the Enterprise and VPN concepts into the SMR Blueprint. By the time you see the ideas in that section, you should be getting them for the third time or even more. Repetition helps, especially when you are under stress from taking the exam.

The SAFE VPN Blueprint

The actual title of this Blueprint is "SAFE VPN: IPSec Virtual Private Networks in Depth" (notice the focus on IPSec VPNs instead of VPNs in general). VPNs are both a major advantage and a major disadvantage for any organization using them, from a one-physician or one-attorney practice, in which the professional sometimes works from home, to a major enterprise with hundreds or even thousands of teleworkers. The advantage lies in letting people work where they are comfortable, when they are traveling, or where (bluntly) the enterprise doesn't have to pay rent and overhead.

The disadvantage is security.

Because there are plenty of VPNs out there, we must conclude that perhaps the advantages outweigh the disadvantages, or security is manageable, or some combination of these holds true. In fact, the last choice is probably closest to the truth for most organizations. Part of the security problem is that not all users connect in the same fashion: dialup, broadband from behind a home router, broadband from a hotel, a branch site with a firewall/router, and so on. The key is getting the security problem to be manageable—and that reminds us to use a modular approach. The SAFE VPN Blueprint separates the different kinds of VPN connections and develops the security to protect each kind. To understand how and why the module choices are made, though, we need to look at the design fundamentals and the axioms.

Design Fundamentals

The SAFE VPN Blueprint is based on the VPN world as it now exists, with a nod to what is developing. Specifically, most organizations do not yet run their own Certificate Authority (CA) to fully implement CA-based authentication. Therefore, even though using certificates—in a properly configured CA-based network, of course—is more secure, SAFE VPN does not require that choice.

Likewise, there is often a choice to be made between setting up and managing your own VPNs or contracting with your service provider for VPNs: provider-provisioned VPNs (PPVPNs). The latter are discussed in RFC 2764. This blueprint, however, is about how you should set up and manage your own VPNs, so it does not reflect the particular characteristics of PPVPNs. Nonetheless, should you want to consider them, understanding what needs to be done for a manageable and secure VPN will help you evaluate your provider's offer. Given that these are locally provided and managed, the assumption is that the equipment is yours (or, at least, your responsibility): You are using CPE. Furthermore, the equipment is located at your facilities (and, of course, you are able to manage it instead of having to get someone else's help).

Quality of service is not considered in setting up these VPNs. If you need to use QoS (perhaps you have different information with greater or lesser importance to transport), that traffic segregation is assumed to be done separately. Also, when remote locations tunnel in, split tunneling is an option. In fact, it's a tradeoff that you should understand.

Split Tunneling

Split tunneling is actually a quite descriptive term for what happens on the remote machine doing the connecting to the headend. If the remote machine makes all its connections to the world beyond via the single VPN tunnel, the VPN devices are carrying the entire outbound/inbound load—Internet surfing as well as the business information exchange that the VPN is nominally there for. However, if the business traffic and the nonbusiness traffic can be separated—split into separate tunnels—the VPN devices need handle only the packets carrying business traffic. This reduces the load on whatever is handling the VPN and its characteristics, such as encryption, at both ends of the connection. It also reduces the load on the headend's WAN link (because the nonbusiness traffic need not be carried over that and can travel instead over the ordinary Internet connection of the remote host).

Split tunneling is not a panacea, however: The second connection, almost certainly to the big, bad Internet, has the potential to allow a hacker to create a back door into the enterprise network via a compromised host. Thus, although split tunneling relieves VPN devices of some of their workload, it can negate the entire purpose of having a VPN—a private network connection (even if it's virtual privacy).

Split tunneling might or might not be a good choice for your network implementation. As with everything else in networking, it depends on both the problem you must solve and the resources with which you must solve it.

Likewise, although the idea of dynamic endpoint discovery is attractive (certainly there is less static configuring to do), it relies on at least partially meshed environments, whereas most VPNs need to operate more on a hub-and-spoke topology. Therefore, dynamic endpoint discovery is not used.

Those constraints leave us plenty of room to be creative. The actual design guidelines are listed here:

- Secure connectivity
- Reliability, performance, and scalability
- Options for high availability (HA)
- Authentication of users and devices in the VPN
- Secure management
- Security and attack mitigation before and after IPSec

The point of these fundamentals is to provide private, ubiquitous communications using public infrastructure, with connections to wherever users and their devices might be. The connections should be as much like private WAN connections as possible to the user. And, of course, the connections must integrate into a network secured according to the SAFE Blueprint—it

does no one any good to try to solve a metric problem with English-measure wrenches. The systems have to fit together.

