SAFE

Extending the Security Blueprint to Small, Midsize, and Remote-User Networks

Table of Contents

Authors

Sean Convery and Roland Saville are the authors of this White Paper, and lead architects for the reference implementation at the Cisco headquarters in San Jose, CA USA. Sean and Roland are both network architects, focused on VPN and security issues.

Abstract

The principal goal of this paper is to provide best-practice information to interested parties on designing and implementing secure networks. The original SAFE white paper is available at the following URL: http://www.cisco.com/warp/public/cc/so/cuso/epso/sqfr/safe_wp.htm. It was written to provide best-practice information on large enterprise network security designs.

This document takes the same principles and sizes them appropriately for smaller networks, including branches of larger enterprises as well as standalone, small to midsize security deployments. It also includes information on remote-user networks such as teleworkers and mobile workers. It is not necessary to read the original SAFE paper prior to reading this document, because all the information covered in the original document is included, as appropriate, in this document.

SAFE serves as a guide to network designers considering the security requirements of their network. SAFE takes a defense-in-depth approach to network security design. This type of design focuses on the expected threats and their methods of mitigation, rather than on "Put the firewall here, put the intrusion detection system there." This strategy results in a layered approach to security where the failure of one security system is not likely to lead to the compromise of network resources. SAFE is based on Cisco products and those of its partners.

Beginning with an overview of the architecture, this document then details the specific designs under consideration. The first two sections of each module describe the key devices, and expected threats with basic mitigation diagrams. Detailed technical analysis of the design follows, along with more detailed threat mitigation techniques and migration strategies. Appendix A details the validation lab for SAFE and includes configuration snapshots. Appendix B is a primer on network security. Readers who are unfamiliar with basic network security concepts are encouraged to read this section before the rest of the document. Appendix C contains glossary definitions of the technical terms used in this document.

This document focuses heavily on threats encountered in networks today. Network designers who understand these threats can better decide where and how to deploy mitigation technologies. Without a full understanding of the threats involved in network security, deployments tend to be incorrectly configured, are too focused on security devices, or lack threat response options. By taking the threat-mitigation approach, this document should provide network designers with information for making sound network security choices.

Audience

Though this document is technical, it can be read at different levels of detail, depending on the reader. A network manager, for example, can read the introductory sections in each area to obtain a good overview of network security design strategies and considerations. A network engineer or designer can read this document in its entirety and gain design information and threat analysis details, which are supported by actual configuration snapshots for the devices involved. This document is also appropriate for readers of the original SAFE white paper who are interested in designing networks on a smaller scale than the architectures presented in the original paper. It may be helpful to read the introductory sections of the paper first and then skip right to the type of network you are interested in deploying.

Caveats

This document presumes that you already have a security policy in place. Cisco Systems does not recommend deploying security technologies without an associated policy. This document directly addresses the needs of small, midsize, and remote-user networks. Readers interested in large enterprise network design should reference the original SAFE white paper available at the following URL: http://www.cisco.com/warp/public/cc/so/cuso/epso/sqfr/safe_wp.htm.

Following the guidelines in this document does not guarantee a secure environment, nor does it guarantee that you will prevent all penetrations. Absolute security can be achieved only by disconnecting a system from the network, encasing it in concrete, and putting it on the bottom floor of Fort Knox. There your data will be very safe, though inaccessible. However, you can achieve reasonable security by establishing a good security policy, following the guidelines in this document, staying up-to-date on the latest developments in the hacker and security communities, and maintaining and monitoring all systems with sound system administration practices. This includes awareness of application security issues that are not comprehensively addressed in this paper.

Though *virtual private networks* (VPNs) are included in this architecture, they are not described in detail. Information such as scaling details, resilience strategies, and other topics related to VPNs are not included. Interested readers should refer to the SAFE VPN White Paper at the following URL: http://www.cisco.com/warp/public/cc/so/cuso/epso/sqfr/safev_wp.htm. Like VPNs, identity strategies (including *certificate authorities* [CAs]) are not discussed at any level of detail in this paper. Similarly, CAs require a level of focus that this document could not provide and still adequately address all the other relevant areas of network security. Also, because most networks have yet to deploy fully functional CA environments, it is important to discuss how to securely deploy networks without them. E-commerce is another area not covered in this document. Network recommendations for e-commerce vary little from small to large organizations, so the original SAFE paper is still appropriate. Finally, certain advanced networked applications and technologies (such as content networking, caching, and server load balancing) are not included in this document. Although their use within SAFE is to be expected, this paper does not cover their specific security needs.

SAFE uses the products of Cisco Systems and its partners. However, this document does not specifically refer to products by name. Instead, components are referred to by functional purpose rather than model number or name. During the validation of SAFE, real products were configured in the exact network implementation described in this document. The lab and results, along with specific configuration snapshots from the lab, are included in Appendix A, “Validation Lab.”

Throughout this document the term “hacker” denotes an individual who attempts to gain unauthorized access to network resources with malicious intent. Although the term “cracker” is generally regarded as the more accurate word for this type of individual, hacker is used here for readability.

Architecture Overview

Design Fundamentals

SAFE emulates as closely as possible the functional requirements of today's networks. Implementation decisions varied, depending on the network functionality required. However, the following design objectives, listed in order of priority, guided the decision-making process.

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

First and foremost, SAFE is a security architecture. It must prevent most attacks from successfully affecting valuable network resources. The attacks that succeed in penetrating the first line of defense, or originate from inside the network, must be accurately detected and quickly contained to minimize their effect on the rest of the network. However, in being secure, the network must continue to provide critical services that users expect. Proper network security and good network functionality can be provided at the same time. The SAFE architecture is not a revolutionary way of designing networks, but merely a blueprint for making networks secure.

This SAFE architecture for small, midsize, and remote networks was designed without resiliency. Readers interested in designing secure networks in a resilient environment should read the original SAFE white paper (hereafter referred to as “SAFE Enterprise”).

At many points in the network design process, you need to choose between using integrated functionality in a network device versus using a specialized functional appliance. The integrated functionality is often attractive because you can implement it on existing equipment, or because the features can interoperate with the rest of the device to provide a better functional solution. Appliances are often used when the depth of functionality required is very advanced or when performance needs require using specialized hardware. Make your decisions based on the capacity and functionality of the appliance versus the integration advantage of the device. For example, sometimes you can choose an integrated higher-capacity Cisco IOS® router with IOS firewall software as opposed to a smaller IOS router with a separate firewall. Throughout this architecture, both types of systems are used. When the design requirements did not dictate a specific choice, the design opted to go with integrated functionality in order to reduce the overall cost of the solution.

Module Concept

Although most networks evolve with the growing IT requirements of an organization, the SAFE architecture uses a green-field modular approach. A modular approach has two main advantages: First, it allows the architecture to address the security relationship between the various functional blocks of the network. Second, it permits designers to evaluate and implement security on a module-by-module basis, instead of attempting the complete architecture in a single phase. The security design of each module is described separately, but is validated as part of the complete design.

Although it is true that most networks cannot be easily dissected into clear-cut modules, this approach provides a guide for implementing different security functions throughout the network. The authors do not expect network engineers to design their networks identical to the SAFE implementation, but rather use a combination of the modules described and integrate them into the existing network.

SAFE Axioms

Routers Are Targets

Routers control access from every network to every network. They advertise networks and filter who can use them, and they are potentially a hacker's best friend. Router security is a critical element in any security deployment. By their nature, routers provide access and, therefore, you should secure them to reduce the likelihood that they can be directly compromised. You can refer to other documents that have been written about router security, which provide more detail on the following subjects:

- Locking down Telnet access to a router
- Locking down *Simple Network Management Protocol* (SNMP) access to a router
- Controlling access to a router through the use of *Terminal Access Controller Access Control System Plus* (TACACS+)
- Turning off unneeded services
- Logging at appropriate levels
- Authentication of routing updates

The most current document on router security is available at the following URL:
http://www.cisco.com/warp/public/707/21.html

Switches Are Targets

Like routers, switches (both Layer 2 and Layer 3) have their own set of security considerations. Unlike routers, not as much public information is available about the security risks in switches and what can be done to mitigate those risks. Most of the security techniques detailed in the preceding section, "Routers Are Targets," apply to switches. In addition, you should take the following precautions:

- Ports without any need to trunk should have any trunk settings set to off, as opposed to auto. This setup prevents a host from becoming a trunk port and receiving all traffic that would normally reside on a trunk port.
- If you are using older versions of software for your Ethernet switch, make sure that trunk ports use a *virtual LAN* (VLAN) number not used anywhere else in the switch. This setup prevents packets tagged with the same VLAN as the trunk port from reaching another VLAN without crossing a Layer 3 device.
  For more information, refer to http://www.sans.org/newlook/resources/IDFAQ/vlan.htm
- Usable all unused ports on a switch. This setup prevents hackers from plugging in to unused ports and communicating with the rest of the network.
- Avoid using VLANs as the sole method of securing access between two subnets. The capability for human error, combined with the understanding that VLANs and VLAN tagging protocols were not designed with security in mind, makes their use in sensitive environments inadvisable. When VLANs are needed in security deployments, be sure to pay close attention to the configurations and guidelines mentioned above.

Within an existing VLAN, private VLANs provide some added security to specific network applications. Private VLANs work by limiting which ports within a VLAN can communicate with other ports in the same VLAN. Isolated ports within a VLAN can communicate only with promiscuous ports. Community ports can communicate only with other members of the same community and promiscuous ports. Promiscuous ports can communicate with any port. This is an effective way to mitigate the effects of a single compromised host. Consider a standard public services segment with a Web, *File Transfer Protocol* (FTP), and *Domain Name System* (DNS) server. If the DNS server is compromised, a hacker can pursue the other two hosts without passing back through the firewall. If private VLANs are deployed, if one system is compromised, it cannot communicate with the other systems. The only targets a hacker can pursue are hosts on the other side of the firewall. Because they restrict layer 2

connectivity, private VLANs make troubleshooting network problems more difficult. Remember that private VLANs are not supported on all Ethernet switches available on the market today. In particular, most low-end switches do not yet support this feature.

Hosts Are Targets

The most likely target during an attack, the host presents some of the most difficult challenges from a security perspective. There are numerous hardware platforms, operating systems, and applications, all of which have updates, patches, and fixes available at different times. Because hosts provide the application services to other hosts that request them, they are extremely visible within the network. For example, many people have visited www.whitehouse.gov, which is a host, but few have attempted to access s2-0.whitehouseisp.net, which is a router. Because of this visibility, hosts are the most frequently attacked devices in any network intrusion attempt. In part because of the security challenges mentioned above, hosts are also the most successfully compromised devices. For example, a given Web server on the Internet might run a hardware platform from one vendor, a network card from another, an operating system from still another vendor, and a Web server that is either open source or from yet another vendor. Additionally, the same Web server might run applications that are freely distributed via the Internet, and might communicate with a database server that starts the variations all over again. That is not to say that the security vulnerabilities are specifically caused by the multisource nature of all of this, but rather that as the complexity of a system increases, so does the likelihood of a failure.

To secure hosts, pay careful attention to each of the components within the systems. Keep any systems up-to-date with the latest patches, fixes, and so forth. In particular, pay attention to how these patches affect the operation of other system components. Evaluate all updates on test systems before you implement them in a production environment. Failure to do so might result in the patch itself causing a denial of service (DoS).

Networks Are Targets

Network attacks are among the most difficult attacks to deal with because they typically take advantage of an intrinsic characteristic in the way your network operates. These attacks include *Address Resolution Protocol* (ARP) and *Media Access Control* (MAC)-based Layer 2 attacks, sniffers, and distributed *denial-of-service* (DDoS) attacks. Some of the ARP and MAC-based Layer 2 attacks can be mitigated through best practices on switches and routers. Sniffers are discussed in the primer at the end of this document. DDoS, however, is a unique attack that deserves special attention.

The worst attack is the one that you cannot stop. When performed properly, DDoS is just such an attack. As outlined in Appendix B, “Network Security Primer,” DDoS works by causing tens or hundreds of machines to simultaneously send spurious data to an IP address. The goal of such an attack is generally not to shut down a particular host, but rather to make the entire network unresponsive. For example, consider an organization with a DS1 (1.5 Mbps) connection to the Internet that provides e-commerce services to its Web site users. Such a site is very security conscious and has intrusion detection, firewalls, logging, and active monitoring. Unfortunately, none of these security devices helps when a hacker launches a successful DDoS attack.

Consider 100 devices around the world, each with DSL (500 Kbps) connections to the Internet. If these systems are remotely told to flood the serial interface of the e-commerce organization’s Internet router, they can easily flood the DS1 with erroneous data. Even if each host is able to generate only 100 Kbps of traffic (lab tests indicate that a stock PC can easily generate 50 Mbps with a popular DDoS tool), that amount is still almost ten times the amount of traffic that the e-commerce site can handle. As a result, legitimate Web requests are lost, and the site appears to be down for most users. The local firewall drops all the erroneous data, but by then the damage is done. The traffic has crossed the WAN connection and filled up the link.

Only through cooperation with its *Internet service provider* (ISP) can this fictitious e-commerce company hope to thwart such an attack. An ISP can configure rate limiting on the outbound interface to the company's site. This rate limiting can drop most undesired traffic when it exceeds a prespecified amount of the available bandwidth. The key is to correctly flag traffic as undesired.

Common forms of DDoS attacks are *Internet Control Message Protocol* (ICMP) floods, TCP SYN floods, or *User Datagram Protocol* (UDP) floods. In an e-commerce environment, this type of traffic is fairly easy to categorize. Only when limiting a TCP SYN attack on port 80 (*Hypertext Transfer Protocol* [HTTP]) does an administrator run the risk of locking out legitimate users during an attack. Even then, it is better to temporarily lock out new legitimate users and retain routing and management connections than to have the router overrun and lose all connectivity.

More sophisticated attacks use port 80 traffic with the ACK bit set so that the traffic appears to be legitimate Web transactions. It is unlikely that an administrator could properly categorize such an attack because acknowledged TCP communications are exactly the sort that you want to allow into your network.

One approach to limiting this sort of attack is to follow filtering guidelines for networks outlined in RFC 1918 and RFC 2827. RFC 1918 specifies the networks that are reserved for private use and should never be seen across the public Internet. RFC 2827 filtering is discussed in the "IP Spoofing" section of Appendix B, "Network Security Primer." For example, for inbound traffic on a router that is connected to the Internet, you employ RFC 1918 and 2827 filtering to prevent this unauthorized traffic from reaching the corporate network. When implemented at the ISP, this filtering prevents DDoS attack packets that use these addresses as sources from traversing the WAN link, potentially saving bandwidth during the attack. Collectively, if ISPs worldwide were to implement the guidelines in RFC 2827, source address spoofing would be greatly diminished. Although this strategy does not directly prevent DDoS attacks, it does prevent such attacks from masking their source, making traceback to the attacking networks much easier. Ask your ISP about which DDoS mitigation options they make available to their customers.

Applications Are Targets

Applications are coded by human beings (mostly) and, as such, are subject to numerous errors. These errors can be benign-for example, an error that causes your document to print incorrectly-or malignant-for example, an error that makes the credit card numbers on your database server available via anonymous FTP. It is the malignant problems, as well as other more general security vulnerabilities, that need careful attention. Care needs to be taken to ensure that commercial and public domain applications are up-to-date with the latest security fixes. Public domain applications, as well as custom developed applications, also require code review to ensure that the applications are not introducing any security risks caused by poor programming. This programming can include scenarios such as how an application makes calls to other applications or the OS itself, the privilege level at which the application runs, the degree of trust that the application has for the surrounding systems, and finally, the method the application uses to transport data across the network. The following section discusses *intrusion detection systems* (IDSs) and how they can help mitigate some of the attacks launched against applications and other functions within the network.

Intrusion Detection Systems

Intrusion detection systems (IDSs) act like an alarm system in the physical world. When an IDS detects something that it considers an attack, it can either take corrective action itself or notify a management system for actions by the administrator. Some systems are more or less equipped to respond and prevent such an attack. Host-based intrusion detection can work by intercepting OS and application calls on an individual host. It can also operate by after-the-fact analysis of local log files. The former approach allows better attack prevention, whereas the latter approach dictates a more passive attack-response role. Because of the specificity of their role, *host-based IDS* (HIDS) systems are often better at preventing specific attacks than *network IDS* (NIDS) systems, which usually issue only an alert upon discovery of an attack. However, that specificity causes a loss of perspective to the overall network. This is where NIDS excels. Ideally, Cisco recommends a combination of the two systems—HIDS on critical hosts and NIDS looking over the whole network—for a complete intrusion detection system. Unfortunately, IT budgets will often dictate a choice of one technology or another. In this case, careful attention should be placed on the overall cost of each technology, the number of devices that need to be monitored, and the personnel required to respond to an attack.

When an IDS is deployed, you must tune its implementation to increase its effectiveness and remove "false positives." False-positives are defined as alarms caused by legitimate traffic or activity. False negatives are attacks that the IDS system fails to see. When the IDS is tuned, you can configure it more specifically as to its threat-mitigation role. As mentioned above, you should configure HIDS to stop most valid threats at the host level because it is well prepared to determine that certain activity is, indeed, a threat.

When deciding on mitigation roles for NIDS, you have two primary options. Remember that the first step prior to implementing any threat-response option is to adequately tune NIDS to ensure that any percieved threat is legitimate.

The first option-and potentially the most damaging if improperly deployed-is to "shun" traffic through the addition of access control filters on routers and firewalls. When a NIDS detects an attack from a particular host over a particular protocol, it can block that host from coming into the network for a predetermined amount of time. Although on the surface this might seem like a great aid to a security administrator, in reality it must be very carefully implemented, if at all. The first problem is that of spoofed addresses. If traffic that matches an attack is seen by the NIDS, and that particular alarm triggers a shun situation, the NIDS will deploy the access list to the device. However, if the attack that caused the alarm used a spoofed address, the NIDS has now locked out an address that never initiated an attack. If the IP address that the hacker used happens to be the IP address of a major ISP's outbound HTTP proxy server, a huge number of users could be locked out. This by itself could be an interesting DoS threat in the hands of a creative hacker.

To mitigate the risks of shunning, you should generally use it only on TCP traffic, which is much more difficult to successfully spoof than UDP. Use it only in cases where the threat is real and the chance that the attack is a false positive is very low. Also consider setting the shun length very short. This setup will block the user long enough to allow the administrator to decide what permanent action (if any) he/she wants to take against that IP address. However, in the interior of a network, many more options exist. With effectively deployed RFC 2827 filtering, spoofed traffic should be very limited. Also, because customers are not generally on the internal network, you can take a more restrictive stance against internally originated attack attempts. Another reason for this is that internal networks do not often have the same level of stateful filtering that edge connections possess. As such, IDS needs to be more heavily relied upon than in the external environment.

The second option for NIDS mitigation is the use of TCP resets. As the name implies, TCP resets operate only on TCP traffic and terminate an active attack by sending TCP reset messages to the attacking and attacked host. Because TCP traffic is more difficult to spoof, you should consider using TCP resets more often than shunning. Keep in mind that TCP resets in a switched environment are more challenging than when a standard hub is used, because all ports don't see all traffic without the use of a *Switched Port Analyzer* (SPAN) or mirror port. Make sure this mirror port supports bidirectional traffic flows and can have SPAN port MAC learning disabled.

Both of these mitigation options require 24x7 staffing to watch the IDS consoles. Because IT staff are often overworked, (particularly in smaller organizations), consider outsourcing your IDS management to a third party.

From a performance standpoint, NIDS observes packets on the wire. If packets are sent faster than the NIDS can process them, there is no degradation to the network because the NIDS does not sit directly in the flows of data. However, the NIDS will lose effectiveness and packets could be missed, causing both false negatives and false positives. Be sure to avoid exceeding the capabilities of IDS so that you can get their benefit. From a routing standpoint, IDS, like many state-aware engines, does not operate properly in an asymmetrically routed environment. Packets sent out from one set of routers and switches and returning through another will cause the IDS systems to see only half the traffic, causing false positives and false negatives.

Secure Management and Reporting

"If you're going to log it, read it." This proposition is so simple that almost everyone familiar with network security has said it at least once. Yet logging and reading information from many devices can be very challenging. Which logs are most important? How do I separate important messages from mere notifications? How do I ensure that logs are not tampered with in transit? How do I ensure that my time stamps match each other when multiple devices report the same alarm? What

information is needed if log data is required for a criminal investigation? How do I deal with the volume of messages that can be generated when a system is under attack? You must address all these questions when considering managing log files effectively. From a management standpoint, a different set of questions needs to be asked: How do I securely manage a device? How can I push content out to public servers and ensure that it is not tampered with in transit? How can I track changes on devices to troubleshoot when attacks or network failures occur?

Although the *"out-of-band"* (OOB) management architecture described in SAFE Enterprise provides the highest levels of security, it is not recommended here because our goal is a cost-effective security deployment. In the OOB environment, each network device and host has its own dedicated management interface, which connects to the private management network. This setup mitigates the risk of passing insecure management protocols such as Telnet, *Trivial File Transfer Protocol* (TFTP), SNMP, and syslog over the production network where it could be intercepted or modified. In the architecture described in this paper, management traffic flows "in band" in all cases, and is made as secure as possible using tunneling protocols and secure variants to insecure management protocols. For example, *Secure Shell Protocol* (SSH) is used whenever possible instead of Telnet. With management traffic flowing "in band" across the production network, it becomes very important to follow more closely the axioms mentioned earlier in this document.

When management of a device on the outside of a firewall is required, you should consider several factors. First, what management protocols does the device support? For devices with *IP Security* (IPSec), devices should be managed by simply creating a tunnel from the management network to the device. This setup allows many insecure management protocols to flow over a single encrypted tunnel. When IPSec is not possible because it is not supported on a device, other less-secure alternatives must be chosen. For configuration of the device, SSH or *Secure Sockets Layer* (SSL) can often be used instead of Telnet to encrypt any configuration modifications made to a device. These same protocols can sometimes also be used to push and pull data to a device instead of insecure protocols such as TFTP and FTP. Often, however, TFTP is required on Cisco equipment to back up configurations or to update software versions. This leads to the second question: Does this management channel need to be active at all times? If not, then temporary holes can be placed in a firewall while the management functions are performed and then later removed. This process does not scale with large numbers of devices, however. If the channel needs to be active at all times, such as with SNMP, the third question should be considered: Do you really need this management tool? Often SNMP managers are used on the inside of a network to ease troubleshooting and configuration. But for a DMZ switch that is providing Layer 2 services to two or three servers, is it really necessary? If not, disable it. If you decide it is required, know that you are introducing a potential vulnerability into your environment. The next several paragraphs discuss in more detail the specific types of management.

From a reporting standpoint, most networking devices can send syslog data that can be invaluable when troubleshooting network problems or security threats. Send this data to your syslog analysis host from any devices whose logs you wish to view. This data can be viewed in real-time or via on-demand and scheduled reports. Depending on the device involved, you can choose various logging levels to ensure that the correct amount of data is sent to the logging device. You also need to flag device log data within the analysis software to permit granular viewing and reporting. For example, during an attack, the log data provided by Layer 2 switches might not be as interesting as the data provided by the intrusion detection system. To ensure that log messages are time synchronized to one another, clocks on hosts and network devices must be in sync. For devices that support it, *Network Time Protocol* (NTP) provides a way to ensure that accurate time is kept on all devices. When dealing with attacks, seconds matter because it is important to identify the order in which a specified attack occurred.

Configuration change management is another issue related to secure management. When a network is under attack, it is important to know the state of critical network devices and when the last known modifications occurred. Creating a plan for change management should be a part of your comprehensive security policy, but, at a minimum, you should record changes using authentication systems on the devices, and archive configurations via FTP or TFTP.

Headend versus Branch Considerations

The small and medium designs that follow can be used in two possible configurations. In the first, the design is the "headend" of an organization's network. This headend may have VPN connections to other offices of the same organization. For example, a large law office may use the medium network design for its headend, and several small network designs for its other locations. Full-time teleworkers might come into the headend over some of the options discussed in the remote network design. In the second configuration, the design is acting as a branch of a larger organization, built in the configuration described in SAFE Enterprise.

Still another example would be a large automotive company that might use the SAFE Enterprise design for its corporate headquarters, and many of the designs in this paper for its remote locations and teleworkers. Where appropriate, the specific design changes that may be required are discussed in each section.

Expected Threats

From a threat perspective, a small or midsize network is like most networks connected to the Internet-there are internal users who need access out and external users who need access in. Several common threats can generate the initial compromise that a hacker needs to further penetrate the network with secondary exploits.

