

SAFE VPN IPSec Virtual Private Networks in Depth

Table of Contents

Authors	2
Abstract.	2
Audience.	3
Caveats.	3
Architecture Overview	4
Design Fundamentals.	4
SAFE VPN Axioms.	4
Branch versus Headend Considerations	18
Remote-User Design	19
Small VPN Design	23
Corporate Internet Module	24
Branch versus Standalone Considerations	27
Medium VPN Design	27
Corporate Internet Module	28
Branch versus Standalone Considerations	31
Large VPN Design	31
VPN and Remote-Access Module	32
Extranet Module	38
Management Module	41
Distribution-Hub Module	44
Migration Strategies	47
Appendix A: Validation Lab	48
Appendix B: VPN Primer	88
Appendix C: Architecture Taxonomy	93



Authors

Jason Halpern is the primary author of this white paper and the lead architect for the reference implementation at Cisco headquarters in San Jose, CA USA. Sean Convery and Roland Saville provided significant contributions to this paper. All three are network architects focusing on VPN and security issues.

Abstract

The principal goal of this paper is to provide best-practice information to interested parties for designing and implementing Enterprise IP Security (IPSec) virtual private networks (VPNs). The SAFE white paper for large enterprises and the SAFE white paper for small, midsize and remote user networks are both available at the SAFE Web site:

<http://www.cisco.com/go/safe>

These documents were written to provide best-practice information on network security designs. They include some element of VPN configurations and design guidance. This document continues the discussion, examining specific design considerations and best-practice recommendations for IPSec VPNs in networks today. Although you can read this document without having read either of the two security design documents, it is recommended that you read the document most appropriate to your network size before you read this document. For example, a business with a large network might look at the enterprise SAFE white paper before reading this document. This exercise will frame the VPN conversation within the context of overall security design. SAFE represents a system-based approach to security and VPN design. This type of approach focuses on overall design goals and translates those goals into specific configurations and topologies. SAFE is based on Cisco products and those of its partners.

This document begins with an overview of the architecture, and then details the specific designs under consideration. The following designs are covered in detail:

- Remote-user VPN designs
- Small-network VPN design
- Medium-network VPN design
- Large-network VPN design (with extranet connectivity)
- Distributed large-network VPN design

Each design may have multiple modules that address different aspects of VPN technology. The concept of modules is addressed in the SAFE security white papers. Topics that are covered in each design or module (where appropriate) include:

- Overall design best practices
- High availability
- Scalability
- Performance
- Identity
- Secure management
- Network Address Translation (NAT)
- Security
- Routing
- Extranet considerations



Following the discussion of the specific designs, Appendix A details the validation lab for SAFE VPN and includes configuration snapshots. Appendix B is a primer on VPNs. Readers who are unfamiliar with basic VPN 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.

Audience

Though this document is technical in nature, 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 VPN 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. Because this document covers a wide range of VPN deployments, it may be helpful to read the introductory sections of the paper first and then skip right to the type of VPN you are interested in deploying.

Caveats

This document presumes that you already have a security policy in place. Cisco Systems does not recommend deploying VPNs or any security technology without an associated policy. It presumes you are aware of what data is sensitive in your network so that it can be properly protected when transported through the Internet. Although the topic of network security is mentioned in this document, it is not described in detail. Security within this document is always mentioned as it pertains to VPN technology. Readers interested in more information on network security should look to the SAFE security documents for detailed design guidance:

<http://www.cisco.com/go/safe>

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 at Fort Knox. 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 and the SAFE security documents, 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.

Though this document contains a large amount of detail on many aspects of VPN technologies, it is not exhaustive in its discussion. In particular, several technologies that relate to VPNs are not covered. First, *certificate-authority* (CA) deployment is not discussed. Identity strategies are addressed, including X509 V3 digital certificates, as well as other identity technologies. Best practices for deployment of CAs in an enterprise are not discussed. CAs and their associated deployment issues require a level of focus that this document cannot provide and still adequately address all the other relevant areas of identity and VPN. Also, because most networks have yet to deploy fully functional CA environments, it is important to discuss how to securely deploy networks without them. Second, the VPN designs in this paper assume the VPN gear exists on the customer premises and is managed by the customer. Though these topologies may not change significantly if the VPN is managed by a service provider, the management and provisioning of that type of network would be very different. As such, this document can be used to evaluate the VPN offerings of a service provider, but should not be used as best-practice recommendations for outsourced VPNs. Third, a detailed analysis of the issues surrounding maintaining QoS in VPNs is not addressed in this document. QoS is an essential component in delivering differentiated service levels and ensuring reliable throughput of mission-critical data across the VPN. This paper addresses many other essential design considerations. However, analysis of these issues alone exhausted the allocated resources for the first release of this paper.



Fourth, this document assumes that when VPN is chosen for connectivity between two sites, it is the only method for these sites to communicate. Split-tunneling may be used for nonsecure site-to-site communications. Using IPSec to back up private WAN links, or in *dial-on-demand routing* (DDR) environments, is not addressed. The topologies provided could be utilized in a backup role, although routing and other considerations in the surrounding network are not addressed. Fifth, 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.

Finally, dynamic tunnel endpoint discovery mechanisms are neither analyzed nor discussed in this paper. Dynamic peer discovery mechanisms best function in partially or fully meshed networks where spoke-to-spoke connectivity is required on an infrequent basis. Since these types of networks have not been deployed in large-scale deployments and the issues are not yet entirely known, this document discusses networks that rely more heavily on the static hub-and-spoke network design than the dynamic partially or fully meshed network design.

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 VPN 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.

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

First and foremost, SAFE VPN needs to provide private, ubiquitous communications to the locations and users that require it. It must do this in a secure manner while maintaining as many of the characteristics of traditional private WAN connections as possible. It must integrate with existing network designs based on the SAFE security architecture.

SAFE VPN Axioms

The following axioms represent overarching design considerations that affect nearly every design within SAFE VPN. They are included at the beginning of this document to limit the amount of redundancy in the rest of the paper. SAFE VPN assumes conformance with the security axioms in the original SAFE white paper. However, in comparison to prior SAFE security papers, this document relies more heavily on the axiom section of the document. Although VPN design differs greatly with the size of enterprises, the underlying best practices remain virtually the same. Therefore,



the design discussions are somewhat similar. In the axioms, it is assumed that the users and sites are members of your enterprise and in your domain of control; a separate design discusses the security implications of extranets. After reading the VPN axioms it is the authors' intent that you would come to the same design conclusions that the authors did in the document.

Identity and IPSec Access Control

In site-to-site and remote-access VPNs today, it is important that devices are identified in a secure and manageable way. In remote-access VPNs, user authentication as well as device authentication occurs. When the remote device is authenticated, some level of access control needs to be in place to permit only the traffic over the tunnel that should be there.

Device authentication uses either a preshared key or digital certificate to provide the identity of a device. There are three types of preshared keys: wildcard, group, and unique. Unique preshared keys are tied to a specific IP address. Group preshared keys are tied to a group name identity; these are applicable only to remote access today. Wildcard preshared keys are not associated with any unique information to determine a peer's identity. Any device that has the wildcard key will successfully authenticate. Therefore, wildcard preshared keys should not be used for site-to-site device authentication. The authors feel doing so is asking for trouble. When using wildcard preshared keys, every device in the network uses the same key. If a single device in your network is compromised and the wildcard preshared key has been determined, a hacker can establish a tunnel with any device in the network. Compromised pre-shared keys, unique or wildcard, are also susceptible to man-in-the-middle attacks during tunnel establishment. Dynamic cryptographic maps facilitate this hacking by accepting *Internet-Key-Exchange* (IKE) requests from any IP address. At an absolute minimum, you should consider using a unique preshared key between two devices. However, obviously this setup would not scale in large networks. Depending on how strong the preshared keys are and how often they are changed, they may not provide strong device authentication.

Digital certificates scale better than unique preshared keys because they allow any device to authenticate to any other device but do not have the security properties of wildcard keys. Digital certificates are not tied to IP addresses but to unique, signed information on the device that is validated by the enterprise's CA. If a hacker compromises or steals a device with a digital certificate, the administrator will revoke the digital certificate and notify all other devices by publishing a new *certificate revocation list* (CRL). The CRL contains a CA-signed list of revoked certificates. When a device receives a request for tunnel establishment and uses a digital certificate for proof of identity, the device checks the peer certificate against the CRL. Devices generating digital certificates or validating received certificates during tunnel authentication and establishment must have configured the correct time of day (preferably *Coordinated Universal Time* [UTC]). Time is also used to determine when the CRL expires so that a new one can be retrieved. Although checking CRLs can be configured as optional, it should always be enabled on remote and headend devices when digital certificates are deployed. A simple certificate revoke command on the CA server carries out the revocation. In comparison, preshared keys are revoked by removing them from each device.

Digital certificates also provide more key entropy (more bits for seeding functions), public/private key pair aging, and nonrepudiation. Digital certificates do, however, require additional administrative resources to deploy and manage, given their feature complexity. Using a third-party-managed CA versus an enterprise-managed CA may help to facilitate deploying an extranet VPN. Consider using digital certificates if the size of the VPN grows beyond 20 devices—or even sooner if there are requirements for strong device authentication. Today, the administrator burden for deploying digital certificates to remote-access clients is significant.



“Typically wildcard or group preshared keys are used with remote-access clients for device authentication since the remote IP address is dynamic. As mentioned previously, this form of device-authentication does not provide strong device authentication. To compensate, an additional layer of user-authentication is provided before granting access to the network. This is known as extended authentication (XAUTH). In addition, by using a strong user-authentication scheme such as one-time passwords (OTPs), the lack of strong device authentication is somewhat reduced.

However, it should be noted that if the group key is compromised, further attack escalation via a man-in-the-middle attack may provide access to data within the IPSec tunnel. In fact, if the key is discovered a hacker could then masquerade as the VPN termination device via various means (DNS poisoning, L2 subversion, and so on) and man-in-the-middle the entire secure establishment of the IPSec tunnel.

For this reason administrators should not only strongly consider the security trade-offs between digital certificates and preshared keys, but also the methods used to store the actual keying material of both options in modern-day operating systems. Many hardware-based solutions exist today for protecting the keying material and these are generally considered more secure than software implementations.”

After the device and user authentications (if applicable) are complete, IPSec access control occurs. Normally the networks, hosts, and ports that are allowed to traverse the tunnels are defined in the *Security Policy Database* (SPD), as defined by the IPSec standard. This database is populated by the use of *access control lists* (ACLs). These ACLs are sometimes referred to as “crypto ACLs” or “network rules.” You might consider using the cryptographic ACLs for rudimentary network security access control, but Cisco does not recommend this scenario because it complicates the configuration significantly. Rather, you should use inbound ACLs on the VPN devices for site-to-site traffic. For remote-access traffic filtering, access-control occurs dynamically by loading the per-user granular authorization information when the user successfully authenticates via XAUTH.

IPSec

IPSec provides numerous security features. The following have configurable values for the administrator to define their behavior:

- Data encryption
- Device authentication and credential
- Data integrity
- Address hiding
- *Security-association* (SA) key aging.

The IPSec standard requires the use of either data integrity or data encryption; using both is optional. Cisco highly recommends using both encryption and integrity. Because single *Data Encryption Standard* (DES) was compromised in the last competition in 1999 in about 22 hours and 15 minutes with US\$50,000 worth of equipment, Cisco recommends that you do not use single DES for data encryption. Instead, Cisco recommends the use of *Triple DES* (3DES). Data integrity comes in two types: 128-bit strength *Message Digest 5* (MD5)-HMAC or 160-bit strength *secure hash algorithm* (SHA)-HMAC. Because the bit strength of SHA is greater, it is considered more secure. Cisco recommends the use of SHA because the increased security outweighs the slight processor increase in overhead (in fact, SHA is sometimes faster than MD5 in certain hardware implementations). Both IPSec phases offer the ability to change the lifetime of the SA. You may consider changing the lifetime from the default when the sensitivity of the tunneled data mandates replacing the encryption keys and reauthenticating each device on a more aggressive basis.



The use of strong encryption algorithm in non-US countries is sometimes regulated by local import and usage laws. These strong encryption algorithms cannot be exported to some countries or some customers. For more information, please see:

<http://www.cisco.com/wwl/export/crypto>

Changing these values increases the level of security; at the same time, however, increases the processor overhead. The default behavior for *quick-mode* (QM) SA rekeying is to base the new key in part on the old key to save processing resources. *Perfect forward secrecy* (PFS) generates a new key based on new seed material altogether by carrying out a *Diffie-Hellman* (DH) exponentiation every time a new QM SA needs new key generation. Again, this option increases the level of security but at the same time increases processor overhead. Cisco does not recommend changing the SA lifetimes or enabling PFS unless the sensitivity of the data mandates it. If you choose to change these values, make sure you include this variable when determining the network design. The strength of the Diffie-Hellman exponentiation is configurable; Groups 1 (768 bits), 2 (1024 bits), and 5 (1536 bits) are supported. Group 2 is recommended. Throughout the SAFE VPN architecture, at a minimum the following modes were used: IKE-3DES, IKE-SHA-HMAC, IKE-DH Group 2, IKE-preshared key, IPSec-3DES, IPSec-SHA-HMAC, IPSec-no PFS, and IPSec-tunnel mode ESP.

IP Addressing

Proper IP addressing is critical for a successful VPN as in any large IP network. In order to maintain scalability, performance, and manageability, it is highly recommended that remote sites use a subnet of the major network to allow for summarization. This way, the cryptographic ACLs will contain a single line for every local network, possibly a single entry if the local networks are themselves summarizable. For example, a remote-site network 10.1.1.0/24 is summarizable into the major network 10.0.0.0/8. If any host on the 10.1.1.0/24 subnet needs to connect to any other subnet in the 10.0.0.0 network via the headend, a single ACL entry will suffice. If you cannot summarize the remote networks in a major network, an ACL entry is required at the remote site for every local network-to-remote network. Increasing ACL entries slows performance, complicates troubleshooting, and hinders the scalability by requiring ACL changes at remote sites constantly to keep up with new networks available at the headend. Each ACL entry will build a separate tunnel (two IPSec SAs). Proper subnetting also allows for simplified router headend configuration to enable spoke-to-spoke intercommunication and requires fewer tunnels on all devices to classify traffic flows. IP addressing also affects many facets of VPNs including remote management and connection of overlapping networks.

Multiprotocol Tunneling

IPSec as a standard supports IP unicast traffic only. For multiprotocol or IP multicast tunneling, you must use another tunneling protocol. Because of its *Point-to-Point Protocol* (PPP) ties, *Layer 2 Tunneling Protocol* (L2TP) is best suited for remote-access VPNs. *Generic routing encapsulation* (GRE) is best suited for site-to-site VPNs. GRE is typically used to tunnel multicast packets such as routing protocols. Neither of these tunneling protocols supports data encryption or packet integrity. GRE and L2TP application with IPSec is purely for non-IP-unicast support—not for additional security. L2TP also simplifies remote-access client address assignment and user authentication. Packets are first encapsulated by the secondary tunneling protocol and then encapsulated by IPSec. GRE and L2TP also allow for a single set of IPSec SAs to tunnel traffic from one site to another. Typically with IPSec you need a unique set of



IPSec SAs to provide tunneling capability for each local network to each remote network. GRE encapsulates all traffic, regardless of its source and destination into a single tunnel. In summary, use point-to-point GRE or L2TP when you need support for tunneling packets other than IP unicast type.

Network Address Translation

NAT can occur before and after IPSec. It is important to realize when NAT will occur since in some cases NAT may interfere with IPSec by blocking tunnel establishment or traffic flow through the tunnel. It is a best practice to avoid the application of NAT to VPN traffic unless it is necessary to provide access.

NAT After IPSec

You may consider applying NAT after IPSec encryption for address hiding. However, this provides no benefit because the actual IP addresses of the devices utilizing the tunnel for transport are hidden via the encapsulation. Only the public IP addresses of the IPSec peers are visible, and address hiding of these addresses provides no real additional security. NAT application after IPSec encapsulation will occur in cases where IP address conservation is taking place. This is, in fact, commonplace in hotels, cable/*digital subscriber line* (DSL) residential deployments, and enterprise networks. In these cases, depending on the type of NAT used, NAT's application may interfere with the IPSec tunnel establishment.

When IPSec uses *Authentication-Header* (AH) mode for packet integrity, if one-to-one address translation occurs it will invalidate the signature checksum. Because the signature checksum is partially derived based on the AH packet's IP header contents, when NAT changes the IP header, the signature checksum is invalidated. In this case, the packet will appear to have been modified in transit and will promptly be discarded when received by the remote peer. However, when IPSec uses ESP, the devices will be able to successfully send packets over the VPN, even when one-to-one address translation occurs after encapsulation. This scenario is possible because ESP does not use the IP header contents to validate the integrity of the packets in comparison to AH. For this reason and the fact that ESP also supports encryption and authentication, AH is rarely used. In cases where many-to-one address translation occurs (aka port address translation), the IP address and source IKE port, normally *User Datagram Protocol* (UDP) port 500, will change. Some VPN devices do not support IKE requests sourced on ports other than UDP 500. Some devices performing many-to-one NAT do not handle ESP or AH correctly and may drop or even corrupt packets. Remember that ESP and AH are higher-layer protocols on top of IP that do not use ports.

Because many-to-one address translation is commonplace in many environments where remote-access clients are deployed, a special mechanism called NAT transparency exists to overcome these NAT issues. NAT transparency reencapsulates the IKE and ESP packets into another transport layer protocol with ports, such as UDP or TCP, which address-translating devices know how to translate correctly. This mechanism also allows the client to bypass access control in the network that allows TCP or UDP but blocks encrypted traffic. Note that this feature does not affect the security of the transport in any way. NAT transparency takes packets already secured by IPSec and then encapsulates them again in TCP or UDP. In summary, use ESP tunnel mode and avoid NAT whenever possible. For remote access specifically, use the NAT transparency mode when PAT is occurring.

NAT Before IPSec

When two networks are connected via IPSec if any of the address ranges overlap, the tunnel will not establish. This occurs because it is not possible for the VPN termination devices to determine the site to which to forward the packets. Utilizing NAT before IPSec overcomes this restriction by translating one set of the overlapping networks into a unique network address range that will not interfere with the IPSec tunnel establishment. This is the only scenario where the application of NAT is recommended. Be aware however that some protocols embed IP addresses



in packet data segments. In general, when address translation occurs, make sure that a protocol-aware device carries out the address translation, not only in the IP header but also in the data segment of the packet. If the packet was not correctly address translated before it entered the tunnel due to embedded IP addresses, when the packet exits the tunnel the remote application will not receive the correct IP address embedded in the data segment. In this case, it is likely that the application will fail to function properly. Many remote-access VPN clients today support the ability to use a virtual address assigned by the headend terminating VPN device. Devices at the remote site may connect to the remote access client using this virtual address. This is actually carried out by one-to-one address translating all packets traversing the tunnel at the client. If the VPN client does not address translate packets correctly or a new application arrives that is not yet supported, the application may fail to function.

In summary, use address ranges for your sites and remote access VPN client virtual address pools that do not overlap with the addresses of other devices you will connect via IPSec. If this is not possible, use NAT only in this scenario to allow for connectivity. Don't address hide the public peer addresses of the VPN devices as this provides no real security value-add and may cause connectivity problems. When you believe that NAT is involved when a remote-access client is not able to successfully establish a tunnel or send packets over an established tunnel, consider enabling the NAT transparency mode. Finally, be aware that the NAT transparency mode will not resolve the connection problems associated with client applications that are not NAT friendly.

Single Purpose versus Multipurpose Devices

At many points in the network design process, you need to choose between using integrated functionality in a networking or security device versus using the specialized functions of a VPN appliance. The integrated functionality is often attractive because you can implement it on existing equipment, because it is cost effective, or because the features can interoperate with the rest of the devices to provide a better functional solution. Dedicated VPN appliances are often used when the depth of functionality required is very advanced or when performance needs require using specialized hardware.

When deciding which option to select, weigh your decision based on the capacity and functionality available on the appliance, versus the functionality advantage of the integrated device. For example, sometimes you can choose an integrated higher-capacity Cisco IOS[®] Router with IPSec encryption software as opposed to a smaller Cisco IOS Router and an associated VPN device. Throughout this architecture, both types of systems are used. Because IPSec is such a demanding function, as the size of the designs increases, so does the likelihood that a VPN appliance was chosen instead of an integrated router or firewall. Note that the notion of a VPN appliance is a difficult one.

Many VPN appliances today offer good performance and VPN management options while at the same time provide limited routing, firewall, or QoS functionality that might be typically associated with an integrated device. If all this advanced functionality were enabled, the appliance would begin to act more and more like an integrated device from a performance and deployment options standpoint. Likewise, a VPN router that supports the full range of VPN functions in addition to a full complement of routing and security features could be configured in a VPN-only role where its characteristics would seem much more like an appliance.

Intrusion Detection, Network Access Control, Trust, and VPNs

When considering the deployment of VPN technology, remember that by doing so, you are extending the security perimeter of your network to include some areas that are not typically considered high security. These include:

- Employee homes
- Airports
- Hotels
- Internet cafés



One of the first questions an organization needs to answer is what is its level of trust for the VPN technology itself and the surrounding applications and hardware that will utilize it. A good way to arrive at a conclusion is by answering the following question: Do you as an organization want to trust an individual or remote site coming in over a VPN just as much as you trust local employees and sites connected via a private WAN link? If you answered that question “yes,” then you should deploy VPN technology in much the same way as you deploy private WAN links and modem pools today. However, it is the position of Cisco and most of its customers that VPN links should be trusted a little bit less. Therefore, IPSec VPNs are often deployed with layers of access control and intrusion detection around them. Even though IPSec with 3DES is very secure, the human potential to store keys insecurely and misconfigure devices creates enough uncertainty to warrant additional security considerations; not to mention stolen laptops or trojans. This document is written primarily for this latter group. If you fall into the former, much of the information here is valuable, though you may find the designs themselves to be too security focused.

Network Intrusion Detection System

Network intrusion detection system (NIDS) is a technology that can be used to reduce the risk associated with extending the security perimeter. NIDS carries out two primary functions in VPN designs. First, NIDS can be used to analyze traffic coming from, or destined to, the VPN device before encryption. Here NIDS will detect attacks coming through the VPN from remote sites or remote users. Because we know the origins of this traffic, and the chances that it is spoofed are low, any attack can be met with a strong response from the NIDS. This response can include shunning or TCP resets as appropriate. NIDS is critical in most VPN environments because most VPN security policies dictate that Layer 3 and Layer 4 access to the network over VPN should be fairly ubiquitous. This setup increases the reliance on NIDS to catch and stop most of the attacks from remote sites or users. Second, NIDS can be used after encryption to validate that only encrypted traffic is sent and received by VPN devices. By tuning a NIDS to alarm on any non-VPN packet, you can validate that only encrypted packets are flowing over the network. This setup guards against any misconfiguration of the VPN devices that could inadvertently allow unencrypted traffic through the device. This functionality is called out in greater detail in the large network VPN design.

Network Access Control

In addition to NIDS, access control, generally through the use of a firewall, should be performed before and after the VPN device. When done on the interior of a VPN device as traffic heads toward the campus, the access control can ensure only that the proper address ranges and protocols are allowed. As was mentioned earlier, most policies for VPN access tend to allow the remote users to use almost any protocol they could on the local LAN. As such, it may be easier to define the protocols you don't want your remote-user communities to be able to access, rather than define the ones that you do want to allow.

In larger deployments, it is helpful to segment the various types of VPNs off of discrete access control points of the network. This can be done through providing a dedicated firewall interface for each VPN type, as was done in the large VPN design. This setup allows different levels of trust for different VPN applications. For example, an organization might decide it trusts site-to-site VPNs a little more than remote-access VPNs. This better trust is a result of the fact that with site-to-site you know the IP address of the remote peer and are potentially using digital certificates, whereas with remote-access VPNs you generally do not know the address of your remote peer and are relying on group preshared keys combined with secondary authentication to allow your users into the network. When deployed in this manner, VPN traffic can be filtered differently based on what interface it arrives on at the access-control device. Using unique address ranges for remote site users and sites also provides additional filtering capabilities.



Filtering outbound from the VPN device (toward the public network) is also important. This filtering can help ensure that the VPN devices see only IPSec traffic coming into and out of their public interfaces. This filtering can generally be done on a router with a standard ACL instead of a firewall, freeing the firewall to sit behind the VPN device as was specified earlier. This setup 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. Most of the benefits that stateful firewalls could provide in front of a VPN device are lost regardless, because IPSec traffic cannot be intelligently filtered by most firewalls. The administrator would need to open a hole through the firewall to allow the traffic (that is, UDP 500 for IKE and IP 50 for ESP) and at that point, it is behaving in much the same way as a standard packet filter on a router. Filtering inbound on the VPN device itself is recommended to allow only IKE and ESP. If the NAT transparency mechanism is enabled, you should allow only the specific UDP or TCP port to the VPN device.

Often this access-control function can exist on the same hardware platform as the IPSec function. It can if your VPN device also has a stateful firewall, or when a remote user connects using a laptop that has both VPN client software and a personal firewall.

Split Tunneling

Split tunneling occurs when a remote VPN user or site is allowed to access a public network (e.g. the Internet) at the same time that it accesses the private network via the VPN. If split tunneling were disabled, the remote VPN user or site would need to pass all Web-destined traffic through the VPN headend, where it could be decrypted and inspected before being sent out to the public network in Web access. For example, a remote-access user who dials his/her local *Internet service provider* (ISP) and connects to the private network over an IPSec client has two options. The first has the user passing only corporate-bound data over the VPN connection. Browsing the Web would occur directly through his/her ISP in the clear. The second option has the user passing all traffic (including Internet traffic) to the headend first, where it is then routed in the clear either to the private network or out to the Internet. Deciding between the two often depends on the amount of trust you can place in the remote sites or users. To increase the level of trust for these users, consider using additional available security technologies, such as personal firewall and virus scanning. Remote sites that wish to utilize split tunneling should have a stateful firewall on their premises to control the cleartext traffic allowed into and out of the remote site. Likewise, a remote user should run a personal firewall to filter traffic and carry out virus scanning while the VPN is connected and when the VPN is not. Even when split tunneling is disabled, a personal firewall is often necessary because the user is not always connected over the VPN. A traveling user may connect in through high-speed Internet access in a hotel and elect to browse the Web while not connected to the corporation. Without a personal firewall, that system is open to attack whenever it is not connected to the VPN.