Axioms

The SAFE VPN Blueprint places more emphasis than most on the material in the axioms. In part, this is because the entire VPN universe is evolving rapidly; implementing technologies or techniques can change—hopefully for the better—rapidly. Another reason for the strong dependence on the axioms is that you might implement multiple VPN technologies in one enterprise network. It's better to understand why you do certain things before you do them; when you must compromise with reality (which always happens in real-world implementations), you will know which compromises are acceptable and which are not.

These axioms might not seem terribly relevant because the exam focuses on the SMR Blueprint. Remember the *R* of SMR? Although the necessary material might be covered in the remote-user portion of the SMR Blueprint, it is addressed in more depth in the VPN Blueprint. Some of the Cisco recommendations that result from these axioms could show up as test questions. It will help if you can put things in context and recognize plausible-but-wrong possible answers when you see them.

The SAFE VPN axioms are supported by a great deal of explanatory matter, so we present them in *italics*, followed by some support. First, there is the matter of *identity and IPSec access control*. In all VPNs, devices are authenticated, but users should be authenticated, too, in remote access scenarios (remote users, not site-to-site). Preshared keys are a practical means to use for authentication until the number of devices grows large (more than 20, as a rule of thumb); at that point, you are better off using certificates. Preshared keys can be unique (mapped to an IP address, which will not work for someone whose IP address cannot be reliably predicted), can be a group (associated with a group name that has its associated users), and can be a wildcard (which works for anyone who presents it). The last is not recommended for site-to-site VPNs because multiple hosts might "know" the key and could compromise it. IPSec access control occurs after both device and user authentication (if user authentication is present). It can be based on XAUTH—extended authentication, in which the device is authenticated based on the preshared key and the user is authenticated by challenge/reply—or by crypto ACLs. XAUTH is considered preferable.

The next axiom is simply called *IPSec*, which, of course, covers a great deal of territory. For the SAFE VPN Blueprint, however, we can reduce the scope of the target a bit. IPSec offers data encryption and/or data integrity (validation); Cisco recommends using both. Although DES is supported, 3DES is

much stronger (and DES has proven vulnerable to brute-force attacks by networked PCs). Likewise, MD5 yields a 128-bit hash, while SHA-1 yields a (stronger) 160-bit hash (and is a more efficient algorithm, compensating for the larger calculation). Cisco recommends using 3DES and SHA-HMAC for your data encryption and data integrity. Although it would be nice to use, perfect forward secrecy (PFS) is not necessarily preferred (depending on the sensitivity of your data exchange), nor does Cisco recommend shortening the default security association (SA) lifetime. Of the three supported group sizes for Diffie-Hellman exponentiation to establish keys, Cisco recommends the middle of the three (group 1 is 768 bits, group 2 is 1024 bits, and group 5 is 1536 bits; group 2 is recommended). VPN devices spend a large—perhaps overwhelmingly large—proportion of their processing on the encryption and decryption processes. You must consider the load you're imposing as well as your need for security when you make your design decisions.

The third axiom is short and sweet: *IP addressing*. As much as possible, you should use an addressing scheme that allows you to summarize the VPN addresses into the fewest possible lines for your crypto access lists. That reduces the load on the IPSec devices, especially at the headend, which is simultaneously processing packets from many connections.

The next axiom is almost more of a reminder of operating limitations: *multiprotocol tunneling*. Remember that IPSec supports only IP unicast—no other protocols and no multicasting. When you need either or both of those, Cisco recommends using GRE for site-to-site tunnels and L2TP for remote access tunnels.

Next, we have *NAT*; specifically, NAT can be done before or after IPSec. However, there are implications with each choice. NAT after encryption gains no privacy, and it will interfere with the AH checksum (integrity), though not the ESP checksum (because of the different content of the two headers). PAT can result in a change in the port number, causing a loss of the port number (UDP 500) needed for IKE. (Bear in mind that not all VPN devices allow that port to be modified.) If you must use NAT or PAT after IPSec, Cisco recommends using NAT Transparency. This causes the IPSec packet to be encapsulated in a UDP or TCP packet, which can then go through NAT or PAT as necessary. Of course, this adds overhead, but it might solve the problems that you run into using NAT after IPSec. NAT before IPSec is useful when you must use overlapping IP address blocks on your remote devices. As with any other case of NAT, you might find a problem with those cases when an application embeds IP addresses in the packets (the solution is to run a protocol-aware NAT function).

