The Medium Network Implementation

Terms you'll need to understand:

- ✓ Headend
- ✓ Man-in-the-middle attack
- ✓ Network topology discovery
- ✓ Password attack
- ✓ Trust exploitation

Techniques you'll need to master:

- ✓ Performing a threat analysis for VPN services
- ✓ Analyzing design alternatives

In this chapter, we do with the medium network what we did with the small network in the last chapter—look at it from these five points of view:

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

In the case of the medium business network, however, we're talking about a larger and more complex edge, although the campus is the same architecturally (even if it has more hosts). The medium network might be a business in its own right (standalone), perhaps serving as the headend for smaller organizations, or it might be a branch operation of an even larger organization (in which case, it is a large branch). Whatever its role, just like the small network, it must be secured as an entity.

Because the medium network is larger and more complex than the small network, you can expect to see more questions on the exam concerning its structure, how it is secured, and what alternatives you might want to implement. As always with Cisco exam questions, the devil will be in the details of the question phrasing.

For review, Figure 11.1 shows the medium business network model in its entirety.

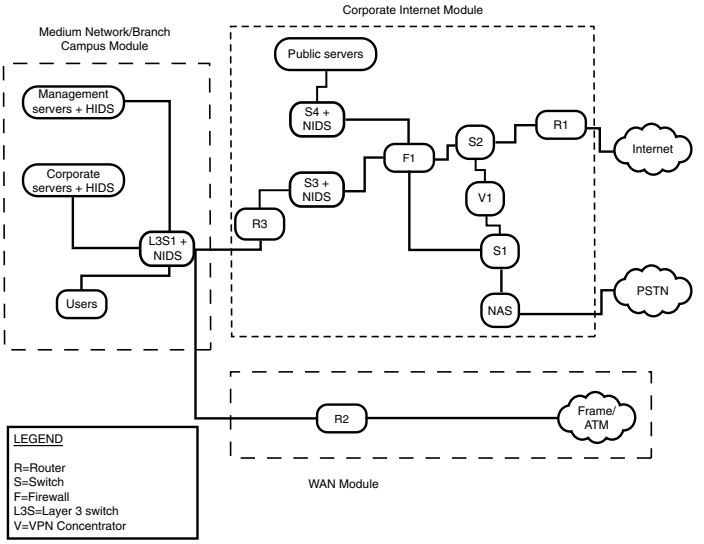

Figure 11.1 The medium business/branch network.

The Medium Network Edge

The medium network's edge can have a WAN module as well as a Corporate Internet module, which is more complex than the one we saw in the last chapter. We focus most of our attention on the Corporate Internet module: It has an Internet connection via a perimeter router and a PSTN connection for dialup data access entering via a NAS. Using switches, both means of access can be sent to the firewall, or incoming traffic can be sent to a VPN concentrator, if it is tunneled traffic, after which it goes to the firewall. From the firewall, traffic for the public-facing servers is sent to a DMZ, while traffic for the organization itself is sent into the campus via a router.

Assets

Assets in this module come in two broad classifications: those for use by the public and networking assets that segregate and sort the traffic (incoming and outgoing). Public-use assets are the servers in the DMZ and the information they contain. These are quite similar to the server set in the small network, but scaled up: There might be redundancy, and there will be a greater capacity to serve more traffic (more CPU, RAM, greater disk size, more information, and so on). Server functions likely to be present include mail, Web, file transfer, and DNS. Some of these functions can be consolidated on one physical system but are redundant across multiple physical systems (especially the Web server, if there is substantial traffic). DNS should not be redundant in this same DMZ because there is a single point of failure in the access via the firewall. The same is true of any other services provided that must be reliable. If a high degree of reliability is truly needed, you must begin to transition to a model more like the Enterprise SAFE Blueprint (with or without the E-Commerce module).

The networking devices in the medium Corporate Internet module include edge routers (one at each ingress to the module), switches, a VPN concentrator, a NAS, and a firewall. The role of these devices generally involves traffic validation and segregation. The switches operate only at Layer 2, so they provide only traffic segregation. The other devices can evaluate traffic based on higher-layer headers and can provide validation and segregation based on that higher-layer information.