First is the threat from internal users. Though statistics vary on the percentage, it is an established fact that most attacks come from the internal network. Disgruntled employees, corporate spies, visiting guests, and inadvertent bumbling users are all potential sources of such attacks. When designing security, you must be aware of the potential for internal threats.

Second is the threat to the publicly addressable hosts that are connected to the Internet. These systems will likely be attacked with application layer vulnerabilities and DoS attacks.

For a complete discussion of threat details, refer to Appendix B, "Network Security Primer."

Small Network Design

The small network design has two modules: the corporate Internet module and the campus module. The corporate Internet module has connections to the Internet and also terminates VPN and public services (DNS, HTTP, FTP, SMTP) traffic. The campus module contains the Layer 2 switching and all the users, as well as the management and intranet servers. Most of the discussion for this design is based on the small network operating as the headend for a corporation. Specific design changes when used as a branch are also included.

Figure 1 Detailed Model of Small Network

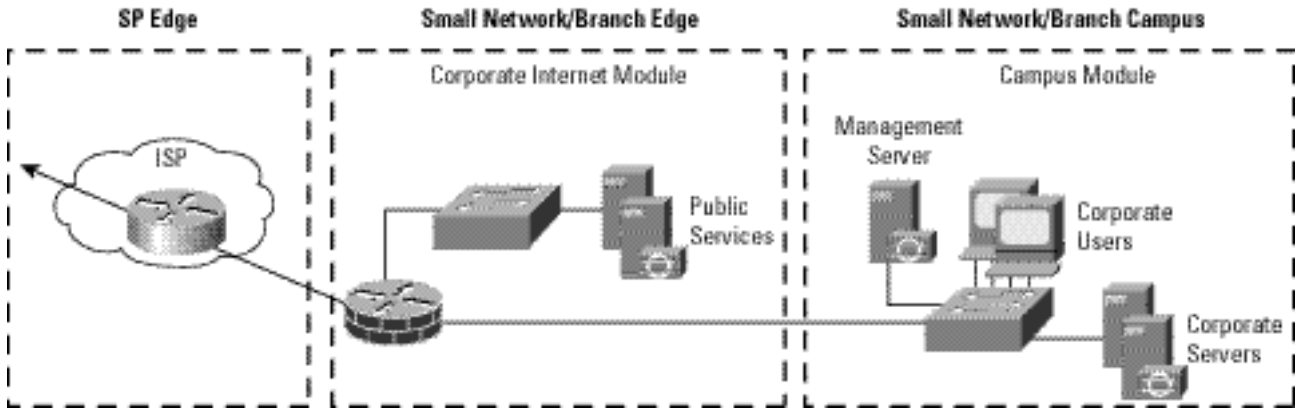

Corporate Internet Module

The corporate Internet module provides internal users with connectivity to Internet services and Internet users access to information on public servers. VPN access is also provided to remote locations and telecommuters. This module is not designed to serve e-commerce type applications. Refer to the section "E-Commerce Module" in SAFE Enterprise for more details on providing Internet commerce.

Key Devices

- *SMTP server*—Acts as a relay between the Internet and the intranet mail servers
- *DNS server*—Serves as authoritative external DNS server for the enterprise; relays internal requests to the Internet
- *FTP/HTTP server*—Provides public information about the organization
- *Firewall or firewall router*—Provides network-level protection of resources, stateful filtering of traffic, and VPN termination for remote sites and users
- *Layer 2 switch (with private VLAN support)*—Ensures that data from managed devices can only cross directly to the IOS firewall

Figure 2 Detailed Model of Small Network Corporate Internet Module

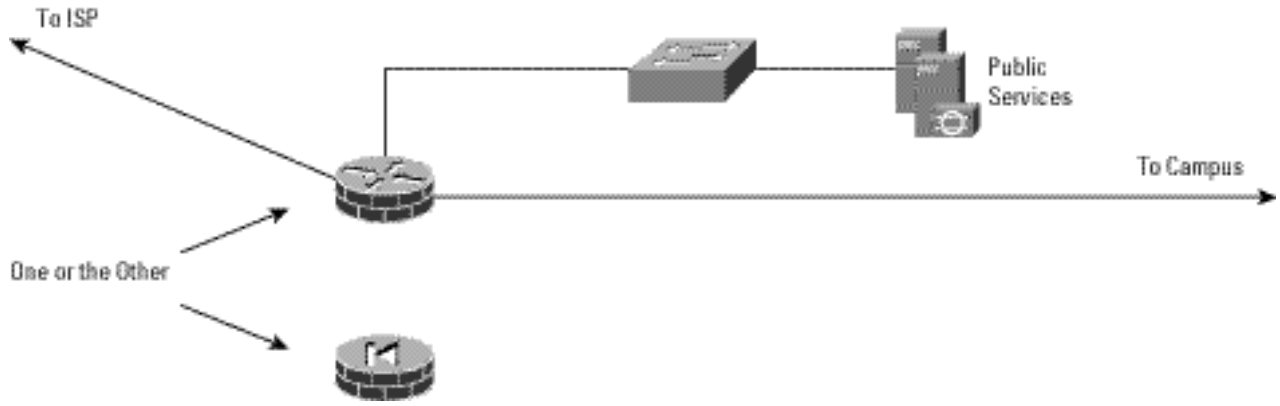

Threat Mitigation

There are publicly addressable servers that are the most likely points of attacks. The following are expected threats:

- *Unauthorized access*—Mitigated through filtering at the firewall
- *Application layer attacks*—Mitigated through HIDS on the public servers
- *Virus and Trojan-horse attacks*—Mitigated through virus scanning at the host level
- *Password attacks*—Limited services available to brute force; OS and IDS can detect the threat
- *Denial of service*—Committed access rate (CAR) at ISP edge and TCP setup controls at firewall to limit exposure
- *IP spoofing*—RFC 2827 and 1918 filtering at ISP edge and local firewall
- *Packet sniffers*—Switched infrastructure and host IDS to limit exposure
- *Network reconnaissance*—HIDS detects recon; protocols filtered to limit effectiveness
- *Trust exploitation*—Restrictive trust model and private VLANs to limit trust-based attacks
- *Port redirection*—Restrictive filtering and host IDS to limit attack

Figure 3 Small Network Attack Mitigation Roles for Corporate Internet Module

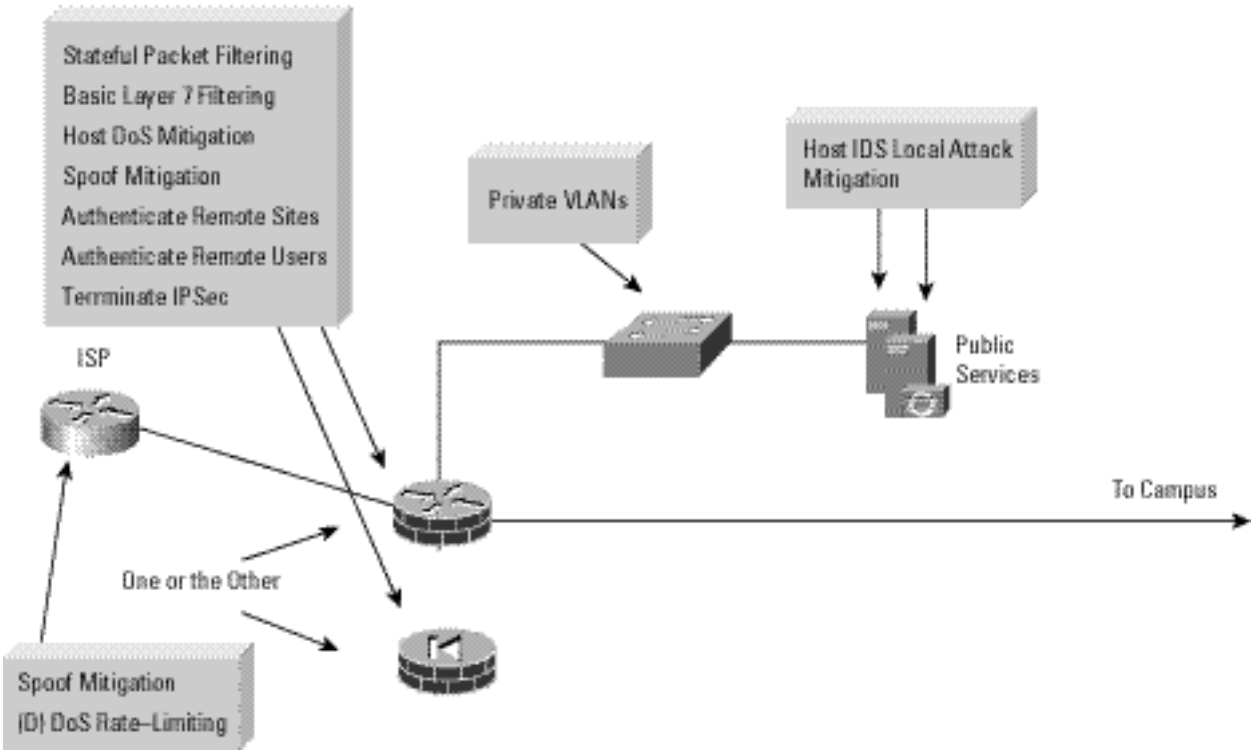

Design Guidelines

This module represents the ultimate in scaled-down security-conscious network design, where all the security and VPN services are compressed into a single box. Two principal alternatives come into play when deciding how to implement this functionality. The first is to use a router with firewall and VPN functionality. This setup yields the greatest flexibility for the small network because the router will support all the advanced services (QoS, routing, multi-protocol support, etc.) that may be necessary in today's networks. As an alternative, a dedicated firewall may be used instead of the router. This setup places some restrictions on the deployment. First, firewalls are generally Ethernet only, requiring some conversion to the appropriate WAN protocol. In today's environments, most cable and *digital-subscriber-line* (DSL) routers/modems are provided by the service provider and can be used to connect to the firewall over Ethernet. If WAN connectivity on the device is required (such as with a DS1 circuit from a telco provider), then a router must be used. Using a dedicated firewall does have the advantage of easier configuration of security services, and a dedicated firewall can provide improved performance when doing firewall functions. Whatever the selection of device, stateful inspection is used to examine traffic in all directions, ensuring that only legitimate traffic crosses the firewall. Before the traffic even reaches the firewall, ideally, some security filtering has already occurred at the ISP. Remember that routers tend to start out permitting traffic, whereas firewalls tend to deny traffic by default.

Starting at the customer-edge router in the ISP, the egress out of the ISP rate limits nonessential traffic that exceeds prespecified thresholds in order to mitigate against DDoS attacks. Also at the egress of the ISP router, RFC 1918 and RFC 2827 filtering mitigate against source-address spoofing of local networks and private address ranges.

At the ingress of the firewall, RFC 1918 and RFC 2827 filtering is first provided as a verification of the ISP's filtering. In addition, because of the enormous security threat that fragmented packets create, the firewall is configured to drop most fragmented packets that should not generally be seen for standard traffic types on the Internet. Any legitimate traffic lost because of this filtering is considered acceptable when compared to the risk of allowing such traffic. Traffic destined to the firewall itself from the outside is limited to IPSec traffic and any necessary protocols for routing.

The firewall provides connection-state enforcement and detailed filtering for sessions initiated through the firewall. Publicly addressable servers have some protection against TCP SYN floods through mechanisms such as the use of half-open connection limits on the firewall. From a filtering standpoint, in addition to limiting traffic on the public services segment to relevant addresses and ports, filtering in the opposite direction also takes place. If an attack compromises one of the public servers (by circumventing the firewall and host-based IDS), that server should not be able to further attack the network. To mitigate against this type of attack, specific filtering prevents any unauthorized requests from being generated by the public servers to any other location. As an example, the Web server should be filtered so that it cannot originate requests of its own, but merely respond to requests from clients. This setup helps prevent a hacker from downloading additional utilities to the compromised box after the initial attack. It also helps stop unwanted sessions from being triggered by the hacker during the primary attack. An attack that generates an xterm from the Web server through the firewall to the hacker's machine is an example of such an attack. In addition, private VLANs on the DMZ switch prevent a compromised public server from attacking other servers on the same segment. This traffic is not even detected by the firewall, a fact that explains why private VLANs are critical.

From a host perspective, each of the servers on the public services segment has host intrusion detection software to monitor against any rogue activity at the OS level, as well as activity in common server applications (HTTP, FTP, SMTP, and so forth). The DNS host should be locked down to respond only to desired commands and eliminate any unnecessary responses that might assist hackers in network reconnaissance. This includes preventing zone transfers from anywhere except legitimate secondary DNS servers. For mail services, the firewall itself filters SMTP messages at Layer 7 to allow only necessary commands to the mail server.

Firewalls and firewall routers generally have some limited NIDS capability within their security functions. This capability will affect the performance of the device, but does provide some additional attack visibility in the event you are under attack. Remember that you are trading performance for attack visibility. Many of these attacks can be dropped without the use of IDS, but the monitoring station will not be aware of the specific attack being launched.

The VPN connectivity is provided through the firewall or firewall/router. Remote sites authenticate each other with pre-shared keys and remote users are authenticated through the access control server in the campus module.

Alternatives

Any deviation from this design would be geared toward increasing the capacity of the network, or separating the various security functions onto distinct devices. In doing this, the design will start to look more and more like the medium network design discussed later in this document. A first step rather than adopting the complete medium design might be the addition of a dedicated remote access VPN concentrator to increase the manageability of the remote-user community.

Campus Module

The campus module contains end-user workstations, corporate intranet servers, management servers, and the associated Layer 2 infrastructure required to support the devices. Within the small network design, this Layer 2 functionality has been combined into a single switch.

Key Devices

- *Layer 2 switching (with private VLAN support)*—Provides Layer 2 services to user workstations
- *Corporate servers*—Provides e-mail (SMTP and POP3) services to internal users, as well as delivering file, print, and DNS services to workstations
- *User workstations*—Provide data services to authorized users on the network
- *Management host*—Provides HIDS, syslog, TACACS+/*Remote Access Dial-In User Service* (RADIUS), and general configuration management

Figure 4 Detailed Model of Small Network Campus Module

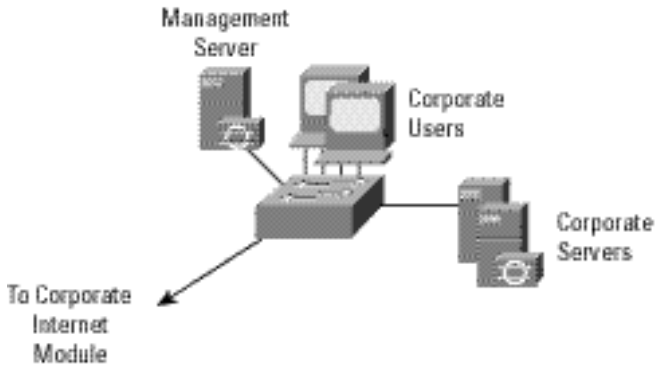

Threats Mitigated

- *Packet sniffers*—A switched infrastructure limits the effectiveness of sniffing
- *Virus and Trojan-horse applications*—Host-based virus scanning prevents most viruses and many Trojan horses
- *Unauthorized access*—This type of access is mitigated through the use of host-based intrusion detection and application access control
- *Application layer attacks*—Operating systems, devices, and applications are kept up-to-date with the latest security fixes, and they are protected by HIDS
- *Trust exploitation*—Private VLANs prevent hosts on the same subnet from communicating unless necessary
- *Port redirection*—HIDS prevents port redirection agents from being installed

Figure 5 Small Network Attack Mitigation Roles for Campus Module

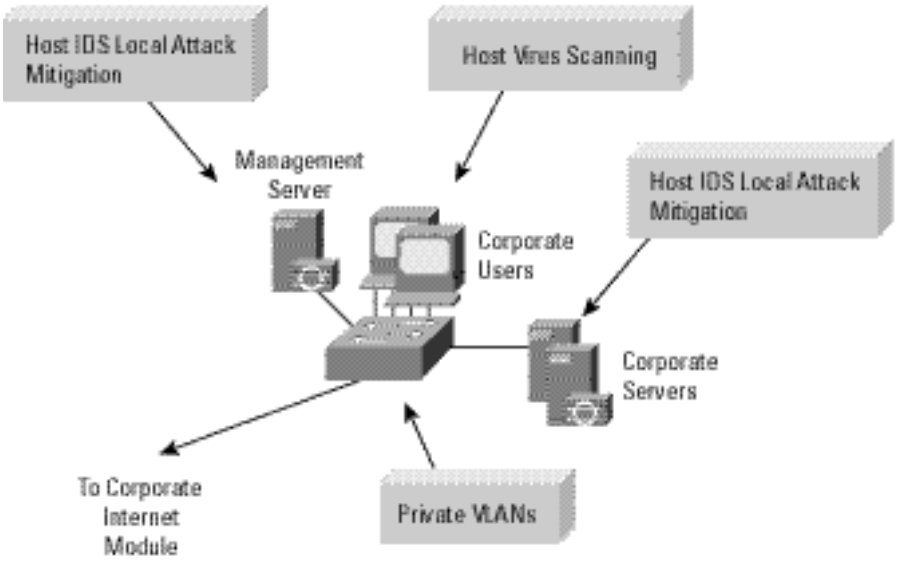

Design Guidelines

The primary functions of the campus switch are to switch production and management traffic and to provide connectivity for the corporate and management servers and users. Within the switch, private VLANs can be enabled in order to mitigate trust-exploitation attacks between the devices. For instance, the corporate users might need to be able to talk to the corporate servers but may not have any requirement to communicate with each other.

Because there are no Layer 3 services within the campus module, it is important to note that this design places an increased emphasis on application and host security because of the open nature of the internal network. Therefore, HIDS was also installed on key systems within the campus, including the corporate servers and management systems.

Alternatives

Setting a small filtering router or firewall between the management stations and the rest of the network can improve overall security. This setup will allow management traffic to flow only in the specific direction deemed necessary by the administrators. If the level of trust within the organization is high, HIDS can potentially be eliminated, though this is not recommended.

Branch versus Standalone Considerations

When configured in a branch role, remote access VPN functionality is not required as this is generally provided by the corporate headquarters. Futher, the management hosts will typically be located at the central site, a setup that requires the management traffic to traverse the site-to-site VPn connection back to corporate headquarters.

Medium Network Design

The SAFE medium network design consists of three modules: the corporate Internet module, the campus module, and the WAN module. As in the small network design, the corporate Internet module has the connection to the Internet and terminates VPN and public-services (DNS, HTTP, FTP, and SMTP) traffic. Dial-in traffic also terminates at the corporate Internet module. The campus module contains the Layer 2 and Layer 3 switching infrastructure along with all the corporate users, management servers, and intranet servers. From a WAN perspective, there are two options for the remote sites connecting into the medium design. The first is a private WAN connection using the WAN module, the second is an IPSec VPN into the corporate Internet module. Most of the discussion in this design is based on the medium network operating as the headend for a corporation. Specific design changes when used as a branch are also included.

Figure 6 Detailed Model of Medium Network

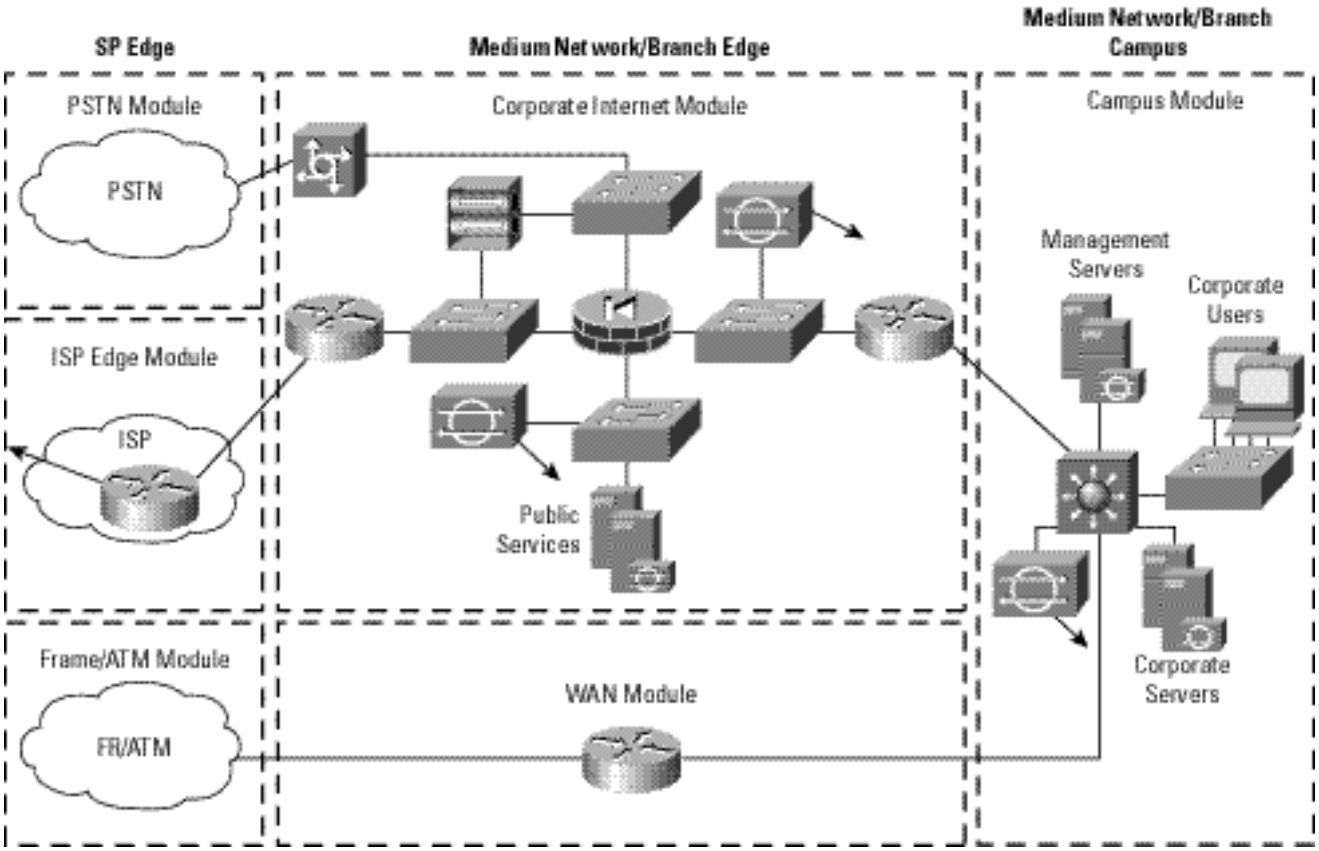

Corporate Internet Module

The goal of the corporate Internet module is to provide internal users with connectivity to Internet services and Internet users access to information on the public servers (HTTP, FTP, SMTP, and DNS). Additionally, this module terminates VPN traffic from remote users and remote sites as well as traffic from traditional dial-in users. The corporate Internet module is not designed to serve e-commerce type applications. Refer to the section "E-Commerce Module" in SAFE Enterprise for more details on providing Internet commerce.