The next axiom repeats what we saw in the Enterprise SAFE Blueprint: *single-purpose vs. multipurpose devices*. You should base your evaluation on the

dedicated device's characteristics rather than the advantages you might gain from employing a multipurpose device.

A larger discussion revolves around *intrusion-detection*, *network access control*, *trust*, *and VPNs*. We've already discussed IDS and network access control (in Chapter 6, "The SAFE Security Blueprint"); the point in SAFE VPN is that you must rethink these issues when using VPNs because your security perimeter has changed. The problem is, you must decide how much you can trust the many and various other sites connecting to you. For instance, on a site-to-site VPN between two of your own sites, a reasonable trust level might be moderately high, but between your headend and an airport or wireless hotspot, the reasonable trust level will be quite a bit lower. One aspect to consider is where you place your VPN device with respect to the firewall and/or switch running NIDS. A VPN creates a hole in your firewall: You must protect the ingress to and egress from that hole.

The next axiom goes into more detail on *split tunneling*. If you choose to use it, you must protect the remote device from that which could enter it from other connections outside your tunnel. That means not only antivirus protection with frequent scanning, but also a (software) personal firewall. If you are operating split tunneling on a site-to-site VPN, you should use a stateful firewall to manage and track the connections. Disabling split tunneling can result in a serious load addition to the headend, so the SAFE VPN Blueprint assumed that it was enabled—and the security precautions that it requires were emplaced.

It is quite possible that you will see a question on the exam related to split tunneling (or you might not—I cannot offer a guarantee either way). However, this is a tricky point. The Enterprise SAFE Blueprint states, "For example, in SAFE, users are prevented from enabling split tunneling, thereby forcing the user to access the Internet via the corporate connection."

However, the SAFE VPN Blueprint, written by a different CCIE at Cisco Systems, states

"When considering the security risks of enabling split tunneling, it is too easy to conclude that it should never be considered. Actually, disallowing split tunneling creates an enormous load on the VPN headend because all Internet-bound traffic needs to travel across the WAN bandwidth of the headend twice. This use of WAN resources is not an optimal one, and it often leads to the decision to implement the appropriate security technologies at the remote sites to allow split tunneling to occur. In SAFE VPN, remote sites were assumed to have split tunneling enabled unless otherwise specified. If split tunneling were disabled, the designs would not change, but the performance and scaling considerations might change slightly because of the increased traffic load on the headend."

Although it is not mentioned here, the headend also has to handle decryption of the packets that need never have been encrypted in the first place, which is a significant burden. Enabling or disabling split tunneling is a design decision that you must make based on the tradeoff you face between the need for security on the remote user's system and the burden you will impose on the VPN process.

These are not necessarily mutually exclusive positions: Note that in the Enterprise SAFE, *users* are prevented from enabling split tunneling, while SAFE VPN refers to enabling it at *remote sites*. Read the question you get (if you get such a question) very carefully.

Next, the axioms address (if you'll pardon the expression) network architecture: *partially meshed*, *fully meshed*, *distributed*, *and hub-and-spoke networks*. Meshes, as we all know, do not scale. Therefore, even if you have VPN traffic between remote sites, you should use a hub-and-spoke topology if you have many sites. Unfortunately, that requires you to investigate your hardware and software for the headend: Not all vendors' devices support a hub-and-spoke topology (they might not enable spoke-to-spoke communications).

That segues nicely into the next axiom: *interoperability and mixed vs. homogeneous deployments*. Mixed deployments offer the benefit of not all sites having the same vulnerabilities. However, they simultaneously cost you in terms of greater management workload and the possibility of more vulnerabilities that might be exploitable. In addition to these obvious tradeoffs, there is another: Sometimes you must "tweak" devices to make them interoperate; those tweaks could lead to a weaker security stance because of suboptimal configuration. *Caveat emptor* again.

