

VPN & Security



Cisco SAFE **Security Blueprint** for Secure E-Business

The principal goal of this paper is to provide best-practice information to interested parties on designing and implementing secure networks. Meant as an introduction to the SAFE concept, this document is heavily based on the SAFE white paper for Small, Midsize and Remote-User Networks. The original SAFE white paper was written to provide best-practice information on large enterprise network security designs and it is recommended that original versions of the SAFE white papers are obtained for an in-depth technical analysis.

This document takes the SAFE principles and sizes them appropriately for smaller networks, including branches of larger enterprises as well as stand-alone, small to midsize security deployments. It also includes information on remote-user networks such as teleworkers and mobile workers.

SAFE serves as a guide to network designers considering the security requirements of their network. SAFE takes a defence-in-depth approach to network security design. This type of design focuses on the expected threats and methods of mitigation, rather than on 'Put the firewall here, put the intrusion detection system there' approach. Such a 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 uses Cisco Systems and its partners' products. However, this document does not specifically refer to products by name. Components are referenced by functional purpose rather than by model name or number. During the validation of SAFE, real products were configured in the exact network implementation described 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.

Throughout this document the term 'hacker' denotes an individual who attempts to gain unauthorised 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 specialised security devices).
- Cost-effective deployment.
- Secure management and reporting.
- Authentication and authorisation 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 defence, or originate from inside the network, must be accurately detected and quickly contained to minimise 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 specialised 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 specialised 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 for 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 organisation, 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. Potentially they are 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 to 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.



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 set-up 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 set-up prevents packets tagged with the same VLAN as the trunk port from reaching another VLAN without crossing a Layer 3 device.
- Disable all unused ports on a switch. This set-up 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. 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. 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 organisation with a DS1 (1.5 Mbps) connection to the Internet that provides e-commerce services to its website 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 organisation'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 pre-specified 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 categorise. Only when limiting a TCP SYN attack on port 80 (Hyper Text Transfer Protocol (HTTP)) does an administrator run the risk of locking out legitimate users during an attack. Even then, it is better to lock out new legitimate users temporarily and retain routing and management connections than have the router over-run 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 categorise 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. For example, for inbound traffic on a router that is connected to the Internet, you employ RFC 1918 and 2827 filtering to prevent unauthorised 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, such as 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 that 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 perceived 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 set-up will block the user long enough to allow the administrator to decide what permanent action (if any) should be taken 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 Analyser (SPAN) or mirror port. Make sure this mirror port supports bi-directional 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 organisations), 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 set-up 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 tunnelling 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 set-up 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 synchronised 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 organisation's network. This headend may have VPN connections to other offices of the same organisation. 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 organisation, 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.

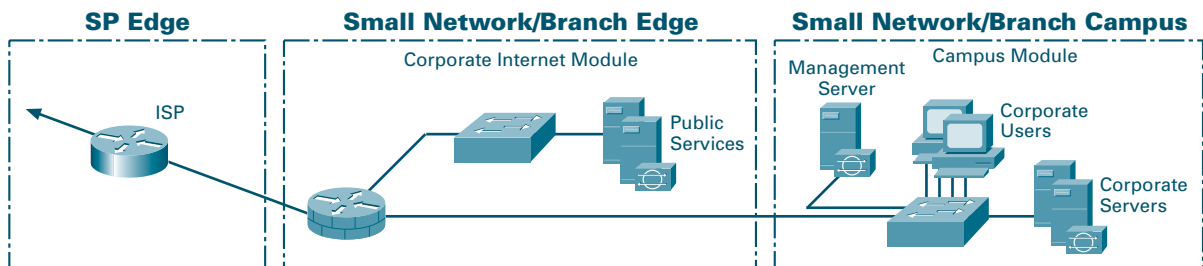
Second is the threat to the publicly addressable hosts that are connected to the Internet. These systems will likely be attacked with application layer vulnerabilities and DoS attacks.

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 organisation.
- **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.