Threats

The medium Corporate Internet Module's description divides the threats by their targets. The public-facing servers are likely to be targeted by these

threats, divided as in Chapter 10, "The Small Network Implementation," into low-profile and forceful. The low-profile threats are listed here:

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

The forceful threats against the public-facing servers are as follows:

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

In addition, the more extensive remote-access and site-to-site VPN services face some of these threats, plus two more:

- Man-in-the-middle attacks
- Network topology discovery
- Packet sniffers
- Password attacks
- Unauthorized access

Man-in-the-middle attacks (sometimes called middleman attacks) occur when a third party interposes itself between two parties that are communicating. A successful man-in-the-middle attack is transparent to the two end parties: They have no idea that it is happening. You can see a conceptualization of this in Figure 11.2.

You can see that Alice and Bob, the quintessential crypto couple, believe that they are communicating directly, as shown in the upper half of the diagram. Instead, Fred (a man-in-the-middle) has interposed himself into their channel so that Alice's information goes to Fred, who then communicates something (the same thing? an altered set of information? who knows?) to Bob, and vice versa.

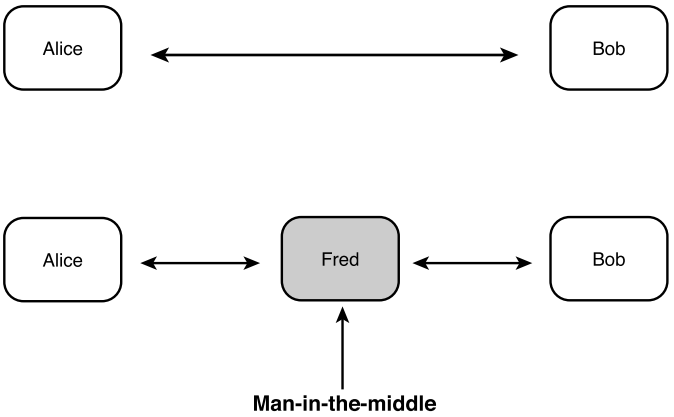

Figure 11.2 A man-in-the-middle attack.

Because you do not control the path between your data's egress and its ingress at its destination, for sensitive communications (including reliability or integrity as well as content), you must protect against these attacks.

The second additional type of attack is network topology discovery, which is not quite the same as network reconnaissance (reconnaissance is more general than just learning the topology). Almost by definition, VPNs carry valuable traffic (otherwise, you would not waste the resources to encrypt it). Therefore, discovering the topology of the VPNs tells a hacker who is communicating with whom and how. With that knowledge, interception and/or disruption becomes easier to achieve.

Devices and Implementation

The server protections in this module will look very much like those in the Corporate Internet module for the small network. As always, a strong antivirus package should be installed and kept current on every host in the entire network, including the servers in the edge. These hosts should also be periodically scanned for viruses and other malware (trojan horses, keystroke loggers, and so on). Each server should be locked down: Only the processes that are required to fulfill its function or functions should be allowed to run, and all applications as well as the OS must be kept current on patches. In addition, unnecessary open ports should be closed (as an example, the NetBIOS ports attacked by the recent MSBlast worm).

Wherever possible, all applications should be configured to accept updates only from specified internal addresses. As described in Chapter 10 (when we discussed sources for zone transfers), this is especially important for the DNS

server because this is how the outside world sends traffic to the public servers. However, it also applies to the other servers. Another lockdown example is the mail server. Locking it down means that, in addition to being current on all patches, it accepts only the minimum set of SMTP commands required to be RFC 821–compliant (`HELO`, `MAIL`, `RCPT`, `DATA`, `RSET`, `NOOP`, `QUIT`). Various mail server packages are generally capable of being so limited; to be truly secure, you might need to use a different package for the mail relay in the DMZ than you use inside the campus.

This degree of detailed knowledge about locking down your servers is not required on Cisco's CSI exam. However, it doesn't hurt to be aware of it for your implementations in the real world.

In addition to being locked down, as with the small network, every server in the medium Corporate Internet module should be running a host IDS; consider that one of the necessary processes. Your HIDS configuration should be aggressive: Expect to send an alarm and to drop offending traffic (sending a reset is debatable, as mentioned in Chapter 10). Obviously, the medium network has more devices whose IDS and logs might report things; you must have someone monitoring and reviewing these, or such precautions will help you understand what went wrong only after the fact—and you'll miss those cases in which you could have seen trouble brewing and did nothing.