Key Devices

- *Dial-in server*—Authenticates individual remote users and terminates their analog connections
- *DNS server*—Serves as authoritative external DNS server for the medium network; relays internal requests to the Internet
- *FTP/HTTP server*—Provides public information about the organization
- *Firewall*—Provides network-level protection of resources and stateful filtering of traffic; provides differentiated security for remote access users; authenticates trusted remote sites and provides connectivity using IPSec tunnels
- *Layer 2 switches (with private VLAN support)*—Provides Layer 2 connectivity for devices
- *NIDS appliance*—Provides Layer 4-to-Layer 7 monitoring of key network segments in the module
- *SMTP server*—Acts as a relay between the Internet and the intranet mail servers; inspects content
- *VPN concentrator*—Authenticates individual remote users and terminates their IPSec tunnels
- *Edge Router*—Provides basic filtering and layer 3 connectivity to the Internet

Figure 7 Detailed Model of Medium Network Corporate Internet Module

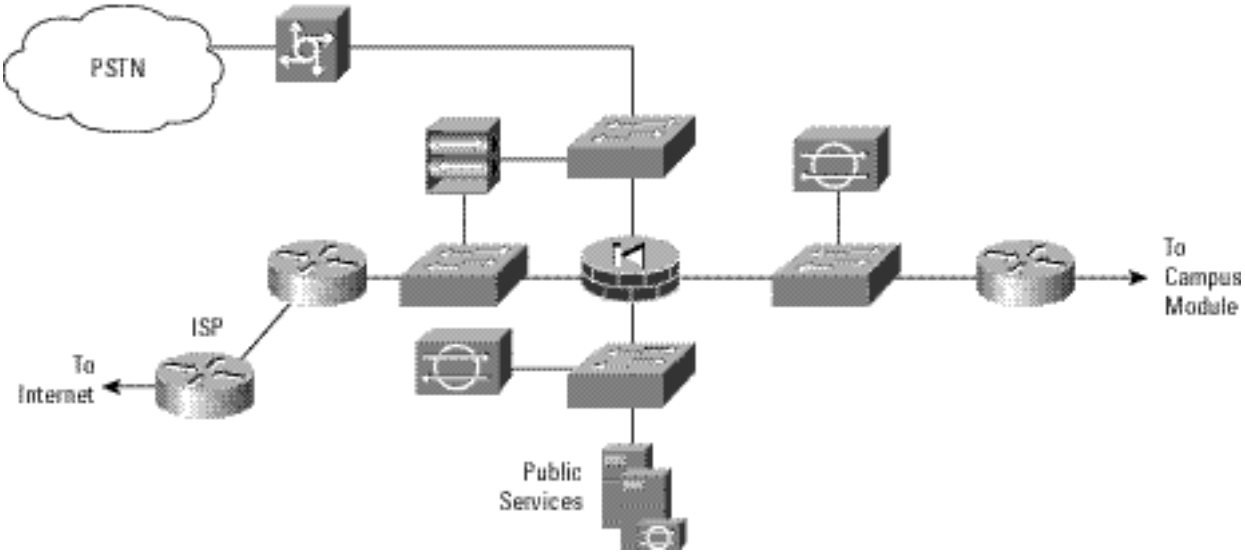

Threats Mitigated

The publicly addressable servers are likely points of attack within this module. The following are expected threats:

- *Unauthorized access*—Mitigated through filtering at the ISP, edge router, and corporate firewall
- *Application layer attacks*—Mitigated through IDS at the host and network levels
- *Virus and Trojan horse attacks*—Mitigated through e-mail content filtering, HIDS, and host-based virus scanning
- *Password attacks*—Limited services available to brute force; OS and IDS can detect the threat
- *Denial of service*—CAR at ISP edge and TCP setup controls at firewall
- *IP spoofing*—RFC 2827 and 1918 filtering at ISP edge and medium network edge router
- *Packet sniffers*—Switched infrastructure and host IDS to limit exposure
- *Network reconnaissance*—IDS detects reconnaissance, protocols filtered to limit effectiveness
- *Trust exploitation*—Restrictive trust model and private VLANs to limit trust-based attacks
- *Port redirection*—Restrictive filtering and host IDS to limit attacks

The remote access and site-to-site VPN services are also points of attack within this module. The following are expected threats:

- *Network topology discovery*—Access control lists (ACLs) on the ingress router limit access to the VPN concentrator and firewall (when used to terminate IPSec tunnels from remote sites) to Internet Key Exchange (IKE) and Encapsulating Security Payload (ESP) from the Internet
- *Password attack*—*One-time passwords* (OTP) mitigate brute force password attacks
- *Unauthorized access*—Firewall services after packet decryption prevent traffic on unauthorized ports
- *Man-in-the-middle attacks*—These attacks are mitigated through encrypted remote traffic
- *Packet sniffers*—A switched infrastructure limits the effectiveness of sniffing

Figure 8 Medium Network Attack Mitigation Roles for Corporate Internet Module

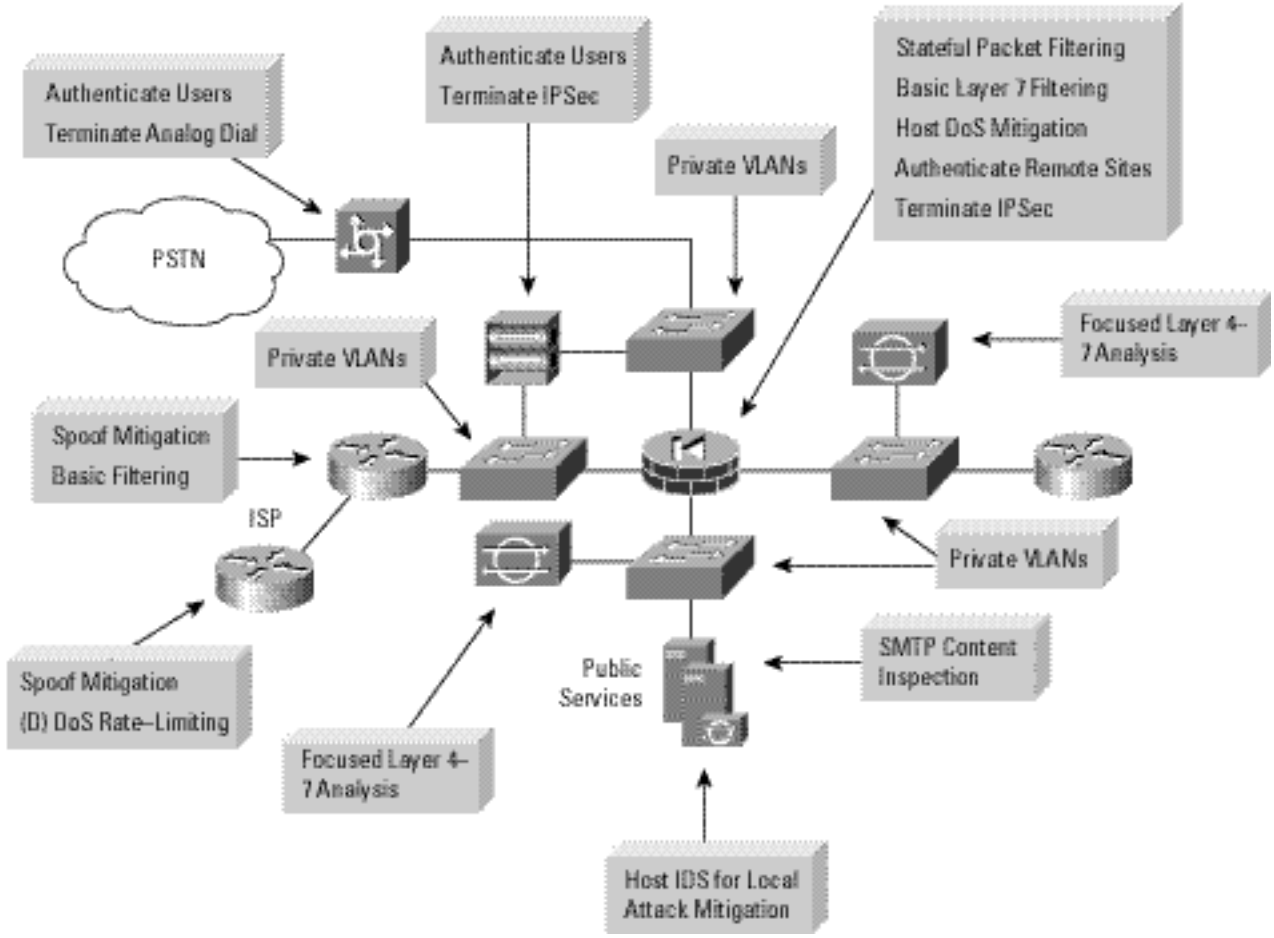

Design Guidelines

The following sections detail the functionality of each of the devices within the corporate Internet module.

ISP Router

The primary function of the customer-edge router in the ISP is to provide connectivity to the Internet or ISP network. The egress out of the ISP router rate limits nonessential traffic that exceeds prespecified thresholds in order to mitigate against DDoS attacks. Finally, at the egress of the ISP router, RFC 1918 and RFC 2827 filtering is configured to mitigate against source-address spoofing of local networks and private address ranges.

Edge Router

The function of the edge router on the medium network is to provide the demarcation point between the ISP network and the medium network. At the ingress of the edge router on the medium network, basic filtering limits access to allow only expected IP traffic, providing a coarse filter for the most basic attacks. RFC 1918 and RFC 2827 filtering is also provided here as a verification of the ISP's filtering. In addition, because of the enormous security threat that they create, the router is configured to drop most fragmented packets that should not generally be seen for standard traffic types on the Internet. Any legitimate traffic lost because of this filtering is considered acceptable when compared to the risk of allowing such traffic. Finally, any IPSec traffic destined for the VPN concentrator or the firewall is allowed through. Filtering on the router is configured to allow only IKE and IPSec traffic to reach the VPN concentrator or firewall. Because with remote access VPNs the IP address of the remote system is not generally known, the filtering can be specified only to the headend peer (VPN concentrator) with which the remote users are communicating. With site-to-site VPNs, the IP address of the remote site is usually known; therefore, filtering may be specified for VPN traffic to and from both peers.

Firewall

The primary function of the firewall is to provide connection-state enforcement and detailed filtering for sessions initiated through the firewall. The firewall also acts as a termination point for site-to-site IPSec VPN tunnels for both remote site production and remote site management traffic. There are multiple segments off the firewall. The first is the public services segment, which contains all the publicly adressable hosts. The second is for remote access VPN and dial-in, which is discussed later. Publicly addressable servers have some protection against TCP SYN floods through mechanisms such as the use of half-open connection limits on the firewall. From a filtering standpoint, in addition to limiting traffic on the public services segment to relevant addresses and ports, filtering in the opposite direction also occurs. If an attack compromises one of the public servers (by circumventing the firewall, HIDS, and NIDS), that server should not be able to further attack the network. To mitigate against this type of attack, specific filtering prevents any unauthorized requests from being generated by the public servers to any other location. As an example, the Web server should be filtered so that it cannot originate requests of its own, but merely respond to requests from clients. This setup helps prevent a hacker from downloading additional utilities to the compromised box after the initial attack. It also helps stop unwanted sessions from being triggered by the hacker during the primary attack. An attack that generates an xterm from the Web server through the firewall to the hacker's machine is an example of such an attack. In addition, private VLANs prevent a compromised public server from attacking other servers on the same segment. This traffic is not even detected by the firewall, a fact that explains why private VLANs are critical.

Intrusion Detection

The public services segment includes a NIDS appliance. Its primary function is to detect attacks on ports that the firewall is configured to permit. These most often are application layer attacks against specific services. The NIDS on the public services segment should be set in a restrictive stance because signatures matched here have successfully passed through the firewall already. Each of the servers has HIDS on it as well. The primary function of HIDS is to monitor against any rogue activity that occurs at the OS level as well as in common server applications (HTTP, FTP, SMTP, and so forth). DNS should be locked down to respond only to desired commands and eliminate any unnecessary responses that might assist hackers in network reconnaissance. This includes preventing zone transfers from anywhere except legitimate secondary DNS servers. The SMTP server includes mail-content inspection services that mitigate against virus and Trojan horse-type attacks generated against the internal network that are usually introduced through the mail system. The firewall itself filters SMTP messages at Layer 7 to allow only necessary commands to the mail server.

The NIDS appliance between the private interface of the firewall and the internal router provides a final analysis of attacks. Very few attacks should be detected on this segment because only responses to initiated requests, a few select ports from the public services segment, and traffic from the remote access segment are allowed to the inside. Only sophisticated attacks should be seen on this segment because they could mean that a system on the public services segment has been compromised and the hacker is attempting to take advantage of this foothold to attack the internal network. For example, if the public

SMTP server were compromised, a hacker might try to attack the internal mail server over TCP port 25, which is permitted to allow mail transfer between the two hosts. If attacks are seen on this segment, the responses to those attacks should be more severe than those on other segments because they probably indicate that a compromise has already occurred. The use of TCP resets or shunning to thwart, for example, the SMTP attack mentioned above, should be seriously considered.

Remote Access VPN

The primary function of the remote access VPN concentrator is to provide secure connectivity to the medium network for remote users. The VPN concentrator initiates a session with an access control server on the internal network to authenticate users before granting them access to the network. The access control server then queries a *one-time password* (OTP) system to validate the user's authentication credentials. Via IPSec policy sent from the concentrator to the client, users are prevented from enabling split tunneling, thereby forcing users to access the Internet via the corporate connection. The IPSec parameters used are *Triple Data Encryption Standard* (3DES) for encryption and *secure hash algorithm/hash-based message authentication code* (SHA/HMAC) for data integrity. Following termination of the VPN tunnel, traffic is sent through a firewall to ensure that VPN users are appropriately filtered. This setup also allows IDS shunning to take place on the firewall. This scenario is in contrast to many deployments today that place the firewall in front of the VPN device. When placed in front, no visibility into the specific types of user traffic is possible because the traffic is still encrypted.

Dial-In Access Users

The traditional dial-in users are terminated on an access router with built-in modems. When the Layer 2 connection is established between the user and the server, three-way *Challenge Handshake Authentication Protocol* (CHAP) is used to authenticate the user. As in the remote access VPN service, the *authentication, authorization, and accounting* (AAA) server is used for authentication. When authenticated, the users are provided with IP addresses from an IP pool.

Layer 2 Switches

The primary function of the switches within the corporate Internet module is to provide Layer 2 connectivity between the various devices within the module. Separate switches, rather that a single switch with multiple VLANs, were chosen in order to provide physical separation between the outside segment, public services segment, VPN segment, and inside segment. This setup mitigates against any potential misconfiguration on a switch that could compromise security. Additionally, each of the switches runs the private VLAN feature, a setup that helps mitigate against attacks based on trust exploitation.

Inside Router

The primary function of the inside router is to provide Layer 3 separation and routing between the corporate Internet module and the campus module. This device functions strictly as a router with no access lists restricting traffic across either interface. Because routing information itself can be used in a DoS attack, authentication of routing updates between devices may be utilized in order to mitigate against such an attack. This router provides a final point of demarcation between the routed intranet and the outside world. Because most firewalls are configured without routing protocols, it is important to provide a point of routing within the corporate Internet module that does not rely on the rest of the network.

Alternatives

This module has several alternative designs. Rather than implementing basic filtering on the edge router to the medium network, a network administrator may choose to implement a stateful firewall on this device as well. Having two stateful firewalls provides more of a defense in depth approach to security within the module. Depending on the network administrator's attitude toward attack awareness, a NIDS appliance might be required in front of the firewall. With the appropriate basic filters, the IDS outside the firewall can provide important alarm information that would otherwise be dropped by the firewall Because the amount of alarms generated on this segment is probably large, alarms generated here should have a lower severity than alarms generated behind a firewall. Also, consider logging alarms from this segment to a separate management station to ensure that legitimate alarms from other segments get the appropriate attention. With the visibility that NIDS outside the firewall provides, evaluation of the attack types your organization is attracting can be better seen. In addition, evaluation of the effectiveness of ISP and enterprise edge filters can be performed.

Two other alternatives are available. First is the elimination of the router between the firewall and the campus module. Although its functions can be integrated into the campus module Layer 3 switch, this setup would eliminate the ability of the corporate Internet module to function without relying on Layer 3 services from another area of the network. Second is the addition of content inspection beyond the mail-content inspection already specified. For example, a URL filtering server could be placed on the public services segment to filter the types of Web pages that employees can access.

Campus Module

The campus module contains end-user workstations, corporate intranet servers, management servers, and the associated Layer 2 and Layer 3 infrastructure required to support the devices. All the campus modules from SAFE Enterprise have been combined into a single module. This setup more accurately reflects the smaller size of medium networks, and reduces the overall cost of the design. As in the corporate Internet module, the redundancy that would normally be found in an enterprise design has been removed from the medium network design.

Key Devices

- *Layer 3 switch*—Route and switch production and management traffic within the campus module, provide distribution layer services to the building switches, and support advanced services such as traffic filtering
- *Layer 2 switches (with private VLAN support)*—Provides Layer 2 services to user workstations
- *Corporate servers*—Provides e-mail (SMTP and POP3) services to internal users, as well as delivering file, print, and DNS services to workstations
- *User workstations*—Provide data services to authorized users on the network
- *SNMP management host*—Provides SNMP management for devices
- *NIDS host*—Provides alarm aggregation for all NIDS devices in the network
- ·*Syslog host(s)*—Aggregates log information for firewall and NIDS hosts
- *Access control server*—Delivers authentication services to the network devices
- *One-time Password (OTP) Server*—Authorizes one-time password information relayed from the access control server
- *System admin host*—Provides configuration, software, and content changes on devices
- *NIDS appliance*—Provides Layer 4-to-Layer 7 monitoring of key network segments in the module

Figure 9 Detailed Model of Medium Network Campus Module

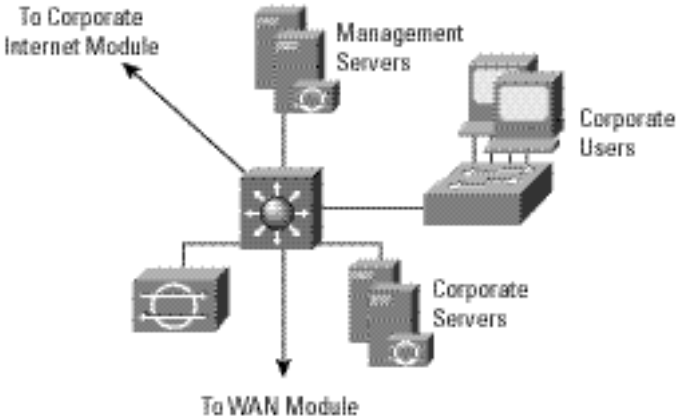

Threats Mitigated

- *Packet sniffers*—A switched infrastructure limits the effectiveness of sniffing
- *Virus and Trojan horse applications*—Host-based virus scanning prevents most viruses and many Trojan horses
- *Unauthorized access*—These types of attacks are mitigated through the use of host-based intrusion detection and access control
- *Password Attacks*—The access control server allows for strong two-factor authentication for key applications
- *Application layer attacks*—Operating systems, devices, and applications are kept up-to-date with the latest security fixes, and they are protected by HIDS
- *IP spoofing*—RFC 2827 filtering prevents source-address spoofing
- *Trust exploitation*—Trust arrangements are very explicit; private VLANs prevent hosts on the same subnet from communicating unless necessary
- *Port redirection*—HIDS prevents port redirection agents from being installed

Figure 10 Medium Network Attack Mitigation Roles for Campus Module

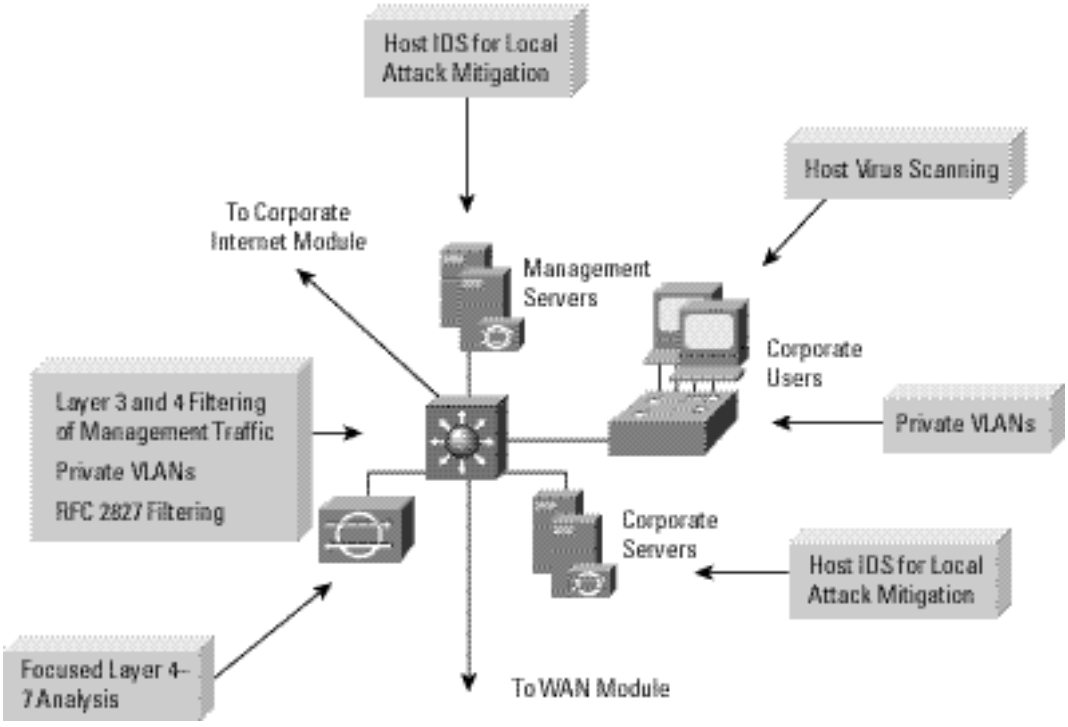

Design Guidelines

The following sections detail the functionality of each of the devices within the campus module.

Core Switch

The primary function of the core switch is to provide routing and switching for production and management traffic, distribution layer services (routing, *quality of service* [QoS], and access control) for the building switches, connectivity for the corporate and management servers, and advanced services such as traffic filtering between the subnets. A Layer 3 switch was chosen instead of a Layer 2 switch in order to provide separate VLANs for the corporate server segment(s), the management server segment, the corporate user segment(s), and connectivity to the WAN module and to the corporate

Internet module. The Layer 3 switch provides a line of defense and prevention against internally originated attacks. It can mitigate the chance of a department accessing confidential information on another department's server through the use of access control. For example, a network that contains marketing and research and development might segment off the R&D server to a specific VLAN and filter access to it, ensuring that only R&D staff have access to it. For performance reasons, it is important that this access control be implemented on a hardware platform that can deliver filtered traffic at near wire rates. This setup generally dictates the use of Layer 3 switching, as opposed to more traditional dedicated routing devices. This same access control can also prevent local source-address spoofing through the use of RFC 2827 filtering. RFC 2827 filtering should be implemented on the corporate user and corporate intranet server VLANs.

Within each of the VLANs, private VLANs can be utilized in order to mitigate trust-exploitation attacks between the devices. For instance, within the corporate server segment, the individual servers may not have any requirement to communicate with each other. They need to communicate only with devices connected to the corporate user segment(s).

In order to provide a further line of defense for the management servers, extensive Layer 3 and Layer 4 filtering is configured outbound on the VLAN interface connecting to the management server segment. The ACL limits connectivity to and from the management servers only to those devices (via IP addresses) under their control, and only for those protocols/services (via port number) that are required. This also includes access control for management traffic destined for the remote site devices. This traffic is encrypted by the firewall and sent to the remote sites. Access to the managed devices is further controlled by allowing only established connections back through the ACL.

Building Switches

The primary function of the building switches within the campus module is to provide Layer 2 services to corporate user workstations. Private VLANs are implemented on the building switches in order to mitigate against a trust-exploitation attack, because individual end-user workstations generally do not have a requirement to communicate with each other. In addition to the network security guidelines described in the switch security axiom, host-based virus scanning is also implemented at the workstation level.

Intrusion Detection

The campus module also includes an NIDS appliance. The switch port that connects to the NIDS appliance is configured such that traffic from all VLANs that require monitoring is mirrored to the monitoring port of the appliance. Very few attacks should be detected here because this NIDS appliance provides analysis against attacks that may originate from within the campus module itself. For instance, if a user workstation were compromised because of an unknown modem connection to that host, the NIDS could detect suspicious activity originating from within the campus. Other internal attacks could originate from disgruntled employees, workstations left where unauthorized people can gain access to them, or Trojan horse applications inadvertently loaded on portable PCs. Each of the corporate intranet and management servers also has HIDS installed.

Alternatives

If the medium network is small enough, the functionality of the building switches can be rolled into the core switch, and the building switches can be eliminated. In this case, the end-user workstations would be connected directly to the core switch. Private VLAN functionality would be implemented on the core switch in order to mitigate against trust-exploitation attacks. If the performance requirements of the internal network are not high, a separate router and Layer 2 switch could be used for the core and distribution instead of the higher-performing Layer 3 switch.

If desired, the separate NIDS appliance can be replaced with an integrated IDS module that fits into the core switch. This setup provides higher traffic throughput into the IDS module because it sits on the backplane of the switch, rather than being connected via a single 10/100-Mbps Ethernet port. ACLs on the switch can be used to control what traffic is sent to the IDS module.

WAN Module

The WAN module is included only when connections to remote locations over a private network are required. This requirement may occur when stringent QoS requirements cannot be met by an IPSec VPN, or when legacy WAN connections are in place without a compelling cost justification to migrate to IPSec.