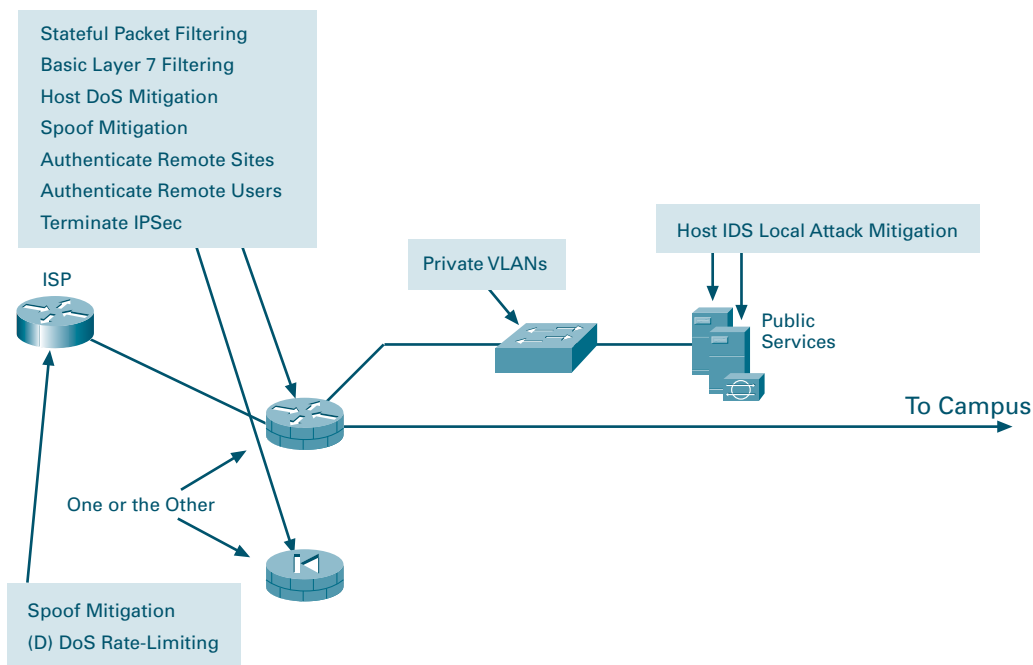


Threat Mitigation

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

- **Unauthorised 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 set-up 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 set-up 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 set-up 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 non-essential traffic that exceeds pre-specified 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 unauthorised 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 set-up 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 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.



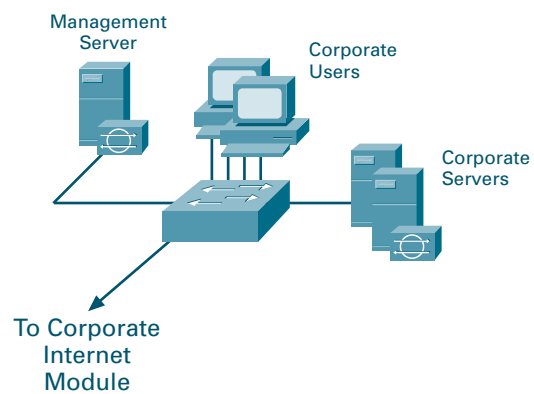
Campus Module

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 authorised users on the network.
- **Management host:** Provides HIDS, syslog, TACACS+/Remote Access Dial-In User Service (RADIUS) and general configuration management.

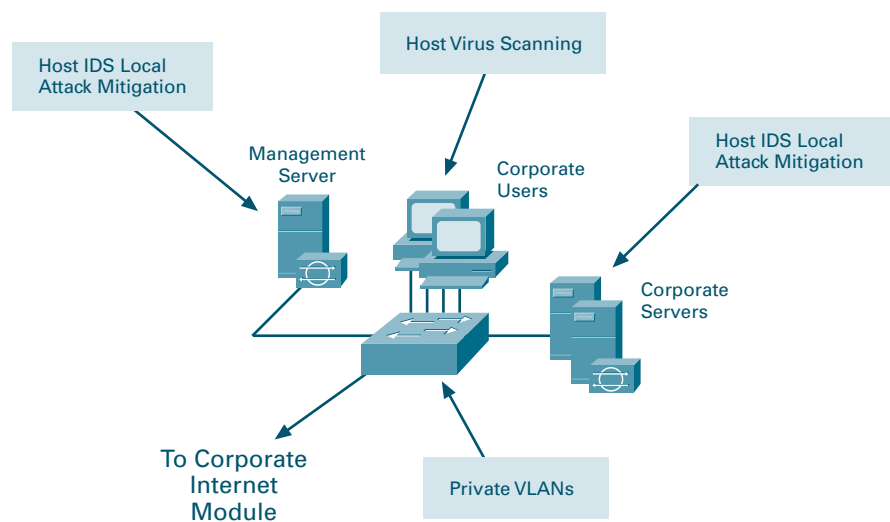
Figure 4: Detailed Model of Small Network Campus Module.