Similarly, many hardware VPN devices utilize NAT and market it as a form of firewalling. In the authors' opinion, NAT is not a security feature and should not be deployed as such. Even though addresses are hidden, no packet filtering or sequence-number checking is occurring, leaving systems protected by NAT open to brute-force attacks against them. A security perimeter that relies solely on NAT is indeed not a security perimeter. When utilizing these devices, it is important that you provide a personal firewall for the PCs residing behind the device. Even if split tunneling were disabled, a personal firewall will be necessary if that host is mobile (as in the case of a laptop).

When considering the security risks of enabling split tunneling, it is too easy to conclude that it should never be considered. Actually, disallowing split tunneling creates an enormous load on the VPN headend because all Internet-bound traffic needs to travel across the WAN bandwidth of the headend twice. This use of WAN resources is not an optimal one, and it often leads to the decision to implement the appropriate security technologies at the



remote sites to allow split tunneling to occur. In SAFE VPN, remote sites were assumed to have split tunneling enabled unless otherwise specified. If split tunneling were disabled, the designs would not change, but the performance and scaling considerations might change slightly because of the increased traffic load on the headend.

Partially Meshed, Fully Meshed, Distributed, and Hub-and-Spoke Networks

When overlaying VPNs on any network topology, many factors affect the scalability and performance of the network. Some of these factors include encrypted versus clear traffic processing, hardware accelerated versus software based IPsec, configuration complexity, high availability, related security features (firewall, IDS, and so on), the number of routing peers and networks to track, and maintaining QoS. Fully meshed networks quickly run into scalability constraints because every device in the network must communicate with every other device in the network via a unique IPsec tunnel. That is $n(n - 1)/2$ tunnels, or for a 50-node network, 1225 tunnels! The configuration complexity is immense, and at some point growing the size of the mesh will not be possible. Keeping state for that many tunnels also has performance implications. Partially meshed networks scale better than fully meshed because inter-spoke connections are established only as needed. Similar to devices in fully meshed networks, the limiting factor in this topology is the number of tunnels that the devices can support at a reasonable CPU utilization. Both of these networks could use a dynamic tunnel endpoint discovery mechanism to simplify the configuration and increase scalability. However, as documented in the caveats section, this mechanism is not covered in this document.

Hub-and-spoke networks scale better because the headend hub site can expand to meet growing spoke capacity requirements. Low-horsepower spokes that need connectivity to other remote sites will be able to connect via the hub site. However, all traffic flows through the hub site, and this setup requires significant bandwidth because it includes all spoke-to-spoke traffic as well as spoke-to-hub traffic. Not all headend VPN devices support spoke-to-spoke intercommunication. Split tunneling may be required at remote sites. Depending on the type of device chosen at the headend, spoke to spoke connections or Web access via the headend may or may not be possible. For instance, the model for firewalls is to enable split tunneling at all sites, thus eliminating the need for the hub firewall to process spoke-to-spoke traffic. If there are regional or other requirements for traffic routing where most traffic does not require access to networks via the hub site, consider a distribution layer to lower the bandwidth requirements at the headend and thus increase the scalability of the network.

Interoperability and Mixed versus Homogenous Device Deployments

Although IPsec is a documented standard, the *Request for Comments* (RFCs) that document it have left room for interpretation. In addition, Internet drafts such as IKE mode-configuration and vendor-proprietary features increase the likelihood of interoperability challenges. For instance, there is no standard mechanism for IPsec to determine tunnel up/down state and remote peer reachability. For these reasons, you should check with vendors of both products for interoperability information and their participation in interoperability bake-offs. Typically a few minor changes to configurations-and sometimes code-are necessary to facilitate interoperability in a reliable fashion. Realize, though, that these changes may affect the security stance of the device, so be aware of the implications of these changes. Also, in order to ensure interoperability between products from a single vendor, it is a best practice to use the same code base across all platforms. This scenario will decrease the likelihood of any interoperability issues with products made by the same vendor as changes are made over time to adhere to the standards and increase interoperability with other vendors.

Issues in addition to interoperability arise in environments where different device types are deployed to build a VPN. These issues usually arise because of interaction between the VPN and other features that complement its operation. For instance, consider the *authentication, authorization, and accounting* (AAA) protocol used to manage remote



users and administrators. The granularity of support for this protocol, say *Terminal Access Controller Access Control System Plus* (TACACS+) or *Remote Access Dial-In User Service* (RADIUS), may differ among the device types. This difference can complicate matters if your user database does not support one of these mechanisms across all the device types deployed.

The mechanisms used for IPsec high-availability and CA support differs for some routers, firewalls, concentrators, and remote-access clients. Finally, consider the additional resources required to train administrators on how to configure, manage, monitor, and troubleshoot multiple device types.

Fragmentation and Path Maximum Transmit Unit Discovery

Fragmentation should be avoided at all costs. Packet reassembly is resource intensive from a CPU and memory allocation perspective, and normally fragmentation can be avoided. Allowing fragmented packets into your network creates security concerns. Fragmented IPsec packets require reassembly before the packets can undergo integrity validation and decryption. Most of the time fragmentation occurs when a packet is sent over a tunnel and the encapsulated packet is too large to fit on the smallest link on the tunnel path. IP supports *path maximum transmission unit discovery* (PMTUD) to deal with this scenario and most modern-day operating systems support this mechanism. As long as filtering does not block the *Internet Control Message Protocol* (ICMP) messages, PMTUD will determine the maximum MTU that a host can use to send a packet through the tunnel without causing fragmentation. To allow PMTUD in your network, do not filter ICMP message Type 3, Code 4. If ICMP filtering is occurring and is out of your administrative control, you will either have to manually set the MTU lower on the VPN termination device and allow PMTUD locally, or clear the *Don't Fragment* (DF) bit and force fragmentation. Avoid the latter at all costs and consider it only as an last resort.

Packets generated by hosts that do support PMTUD will use it locally to match the MTU on the tunnel whether it was set statically or dynamically via PMTUD. When you manually set the MTU on the tunnel, you must set it low enough to allow packets to pass through the smallest link on the path; otherwise the packets that are too large to fit will be dropped, and if ICMP filtering is in place, no feedback will be provided. Remember that multiple layers of encapsulation will add layers of overhead to the packet. For instance, GRE and ESP tunneling protocols are used together frequently. In this scenario, GRE adds 24 bytes of overhead to the packet before it undergoes encapsulation again by ESP. ESP, when using 3DES and SHA, then adds 56 bytes of additional overhead. Because ESP and GRE support PMTUD, the likelihood of fragmentation is reduced. Depending on the VPN termination device, the manner in which you should set the MTU on the tunnel varies. Options include changing the MTU via the tunnel interface (routers), the TCP maximum segment size (firewalls), policy routing (routers), clear/set/copy DF bit (routers), OS application level (VPN clients), and physical/logical interfaces (any VPN device).

Network Operations

Remote devices need to be managed via a VPN from the central site when operating in a centralized IT model. VPN devices support numerous configuration options to determine the tunnel endpoint. Depending on the method chosen, these options may impact the manageability of the network. To be the most effective in managing remote devices, you must use static cryptographic maps at the site where your management applications are located. You should not use dynamic cryptographic maps at the headend. Dynamic cryptographic maps accept only incoming IKE requests, and they cannot initiate them, so it is not always guaranteed that a tunnel exists between the remote device and the headend site. Static cryptographic map configuration includes the static IP addresses of the remote peers. Thus remote sites must use static IP addresses to support remote management.



Some management services, *Trivial File Transfer Protocol* (TFTP) for example, utilize the nearest interface as the source address of the generated packets. For this reason, you should be careful when setting up the cryptographic ACLs to ensure that the traffic will pass through the tunnel and not in the clear to the headend. You should enable read-only *Simple Network Management Protocol* (SNMP) access on VPN devices to trap the information available via an IPsec Management Information Base (MIB). You should allow only SNMP access on secure interfaces. An IPsec MIB can track tunnel statistical information and tunnel status via a tunnel history table and tunnel failure table. The history table archives attribute and statistics information about the tunnel; the failure table archives tunnel failure reasons along with the time that failure occurred. This information is crucial when monitoring and troubleshooting devices in any network size. Consider utilizing the MIBs in addition to the *command-line interface* (CLI) for troubleshooting and proactive monitoring in large deployments. Most configuration tools today assume green-field environments. For this reason, you should deploy these tools first, even in the prototype stage, to ease the configuration burden when going production.

To securely manage remote devices, some form of user authentication must occur beyond the device authentication and tunnel encryption that occurs when establishing the tunnel. Management tools require a static username/password pair to configure remote devices without administrator intervention. Do not use the same username/password pair to access devices for management that you use for normal day-to-day administration. Make sure that fixed username/passwords are aggressively aged. The administrators should use strong authentication to access devices, preferably OTPs. Also consider that the user authentication protocol will likely need to run over the tunnel to query an AAA server at the headend.

If the device fails or is misconfigured, the management tools and remote administrators will not be able to manage the device. Consider using a local static username/password pair just in case, or a static cryptographic map entry defined solely for the purpose of remote management. If digital certificates are required for device authentication for tunnel establishment, the remote device may lose reachability to the CA to validate the received certificate against the CRL. If the default to require CRL checking has not been changed, tunnel establishment cannot occur. When using certificates, time synchronization is mandatory for checking the certificate and CRL lifetimes. All VPN devices should use the NTP protocol (using authenticated NTP) to synchronize the time. To avoid the chicken and the egg issue, ensure that CRL and NTP have a path that does not rely on digital certificates. This is carried out in the simplest case by using a preshared key and separate ACL for CRL and NTP access. For management traffic if you use a different static entry with different ACLs than that used for all other traffic, it is less likely that a configuration error will cause loss of reachability to the device.

Finally, a common administrative task is updating the version of VPN remote-access client software. To ease this burden, some VPN concentrators support a mechanism that pushes new client software to remote-access users the next time they connect to the headend. This scenario allows your enterprise to manage large numbers of clients that are constantly on the run.

HSRP

For router-based solutions, you should consider *Hot Standby Router Protocol* (HSRP) for resilience when routing protocols are not in use for high availability. Today, headend routers can listen for IKE requests on the HSRP virtual IP address. When using IKE keepalives for high availability of remote crypto peers, this mechanism provides resiliency by checking for reachability to the virtual IP address. It also aids in resilience when the next-hop device behind the headend, such as a firewall, needs resilience to two resilient VPN termination devices.



Compression

Layer 2 compression provides no reduction in link bandwidth for VPN traffic. Compression works by finding repeated instances of information in streams of data and replaces them with a smaller representation (a unique set of bits) for transmission. A table maps the two together. On receiving the stream, the remote device replaces the unique sets of bits with the original data. There are no repeatable instances of similar bits in VPN traffic. Encryption randomizes the data stream to such a degree that virtually no compression algorithm can provide any reduction. If this were not the case, the encryption algorithm would not be cryptographically strong. The IP Layer 3 standard for compression is called IPComp. Layer 3 compression, which occurs before encryption, does provide reduction in the amount of data to undergo encryption for moderate to large sized packets. Less data to encrypt actually increases the encryption throughput. Given that Layer 3 compression is not supported in most hardware accelerators today, when it is used it is very CPU intensive. If you choose to enable software compression, be sure to include this variable when determining the network design.

Remote-Access User Requirements


Road warriors and telecommuters alike have the same requirements while on the road as when in the office. Most likely they will need access to the following: *Domain Name System* (DNS) to resolve Internet host and domain names, *Windows Internet Naming Service* (WINS) to resolve hosts in their Windows domain, and a virtual IP address that will allow them to access the corporate Intranet. These values are pushed to a remote client by the ISAKMP configuration method (IKE MODCFG) during tunnel establishment but only after successful authentication. The virtual address can in turn be used by any device on the Intranet to connect to the client. For times when users are out of the office and out of control, Cisco recommends that you control the connection they have both to your Intranet and the Internet. If you choose to allow split tunneling, make sure that personal firewall software is installed, updated, and running, and that it has a valid security policy on the remote-access client. Otherwise, if a rogue applet, Trojan horse, or some other outside source gains control of the client, after it has been compromised it could then be used to attack the private network. Cisco does not recommend enabling split tunneling on clients without firewall capability. Even with an effective personal firewall deployment, you should still deploy a firewall in the headend in case the remote user device is compromised and to guard against mischievous users.

High Availability

Because IPSec tunnels send data without any acknowledgment or feedback mechanism from the remote peer that it has received the data, devices should track the remote peer's state. Otherwise, if the device loses reachability to its peers, the tunnel will turn into a black hole even though the tunnel still appears established. Today two mechanisms are available to determine remote peer availability and tunnel establishment state: routing protocols and IKE keepalives. Routers support both mechanisms, whereas firewalls, concentrators, and remote-access clients support IKE keepalives.




Figure 1
IKE Keepalive High Availability Example



Note: One tunnel is active at any given time.

IKE keepalives are sent over the IKE SA to determine remote-site IKE peer reachability. When the peer is no longer reachable, a new tunnel is established. In case of failure in this environment, a new path establishment will occur in the amount of time that it takes to establish a new tunnel. If the primary device comes back on line, the remote devices will continue to use the secondary device for termination. It will not preempt unless HSRP is used and preempt is enabled. Note that for this mechanism, each remote site has a single path to the headend.

Figure 2
Routing Protocol High Availability Example



Note: Two tunnels are active simultaneously.

In comparison, routing protocols use two paths. Routing protocols are sent over the IPsec-protected GRE tunnels to track remote network reachability. A remote site using routing protocols for high availability will establish two IPsec-protected GRE tunnels, one to each headend. Routing updates traverse both tunnels to the remote site, which will then forward the traffic to the headend that has reachability to the destination network. From the perspective of the remote site, there are two paths to the headend. Consider defining one of the tunnels as the primary tunnel to avoid asymmetric routing. This setup is accomplished by adjusting the routing protocol cost for one of the links. In case of tunnel failure, convergence will occur as soon as the routing protocol realizes the path is no longer available. On failure recovery, remote sites using routing protocols will optionally revert back to their primary preferred path once the primary tunnel automatically recovers.


Concentrator and firewall headends often support fail-over capabilities in an active/standby configuration. When the primary fails, the secondary unit assumes the IP and *Media-Access-Control* (MAC) address of the primary, and the tunnel reestablishment commences. Routers function in an active/active configuration. Both headend devices will allow tunnel establishment. You might consider using IKE keepalives in the headend for heterogeneous remote-site



device support. There is no IETF standard for keepalives today, only proposals, and thus this mechanism will work only with products from a single vendor. If a momentary loss of connectivity occurs at a remote site, it may establish a new tunnel with the secondary (but always active) headend device. Because tunnel establishment does not affect the routing table unless routing protocols are running over the tunnel, the routing state in the headend will not change. When the tunnel switches between the headends because of the remote-site flapping, the next-hop router will not be able to determine which active headend device has a valid path to the remote site. Flapping occurs when the remote site temporarily loses WAN connectivity. In order to avoid this issue consider using HSRP and *reverse-route-injection* (RRI). RRI works by modify the routing table on the device to reflect tunnel SA status. Thus once a tunnel is established to a remote site, its network is injected into the routing table. Should that SA fail or expire, the route is removed.

In summary, when using VPN concentrators or VPN firewalls at the headend, use IKE keepalives for high availability. When using VPN routers at the headend, use routing protocol resilience for high availability.

Figure 3
Load Dispersion on Failure



When a headend tunnel termination device fails, its load should be equally shared among the other remaining headend devices; see above figure. Referred to as load dispersion on failure, this process significantly aids in the resiliency and scalability of the headend. Unfortunately it also adds to the configuration complexity. Load dispersion on failure is applicable only in active/active configurations and not active/standby because partial load distribution cannot occur to a standby device. Regardless of the high-availability mechanism chosen, a headend device should not be deployed in a configuration that results in CPU utilization higher than 50 percent after failure. The 50-percent target includes all overhead incurred by IPSec and any other enabled features (firewall, routing, IDS, logging, and so on). In some environments it may be possible to achieve a higher CPU utilization in a stable and reliable fashion.



However, under failure circumstances when hundreds of sites are deployed, the amount of processing power needed for routing protocol convergence or new tunnel establishment mandates a design that has enough headroom to reliably recover.

One of the most common issues today in large-scale VPNs is the stale SA, which occurs when one device at the end of the tunnel maintains the tunnel state but the other remote end does not. The loss of state could occur during link failure, misconfiguration, troubleshooting, system maintenance, or complete device failure. IKE keepalives resolve this issue by removing the state of the old tunnel and setting up a new tunnel. Routing protocol resilience, however, keeps the tunnels up at all times and, therefore, is more likely to run into a stale SA problem. There is no feedback link between network reachability over a tunnel and tunnel status. In other words, if a network is no longer reachable over a tunnel, the tunnel is not torn down until it times out. When the remote device comes back on line, if it had lost tunnel state, it will attempt to establish a new tunnel. The device that remained active will receive a tunnel-establishment request for a tunnel for which it already has state. It will use the new tunnel to transfer data and the old tunnel will be torn down. Continuing to use the old tunnel to transfer data before the remote peer comes back up is an issue in everyday system administration when devices are taken off line for maintenance.

Cisco recommends that when you are taking headend VPN devices off line for maintenance that you clear the IKE SAs on the remote devices to facilitate IKE reestablishment. Do not change the IKE SA lifetimes. Although this solution will expire the tunnels faster and make administration more simplified in these scenarios, the amount of increased processor utilization does not justify the end to the means. Cisco does not recommend running IKE keepalives in combination with routing protocols for resilience to assist in keeping the state current. Running both mechanisms simultaneously would reduce scalability significantly, given the summed CPU overhead.

A new type of IKE keepalive from Cisco, called *dead-peer detection (DPD)*, offers the same functionality as IKE keepalives at lower CPU overhead. DPD was recently proposed to the IETF and is in the draft stage. It works by sending IKE peer reachability probes only to devices that it has not received data from in a configurable time period. IKE keepalives send updates to all IKE peers regardless, a fact that explains why it is a CPU-intensive process similar to a routing protocol sending hellos. In routing protocol-resilient environments, the VPN router will send routing protocol traffic constantly over all tunnels. Finally, even if you do not require high availability in your network, consider using IKE keepalives-or preferably DPD-to keep tunnel state current. This setup will ease the administration burden.

Branch versus Headend Considerations

The small and medium designs that follow can be used in two possible configurations. In the first, the design is acting as a branch of a larger organization, built in the configuration described in SAFE Enterprise. In the second configuration, 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 telecommuters might come into the headend over some of the options discussed in the remote network design.

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 (remote, small, medium) discussed in this paper for its remote locations and telecommuters. Where appropriate, the specific changes that may be required to a design are discussed in each section.



Remote-User Design

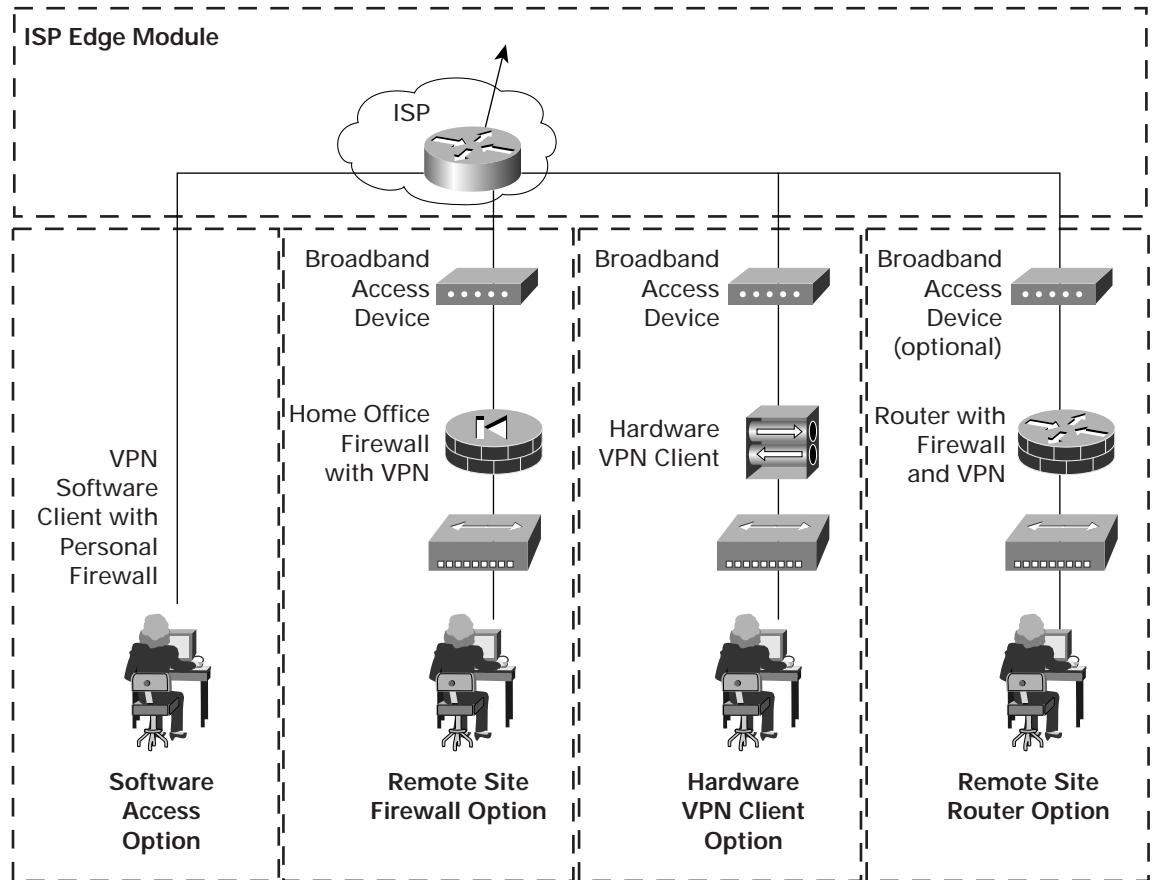
This section discusses four options for providing remote users VPN connectivity to the headend sites within the SAFE design. Remote connectivity applies to both mobile and home-office workers. The primary focus of these designs is to provide connectivity from the remote site to the corporate headquarters. This would occur most likely through broadband access to the Internet. The following four options are available:

- *Software access option*—Remote user with a software VPN client, personal firewall, and virus scanning software on the PC
- *Remote-site firewall option*—Remote site 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 (that is, 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 following design guidelines section.



Figure 4
Detailed Model of Remote-User Module

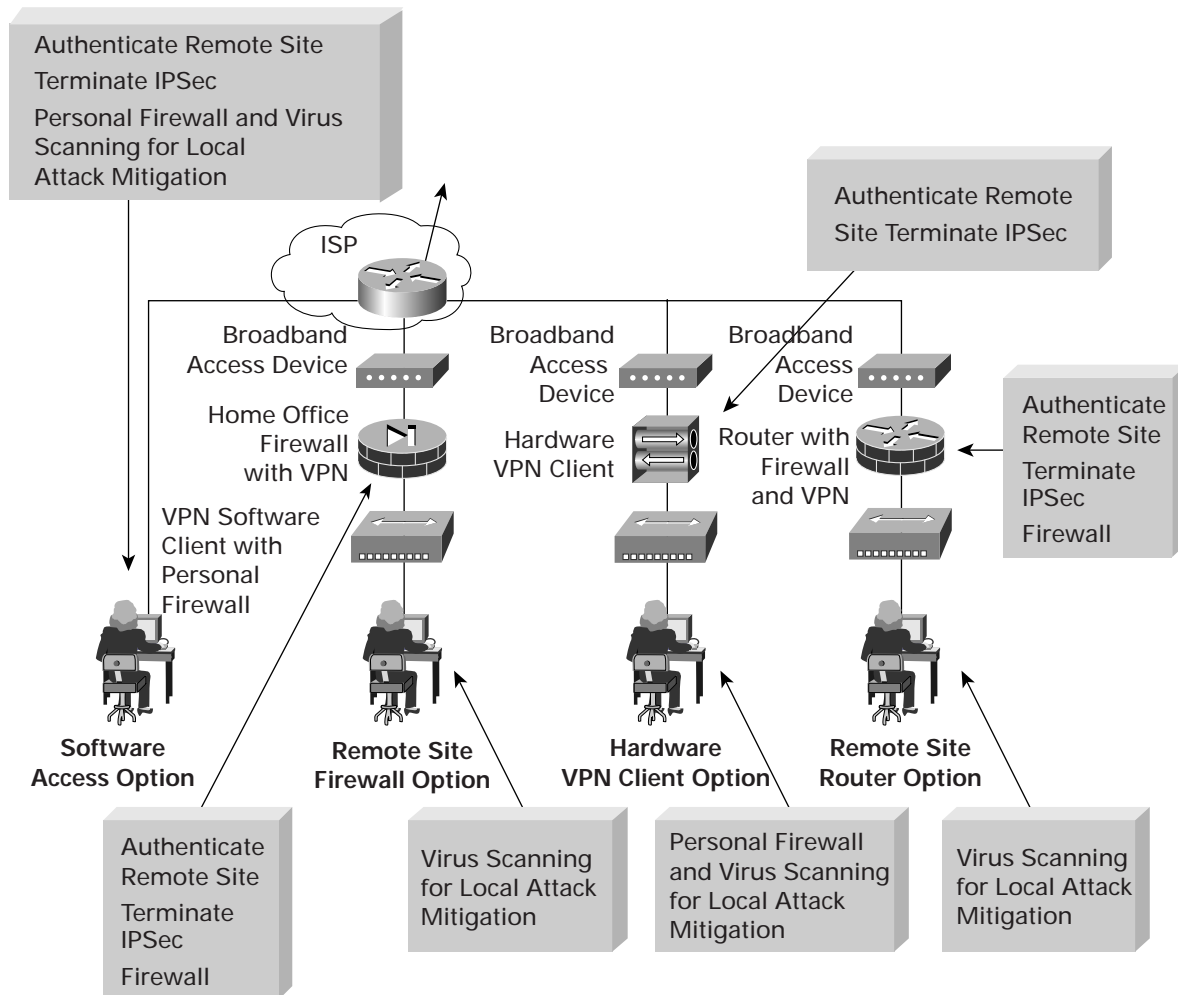


Key VPN Devices

- *Broadband access device*—Provides access to the broadband network (DSL, cable, and so on)
- *VPN firewall*—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
- *Personal firewall software*—Provides device-level protection for individual PCs
- *Virus scanning software*—Provides protection against viruses and worms for individual PCs
- *VPN firewall router option*—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
- *Remote-access VPN client*—A software solution that 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 5
Detailed Model of Remote-User Module: VPN



Design Guidelines

The following sections detail the functionality of each of the remote-user connectivity options. All devices are assumed to have a single local or “flat” network. All VPN devices are configured in these designs with a default route set to the ISP next-hop device. It is assumed that the *customer-premises-equipment* (CPE), e.g. the broadband device shown above, is not performing a NAT function. Performance in these designs is limited by the restrictions normally imposed by the media types or by service-provider restrictions on available upload bandwidth.



Software Access Option

The software access option is geared toward both the mobile and home-office workers. 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. 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 that is used for VPN-destined traffic, and the location of name servers (DNS and 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 tunnel established. Because the remote user may not always want the 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 Trojans infecting the PC.

IKE keepalives are used for the high-availability mechanism to determine headend availability. Device authentication to the headend occurs via a group preshared keys, and OTPs are used for user authentication via XAUTH. Secure management in this option includes policy pushing over the tunnel and notification of new VPN client updates. If many-to-one NAT occurs between the client and the headend, the NAT transparency mode should be enabled. The only alternative to this option is to consider one of the next options, which in comparison allow for Internet access via split tunneling and provides stronger management of security features via an integrated firewall/VPN device.

Remote-Site Firewall Option

The remote-site VPN 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 VPN firewall is installed behind a DSL or cable modem. The VPN firewall establishes a tunnel to a VPN headend device and provides Internet access via NAT, stateful inspection, and filtering. Individual PCs on the remote-site network do not need VPN client software to access corporate resources unless they travel and need access to the corporate Intranet via the Internet. Split tunneling was enabled in this remote-site configuration, given the enterprise-class firewall features available. Virus-scanning software is recommended to mitigate the risks of split tunneling.

Proper address summarization should be implemented in order to ease the administrative burden and allow for remote-site intercommunication if needed. 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) are not required to perform any configuration changes on the home-office firewall. Individual users at this remote site who access the corporate network do not undergo user authentication in this option. It is assumed that the environment is controlled. If the environment is not controlled, consider user authentication at the headend firewall. The VPN utilizes device preshared key authentication. Given a large deployment, digital certificates are recommended. DPD-type IKE keepalives are used for the high-availability mechanism to determine headend availability. NAT was not used over the VPN to translate the local network because it was assumed that the network does not overlap with any other networks.



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 option 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 dependent upon the attacker being unable to circumvent NAT. This dependency results from the fact that, when split tunneling is enabled, connections to the Internet pass through a simple many-to-one address 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.

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 a *Secure-Sockets-Layer* (SSL) connection from the central site. This setup ensures that remote-site user(s) are 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, regardless of the OS type installed, 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. The hardware client operates in two possible modes. In the first, all users behind the hardware client appear as a single user on the corporate Intranet via the use of many-to-one NAT. In the second, all devices access the corporate Intranet without NAT, and hosts in the Intranet may initiate connections to the hosts behind the hardware client once the tunnel is established. In SAFE VPN, the latter mode was deployed. The first mode is simpler to manage and thus more scalable, but the second mode is more versatile. The level of security provided by both modes is the same.

The VPN hardware client undergoes device authentication with the VPN headend concentrator using a statically configured group preshared key. It is assumed that the environment is controlled. If the environment is not controlled, consider user authentication at the headend firewall. Given a large deployment, digital certificates are recommended. DPD-type IKE keepalives are used for the high-availability mechanism to determine headend availability.

Remote-Site Router Option

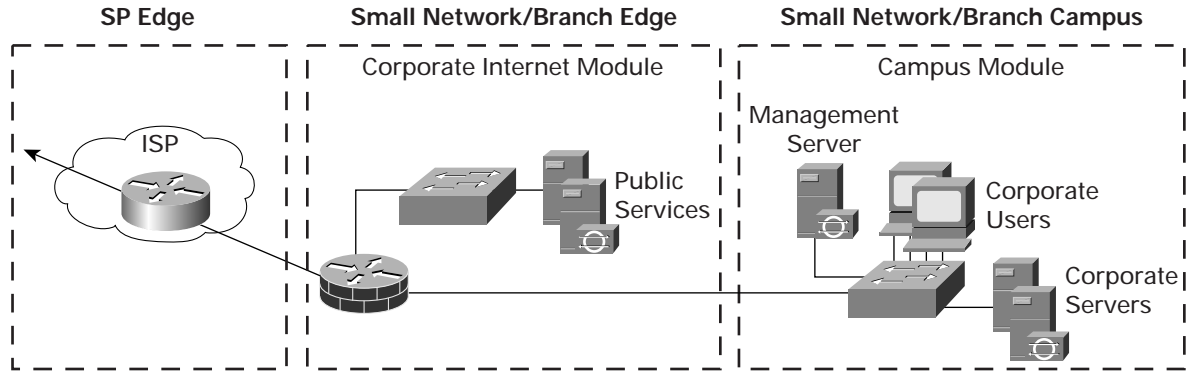
The remote-site router option is nearly identical to the remote-site firewall option with a few exceptions. First, because this router is a full-featured router, advanced applications such as QoS can be supported. QoS could be used to prioritize access to the corporate Intranet over Internet Web surfing. Secondly, there exists an option to integrate the functions of both the VPN firewall and the broadband access device into a single device. This option requires that your ISP allow you to manage the broadband router itself, a scenario that is not common. IKE keepalives or routing protocols could be used for the high-availability mechanism to determine headend availability.

Small VPN Design

The small VPN design utilizes the same topology as the small network design from the SAFE security papers. This design supports both site-to-site and remote-access VPNs, with some caveats included in the configurations section. Most of the discussion in this design is based on the premise that this design will operate as the headend for a corporation. Specific design changes when used as a branch are also included. The small network VPN design is contained within the small network corporate Internet module. The entire small business design is shown here for reference:




Figure 6
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 was also provided to remote locations and telecommuters.

Figure 7
Detailed Model of Small Network Corporate Internet Module

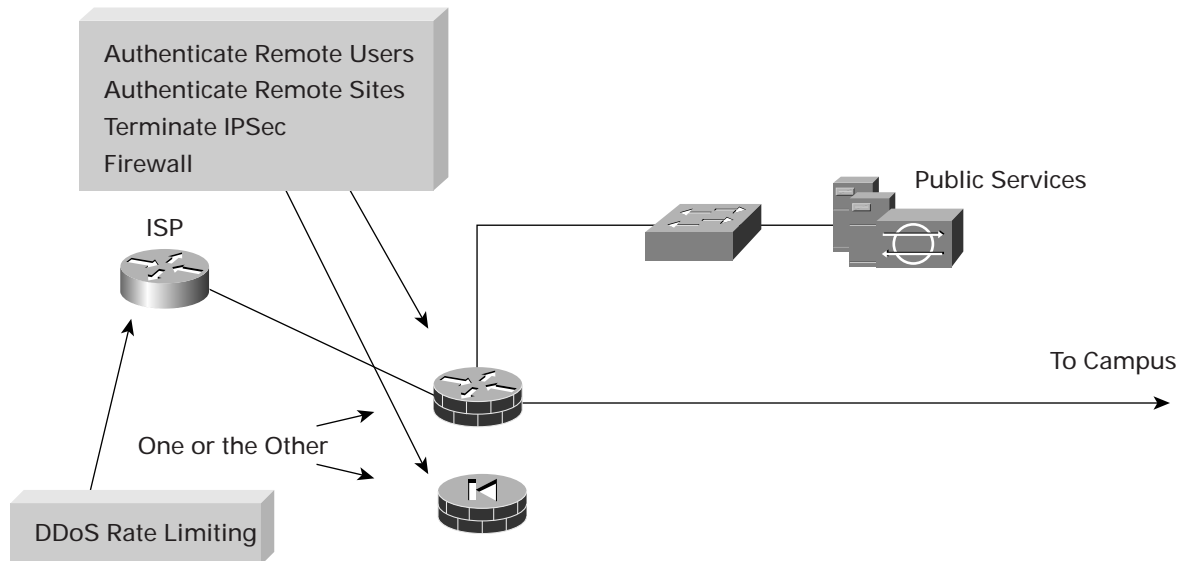


Key VPN Devices

- *Firewall or firewall router*—Provides network-level protection of resources and stateful filtering of traffic
- *RADIUS authentication server*—Provides authentication of remote-access connections (located in campus)



Figure 8
Detailed Model of Small Network Corporate Internet Module: VPN



Design Guidelines

This module represents the ultimate in scaled-down VPN network design, where the VPN functions are compressed into a single box that also performs routing, NAT, IDS, and firewalling. 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 scenario yields the greatest flexibility for the small network because the router will support all the advanced services that may be necessary in today's networks. As an alternative, a dedicated firewall with VPN may be used instead of the router. This setup places some restrictions on the deployment. First, firewalls are generally Ethernet only, a setup that would require some conversion to the appropriate WAN protocol. In today's environments, most cable and 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 being generally easier to configure security and VPN services and can provide improved performance when doing firewall functions. Whatever the selection of device, there are numerous VPN considerations. Remember that routers tend to start out permitting traffic, whereas firewalls tend to deny traffic by default.

The VPN functionality as implemented in this design is similar, regardless of the hardware platform chosen. Both router and firewall support stateful firewalling, basic NIDS, NAT, and IPsec. Because there is no headend resiliency unless the small design is used as a branch of the large design (discussed later), basic tunnel-mode IPsec was chosen without any options. The following sections detail the specific design considerations for the small network.



Identity

For site-to-site VPN connections to remote locations, preshared keys and the IPSec peer IP addresses are used to validate the identity of the IPSec devices. Although this scenario does not have the security of a digital certificate, for a small VPN, preshared keys can be easily managed because the number of sites tends to be fewer than ten.

The remote-access VPN connections employ a two-part authentication scheme that uses a wildcard preshared key on the device coupled with a secondary user authentication through RADIUS. Because this authentication does not include OTP, a greater reliance is placed on a user's selection of a strong password. Passwords should also be aged quickly, and users should be locked out of the VPN after a certain number of failed login attempts. This scenario will aid in the prevention of brute-force attacks if someone is able to steal a user's laptop and IT is not immediately informed. Remember that this fixed-password authentication scheme is in no way strong, and it leaves the organization with a very thin layer of defense if one of the two authentication schemes is compromised.

Security

From a security perspective, the inbound ACL allows only IKE and ESP traffic to terminate on the public interface of the VPN device. From there, traffic was decrypted and then filtered again through the firewall function on the device. This scenario allows the administrators of the small network VPN to define the types of protocols it wants to allow into the network. Remote access VPN users are not allowed to split tunnel but remote site-to-site devices are allowed.

Scalability

This type of design is not particularly scalable. It is designed for fewer than 20 remote sites and fewer than 50 concurrent remote users. However, this design will meet the needs of most small networks.

Secure Management

Secure management of the device itself was done using a mix of secure and nonsecure protocols such as *Secure Shell Protocol* (SSH), SNMP, TFTP, and syslog. With regard to management of the remote sites, the management traffic can be passed over the IPSec VPN connection to that site. This setup enables the management functions to pass encrypted over the Internet. The remote VPN devices themselves will not be part of the IPSec tunnel and will need to be managed via a separate tunnel to allow secure intercommunication between the management hosts and the outside interface of the remote device. Remote access VPN clients are securely managed every time they establish a tunnel to the headend and are pushed the latest policy each time.

NAT

NAT was used in this configuration only for access out to the Internet for the small network. NAT was bypassed for inter-site and remote-user VPN communications, a setup that allows all trusted parties to communicate with their real IP address as non-overlapping RFC-1918 addressing was used. For remote users, this "real" address is the virtual address pushed to their VPN client.

Routing

With the exception of some simple static routes, routing was not needed in this configuration given the flat network. A default route could be attained dynamically from the ISP. All the internal users by default route to the VPN device, which then default routes to the ISP. Static routes to remote sites are not needed because after the packet is routed to the outside interface it is encrypted and sent to the remote peer.



Extranets

This type of design is not particularly conducive to an extranet environment. If an extranet connection is needed, strong consideration should be given to placing that link on a separate VPN device. This device should be located off a new segment on the original VPN device in this design. This setup allows that traffic to be dealt with separately from the corporate VPN traffic and reduces the chance for misconfiguration.

Performance

Typically the WAN connection is the limiting factor in small networks today. When VPN traffic is coupled with standard Internet traffic, care should be taken to avoid overflowing your Internet link with VPN traffic. Allowing remote sites to use split tunneling can aid in this. Because most VPN devices do not yet offer the ability to limit VPN traffic or the number of users on a device, the mix of traffic at the headend will need to be carefully watched.

Alternatives

The most obvious need in the small VPN design is stronger authentication for remote users. This authentication can be achieved through the addition of OTP into the small environment. It was not included in this design because many small networks are unable to make the financial commitment to use OTP technology. Any other deviation from this design would be geared toward increasing the capacity of the network, or separating the various security functions onto distinct devices. If these changes are incorporated, the design will start to look more and more like the medium VPN 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.

Branch versus Standalone Considerations

The choice of the type of VPN device to use for a branch role is driven more by the type of resiliency (if any) located in the headend. If the headquarters is not using IPSec high availability in any form, or if headquarters is using some sort of IKE keepalive mechanism, then the choice of device for the small branch remains grounded in the same selection criteria mentioned above. If, however, the headend is utilizing IPSec over GRE for high availability, the small VPN device needs to support a routing protocol and GRE encapsulation. Many firewalls do not offer this functionality today.

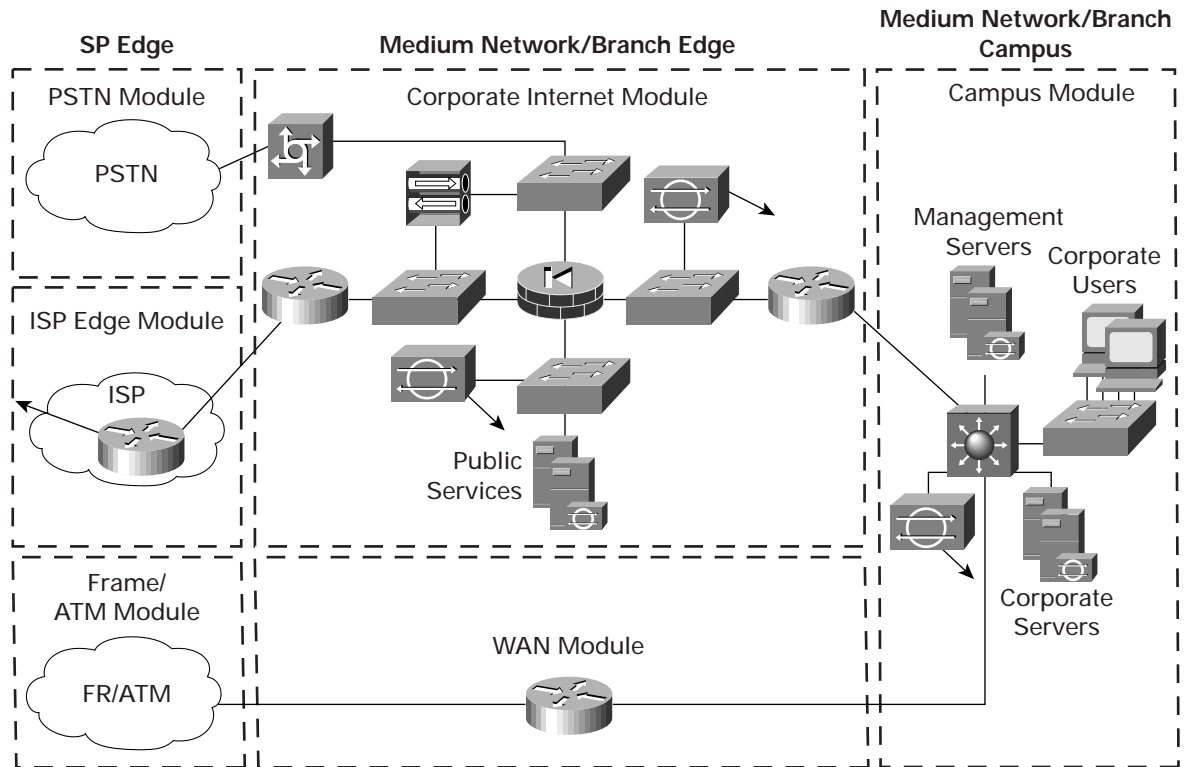
Also, as a branch, the VPN device will most likely be managed from the remote site, making the cryptographic ACLs slightly different.

Medium VPN Design

The medium-network VPN design utilizes the same topology as the medium-network design from the SAFE security papers. This design supports both site-to-site and remote-access VPNs. Most of the discussion in this design is based on the premise that this design will operate as the headend for a corporation. Specific design changes when used as a branch are also included. The medium VPN design is contained within the medium-network corporate Internet module. The entire medium-network design is shown here for reference:



Figure 9
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 (*Hypertext Transfer Protocol [HTTP]*, *FTP*, *Simple Mail Transfer Protocol [SMTP]*, and *DNS*). Additionally, this module terminates VPN traffic from remote users and remote sites as well as traffic from traditional dial-in users.




Figure 10
Detailed Model of Medium-Network Corporate Internet Module



Key VPN Devices

- *VPN 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
- *VPN concentrator*—Authenticates individual remote users and terminates their IPSec tunnels
- *NIDS appliance*—Provides Layer 4-to-Layer 7 monitoring of key network segments in the module

Figure 11
Detailed Model of Medium-Network Corporate Internet Module: VPN





Design Guidelines

The medium VPN design separates site-to-site and remote-access VPN traffic onto two separate devices. This setup allows better performance because each device is concerned with only one type of VPN. By moving to a dedicated remote-access VPN device, manageability of the remote-user community also increases. The site-to-site VPNs are done on the dedicated firewall at the heart of the module. Because resilience was not a part of this design, as in the small design, standard tunnel mode IPSec was used.

Identity

For site-to-site VPN connections to remote locations, digital certificates and the IPSec peer IP addresses are used to validate the identity of the peer devices. This scenario provides better security and manageability over standard preshared keys. The remote-access VPN connections employ a two-part authentication that uses a wildcard preshared key on the device coupled with a secondary user authentication via OTPs.

Security

From a security perspective, the edge router allows only IKE and ESP traffic to terminate on the public interfaces of the VPN devices. From there, traffic is decrypted and then either passed to a firewall if it is remote-access VPN traffic or filtered locally using the firewall function in the case of site-to-site VPNs. After filtering has occurred, the traffic passes a layer of NIDS as traffic is sent to either the public services segment or is routed into the campus. This NIDS is configured to shun on the firewall for certain alarms. Remote access VPN users are not allowed to split tunnel but remote site-to-site devices are allowed.

Scalability

This type of design is much more scalable than the small design. With hardware acceleration, this design can easily support over 500 concurrent remote users and up to 50 remote sites. Depending on the amount of bandwidth each remote site and user consumes, these numbers could be larger or smaller. If you think that you will be pushing the limits of this design but do not have the financial resources for the large design, consider the “alternatives” section for this design below.

Secure Management

Secure management of the VPN devices was done using a mix of secure and nonsecure protocols such as SSH, SSL, SNMP, TFTP, and syslog. With regard to management of the remote sites, the management traffic can be passed over the IPSec VPN connection to that site. This scenario enables the management functions to pass encrypted over the Internet. The remote VPN devices themselves will not be part of the IPSec tunnel and will need to be managed via a separate tunnel to allow secure intercommunication between the management hosts and the outside interface of the remote device. Remote access VPN clients are securely managed every time they establish a tunnel to the headend and are pushed the latest policy each time.

NAT

NAT was used in this configuration only for access out to the Internet for internal users. NAT was bypassed for inter-site VPN communications, a scenario that allows all trusted parties to communicate with their real IP address as non-overlapping RFC 1918 addressing was used. Remote users are assigned a virtual IP address by the concentrator, which is provided by the AAA server. NAT transparency mode is enabled on the VPN concentrator to facilitate remote-access client connections.



Routing

All internal user traffic is routed to the VPN firewall. It then routes via a static-route, remote-access client-bound traffic to the VPN concentrator and routes all other traffic via a default route to the edge router. This scenario results in remote site-bound traffic to trigger the cryptographic ACLs.

Extranets

This type of design could easily support an extranet connection with a separate VPN device connected off a new interface on the primary firewall. Depending on the type of extranet, this connection could route through the firewall or connect directly into the campus (assuming the device had local firewalling features).

Performance

Like the small network, the WAN connection will probably be the limiting factor in this design. The equipment in the design could easily saturate a DS3 link (45 Mbps) or more. Most networks of this size will have less bandwidth, however, and the network will need to be designed carefully to avoid overflowing it. Allowing remote sites to use split tunneling can aid in this overflow prevention. Because most VPN devices do not yet offer the ability to limit VPN traffic or the number of users on a device, the mix of traffic at the headend will need to be carefully watched.

Alternatives

Dedicated Site-to-Site VPN Device

The most common modification to this design would be to dedicate all remote access and site-to-site VPN functionality to the concentrator. Many customers choose this option because they prefer to specialize the functions of their VPN and firewalls with two separate devices. This scenario frees the firewall to perform firewalling only and allows the VPN device to focus on IPSec. Another option would be to deploy a site-to-site VPN device on the same network, or a parallel network, with the concentrator thus providing remote access and site-to-site VPN separation. Both options yield greater scalability and manageability to the medium network and are similar to the large VPN design in the next section.

Branch versus Standalone Considerations

When configured in a branch role, the dedicated remote-access VPN concentrator would probably not be needed because this service would be provided by the headquarters location. The site-to-site function could still exist on the firewall, assuming the headend is not using IPSec over GRE as its primary tunneling choice. If GRE is used, site-to-site VPNs would need to move to a dedicated router-based platform where GRE is supported.


Also, as a branch, the VPN devices will most likely be managed from the remote site, making the cryptographic ACLs slightly different. In addition, some devices, such as the edge router, would be outside the VPN itself, requiring a separate means of management. This management could be via a discrete tunnel or through the use of application level security (SSH). Remember that not all management protocols have a secure variant.

Large VPN Design

The large VPN design utilizes the same topology as the large enterprise network design from the SAFE security papers. The large network VPN design is contained within the large enterprise VPN and remote-access module and supports both site-to-site and remote-access VPNs. This module was redesigned to provide high-speed, highly available VPN termination. The entire large business design is shown here for reference:



Figure 12
Detailed Model of Large Network




VPN and Remote-Access Module

The VPN and remote-access module provides termination of VPN traffic from remote users, VPN traffic from remote sites, and termination of traditional dial-in users. The designer is given an option to choose a VPN router or VPN firewall for site-to-site VPN termination. This device is labeled “VPN” in the module layout shown below. Because of the high-speed requirements for the VPN, purpose-defined devices are deployed throughout the module, providing disparate functions.



Figure 13
Detailed Model of VPN and Remote-Access Module

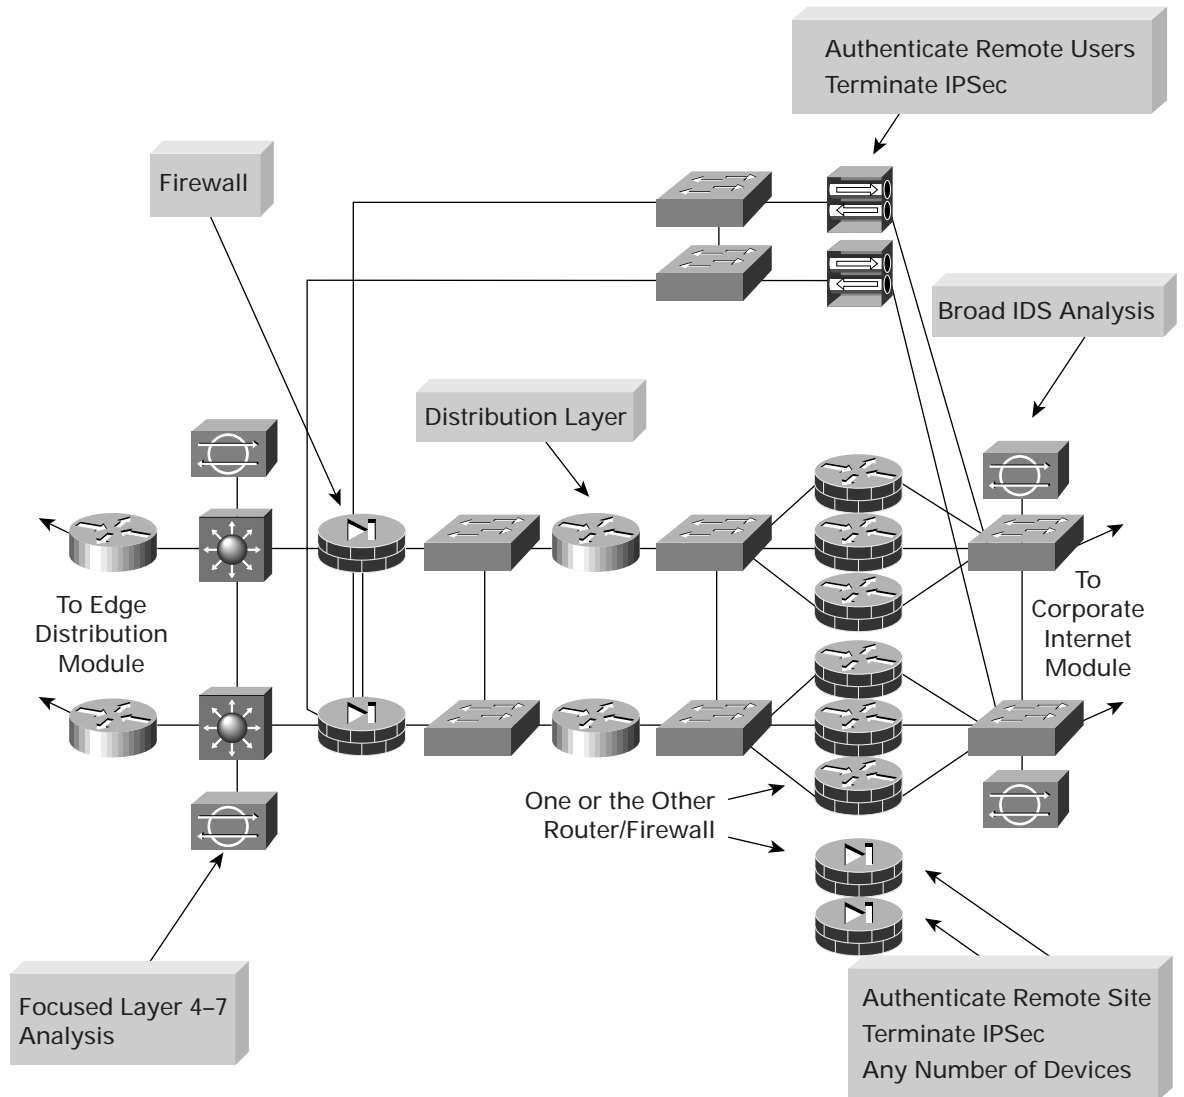


Key VPN Devices

- *Interior firewall*—Provides network-level protection of resources and stateful filtering of traffic
- *Distribution router*—Tracks the availability of remote-site networks across the VPN routers
- *VPN concentrator*—Authenticates individual remote users using XAUTH and terminates their IPsec tunnels
- *VPN router*—Authenticates trusted remote sites and provides connectivity using GRE/IPsec tunnels
- *VPN firewall*—Authenticates trusted remote sites and provides stateful filtering of remote-site traffic
- *NIDS appliance*—Provides Layer 4-to-Layer 7 monitoring of key network segments in the module



Figure 14
Detailed Model of Large Network VPN and Remote-Access Module: VPN



Design Guidelines

The core VPN requirement of this module is to authenticate remote devices and users and terminate IPSec. Multiple VPN termination devices generate gigabit traffic load, driving the need for high-speed Layer 3 switching in the module from the egress of the termination devices to the ingress to the edge distribution module. For this reason, and because of the need for packet classification for differentiated services, the interior routing and distribution layer functionality were carried out using high-speed Layer 3 switching. These factors also drove the requirement of a gigabit-line-rate-capable firewall for stateful inspection and filtering of all traffic in the module and high-speed IDS appliances for attack detection. Because the traffic comes from different sources outside the enterprise network, the



decision was made to provide a separate interface on the firewall for each of these services. The design considerations for the VPN functions of each of these services are addressed below. Extranet VPNs are addressed in the extranet module in a later section of this paper.

Remote-Access VPN

The remote-access VPN traffic is forwarded from the corporate Internet module access routers, where it is first filtered at the egress point to the specific IP addresses and protocols that are part of the VPN segment. These protocols include IKE, ESP, and the UDP/TCP NAT transparency port (e.g. UDP port 10,000 is used throughout the paper). The concentrators in this configuration are configured to support IPSec for tunneling termination only. NAT transparency mode was enabled on the VPN concentrator to facilitate remote-access client connections. Given the high security risk of remote-access clients, split tunneling was not allowed. Once connected, DNS, WINS, and a virtual IP address are pushed to the remote-access clients. After successful authentication, group-level authorization filters are applied to the user's traffic. User-level filters were available for use as well but were not deployed in this module. The concentrators are equipped with hardware acceleration to meet the scalability and performance requirements for a large enterprise. Load balancing was enabled to provide dynamic remote-access client load balancing across the headend. Following termination of the remote-access tunnel by the concentrator, all traffic bound for the Internet or Intranet is sent through the firewall to ensure that VPN user traffic is appropriately inspected, filtered, and logged.

Site-to-Site VPN

The following discussion applies to both VPN router and VPN Firewall termination options. The topology chosen for the VPN was hub and spoke. Spokes used summarizable subnets of the 10.0.0.0/8 network. The site-to-site VPN traffic is forwarded from the corporate Internet module access routers, where it is first filtered at the egress point to the specific IP addresses and protocols that are part of the VPN segment. These protocols include IKE and ESP. Destination IP addresses are limited to the IP addresses of the public interfaces of the headend termination devices, and the source addresses are limited to known static remote-site IP address. Although filtering source addresses provides additional security, this will not be possible if dynamic addressed remote sites are supported. Scalability is a concern. Each static remote site requires four lines of ACLs on each access router (two peers and an ESP and IKE entry for each peer). The VPN devices are configured to support IPSec, and optionally GRE, for tunneling termination. Remote sites in this size of a design must have firewalling capability; otherwise it will be difficult to meet the performance and scalability requirements of the network if all traffic is routed to the hub site.

Remote sites implement firewalling in accordance with the remote, small, and medium network designs documented in the SAFE security papers. The headend VPN devices are equipped with hardware acceleration, given the expected load. Following termination of the site-to-site traffic, all traffic bound for the Intranet is sent through the firewall for stateful inspection, filtering, and logging via a distribution routing layer. This layer exists because in the VPN router option, the site-to-site VPN traffic must be forwarded from the firewall to the VPN router that has reachability to the remote site. Under failure conditions reachability is likely to change. However, because these firewalls do not listen to routing updates, they cannot determine which device has reachability. The distribution routing layer serves the firewall by allowing for packet forwarding to a single IP address via a HSRP interface in ingress thus increasing scalability. On egress, it uses routing protocols to determine remote-site reachability.

VPN Router Option for Site-to-Site VPN

Routing protocol resilience was chosen as the high-availability mechanism for this headend type. The headend VPN routers in this configuration are configured to support IPSec and GRE. Routing protocol updates from the remote sites are advertised over the primary and secondary tunnels to the headend and redistributed on the back end to the



distribution routers. *Enhanced Interior Gateway Routing Protocol (EIGRP)* was chosen over *Open Shortest Path First (OSPF)* because of its lower CPU overhead and, from a headend perspective, only multiple point-to-point links existed. All remote sites are configured to carry out load dispersion in case of failure of a headend VPN router. In any large network with hundreds of nodes, follow Cisco routing protocol best practices.

VPN Firewall Option for Site-to-Site VPN

IKE keepalive resilience was chosen as the high-availability mechanism for this headend type. Remotes are VPN firewalls, VPN routers, and VPN concentrators. The headend VPN firewalls in this configuration are configured to support IPSec only. Each set of VPN firewalls was statically assigned the remote-site networks to trigger outbound tunnel establishment. In case of device failure, the secondary device will use this saved state to allow remote-site connections to continue. Filtering of site-to-site VPN traffic occurs on the interior firewall and not on the VPN firewall in order to provide as much headroom for VPN termination as possible.

Identity

For site-to-site VPN connections to remote locations, digital certificates are used for strong and scalable device authentication. The remote-access VPN connections employ a two-part authentication scheme that uses a group preshared key on the device coupled with a secondary user authentication via OTPs.

Security

Given the high level of access to the corporate Intranet, this module exhibits a high level of security. The stateful firewall software on the VPN firewall allows only IKE and ESP traffic to terminate on the public interface of the VPN firewall. The VPN router option was configured with inbound ACLs to allow only IKE and ESP traffic to terminate on the public interfaces of the VPN routers. The VPN concentrator allows only ESP, IKE, and UDP port 10,000 on its public interface. Filtering that occurs before traffic is passed to the remote-access and VPN module allows only the following traffic flows:

- ESP, IKE, and UDP port 10,000 from any address to the public and virtual cluster IP addresses of the VPN concentrators
- ESP and IKE from known static IP-addressed remote sites to the public IP address of the VPN firewall
- ESP and IKE from known static IP-addressed remote sites to the public IP address of the VPN router

Any flows not listed above will trigger the IDS sensor to high-severity alarm. After decryption, all VPN-sourced traffic is immediately forwarded to the interior firewall, where it is filtered and statefully inspected. While the traffic is being forwarded to the interior router, IDS on that segment performs detailed Layer 4-7 traffic analysis. If IDS detects an attack in a flow, it will shun that flow on the interior firewall. Cisco recommends deploying shunning by the interior IDS because the addresses used in this segment are all private 10.0.0.0 network addresses. The likelihood of a spoofed 10.0.0.0 network address in a VPN is greatly reduced, and even if this were to occur, the remote site or users sending this traffic are easily determined by analyzing the QM SAs.

Scalability

This type of design is extremely scalable. Depending on the remote-site bandwidth requirements, this module can support between 100 and 250 remote-site tunnels per device, possibly even more with very-low-bandwidth remote sites. Because the infrastructure surrounding the VPN devices was designed for the high-speed requirements, this factor was not a limiting one. As with any large-scale VPN design, the primary limiting factor is the number of remote sites that can be terminated given the high-availability mechanism chosen. Remote-access user termination can scale up to 5000 or more of concurrent users.



Secure Management

Secure management of all devices was carried out by using a mix of secure and nonsecure protocols, including SSH, SNMP, TFTP, and syslog. All management in this module occurs via the *out-of-band* (OOB) management network. Remote access VPN clients are securely managed every time they establish a tunnel to the headend and are pushed the latest policy each time.

NAT

NAT was not used in this module. NAT was bypassed for inter-site VPN communications, a scenario that allows all trusted parties to communicate with their real IP address as non-overlapping RFC 1918 addressing was used. NAT was used in the corporate Internet module, however, to address-translate the private addresses used over the VPN by remote-access clients to allow for Internet access as split tunneling was disabled.

Routing

The edge distribution layer is aware of which subnets are used by remote users and sites. The edge distribution routers use a default network to determine reachability to these subnets in the remote-access and VPN module. The default network used was the internal firewall and IDS segment. This way, the edge distribution routers can track the availability of the VPN module via advertisements by its two interior edge routers. The interior edge routers in turn track Intranet reachability via dynamic routing updates and statically route all remote-user and site-directed traffic to the firewall. The interior firewall statically routes remote-access traffic to the VPN concentrator segment and statically routes remote-site traffic to the HSRP virtual address on the distribution routers or VPN Firewalls (whichever is used). The distribution routers run the same routing protocol as the VPN routers terminating the site-to-site VPN traffic and receive updates regarding remote-site network availability. The distribution routers also inject a static route for the major enterprise network into the routing table for redistribution so that remote sites have reachability to the corporate Intranet. Otherwise, reachability to the major network by remote sites would not be possible because routing protocols do not pass through the firewall.

Extranets

Although you could use either the VPN firewalls or VPN routers for extranet termination, it is highly recommended that you do not mix business partner or customer traffic with corporate user or other Intranet traffic. For this reason, an extranet module is described in a following section.

Performance

Large numbers of remote sites and users mandate the use of hardware acceleration at high speed with low latency. Given the bandwidth requirements of the network, high-speed WAN links such as OC-3 (155 Mbps) or greater will be required. Because the remote sites do split tunnel, this setup will alleviate some of the headend bandwidth requirements. However, disallowing the split-tunneling of high-speed DSL/cable users may use significant amounts of data. Care should be taken to not overrun the WAN with VPN traffic.

Alternatives

If the VPN router solution is chosen and the number of devices needed remains small, multiple HSRP groups on the VPN routers could substitute for the distribution layer of routing. However, this solution will not scale as the number of headends grow. If successful personal firewall software deployment occurs as outlined in the axioms, consider enabling remote-access client split tunneling to reduce the performance requirements at the headend. You may consider digital certificates for remote-access user device authentication; however, given the scalability requirements,

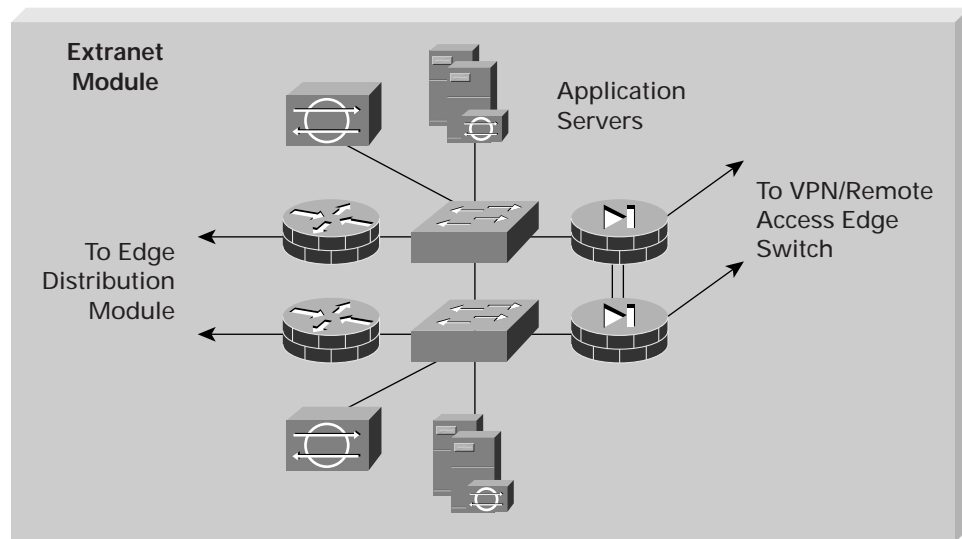


deployment will be extremely difficult. In this design, the same WAN infrastructure is used for Internet access, e-commerce, and VPN. Finally, given the performance requirements for networks, this size of a dedicated VPN WAN infrastructure may be in order. This solution also requires less bandwidth management at the edge because only VPN traffic is routed.

Extranet Module

The extranet module is designed to securely terminate site-to-site and remote-access-based extranets for business-partner access to application servers. Redundant VPN firewalls provide VPN termination, filtering, and stateful firewalling. Both NIDS and *Host IDS* (HIDS) are deployed, given the sensitivity of the data and the users accessing the data that are out of your domain of control.

Figure 15
Detailed Model of Large Extranet Module

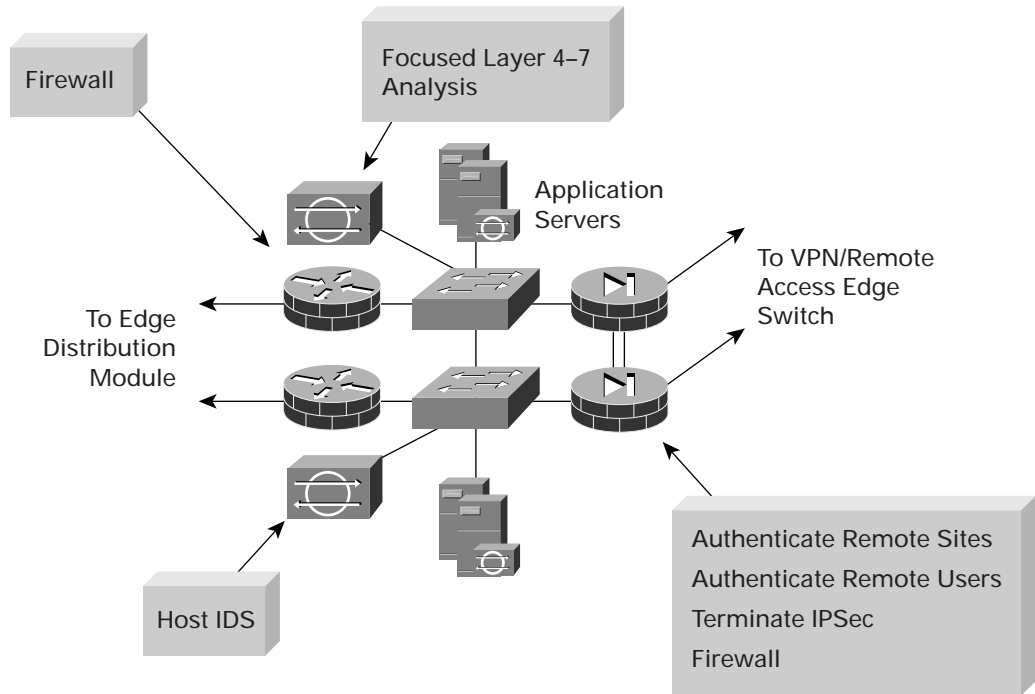


Key VPN Devices

- *Interior firewall router*—Provides access control, stateful firewalling, and access to the corporate Intranet
- *VPN firewall*—Authenticates somewhat trusted remote sites and provides stateful filtering and access control for remote-site traffic
- *NIDS appliance*—Provides Layer 4-to-Layer 7 monitoring of key network segments in the module



Figure 16
Detailed Model of Large Extranet Module: VPN



Design Guidelines

This module must terminate site-to-site and remote-access VPNs in a highly available and secure fashion for mission-critical applications. Security posture is rigid because the remote-site devices terminating on the VPN firewalls are not under administrative control of the enterprise and are granted, sometimes without user authentication, access to application servers one hop away from the corporate Intranet. Guaranteeing high availability without a standards-based high-availability mechanism for IPSec is not possible unless the same vendor's products are used. Therefore, do not consider this module highly available unless this situation exists. Two VPN firewalls are used to guard against their possible device or link failure. Remote VPN devices consist of any IPSec-capable device. The VPN firewalls in this configuration were configured to support IPSec only. IPSec-PFS Group 2 was used in addition to the standard IKE and IPSec policies outlined in the axioms, given the sensitivity of the data and the liability concerns.

Identity

For connections to business partners, a *public-key-infrastructure* (PKI) provider is used to validate the identity of the peer devices. A third-party PKI provider performs as an intermediary between you and your business partners and also provides stronger and more scalable device authentication via the use of digital certificates. Users who come in over the site-to-site VPN do not undergo user authentication. For this reason, strong application level security on the application servers is highly recommended. The remote-access VPN connections employ a two-part authentication scheme that uses a group preshared key on the device coupled with a secondary user authentication via OTPs.



Although OTPs are used, depending on the size of the number of remote users, the associated cost may be high. Cisco does not recommend deploying static username and password pairs. If you do, you should age the passwords aggressively and users should be locked out of the VPN after a certain number of failed login attempts.

Security

Given the high level of access potentially required by partner connections, this module exhibits a high level of security. The stateful firewall software on the VPN firewall allows only IKE and ESP traffic to terminate on its public interface. Filtering in the corporate Internet module allows only ESP and IKE to the public IP address of the VPN firewall. Filtering based on the source IP address is not possible if the IP address of the remote is not static. For remote-access VPN, this is normally the case. Any traffic other than IKE and ESP that is directed to the VPN firewall will trigger the IDS sensors in the remote-access and VPN module to alarm with a high severity. The VPN firewalls should terminate VPN traffic bound only for the local inside subnet. You should implement strict inbound ACLs on the VPN firewall to accomplish this. It is not possible to control remote-device split tunneling unless you provide and control the remote device or software.

If you do not configure any routes to the Intranet on the VPN firewall, it will not forward packets anywhere other than the locally connected network, the OOB management network, or the corporate Internet module edge routers. Packets can enter the Intranet only via the application servers. The interior firewall routers have access control to allow only the IP addresses of the application servers, and the services that they are allowed to use, through to the Intranet. The application servers utilize HIDS for local attack mitigation to guard against being compromised and possibly providing a hacker access to the Intranet. If the NIDS deployed on the application-server subnet detects an attack, it will alarm with high severity.

Scalability and Performance

This module was designed for moderate to high extranet use. It uses hardware acceleration to provide high-speed, low-latency VPNs. The extranet module will support more than 200 remote sites and 500 remote concurrent users. If greater numbers are required, you should add more VPN devices. This addition will require a routing layer to provide a default route for the application servers. Because the devices are supporting only the extranet application, high performance should result.

Secure Management

Using a mix of secure and nonsecure protocols, including SSH, SNMP, TFTP, and syslog, ensures secure management of all the devices. All management in this module occurs via the OOB management network.

NAT

Use NAT in this module on the VPN firewall only when the address space of the connecting remote site overlaps with that of the application servers.

Routing

Dynamic routing protocols are used by the interior firewall routers to advertise the application network to the edge distribution routers. Routes on the VPN firewall should not exist for any internal network other than the OOB management network segment.



Alternatives

The primary alternative is to consider the large design discussed previously if this design does not meet the scalability requirements for an Extranet. Also consider that since this module resembles the remote-access and VPN module, some integration could occur. You might consider terminating partner and internal user VPN traffic on the same device, but this configuration is not recommended. Adding another interface to the firewall in the remote-access and VPN module to segment the traffic is much more secure.

Another alternative is to use SSL and to deploy a different design altogether. IPSec is used under the assumption that “always on” network-to-network access is a requirement. In this design, multiple hosts of a partner’s network could access the application segment. If single client or application-to-application access is all that is necessary, consider SSL.

In this design digital certificates would provide much stronger device authentication for remote sites and users. As stated in the axiom, however, deployment is a significant administrator burden. User authentication for site-to-site traffic could be implemented, in addition to device authentication; this setup would, however, require the user to first connect using a username/password prompt-capable application such as Telnet or Web. This scenario is likely to be no more secure than the application level security—unless OTPs are used—and OTPs will increase the cost. Also, you might consider deploying a purpose-defined Firewall on the interior edge given the rigid security stance.

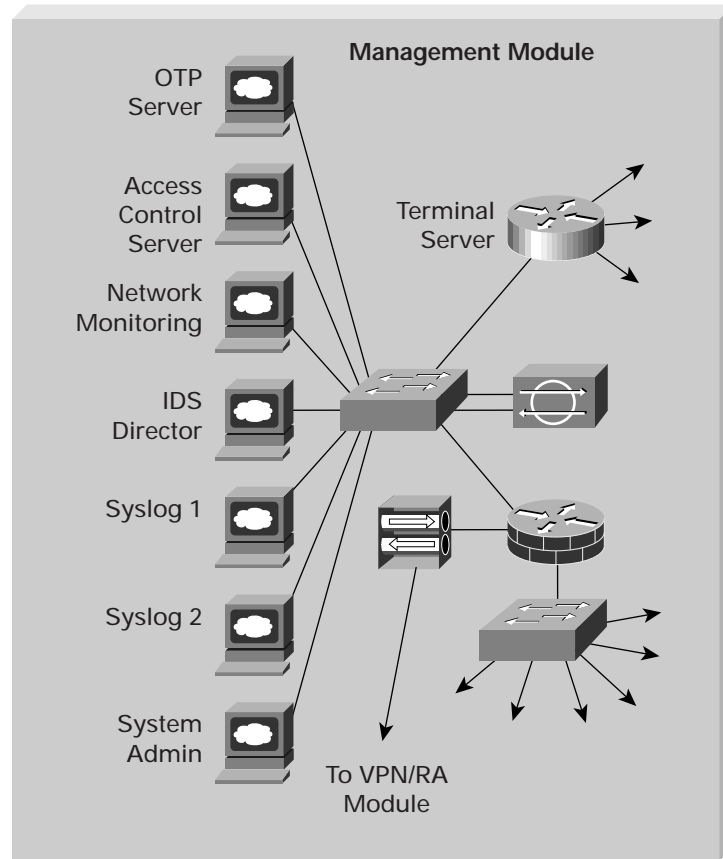
Alternatively, use of VPN routers could replace the VPN firewalls. Since VPN routers support multiple digital certificate identities, you will be able to use both an in-house CA and a third-party or business partner CA simultaneously. The first identity is used for identifying the VPN router as a device that exists in your network infrastructure, and the second identity is used for device authentication with the partner device.

Management Module

The primary goal of the management module is to facilitate the secure management of all devices and hosts within the enterprise SAFE architecture.



Figure 17
Detailed Model of Large Management Module

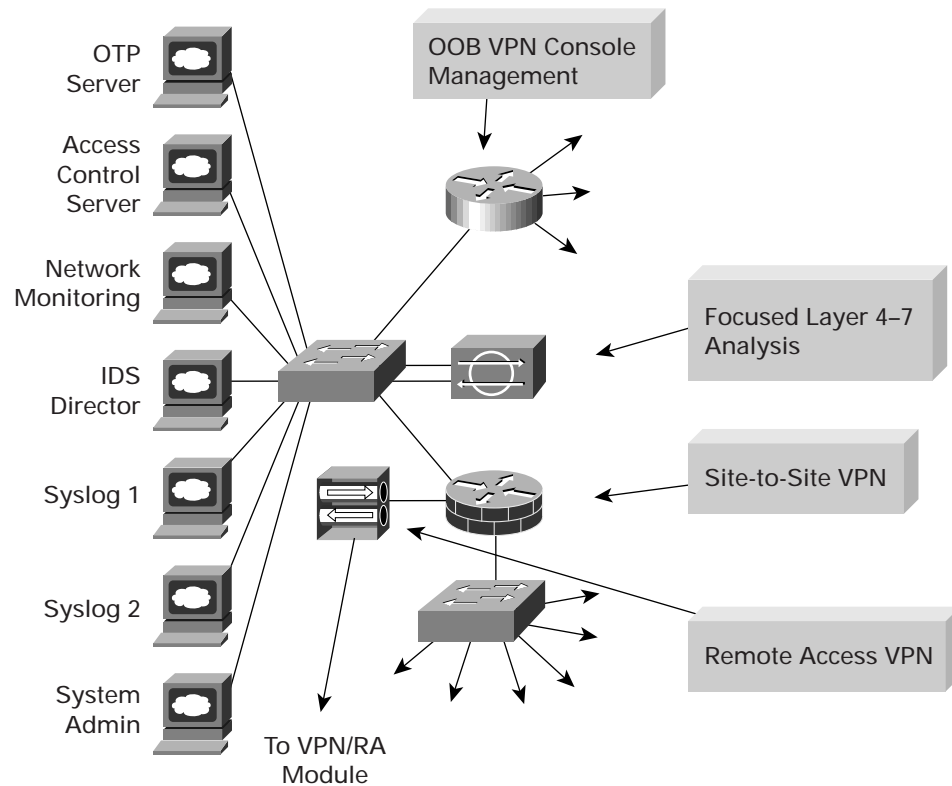


Key VPN Devices

- *Firewall and VPN router*—Authenticates trusted remote sites and provides stateful filtering and access control for OOB management
- *VPN concentrator*—Authenticates remote administrators and terminates their IPSec tunnels
- *NIDS appliance*—Provides Layer 4-to-Layer 7 monitoring of key network segments in the module and VPN terminated traffic



Figure 18
Detailed Model of Large Management Module: VPN



Design Guidelines

The module remains virtually the same as it did in the original SAFE security enterprise paper. Only the changes are discussed. Please refer to the SAFE security papers for information not covered here. The primary change constitutes adding a VPN concentrator to facilitate secure remote access for administrators to the management network. There is a high risk potential, specifically compromising of the network, if this VPN application is not implemented correctly. After all, this is the module that has access to every system in the entire enterprise network. With this modification, it is now connected to the Internet via a VPN concentrator. Therefore, every precaution was taken to implement this design in a secure fashion.

NAT

NAT transparency mode was enabled on the VPN concentrator to facilitate remote-access client connections. NAT was used on the Firewall router to provide the management segment secure access to the Internet. No other forms of NAT were used.

Identity

The remote-access VPN connections employ a two-part authentication scheme that uses a digital certificate for device authentication coupled with a secondary user authentication via OTPs. This design is the only one in the SAFE VPN architecture in which strong device authentication and strong user authentication are required simultaneously.



Security

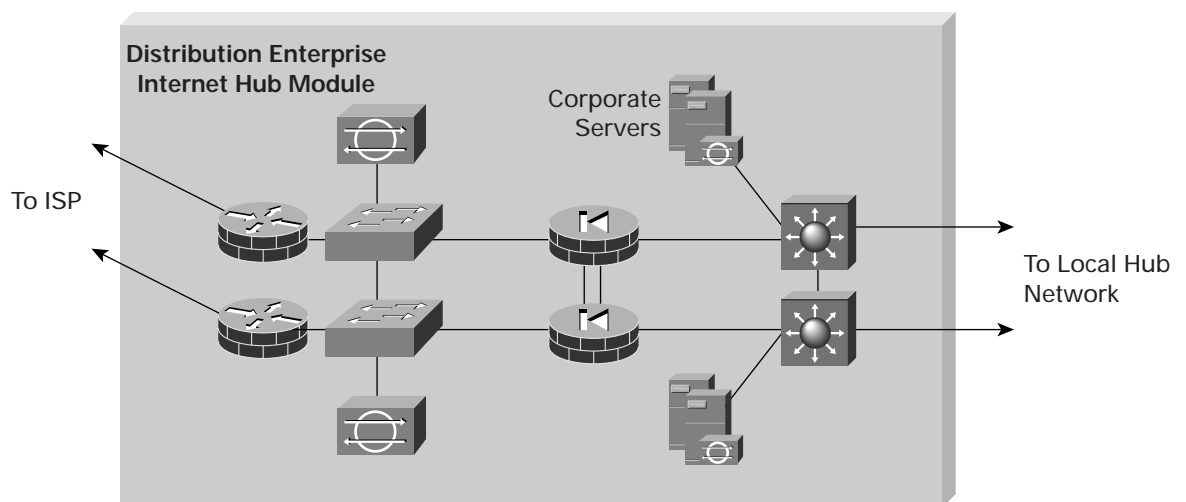
The VPN concentrator was configured to support IPSec only. The public interface of the VPN Concentrator will accept only IKE, ESP, and UDP port 10,000 packets. The corporate Internet module edge router filters out any traffic other than IKE, ESP, and UDP port 10,000 packets destined to the management VPN concentrator. The edge NIDS in the remote-access and VPN module will trigger an alarm with high severity if any traffic other than IKE, ESP, or UDP 10,000 is directed to the management VPN concentrator. There is a high possibility that configurations with preshared keys and other sensitive information will traverse the tunnel. Knowing this, a hacker could use typical information found in management traffic to brute-force attack the encrypted packets. Therefore, PFS Group 2 was mandated in this module. Split tunneling should not be allowed by remote administrators unless personal firewall software has been successfully deployed or the management network has been granted Internet access via a stateful firewall. Even then, you should consider the possible implications.

After successful authentication, per-administrator filters were applied to the tunneled traffic. Administrators should have access only to the applications in the management network that they need to use. The private interface filter of the VPN concentrator's permits only management traffic. This setup helps to mitigate the possibility of misconfiguration of the user authorization privileges that might allow a remote user to tunnel nonmanagement traffic. After packets are decrypted and user and interface filters are applied, the traffic is forwarded to the firewall router for stateful inspection and another round of filtering. Only after this layered security is traversed will the remote administrator traffic have access to the devices on the management network. While accessing devices, the NIDS in the management module will shun any traffic from the VPN concentrator that appears to be an attack or is not management traffic.

Distribution-Hub Module

The distribution-hub module is designed to securely terminate site-to-site-based VPNs and provide an intermediate layer between VPN-enabled small spoke sites and the large enterprise remote-access and VPN module headend. This layer allows for local small spoke sites to intercommunicate without having to send traffic to the enterprise hub site.

Figure 19
Detailed Model of Distribution VPN Module



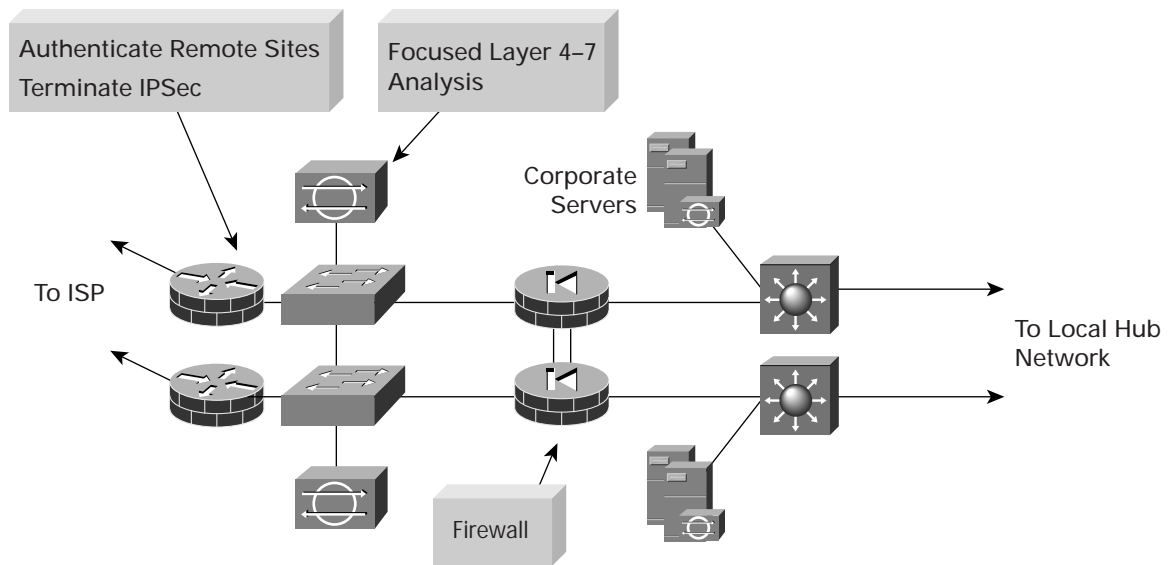


Key VPN Devices

- *VPN router*—Authenticates trusted remote sites and provides stateful filtering and access control for remote-site traffic
- *Firewall*—Provides stateful firewalling and filtering
- *NIDS appliance*—Provides Layer 4-to-Layer 7 monitoring of key network segments in the module

Figure 20

Detailed Model of Distribution VPN Module: VPN



Design Guidelines

The VPN routers terminate traffic from the remote sites and enterprise headend. Most of the traffic is spoke-to-spoke traffic. Firewalls provide stateful firewalling and filtering for all traffic between the remote sites and the local services. NIDS was deployed to detect attack signatures from the remote sites and local services in case they are compromised. Services usually include databases and applications such as mail. The following assumptions were made:

- Spoke high availability to the distribution hub was required.
- Distribution layer high availability to the headend hub was required.
- Multiple distribution hubs may exist, and intercommunication was needed via the headend.
- Spoke-to-spoke intercommunication was needed via the distribution hub.

For the above reasons, VPN routers are needed at the headend and distribution-hub sites to route packets between remote sites. Spokes could be any VPN device. The distribution-hub VPN routers track enterprise headend reachability by the use of routing protocols over the two tunnels to the headend. EIGRP was chosen over OSPF because of its lower CPU overhead. The distribution routers track link status by using two HSRP groups. Two groups are necessary as each VPN router serves as a primary to two disparate sets of remote sites. After failure of a VPN router, the secondary VPN router will take over assuming the load from the secondary group of remote sites. As the tiered firewalls do not listen to routing updates, this allows them to forward packets to the active VPN router that has reachability to the remote site.



Spoke and tiered local networks should use summarizable subnets of the enterprise major network. Split tunneling at the tiered site was not required because it was assumed that both the headend and spoke sites have this capability. If the remote does not, this configuration will put increased performance and scalability requirements on the headend and tiered site. This module assumes dual VPN routers at the distribution layer. If more than two devices are required to meet the loading requirements of the network at any layer, you should deploy load dispersion on failure.

Identity

For site-to-site VPN connections to remote locations, digital certificates are used for strong and scalable device authentication.

Security

It should be noted that no IDS or firewalling occurs when traffic travels from spoke to spoke or spoke to headend. Because firewall and IDS facilities exist in the headend, there is no reason to carry out this function twice for spoke to headend traffic. It is assumed that the remotes support split tunneling. Stateful firewalling and filtering does occur when any of the remote sites access the local services. Filtering inbound on the public interfaces of the VPN routers allows ESP and IKE only from known static IP-addressed remote sites and headends.

Scalability

This module was designed to increase the scalability of large enterprise hub-and-spoke networks. This distribution layer has the capability to scale up to 200 or more remote sites.

Secure Management

Secure management of the all devices was carried out by using a mix of secure and nonsecure protocols, including SSH, SNMP, TFTP, and syslog. All management in this module occurs via in-band VPN tunnels to the management module in the headend.

NAT

NAT was only used in this module to allow local servers Internet access. All other addressing handled was private to the VPN and thus NAT was not necessary.

Routing

The VPN routers have a default route to the Internet. This route encompasses all public IKE peers and remote and headend-site private networks. Only the local networks are statically routed to the firewall pair. The local tiered firewall forwards half of the remote site traffic to each virtual HSHP address. Traffic destined to the headend is forwarded to the HSRP groups and load balanced based on IP address.

Performance

Numerous spokes and the need for spoke-to-spoke intercommunication drive the requirement for high levels of VPN device performance at low latencies. Given the bandwidth capacity of these devices and the number of remote sites, a high-speed WAN link such as a DS3 (45 Mbps) or greater is required.



Alternatives

If the remote sites do not allow split-tunneling and a distribution layer needs to perform this function, consider the approach to split-tunneling as outlined in the small- and medium-business designs. The existing distribution-hub module would not be appropriate for split-tunneling because it would mix decrypted and clear Internet traffic on the same wire between the firewalls and VPN routers. If more than two VPN routers are needed to meet performance needs, add more VPN routers and consider adding a WAN routing layer in front of the existing VPN router segment. Note that this is a similar design to the large enterprise headend and, therefore, load dispersion of failure should be configured to aid in scalability. In order to accomplish this routing, protocols should be used in place of multiple HSRP groups to aid in scalability.

Migration Strategies

SAFE VPN is a guide for implementing VPNs. The designs proposed here are not meant to serve as the all-encompassing designs for providing VPNs for all existing networks. Rather, SAFE VPN is a template that enables network designers to consider how they design and implement their enterprise VPN in order to meet their security and connectivity requirements.

Cataloging of the number of remote users and sites, and their associated performance, application, and reliability requirements should be the first activity in migrating the existing network to a VPN. The next activity is to design a network that supports these requirements in a secure and scalable fashion. Basic recommendations for VPN-specific security considerations and VPN deployment are referenced in depth in the axioms of this document. After the effect of applying these considerations to the existing network is determined, the network designer should consider how these may be applied to the existing network infrastructure.

The architecture has enough flexibility to enable SAFE VPN to be adapted to most networks. SAFE VPN allows the designer to address the VPN application to each network function, almost independently of each other. Each module is generally self-contained and can connect to any other VPN-enabled module in the architecture. Because one application of VPNs is a replacement technology for network access, they can be deployed in parallel with already existing private networks. Single-purpose VPN devices simplify the migration by segmenting the new functionality from existing infrastructure thus reducing the possibility of adversely effecting the network. When ready for production, a simple routing cutover can occur.

This is the third white paper detailing the specifics of the SAFE architecture. When this paper is combined with SAFE Enterprise and SAFE SMB, the sum of the documents address the VPN and security design aspects 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 VPN management analysis and implementation
- QoS and voice analysis and implementation
- VPN and DDR coexistence; VPNs as a backup to the traditional WAN
- In-depth identity, directory services, AAA technologies, and CA analysis and implementation considerations
- In-depth campus and wireless VPN design, management, and implementation considerations



Appendix A: Validation Lab

A reference SAFE VPN implementation exists to validate the functionality described in this document. This appendix details the configurations of the specific devices as they relate to VPN functionality within each module. It 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. Note that access control for VPN devices to ensure only IKE (UDP port 500) and ESP (IP protocol 50) are not shown in this document; for these configurations, please refer to the SAFE security papers.

Overall Guidelines

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



SAFE VPN Standard Configuration Parameters

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

```
!  
! SAFE VPN standard IKE policy  
!  
crypto isakmp policy 10  
  encr 3des  
  authentication rsa-sig  
  group 2  
  hash sha  
!  
! Typical CA identity  
!  
crypto ca identity safevpn  
  enrollment mode ra  
  enrollment url http://172.16.128.50:80/certsrv/mscep/mscep.dll  
!  
!  
!  
! High entropy and unique per-address pre-shared key  
!  
crypto isakmp key 7Q!r$y$+xE address 172.16.144.3  
!  
! Digital Certificates (CA,RA, encryption and signing)  
!  
crypto ca certificate chain safevpn  
certificate 613500AA000000000007  
  308203EB 30820395 A0030201 02020A61 3500AA00 00000000 07300D06 092A8648  
  86F70D01 01050500 306C310B 30090603 55040613 02555331 0B300906 03550408  
  13024341 3111300F 06035504 07130853 616E204A 6F736531 1B301906 0355040A  
  13124369 73636F20 53797374 656D732C 20496E63 310D300B 06035504 0B130456  
  53454331 11300F06 03550403 1308496E 7465726E 6574301E 170D3031 30363230  
  32333230 35385A17 0D303230 36323032 33333035 385A3027 31253023 06092A86  
  4886F70D 01090213 16523236 32312D31 2E736166 652D736D 616C6C2E 636F6D30  
  819F300D 06092A86 4886F70D 01010105 0003818D 00308189 02818100 9F88ECFF  
  D3213656 C027D6AC 08076401 16D75A25 643E8881 CA2BA5B1 6215C0A7 9C80C831  
  6A469DB5 72B1A530 72492649 D42812B2 AB26E536 61CEFFFE 468CADE4 A9A498F2  
  F2E134F4 F2780C81 B5C1B1CB 45EE5DBD 336F5842 954F37CE E81FACCD 384CB388  
  141BD5E1 1015DB15 AF2BDD5F 2D67BB19 D9708F68 58A99A8E 5DEC20F1 02030100  
  01A38202 18308202 14300B06 03551D0F 04040302 05A0301D 0603551D 0E041604  
  1442A32E 697145AF 42211881 2396DA0B 96C39C74 003081A5 0603551D 2304819D  
  30819A80 14BFB4E5 C7D8D6B6 55FCC1CB 6F5B2C48 1C1C1A34 02A170A4 6E306C31  
  0B300906 03550406 13025553 310B3009 06035504 08130243 41311130 0F060355  
  04071308 53616E20 4A6F7365 311B3019 06035504 0A131243 6973636F 20537973  
  74656D73 2C20496E 63310D30 0B060355 040B1304 56534543 3111300F 06035504  
  03130849 6E746572 6E657482 102E9E46 057DECA2 8C42F3BF 9C90639D 3F302406  
  03551D11 0101FF04 1A301882 16523236 32312D31 2E736166 652D736D 616C6C2E  
  636F6D30 73060355 1D1F046C 306A3032 A030A02E 862C6874 74703A2F 2F696E74  
  65726E65 742D6D73 61632F43 65727445 6E726F6C 6C2F496E 7465726E 65742E63  
  726C3034 A032A030 862E6669 6C653A2F 2F5C5C69 6E746572 6E65742D 6D736163  
  5C436572 74456E72 6F6C6C5C 496E7465 726E6574 2E63726C 3081A206 082B0601  
  05050701 01048195 30819230 4606082B 06010505 07300286 3A687474 703A2F2F  
  696E7465 726E6574 2D6D7361 632F4365 7274456E 726F6C6C 2F696E74 65726E65  
  742D6D73 61635F49 6E746572 6E65742E 63727430 4806082B 06010505 07300286  
  3C66696C 653A2F2F 5C5C696E 7465726E 65742D6D 7361635C 43657274 456E726F  
  6C6C5C69 6E746572 6E65742D 6D736163 5F496E74 65726E65 742E6372 74300D06  
  092A8648 86F70D01 01050500 03410024 8B79077E 37C7C8EA 1C53FAAB 92264274
```



```
1E875C7A 809618B0 5A5B1719 5F4FC690 9B8D8320 14ACD7DD 0F8035F8 CA18644D
79588D7F 156F4EA4 805952FA B39EC8
quit
certificate ra-sign 47364E9E0000000000002
30820456 30820400 A0030201 02020A47 364E9E00 00000000 02300D06 092A8648
86F70D01 01050500 306C310B 30090603 55040613 02555331 0B300906 03550408
13024341 3111300F 06035504 07130853 616E204A 6F736531 1B301906 0355040A
13124369 73636F20 53797374 656D732C 20496E63 310D300B 06035504 0B130456
53454331 11300F06 03550403 1308496E 7465726E 6574301E 170D3031 30363230
32323232 31345A17 0D303230 36323032 32333231 345A3081 9D312030 1E06092A
864886F7 0D010901 16116164 6D696E40 73616665 76706E2E 636F6D31 0B300906
03550406 13025553 310B3009 06035504 08130243 41311330 11060355 0407130A
54686520 56616C6C 65793116 30140603 55040A13 0D534146 45205650 4E20496E
632E311C 301A0603 55040B13 13534146 45205650 4E204272 61696E74 72757374
31143012 06035504 03130B53 41464520 56504E20 43413081 9F300D06 092A8648
86F70D01 01010500 03818D00 30818902 818100BB 73EF0897 57DFDC7C 0F72482D
39EE9562 E1291155 AB6F5627 338BE5A7 BF2F2904 BD643F7A 63EF9EDB 75ED8C44
D006503E 1A88D16D 45AF4E31 E5B01EBE 1BC829E0 4A5A3701 E3CC67B5 270BB2DE
80561B60 96732CD6 D6FF7601 4920A82B AD8A0EF0 F3AA24D2 D5D2D64F 21EDB990
3E51A64F 3C1DCBA6 94AA6B3F 21DED3BC FB392102 03010001 A382020C 30820208
300E0603 551D0F01 01FF0404 030206C0 30150603 551D2504 0E300C06 0A2B0601
04018237 14020130 1D060355 1D0E0416 0414726C 201817B0 D5C56A58 DEEB6024
5FF2DED6 B6BE3081 A5060355 1D230481 9D30819A 8014BFB4 E5C7D8D6 B655FCC1
CB6F5B2C 481C1C1A 3402A170 A46E306C 310B3009 06035504 06130255 53310B30
09060355 04081302 43413111 300F0603 55040713 0853616E 204A6F73 65311B30
19060355 040A1312 43697363 6F205379 7374656D 732C2049 6E63310D 300B0603
55040B13 04565345 43311130 0F060355 04031308 496E7465 726E6574 82102E9E
46057DEC A28C42F3 BF9C9063 9D3F3073 0603551D 1F046C30 6A3032A0 30A02E86
2C687474 703A2F2F 696E7465 726E6574 2D6D7361 632F4365 7274456E 726F6C6C
2F496E74 65726E65 742E6372 6C3034A0 32A03086 2E66696C 653A2F2F 5C5C696E
7465726E 65742D6D 7361635C 43657274 456E726F 6C6C5C49 6E746572 6E65742E
63726C30 81A20608 2B060105 05070101 04819530 81923046 06082B06 01050507
3002863A 68747470 3A2F2F69 6E746572 6E65742D 6D736163 2F436572 74456E72
6F6C6C2F 696E7465 726E6574 2D6D7361 635F496E 7465726E 65742E63 72743048
06082B06 01050507 3002863C 66696C65 3A2F2F5C 5C696E74 65726E65 742D6D73
61635C43 65727445 6E726F6C 6C5C696E 7465726E 65742D6D 7361635F 496E7465
726E6574 2E637274 300D0609 2A864886 F70D0101 05050003 41006094 1ACE2B69
CF9729A0 6325E48B 2C2D1C21 493748A6 9CCAD32E F1CE0C5E A1C51CD0 3C5404A6
AD6CACF6 4381884D 1128C62A CD8C6DB5 76D205BA 68FDC8E8 86CA
quit
certificate ca 2E9E46057DECA28C42F3BF9C90639D3F
308202A0 3082024A A0030201 0202102E 9E46057D ECA28C42 F3BF9C90 639D3F30
0D06092A 864886F7 0D010105 0500306C 310B3009 06035504 06130255 53310B30
09060355 04081302 43413111 300F0603 55040713 0853616E 204A6F73 65311B30
19060355 040A1312 43697363 6F205379 7374656D 732C2049 6E63310D 300B0603
55040B13 04565345 43311130 0F060355 04031308 496E7465 726E6574 301E170D
30313036 30363139 35373538 5A170D30 33303630 36323030 3632355A 306C310B
30090603 55040613 02555331 0B300906 03550408 13024341 3111300F 06035504
07130853 616E204A 6F736531 1B301906 0355040A 13124369 73636F20 53797374
656D732C 20496E63 310D300B 06035504 0B130456 53454331 11300F06 03550403
1308496E 7465726E 6574305C 300D0609 2A864886 F70D0101 01050003 4B003048
024100BF 6EDF974C 2BAB7BD1 7146096A 11413145 663A67F9 3B8893B0 585F188E
41CBBEB4 24C2C154 EAC65101 CF43AC28 D970A6CF 4448E9E2 CA0B7288 76AF561C
871B4102 03010001 A381C730 81C4300B 0603551D 0F040403 0201C630 0F060355
1D130101 FF040530 030101FF 301D0603 551D0E04 160414BF B4E5C7D8 D6B655FC
C1CB6F5B 2C481C1C 1A340230 73060355 1D1F046C 306A3032 A030A02E 862C6874
74703A2F 2F696E74 65726E65 742D6D73 61632F43 65727445 6E726F6C 6C2F496E
7465726E 65742E63 726C3034 A032A030 862E6669 6C653A2F 2F5C5C69 6E746572
```



```
6E65742D 6D736163 5C436572 74456E72 6F6C6C5C 496E7465 726E6574 2E63726C
30100609 2B060104 01823715 01040302 0100300D 06092A86 4886F70D 01010505
00034100 7E74B2F7 15D185EC 5C89DE9C 0F0E0E12 6F90397F 4AAE9E26 6D0025F6
F0A06935 F3D842F6 98689B35 FDF175F7 8CBDDEE6 6201B69A 415624A5 6D130AEE
ACA5B1F1 certificate ra-encrypt 47364F8800000000000003
30820456 30820400 A0030201 02020A47 364F8800 00000000 03300D06 092A8648
86F70D01 01050500 306C310B 30090603 55040613 02555331 0B300906 03550408
13024341 3111300F 06035504 07130853 616E204A 6F736531 1B301906 0355040A
13124369 73636F20 53797374 656D732C 20496E63 310D300B 06035504 0B130456
53454331 11300F06 03550403 1308496E 7465726E 6574301E 170D3031 30363230
32323232 31345A17 0D303230 36323032 32333231 345A3081 9D312030 1E06092A
864886F7 0D010901 16116164 6D696E40 73616665 76706E2E 636F6D31 0B300906
03550406 13025553 310B3009 06035504 08130243 41311330 11060355 0407130A
54686520 56616C6C 65793116 30140603 55040A13 0D534146 45205650 4E20496E
632E311C 301A0603 55040B13 13534146 45205650 4E204272 61696E74 72757374
31143012 06035504 03130B53 41464520 56504E20 43413081 9F300D06 092A8648
86F70D01 01010500 03818D00 30818902 818100BC 14978E5B 9522DF96 E75DB97B
2556553C 9D9E78C6 A1B634CF D49D05A8 C45D9483 E5EC53F4 6FBA51AC 186FA67C
F2320FE4 B6BDA64C 28D1E646 5298A5BC 968132AD 222D99BE 76EB1EC7 BC8076ED
88F44D24 8F9A24FF E2161187 4CFA012F 5E309430 286D77FF 6A920E61 C8325711
1FFB19D3 51EB83C3 6157DA98 3F25104B EF62BF02 03010001 A382020C 30820208
300E0603 551D0F01 01FF0404 03020430 30150603 551D2504 0E300C06 0A2B0601
04018237 14020130 1D060355 1D0E0416 04146BDB 14526D6C 833A6F9B A033FE55
6B2D3E80 84723081 A5060355 1D230481 9D30819A 8014BFB4 E5C7D8D6 B655FCC1
CB6F5B2C 481C1C1A 3402A170 A46E306C 310B3009 06035504 06130255 53310B30
09060355 04081302 43413111 300F0603 55040713 0853616E 204A6F73 65311B30
19060355 040A1312 43697363 6F205379 7374656D 732C2049 6E63310D 300B0603
55040B13 04565345 43311130 0F060355 04031308 496E7465 726E6574 82102E9E
46057DEC A28C42F3 BF9C9063 9D3F3073 0603551D 1F046C30 6A3032A0 30A02E86
2C687474 703A2F2F 696E7465 726E6574 2D6D7361 632F4365 7274456E 726F6C6C
2F496E74 65726E65 742E6372 6C3034A0 32A03086 2E66696C 653A2F2F 5C5C696E
7465726E 65742D6D 7361635C 43657274 456E726F 6C6C5C49 6E746572 6E65742E
63726C30 81A20608 2B060105 05070101 04819530 81923046 06082B06 01050507
3002863A 68747470 3A2F2F69 6E746572 6E65742D 6D736163 2F436572 74456E72
6F6C6C2F 696E7465 726E6574 2D6D7361 635F496E 7465726E 65742E63 72743048
06082B06 01050507 3002863C 66696C65 3A2F2F5C 5C696E74 65726E65 742D6D73
61635C43 65727445 6E726F6C 6C5C696E 7465726E 65742D6D 7361635F 496E7465
726E6574 2E637274 300D0609 2A864886 F70D0101 05050003 4100A0ED 49063B8B
320DCEC8 F2FC6A7A 0D5F6BE3 C1772559 FE914CE8 C681C685 B7D6B4F7 4785383A
98E8E280 0BAEB8C9 499F44FB 014FB4C2 DB4AEAAD 4B2B82E0 35A6
quit
!
! SAFE VPN standard IPsec SA transform set
!
crypto ipsec transform-set strong esp-3des esp-sha-hmac
!
! Example crypto ACLs for remote management
!
access-list 150 permit ip host 172.16.128.2 host 172.16.144.51
!
! crypto map definition for remote management traffic
!
crypto map main_map 150 ipsec-isakmp
 set peer 172.16.128.2
 set transform-set strong
 match address 150
!
! CEF is enabled on all VPN Devices as this is the optimized path
```