Key Devices

- *IOS Router*—Provides routing, access-control, and QoS mechanisms to remote locations

Figure 11 Detailed Model of Medium Network WAN Module

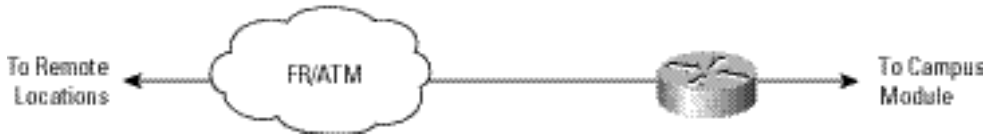

Threats Mitigated

- *IP spoofing*—IP spoofing can be mitigated through Layer 3 filtering
- *Unauthorized access*—Simple access control on the router can limit the types of protocols to which branches have access

Figure 12 Attack Mitigation Roles for WAN Module

Design Guidelines

The amount of security placed in the WAN module will depend on the level of trust for the remote sites and the ISP to which you are connecting. Security is provided by using IOS security features. In this design, inbound access lists applied to the serial interface are used to block all unwanted traffic from accessing the medium network. Inbound access lists applied to the Ethernet interface can be used to further limit what traffic passes from the medium network back to the remote sites.

Alternatives

Some organizations that are very concerned about information privacy encrypt traffic across their classic WAN links. Similar to site-to-site VPNs, IPSec can be used to achieve this level of information privacy. Additionally, running a firewall on the WAN router can provide additional access control options when compared with the basic ACLs used in the SAFE design.

Branch versus Headend Considerations

When configured as a branch, several components within the medium design can be eliminated. The first consideration is whether or not an organization wants to connect to the corporate headquarters over a private WAN link or an IPSec VPN. Some reasons for choosing private WAN include more granular QoS support, multicast support, reliability of the network infastructure, or the requirement for non-IP traffic. Remember that when you are using IPSec over *generic routing encapsulation* (GRE) (discussed in SAFE Enterprise), multicast and non-IP traffic can be supported in a VPN environment. There are several reasons for choosing IPSec VPNs instead of a private WAN connection. First, an IPSec VPN over the Internet can provide local Internet access for all remote sites, thus saving bandwidth (and cost) at the headend. Also, in many domestic and most international applications, IPSec VPNs provide a significant cost savings over private WAN connections.

If a private WAN link is chosen for the medium network when operating as a branch, the entire corporate Internet module is not needed (unless local Internet access is desired from the branch). On the other hand, if an IPSec VPN is chosen, the WAN module is not needed. In addition to the WAN module, a branch medium design may not need a VPN concentrator or dial-access router for remote access services if the services are provided by the corporate headquarters.

From a management perspective, configuration and security management of the medium network would be done from the corporate headquarters management module (assuming centralized IT resources). If a private WAN link is chosen for the intersite connectivity, management traffic can easily flow across the WAN module and into the devices that require management. When an IPSec VPN is chosen for intersite connectivity, most management traffic can flow as it did when a private WAN link was used. Some devices, such as the edge router on the outside of the firewall, will not be part of the IPSec tunnel and will need to be managed in another way. This setup could include a separate IPSec tunnel to the device, or relying on application layer encryption (SSH) for configuration changes to those devices. As was mentioned in the axioms, not all management protocols have an associated secure variant.

Remote-User Design

This section discusses four different options for providing remote-user connectivity within the SAFE design. Remote connectivity applies to both mobile workers and home-office workers. The primary focus of these designs is providing connectivity from the remote site to the corporate headquarters and through some means, the Internet. The following four options include software-only, software-with-hardware, and hardware-only solutions:

- *Software access*—Remote user with a software VPN client and personal firewall software on the PC
- *Remote-site firewall option*—Remote site is protected with a dedicated firewall that provides firewalling and IPSec VPN connectivity to corporate headquarters; WAN connectivity is provided via an ISP-provided broadband access device (i.e. DSL or cable modem).
- *Hardware VPN client option*—Remote site using a dedicated hardware VPN client that provides IPSec VPN connectivity to corporate headquarters; WAN connectivity is provided via an ISP-provided broadband access device
- *Remote-site router option*—Remote site using a router that provides both firewalling and IPSec VPN connectivity to corporate headquarters. This router can either provide direct broadband access or go through and ISP-provided broadband access device.

Each of these designs is discussed further in the design guidelines section below. All discussions assume that the connectivity is through the Internet. If private WAN connectivity (ISDN, private DSL, etc.) is used instead, encryption of the traffic may not be required. Keep in mind that with any remote site option, the security perimeter of your organization is extended to include those remote sites.

Key Devices

- *Broadband access device*—Provides access to the broadband network (DSL, cable, and so on)
- *Firewall with VPN support*—Provides secure end-to-end encrypted tunnels between the remote site and the corporate headend; provides network-level protection of remote-site resources and stateful filtering of traffic
- *Layer 2 hub*—Provides connectivity for devices within the remote site (can be integrated into the firewall or hardware VPN client)
- *Personal firewall software*—Provides device-level protection for individual PCs
- *Router with firewall and VPN support*—Provides secure end-to-end encrypted tunnels between the remote site and the corporate headend; provides network-level protection of remote-site resources and stateful filtering of traffic; can provide advanced services such as voice or QoS.
- *VPN software client*—Provides secure end-to-end encrypted tunnels between individual PCs and the corporate headend
- *VPN hardware client*—Provides secure end-to-end encrypted tunnels between the remote site and the corporate headend

Figure 13 Detailed Model of Remote-User Configuration

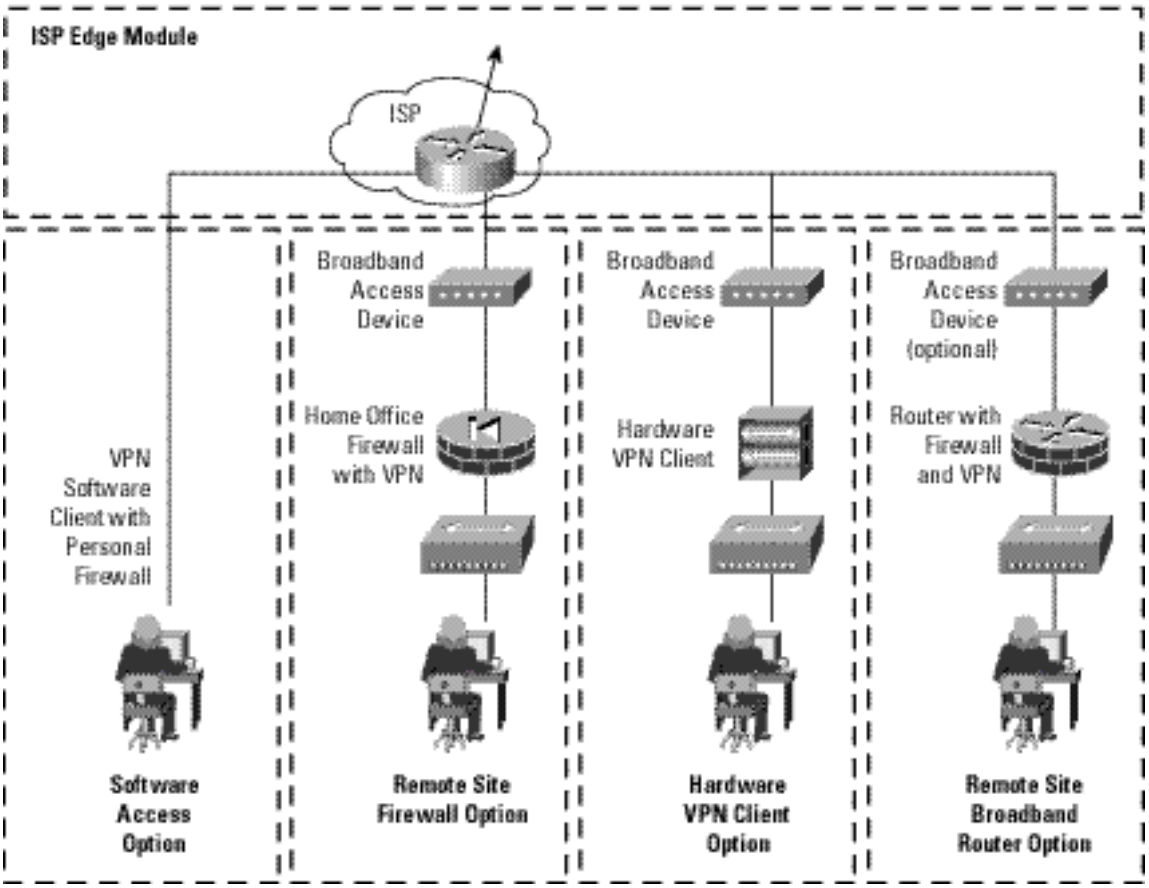

Threats Mitigated

- *Unauthorized access*—Mitigated through filtering and stateful inspection of sessions at the remote-site firewall or router, or through application access control via personal firewall software.
- *Network reconnaissance*—Protocols filtered at remote-site device to limit effectiveness
- *Virus and Trojan horse attacks*—Mitigated through virus scanning at the host level
- *IP spoofing*—Mitigated through RFC 2827 and 1918 filtering at ISP edge and remote-site device
- *Man-in-the-middle attacks*—Mitigated through encrypted remote traffic

Figure 14 Remote-User Design Attack Mitigation Roles

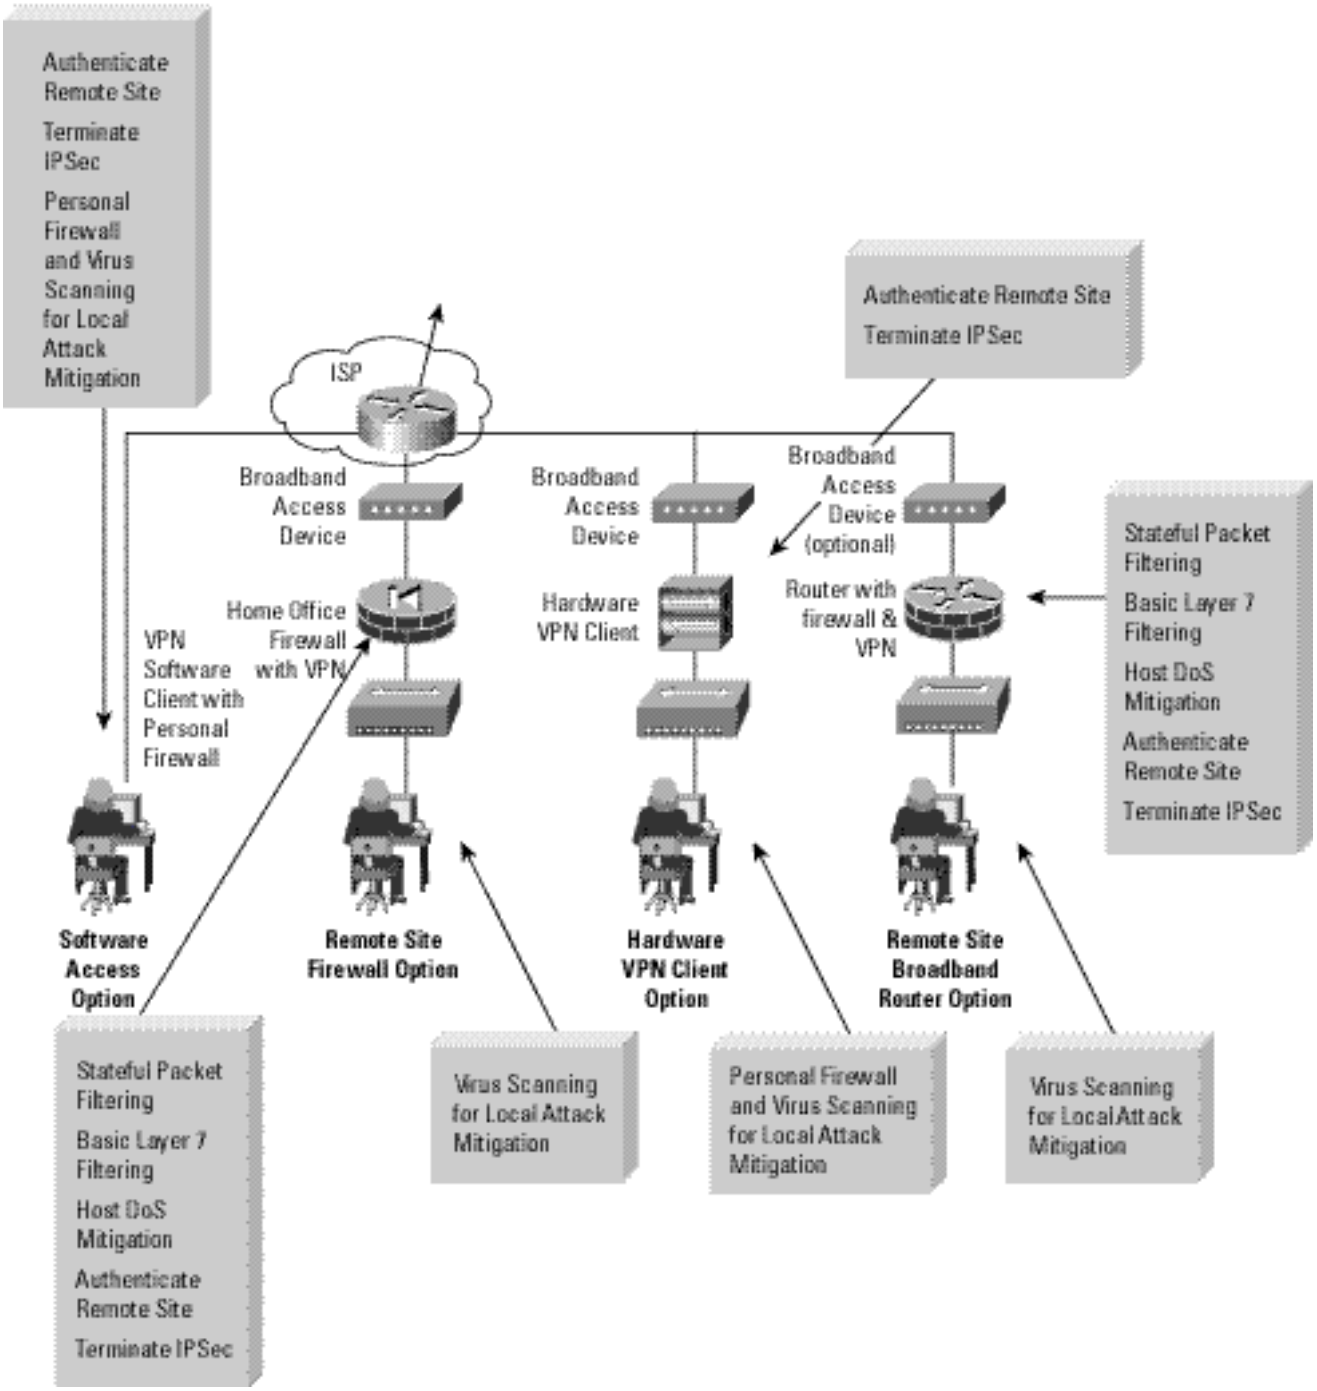

Design Guidelines

The following sections detail the functionality of each of the remote-user connectivity options.

Software Access Option

The software access option is geared toward the mobile worker as well as the home-office worker. All the remote user requires is a PC with VPN client software and connectivity to the Internet or ISP network via a dial-in or Ethernet connection. The primary function of the VPN software client is to establish a secure, encrypted tunnel from the client device to a VPN

headend device. Access and authorization to the network are controlled from the headquarters location when filtering takes place on the firewall and on the client itself if access rights are pushed down via policy. The remote user is first authenticated, and then receives IP parameters such as a virtual IP address, which is used for all VPN traffic, and the location of name servers (DNS and *Windows Internet Name Service* [WINS]). Split tunneling can also be enabled or disabled via the central site. For the SAFE design, split tunneling was disabled, making it necessary for all remote users to access the Internet via the corporate connection when they have a VPN tunnel established. Because the remote user may not always want the VPN tunnel established when connected to the Internet or ISP network, personal firewall software is recommended to mitigate against unauthorized access to the PC. Virus-scanning software is also recommended to mitigate against viruses and Trojan horse programs infecting the PC.

Remote-Site Firewall Option

The remote-site firewall option is geared toward the home-office worker, or potentially a very small branch office. With this option, it is expected that the remote site has some form of broadband access available from a service provider. The firewall is installed behind the DSL or cable modem.

The primary function of the firewall is to establish the secure, encrypted tunnel between itself and a VPN headend device, as well as providing connection-state enforcement and detailed filtering for sessions initiated through it. Individual PCs on the remote-site network do not need VPN client software to access corporate resources. Additionally, because the stateful firewall protects access to the Internet, personal firewall software isn't necessarily required on the individual PCs. However, if the network administrator wants an additional level of security, personal firewall software can also be implemented on remote-site PCs. This setup may be useful if the home worker also travels and connects to the Internet directly over some public network. Because we have a stateful firewall protecting the hosts, the remote site can have direct access to the Internet, rather than passing all traffic back through the corporate headquarters. Unless NAT is used when communicating with the headquarters, the IP addresses of the remote-site devices should be assigned in such a manner as to not overlap addressing space in the headquarters location or another remote site. Remote-site devices that require direct access to the Internet will require address translation to a registered address. This address translation can be achieved by translating all Internet-bound sessions to the public IP address of the firewall itself.

Access and authorization to the corporate network and the Internet are controlled by the configuration of both the remote-site firewall and the VPN headend device. Configuration and security management of the remote-site firewall can be achieved via an IPSec tunnel from the public side of the firewall back to the corporate headquarters. This setup ensures that the remote-site user(s) is not required to perform any configuration changes on the home-office firewall. Authentication should be set up on the firewall to prevent a local user from inadvertently modifying his/her firewall configuration and thereby compromise the security policy of that device. Individual users at the remote site who access the corporate network are not authenticated with this option. Instead, the remote-site firewall and VPN headend utilize device authentication.

Virus-scanning software is still recommended to mitigate against viruses and Trojan horse programs infecting individual PCs at the remote site-just like all the PCs in the entire corporation.

Hardware VPN Client Option

The hardware VPN client option is identical to the remote-site firewall option except that the hardware VPN client does not have a resident stateful firewall. This setup requires use of a personal firewall on the individual hosts, particularly when split tunneling is enabled. Without the personal firewall, the security of the individual hosts behind the VPN device is dependant upon the attacker being unable to circumvent *Network Address Translation* (NAT). This is because when split tunneling is enabled, connections to the Internet pass through a simple many-to-one NAT translation and do not undergo any filtering at Layer 4 and above. With split tunneling disabled, all access to the Internet must be through the corporate headquarters. This setup partially mitigates the requirement for personal firewalls on the end systems.

Using a hardware VPN client offers two primary advantages. First, as with the VPN software client, access and authorization to the corporate network and the Internet are controlled centrally from the headquarters location. Configuration and security management of the VPN hardware client device itself is done via an SSL connection from the central site. This setup ensures that the remote-site user(s) is not required to perform any configuration changes on the hardware VPN client. The second advantage of the hardware VPN client option is that individual PCs on the remote-site network do not need VPN client software to access corporate resources. However, individual users at the remote site who access the corporate network are not authenticated with this option. Instead, the VPN hardware client and VPN headend concentrator authenticate each other.

Remote-Site Router Option

The remote-site router option is nearly identical to the remote-site firewall option with a few exceptions. When deployed behind a stand-alone broadband access device, the only difference is the router can support advanced applications such as QoS, routing, and more encapulation options. Additionaly, if the broadband capability is integrated into the router, a stand-alone broadband access device is not needed. This option requires that your ISP allow you to manage the broadband router itself, an uncommon scenario.

Migration Strategies

SAFE is a guide for implementing security on a network. It is not meant to serve as a security policy for networks, nor is it meant to serve as the all-encompassing design for providing full security for all existing networks. Rather, SAFE is a template that enables network designers to consider how they design and implement their enterprise network in order to meet their security requirements.

Establishing a security policy should be the first activity in migrating the network to a secure infrastructure. Basic recommendations for a security policy can be found at the end of the document in Appendix B, "Network Security Primer." After the policy is established, the network designer should consider the security axioms described in the first section of this document and see how they provide more detail to map the policy onto the existing network infrastructure.

The architecture has enough flexibility to enable SAFE to be adapted to most networks. SAFE allows the designer to address the security requirements of each network function almost independently of each other. Each module is generally self-contained and assumes that any interconnected module is at only a basic security level. This setup allows network designers to use a phased approach to securing the enterprise network. They can address securing the most critical network functions as determined by the policy without redesigning the entire network.

This is the second white paper that details the specifics of the SAFE architecture. When this paper is combined with SAFE Enterprise, the documents address the security requirements and implementations for networks of varying sizes. The authors know that many other areas need further research, exploration, and improvement. Some of these areas include, but are not limited to, the following:

- In-depth security management analysis and implementation
- In-depth identity, directory services, AAA technologies, and *certificate authority* (CA) analysis and implementation
- In-depth wireless design, management, and implementation considerations

Appendix A: Validation Lab

A reference implementation exists to validate the functionality described in this document. This appendix details the configurations of the specific devices within each module as well as the overall guidelines for general device configuration. The following are configuration snapshots from the live devices in the lab. The authors do not recommend applying these configurations directly to a production network.

Overall Guidelines

The sample commands presented in this section correspond in part to the SAFE axioms presented earlier in this document.

Routers

The following are sample commands that enable most of the basic configuration options present on most routers in the lab.

```
! Turn off unnecessary services
!
no ip domain-lookup
no cdp run
no ip http server
no ip source-route
no service finger
no ip bootp server
no service udp-small-servers
no service tcp-small-servers


! Turn on logging and read-only snmp
!
service timestamp log datetime localtime
logging 10.3.8.254
logging 10.3.8.253
snmp-server community Txo~QbW3XM ro 98


! Generate RSA keys and enable SSH access. This requires the router support encryption.
! You will be prompted for the size of the RSA key. 1024 bits was chosen for the SAFE
! lab implementation
!
crypto key generate rsa
ip ssh timeout 120
ip ssh authentication-retries 5


! Set passwords and access restrictions
!
service password-encryption
enable secret %Z<)|z9~zq
no enable password
!
!
access-list 99 permit host 10.3.8.254
access-list 99 deny any log
```

```
!
access-list 98 permit host 10.3.8.253
access-list 98 permit host 10.3.8.254
access-list 98 deny any log
!
line vty 0 4
 access-class 99 in
 login authentication default
 password 0 X)[^j+#T98
 exec-timeout 2 0
 transport input ssh
 transport output none
line con 0
 login authentication no_tacacs
 password 0 X)[^j+#T98
 exec-timeout 2 0
 transport input none
line aux 0
 transport input none
 password 0 X)[^j+#T98
 no exec
!
banner motd #

        This is a private system operated for and by Cisco VSEC BU.
    Authorization from Cisco VSEC management is required to use this system.
            Use by unauthorized persons is prohibited.

#


! Turn on NTP with authentication and access control
!
clock timezone PST -8
clock summer-time PST recurring
!
ntp authenticate
ntp authentication-key 1 md5 -UN&/6[oh6
ntp trusted-key 1
ntp access-group peer 96
ntp server 10.3.4.4 key 1
!
access-list 96 permit host 10.3.4.4
access-list 96 deny any log


! Turn on AAA
!
aaa new-model
```

```
aaa authentication login default tacacs+
aaa authentication login no_tacacs line
aaa authorization exec tacacs+
aaa authorization network tacacs+
aaa accounting network start-stop tacacs+
aaa accounting exec start-stop tacacs+
!
tacacs-server host 10.3.8.253 single-connection
tacacs-server key SJj)j~t]6-
```

The following sample commands define the *Open Shortest Path First* (OSPF) authentication parameters for routers within the network. Notice that *message Digest 5* (MD5) authentication is being used in the configurations.

```
interface FastEthernet1/0
 ip address 10.3.3.3 255.255.255.0
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 8R%xi!0eUUxF
!
router ospf 1
 log-adjacency-changes
 area 0 authentication message-digest
 network 10.3.3.0 0.0.0.255 area 1
 network 10.3.4.0 0.0.0.255 area 1
```

Switches

Following are the sample commands to enable most of the basic configuration options present on most Catalyst® OS Switches in the lab. Cisco IOS® Switches use a configuration nearly identical to the router configuration.

```
! Turn on NTP
!
set timezone PST -8
set summertime PST
set summertime recurring
set ntp authentication enable
set ntp key 1 trusted md5 -UN&/6[oh6
set ntp server 10.3.4.4 key 1
set ntp client enable


! Turn off un-needed services
!
set cdp disable
set ip http server disable


! Turn on logging and snmp
!
set logging server 10.3.8.253
set logging server 10.3.8.254
set logging timestamp enable
set snmp community read-only Txo~QbW3XM
```

```
set ip permit enable snmp
set ip permit 10.3.8.254 snmp


! Turn on AAA
!
set tacacs server 10.3.8.253 primary
set tacacs key SJj)j~t]6-
set authentication login tacacs enable telnet
set authentication login local disable telnet
set authorization exec enable tacacs+ deny telnet
set accounting exec enable start-stop tacacs+
set accounting connect enable start-stop tacacs+


! Set passwords and access restrictions
!
set banner motd <c>

        This is a private system operated for and by Cisco VSEC BU.

    Authorization from Cisco VSEC management is required to use this system.

            Use by unauthorized persons is prohibited.

<c>


! Console password is set by 'set password'
! Enter old password followed by new password
! Console password = X)[^j+#T98
!
! Enable password is set by 'set enable'
! Enter old password followed by new password
! Enable password = %Z<)|z9~zq
!
! The following password configuration only works the first time
!
set password

X)[^j+#T98
X)[^j+#T98
set enable
cisco
%Z<)|z9~zq
%Z<)|z9~zq
!
! The above password configuration only works the first time
!
```

```
set logout 2
set ip permit enable telnet
set ip permit 10.3.8.253 255.255.255.255 telnet
set ip permit 10.3.8.254 255.255.255.255 telnet
```

Hosts

As discussed in the Axioms sections, the host operating systems and applications ran HIDS and had the latest patches and fixes applied. The HIDS application used in the lab is the Entercept application from Entercept Security Technologies. More information is available at http://www.entercept.com.