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.
- **Unauthorised 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 set-up will allow management traffic to flow only in the specific direction deemed necessary by the administrators. If the level of trust within the organisation is high, HIDS can potentially be eliminated, though this is not recommended.

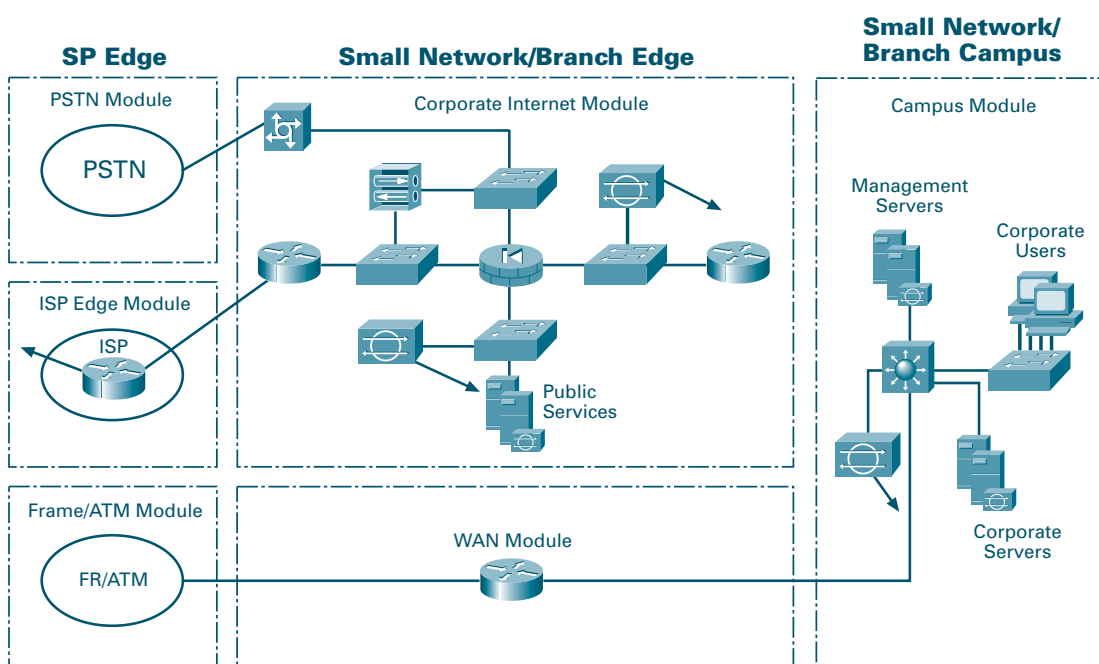
Branch versus Stand-alone Considerations

When configured in a branch role, remote access VPN functionality is not required as this is generally provided by the corporate headquarters. Management hosts are typically located at the central site, a set-up that requires 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 and 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.