The switches in this module are segregating traffic far more than in the small network's Corporate Internet module. This is not so much because of a greater traffic load as it is a greater separation of functionality. There is more likely to be a greater range of types of incoming traffic: Some might need to go directly to the firewall, while other traffic might be an incoming user VPN to be sent to the concentrator. Even on the switches that are not directly connected to the servers, private VLANs are advisable. If you look again at which devices are being connected via switches, these are devices that you do not want incoming traffic to be capable of bypassing. Use of community ports should be limited—even more so for promiscuous ports.

Port management should also be proactive regarding unused ports and trunking: All unused ports should be disabled rather than left available for device connectivity. Trunking should be reserved for links in which you are certain you will not need to force any traffic to Layer 3 for inspection and filtering.

In addition, the switches to the inside of the firewall (switches 3 and 4 in Figure 11.1) have NIDS to monitor (alarm) for intrusions that succeed in

penetrating the outer defense; drops and resets are still left to the HIDS. It is relatively safe to leave the NIDS as alarm-only: Every significant asset is protected by HIDS (NIDS offers forensic value in tracing attack paths and in warning of attacks that HIDS might prevent).

Again, the routers and firewall are the hardest-working elements of the edge, although some of their work is offloaded to the VPN concentrator. The latter is sensible only if sufficient tunnels are being terminated here (a rule of thumb is more than 20 tunnels, but, like all things in networking, it depends). We address these devices in turn from the viewpoint of traffic entering from the Internet.

Ingressing traffic at the perimeter router (or edge router) is inspected for address spoofing (RFC 2827 and RFC 1918 filtering). Fragmented packets are dropped, as they were in the small network's edge, for the same reason. Connection limits to protect against DoS attacks are set; further rate limiting can be configured here, but (if possible) it is better to have the upstream do this and not waste bandwidth on traffic destined for the bit bucket.

One further item to be configured on the perimeter router is to ensure that the incoming VPN tunnels can connect to the VPN concentrator or firewall by allowing passage of IKE traffic (UDP 500) and ESP (protocol 50) or AH (protocol 51). If a tunnel is being encapsulated in a Layer 4 protocol, be sure to open the appropriate port for that (such as UDP 5000) as well. If the other end of the tunnel is not a predictable IP address (often the case with users, but not likely the case in a site-to-site VPN), you can use only the destination address in the ACL.

The VPN concentrator acts as a headend for users tunneling in from elsewhere. Therefore, it needs to be capable of authenticating users, preferably against a AAA database in the campus. If desired, policy pushed from the concentrator to the clients can disable split tunneling.

You might recall that the SMR Blueprint says to disable split tunneling, while the VPN Blueprint says to enable it. Because this exam is focused primarily on the SMR Blueprint, assume that split tunneling should be disabled unless something in the question suggests that you should take the VPN Blueprint perspective.

Because you need to protect the data (the reason it's in the tunnel), you should use the stronger settings: 3DES and SHA-HMAC (instead of DES and MD5-HMAC). The default setting for Diffie-Hellman key creation should be Group 2 (1024 bits).

One point regarding placement of the VPN termination should be mentioned: If you place it after the firewall, either encrypted traffic is allowed in (and you don't know what was in that traffic) or the decryption must be performed on the firewall, which means that you don't need the concentrator. The point of using a concentrator is to offload the decryption/encryption workload from the firewall, so the concentrator must be placed between the firewall and the egress to/ingress from the outside world. Note that this placement allows the firewall to inspect and accept or reject traffic based on its actual (plaintext) content.

Also potentially passing through the VPN concentrator are VPNs coming from dialup clients via the NAS. The NAS is linked to the PSTN instead of the Internet. The authentication for the Layer 2 connection should be based on CHAP (instead of PAP, to prevent sending passwords in plain text), and the remote user should then be authenticated via AAA.