Small Network Configurations

The following are configuration snapshots from the SAFE small network.

Corporate Internet Module

Figure 15 Detailed Model of Small Network Corporate Internet Module

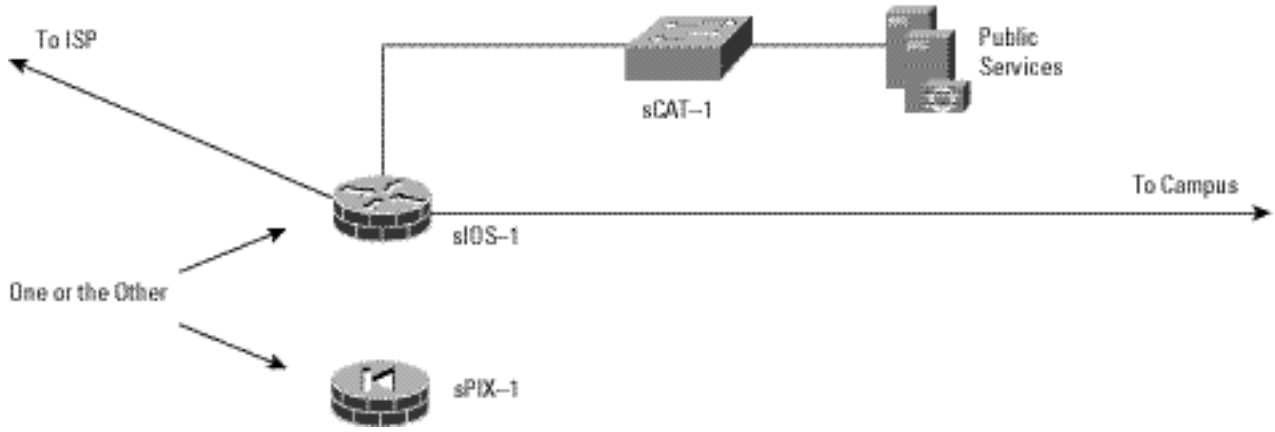

Products Used

- Cisco Catalyst Layer 2 Switch (sCAT-1)
- Cisco IOS Router with *Triple Data Encryption Standard* (3DES) encryption support (sIOS-1)
- Cisco Secure PIX Firewall (sPIX-1)
- Entercept HIDS

sIOS-1

The following configuration snapshot details the access lists on the small network edge router, controlling traffic inbound and outbound to the small network. Note: The small router configurations do not show the remote access VPN configurations; this functionality will be available in Cisco IOS Software soon.

```
! Basic IOS IDS configuration using syslog for reporting
!
ip audit attack action alarm drop reset
ip audit notify log
ip audit name alarm1 info action alarm
ip audit name alarm1 attack action alarm drop
!
! IPSec crypto configuration to remote branches of the small network
!
```

```
crypto isakmp policy 1
 encr 3des
 authentication pre-share
 group 2
crypto isakmp key 7Q!r$y$+xE address 172.16.128.2
crypto isakmp key 7Q!r$y$+xE address 172.16.128.5
!
!
crypto ipsec transform-set remote1 esp-3des esp-sha-hmac
!
crypto map ent1 30 ipsec-isakmp
 set peer 172.16.128.2
 set transform-set remote1
 match address 107
crypto map ent1 40 ipsec-isakmp
 set peer 172.16.128.5
 set transform-set remote1
 match address 108
!
! The access-lists below allow both user traffic as well as management
! traffic to be encrypted
!
access-list 107 permit ip 10.4.0.0 0.0.255.255 10.5.0.0 0.0.255.255
access-list 107 permit ip host 10.4.1.253 host 172.16.128.2
access-list 108 permit ip 10.4.0.0 0.0.255.255 10.6.0.0 0.0.255.255
access-list 108 permit ip host 10.4.1.253 host 172.16.128.5
!
! Interface settings for the inside interface of the router. NAT, IOS IDS, and
! IOS firewall are enabled.
!
interface FastEthernet0/0
 description Inside Interface
 ip address 10.4.1.1 255.255.255.0
 ip access-group 109 in
 ip nat inside
 ip inspect smbranch_fw in
 ip audit alarm1 in
!
! Allow ICMP from within the small network out to the Internet
!
access-list 109 permit icmp any any echo
!
! Allow the internal DNS server to communicate with the public DNS server
!
access-list 109 permit udp host 10.4.1.201 host 10.4.2.50 eq domain
!
! Allow internal users access to the public services server for HTTP,
! SSL and FTP traffic.
!
```

```
access-list 109 permit tcp 10.4.0.0 0.0.255.255 host 10.4.2.50 eq www
access-list 109 permit tcp 10.4.0.0 0.0.255.255 host 10.4.2.50 eq 443
access-list 109 permit tcp 10.4.0.0 0.0.255.255 host 10.4.2.50 eq ftp
!
! Allow the internal mail server to communicate with the public mail server
!
access-list 109 permit tcp host 10.4.1.201 host 10.4.2.50 eq smtp
!
! Allow Telnet access from the management host to the sCAT-1 switch
!
access-list 109 permit tcp host 10.4.1.253 host 10.4.2.4 eq telnet
!
! Deny all other access to the public services segment
!
access-list 109 deny  ip any 10.4.2.0 0.0.0.255
!
! Allow the sIOS-1 router and sCAT-2 switch to synchronize time
!
access-list 109 permit udp host 10.4.1.4 host 10.4.1.1 eq ntp
!
! Allow SSH access from the management host to the sIOS-1 router
!
access-list 109 permit tcp host 10.4.1.253 host 10.4.1.1 eq 22
!
! Allow established connections to the management host back to the sIOS-1 router
!
access-list 109 permit tcp host 10.4.1.253 eq tacacs host 10.4.1.1 established
!
! Necessary for TFTP access from the management host to the sIOS-1 router
!
access-list 109 permit udp host 10.4.1.253 gt 1023 host 10.4.1.1 gt 1023
!
! Block all other access to the inside interface on the sIOS-1 router from the
! internal network
!
access-list 109 deny  ip 10.4.0.0 0.0.255.255 host 10.4.1.1
!
! Block all other access to the outside interface on the sIOS-1 router from the
! internal network
!
access-list 109 deny  ip 10.4.0.0 0.0.255.255 host 172.16.132.2
!
! Allow all other internal devices access to the Internet
!
access-list 109 permit ip 10.4.0.0 0.0.255.255 any
!
! Block and log all other traffic
!
access-list 109 deny  ip any any log
```

```
!
! Interface settings for the public services interface of the router. NAT, IOS IDS, and IOS firewall are enabled.
!
interface FastEthernet0/1
 description DMZ Interface
 ip address 10.4.2.1 255.255.255.0
 ip access-group 105 in
 no ip redirects
 ip nat inside
 ip inspect smbranch_fw in
 ip audit alarm1 in
!
! Allow sCAT-1 switch to synchronize time with sIOS-1 router
!
access-list 105 permit udp host 10.4.2.4 host 10.4.2.1 eq ntp
!
! Allow TACACS+, TFTP, and syslog from the SCAT-1 switch to the
! management host
!
access-list 105 permit tcp host 10.4.2.4 host 10.4.1.253 eq tacacs
access-list 105 permit udp host 10.4.2.4 host 10.4.1.253 eq tftp
access-list 105 permit udp host 10.4.2.4 host 10.4.1.253 eq syslog
!
! Allow HIDS traffic from the public services server to
! the management host
!
access-list 105 permit tcp host 10.4.2.50 host 10.4.1.253 eq 5000
!
! Allow the public email server to send mail to the internal mail server
!
access-list 105 permit tcp host 10.4.2.50 host 10.4.1.201 eq smtp
!
! Deny all other connections originating from the public services segment
! to the internal network
!
access-list 105 deny  ip any 10.4.0.0 0.0.255.255
!
! Permit all mail and DNS traffic originating from the public services
! server
!
access-list 105 permit tcp host 10.4.2.50 any eq smtp
access-list 105 permit udp host 10.4.2.50 any eq domain
!
! Deny all other traffic and log
!
access-list 105 deny  ip any any log
!
! Interface settings for the public interface of the router. NAT, IOS IDS,
! IOS firewall, and IPSec are enabled.
```

```
!
interface Serial1/0
 description Outside Interface
 ip address 172.16.132.2 255.255.255.0
 ip access-group 103 in
 no ip redirects
 ip nat outside
 ip inspect smbranch_fw in
 ip audit alarm1 in
 crypto map ent1
!
! Allows traffic from remote sites to the small network. Only needed
! when small network functions as a headend for remote sites.
!
access-list 103 permit ip 10.5.0.0 0.0.255.255 10.4.0.0 0.0.255.255
access-list 103 permit ip 10.6.0.0 0.0.255.255 10.4.0.0 0.0.255.255
!
! RFC 1918 filtering. Note network 172.16.x.x was not included in the
! filter here since it is used to simulate the ISP in the lab.
!
access-list 103 deny   ip 10.0.0.0 0.255.255.255 any
access-list 103 deny   ip 192.168.0.0 0.0.255.255 any
!
! Allow any echo replies which originate from the 172.16.132.0
! network (NAT translated internal addresses) back.
!
access-list 103 permit icmp any 172.16.132.0 0.0.0.255 echo-reply
!
! Allow path MTU discovery (PMTUD) traffic.
!
access-list 103 permit icmp any 172.16.132.0 0.0.0.255 unreachable
!
! Allows IPSec traffic from remote sites to terminate on the sIOS-1
! router. Only needed when small network functions as a headend for
! remote sites.
!
access-list 103 permit esp host 172.16.128.2 host 172.16.132.2
access-list 103 permit udp host 172.16.128.2 host 172.16.132.2 eq isakmp
access-list 103 permit esp host 172.16.128.5 host 172.16.132.2
access-list 103 permit udp host 172.16.128.5 host 172.16.132.2 eq isakmp
!
!
! Allows management of remote sites.  Only needed when small network
! functions as a headend for remote sites.
!
access-list 103 permit tcp host 172.16.128.2 host 10.4.1.253 eq tacacs
access-list 103 permit udp host 172.16.128.2 host 10.4.1.253 eq syslog
access-list 103 permit udp host 172.16.128.2 host 10.4.1.253 eq tftp
access-list 103 permit tcp host 172.16.128.5 host 10.4.1.253 eq tacacs
access-list 103 permit udp host 172.16.128.5 host 10.4.1.253 eq syslog
```

```
access-list 103 permit udp host 172.16.128.5 host 10.4.1.253 eq tftp
!
! Allow access to the public services server (via the NAT
! address of the server) for DNS, FTP, HTTP, SSL, and mail
! traffic
!
access-list 103 permit udp any host 172.16.132.50 eq domain
access-list 103 permit tcp any host 172.16.132.50 eq ftp
access-list 103 permit tcp any host 172.16.132.50 eq www
access-list 103 permit tcp any host 172.16.132.50 eq 443
access-list 103 permit tcp any host 172.16.132.50 eq smtp
!
! Deny all other traffic and log
!
access-list 103 deny   ip any any log

!
! The following NAT configuration creates a pool of public addresses which
! are used by internal devices when they access the Internet
!
ip nat pool small_pool 172.16.132.101 172.16.132.150 netmask 255.255.255.0
ip nat inside source route-map nat_internet pool small_pool
!
! Static translation of the public services server to a registered
! address accessible from the Internet
!
ip nat inside source static 10.4.2.50 172.16.132.50
!
route-map nat_internet permit 10
 match ip address 104
!
! Do not use NAT for internal devices communicating with other network
! 10.0.0.0 devices, or for management traffic.  Use NAT for all internal
! devices communicating with the Internet.
!
access-list 104 deny   ip 10.4.0.0 0.0.255.255 10.0.0.0 0.255.255.255
access-list 104 deny   ip host 10.4.1.253 host 172.16.128.2
access-list 104 deny   ip host 10.4.1.253 host 172.16.128.5
access-list 104 permit ip 10.4.1.0 0.0.0.255 any
!
```

Branch versus Headend Configuration Changes

The following configuration snapshot details the changes necessary to make the small network a branch of a larger network using a redundant IPSec-over-GRE VPN connection.

```
!
! Crypto Policy Settings
!
crypto isakmp policy 1
 encr 3des
 authentication pre-share
```

```
 group 2
crypto isakmp key 7Q!r$y$+xE address 172.16.226.28
crypto isakmp key 7Q!r$y$+xE address 172.16.226.27
!
!
crypto ipsec transform-set 3dessha esp-3des esp-sha-hmac
 mode transport
!
crypto map ent1 10 ipsec-isakmp
 set peer 172.16.226.28
 set transform-set 3dessha
 match address 101
crypto map ent1 20 ipsec-isakmp
 set peer 172.16.226.27
 set transform-set 3dessha
 match address 102
!
access-list 101 permit gre host 172.16.132.2 host 172.16.226.28
access-list 102 permit gre host 172.16.132.2 host 172.16.226.27
!
! GRE Tunnel Settings
!
interface Tunnel0
 bandwidth 8
 ip address 10.1.249.2 255.255.255.0
 tunnel source 172.16.132.2
 tunnel destination 172.16.226.27
 crypto map ent1
!
interface Tunnel1
 ip address 10.1.248.2 255.255.255.0
 tunnel source 172.16.132.2
 tunnel destination 172.16.226.28
 crypto map ent1
!
! Crypto Map Application to Physical Interface
!
interface Serial1/0
 ip address 172.16.132.2 255.255.255.0
 ip access-group 103 in
 crypto map ent1
!
! Access-list 103 would need to be modified to both the IPSec connections
! from the corporate headquarters, as well as the GRE traffic.
!
access-list 103 permit gre host 172.16.226.28 host 172.16.132.2
access-list 103 permit gre host 172.16.226.27 host 172.16.132.2
access-list 103 permit esp host 172.16.226.27 host 172.16.132.2
access-list 103 permit udp host 172.16.226.27 host 172.16.132.2 eq isakmp
```

```
access-list 103 permit esp host 172.16.226.28 host 172.16.132.2
access-list 103 permit udp host 172.16.226.28 host 172.16.132.2 eq isakmp
!
! Note that all configurations pertaining to the remote sites is removed
!
access-list 103 deny  ip 10.0.0.0 0.255.255.255 any
access-list 103 deny  ip 192.168.0.0 0.0.255.255 any
access-list 103 permit udp any host 172.16.132.50 eq domain
access-list 103 permit tcp any host 172.16.132.50 eq ftp
access-list 103 permit tcp any host 172.16.132.50 eq www
access-list 103 permit tcp any host 172.16.132.50 eq 443
access-list 103 permit tcp any host 172.16.132.50 eq smtp
access-list 103 permit icmp any 172.16.132.0 0.0.0.255 echo-reply
access-list 103 permit icmp any 172.16.132.0 0.0.0.255 unreachable
access-list 103 deny  ip any any log
```

Minor modifications to the other access lists are required as well, but they are not shown here.

sPIX-1

The following configuration snapshot details the access lists and cryptographic configuration when using the PIX Firewall as the headend device in the small network. The PIX Firewall as configured is capable of communicating with remote sites and terminating dial-in IPSec VPN connections.

```
!
! Interface settings for the public interface of the firewall
!
ip address outside 172.16.144.3 255.255.255.0
access-group 103 in interface outside
!
! Allow encrypted traffic from remote sites and remote access users.
!
access-list 103 permit ip 10.5.0.0 255.255.0.0 10.4.0.0 255.255.0.0
access-list 103 permit ip 10.6.0.0 255.255.0.0 10.4.0.0 255.255.0.0
access-list 103 permit ip 10.4.3.0 255.255.255.0 10.4.0.0 255.255.0.0

! RFC 1918 filtering. Note network 172.16.x.x was not included in the
! filter here since it is used to simulate the ISP in the lab.
!
access-list 103 deny ip 10.0.0.0 255.0.0.0 any
access-list 103 deny ip 192.168.0.0 255.255.0.0 any
!
! Allow access to the public services server (via the NAT
! address of the server) for DNS, FTP, HTTP, SSL, and mail
! traffic
!
access-list 103 permit udp any host 172.16.144.50 eq domain
access-list 103 permit tcp any host 172.16.144.50 eq ftp
access-list 103 permit tcp any host 172.16.144.50 eq www
access-list 103 permit tcp any host 172.16.144.50 eq 443
access-list 103 permit tcp any host 172.16.144.50 eq smtp
!
```

```
! Allow echo reply generated from the internal network (via NAT translated
! addresses) back into the firewall
!
access-list 103 permit icmp any 172.16.144.0 255.255.255.0 echo-reply
!
! Allow path MTU discovery (PMTUD) traffic through the firewall.
!
access-list 103 permit icmp any 172.16.144.0 255.255.255.0 unreachable
!
! Allow syslog, TFTP, and TACACS+ management traffic in from remote sites.
!
access-list 103 permit udp host 172.16.128.2 host 172.16.144.51 eq syslog
access-list 103 permit udp host 172.16.128.2 host 172.16.144.51 eq tftp
access-list 103 permit tcp host 172.16.128.2 host 172.16.144.51 eq tacacs
access-list 103 permit udp host 172.16.128.5 host 172.16.144.51 eq syslog
access-list 103 permit udp host 172.16.128.5 host 172.16.144.51 eq tftp
access-list 103 permit tcp host 172.16.128.5 host 172.16.144.51 eq tacacs

!
! Interface settings for the private interface of the firewall
!
ip address inside 10.4.1.1 255.255.255.0
access-group 109 in interface inside
!
! Allow echo from internal devices
!
access-list 109 permit icmp any any echo
!
! Allow the internal DNS and mail server to communicate with the
! public DNS and mail server
!
access-list 109 permit udp host 10.4.1.201 host 10.4.2.50 eq domain
access-list 109 permit tcp host 10.4.1.201 host 10.4.2.50 eq smtp
!
! Allow internal devices to access the public services server for web, FTP
! and SSL access
!
access-list 109 permit tcp 10.4.0.0 255.255.0.0 host 10.4.2.50 eq www
access-list 109 permit tcp 10.4.0.0 255.255.0.0 host 10.4.2.50 eq ftp
access-list 109 permit tcp 10.4.0.0 255.255.0.0 host 10.4.2.50 eq 443
!
! Allow Telnet access from the management host to the mCAT-1 switch
!
access-list 109 permit tcp host 10.4.1.253 host 10.4.2.4 eq telnet
!
! Block all other access to the public services segment
!
access-list 109 deny ip any 10.4.2.0 255.255.255.0
!
! Permit internal devices access to the Internet
```

```
!
access-list 109 permit ip 10.4.0.0 255.255.0.0 any
!
!
! Interface settings for the public services (DMZ) interface of the firewall
!
ip address pss 10.4.2.1 255.255.255.0
access-group 105 in interface pss
!
! Allow echo-replies from internal network back through firewall
!
access-list 105 permit icmp 10.4.2.0 255.255.255.0 10.4.1.0 255.255.255.0 echo-reply
!
! Allow TACACS+, TFTP, and syslog from the sCAT-1 switch to the management server
!
access-list 105 permit tcp host 10.4.2.4 host 10.4.1.253 eq tacacs
access-list 105 permit udp host 10.4.2.4 host 10.4.1.253 eq tftp
access-list 105 permit udp host 10.4.2.4 host 10.4.1.253 eq syslog
!
! Allow HIDS traffic from the public services server to the
! management server
!
access-list 105 permit tcp host 10.4.2.50 host 10.4.1.253 eq 5000
!
! Allow the public mail server to communicate with the internal mail server
!
access-list 105 permit tcp host 10.4.2.50 host 10.4.1.201 eq smtp
!
! Block all other access from this segment to the internal network
!
access-list 105 deny ip any 10.4.0.0 255.255.0.0
!
! Allow access from the public services server to the Internet for
! mail and DNS
!
access-list 105 permit tcp host 10.4.2.50 any eq smtp
access-list 105 permit udp host 10.4.2.50 any eq domain
!
! IDS Settings
!
ip audit name full info action alarm
ip audit name fullb attack action alarm drop
ip audit interface outside full
ip audit interface outside fullb
ip audit interface inside full
ip audit interface inside fullb
ip audit interface pss full
ip audit interface pss fullb
!
! The following NAT configuration creates a pool of public addresses which
```

```
! are used by internal devices when they access the Internet
!
global (outside) 1 172.16.144.201-172.16.144.220
!
nat (inside) 0 access-list nonat
nat (inside) 1 0.0.0.0 0.0.0.0 0 0
nat (pss) 0 access-list nonat
!
! Static translation of the public services server to a registered address
! accessible from the Internet
!
static (pss,outside) 172.16.144.50 10.4.2.50 netmask 255.255.255.255 0 0
!
static (inside,pss) 10.4.1.253 10.4.1.253 netmask 255.255.255.255 0 0
static (inside,pss) 10.4.1.201 10.4.1.201 netmask 255.255.255.255 0 0
static (inside,outside) 172.16.144.51 10.4.1.253 netmask 255.255.255.255 0 0
!
! Access-list nonat determines what addresses use address translation
!
access-list nonat permit ip 10.4.0.0 255.255.0.0 10.5.0.0 255.255.0.0
access-list nonat permit ip 10.4.0.0 255.255.0.0 10.6.0.0 255.255.0.0
access-list nonat permit ip 10.4.1.0 255.255.255.0 10.4.3.0 255.255.255.0
access-list nonat permit ip 10.4.2.0 255.255.255.0 10.4.3.0 255.255.255.0
access-list nonat permit ip 10.4.1.0 255.255.255.0 10.4.2.0 255.255.255.0
!
! The following crypto settings are used when the firewall terminates VPN connections from the remote sites
!
no sysopt route dnat
crypto ipsec transform-set 3dessha esp-3des esp-sha-hmac
crypto ipsec transform-set remote1 esp-3des esp-sha-hmac
crypto dynamic-map vpnuser 20 set transform-set remote1
crypto map ent1 30 ipsec-isakmp
crypto map ent1 30 match address 107
crypto map ent1 30 set peer 172.16.128.2
crypto map ent1 30 set transform-set remote1
crypto map ent1 40 ipsec-isakmp
crypto map ent1 40 match address 108
crypto map ent1 40 set peer 172.16.128.5
crypto map ent1 40 set transform-set remote1
crypto map ent1 50 ipsec-isakmp dynamic vpnuser
crypto map ent1 client configuration address initiate
crypto map ent1 client authentication vpnauth
crypto map ent1 interface outside
!
access-list 107 permit ip 10.4.0.0 255.255.0.0 10.5.0.0 255.255.0.0
access-list 107 permit ip host 172.16.144.51 host 172.16.128.2
access-list 108 permit ip 10.4.0.0 255.255.0.0 10.6.0.0 255.255.0.0
access-list 108 permit ip host 172.16.144.51 host 172.16.128.5
!
```

```
isakmp enable outside
isakmp key 7Q!r$y$+xE address 172.16.128.5 netmask 255.255.255.255
isakmp key 7Q!r$y$+xE address 172.16.128.2 netmask 255.255.255.255
isakmp key 7Q!r$y$+xE address 172.16.226.28 netmask 255.255.255.255
isakmp key 7Q!r$y$+xE address 172.16.226.27 netmask 255.255.255.255
isakmp policy 1 authentication pre-share
isakmp policy 1 encryption 3des
isakmp policy 1 hash sha
isakmp policy 1 group 2
isakmp policy 1 lifetime 86400
!
! The following configuration block configures the PIX to terminate remote access
! VPN users
!
vpngroup VPN1 address-pool vpnpool
vpngroup VPN1 dns-server 10.4.1.201
vpngroup VPN1 default-domain safe-small.com
vpngroup VPN1 idle-time 1800
vpngroup VPN1 password Y0eS)3/i6y
ip local pool vpnpool 10.4.3.1-10.4.3.254
```

Campus Module

Figure 16 Detailed Model of Small Network Campus Module

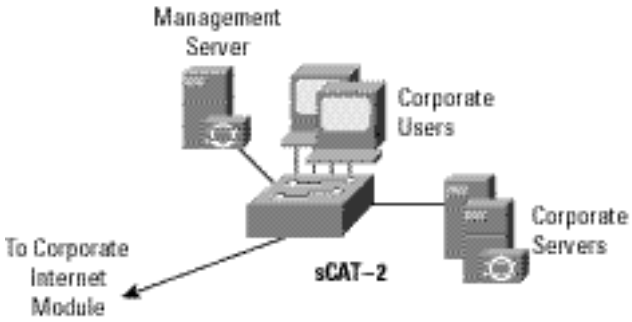

Products Used

- Cisco Catalyst Layer 2 Switch (sCAT-2)
- Entercept HIDS
- Cisco Secure Access Control Server
- Cisco Secure Policy Manager Lite
- OpenSystems Private I syslog analysis tool
- F-Secure *Secure Shell Protocol* (SSH) Client

Medium Network Configurations

The following are configuration snapshots from the SAFE medium network. Unless otherwise specified, they are configurations for the medium network operating as a headend.