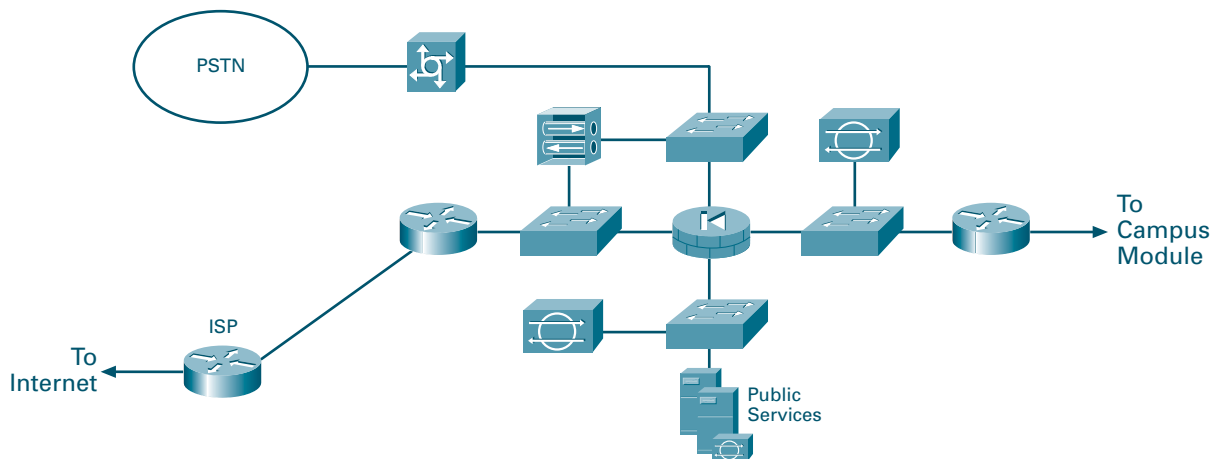
Corporate Internet Module

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 analogue 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 organisation.
- **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:

- **Unauthorised 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 set-up 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.
- **Unauthorised access:** Firewall services after packet decryption prevent traffic on unauthorised ports.
- **Man-in-the-middle attacks:** These attacks are mitigated through encrypted remote traffic.
- **Packet sniffers:** A switched infrastructure limits the effectiveness of sniffing.



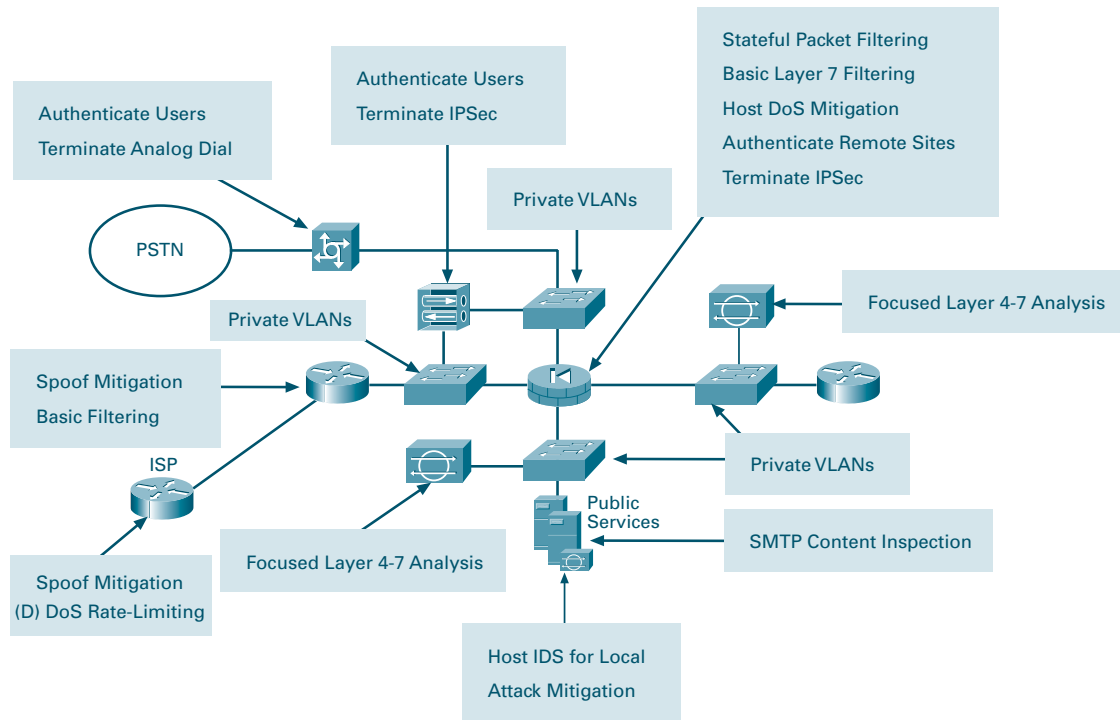


Figure 8: Medium Network Attack Mitigation Roles for Corporate Internet Module.