Another axiom related to how networks operate is *fragmentations and path maximum transmission unit discovery* (PMTUD). Packet fragmentation always imposes workload on the receiving device, but with encryption, the problem is worsened: Until the packet is reassembled, evaluation against the crypto map cannot begin. PMTUD uses an ICMP type 3, code 4 message, which you should allow as much as possible to avoid this problem. The alternative is to evaluate the data path, count all the various headers and overhead (all of them, not just those caused by tunneling), and manually set the origin device's MTU to ensure that fragmentation does not occur en route. PMTUD is a much more scalable solution, if you are able to use it.

Speaking of network operations axioms, that is the name of the next one: *network operations*. In this case, though, the meaning is more specific: Many network devices, especially at remote locations, are managed via VPNs. Remember that dynamic crypto maps do not initiate connections; they merely respond to a call from the other end (to put it simply). Therefore, you should use static maps for your device-management VPNs, to ensure that you can make the call rather than having to wait for the remote device to phone home. Of course, static maps require the remote devices to have static IP addresses; and those devices should require authentication and authorization before management activity is permitted. Certificates simplify this, but they require checking against the current time; if you use certificates, use NTP on your network devices to be sure that your times are synchronized. One advantage of using VPNs for device management is the possibility of pushing client configurations and updates from the headend to the remote

sites/users (reducing the opportunity for misconfiguration by well-meaning or not-so-well-meaning users).

Reliability should be considered an aspect of implementing VPNs. The next axiom reflects that: *HSRP* can be used when deploying VPNs using routers. The tradeoff here, of course, is that routers are often less capable (in terms of the number of VPNs they can handle) than dedicated devices.

Compression is often used as an approach to get the best use of bandwidth, but it does not work well with VPNs. Layer 2 compression operates by eliminating repeating patterns of bits (a simplification), but encryption leaves no apparent patterns (if it did, it wouldn't be good encryption). Therefore, using Layer 2 compression in conjunction with VPNs gains essentially nothing. Layer 3 compression might reduce the amount of data requiring encryption (saving processor load), but at the cost of processor load to perform the compression function. Again, little to nothing is gained.

The penultimate axiom is called *remote access user requirements*. This is both a reminder that remote users have the same needs as LAN users—DNS, WINS, a virtual IP address for the intranet—and a recommendation: Use ISAKMP MODCFG to push the information to them during tunnel establishment. This is an extension of IKE, and it provides authorization services to control the remote user's access. Among the authorizations that can be passed is enabling split tunneling.

The last SAFE VPN axiom is known as *high availability*. In this case, HA refers more to high availability of the IPSec tunnels than to routing in general. Tunnel endpoints send traffic without acknowledgement (nor should they expect one—IPSec tunnels operate at Layer 3, and acknowledgements come from Layer 4 and above). If one end of a tunnel goes down, though, traffic can wind up being sent to a black hole (or the bit bucket, if you prefer). If you are using routers for your tunnel endpoints, routing protocol keepalives might handle this for you. However, most high-density VPN deployments use more specialized devices, such as VPN concentrators and firewalls, and the clients do not run routing protocols, either. In this case, you are limited to IKE keepalives. However, Cisco has a proprietary means of making the keepalive process between specialized devices more efficient: the *dead peer detection* mechanism (DPD). In this mechanism, keepalives are sent only over tunnels that have not exchanged traffic within a specified interval. This reduces the use of keepalives to their intended purpose (maintaining state when it would otherwise end) while preventing *stale SAs*, which occur when one end of the connection has maintained tunnel state while the other has not.

During the discussion of high availability (HA), Cisco reminds us that the failure of one element should not cause an overload on another. To that end, Cisco offers the simple example of three headend locations and six remote sites, as shown in Figure 7.1.