Corporate Internet Module

Figure 17 Detailed Model of Medium Corporate Internet Module

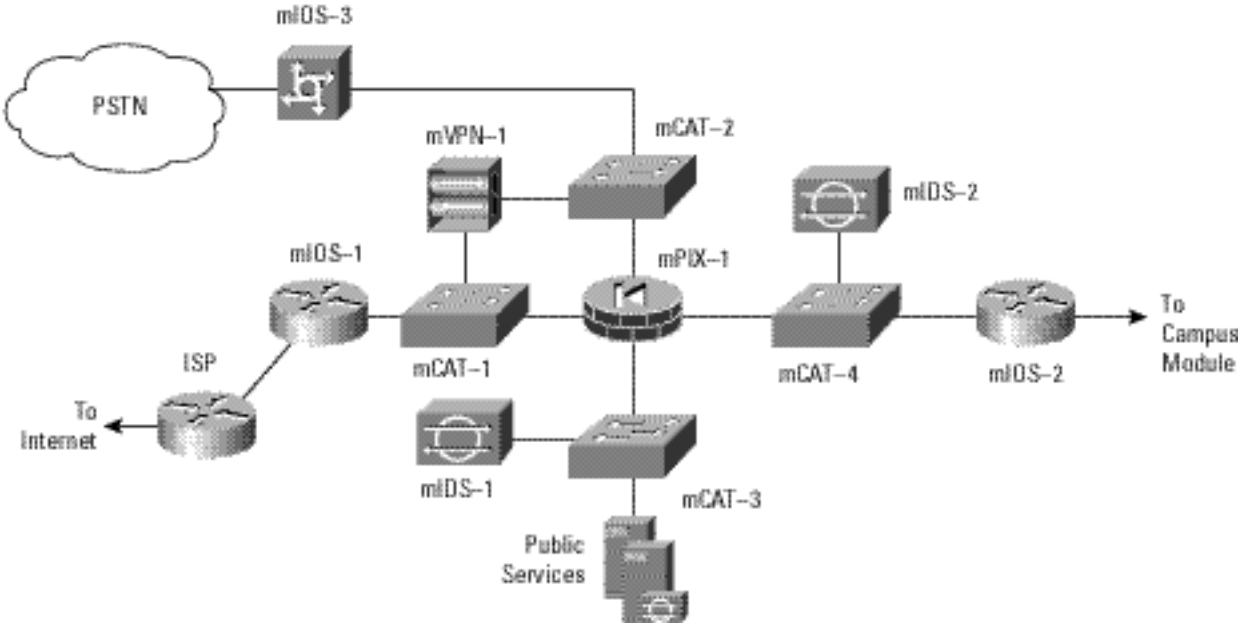

Products Used

- Cisco Catalyst Layer 2 Switches (mCAT-1 through mCAT-4)
- Cisco IOS Routers with 3DES encryption support (mIOS-1 and mIOS-2)
- Cisco IOS Dial-Access Router (mIOS-3)
- Cisco VPN 3000 Series Concentrator (mVPN-1)
- Cisco Secure PIX Firewall (mPIX-1)
- Cisco Secure IDS Sensors (mIDS-1 and mIDS-2)
- Entercept HIDS
- Baltimore MIMESweeper Email Filtering

mIOS-1

The following configuration snapshot details the access lists on the medium network edge router, mIOS-1, controlling traffic inbound from the *Internet service provider* (ISP) to the medium network.

```
interface FastEthernet0/0
 ip address 172.16.240.2 255.255.255.0
 ip access-group 112 in
 no ip redirects

 no cdp enable
!
interface Serial1/0
 ip address 172.16.131.2 255.255.255.0
 ip access-group 150 in
 dsu bandwidth 44210
 framing c-bit

 no cdp enable
!
! RFC 1918 filtering. Note that network 172.16.0.0 was used for the simulated SAFE ISP
! network, and therefore has not been included in the RFC 1918 filtering here.
!
access-list 150 deny  ip 10.0.0.0 0.255.255.255 any
access-list 150 deny  ip 192.168.0.0 0.0.255.255 any
!
!
! Prevent any outside devices from spoofing an address that appears to originate from
! within the medium network.
!
access-list 150 deny  ip 172.16.240.0 0.0.0.255 any
!
! Allow relevant IKE and ESP traffic to reach the VPN devices.
!
access-list 150 permit esp any host 172.16.240.3
access-list 150 permit udp any host 172.16.240.3 eq isakmp
access-list 150 permit esp host 172.16.128.2 host 172.16.240.1
access-list 150 permit udp host 172.16.128.2 host 172.16.240.1 eq isakmp
access-list 150 permit esp host 172.16.128.5 host 172.16.240.1
access-list 150 permit udp host 172.16.128.5 host 172.16.240.1 eq isakmp
!
!
! Restrict any other conversations to mIOS-1, mVPN-1, mPIX-1 and mCAT-1.
!
access-list 150 deny  ip any host 172.16.240.3
access-list 150 deny  ip any host 172.16.240.4
access-list 150 deny  ip any host 172.16.240.2
access-list 150 deny  ip any host 172.16.240.1
!
!
! Allow all other connections to the 172.16.240 subnet, since internal user devices
```

```
! translate to 172.16.240.0 addresses at the firewall as they access the Internet.
!
access-list 150 permit ip any 172.16.240.0 0.0.0.255
!
!
! Block and log all other attempted access.
!
access-list 150 deny  ip any any log
!
!
```

The following configuration snapshot details the access lists that control traffic inbound from the medium network to the ISP on the edge router.

```
! Allow TCP sessions that originated from the router to the management hosts
! (TACACS+, etc.). Management hosts 172.16.240.151 and 172.16.240.152 are the
! NAT translated addresses at the firewall.
!
access-list 112 permit tcp host 172.16.240.151 host 172.16.240.2 established
access-list 112 permit tcp host 172.16.240.152 host 172.16.240.2 established
!
!
! Allow SSH connections originated from the management hosts to the router.
!
access-list 112 permit tcp host 172.16.240.151 host 172.16.240.2 eq 22
access-list 112 permit tcp host 172.16.240.152 host 172.16.240.2 eq 22
!
!
! Necessary for allowing TFTP back from the management host to the router.
!
access-list 112 permit udp host 172.16.240.151 host 172.16.240.2 gt 1024
!
!
! Allow other devices on the 172.16.240 subnet to synchronize clocks to this device.
!
access-list 112 permit udp 172.16.240.0 0.0.0.255 host 172.16.240.2 eq ntp
!
!
! Allow internal devices to ping the Internet.
!
access-list 112 permit icmp 172.16.240.0 0.0.0.255 any
!
!
! Block all other attempts to access this router and log.
!
access-list 112 deny  ip any host 172.16.240.2 log
!
!
! Permit all access to the Internet from hosts with 172.16.240.0 addresses.
!
```

```
access-list 112 permit ip 172.16.240.0 0.0.0.255 any
!
!
```

mPIX-1

The following configuration snapshot of the firewall details the security levels of the interfaces of the PIX Firewall, mPIX-1, as well as the addressing of each interface.

```
nameif ethernet0 outside security0
nameif ethernet1 inside security100
nameif ethernet2 pss security10
nameif ethernet3 vpn security15
!
ip address outside 172.16.240.1 255.255.255.0
ip address inside 10.3.4.1 255.255.255.0
ip address pss 10.3.6.1 255.255.255.0
ip address vpn 10.3.5.1 255.255.255.0
```

The following configuration snapshot of the firewall details the *Network Address Translation* (NAT) configuration of the PIX Firewall.

```
! In combination with the nonat access list, the below configuration does not allow
! NAT for sessions between internal devices (network 10.x.x.x to network 10.x.x.x), but
! allows NAT for sessions between internal or remote access devices and the Internet.
!
global (outside) 100 172.16.240.101-172.16.240.150 netmask 255.255.255.0
global (outside) 200 172.16.240.201-172.16.240.250 netmask 255.255.255.0
nat (inside) 0 access-list nonat
nat (inside) 100 10.0.0.0 255.0.0.0 0 0
nat (pss) 0 access-list nonat
nat (vpn) 200 10.3.7.0 255.255.255.0 0 0
static (inside,vpn) 10.3.0.0 10.3.0.0 netmask 255.255.0.0 0 0
static (inside,pss) 10.3.8.253 10.3.8.253 netmask 255.255.255.255 0 0
static (inside,pss) 10.3.8.254 10.3.8.254 netmask 255.255.255.255 0 0
!
!
! Translates the non-registered address of the public services server to a registered
! address, which can be accessed from the Internet.
!
static (pss,outside) 172.16.240.50 10.3.6.50 netmask 255.255.255.255 0 0
!
!
! Translates the non-registered addresses of the management hosts to registered
! addresses so that managed devices outside the firewall can initiate sessions
! to the management servers.
!
static (inside,outside) 172.16.240.151 10.3.8.254 netmask 255.255.255.255 0 0
static (inside,outside) 172.16.240.152 10.3.8.253 netmask 255.255.255.255 0 0
!
!
!
```

```
access-list nonat permit ip 10.0.0.0 255.0.0.0 10.0.0.0 255.0.0.0
access-list nonat deny ip 10.0.0.0 255.0.0.0 any
```

The following configuration snapshot details the access control in place on the PIX Firewall. The name of the access list denotes the location at which the inbound access control list (ACL) is placed.

Access-list "out" is placed inbound on the outside (public) interface of the firewall.

```
! Allow encrypted traffic from remote sites.
!
access-list out permit ip 10.5.0.0 255.255.0.0 10.3.0.0 255.255.0.0
access-list out permit ip 10.6.0.0 255.255.0.0 10.3.0.0 255.255.0.0
!
! RFC 1918 filtering. Note that network 172.16.0.0 was used for the simulated SAFE
! ISP network, and therefore has not been included in the RFC 1918 filtering here.
!
access-list out deny ip 10.0.0.0 255.0.0.0 any
access-list out deny ip 192.168.0.0 255.255.0.0 any
!
! Allow external hosts access to the public services server for HTTP, SSL, FTP
! SMTP, and DNS.
!
access-list out permit tcp any host 172.16.240.50 eq www
access-list out permit tcp any host 172.16.240.50 eq 443
access-list out permit tcp any host 172.16.240.50 eq ftp
access-list out permit tcp any host 172.16.240.50 eq smtp
access-list out permit udp any host 172.16.240.50 eq domain
!
! Allow echo reply generated from the internal network (via NAT translated
! addresses) back into the firewall
!
access-list out permit icmp any 172.16.240.0 255.255.255.0 echo-reply
!
! Allow path MTU discovery (PMTUD) traffic through the firewall.
!
access-list out  permit icmp any 172.16.240.0 255.255.255.0 unreachable
!
! Allow syslog, TFTP, and TACACS+ management traffic in from remote sites.
!
access-list out permit udp host 172.16.128.2 host 172.16.240.151 eq syslog
access-list out permit udp host 172.16.128.2 host 172.16.240.152 eq syslog
access-list out permit udp host 172.16.128.2 host 172.16.240.151 eq tftp
access-list out permit tcp host 172.16.128.2 host 172.16.240.152 eq tacacs
access-list out permit udp host 172.16.128.5 host 172.16.240.151 eq syslog
access-list out permit udp host 172.16.128.5 host 172.16.240.152 eq syslog
access-list out permit udp host 172.16.128.5 host 172.16.240.151 eq tftp
access-list out permit tcp host 172.16.128.5 host 172.16.240.152 eq tacacs
!
! Permit syslog, TFTP, and TACACS+ which originates from the mIOS-1 router
! and the mCAT-1 switch to the management hosts.
```

```
!
access-list out permit udp host 172.16.240.2 host 172.16.240.151 eq syslog
access-list out permit udp host 172.16.240.2 host 172.16.240.152 eq syslog
access-list out permit udp host 172.16.240.2 host 172.16.240.151 eq tftp
access-list out permit tcp host 172.16.240.2 host 172.16.240.152 eq tacacs
access-list out permit udp host 172.16.240.4 host 172.16.240.151 eq syslog
access-list out permit udp host 172.16.240.4 host 172.16.240.152 eq syslog
access-list out permit udp host 172.16.240.4 host 172.16.240.151 eq tftp
access-list out permit tcp host 172.16.240.4 host 172.16.240.152 eq tacacs
!
```

Access-list “in” is placed inbound on the inside (private) interface of the firewall.

```
! Allow echo from the inside network.
!
access-list in permit icmp any any echo
!
!
! Allow the internal DNS server to query the external DNS server for name
! translation.
!
access-list in permit udp host 10.3.2.50 host 10.3.6.50 eq domain
!
!
! Allow internal corporate users web, SSL, and FTP access to the external public
! services server.
!
access-list in permit tcp 10.0.0.0 255.0.0.0 host 10.3.6.50 eq www
access-list in permit tcp 10.0.0.0 255.0.0.0 host 10.3.6.50 eq 443
access-list in permit tcp 10.0.0.0 255.0.0.0 host 10.3.6.50 eq ftp
!
!
! Allow mail transfer from the external mail server to the internal mail
! server.
!
access-list in permit tcp host 10.3.2.50 host 10.3.6.50 eq smtp
!
!
! Allow Telnet access from the mgmt hosts to the mCAT-2 switch (does not support  ! SSH) on the public services segment.
!
access-list in permit tcp host 10.3.8.253 host 10.3.6.4 eq telnet
access-list in permit tcp host 10.3.8.254 host 10.3.6.4 eq telnet
!
!
! Deny all other access to the public services segment from the internal network.
!
access-list in deny ip any 10.3.6.0 255.255.255.0
!
!
! Permit all internal users access to the Internet.
```

```
!
access-list in permit ip 10.0.0.0 255.0.0.0 any
!
```

Access-list “pss” is placed inbound on the public services segment interface of the firewall.

```
! Allow syslog, TACAS+, and TFTP originated from the mCAT-2 switch to
! the management hosts.
!
access-list pss permit udp host 10.3.6.4 host 10.3.8.254 eq syslog
access-list pss permit udp host 10.3.6.4 host 10.3.8.253 eq syslog
access-list pss permit tcp host 10.3.6.4 host 10.3.8.253 eq tacacs
access-list pss permit udp host 10.3.6.4 host 10.3.8.254 eq tftp
!
!
! Allow the mCAT-2 switch to synchronize time with the internal router
! mIOS-2.
!
access-list pss permit udp host 10.3.6.4 host 10.3.4.4 eq ntp
!
!
! Allow HIDS traffic from the public services host to the
! network management host.
!
access-list pss permit tcp host 10.3.6.50 host 10.3.8.253 eq 5000
!

! Allow mail originated from the public services host (external mail server)
! to the corporate intranet services host (internal mail server).
!
access-list pss permit tcp host 10.3.6.50 host 10.3.2.50 eq smtp
!
!
! Deny all other traffic destined for addresses on the internal network.
!
access-list pss deny ip any 10.3.0.0 255.255.0.0
!
.
!
! Allow mail and DNS generated by the public services host to the Internet.
!
access-list pss permit tcp host 10.3.6.50 any eq smtp
access-list pss permit udp host 10.3.6.50 any eq domain
!
```

Access list “vpn” is placed inbound on the remote access VPN segment interface of the firewall. Remote VPN users are assigned addresses in the 10.3.7.0 subnet via an address pool defined in the access control server. Remote dial-in users are assigned addresses in the 10.3.8.0 subnet.

```
! Allow remote users web, SSL, and FTP access only to the public services server.
!
```

```
access-list vpn permit tcp 10.3.7.0 255.255.255.0 host 10.3.6.50 eq www
access-list vpn permit tcp 10.3.8.0 255.255.255.0 host 10.3.6.50 eq www
access-list vpn permit tcp 10.3.7.0 255.255.255.0 host 10.3.6.50 eq 443
access-list vpn permit tcp 10.3.8.0 255.255.255.0 host 10.3.6.50 eq 443
access-list vpn permit tcp 10.3.7.0 255.255.255.0 host 10.3.6.50 eq ftp
access-list vpn permit tcp 10.3.8.0 255.255.255.0 host 10.3.6.50 eq ftp
access-list vpn deny ip 10.3.7.0 255.255.255.0 10.3.6.0 255.255.255.0
access-list vpn deny ip 10.3.8.0 255.255.255.0 10.3.6.0 255.255.255.0
!
!
! Allow remote users access to the rest of the internal network and the Internet.
!
access-list vpn permit ip 10.3.7.0 255.255.255.0 any
access-list vpn permit ip 10.3.8.0 255.255.255.0 any
!
!
! Permit syslog, TFTP, and TACAS+ which originates from the VPN concentrator,
! mVPN-1, to the management hosts.
!
access-list vpn permit udp host 10.3.5.5 host 10.3.8.254 eq tftp
access-list vpn permit udp host 10.3.5.5 host 10.3.8.254 eq syslog
access-list vpn permit udp host 10.3.5.5 host 10.3.8.253 eq syslog
access-list vpn permit tcp host 10.3.5.5 host 10.3.8.253 eq tacacs
!
!
! Allow the VPN concentrator to synchronize time with the
! internal router, mIOS-2.
!
access-list vpn permit udp host 10.3.5.5 host 10.3.4.4 eq ntp
!
!
! Permit RADIUS authentication data which originates from the VPN concentrator to
! the management host.
!
access-list vpn permit udp host 10.3.5.5 host 10.3.8.253 eq 1645
!
!
! Permit syslog, TFTP, and TACAS+ which originate from the dial-in access
! router, mIOS-3, to the management hosts.
!
access-list vpn permit udp host 10.3.5.2 host 10.3.8.254 eq tftp
access-list vpn permit udp host 10.3.5.2 host 10.3.8.254 eq syslog
access-list vpn permit udp host 10.3.5.2 host 10.3.8.253 eq syslog
access-list vpn permit tcp host 10.3.5.2 host 10.3.8.253 eq tacacs
!
!
! Allow the dial-in access router, mIOS-3, to synchronize time with the
! internal router, mIOS-2.
!
access-list vpn permit udp host 10.3.5.2 host 10.3.4.4 eq ntp
```

```
!
!
! Permit syslog, TFTP, and TACAS+ which originates from the mcAT-3 switch
! to the management hosts.
!
access-list vpn permit udp host 10.3.5.4 host 10.3.8.254 eq tftp
access-list vpn permit udp host 10.3.5.4 host 10.3.8.254 eq syslog
access-list vpn permit udp host 10.3.5.4 host 10.3.8.253 eq syslog
access-list vpn permit tcp host 10.3.5.4 host 10.3.8.253 eq tacacs
!
!
! Allow the mCAT-3 switch to synchronize time with the
! internal router mIOS-2.
!
access-list vpn permit udp host 10.3.5.4 host 10.3.4.4 eq ntp
!
!
```

VPN Considerations

The following configurations are added to enable the site-to-site IPSec VPNs to the remote sites.

```
! Ensures both user and management traffic is encrypted to the first remote device.
!
access-list remote1 permit ip 10.3.0.0 255.255.0.0 10.5.0.0 255.255.0.0
access-list remote1 permit ip host 172.16.240.151 host 172.16.128.2
access-list remote1 permit ip host 172.16.240.152 host 172.16.128.2
!
!
! Ensures both user and management traffic is encrypted to the second remote device.
!
access-list remote2 permit ip 10.3.0.0 255.255.0.0 10.6.0.0 255.255.0.0
access-list remote2 permit ip host 172.16.240.151 host 172.16.128.5
access-list remote2 permit ip host 172.16.240.152 host 172.16.128.5
!
!
! Defines crypto map and applies it to the outside interface.
!
crypto ipsec transform-set 3dessha esp-3des esp-sha-hmac
crypto map remote1 10 ipsec-isakmp
crypto map remote1 10 match address remote1
crypto map remote1 10 set peer 172.16.128.2
crypto map remote1 10 set transform-set 3dessha
crypto map remote1 20 ipsec-isakmp
crypto map remote1 20 match address remote2
crypto map remote1 20 set peer 172.16.128.5
crypto map remote1 20 set transform-set 3dessha
crypto map remote1 interface outside
!
!
! Defines the use of IKE using pre-shared keys.
```

```
!
isakmp enable outside
isakmp key 7Q!r$y$+xE address 172.16.128.2 netmask 255.255.255.255
isakmp key 7Q!r$y$+xE address 172.16.128.5 netmask 255.255.255.255
isakmp identity address
isakmp policy 10 authentication pre-share
isakmp policy 10 encryption 3des
isakmp policy 10 hash sha
isakmp policy 10 group 2
isakmp policy 10 lifetime 86400
!
```

Campus Module

Figure 18 Detailed Model of Medium Campus Module

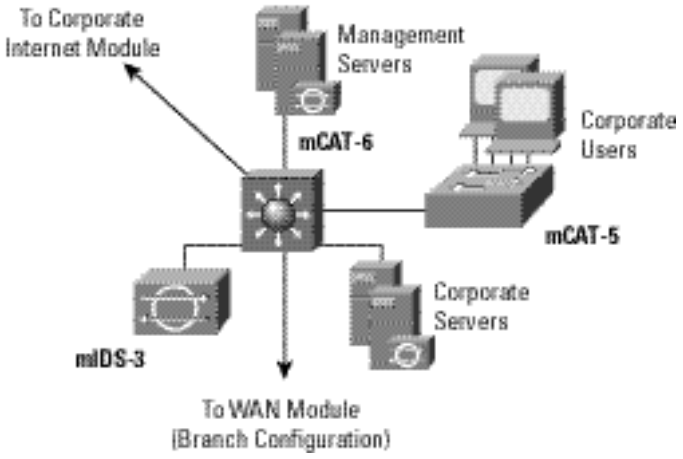

Products Used

- Cisco Catalyst Layer 3 Switch (mCAT-6)
- Cisco Catalyst Layer 2 Switches (mCAT-5)
- Cisco Secure IDS Sensors (mIDS-3)
- Entercept HIDS
- Cisco Secure Policy Manager
- Cisco Secure Access Control Server
- CiscoWorks 2000
- OpenSystems Private I syslog analysis tool
- F-Secure SSH Client
- RSA SecureID OTP System

The following configuration snapshot details the access lists in place on the Catalyst Layer 3 Switch that controls access to the management host virtual LAN (VLAN), as well as the RFC 2827 filtering on the public server VLAN and the building switch VLAN. VLAN 10 defines the corporate user subnet. VLAN 11 defines the corporate intranet server subnet. VLANs 12 and 13 connect to the corporate Internet and WAN modules, respectively. Finally, VLAN 99 defines the management host subnet.

mCAT-6

```
! Corporate user VLAN.
!
interface Vlan10
 ip address 10.3.1.1 255.255.255.0
 ip access-group 101 in
 no ip redirects
 no cdp enable
!
!
! Corporate intranet server VLAN.
!
interface Vlan11
 ip address 10.3.2.1 255.255.255.0
 ip access-group 102 in
 no ip redirects
 no cdp enable
!
!
interface Vlan12
 ip address 10.3.3.1 255.255.255.0
 no ip redirects
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 7 134E031F4158140119
 no cdp enable
!
interface Vlan13
 ip address 10.3.9.1 255.255.255.0
 no ip redirects
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 7 024D105641521F0A7E
 no cdp enable
!
! Management host VLAN.
!
interface Vlan99
 ip address 10.3.8.1 255.255.255.0
 ip access-group 103 out
 no ip redirects
 no cdp enable
!
!
! RFC 2827 filtering on corporate user VLAN.
!
access-list 101 permit ip 10.3.1.0 0.0.0.255 any
access-list 101 deny  ip any any
!
!
! RFC 2827 filtering on corporate intranet server VLAN.
```

```
!
access-list 102 permit ip 10.3.2.0 0.0.0.255 any
access-list 102 deny  ip any any log
!
!
! Example filtering for access to the management subnet (not complete).
!
access-list 103 permit udp host 10.3.2.50 eq domain host 10.3.8.253
access-list 103 permit udp host 10.3.2.50 eq domain host 10.3.8.254
access-list 103 permit tcp host 10.3.2.50 eq www host 10.3.8.253 established
access-list 103 permit tcp host 10.3.2.50 eq www host 10.3.8.254 established
access-list 103 permit tcp host 10.3.2.50 eq ftp host 10.3.8.253 established
access-list 103 permit tcp host 10.3.2.50 eq ftp host 10.3.8.254 established
access-list 103 permit tcp host 10.3.2.50 eq ftp-data host 10.3.8.253
access-list 103 permit tcp host 10.3.2.50 eq ftp-data host 10.3.8.254
access-list 103 permit tcp host 10.3.2.50 host 10.3.8.253 eq 5000
access-list 103 permit udp host 10.3.1.4 host 10.3.8.253 eq syslog
access-list 103 permit udp host 10.3.1.4 host 10.3.8.254 eq syslog
access-list 103 permit tcp host 10.3.1.4 host 10.3.8.253 eq tacacs
access-list 103 permit udp host 10.3.1.4 host 10.3.8.254 eq tftp
access-list 103 permit udp host 10.3.1.4 host 10.3.8.254 gt 1023
access-list 103 permit udp host 10.3.1.4 eq snmp host 10.3.8.254
access-list 103 permit tcp host 10.3.1.4 eq telnet host 10.3.8.253 established
access-list 103 permit tcp host 10.3.1.4 eq telnet host 10.3.8.254 established
access-list 103 deny  ip any any
!
```

Branch versus Headend Considerations

The following configuration is added to the access list on the core switch to allow configuration and security management traffic from the remotely managed devices to the management hosts, when the medium network is configured as a headend.

```
access-list 103 permit udp host 172.16.128.5 host 10.3.8.253 eq syslog
access-list 103 permit udp host 172.16.128.5 host 10.3.8.254 eq syslog
access-list 103 permit tcp host 172.16.128.5 host 10.3.8.253 eq tacacs
access-list 103 permit udp host 172.16.128.5 host 10.3.8.254 eq tftp
access-list 103 permit udp host 172.16.128.5 host 10.3.8.254 gt 1023
access-list 103 permit tcp host 172.16.128.5 eq 22 host 10.3.8.253 established
access-list 103 permit tcp host 172.16.128.5 eq 22 host 10.3.8.254 established
access-list 103 permit tcp host 172.16.128.5 eq 443 host 10.3.8.253 established
access-list 103 permit tcp host 172.16.128.5 eq 443 host 10.3.8.254 established
access-list 103 permit udp host 172.16.128.2 host 10.3.8.253 eq syslog
access-list 103 permit udp host 172.16.128.2 host 10.3.8.254 eq syslog
access-list 103 permit tcp host 172.16.128.2 host 10.3.8.253 eq tacacs
access-list 103 permit udp host 172.16.128.2 host 10.3.8.254 eq tftp
access-list 103 permit udp host 172.16.128.2 host 10.3.8.254 gt 1023
access-list 103 permit tcp host 172.16.128.2 eq 22 host 10.3.8.253 established
access-list 103 permit tcp host 172.16.128.2 eq 22 host 10.3.8.254 established
```

mCAT-5

The following configuration snapshot shows some of the VLAN settings on the Layer 2 switch in this module. Notice that unused ports are disabled. Private VLANs are used on all ports except the uplink to the core switch.

```
! User workstation ports.
!
interface FastEthernet0/1
 port protected
 switchport access vlan 99
 no cdp enable
!
interface FastEthernet0/2
 port protected
 switchport access vlan 99
 no cdp enable
!
!
! Unused ports.
!
interface FastEthernet0/3
 port protected
 shutdown
 no cdp enable
!
interface FastEthernet0/4
 port protected
 shutdown
 no cdp enable
!
!
! Uplink to core switch mCAT-1
!
interface GigabitEthernet0/1
 switchport access vlan 99
 no cdp enable
!
!
!
! Management interface to the switch.
!
interface VLAN99
 ip address 10.3.1.4 255.255.255.0
 no ip directed-broadcast
 no ip route-cache
!
```

Figure 19 Detailed Model of WAN Module

Products Used

• Cisco IOS Router (mIOS-4)

Remote User Design

Figure 20 Detailed Model of Remote-User Designs

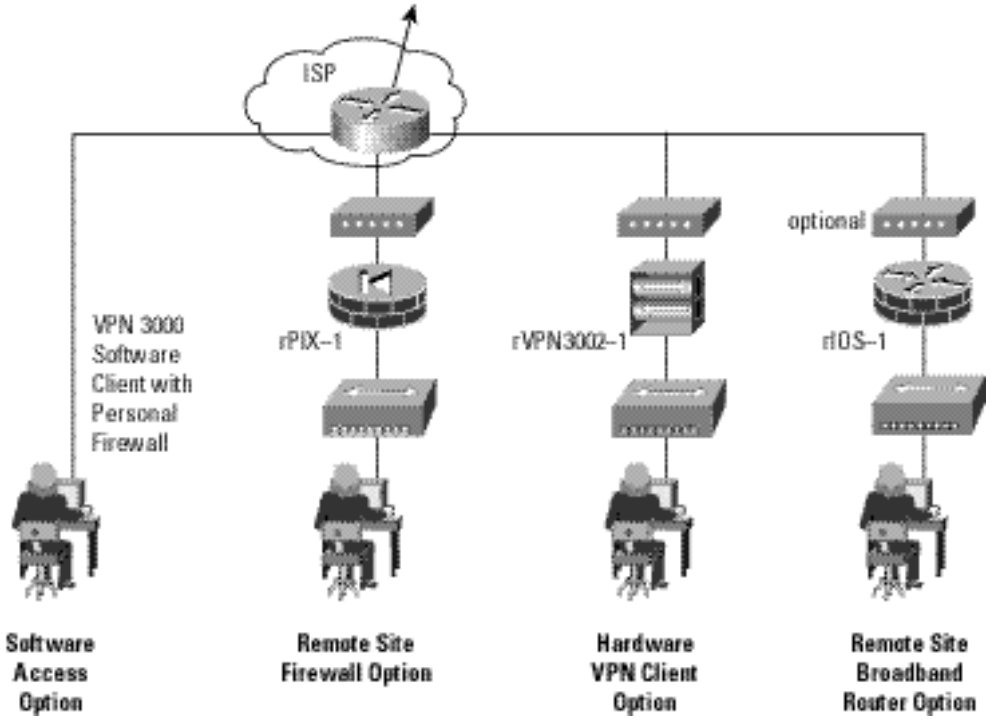

Products Used

- Cisco IOS Router with 3DES encryption support (rIOS-1)
- Cisco VPN 3002 Hardware Client (rVPN3002-1)
- Cisco Secure PIX Firewall (rPIX-1)
- Cisco VPN 3000 Software Client
- Cisco MicroHub (or integrated into Layer 3 device)
- Zone Alarm Pro Personal Firewall

The following are configuration snapshots from some of the SAFE remote-user designs.

rIOS-1 (Remote-Site Router Option)
The following shows the configuration of the IPSec tunnels back to the corporate headquarters.

```
crypto isakmp policy 1
 encr 3des
 authentication pre-share
 group 2
crypto isakmp key 7Q!r$y$+xE address 172.16.240.1
!
!
crypto ipsec transform-set 3dessha esp-3des esp-sha-hmac
!
crypto map remote1 10 ipsec-isakmp
 set peer 172.16.240.1
 set transform-set 3dessha
 match address 101
!
!
! The first lines of the following access-list specifies all traffic from
! network 10.5.0.0 to networks 10.3.0.0 will be encrypted.
!
! The last two lines of the access-list allow encryption of all configuration and
! security management traffic from the remote site router to the headquarters
! management hosts.
!
access-list 101 permit ip 10.5.0.0 0.0.255.255 10.3.0.0 0.0.255.255
access-list 101 permit ip host 172.16.128.2 host 172.16.240.151
access-list 101 permit ip host 172.16.128.2 host 172.16.240.152
!
!
```

The following shows the access control on the private side (FastEthernet0/0) and the public side (FastEthernet0/1) of the router, as well as the application of the Cisco IOS Firewall to the interfaces.

```
interface FastEthernet0/0
 ip address 10.5.1.2 255.255.255.0
 ip access-group 105 in
 ip nat inside
 ip inspect remote_fw in
!
interface FastEthernet0/1
 ip address 172.16.128.2 255.255.255.0
 ip access-group 102 in
 ip nat outside
 crypto map remote1
!
! IKE and ESP traffic must be allowed from the headquarters IPSec peer. All
! traffic from network 10.5.0.0 to networks 10.3.0.0 must also
! be allowed. Finally, traffic from the management hosts is allowed.
!
access-list 102 permit ip 10.3.0.0 0.0.255.255 10.5.0.0 0.0.255.255
access-list 102 deny ip 10.0.0.0 0.255.255.255 any
access-list 102 deny ip 192.168.0.0 0.0.255.255 any
access-list 102 permit icmp any host 172.16.128.2 echo-reply
access-list 102 permit icmp any host 172.16.128.2 unreachable
access-list 102 permit esp host 172.16.240.1 host 172.16.128.2
access-list 102 permit udp host 172.16.240.1 host 172.16.128.2 eq isakmp
access-list 102 permit tcp host 172.16.240.151 host 172.16.128.2 eq 22
access-list 102 permit tcp host 172.16.240.152 host 172.16.128.2 eq 22
access-list 102 permit tcp host 172.16.240.152 eq tacacs host 172.16.128.2
access-list 102 permit udp host 172.16.240.151 host 172.16.128.2 gt 1023
access-list 102 deny ip any any log
!
!
! RFC 2827 filtering only allows 10.5.0.0 addresses to access both the corporate
! headquarters and the Internet.
!
access-list 105 permit ip 10.5.0.0 0.0.255.255 any
access-list 105 deny  ip any any log
!
```

The following shows the configuration of many-to-one NAT on the router. All devices within the remote site that access the Internet will use the public address of the router.

```
ip nat pool remote_pool 172.16.128.2 172.16.128.2 netmask 255.255.255.0
ip nat inside source route-map nat_internet pool remote_pool
!
route-map nat_internet permit 10
 match ip address 104
!
access-list 104 deny  ip 10.5.0.0 0.0.255.255 10.0.0.0 0.255.255.255
access-list 104 permit ip 10.5.0.0 0.0.255.255 any
```

```
!
```

rPIX-1 (Remote-Site Firewall Option)
The following shows the configuration of the IPSec tunnels back to the corporate headquarters.

```
crypto ipsec transform-set 3dessha esp-3des esp-sha-hmac
crypto map remote1 10 ipsec-isakmp
crypto map remote1 10 match address remote1
crypto map remote1 10 set peer 172.16.240.1
crypto map remote1 10 set transform-set 3dessha
crypto map remote1 interface outside
isakmp enable outside
isakmp key 7Q!r$y$+xE address 172.16.240.1 netmask 255.255.255.255
isakmp identity address
isakmp policy 10 authentication pre-share
isakmp policy 10 encryption 3des
isakmp policy 10 hash sha
isakmp policy 10 group 2
isakmp policy 10 lifetime 86400
!
! The first line of the following access-list specifies all traffic from
! network 10.6.0.0 to networks 10.3.0.0 will be encrypted.
!
! The last two lines of the access-list allow encryption of all configuration and
! security management traffic from the remote site firewall to the headquarters
! management hosts.
!
access-list remote1 permit ip 10.6.0.0 255.255.0.0 10.3.0.0 255.255.0.0
access-list remote1 permit ip host 172.16.128.5 host 172.16.240.151
access-list remote1 permit ip host 172.16.128.5 host 172.16.240.152
!
```

The following shows the addressing and access control on the private side (inside) and the public side (outside) of the firewall.

```
nameif ethernet0 outside security0
nameif ethernet1 inside security100
!
ip address outside 172.16.128.5 255.255.255.0
ip address inside 10.6.1.1 255.255.255.0
!
access-group out in interface outside
access-group in in interface inside
!
! RFC 2827 filtering only allows 10.6.0.0 addresses to access both the corporate
! headquarters and the Internet.
!
access-list in permit ip 10.6.0.0 255.255.0.0 any
!
! Allows encrypted traffic from corporate headquarters.
!
access-list out permit ip 10.3.0.0 255.255.0.0 10.6.0.0 255.255.0.0
```

```
!
! RFC 1918 filtering. Note network 172.16.x.x was not included in the
! filter here since it is used to simulate the ISP in the lab.
!
access-list out deny ip 10.0.0.0 255.0.0.0 any
access-list out deny ip 192.168.0.0 255.255.0.0 any
!
! Allow echo replies and path MTU discovery (PMTU) traffic.
!
access-list out permit icmp any host 172.16.128.5 echo-reply
access-list out permit icmp any host 172.16.128.5 unreachable
!
! Allow ESP and IKE traffic from the corporate headquarters peer.
!
access-list out permit esp host 172.16.240.1 host 172.16.128.5
access-list out permit udp host 172.16.240.1 host 172.16.128.5 eq isakmp
```

The following shows the configuration of many-to-one NAT on the firewall. All devices within the remote site that access the Internet will use the public address of the firewall.

```
global (outside) 100 interface
nat (inside) 0 access-list nonat
nat (inside) 100 10.6.1.0 255.255.255.0 0 0
!
!  The access-list prevents any traffic destined for the corporate site from using
!  address translation.
!
access-list nonat permit ip 10.6.0.0 255.255.0.0 10.0.0.0 255.0.0.0
access-list nonat deny ip 10.6.0.0 255.255.0.0 any
!
```

Appendix B: Network Security Primer

The Need for Network Security

The Internet is constantly changing the way we live and conduct business. These changes are occurring both in the ways that we currently experience (e-commerce, real-time information access, e-learning, expanded communication options, and so forth), and in ways we have yet to experience. Imagine a day when your enterprise can make all its telephone calls over the Internet for free. Or perhaps on a more personal note, consider logging on to a daycare provider's Web site to check how your child is doing throughout the day. As a society, we are just beginning to unlock the potential of the Internet. But with the unparalleled growth of the Internet comes unprecedented exposure of personal data, critical enterprise resources, government secrets, and so forth. Every day hackers pose an increasing threat to these entities with several different types of attacks. These attacks, outlined in the next section, have become both more prolific and easier to implement. There are two primary reasons for this problem.

First is the ubiquity of the Internet. With millions of devices currently connected to the Internet-and millions more on the way-a hacker's access to vulnerable devices will continue to increase. The ubiquity of the Internet has also allowed hackers to share knowledge on a global scale. A simple Internet search on the words "hack," "crack," or "phreak" yields thousands of sites, many of which contain malicious code, or the means with which to use that code.

Second is the pervasiveness of easy-to-use operating systems and development environments. This factor has reduced the overall ingenuity and knowledge required by hackers. A truly remarkable hacker can develop easy-to-use applications that can be distributed to the masses. Several hacker tools that are available in the public domain require merely an IP address or host name and a click of a mouse button to execute an attack.

Network Attack Taxonomy

Network attacks can be as varied as the systems that they attempt to penetrate. Some attacks are elaborately complex, whereas others are performed unknowingly by a well-intentioned device operator. It is important to understand some of the inherent limitations of the TCP/IP protocol when evaluating the types of attacks. When the Internet was formed, it linked various government entities and universities to one another with the express purpose of facilitating learning and research. The original architects of the Internet never anticipated the kind of widespread adoption the Internet has achieved today. As a result, in the early days of the *Internet Protocol* (IP), security was not designed into the specification. For this reason, most IP implementations are inherently insecure. Only after many years and thousands of Requests for Comments (RFCs) do we have the tools to begin to deploy IP securely. Because specific provisions for IP security were not designed from the onset, it is important to augment IP implementations with network security practices, services, and products to mitigate the inherent risks of the Internet Protocol. Following is a brief discussion of the types of attacks commonly seen on IP networks and how these attacks can be mitigated.

Packet Sniffers

A packet sniffer is a software application that uses a network adapter card in promiscuous mode (a mode in which the network adapter card sends all packets received on the physical network wire to an application for processing) to capture all network packets that are sent across a particular collision domain. Sniffers are used legitimately in networks today to aid in troubleshooting and traffic analysis. However, because several network applications send data in cleartext (Telnet, *File Transfer Protocol* [FTP], *Simple Message Transfer Protocol* [SMTP], *Post Office Protocol* [POP3], and so forth), a packet sniffer can provide meaningful and often sensitive information, such as usernames and passwords.

One serious problem with acquiring usernames and passwords is that users often reuse their login names and passwords across multiple applications and systems. In fact, many users employ a single password for access to all accounts and applications. If an application is run in client-server mode and authentication information is sent across the network in cleartext, then it is likely that this same authentication information can be used to gain access to other corporate or external resources. Because hackers know and use human characteristics (attack methods known collectively as social engineering

attacks), such as using a single password for multiple accounts, they are often successful in gaining access to sensitive information. In a worst-case scenario, a hacker gains access to a system-level user account, which the hacker uses to create a new account that can be used at any time as a back door to break into a network and its resources.

You can mitigate the threat of packet sniffers in several ways:

- *Authentication*—Using strong authentication is a first option for defense against packet sniffers. Strong authentication can be broadly defined as a method of authenticating users that cannot easily be circumvented. A common example of strong authentication is *one-time passwords* (OTPs). An OTP is a type of two-factor authentication. Two-factor authentication involves using something you have combined with something you know. *Automated teller machines* (ATMs) use two-factor authentication. A customer needs both an ATM card and a *personal identification number* (PIN) to make transactions. With OTP you need a PIN and your token card to authenticate to a device or software application. A token card is a hardware or software device that generates new, seemingly random, passwords at specified intervals (usually 60 seconds). A user combines that random password with a PIN to create a unique password that works only for one instance of authentication. If a hacker learns that password by using a packet sniffer, the information is useless because the password has already expired. Note that this mitigation technique is effective only against a sniffer implementation that is designed to grab passwords. Sniffers deployed to learn sensitive information (such as mail messages) will still be ineffective.
- *Switched infrastructure*—Another method to counter the use of packet sniffers in your environment is to deploy a switched infrastructure. For example, if an entire organization deploys switched Ethernet, hackers can gain access only to the traffic that flows on the specific port to which they connect. A switched infrastructure obviously does not eliminate the threat of packet sniffers, but it can greatly reduce their effectiveness.
- *Antisniffer tools*—A third method used against sniffers is to employ software and hardware designed to detect the use of sniffers on a network. Such software and hardware does not completely eliminate the threat, but like many network security tools, they are part of the overall system. These so-called “antisniffers” detect changes in the response time of hosts to determine if the hosts are processing more traffic than their own. One such network security software tool, which is available from Security Software Technologies, is called AntiSniff. For more information, refer to the URL http://www.securitysoftwaretech.com/antisniff/
- *Cryptography*—The most effective method for countering packet sniffers does not prevent or detect packet sniffers, but rather renders them irrelevant. If a communication channel is cryptographically secure, the only data a packet sniffer will detect is cipher text (a seemingly random string of bits) and not the original message. The Cisco deployment of network-level cryptography is based on *IP Security* (IPSec), which is a standard method for networking devices to communicate privately using IP. Other cryptographic protocols for network management include *Secure Shell Protocol* (SSH) and *Secure Sockets Layer* (SSL).

IP Spoofing

An IP spoofing attack occurs when a hacker inside or outside a network impersonates the conversations of a trusted computer. A hacker can do this in one of two ways. The hacker uses either an IP address that is within the range of trusted IP addresses for a network or an authorized external IP address that is trusted and to which access is provided to specified resources on a network. IP spoofing attacks are often a launch point for other attacks. The classic example is to launch a *denial-of-service* (DoS) attack using spoofed source addresses to hide the hacker's identity.

Normally, an IP spoofing attack is limited to the injection of malicious data or commands into an existing stream of data that is passed between a client and server application or a peer-to-peer network connection. To enable bidirectional communication, the hacker must change all routing tables to point to the spoofed IP address. Another approach hackers sometimes take is to simply not worry about receiving any response from the applications. If a hacker tries to obtain a sensitive file from a system, application responses are unimportant.

However, if a hacker manages to change the routing tables to point to the spoofed IP address, the hacker can receive all the network packets that are addressed to the spoofed address and reply just as any trusted user can.

The threat of IP spoofing can be reduced, but not eliminated, through the following measures.

- *Access control*—The most common method for preventing IP spoofing is to properly configure access control. To reduce the effectiveness of IP spoofing, configure access control to deny any traffic from the external network that has a source address that should reside on the internal network. Note that this helps prevent spoofing attacks only if the internal addresses are the only trusted addresses. If some external addresses are trusted, this method is not effective.
- *RFC 2827 filtering*—You can also prevent users of a network from spoofing other networks (and be a good Internet citizen at the same time) by preventing any outbound traffic on your network that does not have a source address in your organization's own IP range. Your *Internet service provider* (ISP) can also implement this type of filtering, which is collectively referred to as RFC 2827 filtering. This filtering denies any traffic that does not have the source address that was expected on a particular interface. For example, if an ISP is providing a connection to the IP address 15.1.1.0/24, the ISP could filter traffic so that only traffic sourced from address 15.1.1.0/24 can enter the ISP router from that interface. Note that unless all ISPs implement this type of filtering, its effectiveness is significantly reduced. Also, the further you get from the devices you want to filter, the more difficult it becomes to do that filtering at a granular level. For example, performing RFC 2827 filtering at the access router to the Internet requires that you allow your entire major network number (that is, 10.0.0.0/8) to traverse the access router. If you perform filtering at the distribution layer, as in this architecture, you can achieve more specific filtering (that is, 10.1.5.0/24).

The most effective method for mitigating the threat of IP spoofing is the same as the most effective method for mitigating the threat of packet sniffers: namely, eliminating its effectiveness. IP spoofing can function correctly only when devices use IP address-based authentication. Therefore, if you use additional authentication methods, IP spoofing attacks are irrelevant. Cryptographic authentication is the best form of additional authentication, but when that is not possible, strong two-factor authentication using OTP can also be effective.

Denial of Service

Certainly the most publicized form of attack, *denial of service* (DoS) attacks are also among the most difficult to completely eliminate. Even among the hacker community, DoS attacks are regarded as trivial and considered bad form because they require so little effort to execute. Still, because of their ease of implementation and potentially significant damage, DoS attacks deserve special attention from security administrators. If you are interested in learning more about DoS attacks, researching the methods employed by some of the better known attacks can be useful. These attacks include the following:

- TCP SYN Flood
- Ping of Death
- Tribe Flood Network (TFN) and Tribe Flood Network 2000 (TFN2K)
- Trinoo
- Stacheldraht
- Trinity

Another excellent resource on the topic of security is the *Computer Emergency Response Team* (CERT). They have published an excellent paper on dealing with DoS attacks; you can find it at the following URL: http://www.cert.org/tech_tips/denial_of_service.html.

DoS attacks are different from most other attacks because they are generally not targeted at gaining access to your network or the information on your network. These attacks focus on making a service unavailable for normal use, and they are typically accomplished by exhausting some resource limitation on the network or within an operating system or application.

When involving specific network server applications, such as a Web server or an FTP server, these attacks can focus on acquiring and keeping open all the available connections supported by that server, effectively locking out valid users of the server or service. DoS attacks can also be implemented using common Internet protocols, such as TCP and *Internet Control Message Protocol* (ICMP). Most DoS attacks exploit a weakness in the overall architecture of the system being attacked rather than a software bug or security hole. However, some attacks compromise the performance of your network by flooding the network with undesired—and often useless—network packets and by providing false information about the status of network resources. This type of attack is often the most difficult to prevent because it requires coordination with your upstream network provider. If traffic meant to consume your available bandwidth is not stopped there, denying it at the point of entry into your network will do little good because your available bandwidth has already been consumed. When this type of attack is launched from many different systems at the same time, it is often referred to as a *distributed denial of service* attack (DDoS).