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 non-essential traffic that exceeds pre-specified 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. 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 addressable 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 unauthorised 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 set-up 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 was compromised, a hacker may 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 incidents such as the SMTP attack 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 tunnelling, 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 set-up 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, authorisation, 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 than 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 set-up mitigates against any potential misconfiguration on a switch that could compromise security. Additionally, each of the switches runs the private VLAN feature, a set-up 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 utilised 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 defence-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. It is likely that the amount of alarms generated on this segment will probably be large and 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 organisation 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 set-up 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

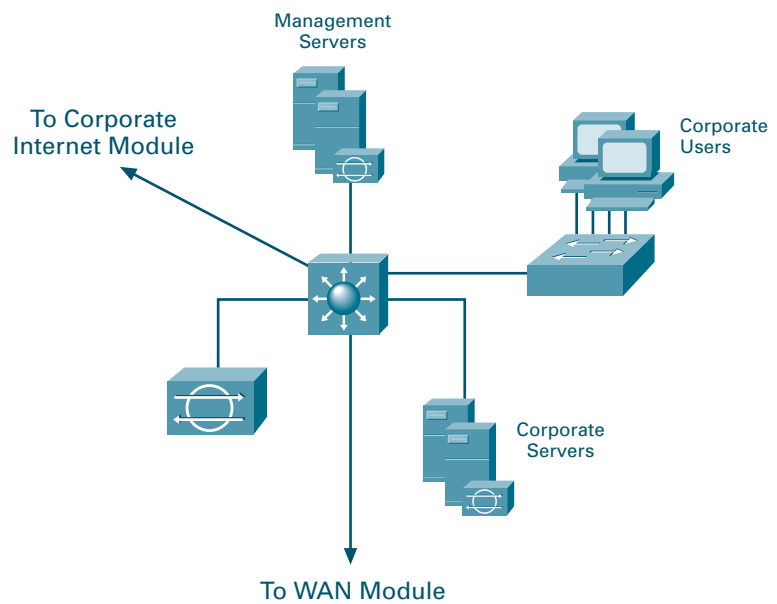
The campus module contains end-user workstations, corporate intranet servers, management servers and the associated Layer 2 and Layer 3 infrastructure required to support the devices. All the campus modules from SAFE Enterprise have been combined into a single module. This set-up more accurately reflects 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 authorised 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:** Authorises 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.

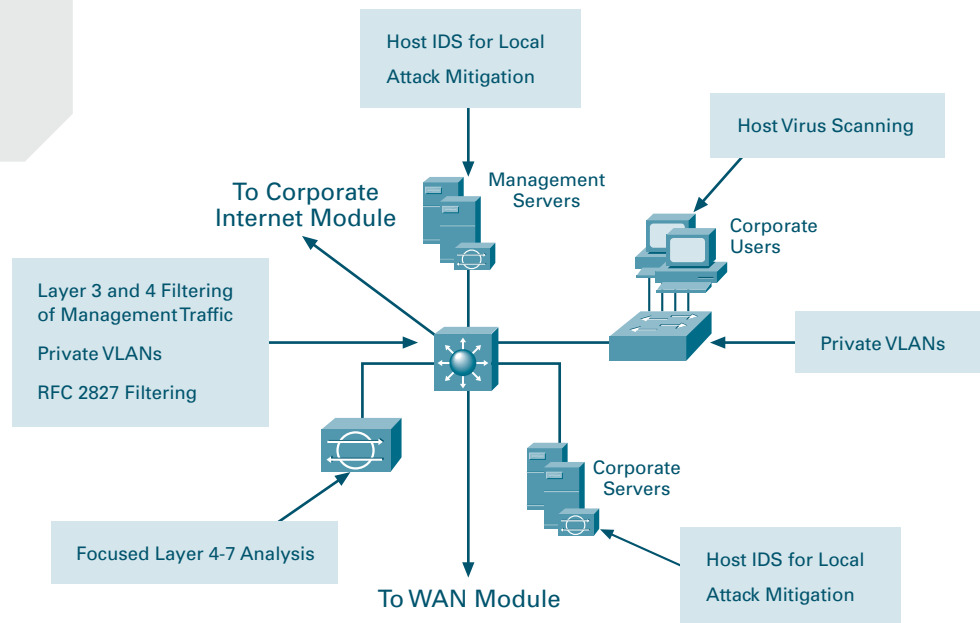


Threats Mitigated

- **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.
- **Unauthorised 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 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 defence 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 set-up 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 utilised 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 defence 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 was 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 unauthorised people could 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 fitted into the core switch. This set-up 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.