The firewall itself is there to do what firewalls are primarily for: manage the connections permitted into and out of the network. The SMR Blueprint also lists the firewall as the termination point for site-to-site VPNs for both device-management and production traffic. This is not necessarily a contradiction of what we described earlier as the role of the VPN concentrator: VPNs from individual users terminate on the concentrator, but those from sites and from networking devices to be managed terminate on the firewall. The firewall permits incoming traffic to the public servers and into the campus, depending on its ACLs (or its conduits and statics). It can permit or deny outbound traffic as well. For instance, the public servers (except the mail server) should have no reason to initiate a session with any other device; they are isolated in the DMZ for a reason.

The final networking device in the Corporate Internet module is another router, this one at the inside interface (just before traffic exits the edge and enters the campus). That description is a strong hint concerning the role played by this router: network segmentation between the edge and the campus. Filtering is probably not required here because the firewall and perimeter router should have taken care of that before traffic reaches this point. Because this inside router is in the edge (if only barely), all routing protocol updates should be authenticated.

Threats Mitigated

We have already listed the threats. Table 11.1 lists the threats and the technologies used to mitigate them, to help you pair them up. This time we've added columns to indicate which threats apply to the servers and networking devices in general, and/or to the VPN services:

Table 11.1 Medium Network Edge Threats and Their Mitigation

Threat	Servers and Networking Devices	VPN Services	Mitigated By
Network reconnaissance	Y		HIDS, protocol filtering
IP spoofing	Y		RFC 2827 and RFC 1918 filtering at ingress (and by ISP)
Trust exploitation	Y		Restrictive trust model, private VLANs
Port redirection	Y		Router/firewall filtering, HIDS
Application layer attacks	Y		HIDS, OS and applications locked down
Virus and trojan horse attacks	Y		Antivirus on every host
Denial of service (DoS)	Y		Rate limiting at edge (preferably by the ISP), TCP setup limits
Packet sniffers	Y	Y	Switched network, HIDS
Unauthorized access	Y	Y	Filtering at the router/firewall, AAA
Password attacks	Y	Y	Strong password policy, OS lockdown, IDS detection
Man-in-the-middle attacks		Y	IPSec tunneling, with content integrity validation
Network topology discovery		Y	ACLs on the perimeter router: IPSec traffic destination control

The two additional threats (compared to those faced by the small network) are man-in-the-middle attacks and network topology discovery. These are mitigated by using IPSec and configuring the ACLs on the perimeter router, which you should do anyway in accepting traffic.

Design Alternatives

Four alternatives are listed in the SMR Blueprint for this module:

- Replace the edge (perimeter) router with a stateful firewall
- Add a NIDS between the edge router and the firewall
- Eliminate the inside router
- Add content inspection (such as URL filtering)

Discussing each of these briefly, you could replace the edge router with a stateful firewall if you want a stronger defense posture than that provided by a router. Related to a stronger security posture is the second alternative. Any firewall might be filtering attack attempts that it might be useful to know about. In that case, adding a NIDS (with a very low alarm level, and possibly with logging separated from that of the other IDS in the network) between the ingress and the firewall would allow such attacks to be detected. The separated logging is recommended because of the volume of alarms this NIDS would generate.

If the layering of the defense in the Blueprint is more than the organization feels is necessary, the inside router could be eliminated. The Layer 3 switch inside the campus would provide the Layer 3 heading manipulation required on traffic, but at a price of moving traffic into the campus and back.

Finally, content inspection would serve to monitor and restrict the URLs visited by hosts inside the organization. Especially if there has been a problem with offensive traffic (of any kind) or if there is some indication of misuse of the Internet connection (even by only a few users), such monitoring and filtering could be used to ensure that business assets are being used for business purposes.

The Medium Network WAN

The SMR Blueprint actually does not address this module in much detail, but there isn't a great deal to the module: an ingress router that connects directly from the Frame Relay/ATM network to the Campus module. It serves to provide ingress/egress filtering, QoS management, and routing protocol updates with peers beyond this network. The filtering mitigates against IP spoofing, and access controls on the router protect it from unauthorized access.

If there is some concern regarding the confidentiality or privacy that is actually obtained over the private circuits, an alternative would be to encrypt data transiting the WAN (use IPSec tunnels). Likewise, if the ACLs seem insufficient for the traffic encountered, a stateful firewall could be added to the router (via the Firewall Feature Set).