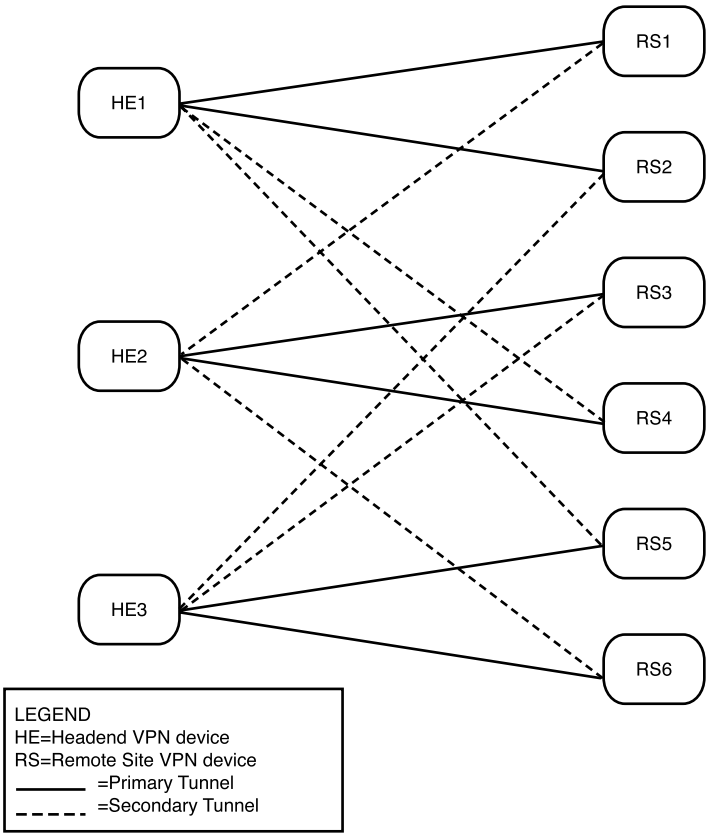

Figure 7.1 Load dispersion on failure of one headend device.

Notice that each headend device supports two remote site devices, but it backs up two, one from each of the other two headends. Therefore, if HE2 fails, RS3 falls back to HE3, but RS4 falls back to HE1. Neither headend's load is potentially doubled while the other remains unaffected. The SAFE VPN reminds us that such fallback, in addition to adding to configuration complexity (surprise), requires an active/active configuration vs. active/standby.

SAFE VPN Network Designs

When we get to the SAFE VPN network designs, they are the same as those of the SMR SAFE Blueprint, so we cover them when we pull the pieces together. For now, we briefly cover the other two blueprints: those for IP telephony and wireless.

The SAFE IP Telephony Blueprint

IPT (IP telephony) is a relatively new technology, at least in production networks (testing and experimentation have gone on for quite a bit longer, of course). Data exchange has had time to develop security protocols that have made it through the standards-development process, but protocols comparable to IPSec in the data arena have not yet been developed for IPT. This is a problem because you could be in a position where there is interest in deploying IPT—provided, of course, that you can do so and deliver ubiquitous telephony service to users and sites that require it, and that it will be (from their perspective) indistinguishable from traditional telephony services. And, of course, you must do it securely, interfacing with a secured data network without creating new security vulnerabilities in that.

Design Fundamentals

The IPT design fundamentals are quite straightforward:

- Security and attack mitigation based on policy
- Quality of service (QoS)
- Reliability, performance, and scalability
- Authentication of users and devices
- Options for HA
- Secure management

The first priority design fundamental is that everything is based on the policy chosen; this is a part of the implementation of that, and, of course, you can't implement effectively if you don't have a good plan. The next highest is the performance make-or-break factor in IPT: QoS. With those two covered, you can then consider reliability, performance, and scalability—the factors that will make this work over a larger network (one which is more than a test lab rollout). Authentication is nothing new, although, again, note that

it is of users as well as devices. It is worth considering options for HA, although you might find that HA is not considered "optional" when it comes to voice communications. When it comes to telephony, HA is part of the requirements, with the discussion limited to means and costs (that is what is meant by the term *options*). Finally, of course, as with every other SAFE network architecture, secure management (and reporting, though not listed) is fundamental.

Axioms

The SAFE IPT Blueprint has axioms, of course, although not with the depth of the SAFE VPN Blueprint. The first one should hardly surprise you:

- Voice networks are targets.
- Data and voice segmentation are key.
- Telephony devices don't support confidentiality.
- IP phones provide access to the data-voice segments.
- Soft phones require open access.
- Soft phones are especially susceptible to attacks.
- Establishing identity is key.
- Rogue devices pose serious threats.
- Secure and monitor all voice services and segments.

The last axiom is also no surprise; the ones between, however, are unique to the IPT environment (remember, the Enterprise SAFE model included IP phones and call managers). In ordinary telephony, the two networks remain separate, so they cannot (without great effort) cross-contaminate one another. However, with IPT, the two networks converge, so a problem can enter via one to threaten the other (and this can go either way). Therefore, it is key to keep the two as logically separated as possible. Soft phones (an unofficial term for PC-based phones) are especially problematic in this regard because the PC is typically not that well protected (there is no IDS, for instance). Disabling unused ports and the "new device-friendly" features of call managers are especially important in keeping intruders out.