Figure 12: Attack Mitigation Roles for WAN Module.

Threats Mitigated

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



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

Organisations 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 organisation 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 infrastructure, 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 centralised IT resources). If a private WAN link is chosen for the intersite connectivity, management traffic can flow easily across the WAN module and into the devices requiring 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 set-up 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 an ISP-provided broadband access device.

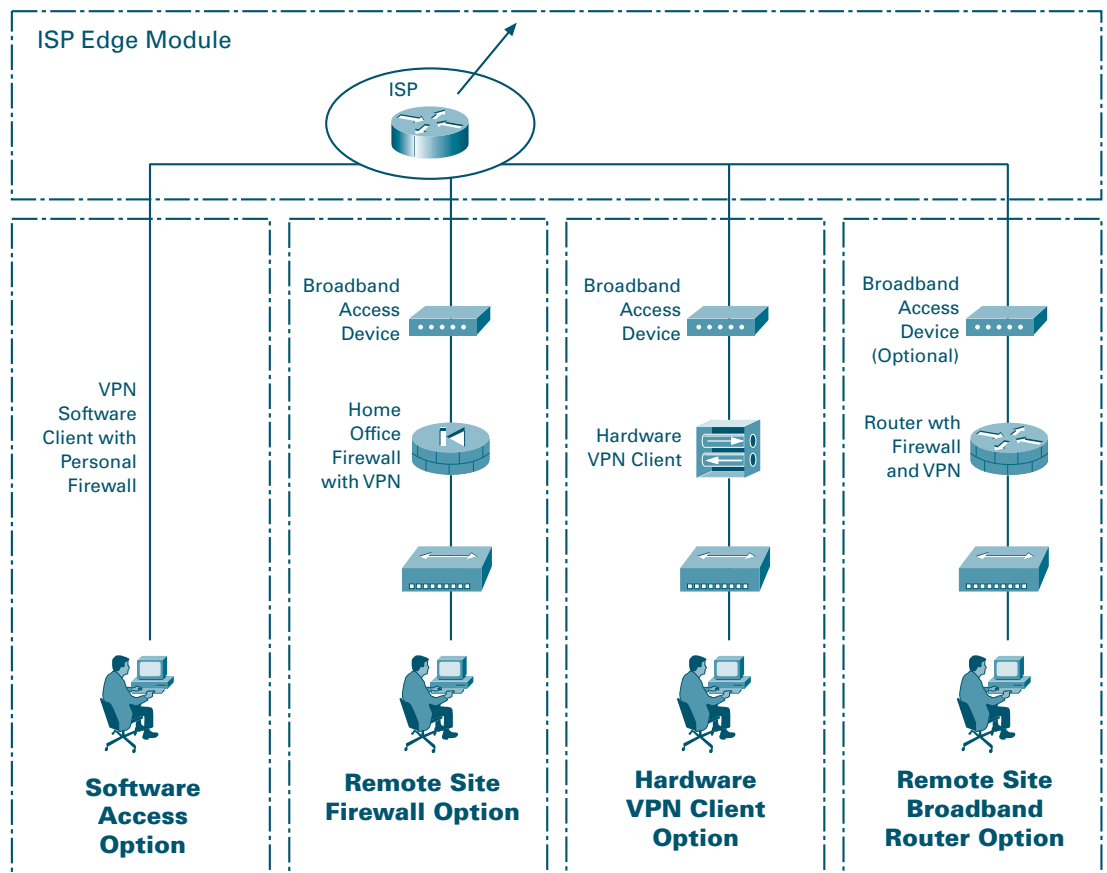
Each of these designs is discussed further in the design guidelines section below. All discussions assume that the connectivity is through the Internet. If private WAN connectivity (ISDN, private DSL, etc.) is used instead, encryption of the traffic may not be required. Keep in mind that with any remote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