The Medium Network Campus

The medium network's Campus module has the same three types of hosts—corporate servers, management servers (possibly more than one), and users—as

the small network did. One difference, however, based on the number of hosts and the amount of traffic, is to use multiple Layer 2 switches here and to complement them with a Layer 3 switch working with a separate NIDS appliance.

Assets

The lowest-value assets in the campus remain the users' systems. Again, few, if any, corporate information assets should reside on any user's system, but working copies and temporary files will remain, containing sensitive information that should be protected. Users here are no different, except that, in the larger population of a medium network, there are likely to be more technically adept users who make more extensive changes (utilities, IRC and/or IM traffic, remote access to their desktops, and so forth).

As before, the corporate servers are the gold mine for a hacker after monetary gain (if the hacker is there for ego reasons, the corporate servers are still certainly a target, but other devices might be targeted just as much). Whether the organization provides products or services, proprietary data and confidential information concerning the organization and its customers must be protected. The information stored in these servers must be considered among the most valuable information assets the organization possesses (if not the most valuable). If this is a branch office of an even larger organization, these servers will have trust relationships with their corollaries at the corporate headquarters, increasing the local (and too often less well-protected) servers' value to a hacker.

The management servers in this module provide AAA (possibly including a server for OTP), logging, the IDS Director function, and general configuration management. Using redundant AAA servers is good practice in this size organization. The knowledgeable hacker will target these servers to facilitate future activities (which makes authorization settings critically important). In addition, the logs stored on this server are always a target for creative editing. If the hacker can access the IDS management function, he can remove the signature file for a particular attack, resulting in no alarm from the IDS when that attack is launched. These servers are not a place to economize.

The Layer 3 switch handles a larger volume of traffic among more hosts and servers. Employing a Layer 3 switch facilitates rapid traffic distribution, with QoS capability. The Layer 2 switches distribute traffic within the department or physical area as before.

Threats

The campus is reasonably well protected from an external threat: The edge modules (Corporate Internet and WAN) are both well filtered, and traffic does not pass through them to the campus lightly. The campus is much more likely to suffer an attack from an internal source than an external one. However, as we noted in the case of outsourced IT in the small network, the origin can be moot. In this case, you could easily see an external source persuading someone on the inside to "use my copy" of a seemingly innocuous application (which is actually malware), or a laptop that is not well secured could come back from travel infected with a worm that propagates through the campus. The real source was external in both cases, but the entry into the network was directly into the campus (an internal attack, strictly speaking) rather than one that penetrated through the protections in place in the edge.

Either way—or via any other path, for that matter—the campus's assets must be protected from threats that arrive inside the external protections. With greater assets and more people inside, you must expect to see more threats than you saw in the small network. The threats in the medium network campus include these:

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

This list includes all the threats to the small network campus, plus IP spoofing and password attacks. Why the two additions? IP spoofing can allow access to devices that trust other network devices (an argument for strong AAA, including on strictly internal resources). Password attacks can be used for the same reason: to gain access to valuable resources. If you dutifully read your logs from the perimeter but seldom get around to logs from the valuable internal hosts, you could be missing the early warnings of an internal threat.

Devices and Implementation

Device implementation in the medium campus is quite similar in principle to that in the small campus. Antivirus, OS, and application maintenance are no different. Configuration management is likely to be more difficult, in part because of the greater number of users and the presence of more applications, and in part because of the greater personal distance between IT personnel and the rest of the organization. In the larger organization, less personal knowledge of who is doing what makes nonstandard (and, therefore, less protected) configurations more likely.

All servers in this module should have HIDS installed and configured to operate very aggressively: False positives in this environment, as well as the small network environment, are better than false negatives. This is likely to be a more aggressive HIDS posture than you used in the edge and much more aggressive than you might use at the very ingress (where you might place the optional NIDS). The NIDS inside the campus should be more aggressively configured than that in the edge; it will monitor mirrored traffic from the sensitive VLANs. This device is likely to detect the first signs of a compromised host (with the compromise coming via an unauthorized external connection, imported malware, or a laptop exposed when outside the network for legitimate reasons).