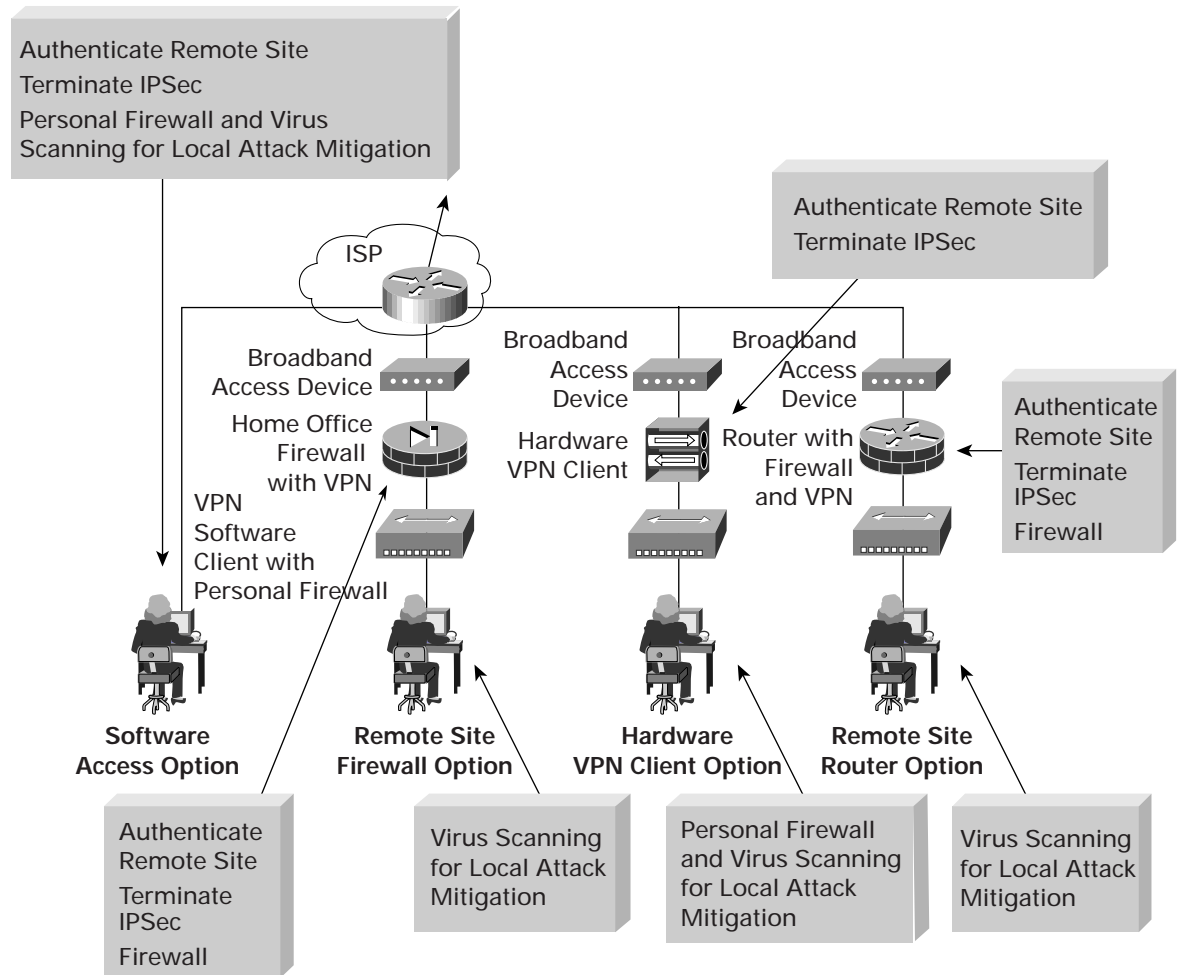


```
!  
ip cef  
!  
The following are sample commands that enable most of the basic  
configuration options present on all VPN firewalls in the SAFE VPN lab.  
!  
! IKE must be explicitly enabled on VPN Firewalls  
!  
isakmp enable outside  
!  
! SAFE VPN standard IKE policy  
!  
isakmp policy 10 authentication rsa-sig  
isakmp policy 10 encryption 3des  
isakmp policy 10 hash sha  
isakmp policy 10 group 2  
!  
! Typical CA identity  
!  
ca identity safevpn 172.16.128.50:/certsrv/mscep/mscep.dll  
ca configure safevpn ra 1 20  
!  
! High entropy and unique per-address pre-shared key, hidden  
!  
isakmp key ***** address 172.16.144.3 netmask 255.255.255.255  
!  
! SAFE VPN standard IPsec SA transform set  
!  
crypto ipsec transform-set strong esp-3des esp-sha-hmac  
!  
! Example crypto ACLs for remote management  
!  
access-list 103 permit ip host 172.16.128.5 host 172.16.144.51  
!  
! Crypto map definition for remote management traffic  
!  
crypto map main_map 150 ipsec-isakmp  
crypto map main_map 150 match address 103  
crypto map main_map 150 set peer 172.16.144.3  
crypto map main_map 150 set transform-set strong  
!
```



Remote-User Design Module Configurations

Figure 21
Attack Mitigation Roles for Remote User Networks



Products Used

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



rIOS-1

The following configuration example is for a remote broadband VPN router connecting to the small enterprise.

```
! Crypto ACL for local network to head-end
!
access-list 103 permit ip 10.5.0.0 0.0.0.255 10.0.0.0 0.255.255.255
!
! Single crypto map entry, no HA
!
crypto map main_map 10 ipsec-isakmp
  set peer 172.16.144.3
  set transform-set strong
  match address 103
!
! Crypto map attached to public interface, we used a LAN interface although
a broadband interface could have been used. 3des-sha encryption used.
!
interface FastEthernet0/1
  ip address 172.16.128.2 255.255.255.0
  crypto map main_map
!
interface FastEthernet0/0
  ip address 10.5.1.2 255.255.255.0
!
! Default route triggers all traffic to hit crypto ACL
!
ip route 0.0.0.0 0.0.0.0 172.16.128.1
!
```



rPIX-1

The following configuration example is for a remote VPN firewall connecting to the small enterprise.

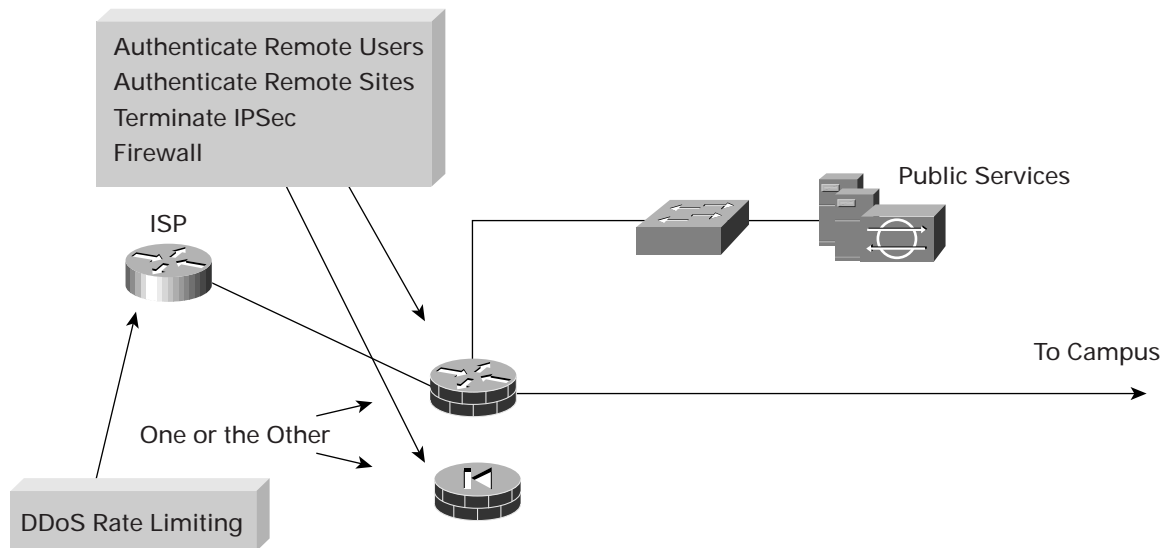
```
!  
! Crypto ACL for local network to head-end  
!  
access-list 101 permit ip 10.6.0.0 255.255.255.0 10.0.0.0 255.0.0.0  
!  
ip address outside 172.16.128.5 255.255.255.0  
ip address inside 10.6.1.1 255.255.255.0  
!  
! Default route to forward packets to IKE peers and remote VPNs  
!  
route outside 0.0.0.0 0.0.0.0 172.16.128.3 1  
!  
! Crypto map attached to public interface, 3des-sha encryption used.  
!  
crypto ipsec transform-set 3dessha esp-3des esp-sha-hmac  
crypto map remotel 20 ipsec-isakmp  
crypto map remotel 20 match address 101  
crypto map remotel 20 set peer 172.16.144.3  
crypto map remotel 20 set transform-set strong  
crypto map remotel interface outside  
!  
! Bypass NAT engine on firewall for traffic bound for VPN  
!  
access-list nonat permit ip 10.6.0.0 255.255.255.0 10.0.0.0 255.0.0.0  
access-list nonat deny ip 10.6.0.0 255.255.255.0 any  
!  
!nat (inside) 0 access-list nonat  
!
```



Small Enterprise Configurations

Caveat: These configurations do not show remote access VPN termination.

Figure 22
Small Business Corporate Internet Module VPN



Products Used

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



SIOS-1

The following configuration example is for the VPN router option in the small enterprise serving as a headend for remotes.

```
!  
! Crypto ACLs for local networks to remote sites  
!  
access-list 107 permit ip 10.4.0.0 255.255.0.0 10.5.0.0 255.255.255.0  
!  
access-list 108 permit ip 10.4.0.0 255.255.0.0 10.6.0.0 255.255.255.0  
!  
!ip address outside 172.16.144.3 255.255.255.0  
ip address inside 10.4.1.1 255.255.255.0  
!  
! Default route to forward packets to IKE peers and remote VPNs  
!  
route outside 0.0.0.0 0.0.0.0 172.16.144.2 1  
!  
! Bypass firewall access-control for traffic inbound on outside interface  
!  
crypto map main_map 30 ipsec-isakmp  
crypto map main_map 30 match address 107  
crypto map main_map 30 set peer 172.16.128.2  
crypto map main_map 30 set transform-set main_map  
!  
crypto map main_map 40 ipsec-isakmp  
crypto map main_map 40 match address 108  
crypto map main_map 40 set peer 172.16.128.5  
crypto map main_map 40 set transform-set main_map  
crypto map main_map interface outside  
!  
! Dynamic crypto map entry for remote access clients (note: no set peer or ACL)  
!  
crypto dynamic-map vpnuser 20 set transform-set main_map  
crypto map main_map 50 ipsec-isakmp dynamic vpnuser  
!  
! Commands needed to enable MODCFG and XAUTH  
!  
crypto map main_map client configuration address initiate  
crypto map main_map client authentication vpnauth  
!  
! List of the networks to bypass NAT when going into the VPN  
!  
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  
!  
nat (inside) 0 access-list nonat  
!
```



sPIX-1


The following configuration example is for the VPN firewall option in the small enterprise serving as a headend for remotes.

```
!  
! Crypto ACLs for local networks to remote sites  
!  
access-list 107 permit ip 10.4.0.0 0.0.255.255 10.5.0.0 0.0.0.255  
!  
access-list 108 permit ip 10.4.0.0 0.0.255.255 10.6.0.0 0.0.0.255  
!  
! Crypto map entries for VPN Router and VPN Firewall remote sites  
!  
crypto map main_map 10 ipsec-isakmp  
  set peer 172.16.128.2  
  set transform-set strong  
  match address 107  
crypto map main_map 20 ipsec-isakmp  
  set peer 172.16.128.5  
  set transform-set strong  
  match address 108  
!  
! Crypto map enabled on public interface  
!  
interface FastEthernet0/0  
  ip address 10.4.1.1 255.255.255.0  
!  
interface Serial1/0  
  ip address 172.16.132.2 255.255.255.0  
  crypto map ent1  
!  
! Default route to forward packets to IKE peers and remote VPNs  
!  
ip route 0.0.0.0 0.0.0.0 172.16.132.1  
!
```



Medium-Enterprise Configurations

Figure 23
Medium Business Corporate Internet Module VPN



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 PIX Firewall (mPIX-1)
- Cisco IDS Sensors (mIDS-1 and mIDS-2)
- Enterscept HIDS



mPIX-1

The following configuration example is for the VPN firewall in the medium enterprise when serving as a headend to remotes.

```
!  
! Crypto ACLs for local networks to remote sites  
!  
access-list remote_concentrators permit ip 10.0.0.0 255.0.0.0 10.9.0.0 255.255.255.0  
!  
access-list remote_routers permit ip 10.0.0.0 255.0.0.0 10.12.0.0 255.255.255.0  
!  
access-list remote_firewalls permit ip 10.0.0.0 255.0.0.0 10.11.0.0 255.255.255.0  
!  
ip address outside 172.16.240.1 255.255.255.0  
ip address inside 10.3.4.1 255.255.255.0  
!  
! Default route to forward packets to IKE peers and remote VPNs local static routes  
!  
route outside 0.0.0.0 0.0.0.0 172.16.240.2 1  
route inside 10.3.1.0 255.255.255.0 10.3.4.2 1  
route inside 10.3.2.0 255.255.255.0 10.3.4.2 1  
route inside 10.3.3.0 255.255.255.0 10.3.4.2 1  
route inside 10.3.8.0 255.255.255.0 10.3.4.1 1  
!  
! Static route for remote access client virtual IP addresses  
!  
route vpn 10.3.7.0 255.255.255.0 10.3.5.5 1  
!  
! Crypto map entries for multiple remote device types  
!  
crypto map main_map 10 ipsec-isakmp  
crypto map main_map 10 match address remote_firewalls  
crypto map main_map 10 set peer 172.16.144.5  
crypto map main_map 10 set transform-set strong  
!  
crypto map main_map 20 ipsec-isakmp  
crypto map main_map 20 match address remote_routers  
crypto map main_map 20 set peer 172.16.144.6  
crypto map main_map 20 set transform-set strong  
!  
crypto map main_map 30 ipsec-isakmp  
crypto map main_map 30 match address remote_concentrators  
crypto map main_map 30 set peer 172.16.144.7  
crypto map main_map 30 set transform-set strong  
crypto map main_map interface outside  
!  
! Bypass NAT for any local network to any remote VPN network  
!  
access-list nonat permit ip 10.0.0.0 255.0.0.0 10.9.0.0 255.255.0.0  
access-list nonat permit ip 10.0.0.0 255.0.0.0 10.12.0.0 255.255.0.0  
access-list nonat permit ip 10.0.0.0 255.0.0.0 10.11.0.0 255.255.0.0  
!  
nat (inside) 0 access-list nonat  
!
```



mPIX-1

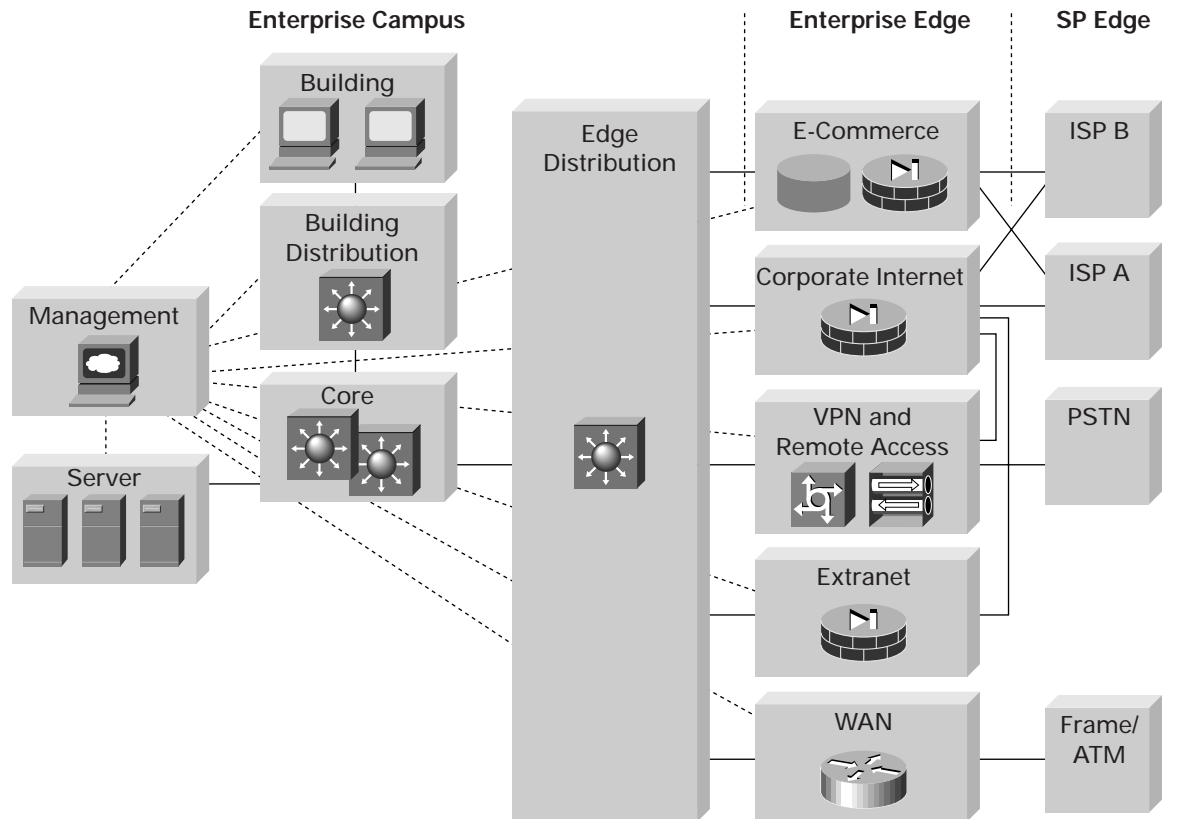
The following configuration example is for the VPN firewall in the medium enterprise when serving as a branch to the large enterprise headend.

```
!  
! Crypto ACLs for local networks to large enterprise head-end  
!  
access-list branch_acl permit ip 10.3.0.0 255.255.0.0 10.0.0.0 255.0.0.0  
!  
ip address outside 172.16.240.1 255.255.255.0  
ip address inside 10.3.4.1 255.255.255.0  
!  
! Default route to forward packets to IKE peers and remote VPNs local static routes  
!  
route outside 0.0.0.0 0.0.0.0 172.16.240.2 1  
!  
route inside 10.3.1.0 255.255.255.0 10.3.4.2 1  
route inside 10.3.2.0 255.255.255.0 10.3.4.2 1  
route inside 10.3.3.0 255.255.255.0 10.3.4.2 1  
route inside 10.3.8.0 255.255.255.0 10.3.4.1 1  
!  
! -type or Cisco-type IKE Keepalives for high availability enabled  
!  
crypto isakmp keepalive 10  
!  
! Crypto map entry for head-end connection  
!  
crypto map main_map 10 ipsec-isakmp  
crypto map main_map 10 match address branch_acl  
crypto map main_map 10 set peer 172.16.226.102  
crypto map main_map 10 set transform-set strong  
crypto map main_map interface outside  
!  
! Bypass NAT for any local network to any remote VPN network  
!  
access-list nonat permit ip 10.3.0.0 255.255.0.0 10.0.0.0 255.0.0.0  
!  
nat (inside) 0 access-list nonat  
!
```



Large-Enterprise Configurations

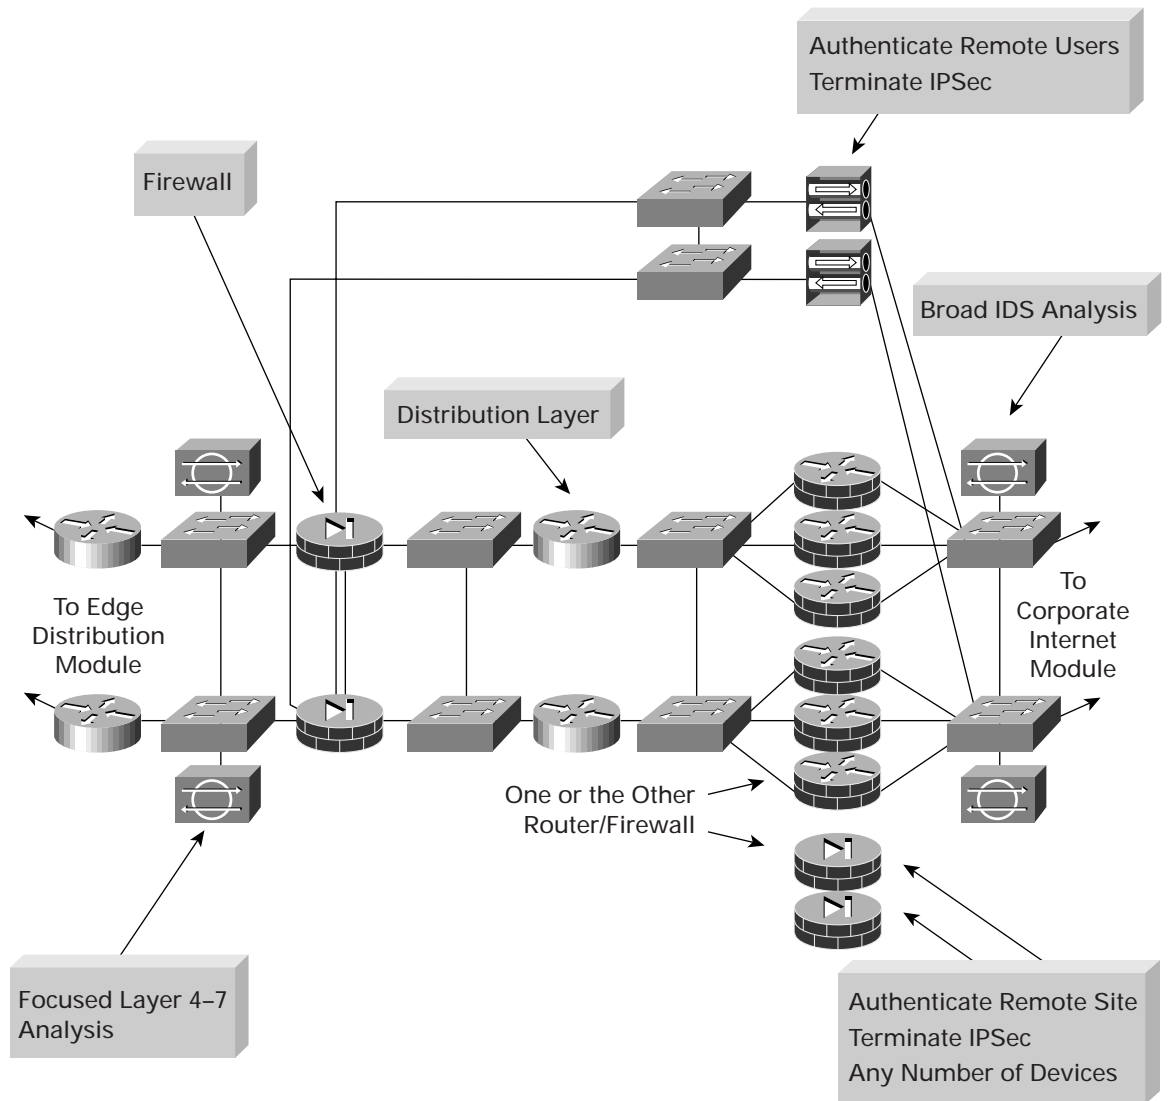
Figure 24
SAFE Enterprise





Remote-Access and VPN Module

Figure 25
Large Enterprise VPN/Remote Access Module



Products Used

- Cisco Catalyst Layer 2 Switches (eCAT-90 through eCAT-93)
- Cisco IOS Routers with 3DES encryption support (eVPN-96 through eVPN-99)
- Cisco VPN 3000 Series Concentrator (eVPN-47 and eVPN-48)
- Cisco PIX Firewall (ePIX-32 and ePIX-32-s; eVPN-101 and eVPN-101-s)
- Cisco IDS Sensors (eIDS-87 though eIDS-90)



- Intercept HIDS
- Baltimore MIMESweeper Email Filtering

eCAT-92

The following configuration example is for the primary distribution router.

```
!  
! Internal interface, HSRP running for next-hop firewall  
!  
interface GigabitEthernet49  
  ip address 10.1.148.92 255.255.255.0  
  standby 1 timers 5 15  
  standby 1 priority 110  
  standby 1 preempt  
  standby 1 authentication k&9ng@6  
  standby 1 ip 10.1.148.100  
!  
! External interface facing VPN devices, HSRP running for VPN Firewall option  
!  
interface GigabitEthernet50  
  ip address 172.16.227.92 255.255.255.0  
  standby 1 timers 5 15  
  standby 1 priority 110  
  standby 1 preempt  
  standby 1 authentication a(4ir#3  
  standby 1 ip 172.16.227.200  
!  
! Routing protocol running on external interface to track VPN devices  
! Redistribution of statics to inject 10 network  
!  
router eigrp 10  
  redistribute static  
  network 172.16.0.0  
  no auto-summary  
!  
! Static routes to inject entire 10 net into table  
!  
ip route 10.0.0.0 255.0.0.0 10.1.148.32  
!  
! Static routes for sites terminated by VPN Firewall option  
!  
ip route 10.3.0.0 255.255.0.0 172.16.227.102  
ip route 10.7.0.0 255.255.0.0 172.16.227.102  
ip route 10.8.0.0 255.255.0.0 172.16.227.102  
!
```




eCAT-93

The following configuration example is for the secondary distribution router.

```
!  
! Internal interface, HSRP running for next-hop firewall  
!  
interface GigabitEthernet49  
  ip address 10.1.148.93 255.255.255.0  
  standby 1 timers 5 15  
  standby 1 priority 100  
  standby 1 preempt  
  standby 1 authentication k&9ng@6  
  standby 1 ip 10.1.148.100  
!  
! External interface facing VPN devices, HSRP running for VPN Firewall option  
!  
interface GigabitEthernet50  
  ip address 172.16.227.93 255.255.255.0  
  standby 1 timers 5 15  
  standby 1 priority 100  
  standby 1 preempt  
  standby 1 authentication a(4ir#3  
  standby 1 ip 172.16.227.200  
!  
! Routing protocol running on external interface to track VPN devices  
! Redistribution of statics to inject 10 network  
!  
router eigrp 10  
  redistribute static  
  network 172.16.0.0  
  no auto-summary  
!  
! Static routes to inject entire 10 net into table  
!  
ip route 10.0.0.0 255.0.0.0 10.1.148.32  
!  
! Static routes for sites terminated by VPN Firewall option  
!  
ip route 10.3.0.0 255.255.0.0 172.16.227.102  
ip route 10.7.0.0 255.255.0.0 172.16.227.102  
ip route 10.8.0.0 255.255.0.0 172.16.227.102  
!
```



ePIX-101

The following configuration example is for the VPN firewall option (the secondary is not shown because it has an identical configuration).

```
!  
! Crypto ACLs for local networks to remote sites  
!  
access-list remote-firewall-1 permit ip 10.0.0.0 255.0.0.0 10.3.0.0 255.255.0.0  
!  
access-list remote-concentrator-1 permit ip 10.0.0.0 255.0.0.0 10.7.0.0 255.255.0.0  
!  
access-list remote-router-1 permit ip 10.0.0.0 255.0.0.0 10.8.0.0 255.255.0.0  
!  
ip address outside 172.16.226.102 255.255.255.0  
ip address inside 172.16.227.102 255.255.255.0  
!  
! Default route to forward packets to IKE peers and remote VPNs local static routes  
! All 10.1.0.0 large enterprise remote networks statically routed  
!  
route outside 0.0.0.0 0.0.0.0 172.16.226.200 1  
route inside 10.1.0.0 255.255.0.0 172.16.227.200 1  
!  
! DPD-type or Cisco-type IKE Keepalives for high availability enabled  
!  
crypto isakmp keepalive 10  
!  
! Crypto map entries for multiple remote device types  
!  
crypto map main_map 10 ipsec-isakmp  
crypto map main_map 10 match address remote-firewall-1  
crypto map main_map 10 set peer 172.16.240.1  
crypto map main_map 10 set transform-set strong  
!  
crypto map main_map 20 ipsec-isakmp  
crypto map main_map 20 match address remote-concentrator-1  
crypto map main_map 20 set peer 172.16.128.4  
crypto map main_map 20 set transform-set strong  
!  
crypto map main_map 30 ipsec-isakmp  
crypto map main_map 30 match address remote-router-1  
crypto map main_map 30 set peer 172.16.128.6  
crypto map main_map 30 set transform-set strong  
!  
crypto map main_map interface outside  
!  
! Bypass NAT for any local network to any remote VPN network  
!  
access-list nonat permit ip 10.1.0.0 255.255.0.0 10.7.0.0 255.255.0.0  
access-list nonat permit ip 10.1.0.0 255.255.0.0 10.8.0.0 255.255.0.0  
access-list nonat permit ip 10.1.0.0 255.255.0.0 10.3.0.0 255.255.0.0  
!  
nat (inside) 0 access-list nonat
```

The following configuration example is for the VPN router option. These configurations show three headend and three remote-site configurations with load dispersion on failure.



eVPN-96

VPN Router Headend Device 1

```
!  
! Crypto ACLs to protect each GRE flow between each peer  
!  
access-list 100 permit gre host 172.16.226.96 host 172.16.144.101  
access-list 101 permit gre host 172.16.226.96 host 172.16.144.103  
!  
! Crypto map entries for each peer  
!  
crypto map main_map 10 ipsec-isakmp  
  set peer 172.16.144.101  
  set transform-set strong  
  match address 100  
crypto map main_map 20 ipsec-isakmp  
  set peer 172.16.144.103  
  set transform-set strong  
  match address 101  
!  
! GRE Tunnels for each remote peer  
!  
!  
interface Tunnel0  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.144.101  
  crypto map main_map  
!  
! Note the bandwidth statement on the next tunnel interface  
!  
interface Tunnel1  
  band 5  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.144.103  
  crypto map main_map  
!  
!  
!  
interface FastEthernet0/0  
  ip address 172.16.226.96 255.255.255.0  
  crypto map main_map  
!  
interface FastEthernet0/1  
  ip address 172.16.227.96 255.255.255.0  
!  
! Routing protocol configuration, updates are not sent on the  
! public interface - only on the inside and tunnel interfaces.  
!  
router eigrp 10  
  passive-interface FastEthernet0/0  
  network 172.16.0.0  
  distribute-list 1 out  
  distribute-list 1 in  
  no auto-summary  
!  
access-list 1 permit 10.0.0.0 0.255.255.255
```



```
!  
! Default route to forward packets to IKE peers and remote VPNs local static routes  
!  
ip route 0.0.0.0 0.0.0.0 172.16.226.200  
!
```

A sample routing table for the above configuration follows.

```
show ip route  
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP  
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area  
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2  
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP  
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area  
* - candidate default, U - per-user static route, o - ODR  
P - periodic downloaded static route  
Gateway of last resort is 172.16.226.200 to network 0.0.0.0  
172.16.0.0/24 is subnetted, 2 subnets  
C 172.16.226.0 is directly connected, FastEthernet0/0  
C 172.16.227.0 is directly connected, FastEthernet0/1  
10.0.0.0/8 is variably subnetted, 24 subnets, 2 masks  
D 10.10.1.0/24 [90/297246976] via 10.10.1.1, 00:33:24, Tunnel0  
D 10.10.2.0/24  
[90/297249536] via 172.16.227.97, 00:33:32, FastEthernet0/1  
D 10.10.3.0/24  
[90/297249536] via 172.16.227.98, 00:33:32, FastEthernet0/1  
D EX 10.0.0.0/8 [170/30720] via 172.16.227.93, 00:00:00, FastEthernet0/1  
[170/30720] via 172.16.227.92, 00:00:00, FastEthernet0/1  
D 10.1.148.0/24 [90/30720] via 172.16.227.93, 00:33:37, FastEthernet0/1  
[90/30720] via 172.16.227.92, 00:33:37, FastEthernet0/1  
S* 0.0.0.0/0 [1/0] via 172.16.226.200
```



eVPN-97

VPN Router Headend Device 2 (abbreviated)

```
!  
access-list 100 permit gre host 172.16.226.97 host 172.16.144.101  
access-list 101 permit gre host 172.16.226.97 host 172.16.144.102  
!  
crypto map main_map 10 ipsec-isakmp  
  set peer 172.16.144.101  
  set transform-set strong  
  match address 100  
crypto map main_map 20 ipsec-isakmp  
  set peer 172.16.144.102  
  set transform-set strong  
  match address 101  
!  
interface Tunnel0  
  band 5  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.144.101  
  crypto map main_map  
!  
interface Tunnel1  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.144.102  
  crypto map main_map  
!  
interface FastEthernet0/0  
  ip address 172.16.226.97 255.255.255.0  
  crypto map main_map  
!  
interface FastEthernet0/1  
  ip address 172.16.227.97 255.255.255.0  
!
```



eVPN-98

VPN Router Headend Device 3 (abbreviated)

```
!  
access-list 100 permit gre host 172.16.226.98 host 172.16.144.102  
access-list 101 permit gre host 172.16.226.98 host 172.16.144.103  
!  
crypto map main_map 10 ipsec-isakmp  
  set peer 172.16.144.102  
  set transform-set strong  
  match address 100  
crypto map main_map 20 ipsec-isakmp  
  set peer 172.16.144.103  
  set transform-set strong  
match address 101  
!  
interface Tunnel0  
  band 5  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.144.102  
  crypto map main_map  
!  
interface Tunnel1  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.144.103  
crypto map main_map  
!  
interface FastEthernet0/0  
  ip address 172.16.226.98 255.255.255.0  
  crypto map main_map  
!  
interface FastEthernet0/1  
  ip address 172.16.227.98 255.255.255.0  
!
```

Sample Remote Router Configurations for the above headends

VPN Router Remote-Site Device 1

```
!  
! Crypto ACLs for the primary and secondary GRE tunnels  
!  
access-list 100 permit gre host 172.16.144.101 host 172.16.226.96  
access-list 101 permit gre host 172.16.144.101 host 172.16.226.97  
!  
! Crypto map for the primary and secondary GRE tunnels  
!  
crypto map main_map 1 ipsec-isakmp  
  set peer 172.16.226.96  
  set transform-set main_map  
  match address 100  
crypto map main_map 2 ipsec-isakmp  
  set peer 172.16.226.97  
  set transform-set main_map  
  match address 101
```



```
!  
! GRE tunnel interfaces, note the second has  
! a bandwidth statement to make it secondary  
!  
interface Tunnel0  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/1  
  tunnel destination 172.16.226.96  
  crypto map main_map  
!  
interface Tunnel1  
  bandwidth 5  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/1  
  tunnel destination 172.16.226.97  
  crypto map main_map  
!  
interface FastEthernet0/0  
  ip address 10.10.1.1 255.255.255.0  
!  
interface FastEthernet0/1  
  ip address 172.16.144.101 255.255.255.0  
  crypto map main_map  
!  
! Routing protocol configuration, updates  
! are only sent on the tunnel interfaces.  
!  
router eigrp 10  
  passive-interface FastEthernet0/0  
  passive-interface FastEthernet0/1  
  network 10.0.0.0  
  no auto-summary  
  no eigrp log-neighbor-changes  
!  
ip route 0.0.0.0 0.0.0.0 172.16.144.2  
!
```

A sample routing table for the above configuration follows.

```
show ip route  
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP  
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area  
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2  
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP  
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area  
       * - candidate default, U - per-user static route, o - ODR  
       P - periodic downloaded static route  
Gateway of last resort is 172.16.144.2 to network 0.0.0.0  
172.16.0.0/24 is subnetted, 1 subnets  
C 172.16.144.0 is directly connected, FastEthernet0/1  
  10.0.0.0/8 is variably subnetted, 24 subnets, 2 masks  
C 10.10.1.0/24 is directly connected, FastEthernet0/0  
D 10.10.2.0/24 [90/310049536] via 172.16.226.96, 00:43:42, Tunnel0  
D 10.10.3.0/24 [90/310049536] via 172.16.226.96, 00:43:42, Tunnel0  
D EX 10.0.0.0/8 [170/297249536] via 172.16.226.96, 00:00:00, Tunnel0  
D 10.1.148.0/24 [90/297249536] via 172.16.226.96, 00:43:46, Tunnel0  
S* 0.0.0.0/0 [1/0] via 172.16.144.2
```



VPN Router Remote-Site Device 2 (abbreviated)

```
!  
access-list 100 permit gre host 172.16.144.102 host 172.16.226.97  
access-list 101 permit gre host 172.16.144.102 host 172.16.226.98  
!  
crypto map main_map 1 ipsec-isakmp  
  set peer 172.16.226.97  
  set transform-set main_map  
  match address 100  
crypto map main_map 2 ipsec-isakmp  
  set peer 172.16.226.98  
  set transform-set main_map  
  match address 101  
!  
interface Tunnel0  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/1  
  tunnel destination 172.16.226.97  
  crypto map main_map  
!  
interface Tunnel1  
  bandwidth 5  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/1  
  tunnel destination 172.16.226.98  
  crypto map main_map  
!  
interface FastEthernet0/0  
  ip address 10.10.2.1 255.255.255.0  
!  
interface FastEthernet0/1  
  ip address 172.16.144.102 255.255.255.0  
  crypto map main_map  
!
```



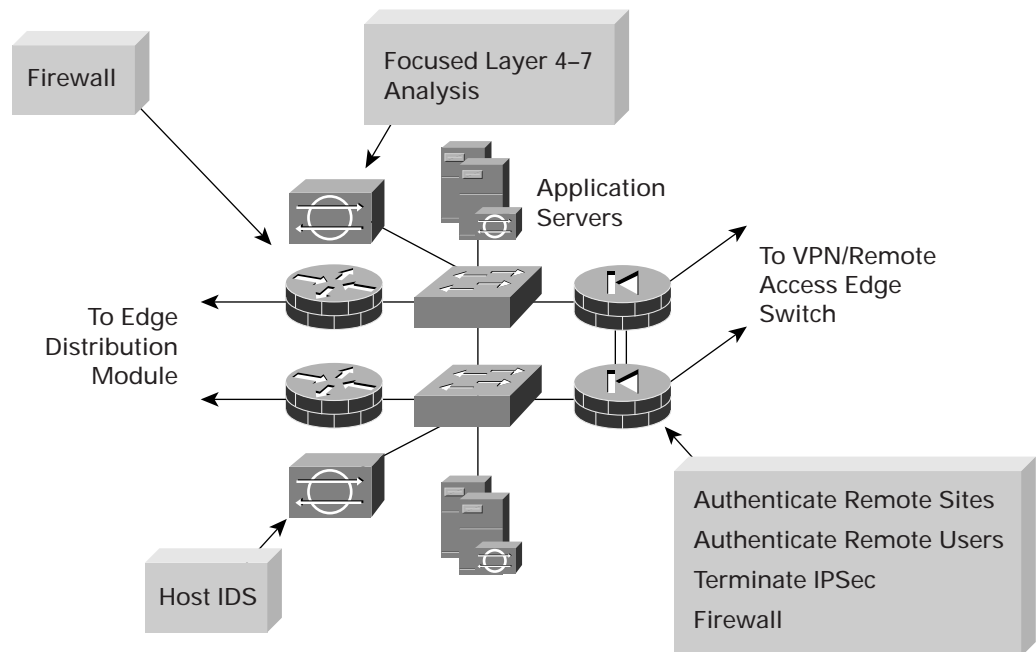

VPN Router Remote-Site Device 3 (abbreviated)

```
!  
access-list 100 permit gre host 172.16.144.103 host 172.16.226.98  
access-list 101 permit gre host 172.16.144.103 host 172.16.226.96  
!  
crypto map main_map 1 ipsec-isakmp  
  set peer 172.16.226.98  
  set transform-set main_map  
  match address 100  
crypto map main_map 2 ipsec-isakmp  
  set peer 172.16.226.96  
  set transform-set main_map  
  match address 101  
!  
interface Tunnel0  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/1  
  tunnel destination 172.16.226.98  
  crypto map main_map  
!  
interface Tunnel1  
  bandwidth 5  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/1  
  tunnel destination 172.16.226.96  
  crypto map main_map  
!  
interface FastEthernet0/0  
  ip address 10.10.3.1 255.255.255.0  
!  
interface FastEthernet0/1  
  ip address 172.16.144.103 255.255.255.0  
  crypto map main_map  
!
```



Distributed Enterprise Configurations

Figure 26
Distribution Enterprise Hub Module VPN



The following configurations include the head-end, two sets of highly available distribution layers and a sample remote-site configuration for each layer.

- Cisco Catalyst Switches (dCAT-1 and dCAT-2)
- Cisco IOS Routers with 3DES encryption support (dVPN-1 and dVPN-2)
- Cisco PIX Firewall (dPIX-1 and dPIX-1-s)
- Cisco IDS Sensors (dIDS-1 and dIDS-2)
- Enterscept HIDS



eVPN-96

Enterprise Headend, Primary Unit

These configurations are provided as sample headend configurations to the distribution module.

```
!  
! Crypto ACLs to protect GRE traffic from distribution layer  
!  
access-list 100 permit gre host 172.16.226.96 host 172.16.144.8  
access-list 101 permit gre host 172.16.226.96 host 172.16.144.9  
access-list 102 permit gre host 172.16.226.96 host 172.16.144.10  
access-list 103 permit gre host 172.16.226.96 host 172.16.144.11  
!  
! Crypto maps for highly available distribution layer  
!  
crypto map main_map 10 ipsec-isakmp  
  set peer 172.16.144.8  
  set transform-set strong  
  match address 100  
crypto map main_map 20 ipsec-isakmp  
  set peer 172.16.144.9  
  set transform-set strong  
  match address 101  
crypto map main_map 30 ipsec-isakmp  
  set peer 172.16.144.10  
  set transform-set strong  
  match address 102  
crypto map main_map 40 ipsec-isakmp  
  set peer 172.16.144.11  
  set transform-set strong  
  match address 103  
!  
! GRE tunnels to distribution layer  
!  
interface Tunnel1  
  bandwidth 20  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.144.8  
  crypto map main_map  
!  
interface Tunnel2  
  bandwidth 5  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.144.9  
  crypto map main_map  
!  
interface Tunnel3  
  bandwidth 20  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.144.10  
  crypto map main_map  
!  
interface Tunnel4  
  bandwidth 5  
  ip unnumbered FastEthernet0/0
```



```
tunnel source FastEthernet0/0
tunnel destination 172.16.144.11
crypto map main_map
!
!
!
interface FastEthernet0/0
ip address 172.16.226.96 255.255.255.0
crypto map main_map
!
interface FastEthernet0/1
ip address 172.16.227.96 255.255.255.0
!
interface FastEthernet4/0
no ip address
shutdown
duplex half
!
! Routing protocols running over tunnels
!
router eigrp 10
passive-interface FastEthernet0/0
network 172.16.0.0
distribute-list 1 out
distribute-list 1 in
no auto-summary
no eigrp log-neighbor-changes
!
access-list 1 permit 10.0.0.0 0.255.255.255
!
!
!
ip route 0.0.0.0 0.0.0.0 172.16.226.200
!
```



eVPN-97

Enterprise Headend, Secondary Unit

```
!  
! Crypto ACLs to protect GRE traffic from distribution layer  
!  
access-list 100 permit gre host 172.16.226.97 host 172.16.144.8  
access-list 101 permit gre host 172.16.226.97 host 172.16.144.9  
access-list 102 permit gre host 172.16.226.97 host 172.16.144.10  
access-list 103 permit gre host 172.16.226.97 host 172.16.144.11  
!  
! Crypto maps for highly available distribution layer  
!  
crypto map main_map 10 ipsec-isakmp  
  set peer 172.16.144.8  
  set transform-set strong  
  match address 100  
crypto map main_map 20 ipsec-isakmp  
  set peer 172.16.144.9  
  set transform-set strong  
  match address 101  
crypto map main_map 30 ipsec-isakmp  
  set peer 172.16.144.10  
  set transform-set strong  
  match address 102  
crypto map main_map 40 ipsec-isakmp  
  set peer 172.16.144.11  
  set transform-set strong  
  match address 103  
!  
! GRE tunnels to distribution layer  
!  
interface Tunnel1  
  bandwidth 20  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.144.8  
  crypto map main_map  
!  
interface Tunnel2  
  bandwidth 5  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.144.9  
  crypto map main_map  
!  
interface Tunnel3  
  bandwidth 20  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.144.10  
  crypto map main_map  
!  
interface Tunnel4  
  bandwidth 5  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.144.11
```



```
crypto map main_map
!
!
!
interface FastEthernet0/0
 ip address 172.16.226.97 255.255.255.0
 crypto map main_map
!
interface FastEthernet0/1
 ip address 172.16.227.97 255.255.255.0
!
interface FastEthernet4/0
 no ip address
 shutdown
 duplex half
!
! Routing protocols running over tunnels
!
router eigrp 10
 passive-interface FastEthernet0/0
 network 172.16.0.0
 distribute-list 1 out
 distribute-list 1 in
 no auto-summary
 no eigrp log-neighbor-changes
!
access-list 1 permit 10.0.0.0 0.255.255.255
!
!
!
ip route 0.0.0.0 0.0.0.0 172.16.226.200
!
```



dVPN-1

Distribution Layer Group 1, Primary Unit

```
!  
! Crypto ACLs to protect GRE traffic from head-ends and remote site  
!  
access-list 100 permit gre host 172.16.144.8 host 172.16.226.96  
access-list 101 permit gre host 172.16.144.8 host 172.16.226.97  
access-list 102 permit gre host 172.16.144.8 host 172.16.128.10  
!  
! Crypto maps for head-ends and remote sites  
!  
crypto map main_map 10 ipsec-isakmp  
  set peer 172.16.226.96  
  set transform-set strong  
  match address 100  
crypto map main_map 20 ipsec-isakmp  
  set peer 172.16.226.97  
  set transform-set strong  
  match address 101  
crypto map main_map 30 ipsec-isakmp  
  set peer 172.16.128.10  
  set transform-set strong  
  match address 102  
!  
! GRE tunnel interfaces for head-end and remote sites  
!  
interface Tunnel1  
  bandwidth 20  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.226.96  
  crypto map main_map  
!  
interface Tunnel2  
  bandwidth 5  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.226.97  
  crypto map main_map  
!  
interface Tunnel3  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.128.10  
  crypto map main_map  
!  
interface FastEthernet0/0  
  ip address 172.16.144.8 255.255.255.0  
  crypto map main_map  
!  
! Device serves as primary for first remote site  
! group and secondary for second remote site group  
!  
interface FastEthernet0/1  
  ip address 10.30.0.1 255.255.255.0  
  standby 1 timers 5 15  
  standby 1 priority 100 preempt delay 2
```



```
standby 1 ip 10.30.0.100
standby 1 track FastEthernet0/0
standby 2 timers 5 15
standby 2 priority 90 preempt delay 2
standby 2 ip 10.30.0.101
standby 2 track FastEthernet0/0
!
! Routing protocol running on tunnels
router eigrp 10
  passive-interface FastEthernet0/0
  network 172.16.144.0 0.0.0.255
  distribute-list 1 out
  distribute-list 1 in
  no auto-summary
!
access-list 1 permit 10.0.0.0 0.255.255.255
!
! Catch all default route
!
ip route 0.0.0.0 0.0.0.0 172.16.144.2
!
```




dVPN-2

Distribution Layer Group 1, Secondary Unit

```
!  
! Crypto ACLs to protect GRE traffic from head-ends and remote site  
!  
access-list 100 permit gre host 172.16.144.9 host 172.16.226.96  
access-list 101 permit gre host 172.16.144.9 host 172.16.226.97  
access-list 102 permit gre host 172.16.144.9 host 172.16.128.10  
!  
! Crypto maps for head-ends and remote sites  
!  
crypto map main_map 10 ipsec-isakmp  
  set peer 172.16.226.96  
  set transform-set strong  
  match address 100  
crypto map main_map 20 ipsec-isakmp  
  set peer 172.16.226.97  
  set transform-set strong  
  match address 101  
crypto map main_map 30 ipsec-isakmp  
  set peer 172.16.128.10  
  set transform-set strong  
  match address 102  
!  
! GRE tunnel interfaces for head-end and remote sites  
!  
interface Tunnel1  
  bandwidth 20  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.226.96  
  crypto map main_map  
!  
interface Tunnel2  
  bandwidth 5  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.226.97  
  crypto map main_map  
!  
interface Tunnel3  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.128.10  
  crypto map main_map  
!  
interface FastEthernet0/0  
  ip address 172.16.144.9 255.255.255.0  
  crypto map main_map  
!  
interface FastEthernet0/1  
  ip address 10.30.0.2 255.255.255.0  
  ! Device serves as secondary for first remote  
  ! site group and primary for second remote site group  
!  
interface FastEthernet0/1  
  ip address 10.30.0.2 255.255.255.0
```



```
standby 1 timers 5 15
standby 1 priority 90 preempt delay 2
standby 1 ip 10.30.0.100
standby 1 track FastEthernet0/0
standby 2 timers 5 15
standby 2 priority 100 preempt delay 2
standby 2 ip 10.30.0.101
standby 2 track FastEthernet0/0
!
! Routing protocol running on tunnels
!
router eigrp 10
  passive-interface FastEthernet0/0
  network 172.16.144.0 0.0.0.255
  distribute-list 1 out
  distribute-list 1 in
  no auto-summary
!
access-list 1 permit 10.0.0.0 0.255.255.255
!
! Catch all default route
!
ip route 0.0.0.0 0.0.0.0 172.16.144.2
!
```

dVPN-3

Distribution Layer Group 2, Secondary Unit!

This is another distribution group configuration identical in topology to the above figure.

```
! Crypto ACLs to protect GRE traffic from head-ends and remote site
!
access-list 100 permit gre host 172.16.144.10 host 172.16.226.96
access-list 101 permit gre host 172.16.144.10 host 172.16.226.97
access-list 102 permit gre host 172.16.144.10 host 172.16.128.11
!
! Crypto maps for head-ends and remote sites
!
crypto map main_map 10 ipsec-isakmp
  set peer 172.16.226.96
  set transform-set strong
  match address 100
crypto map main_map 20 ipsec-isakmp
  set peer 172.16.226.97
  set transform-set strong
  match address 101
crypto map main_map 30 ipsec-isakmp
  set peer 172.16.128.10
  set transform-set strong
  match address 102
!
! GRE tunnel interfaces for head-end and remote sites
!
interface Tunnell
  bandwidth 20
  ip unnumbered FastEthernet0/0
  tunnel source FastEthernet0/0
```



```
tunnel destination 172.16.226.96
crypto map main_map
!
interface Tunnel2
bandwidth 5
ip unnumbered FastEthernet0/0
tunnel source FastEthernet0/0
tunnel destination 172.16.226.97
crypto map main_map
!
interface Tunnel3
ip unnumbered FastEthernet0/0
tunnel source FastEthernet0/0
tunnel destination 172.16.128.11
crypto map main_map
!
interface FastEthernet0/0
ip address 172.16.144.10 255.255.255.0
crypto map main_map
!
interface FastEthernet0/1
ip address 10.30.1.1 255.255.255.0
! Device serves as primary for first remote site
group and secondary for second remote site group
!
interface FastEthernet0/1
ip address 10.30.1.1 255.255.255.0
standby 1 timers 5 15
standby 1 priority 100 preempt delay 2
standby 1 ip 10.30.1.100
standby 1 track FastEthernet0/0
standby 2 timers 5 15
standby 2 priority 90 preempt delay 2
standby 2 ip 10.30.1.101
standby 2 track FastEthernet0/0
!
! Routing protocol running on tunnels
!
router eigrp 10
passive-interface FastEthernet0/0
network 172.16.144.0 0.0.0.255
distribute-list 1 out
distribute-list 1 in
no auto-summary
!
access-list 1 permit 10.0.0.0 0.255.255.255
!
! Catch all default route
!
ip route 0.0.0.0 0.0.0.0 172.16.144.2
!
```



dVPN-4

Distribution Layer Group 2, Secondary Unit!

This is another distribution group configuration.

```
!  
! Crypto ACLs to protect GRE traffic from head-ends and remote site  
!  
access-list 100 permit gre host 172.16.144.11 host 172.16.226.96  
access-list 101 permit gre host 172.16.144.11 host 172.16.226.97  
access-list 102 permit gre host 172.16.144.11 host 172.16.128.11  
!  
! Crypto maps for head-ends and remote sites  
!  
crypto map main_map 10 ipsec-isakmp  
  set peer 172.16.226.96  
  set transform-set strong  
  match address 100  
crypto map main_map 20 ipsec-isakmp  
  set peer 172.16.226.97  
  set transform-set strong  
  match address 101  
crypto map main_map 30 ipsec-isakmp  
  set peer 172.16.128.11  
  set transform-set strong  
  match address 102  
!  
! GRE tunnel interfaces for head-end and remote sites  
!  
interface Tunnel1  
  bandwidth 20  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.226.96  
  crypto map main_map  
!  
interface Tunnel2  
  bandwidth 5  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.226.97  
  crypto map main_map  
!  
interface Tunnel3  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.128.11  
  crypto map main_map  
!  
interface FastEthernet0/0  
  ip address 172.16.144.11 255.255.255.0  
  crypto map main_map  
!  
interface FastEthernet0/1  
  ip address 10.30.1.2 255.255.255.0  
!  
! Device serves as secondary for first remote  
! site group and primary for second remote site group  
!
```



```
interface FastEthernet0/1
 ip address 10.30.1.2 255.255.255.0
 standby 1 timers 5 15
 standby 1 priority 90 preempt delay 2
 standby 1 ip 10.30.1.100
 standby 1 track FastEthernet0/0
 standby 2 timers 5 15
 standby 2 priority 100 preempt delay 2
 standby 2 ip 10.30.1.101
 standby 2 track FastEthernet0/0
 !
 ! Routing protocol running on tunnels
 !
router eigrp 10
 passive-interface FastEthernet0/0
 network 172.16.144.0 0.0.0.255
 distribute-list 1 out
 distribute-list 1 in
 no auto-summary
 !
access-list 1 permit 10.0.0.0 0.255.255.255
 !
 ! Catch all default route
 !
ip route 0.0.0.0 0.0.0.0 172.16.144.2
 !
```

dPIX-1

Distribution Group 1 Firewall (abbreviated)

```
!
 ! Two routes, one for each primary HSRP group
 !
route outside 10.20.0.0 255.255.255.0 10.30.0.100 1
route outside 10.20.1.0 255.255.255.0 10.30.0.101 1
 !
```



Remote Site in Distribution Group 1

```
!  
! Redundant crypto ACLs for each tunnel to the redundantdistribution layer  
!  
access-list 100 permit gre host 172.16.128.10 host 172.16.144.8  
access-list 101 permit gre host 172.16.128.10 host 172.16.144.9  
!  
!  
crypto map main_map 10 ipsec-isakmp  
  set peer 172.16.144.8  
  set transform-set strong  
  match address 100  
crypto map main_map 20 ipsec-isakmp  
  set peer 172.16.144.9  
  set transform-set strong  
  match address 101  
!  
!  
!  
interface Tunnel1  
  bandwidth 20  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.144.8  
  crypto map main_map  
!  
interface Tunnel2  
  bandwidth 5  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.144.9  
  crypto map main_map  
!  
interface FastEthernet0/0  
  ip address 172.16.128.10 255.255.255.0  
  crypto map main_map  
!  
interface FastEthernet0/1  
  ip address 10.20.0.1 255.255.255.0  
!  
!  
!  
access-list 1 permit 10.0.0.0 0.255.255.255  
!  
router eigrp 10  
  passive-interface FastEthernet0/0  
  passive-interface FastEthernet0/1  
  network 10.20.0.0 0.0.0.255  
  network 172.16.128.0 0.0.0.255  
  distribute-list 1 out  
  no auto-summary  
!  
!  
!  
ip route 0.0.0.0 0.0.0.0 172.16.128.1  
!
```



What follows is a sample remote configuration that connects to one of the distribution layers.

Remote Site in Distribution Group 2

```
!  
!  
!  
access-list 100 permit gre host 172.16.128.11 host 172.16.144.10  
access-list 101 permit gre host 172.16.128.11 host 172.16.144.11  
!  
!  
!  
crypto map main_map 10 ipsec-isakmp  
  set peer 172.16.144.10  
  set transform-set strong  
  match address 100  
crypto map main_map 20 ipsec-isakmp  
  set peer 172.16.144.11  
  set transform-set strong  
  match address 101  
!  
!  
!  
interface Tunnel1  
  bandwidth 20  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.144.10  
  crypto map main_map  
!  
interface Tunnel2  
  bandwidth 5  
  ip unnumbered FastEthernet0/0  
  tunnel source FastEthernet0/0  
  tunnel destination 172.16.144.11  
  crypto map main_map  
!  
interface FastEthernet0/0  
  ip address 172.16.128.11 255.255.255.0  
  crypto map main_map  
!  
interface FastEthernet0/1  
  ip address 10.20.128.1 255.255.255.0  
!  
!  
!  
access-list 1 permit 10.0.0.0 0.255.255.255  
!  
router eigrp 10  
  passive-interface FastEthernet0/0  
  passive-interface FastEthernet0/1  
  network 10.20.128.0 0.0.0.255  
  network 172.16.128.0 0.0.0.255  
  distribute-list 1 out  
  no auto-summary  
!  
!  
!  
ip route 0.0.0.0 0.0.0.0 172.16.128.1  
!
```



Appendix B: VPN Primer

The Need for VPNs

Virtual private networks (VPNs) provide an alternative to building a private network for site-to-site communication or dial-in access. Because they operate across a shared infrastructure rather than a private network, companies can cost-effectively extend the corporate network to locations that may not have been justified before. For instance, in many domestic applications and most international applications, VPNs provide significant cost savings over private WAN connections. Additionally, rather than having multiple independent circuits terminating at the corporate headend, VPNs allow all the traffic to be aggregated into a single connection. This scenario results in potential bandwidth and cost savings at the headend. Further cost savings result from not having to maintain a private network.

VPNs also provide opportunities for increased productivity within the company. For example, rather than using slow dial-in links, home-office workers can take advantage of VPN technologies over higher-speed access such as *digital subscriber line* (DSL) and cable modem to increase productivity while working from home. Mobile workers can also take advantage of higher-speed Ethernet connections found in many hotels today for access to corporate resources while traveling. The cost savings alone, from not having to pay long-distance telephone charges, may justify the use of VPNs in such cases. Finally, companies can take advantage of VPN technologies to enable new applications and business processes. For example, new e-commerce and supply-chain management business models have been implemented by the automotive industry through the use of the *Automotive Network Exchange* (ANX), which is based on VPN technology.

Types of VPNs

A wide variety of VPN technologies are deployed today. The following sections provide an overview of some of the technologies and terms that a network administrator will require in order to further understand VPNs. VPNs are often looked at from one of two perspectives—either from a functional view or from a technology view. A functional viewpoint emphasizes more of the purpose the VPN, whereas a technology viewpoint emphasizes more of the particular technology used to implement the VPN.

Functional View

From a functional viewpoint, VPNs are often categorized as either remote-access VPNs or site-to-site VPNs. Remote-access VPNs refer to implementations in which individual remote users, often referred to as mobile workers, access the corporate network via their PCs. Mobile workers may use traditional dial-in connections to a local service provider, and then initiate tunnels back to the corporation. An alternative is to initiate the tunnel across higher-speed media, such as Ethernet, found in many hotels today. A relatively recent variation of this is the wireless remote-access VPN, in which a mobile worker accesses the corporate network via a wireless connection on a *personal digital assistant* (PDA). In all these cases, software on the PC or PDA provides a secure connection, often referred to as a tunnel, back to the corporation. Individual mobile workers should be authenticated before being allowed access to the corporate network. The appropriate access control to corporate resources should be applied to mobile workers based upon corporate security policy. For instance, access for business partners may need to be more restrictive than access for employees.

Site-to-site VPNs refer to implementations in which the network of one location is connected to the network of another location via a VPN. Network devices authenticate each other and then establish the VPN connection between the sites. These devices then act as gateways, securely passing traffic destined for the other site. Routers or



firewalls with VPN support and dedicated VPN concentrators all provide this functionality. Site-to-site VPNs can be further viewed as Intranet VPNs or extranet VPNs. Intranet VPNs refer to connections between sites that are all part of the same company. As such, access between sites is generally less restrictive. Extranet VPNs refer to connections between a company and its business partners. Access between sites should be tightly controlled by both entities at their respective sites.

The distinction between remote-access and site-to-site VPNs will become blurred as new devices, such as hardware VPN clients, become more widespread in use. Such devices can appear as if they are a single device accessing the network, although there may be a network with several devices behind them.

Technology View

From a technology perspective, VPNs can be categorized based upon whether they operate across a Layer 2 network or a Layer 3 network. It should be noted that ATM and Frame Relay networks are sometimes referred to as VPN networks because they use a shared infrastructure to provide network services to a large number of users. However, this document does not refer to ATM or Frame Relay networks as VPNs. Instead they are viewed as private network implementations.

Layer 2 VPNs

Layer 2 VPN technologies are designed to run at the data link layer of the *Open System Interconnection* (OSI) model, as opposed to Layer 3 VPN technologies, such as *IP Security* (IPSec), which runs at the network layer of the OSI model. Layer 2 VPNs include the *Point-to-Point Tunneling Protocol* (PPTP) and the *Layer 2 Tunneling Protocol* (L2TP).

PPTP is an older protocol used primarily for dial-in remote-access VPNs. PPTP operates in client/server mode. The client piece can be a remote PC with PPTP software, or an *Internet service provider's* (ISP's) PPTP-enabled *network access server* (NAS). The server piece can be a dial-in router, specialized VPN concentrator, or actual application server. When the initiation of the PPTP tunnel is done by the remote PC, it is often referred to as voluntary mode. When the initiation of the tunnel is done by the NAS, it is often referred to as compulsory mode. It should be noted that with compulsory mode, the end user's PC is not required to have a PPTP client on it. However, this also means the dial-in connection from the end user's PC to the ISP has none of the security services of PPTP, such as encryption.

PPTP encapsulates *Point-to-Point Protocol* (PPP) packets in a modified version of *generic routing encapsulation* (GRE) and transports them across the network. GRE, defined in RFCs 1701 and 1702, is simply a mechanism for performing encapsulation of an arbitrary network layer protocol over another arbitrary network layer protocol. Therefore, PPTP can be used to transport various Layer 3 protocols such as IP, *Internetwork Packet Exchange* (IPX), NetBEUI, and so on. PPTP relies on the authentication mechanisms of PPP—*Password Authentication Protocol* (PAP) and *Challenge Handshake Authentication Protocol* (CHAP)—which are not considered particularly strong. PAP sends passwords across the link in cleartext, and is not secure. CHAP is somewhat more secure than PAP. Rather than sending the password in cleartext, CHAP issues a challenge to which the other side must respond in order to authenticate. Microsoft has created an enhanced version of CHAP, called MS-CHAP, which utilizes information within NT domains for security. The *Internet Engineering Task Force* (IETF) has also defined the PPP *Extensible Authentication Protocol* (EAP) in RFC 2284, in order to accommodate more robust methods of authentication. However, not all implementations support this protocol at this time. Microsoft has also incorporated a protocol called Microsoft Point-to-Point Encryption (MPPE), in order to provide encryption of traffic across a PPTP link. MPPE is based on the RSA RC4 encryption algorithm.



L2TP is widely regarded as the replacement to PPTP, and is considered more scalable than PPTP. L2TP also operates in client/server mode. Similar to PPTP, the L2TP tunnel can be initiated from the remote PC back to the *L2TP network server* (LNS); or from the *L2TP-enabled access concentrator* (LAS) to the LNS. Although L2TP still makes use of PPP, it defines its own tunneling protocol, depending upon the transport media rather than using GRE. Therefore, L2TP can also be used to transport various Layer 3 protocols, other than IP, although not all implementations support this today. L2TP can make use of PAP, CHAP, and EAP for authentication. A major difference with L2TP, however, is that it supports the use of IPSec, which can be used to secure traffic all the way from the end user's PC to the corporate network.

Layer 3 VPNs

Layer 3 VPN technologies are designed to run at the network layer of the OSI model. Typically these VPNs utilize the IP protocol as the network layer protocol. Layer 3 VPNs include *Multiprotocol Label Switching* (MPLS), and IPSec.

MPLS is typically offered as a site-to-site VPN service from a service provider. The service provider builds a private IP-based network, and offers multiple customers IP connectivity between their sites across this network. The technology allows individual customers to view the MPLS service as if they had a private IP network connecting their sites. This scenario offers customers the same advantages of a Layer 2 private network such as Frame Relay or ATM, but with the scalability and ease of management of a Layer 3 network. Additionally, because MPLS runs across a private IP-based network rather than the Internet, the service provider may be able to provide differentiated levels of service (*quality of service* [QoS]) and *service-level agreements* (SLAs) to its customers. However, because MPLS is based on a service provider's private network, the reach of the service is limited to the locations in which the service provider operates. Typically there is little to no inter-service provider MPLS service today.

IPSec VPNs are covered in depth in the following sections.

IPSec

IPSec is a framework of open standards for ensuring secure private communications over IP networks. IPSec VPNs use the services defined within IPSec to ensure confidentiality, integrity, and authenticity of data communications across public networks, such as the Internet. The security services within IPSec are provided by one of two protocols, the *Authentication Header* (AH) and the *Encapsulating Security Payload* (ESP). Each protocol provides certain services and may be used separately or together, although it is not usually necessary to use both protocols together.

Authentication Header

The AH provides connectionless data integrity and data origin authentication for IP packets. Connectionless data integrity means the original IP packet was not modified in transit from the source to the destination. Data origin authentication verifies the source of the data. Together these joint services are referred to as authentication. The AH is inserted into the IP packet between the IP header and the rest of the packet contents. The AH contains a cryptographic checksum of the packet contents, including the parts of the IP header itself that are immutable in transit. The default cryptographic algorithms for calculating the checksum are *hashed-based message authentication code* (HMAC) coupled with the *Message Digest 5* (MD5) hash function and HMAC coupled with the SHA-1 hash function. A hash algorithm is a one-way mathematical function that takes a variable-length message and produces a unique fixed-length value. SHA-1 is considered to be a stronger hash function because it produces a 160-bit authenticator value (cryptographic checksum), versus the 128-bit authenticator produced by MD5. By taking a received message, calculating the same cryptographic checksum, and comparing it with the value received, the receiver can verify that the message has not been altered in transit. AH also provides an anti-replay service that can



be used to counter a *denial-of-service* (DoS) attack based on an attacker's intercepting a series of packets and then replaying them. It should be noted that the anti-replay service could affect performance if the network reorders packets in order to provide higher QoS for certain types of traffic. If the arriving packets fall outside the anti-replay window, IPSec will reject them. Because AH does nothing to keep the contents of the packets confidential, it is not widely used for IPSec implementations across the Internet. For confidentiality, the ESP must be used.

Encapsulating Security Payload

The Encapsulating Security Payload provides for confidentiality of IP traffic, as well as authentication and anti-replay capabilities. Confidentiality is achieved through encryption. Encryption is the process of taking a message, referred to as cleartext, and passing it through a mathematical algorithm to produce what is known as ciphertext. Decryption is the reverse of the process. Encryption algorithms typically rely on a value, called a key, in order to encrypt and decrypt the data. Two major forms of encryption are used today-symmetric encryption (also known as shared-key encryption) and asymmetric encryption (also known as public/private encryption). Symmetric encryption is about 1000 times faster than asymmetric encryption, and is, therefore, used for the bulk encryption of data. Generally with well-designed encryption algorithms, longer keys result in a higher degree of security, because more brute force is required to try every possible key (known as the key space) in order to decrypt a message. ESP supports a variety of symmetric encryption algorithms for the encryption of data. The default algorithm, Data Encryption Standard (DES), has been in use for about 20 years. DES uses a 56-bit key. However, because DES has been shown to be susceptible to brute-force attacks, *Triple DES* (3DES), which encrypts the data three times with up to three different keys, is the standard algorithm recommended for most business use. It should be noted that because of U.S. government restriction on the export of encryption technology, 3DES may not be available in certain countries. Work is currently being done by the *National Institute of Standards and Technology* (NIST), <http://www.nist.gov>

to define a new, faster, and more secure standard encryption algorithm, referred to as the *Advanced Encryption Standard* (AES). ESP encrypts the higher-level protocol information (the TCP header, for instance) and the actual data itself. Unlike AH, the authentication services of ESP do not protect the IP header of the packet. Most IPSec VPN implementations today use ESP.

IPSec Modes

IPSec can be used to provide secure communication between two hosts, from a host to a security gateway, or between two security gateways. A security gateway is a device, such as a router, firewall, or a dedicated VPN concentrator, which provides the IPSec services (that is, terminates the IPSec connection) and passes traffic through the tunnel to the other side. IPSec can operate in one of two modes to accommodate the different types of connections-tunnel mode and transport mode. With tunnel mode, the entire original IP packet is encapsulated within AH or ESP, and then a new IP header is placed around it. With transport mode, the AH or ESP is placed after the original IP header (as described in the sections earlier). Tunnel mode is used when one or both sides of the IPSec connection is a security gateway and the actual destination hosts behind it do not support IPSec. Therefore, the new IP header has the source address of the gateway itself. With transport mode, operating between two security gateways, the original source and destination addresses can be hidden through the use of encryption. Transport mode can be used when both end hosts support IPSec.

Security Associations

IPSec operates in a peer-to-peer relationship, rather than a client/server relationship. In order for two devices to exchange secured data, they need to agree on which cryptographic algorithms to use. This agreement between the peers is known as a *security association* (SA). Security associations specify information such as what authentication



and encryption algorithms are to be used, the shared session keys, the key lifetimes, the lifetime of the SA itself; as well as other information. There are two types of SAs-Internet Security Association Key Management Protocol (ISAKMP) SAs (also known as IKE SAs) and IPSec SAs. The IKE SA is bidirectional and provides a secure communication channel between the two parties that can be used to negotiate further communications. The IPSec SA is unidirectional, and is used for the actual communication between devices. It should be noted that for two-way communication between devices there must be at least two IPSec SAs-one in each direction.

Authentication and Key Management

Because of all the keys that have to be exchanged in order for the parties to communicate securely, key exchange and management is an important part of IPSec. The two methods of handling key exchange and management specified within IPSec are manual keying and *Internet Key Exchange* (IKE). IKE is based on ISAKMP/Oakley. The terms phases and modes are used to denote the steps involved in setting up an IPSec connection. IKE provides three modes for exchanging key information and setting up SAs. The first two modes are Phase 1 exchanges, which are used to set up the initial secure channel, the IKE SA. The other mode is the Phase 2 exchange, which negotiates IPSec SAs. The two modes in Phase 1 are main mode and aggressive mode, and the Phase 2 mode is called quick mode.

A *Diffie-Hellman* (DH) exponentiation is used to assist in generating a strong initial key. Before IKE will proceed, the potential parties must agree on a way to authenticate themselves to each other. This authentication method is negotiated during the IKE phase main mode exchange. The following mechanisms are in use today: preshared keys, encrypted nonces, and digital certificates. Preshared keys involve the manual installation of the same key on each peer. Encrypted nonces involve the generation of an asymmetric encryption public/private key pair on each peer, and then the manual copying of the public key of each peer to every other peer. Digital certificates involve the use of a trusted third party, called a *certificate authority* (CA), to validate the authenticity of each peer. Digital certificates offer the added benefit of nonrepudiation, meaning that a peer can verify that communication actually took place. After device authentication occurs via the IKE SA, and only when remote-access clients are involved, a second level of user authentication occurs. The headend initiates an *extended authentication* (XAUTH) request to the remote user, prompting him/her for his/her username-password/passcode pair. After this authentication completes in the first phase, the second phase connects the remote and local networks.

Other Technologies that Interact with IPSec

NAT

Network Address Translation (NAT) is commonplace and is used in networking today primarily for three reasons. First, address translation allows for public IP address conservation by sharing a limited set of public addresses among numerous privately addressed devices. Second, in case two enterprises choose the same private address space (say RFC 1918), address translation will allow the devices in each disparate network to intercommunicate with one another. Third, it provides address hiding. NAT takes two forms. The first is a one-to-one translation. For example, 10.10.89.45 is translated into 171.69.235.45. Throughout this document, this type of address translation is simply referred to as Network Address Translation, or NAT. The second form is a many-to-one translation. Each connection initiated by the host is assigned a static port assignment associated with the translated IP address. For example, 10.10.89.45 is translated into 171.69.235.45.4084, where 4084 is the port. Throughout this document this type of translation is referred to as many-to-one address translation.



PMTUD

The *path maximum transmission unit discovery* (PMTUD) mechanism determines what the maximum packet *maximum transmission unit* (MTU) is for a given path and for our case, the maximum MTU handled by the tunnel. Hosts that support PMTUD set the *Don't Fragment* (DF) bit in the IP header. ESP copies the DF bit from the original IP header into the new outer IP header. As the packet travels along the path to its destination, if the encrypted packet is too large to fit over the next link, the router attempting to forward the encrypted packet will send an *Internet-Control-Message-Protocol* (ICMP) message (Type 3-destination unreachable, Code 4-fragmentation needed but DF bit set) to the sending host. The router will also discard the packet. When the host receives the ICMP message, it lowers its MTU so that the packet can successfully traverse the link. This process continues until the packet reaches its destination.

Appendix C: Architecture Taxonomy

VPN termination 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.

Firewall—Stateful packet filtering device that 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.

VPN firewall—Identical to the firewall above this device, this firewall also offers remote-access and site-to-site VPN termination.

VPN concentrator—This VPN-purpose-defined device carries out remote-access and site-to-site VPN termination.

Router—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.

VPN router—Identical to the router above, this device offers site-to-site and remote-access VPN termination.

Remote-access VPN client—A software VPN client that can be installed on a variety of different OSs; it is capable of establishing a tunnel with a VPN termination device to access network resources.

Hardware VPN client—A hardware device that emulates a software VPN client without the need of installation on the users workstation. Based on local user demand, it establishes a tunnel to a VPN termination device to access network resources.

Host IDS—*Host Intrusion Detection System 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 Intrusion Detection System.* 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.

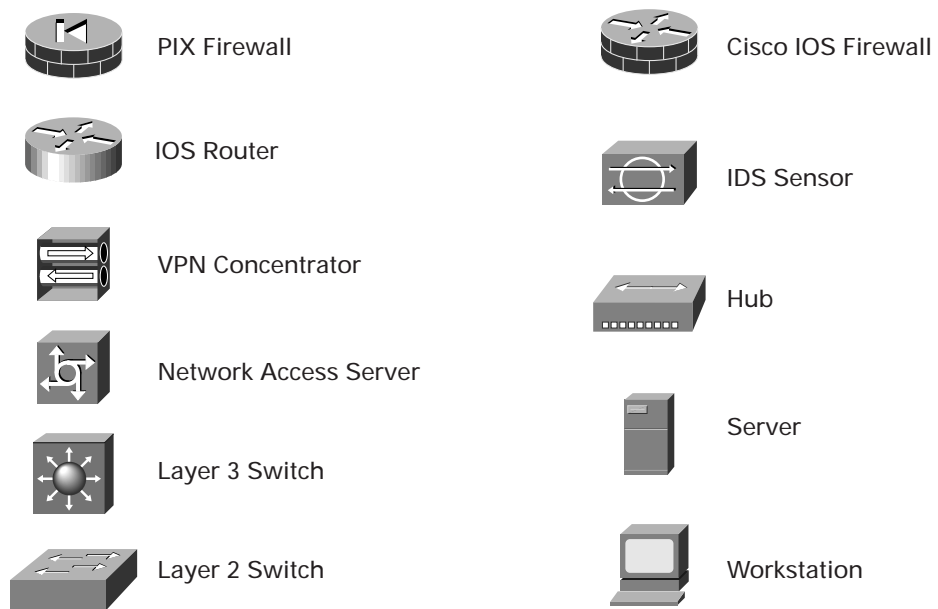
Application server—Provides application services directly or indirectly for enterprise end users. Services can include workflow, general office, and security applications.



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.

Diagram Legend

Figure 27
Diagram Legend



List of Figures

Figure 1	IKE Keepalive High Availability Example	15
Figure 2	Routing Protocol High Availability Example	16
Figure 3	Load Dispersion on Failure	17
Figure 4	Detailed Model of Remote-User Module	19
Figure 5	Detailed Model of Remote-User Module: VPN	20
Figure 6	Detailed Model of Small Network	23
Figure 7	Detailed Model of Small Network Corporate Internet Module	23
Figure 8	Detailed Model of Small Network Corporate Internet Module: VPN	24
Figure 9	Detailed Model of Medium Network	27
Figure 10	Detailed Model of Medium-Network Corporate Internet Module	28
Figure 11	Detailed Model of Medium-Network Corporate Internet Module: VPN	28
Figure 12	Detailed Model of Large Network	31
Figure 13	Detailed Model of VPN and Remote-Access Module	32



List of Figures

Figure 14	Detailed Model of Large Network VPN and Remote-Access Module: VPN	33
Figure 15	Detailed Model of Large Extranet Module	37
Figure 16	Detailed Model of Large Extranet Module: VPN	38
Figure 17	Detailed Model of Large Management Module	41
Figure 18	Detailed Model of Large Management Module: VPN	42
Figure 19	Detailed Model of Distribution VPN Module	43
Figure 20	Detailed Model of Distribution VPN Module: VPN	44
Figure 21	Attack Mitigation Roles for Remote User Networks	51
Figure 22	Small Business Corporate Internet Module VPN	53
Figure 23	Medium Business Corporate Internet Module VPN	55
Figure 24	SAFE Enterprise	58
Figure 25	Large Enterprise VPN/Remote Access Module	59
Figure 26	Distribution Enterprise Hub Module VPN	68
Figure 27	Diagram Legend	87

References

RFCs and Drafts

- RFC 2401 “Security Architecture for the Internet Protocol”
- RFC 2402 “IP Authentication Header”
- RFC 2403 “The Use of HMAC-MD5-96 within ESP and AH”
- RFC 2404 “The Use of HMAC-SHA-1-96 within ESP and AH”
- RFC 2405 “The ESP DES-CBC Cipher Algorithm with Explicit IV”
- RFC 2406 “IP Encapsulating Security Payload (ESP)”
- RFC 2407 “The Internet IP Security Domain of Interpretation for ISAKMP”
- RFC 2408 “Internet Security Association and Key Management Protocol (ISAKMP)”
- RFC 2409 “The Internet Key Exchange (IKE)”
- RFC 2410 “The NULL Encryption Algorithm and Its Use with IPsec”
- RFC 2411 “IP Security Document Roadmap”
- RFC 2412 “The OAKLEY Key Determination Protocol”
- RFC 1918 “Address Allocation for Private Internets”
<http://www.ietf.org/rfc/rfc1918.txt>
- RFC 1191 “Path MTU Discovery”
<http://www.ietf.org/rfc/rfc1191.txt>

Miscellaneous References

- VoIP Bandwidth Consumption
http://www.cisco.com/warp/public/788/pkt-voice-general/bwidth_consume.htm
- QoS
http://www.cisco.com/warp/customer/cc/pd/iosw/prodlit/iosq_ds.htm
- Deployment Guide: Deploying Cisco IOS Security with a Public-Key Infrastructure
http://www.cisco.com/en/US/tech/tk583/tk849/technologies_white_paper09186a00800e79cb.shtml

Partner Product References

- RSA SecureID OTP System
<http://www.rsasecurity.com/products/secureid/>

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. Many thanks to Rahimulah Rahimi who helped to build out much of the equipment in these designs. Thank you all for your special effort.



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 International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

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, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the
Cisco 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

All contents are Copyright © 1992–2004 Cisco Systems, Inc. All rights reserved. Catalyst, Cisco, Cisco IOS, Cisco Systems, the Cisco Systems logo, EtherChannel, SMARTnet, and SwitchProbe are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company.
(0304R) DB/LW5665 02/04