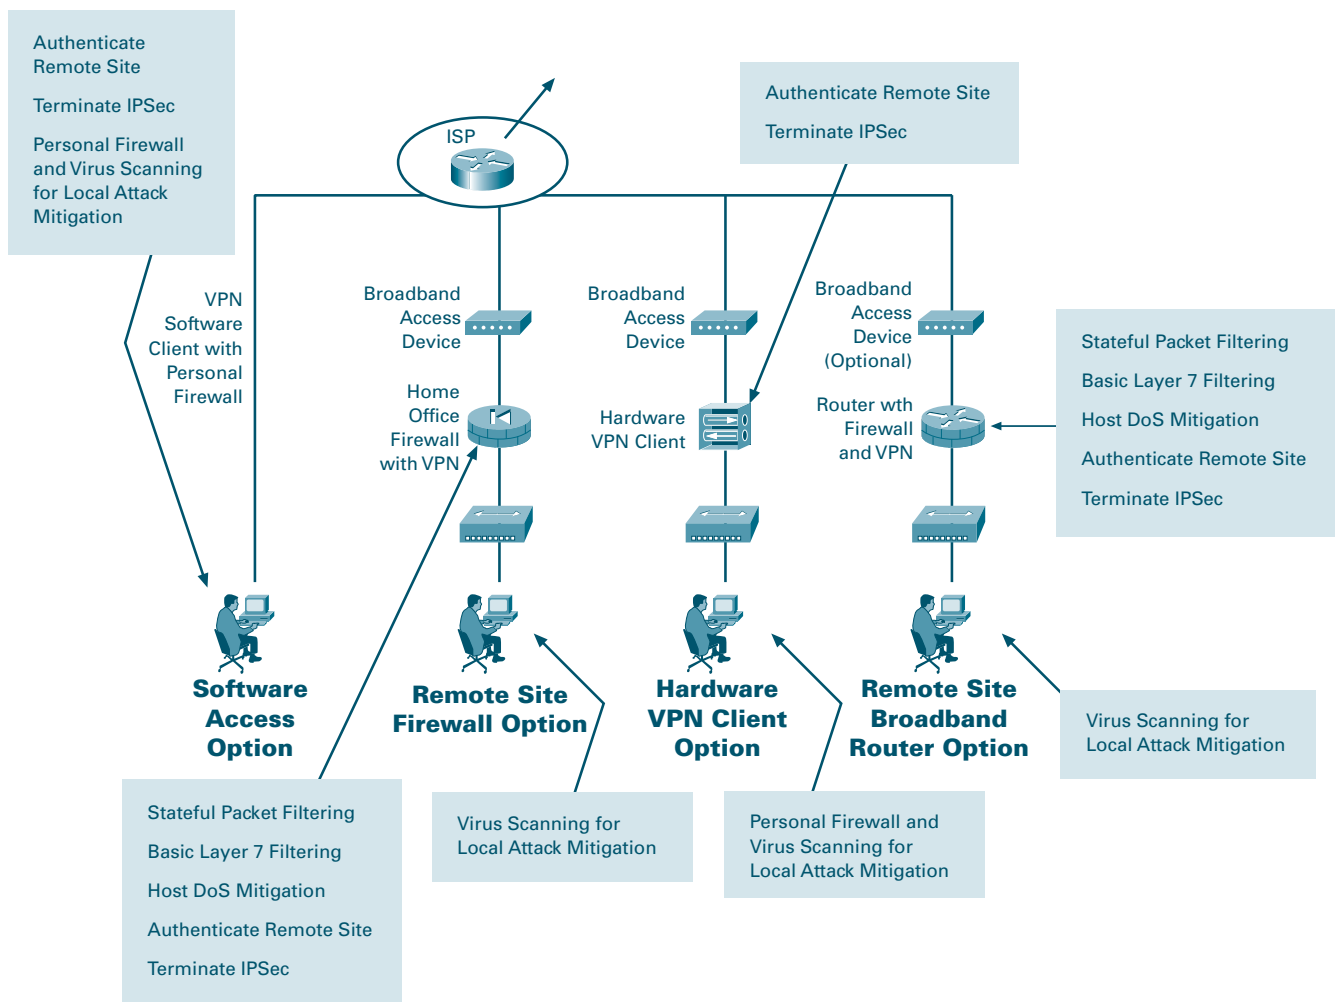
Figure 13: Detailed Model of Remote-User Configuration.