The switches should be configured to have all servers on private VLANs. Management (systems-management) workstations should also be on private VLANs. Departmental VLANs should be filtered from one another (at Layer 3) unless there is a demonstrated need to communicate (Engineering need not be able to access Finance without filtering, for instance). Within these departmental VLANs, you can implement private VLANs to further protect sensitive hosts. As in the edge, all unused ports on all switches should be disabled to prevent unknown and unauthorized devices from connecting to the network. In addition, ports that need to trunk should be specified, and all others should have their trunking autonegotiation set to Off.

As you can see, although the medium campus is somewhat more complicated because of its greater size, the principles behind its secure implementation remain largely the same as those in the small campus.

Threats Mitigated

Having described the threats and the security measures to be taken in the medium campus, it's time to summarize them. Table 11.2 presents the threats and security measures taken to mitigate them.

Table 11.2 Medium Network Campus Threats and Their Mitigation	
Threat	**Mitigated By**
Application-layer attacks	OS and applications locked down, HIDS
IP spoofing	RFC 2827 filtering by segment
Packet sniffers	Switched network, HIDS on servers
Password attacks	Strong authentication required for access to key applications and data
Port redirection	HIDS
Trust exploitation	Private VLANs, restrictive trust model, where appropriate
Unauthorized access	HIDS, strict authorization control on applications
Virus and trojan horse attacks	Antivirus on every host

The two new threats (compared to the small network) are IP spoofing and password attacks. They are mitigated by filtering according to RFC 2827 on a segment basis (traffic entering a segment should not have a source IP within that segment) and through the use of strong authentication requirements on important applications and data.

Design Alternatives

Three alternative designs can be implemented, all based on the traffic load that the medium network must bear:

- Eliminate the Layer 2 switches and perform all switching on the core Layer 3 switch
- Replace the Layer 3 switch with a router and Layer 2 switch
- Replace the separate NIDS appliance with a NIDS module on the Layer 3 switch

Depending on the number of ports needed and the volume of traffic, either all switching can be done by the Layer 3 switch, eliminating the need for Layer 2 devices, or the Layer 3 segmentation can be offloaded to a router and the higher-speed Layer 3 switch can be eliminated. Another way of looking at these two alternatives is to collapse all of Layer 2 and 3 into one large switch, or separate Layer 3 into a router and continue to switch only at Layer 2. The third design alternative is somewhat different, relating to the NIDS. In this alternative, the NIDS module has the advantage of higher throughput via the Layer 3 switch's backplane (versus one Ethernet connection, hopefully 100Mbps); traffic selection for processing by it is done through

ACLs. This tradeoff should be examined just like any other dedicated appliance versus integrated multifunction application: It should be based on the performance of the dedicated appliance versus the technical need, and then should be compared to the advantages inherent in the multifunction solution.

Branch Versus Headend

The medium network can be a standalone organization, in which case it might have subordinate branches for which it acts as the headend, or it can be a branch of a larger enterprise. The medium network that we have described to this point is operating as a headend. As a branch under a headquarters at another location, there would be changes:

- Device management would probably be done via private connection from headquarters (note that the perimeter router, with an externally exposed interface, would need separate management because it would be external to the private connection terminus).
- AAA and other security management functions would likewise be done from headquarters via a private connection.
- The Corporate Internet module would be scaled down, commensurate with the degree of Internet access allowed from the branch and the need (if any) for a DMZ.

The private connection supporting the branch from the headend might be a leased circuit (Frame Relay or ATM), in which case the Corporate Internet module's role would be greatly reduced, or it could be an IPSec tunnel. IPSec offers the opportunity to employ an existing Internet connection and thus eliminates the expense of a leased circuit. However, it supports IP only unless GRE tunneling is added to encapsulate the IPSec and other protocols alike, such as multicast. If IPSec is used, the WAN module might not even be present.

Summary

The medium network is more complex than the small network, mostly because of scale: There are simply more possibilities to go with the larger number of hosts and greater amount of traffic. Having discussed the small and medium centralized operation, it's time to turn to the third piece of the SMR Blueprint: the remote users and their needs. That's the subject of Chapter 12, "The Remote-User Design."