The threat of DoS attacks can be reduced through the following three methods:

- *Antispoof features*—Proper configuration of antispoof features on your routers and firewalls can reduce your risk. This configuration includes RFC 2827 filtering at a minimum. If hackers cannot mask their identities, they might not attack.
- *Anti-DoS features*—Proper configuration of anti-DoS features on routers and firewalls can help limit the effectiveness of an attack. These features often involve limits on the amount of half-open connections that a system allows open at any given time.
- *Traffic rate limiting*—An organization can implement traffic rate limiting with its Internet service provider (ISP). This type of filtering limits the amount of nonessential traffic that crosses network segments to a certain rate. A common example is to limit the amount of ICMP traffic allowed into a network because it is used only for diagnostic purposes. ICMP-based DDoS attacks are common.

Password Attacks

Hackers can implement password attacks using several different methods, including brute-force attacks, Trojan horse programs, IP spoofing, and packet sniffers. Although packet sniffers and IP spoofing can yield user accounts and passwords, password attacks usually refer to repeated attempts to identify a user account or password. These repeated attempts are called brute-force attacks.

Often, a brute-force attack is performed using a program that runs across the network and attempts to log in to a shared resource, such as a server. When hackers successfully gain access to resources, they have the same rights as the users whose accounts have been compromised to gain access to those resources. If the compromised accounts have sufficient privileges, the hackers can create back doors for future access without concern for any status and password changes to the compromised user accounts.

Another problem exists whereby users have the same (possibly strong) password on every system they connect to. Often, this includes personal systems, corporate systems, and systems on the Internet. Because that password is only as secure as the most weakly administered host that contains it, if that host is compromised, hackers have a whole range of hosts on which they can try the same password.

You can most easily eliminate password attacks by not relying on plaintext passwords in the first place. Using OTP or cryptographic authentication can virtually eliminate the threat of password attacks. Unfortunately, not all applications, hosts, and devices support these authentication methods. When standard passwords are used, it is important to choose a password that is difficult to guess. Passwords should be at least eight characters long and contain uppercase letters, lowercase letters, numbers, and special characters (#, %, $, and so forth). The best passwords are randomly generated but are very difficult to remember, often leading users to write their passwords down. Such attacks can also be mitigated by disabling the account after a certain number of incorrect attempts.

Several advances have been made relative to password maintenance both for the user and the administrator. Software applications are now available that encrypt a list of passwords to be stored on a handheld computer. This scenario allows the user to remember only one complex password and have the remaining passwords stored securely within the application. From the standpoint of the administrator, several methods exist to brute-force attack your own users' passwords. One such method involves a tool used by the hacker community called LC3 (formerly L0phtCrack). L3 brute-force attacks Windows NT passwords and can point out when a user has chosen a password that is very easy to guess. For more information, refer to the following URL: http://www.atatake.com/research/lc3/index.html.

Man-in-the-Middle Attacks

A man-in-the-middle attack requires that the hacker have access to network packets that come across a network. An example of such a configuration could be someone who is working for an ISP, who has access to all network packets transferred between his/her employer's network and any other network. Such attacks are often implemented using network packet sniffers and routing and transport protocols. The possible uses of such attacks are theft of information, hijacking of an ongoing session to gain access to private network resources, traffic analysis to derive information about a network and its users, denial of service, corruption of transmitted data, and introduction of new information into network sessions.

Man-in-the-middle attacks can be effectively mitigated only through the use of cryptography. If someone hijacks data in the middle of a cryptographically private session, all the hacker will see is cipher text, and not the original message. Note that if a hacker can learn information about the cryptographic session (such as the session key), man-in-the-middle attacks are still possible.

Application Layer Attacks

Application layer attacks can be implemented using several different methods. One of the most common methods is exploiting well-known weaknesses in software that are commonly found on servers, such as sendmail, *Hypertext Transfer Protocol* (HTTP), and FTP. By exploiting these weaknesses, hackers can gain access to a computer with the permissions of the account running the application, which is usually a privileged system-level account. These application layer attacks are often widely publicized in an effort to allow administrators to rectify the problem with a patch. Unfortunately, many hackers also subscribe to these same mailing lists, resulting in their learning about the attack at the same time (if they haven't discovered it already).

The primary problem with application layer attacks is that they often use ports that are allowed through a firewall. For example, a hacker executing a known vulnerability against a Web server often uses TCP port 80 in the attack. Because the Web server serves pages to users, a firewall needs to allow access on that port. From the perspective of the firewall, it is merely standard port 80 traffic.

Application layer attacks can never be completely eliminated. New vulnerabilities are always being discovered and publicized to the Internet community. The best way to reduce your risk is by practicing good system administration. A few measures you can take to reduce your risks follow:

- Read OS and network log files or have them analyzed by log analysis applications.
- Subscribe to mailing lists that publicize vulnerabilities such as Bugtraq (http://www.securityfocus.com) and the *Computer Emergency Response Team* (CERT) (http://www.cert.org).
- Keep your OS and applications current with the latest patches

In addition to proper system administration, using *intrusion detection systems* (IDSs) can aid in this effort. There are two complementary IDS technologies:

- *Network-based IDS* (NIDS) operates by watching all packets traversing a particular collision domain. When NIDS sees a packet or series of packets that match a known or suspect attack, it can flag an alarm or terminate the session.
- *Host-based IDS* (HIDS) operates by inserting agents into the host to be protected. It is then concerned only with attacks generated against that one host.

IDS systems operate by using attack signatures. Attack signatures are the profile for a particular attack or kind of attack. They specify certain conditions that must be met before traffic is deemed to be an attack. In the physical world, IDS can be most closely compared to an alarm system or security camera. The greatest limitation of the IDS system is the amount of false-positive alarms a particular system generates. Tuning IDS to prevent such false alarms is critical to the proper operation of IDS in a network.

Network Reconnaissance

Network reconnaissance refers to the overall act of learning information about a target network by using publicly available information and applications. When hackers attempt to penetrate a particular network, they often need to learn as much information as possible about the network before launching attacks. This can take the form of *Domain Name System* (DNS) queries, ping sweeps, and port scans. DNS queries can reveal such information as who owns a particular domain and what addresses have been assigned to that domain. Ping sweeps of the addresses revealed by the DNS queries can present a picture of the live hosts in a particular environment. After such a list is generated, port-scanning tools can cycle through all well-known ports to provide a complete list of all services running on the hosts discovered by the ping sweep. Finally, the hackers can examine the characteristics of the applications that are running on the hosts. This scenario can lead to specific information that is useful when the hacker attempts to compromise that service.

Network reconnaissance cannot be prevented entirely. If ICMP echo and echo reply is turned off on edge routers, for example, ping sweeps can be stopped, but at the expense of network diagnostic data. However, port scans can easily be run without full ping sweeps; they simply take longer because they need to scan IP addresses that might not be live. IDS at the network and host levels can usually notify an administrator when a reconnaissance gathering attack is under way. This allows the administrator to better prepare for the coming attack or to notify the ISP who is hosting the system that is launching the reconnaissance probe.

Trust Exploitation

Although not an attack in and of itself, trust exploitation refers to an attack where an individual takes advantage of a trust relationship within a network. The classic example is a perimeter network connection from a corporation. These network segments often house DNS, *Simple Message Transfer Protocol* (SMTP), and HTTP servers. Because they all reside on the same segment, a compromise of one system can lead to the compromise of other systems because they might trust other systems attached to their same network. Another example is a system on the outside of a firewall that has a trust relationship with a system on the inside of a firewall. When the outside system is compromised, it can take advantage of that trust relationship to attack the inside network.

You can mitigate trust exploitation-based attacks through tight constraints on trust levels within a network. Systems on the outside of a firewall should never be absolutely trusted by systems on the inside of a firewall. Such trust should be limited to specific protocols and should be authenticated by something other than an IP address where possible.

Port Redirection

The port-redirection attack is a type of trust-exploitation attack that uses a compromised host to pass traffic through a firewall that would otherwise be dropped. Consider a firewall with three interfaces and a host on each interface. The host on the outside can reach the host on the public services segment (commonly referred to as a DMZ), but not the host on the inside. The host on the public services segment can reach the host on both the outside and the inside. If hackers were able to compromise the public services segment host, they could install software to redirect traffic from the outside host directly to the inside host. Though neither communication violates the rules implemented in the firewall, the outside host has now achieved connectivity to the inside host through the port-redirection process on the public services host. An example of an application that can provide this type of access is netcat. For more information, refer to http://www.avian.org.

Port redirection can be mitigated primarily through the use of proper trust models (as mentioned earlier). If we assume that a system is under attack, host-based IDS can help detect and prevent a hacker installing such utilities on a host.

Unauthorized Access

Although unauthorized-access attacks are not a specific type of attack, they refer to most attacks executed in networks today. In order for someone to brute-force attack a Telnet login, he/she must first get the Telnet prompt on a system. Upon connection to the Telnet port, a message might indicate: “authorization required to use this resource.” If the hacker continues to attempt access, his/her actions become “unauthorized.” These kinds of attacks can be initiated on both the outside and inside of a network.

Mitigation techniques for unauthorized-access attacks are very simple. They involve reducing or eliminating the ability of a hacker to gain access to a system using an unauthorized protocol. An example would be preventing hackers from having access to the Telnet port on a server that needs to provide Web services to the outside. If a hacker cannot reach that port, it is very difficult to attack it. The primary function of a firewall in a network is to prevent simple unauthorized-access attacks.

One of the most popular types of firewall is a stateful firewall. These firewalls inspect traffic in both directions and dynamically open ports as applications require them. For example, active FTP negotiates a specific port for the data transfer. The stateful firewall will see this information in the packet and will allow that port to communicate between the server and client. This is very different than a standard packet filtering device which has no application awareness. These devices merely look at layer 3 and 4 data when making an access control decision. In the FTP example above, the administrator would need to manually open all TCP high ports (>1023) from the outside in order for FTP to be successful.

Virus and Trojan Horse Applications

The primary vulnerabilities for end-user workstations are viruses and Trojan horse attacks. Viruses refer to malicious software that is attached to another program to execute a particular unwanted function on a user's workstation. An example of a virus is a program that is attached to command.com (the primary interpreter for windows systems), which deletes certain files and infects any other versions of command.com that it can find. A Trojan horse is different only in that the entire application was written to look like something else, when in fact it is an attack tool. An example of a Trojan horse is a software application that runs a simple game on the user’s workstation. While the user is occupied with the game, the Trojan horse mails a copy of itself to every user in the user’s address book. Then other users get the game and play it, thus spreading the Trojan horse.

These kinds of applications can be contained through the effective use of antivirus software at the user level and potentially at the network level. Antivirus software can detect most viruses and many Trojan horse applications and prevent them from spreading in the network. Keeping up-to-date with the latest developments in these sorts of attacks can also lead to a more effective posture against these attacks. As new virus or Trojan applications are released, enterprises need to keep up-to-date with the latest antivirus software, and application versions.

What Is a "Security Policy?"

A security policy can be as simple as an acceptable use policy for network resources or it can be several hundred pages in length and detail every element of connectivity and associated policies. Although somewhat narrow in scope, RFC 2196 suitably defines a security policy as follows:

A security policy is a formal statement of the rules by which people who are given access to an organization's technology and information assets must abide.

This document does not attempt to go into detail on the development of a security policy. RFC 2196 has some good information on the subject, and numerous locations on the Web have example policies and guidelines. The following Web pages may assist the interested reader:

- RFC 2196 "Site Security Handbook" http://www.ietf.org/rfc/rfc2196.txt
- An sample security policy for the University of Illinois http://www.aits.uillinois.edu/security/securestandards.html
- Design and Implementation of the Corporate Security Policyhttp://www.knowcisco.com/content/1578700434/ch06.shtml

The Need for a Security Policy

It is important to understand that network security is an evolutionary process. No one product can make an organization "secure." True network security comes from a combination of products and services, combined with a comprehensive security policy and a commitment to adhere to that policy from the top of the organization down. In fact, a properly implemented security policy without dedicated security hardware can be more effective at mitigating the threat to enterprise resources than a comprehensive security product implementation without an associated policy.

Management Protocols and Functions

- *SSH and SSL*—Provide encrypted and authenticated remote access to the managed device
- *Telnet*—Provides remote access to the managed device in cleartext
- *Syslog*—Provides device logging and alarm information to management servers
- *Trivial File Transfer Protocol (TFTP)*—Allows administrators to transport configuration files of managed devices to management servers
- *Simple Network Management Protocol (SNMP)*—Provides transport of device information to management servers
- *Network Time Protocol (NTP)*—Provides synchronization of clocks within devices

The management functions enabled on the devices are discussed in the following sections:

Configuration management (SSH, SSL, Telnet)—When possible, the use of IPSec, SSH, SSL, or any other encrypted and authenticated transport that allows management information to traverse it should be used for remote access to devices. However, if the device does not support any of these protocols, Telnet may have to be used, although this protocol is not highly recommended. The network administrator should recognize that the data within a Telnet session is sent as cleartext, and may be intercepted by anyone with a packet sniffer located along the data path between the device and the management server. The data may include sensitive information, such as the configuration of the device itself, passwords, and so on. Regardless of whether SSH, SSL, or Telnet is used for remote access to the device, access control lists (ACLs) should be configured to allow only management servers to connect to the device. All attempts from other IP addresses should be denied and logged. RFC 2827 filtering at the ingress router should also be implemented to mitigate the chance of an attacker from outside the network spoofing the addresses of the management hosts. SSH uses TCP port 22, Telnet uses TCP port 23, and SSH uses TCP port 443.

Logging—Syslog is also sent as cleartext between the managed device and the management host. Syslog has no packet-level integrity checking to ensure that the packet contents have not been altered in transit. An attacker may alter syslog data in order to confuse a network administrator during an attack. Where possible, syslog traffic may be encrypted within an IPSec tunnel in order to mitigate the chance of its being altered in transit. Where the syslog data cannot be encrypted within an IPSec tunnel because of cost or the capabilities of the device itself, the network administrator should note that there is a potential for the syslog data to be falsified by an attacker. When allowing syslog access from devices on the outside of a firewall, RFC 2827 filtering at the egress router should be implemented. This scenario will mitigate the chance of an attacker from outside the network spoofing the address of the managed device, and sending false syslog data to the management hosts. ACLs should also be implemented on the firewall in order to allow syslog data from only the managed devices themselves to reach the management hosts. This scenario prevents an attacker from sending large amounts of false syslog data to a management server in order to confuse the network administrator during an attack. Syslog uses UDP port 514.

TFTP—Many network devices use TFTP for transferring configuration or system files across the network. TFTP uses *User Datagram Protocol* (UDP) port 69 as well as high UDP ports (>1023) for the data stream between the device and the TFTP server, and also sends data in cleartext. The network administrator should recognize that the data within a TFTP session may be intercepted by anyone with a packet sniffer located along the data path between the device and the management server. The data may include sensitive information, such as the configuration of the device itself, and so on. Where possible, TFTP traffic should be encrypted within an IPSec tunnel in order to mitigate the chance of its being intercepted.

SNMP—SNMP is a network management protocol that can be used to retrieve information from a network device (commonly referred to as read-only access) or to remotely configure parameters on the device (commonly referred to as read-write access). SNMP agents listen on UDP port 161. SNMP uses passwords, called community strings, within each message as a very simple form of security. Unfortunately, most implementations of SNMP on networking devices today send the community string in cleartext along with the message. Therefore, SNMP messages may be intercepted by anyone with a packet sniffer located along the data path between the device and the management server, and the community string may be compromised. When the community string is compromised, an attacker could reconfigure the device if read-write access via SNMP is allowed. Therefore, it is recommend that you configure SNMP with only read-only community strings. You can further protect yourself by setting up access control on the device you wish to manage via SNMP to allow only the appropriate management hosts access.

NTP—Network Time Protocol is used to synchronize the clocks of various devices across a network. Synchronization of the clocks within a network is critical for digital certificates, and for correct interpretation of events within syslog data. A secure method of providing clocking for the network is for the network administrator to implement his/her own master clock for the private network synchronized to *Coordinated Universal Time* (UTC) via satellite or radio. However, clock sources are available to synchronize to via the Internet, if the network administrator does not wish to implement his/her own master clock because of costs or other reasons. An attacker could attempt a DoS attack on a network by sending bogus NTP data across the Internet in an attempt to change the clocks on network devices in such a manner that digital certificates are considered invalid. Further, an attacker could attempt to confuse a network administrator during an attack by disrupting the clocks on network devices. This scenario would make it difficult for the network administrator to determine the order of syslog events on multiple devices. Version 3 and above of NTP supports a cryptographic authentication mechanism between peers. The use of the authentication mechanism as well as ACLs that specify which network devices are allowed to synchronize with other network devices is recommended to help mitigate against such a scenario. The network administrator should weigh the cost benefits of pulling clock from the Internet with the possible risk of doing so and allowing it through the firewall. Many NTP servers on the Internet do not require any authentication of peers. Therefore, the network administrator must trust that the clock itself is reliable, valid, and secure. NTP uses UDP port 123.

Appendix C: Architecture Taxonomy

Application server—The application server provides application services directly or indirectly for enterprise end users. Services can include work-flow, general office, and security applications.

Firewall (stateful)—This stateful packet-filtering device maintains state tables for IP-based protocols. Traffic is allowed to cross the firewall only if it conforms to the access-control filters defined, or if it is part of an already established session in the state table.

Host IDS—Host intrusion detection system (HIDS) is a software application that monitors activity on an individual host. Monitoring techniques can include validating operating system and application calls, and checking log files, file system information, and network connections.

Network IDS—Network IDS (NIDS) is typically used in a nondisruptive manner. This device captures traffic on a LAN segment and tries to match the real-time traffic against known attack signatures. Signatures range from atomic (single packet and direction) signatures to composite (multipacket) signatures that require state tables and Layer 7 application tracking.

Cisco IOS Firewall—The Cisco IOS Firewall is a stateful packet-filtering firewall that runs natively on Cisco IOS Software.

Cisco IOS Router—The Cisco IOS Router constitutes a wide spectrum of flexible network devices that provide many routing and security services for all performance requirements. Most devices are modular and have a range of LAN and WAN physical interfaces.

Layer 2 Switch—A Layer 2 Switch provides bandwidth and *virtual LAN* (VLAN) services to network segments at the Ethernet level. Typically these devices offer 10/100 individual switched ports, Gigabit Ethernet uplinks, VLAN trunking, and Layer 2 filtering features.

Layer 3 Switch—A Layer 3 switch provides high-throughput functions similar to those of a Layer 2 switch with added routing, quality-of-service (QoS), and security features. These switches often have the capability of special function processors.

Management server—The management server provides network management services for the operators of enterprise networks. Services can include general configuration management, monitoring of network security devices, and operation of the security functions.

SMTP content-filtering server—This server application typically runs on an external SMTP server that monitors the content (including attachments) of incoming and outgoing mail in order to decide whether that mail is authorized to be forwarded as is, altered and forwarded, or dropped.

URL filtering server—This server application typically runs on a standalone server that monitors URL requests forwarded to it by a network device and informs the network device whether the request should be forwarded on to the Internet. This setup allows an enterprise to implement a security policy that dictates the categories of Internet sites that are unauthorized.

VPN termination device—This device terminates IPSec tunnels for either site-to-site or remote access VPN connections. The device should provide additional services in order to offer the same network functionality as a classic WAN or dial-in connection.

Workstation or user terminal—A workstation or user terminal is any device on the network that is used directly by the end user, including PCs, IP phones, wireless devices, and so forth.

Figure 21 Legend

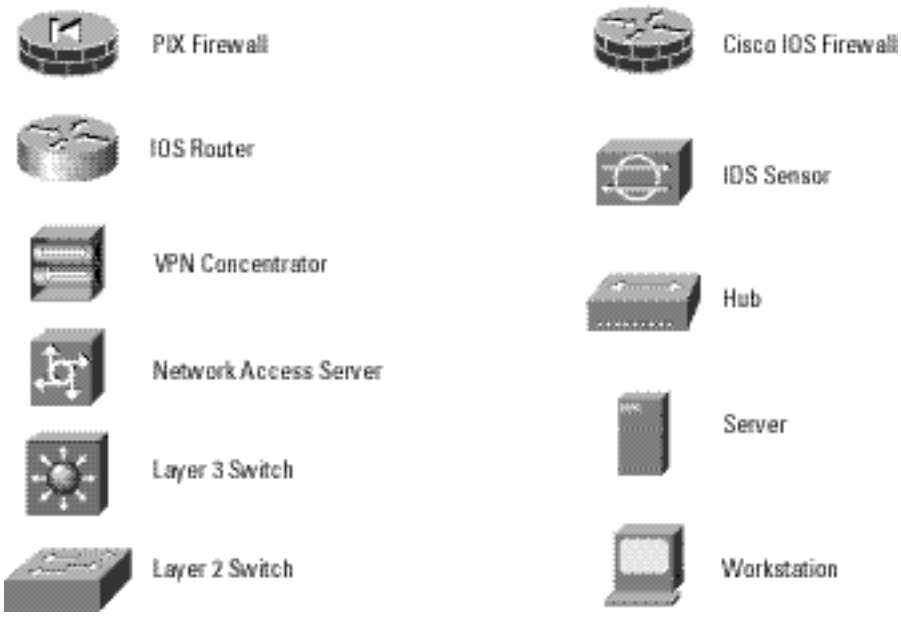

Diagram Legend

References

RFCs

RFC 2196 “Site Security Handbook”	http://www.ietf.org/rfc/rfc2196.txt
RFC 1918 “Address Allocation for Private Internets”	http://www.ietf.org/rfc/rfc1918.txt
RFC 2827 “Network Ingress Filtering: Defeating Denial of Service Attacks which Employ IP Source Address Spoofing”	http://www.ietf.org/rfc/rfc2827.txt

Miscellaneous References

“Improving Security on Cisco Routers”	http://www.cisco.com/warp/public/707/21.html
“VLAN Security Test Report”	http://www.sans.org/newlook/resources/IDFAQ/vlan.htm
“AntiSniff”	http://www.securitysoftwaretech.com/antisniff
“LC3”	http://www.atstake.com/research/lc3/index.html
“Denial of Service Attacks”	http://www.cert.org/tech_tips/denial_of_service.html
“Computer Emergency Response Team”	http://www.cert.org
“Security Focus (Bugtraq)”	http://www.securityfocus.com
“Avian Research (netcat)”	http://www.avian.org
“University of Illinois Security Policy”	http://www.aits.uillinois.edu/security/securestandards.html
“Design and Implementation of the Corporate Security Policy”	http://www.knowcisco.com/content/1578700434/ch06.shtml

Partner Product References

Entercept Host-Based IDS:	http://www.entercept.com
RSA SecureID OTP System:	http://www.rsasecurity.com/products/securid/
Baltimore MIMESweeper Email Filtering System:	http://www.mimesweeper.com
Websense URL Filtering:	http://www.websense.com/products/integrations/ciscopix.cfm
F-Secure SSH Client:	http://www.fsecure.com/products/ssh/
OpenSystems PrivateI Syslog Analysis Tool:	http://www.opensystems.com/products/index.asp
Zone Alarm Pro Personal Firewall:	http://www.zonelabs.com/products/index.html
General Cisco AVVID Security and VPN Solution Partners Information:	http://www.cisco.com/go/securitypartners

Acknowledgments

The authors would like to publicly thank all the individuals who contributed to the SAFE architecture and the writing of this document. Certainly, the successful completion of this architecture would not have been possible without the valuable input and review feedback from all of the Cisco employees both in corporate headquarters and in the field. In addition, many individuals contributed to the lab implementation and validation of the architecture. The core of this group included Rahimulah Rahimi, Jason Halpern, Mark Doering, Tom Hunter, Masamichi Kaneko, Alok Mittal, and Mike Steinkoenig. Thank you all for your special efforts.

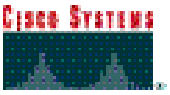

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe
11, Rue Camille Desmoulins
92782 Issy Les Moulineaux
Cedex 9
France
www.cisco.com
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems Australia, Pty., Ltd
Level 17, 99 Walker Street
North Sydney
NSW 2059 Australia
www.cisco.com
Tel: +61 2 8448 7100
Fax: +61 2 9957 4350

Cisco Systems has more than 190 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the Cisco.com Web site at www.cisco.com/go/offices.

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China • Colombia • Costa Rica • Croatia • Czech Republic • Denmark • Dubai, UAE Finland • France • Germany • Greece • Hong Kong • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

DA/0601