The SAFE Wireless Blueprint

Once again, we are working with a newer technology that is only beginning to be rolled out in large implementations and for which security standards are not yet well developed. Nonetheless, basic ideas remain: If you haven't seen a pattern before, the design fundamentals and axioms of the SAFE Wireless Blueprint should help.

Design Fundamentals

The Wireless design fundamentals are also quite straightforward:

- Security and attack mitigation based on policy
- Authentication and authorization of wireless networks to wired network resources
- Wireless data confidentiality
- Access point (AP) management
- Authentication of users to network resources
- Options for HA (in large networks only)

The design fundamentals here revolve around knowing who's on the network and limiting what they are allowed to do: authentication not only of devices, but, again, of users. AP management is the same thing: Ensure that your APs are not open to anyone who happens to have a wireless interface. Authorization applies when it comes to using the wired network's resources. Data confidentiality is important because you are dealing with a broadcast medium. In short, wireless is potentially a wide-open network, which offers an ingress for an unauthorized user into your (main) wired network. That is reflected in the axioms as well.

Axioms

The Wireless axioms are few but are no less important:

- Wireless networks are targets.
- Wireless networks are weapons.
- 802.11b is insecure.
- Security extensions are required.

All networks are targets (surely you've gotten that point by now), but wireless networks can be easier than most for an attacker. When wireless devices must connect through an AP, it is known as infrastructure mode. However, wireless devices can become aware of each other and form an ad hoc network, an informal peer network of wireless devices. If one of those devices is an authenticated device accessing your wired network, you can imagine the consequences. That, coupled with the weaknesses of the initial security implementations (such as WEP, the Wired Equivalent Privacy standard), contributes to hackers being able to use wireless as an attack means. Cisco's recommended security extensions are IPSec, EAP/802.1*x*, and LEAP, Cisco's proprietary extension to EAP.

The SAFE SMR Blueprint

Finally, we come to the one that the CSI Exam is officially about—"SAFE: Extending the Security Blueprint to Small, Midsize, and Remote User Networks." In fact, we've seen the foundation of the SAFE SMR Blueprint in the Enterprise and the VPN Blueprints. The SMR Blueprint scales things down a bit and pays attention to remote connectivity: Where a large enterprise might need to separate things out for manageability, small and medium organizations don't have that luxury. Therefore, the security needs of these organizations as a whole and their remote connectivity are combined in the SAFE SMR Blueprint.

Design Fundamentals

The SAFE SMR design fundamentals will seem familiar, and they bear obvious relationships to the design fundamentals you've seen before:

- Security and attack mitigation based on policy
- Security implementation through the network (not just on specialized 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

These are really very similar to the design fundamentals for Enterprise SAFE, discussed in Chapter 6. A major difference is that SAFE SMR is not

intended to incorporate the resiliency and scalability of the Enterprise model; the focus instead is on cost-effective deployment (and there is no e-commerce in the SMR model). For small enterprises, cost containment is a major issue; by itself, that limits the degree of resiliency that can be implemented. However, the consequence of that lack of resiliency is to make controlling access to critical network resources even more important. It also increases the importance of securing your reporting and network-management functions, along with tuning your IDS as closely as you can manage (to avoid false negatives, even at the expense of dealing with a few more false positives). In short, the design fundamentals of the SAFE SMR Blueprint are almost those of the Enterprise, without the budget or the resources. This actually reflects the case in many, if not most, real networks.

Axioms

The SAFE SMR axioms will sound awfully familiar by now (on purpose—you'll want to know them):

- Routers are targets.
- Switches are targets.
- Hosts are targets.
- Networks are targets.
- Applications are targets.
- Secure management and reporting.

In fact, these are exactly the same axioms as those used for the Enterprise SAFE Blueprint. They were discussed in detail in Chapter 6; if you need to review that, do so now.

Headend vs. Branch Considerations

We've thrown around the term *headend* quite a bit in this chapter, mostly when dealing with VPNs. The concept is simple: When establishing a link between two locations, one location is bigger or more important in the great scheme of things than the other. This is the headend, the entry to the larger network or the greater set of resources. Thus, if you have a large enterprise with a branch campus (which might be designed according to the medium SAFE model), the large enterprise end of the connection between them is the headend, and the branch office is the branch. Likewise, that same branch

might have VPNs to remote users—individuals. For those VPNs, the branch office is the headend and the remote user is the branch.

The SAFE SMR network designs—the small network design and the medium network design—can be used for locations that act as either the headend or the branch, depending on their relationship to the network and users at the other end of the connection.

So what do these SMR networks look like? They're much less complex than the Enterprise model, partly as a matter of scale and partly because of the absence of the heavy resilience in the Enterprise model. Even the medium model can fit into one diagram, as shown in Figure 7.2.

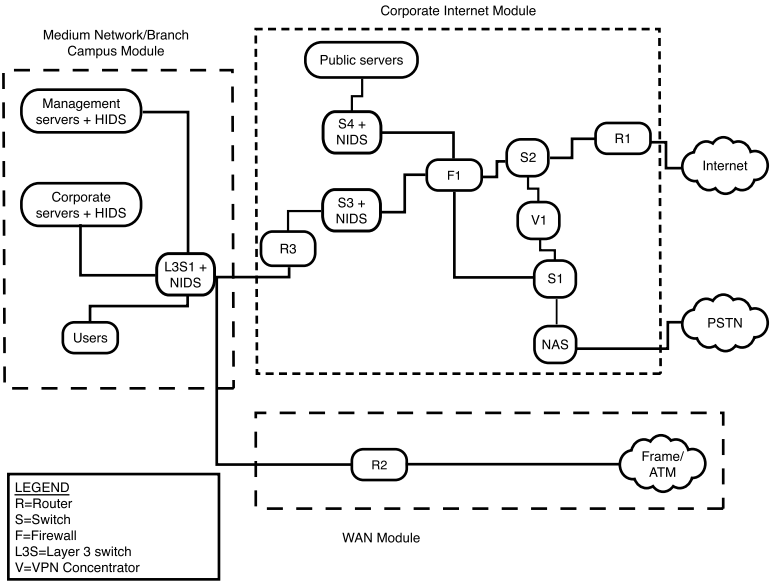

Figure 7.2 The SAFE SMR medium network model.

Working from left to right, you can see that the servers are all protected with HIDS, and the switch connecting this module to the Corporate Internet module is also employing NIDS. (Server traffic is inspected twice because HIDS are often tuned more tightly, given the limited applications present on most servers, even multipurpose ones.) The WAN module is quite straightforward because it was on the Enterprise model: a router to perform basic filtering and traffic forwarding to and from the Frame/ATM cloud.

The most complex module, though much less complex than its Enterprise cousin, is the Corporate Internet module. The public servers are isolated in a "mini-DMZ" that branches off the firewall. Like the public servers in the

Enterprise model, the various servers here are set on private VLANs from a switch with NIDS. The incoming Internet connection passes through a perimeter router and then a switch.

If the traffic is an incoming VPN, it is redirected to a VPN concentrator. After the VPN tunneling header is validated and stripped off, the traffic is switched on into the firewall for inspection. Then (assuming that it passes, of course) it is switched on into the Campus module. If the traffic is incoming public traffic, it goes from the switch to the firewall and then to the public servers. Some VPN incoming traffic might be dialup; that comes in from the PSTN and passes through a NAS, after which it can be switched through to the VPN concentrator or directly to the firewall, as appropriate.

The Corporate Internet module in the SMR model is thus a combination of the Corporate Internet module and the VPN module from the Enterprise model, again on a smaller scale.

The small network is even easier to see, as in Figure 7.3.

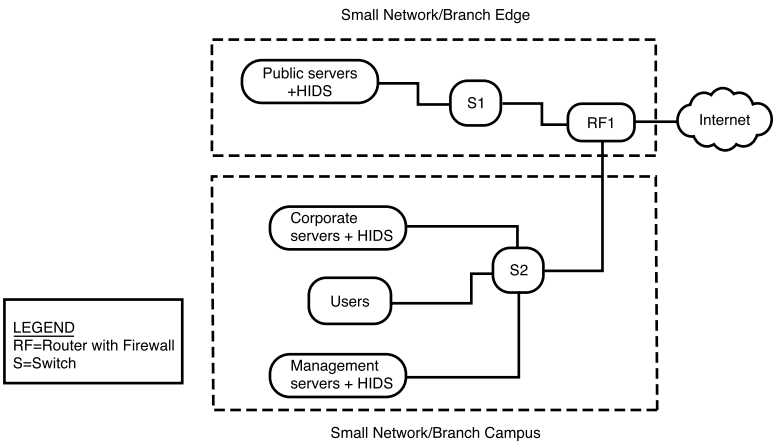

Figure 7.3 The SAFE SMR small network model.

In this much smaller network, you see Internet traffic filtered by a router with firewall software on ingress, where it can be directed to a mini-DMZ, again protected by private VLANs (if multiple physical devices are used as servers). The Campus module is functionally the same as that of the medium model; the only difference is that the switch is scaled back to a Layer 2 switch without NIDS instead of a Layer 3 switch with NIDS.

That leaves us with the remote-user model, which, as we noted earlier, uses the same architecture set as the remote end of the SAFE VPN Blueprint. This set is shown in Figure 7.4.

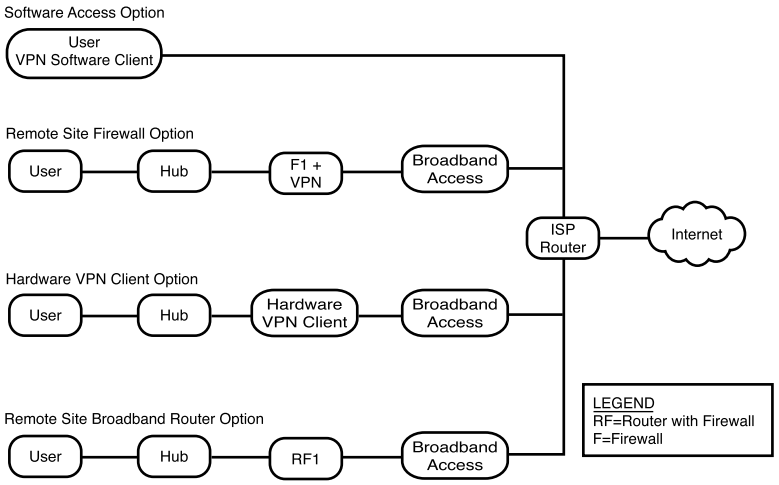

Figure 7.4 The SAFE SMR remote-user model.

In this case, we actually have four possible models, depending on how the remote user accesses the Internet. Working down from the top, the simplest model is when the user connects directly to the ISP, typically dialing into the ISP. In this case, the ISP router might be a NAS, but we know that the client's traffic will enter the Internet via a router, so this is still representational. In this simple case, the VPN software client is the "device" that creates the user end of the tunnel. Note that, with no other device intervening between this host and the big, bad Internet, all protective functions must reside locally on the PC. This host should have a software firewall (also called a personal firewall) along with antivirus software—and both must be kept current.

In the remaining cases, the user is connecting via broadband (xDSL or cable modem service). In the second case, the user's system is protected with a firewall that can terminate VPN tunnels (*caveat emptor* one more time—not all broadband firewalls have VPN-termination capability). However, assuming that multiple systems are using the broadband connection, a hub can be used to distribute the secured connection as needed among them.

Next is a similar situation, except that a hardware VPN client is providing the protection as well as VPN termination. This is a subtle difference—the firewall offers security with VPN, while the hardware client offers VPN plus security. Which is more appropriate depends on the nature of the outbound traffic from the remote user(s) in question.

Finally comes the case in which the protection comes in the form of a router with firewall software, which can also terminate VPNs. This device is often

advertised as a broadband router: It connects to broadband, provides router functionality (including DHCP and NAT), and might or might not offer multiple switched interfaces. If the switched interfaces are not available, you can again use a hub to distribute traffic as needed. Broadband routers, too, often—but not always—come with VPN termination.

Summary

We've now covered all the various versions of the SAFE network designs. Most important to you on the exam is the SAFE SMR Blueprint, which is actually a blend of the Enterprise SAFE and the SAFE VPN models. The other two models, IPT and Wireless, might be mentioned on the exam (or might not be), but they are newer and deal with newer, less settled technologies. Nevertheless, it's a good idea to be familiar with their design fundamentals and axioms.

Next we look at the Cisco products used in the validation of the SAFE Blueprints, starting with the Campus module devices (Chapter 8) and then the Edge module devices (Chapter 9). After that, we finish by putting the pieces back together for each of the SMR networks.