Threats Mitigated

- **Unauthorised 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 authorisation 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 tunnelling can also be enabled or disabled via the central site. For the SAFE design, split tunnelling 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 unauthorised 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 set-up 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 authorisation 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 set-up 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 their firewall configuration and thereby compromising 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 utilise 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 set-up requires use of a personal firewall on the individual hosts particularly when split tunnelling is enabled. Without the personal firewall the security of the individual hosts behind the VPN device is dependent upon the attacker being unable to circumvent Network Address Translation (NAT). This is because when split tunnelling 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 tunnelling disabled, all access to the Internet must be through the corporate headquarters. This set-up 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 authorisation 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 set-up 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 encapsulation options. Additionally, 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.

Conclusions

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. 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 set-up 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.



Architecture Taxonomy and Legend

Application server: The application server provides application services directly or indirectly for enterprise end users. Services can include workflow, 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 non-disruptive 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 (multi-packet) 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 authorised to be forwarded as is, altered and forwarded, or dropped.



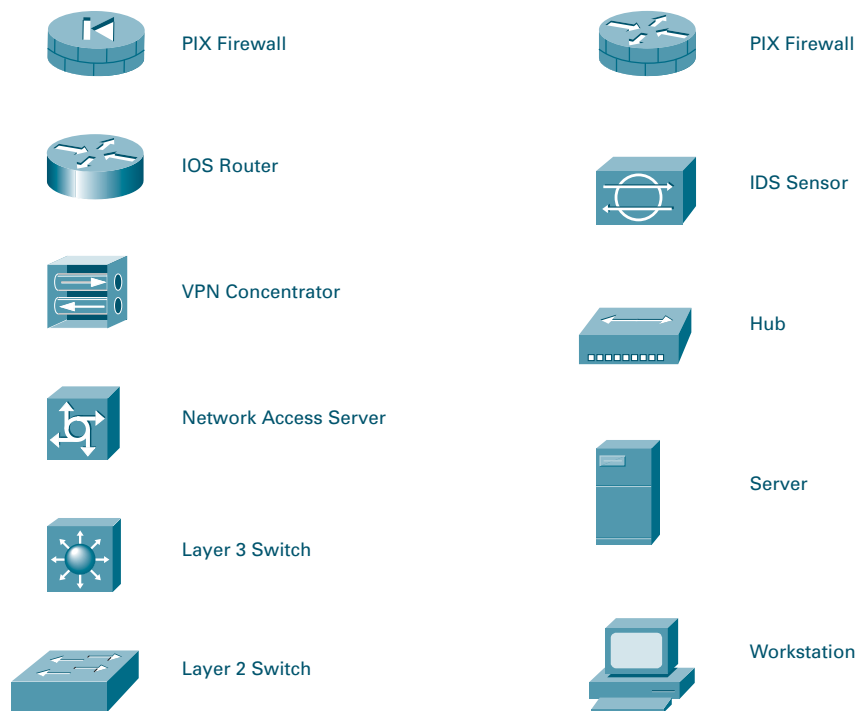
URL filtering server: This server application typically runs on a stand-alone server that monitors URL requests forwarded to it by a network device and informs the network device whether the request should be forwarded on to the Internet.

This set-up allows an enterprise to implement a security policy that dictates the categories of Internet sites that are unauthorised.

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.

Figure 15: Legend.

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.





Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA

Tel: +1 408 526 4000 800
553 NETS (6387)
Fax: +1 408 526 4100

European Headquarters
Cisco Systems Europe s.a.r.l.
11, Rue Camille Desmoulins
92782 Issy-les-Moulineaux
Cedex 9
France

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

Tel: +1 408 526-7660
Fax: +1 408 527-0883

Asia Pacific Headquarters
Cisco Systems Australia, Pty., Ltd
Level 9, 80 Pacific Highway
P.O. Box 469
North Sydney
NSW 2060 Australia

Tel: +61 2 8448 7100
Fax: +61 2 9957 4350

Cisco Systems has more than 200 offices in the following countries and regions. 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 PRC • Colombia • Costa Rica • Croatia • Czech Republic • Denmark • Dubai, UAE
Finland • France • Germany • Greece • Hong Kong SAR • 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

Copyright © 2002, Cisco Systems, Inc. All rights reserved. Cisco Systems, the Cisco Systems logo and Empowering the Internet Generation are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the US and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners.