



Network Design Considerations

This chapter provides an overview of the business scenarios covered in this guide, items you should consider when configuring a Virtual Private Network (VPN) on your Cisco VPN gateway, and the assumptions this guide makes.

This chapter includes the following sections:

- Overview of Business Scenarios, page 2-1
- Assumptions, page 2-3
- Cisco SAFE Blueprint, page 2-4
- Hybrid Network Environments, page 2-5
- Integrated versus Overlay Design, page 2-6
- Network Traffic Considerations, page 2-7
- Network Resiliency, page 2-15
- VPN Performance Optimization Considerations, page 2-19
- Network Management Considerations, page 2-24

Overview of Business Scenarios

The site-to-site and extranet scenarios explained in this guide provide a remote office and a business partner access to a corporate headquarters network through Generic Routing Encapsulation (GRE) or IP Security Protocol (IPSec) tunnels.


The remote access scenario provides a remote user access to a corporate headquarters network through secure IPSec, Point-to-Point Tunneling Protocol (PPTP), or Layer 2 Tunnel Protocol (L2TP) tunnels. (See Figure 2-1.)

**Note**

For detailed information on configuring network access server (NAS)-initiated access VPNs using the Layer 2 Forwarding (L2F) tunneling protocol, refer to the *Access VPN Solutions Using Tunneling Technology* publication.

In each scenario, a tunnel is constructed, encryption is applied on the tunnel, and different traffic types (for example, IP, User Datagram Protocol [UDP], and Transmission Control Protocol [TCP]) are either permitted or denied access to the tunnel. This controls the level of access the remote office and business partner have to the corporate intranet and secures the data exchanged between the sites.

Figure 2-1 Business Scenarios



The site-to-site VPN business scenario explained in Chapter 3, “Site-to-Site and Extranet VPN Business Scenarios” links the corporate headquarters to a remote office using connections across the Internet. Users in the remote office are able to access resources as if they were part of the private corporate intranet.

The extranet VPN business scenario explained in Chapter 3, “Site-to-Site and Extranet VPN Business Scenarios” builds on the VPN scenario by linking the same corporate headquarters to a business partner using connections across the Internet; however, the business partner is given limited access to the headquarters network—the business partner can access only the headquarters public server.

The remote access VPN business scenario, explained in Chapter 4, “Remote Access VPN Business Scenario” provides a remote user access to the corporate headquarters network through a secure IPSec, PPTP, or L2TP tunnel that is initiated by the remote user running VPN client software on a PC. In this scenario, the user can access the corporate network remotely.

**Note**

Although supported by Cisco VPN gateways, this guide does not explain how to configure your gateway for use with the Cisco Secure VPN Client. For detailed information on client-initiated VPNs using Cisco Secure VPN Client software, refer to the *Cisco Secure VPN Client Solutions Guide* publication. You can access the publication by logging on to Cisco.com and selecting **Technical Documents: Network Security: Cisco Secure VPN Client: Cisco Secure VPN Client Solutions Guide**.

Assumptions

This guide assumes the following:

- You are configuring a service provider transparent VPN, whereby the tunnel endpoints are outside of the service provider network (on the headquarters and remote site routers).
- You are configuring your VPN based on IP, a routing mechanism, cryptography, and tunneling technologies, such as IPSec and GRE.

**Note**

The scenarios in this guide do not explain how to configure certification authority (CA) interoperability on your Cisco VPN gateway. For detailed configuration information on CA interoperability, refer to the “Configuring Certification Authority Interoperability” chapter in the *Security Configuration Guide*.

- You have identified the Cisco IOS firewall features that you plan to configure on your Cisco VPN gateway features. When considering IOS firewall features, you may find it useful to review the “Network Traffic Considerations” section on page 2-7. The business scenarios in this guide explain how to configure extended access lists, which are sequential collections of permit and deny conditions that apply to an IP address.

**Note**

For advanced firewall configuration information, refer to the “Traffic Filtering and Firewalls” section of the *Security Configuration Guide*.

Cisco SAFE Blueprint

Cisco's secure blueprint for enterprise networks (SAFE) primary goal is to provide best practice information to interested parties on designing and implementing secure networks. SAFE serves as a guide to network designers considering the security requirements of their network. SAFE takes a defense-in-depth approach to network security design. This type of design focuses on the expected threats and their methods of mitigation. This strategy results in a layered approach to security, where the failure of one security system is not likely to lead to the compromise of network resources. SAFE is based on Cisco products and those of its partners.

Cisco encourages the audience of this configuration guide to reference the SAFE Blueprint, which is available at the following URL:

www.cisco.com/go/safe.

Refer to the white paper, *SAFE VPN: IPSec Virtual Private Networks in Depth*, for information relevant to network design considerations. While this configuration guide incorporates several key components of the white paper, Cisco recommends referencing it for an expanded discussion in a context relevant to your specific network, such as small, medium, or large network designs, and remote access and VPN modules. It is available at the following URL:

http://www.cisco.com/warp/public/cc/so/cuso/epso/sqfr/safev_wp.htm

In addition to network topology, network design considerations, and configuration examples, the white paper discusses the following topics:

- Overall design best practices
- High availability (failover)
- Scalability
- Performance
- Identity (authentication methods)
- Secure Management
- NAT (Network Address Translation)
- Security
- Quality of Service
- Routing
- Extranet Considerations

Hybrid Network Environments

While Cisco IOS devices are interoperable with non-IOS devices, such as the PIX Firewall, the Cisco VPN 5000, and the Cisco VPN 3000, this configuration guide focuses on IOS headend VPN configurations. For information on configuring a hybrid VPN, refer to the configuration guide for your particular device.

Mixed Device Deployments

In considering a VPN design, it is critical to ascertain interoperability information about all devices. Networking standards exist, but each manufacturer may or may not utilize the standard in the same way.

For example, although IPSec is a documented standard, the Request for Comments (RFCs) that document it has 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, no standard mechanism for IPSec exists to determine tunnel up or down state, and remote peer reachability. For these reasons, check with vendors of both products for Cisco

product 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, and consider the implications of these changes.

Also, in order to ensure interoperability between products from a single vendor, use the same code base across all platforms. Doing so decreases the likelihood of any interoperability issues with products made by the same vendor as changes occur and interoperability with other vendors increases.

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, for example 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.

Also consider the additional resources required to train administrators on how to configure, manage, monitor, and troubleshoot multiple device types.

Integrated versus Overlay Design

An integrated network design is one in which the WAN, VPN, and IOS firewall functions are run on the same device, for example, on a remote site gateway. Integrated network designs are common in remote offices because of their simplicity and manageability.

An overlay design is one in which any single function, or all functions, are separated, as in headend designs. Firewall functionality is usually separate, the WAN and VPN functions are often integrated (meaning that the functions run on the same device), and VPN functionality is frequently separate from the WAN and firewall functions.

The primary advantage of an overlay design in the headend configuration is that the separation of tasks optimizes network performance. Each device may be dedicated to one or two tasks, rather than all three, in a heavy traffic environment. For example, ACLs (Access Control Lists) require a fair amount of CPU utilization. Therefore, performing ACL tasks on a device other than the gateway allows the gateway more power to support network traffic.

Network Traffic Considerations

Cisco IOS is feature-rich software. However, if improperly used, these features can degrade the flow of VPN traffic. This section provides a discussion of when and how to use several Cisco IOS options to maximize VPN performance, and includes the following topics:

- Dynamic versus Static Crypto Maps
- Digital Certificates versus Pre-shared Keys
- Generic Routing Encapsulation Inside IPsec
- Network Address Translation
- Quality of Service
- Network Intrusion Detection System

Dynamic versus Static Crypto Maps

Cisco recommends using static crypto maps on headend devices whenever possible. Remember that a tunnel being established from a dynamic crypto map can only be originated from the remote end. If devices must be remotely managed, static maps should be used, as the headend cannot establish a tunnel when using dynamic crypto maps.

In network environments in which the remote IP addresses are unknown (such as remote users using dial-up, cable, or DSL), however, dynamic maps must be used. Additionally, dynamic maps can be used for configuration simplicity. They simplify configuration because a crypto map statement is not required for each IP address range. Digital certificates are also highly recommended with the use of

dynamic crypto maps. Dynamic cryptographic maps accept only incoming IKE requests. Because dynamic maps cannot initiate IKE requests, it is not always guaranteed that a tunnel exists between the remote device and the headend site.

This problem can be mitigated by configuring a protocol like Network Time Protocol (NTP) on remote peers to ensure that the tunnel has been established. When a protocol such as NTP or SNMP generates traffic to the headend, it forces IPSec tunnel establishment from the remote end, since the time server is at the headend. Forcing tunnel establishment from the remote end allows the use of dynamic crypto maps, while ensuring that an IPSec tunnel exists. If you use static crypto maps, you are assured that an IPSec tunnel exists, and do not need to configure establishment from the remote end.

Another consideration is that dynamic crypto maps decrease VPN security, as they accept IKE requests from any IP address.

Static cryptographic map configurations include the static IP addresses of the remote peers, and are therefore more secure. The lack of ambiguity associated with static maps also allows a faster traffic flow.

Digital Certificates versus Pre-shared Keys

Digital certificates (DCs) simplify authentication, and increases VPN performance. You need only enroll each peer with the CA, rather than manually configuring each peer to exchange keys. Cisco recommends using digital certificates especially in site-to-site networks of more than 50 peers. Digital certificates offer the added security and network management benefit of nonrepudiation, meaning that a peer can verify that communication actually took place.

In addition to easing the flow of network traffic, digital certificates offer inherent benefits over pre-shared keys. Compromised pre-shared keys are susceptible to man-in-the-middle attacks. With the key, a hacker can connect to any device in your network allowed by the remote-site access policy. Digital certificates scale better than unique pre-shared keys because they allow any device to authenticate to any other device. Digital certificates are not tied to IP addresses, but to unique, signed information on the device that is validated by the enterprise 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.

Wildcard pre-shared keys should not be used for site-to-site device authentication. When using wildcard pre-shared keys, every device in the network uses the same key. If a single device in your network is compromised and the wildcard pre-shared key has been determined, all the devices are then compromised.

Devices generating digital certificates or validating received certificates during tunnel authentication and establishment must know the correct time of day (preferably Coordinated Universal Time [UTC]). Time also determines 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. This is the only revocation scheme for digital certificates compared to pre-shared keys that are simply removed from the uncompromised devices.

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 rather than an enterprise managed CA might facilitate deploying an extranet VPN. Third party CAs include Microsoft, Verisign, Baltimore, and Entrust.

If you specify digital certificates as the authentication method in a policy, the CA must be properly configured to issue certificates. You must also configure the peers to obtain certificates from the CA. Configure this certificate support as described in the “Configuring Certification Authority Interoperability” chapter of the *Security Configuration Guide*.

Generic Routing Encapsulation Inside IPsec

Generic routing encapsulation (GRE) is best suited for site-to-site VPNs because it supports routing updates, multiprotocol, and multicast traffic. Packets are first encapsulated by GRE, and then encapsulated by IPsec. GRE also allows for a single set of IPsec security associations (SAs) to tunnel traffic from one site to another. Typically, IPsec requires 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, and does not encrypt packets. Use GRE when you need support for tunneling packets other than IP unicast type.

Cisco recommends using GRE tunnels with IPsec in tunnel mode to improve the flow of network traffic. IPsec in tunnel mode can be used as a tunneling protocol itself for unicast traffic, but not for multicast traffic. Multicast IPsec traffic requires a GRE tunnel, and that IPsec be used in either transport or tunnel mode. Cisco recommends using IPsec in tunnel mode for the best network traffic performance.

Changing these values increases the level of security; at the same time, however, it increases the processor overhead. The default behavior for 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 by carrying out a Diffie-Hellman (DH) exponentiation every time a new quick-mode (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.

IPsec Considerations

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, and 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. Cisco recommends the use of Triple DES (3DES), rather than DES, as it provides stronger encryption. 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 might 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. Keep in mind that the shorter the SA lifetime, the greater the impact on network traffic (see the “IKE Key Lifetimes” section on page 2-20). The use of strong encryption algorithms 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 about the exportation of encryption algorithms, please see <http://www.cisco.com/wwl/export/crypto>.

- Keep in mind the following when configuring IPsec:
 - IPsec works with the following serial encapsulations: High-Level Data Link Control (HDLC), Point-to-Point Protocol (PPP), and Frame Relay. IPsec also works with the GRE and IPinIP Layer 3, L2F, and L2TP tunneling protocols; however, multipoint tunnels are not supported.
 - IPsec and Internet Key Exchange (IKE) must be configured on the router and a crypto map must be assigned to all interfaces that require encryption services of Cisco IOS VPN gateways.
 - When using tunnel mode, IPsec can be applied to unicast IP datagrams only. Because the IPsec Working Group has not yet addressed the issue of group key distribution, IPsec does not currently work with multicasts or broadcast IP datagrams. When using IPsec with GRE or L2TP, this restriction does not apply.

If you use NAT, you should configure static NAT as redundant so that IPsec works properly. Preferably, NAT should occur before the router performs IPsec encapsulation; in other words, IPsec should be working with global addresses. The following section discusses NAT in further detail.

Network Address Translation

Network Address Translation (NAT) can occur before or after IPsec. It is important to realize when NAT will occur, since in some cases NAT might 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, as NAT can have an adverse effect on network traffic flow.

NAT After IPsec

You might 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 through encryption. 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 occurs in cases where IP address conservation is taking place. This is, in fact, commonplace in hotels, cable and digital subscriber line (DSL) residential deployments, and enterprise networks. In these cases, depending on the type of NAT used, its application might 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 IP header contents, when the IP header changes, the signature checksum is invalidated. In this case, the packet will appear to have been modified in transit and is promptly 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 cases where many-to-one address translation occurs (as in 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, and some devices performing many-to-one NAT do not handle ESP or AH correctly. Remember that ESP and AH are higher-layer protocols on top of IP that do not use ports.

NAT Before IPSec

When two sites are connected through an IPSec tunnel, if any of the network address ranges at each site 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. If the VPN client does not address translate packets correctly or a new application arrives that is not yet supported, the application might not function.

Use address ranges at 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. Do not address hide the public peer addresses of the VPN devices because it provides no real security value-add and may cause connectivity problems.

Quality of Service

The goal of quality of service (QoS) is to provide more efficient and predictable network service by providing dedicated bandwidth, controlled jitter and latency, and improved loss characteristics. QoS achieves these goals by providing tools for managing network congestion, shaping network traffic, using expensive wide-area links more efficiently, and setting traffic policies across the network. QoS prioritizes voice, data, and web traffic to ensure that mission-critical applications get the service they require. Benefits to be derived from QoS include the following:

- Control over resources—You have control over which resources (bandwidth, equipment, wide-area facilities, and so on) are being used. As an example, you can limit the bandwidth consumed over a backbone link by FTP transfers or give priority to an important database access.
- More efficient use of network resources—Using Cisco's network analysis management and accounting tools, you will know what your network is being used for and that you are servicing the most important traffic to your business.
- Tailored services—The control and visibility provided by QoS enables Internet service providers to offer carefully tailored grades of service to their customers.
- Coexistence of mission-critical applications—Cisco's QoS technologies make certain that your WAN is used efficiently by mission-critical applications that are most important to your business; that bandwidth and

minimum delays required by time-sensitive multimedia and voice applications are available; and that other applications using the link get their fair service without interfering with mission-critical traffic.

- Foundation for a fully integrated network in the future—Implementing Cisco QoS technologies in your network now is a good first step toward the fully integrated multimedia network needed in the near future. For example, you can implement weighted fair queuing today and get its immediate benefit of increasing service predictability and IP Precedence signaling for traffic differentiation. You reap additional benefits in the future, because weighted fair queuing is Resource Reservation Protocol (RSVP) enabled, thereby allowing you to take advantage of dynamically signaled QoS from the inevitable coming wave of RSVP-enabled applications.

For a detailed overview of Cisco IOS QoS benefits, features, and application examples, refer to the white paper entitled, “Cisco IOS Software Quality of Service Solutions,” available at the following URL:

http://www.cisco.com/warp/public/cc/techno/protocol/tech/qosio_wp.htm

Network Intrusion Detection System

A Network Intrusion Detection Systems (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 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 guards against any misconfiguration of the VPN devices that could inadvertently allow unencrypted traffic through the device.

Second, NIDS can be used to analyze traffic coming from, or destined to, the VPN device. Here NIDS will detect attacks coming through the VPN from remote sites or remote users. Since we know the traffic origin, and the chances it is spoofed are low, any attack can be met with a strong response from the NIDS. This can include shunning, or TCP resets, as appropriate. NIDS is critical in most VPN environments as most VPN security policies dictate that L3 and L4 access over a VPN should be fairly ubiquitous. This increases the reliance on NIDS to catch and stop most of the attacks from remote sites.

While the benefits of NIDS are compelling, NIDS significantly decreases network throughput, because it inspects every single packet. In a headend environment, consider using alternatives to NIDS. For example, in an overlay network environment (see the “Integrated versus Overlay Design” section on page 2-6), the decrease in performance associated with NIDS can be mitigated by designating a device other than the gateway, such as the Cisco Intrusion Detection System (CIDS), to perform NIDS functions.

Information about CIDS can be found at the following URL:

<http://www.cisco.com/warp/public/cc/pd/sqsw/sqidsz/index.shtml>

Split Tunneling

Split tunneling occurs when a remote VPN user or site is allowed to access a public network (the Internet) at the same time that they access the private VPN network without placing the public network traffic inside the tunnel first. If split tunneling were disabled, the remote VPN user or site would need to pass all traffic through the VPN headend where it could be decrypted and inspected before being sent out to the public network. Therefore, enabling split tunneling can increase the traffic throughput of your VPN, but poses a security risk if the remote user does not have a personal firewall. Despite the benefit of sending less traffic through the VPN gateway, Cisco does not recommend enabling split tunneling unless the remote user has sufficient firewall protection.

Network Resiliency

Network resiliency, or redundancy, enables remote sites to locate another tunneling peer if the primary headend peer is unreachable, or if there is a permanent loss of IP connectivity between peers. Consider network resiliency in both the network configuration and in the decision to use GRE tunnels, IPSec tunnels, or tunnels which utilize IPSec inside GRE. Resiliency can be achieved by properly utilizing and configuring GRE tunnels, IKE keepalives, and Hot Standby Routing Protocol (HSRP) with Reverse Route Injection (RRI).

This section contains the following topics:

- Headend Failover
- GRE

- IKE Keepalives
- RRI with HSRP

Headend Failover

Headend failover ensures that network traffic will be routed through a backup gateway if the primary gateway should fail. GRE and IKE keepalives are the two primary means of attaining headend failover in Cisco IOS VPNs.

GRE

For VPN resilience, the remote site should be configured with two GRE tunnels, one to the primary headend VPN gateway, and the other to the backup headend VPN gateway. If the GRE tunnels are secured with IPSec, each tunnel has its own IKE SA and a pair of IPSec SAs. Since GRE can carry multicast and broadcast traffic, it is possible and very desirable to configure a routing protocol for these virtual links. Once a routing protocol is configured, the failover mechanism comes automatically. The hello/keepalive packets, such as IKE keepalives, sent by the routing protocol over the GRE tunnels provide a mechanism to detect the loss of connectivity. In other words, if the primary GRE tunnel is lost, the remote site will detect this event by the loss of the routing protocol hello packets.

Once virtual-link loss is detected, the routing protocol will choose the next best route; the backup GRE tunnel will be chosen. Hence, the second part of VPN resilience is obtained by the automatic behavior of the routing protocol. Since the backup GRE tunnel is already up and secured, the failover time is determined by the hello packet mechanism and the convergence time of the routing protocol.

Aside from providing a failover mechanism, GRE tunnels provide the ability to encrypt multicast and broadcast packets and non-IP protocols with IPSec. They also provide enhanced performance and scalability for site-to-site VPN services. Since GRE tunnels are unique interfaces, they can each be assigned their own crypto maps. When the headend router needs to send a packet on the VPN, it first makes a routing decision to send it out an interface and then does a search of the SPI table to find the corresponding SA. With GRE tunnels, the router must make a routing decision across a multitude of GRE interfaces. Once the GRE tunnel is chosen, there are only a few SAs to choose from.

GRE tunnels can encapsulate clear text traffic, which enables the passage of routing updates to peer routers. Passage of routing updates provides reachability information between peers. It also enables detection of a secondary peer in the case of a loss of reachability for the primary peer. IPSec can be applied to the GRE tunnel packet to provide encryption for transport security.

IKE Keepalives

IKE keepalives, or hello packets, are a component of IPSec that tracks reachability of peers by sending hello packets between peers. In the case of loss of reachability to a peer, a tunnel is established with a predefined backup or secondary peer.

During the typical life of the IKE Security Association (SA), as defined by the RFCs, packets are only exchanged over this SA when an IPSec quick mode (QM) negotiation is required at the expiration of the IPSec SAs. For a Cisco IOS device, the default lifetime of an IKE SA is 24 hours and that of an IPSec SA is one hour. There is no standards-based mechanism for either type of SA to detect the loss of a peer, except when the QM negotiation fails. These facts imply that for IOS defaults, an IPSec termination point could be forwarding data into a black hole for as long as one hour before the protocol detects a loss of connectivity.

By implementing a keepalive feature over the IKE SA in Cisco IOS software, Cisco has provided network designers with a simple and non-intrusive mechanism for detecting loss of connectivity between two IPSec peers. The keepalive packets are sent every 10 seconds by default. Once three packets are missed, an IPSec termination point concludes that it has lost connectivity with its peer.

To reestablish connectivity, the IPSec termination point must have at least two IPSec peer addresses in its crypto map statement. The IPSec termination point will send out a main mode (MM) request to initiate the MM and quick mode (QM) negotiations with the second peer in its list. This type of functionality is available in all IOS devices that support the IPSec feature set.

IKE keepalives are suggested for use with devices that do not support GRE.

RRI with HSRP

In environments where redundant VPN devices using IKE keepalives for resiliency are present, be sure to track which device has the active IPsec connection with a remote peer to ensure tunnels are not duplicated across devices. Duplication of tunnels results in a mismatch of IPsec policy and the dropping of traffic. RRI and HSRP are two IOS features which, when used together, increase the resiliency of networks using IKE keepalives.

VPN Reverse Route Injection (RRI) is a new IOS feature that resolves the duplicate tunnel problem by injecting a static route for advertisement on the network. It is based on which device currently holds the IPsec session for a specific peer. Advertising this route ensures return IPsec traffic associated with the specific session will be routed through the device that has the active IPsec session.

The primary benefits of RRI are that it enables the routing of IPsec traffic to a specific VPN headend device in environments with multiple (redundant) VPN headend devices, and ensures predictable failover time of remote sessions between headend devices when using IKE keepalives.

HSRP complements the new RRI feature in attaining network resiliency. Using HSRP, a set of routers work in concert to present the illusion of a single virtual router with a virtual IP address that is linked to real IP addresses. The hosts on the network recognize the virtual router and IP address as the only router and IP address. The set of routers that comprises the virtual router is known as an HSRP group, or a standby group. A single router elected from the group is responsible for forwarding the packets that hosts send to the virtual router. This router is known as the active router. Another router is elected as the standby router. In the event that the active router fails, the standby router assumes the packet forwarding duties of the active router. Although an arbitrary number of routers may run HSRP, only the active router forwards the packets sent to the virtual router.

To minimize network traffic, only the active and the standby routers send periodic HSRP messages once the protocol has completed the election process. If the active router fails, the standby router takes over as the active router. If the standby router fails or becomes the active router, another router is elected as the standby router. RRI then informs peers of the active router, ensuring that peers use the active tunnel that HSRP has established.

While HSRP and RRI can be used in conjunction with each other for maximum network resiliency, they can also be used separately.

VPN Performance Optimization Considerations

Several key considerations can maximize the performance of your VPN. For a further discussion of each subject, you can read the referenced documentation.

This section contains the following topics:

- Generic Switching Paths
- Fragmentation
- IKE Key Lifetimes
- IKE Keepalives

Generic Switching Paths

Choose the best switching path available (from fastest to slowest): CEF, optimum, or fast. Enabling CEF will lead to the best performance. If you configure multiple switching paths such as fast-switching and CEF on the same interface, the router will try all of them from best to worst (starting from CEF and ending with process-switching). Choosing one switching path will increase network performance by eliminating the CPU overhead associated with trying all of them.

Fragmentation

Avoid fragmentation at all costs. Packet reassembly is resource intensive from a CPU and memory allocation perspective, and decreases network performance. Allowing fragmented packets into your network also creates security concerns. Fragmented IPSec packets require reassembly before the packets can undergo integrity validation and decryption.

Fragmentation can typically be avoided, as it usually occurs when an encapsulated packet, sent over a tunnel, is too large to fit on the smallest link on the tunnel path. As long as filtering does not block the Internet Control Message Protocol (ICMP) messages, path maximum transmission unit discovery (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 occurs and is out of your administrative control, you will have to either 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. In this scenario, packets generated by hosts that do not support PMTUD, and have not set the DF bit in the IP header, will undergo fragmentation before IPSec encapsulation. Packets generated by hosts that do support PMTUD will use it locally to match the statically configured MTU on the tunnel. If 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 example, 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. Use of ESP and GRE to support PMTUD reduces the likelihood of fragmentation.

Depending on the VPN termination device, the manner in which you should set the MTU on the tunnel varies. Options include changing the MTU through 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).

IKE Key Lifetimes

When IKE begins negotiations, the first thing it does is agree upon the security parameters for its own session. The agreed-upon parameters are then referenced by an SA at each peer. The SA is retained by each peer until the SA's lifetime expires. Before an SA expires, it can be reused by subsequent IKE negotiations, which can save time when setting up new IPSec SAs. New SAs are negotiated before current SAs expire.

To save setup time for IPSec, and thereby optimize VPN performance, configure a longer IKE SA lifetime. However, the shorter the lifetime, the more secure the IKE negotiation is likely to be.

Note that when your local peer initiates an IKE negotiation between itself and a remote peer, an IKE policy can be selected only if the lifetime of the remote peer's policy is shorter than or equal to the lifetime of the local peer's policy. Then, if the lifetimes are not equal, the shorter lifetime will be selected. To restate this behavior:

If the two peer's policies' lifetimes are not the same, the initiating peer's lifetime must be longer and the responding peer's lifetime must be shorter, and the shorter lifetime will be used.

IKE Keepalives

IKE keepalive settings can aid in optimizing VPN performance. By Cisco IOS default, keepalives are sent in 10 second intervals. A longer interval between keepalives reduces CPU usage, thereby increasing network performance. There is, however, a trade-off. The longer the interval, the longer it will take to detect a loss of connectivity. This risk can be mitigated by implementing RRI and/or HSRP. Refer to the “Network Resiliency” section on page 2-15, for a discussion of RRI and HSRP failover mechanisms.

Practical VPN Suggestions

The following are additional considerations you might implement when configuring a VPN on your Cisco VPN gateway:

- **Syslog**—Set up a syslog host, such as a CiscoWorks Essentials Workstation, and configure all the routers in the network to use the syslog host. By logging all syslog messages from the routers, you can determine when significant events, like configuration changes, occurred.
- **Telnet and console access**—In client-initiated or NAS-initiated access VPN environments, implement TACACS+ or Remote Access Dial-In User Service (RADIUS) security for Telnet and console access to the router. Doing so logs all access to the router. The addition of access lists to only allow Telnet access from particular source IP addressees helps to secure the router.
- **Access lists**—Use access list numbers and names consistently to help manage and troubleshoot configurations.

- Template configurations—Use a configuration template when deploying many routers that require consistent configurations.
- Tunneling—Observe the following when configuring tunneling:
 - To avoid anomalies that occur on physical interfaces, configure each tunnel source and destination on a loopback interface. A loopback interface is a virtual interface that is always up and allows routing protocols to stay up even if the physical interface is down.
 - Process switching and fast switching of the GRE, IPsec, L2F, and L2TP tunneling protocols, and Cisco Express Forwarding (CEF) of the IPsec tunneling protocol is supported on Cisco 7100 series gateways in Cisco IOS Release 12.0(4)XE or a later 12.1E software release, or Cisco IOS Release 12.0(6)T or a later 12.0 T software release.
 - Be careful not to violate access control lists. You can configure a tunnel with a source and destination that are not restricted by firewall routers.
 - Routing protocols that make their decisions based solely on hop count will often prefer a tunnel over a multipoint real link. A tunnel might appear to be a one-hop, point-to-point link and have the lowest-cost path, but may actually cost more.
- Firewall—Observe the following when configuring Cisco IOS firewall features (when configuring your Cisco VPN gateway as a firewall):
 - When setting passwords for privileged access to the firewall, use the **enable secret** command rather than the **enable password** command, which does not have as strong an encryption algorithm.
 - Configure a password on the console port. In authentication, authorization, and accounting (AAA) environments, use the same authentication for the console as for elsewhere. In a non-AAA environment, at a minimum, configure the **login** and **password password** commands.
 - Think about access control *before* you connect a console port to the network in any way, including attaching a modem to the port. Be aware that a *break* on the console port might render total control of the firewall, even with access control configured, to a hacker.
 - Apply access lists and password protection to all virtual terminal ports. Use access lists to limit who can Telnet into your router.

- Do not enable any local service (such as Simple Network Management Protocol [SNMP] or Network Time Protocol [NTP]) that you do not plan to use. Cisco Discovery Protocol (CDP) and NTP are on by default, and you should turn these off if you do not need them.

To turn off CDP, enter the **no cdp run** global configuration command. To turn off NTP, enter the **ntp disable** interface configuration command on each interface not using NTP.

If you must run NTP, configure NTP only on required interfaces, and configure NTP to listen only to certain peers.

Any enabled service could present a potential security risk. As determined, hostile party might be able to find creative ways to misuse the enabled services to access the firewall or the network.

For local services that are enabled, protect against misuse. Protect by configuring the services to communicate only with specific peers, and protect by configuring access lists to deny packets for the services at specific interfaces.

- Protect against spoofing: protect the networks on both sides of the firewall from being spoofed from the other side. You could protect against spoofing by configuring input access lists at all interfaces to pass only traffic from expected source addresses and to deny all other traffic.

You should also disable source routing. For IP, enter the **no ip source-route** global configuration command. If you disable source routing at *all* routers, it helps prevent spoofing.

You should also disable minor services. For IP, enter the **no service tcp-small-servers** and **no service udp-small-servers** global configuration commands.

- Prevent the firewall from being used as a relay by configuring access lists on any asynchronous Telnet ports.
- Normally, you should disable directed broadcasts for all applicable protocols on your firewall and on all your other routers. For IP, use the **no ip directed-broadcast** command. Rarely, some IP networks do require directed broadcasts; if this is the case, do not disable directed broadcasts.

Directed broadcasts can be misused to multiply the power of denial-of-service attacks, because every denial-of-service packet sent is broadcast to every host on a subnet. Furthermore, some hosts have other intrinsic security risks present when handling broadcasts.

- Configure the **no proxy-arp** command to prevent internal addresses from being revealed. (This is important to do if you do not already have NAT configured to prevent internal addresses from being revealed).
- Whenever possible, keep the firewall in a secured (locked) room.

To access the documentation for the applications discussed in this section on Cisco.com, refer to the following URL:

<http://cisco.com/univercd/cc/td/doc/product/rtrmgmt/index.htm>

Network Management Considerations

This section contains the following topics:

- Tunnel Endpoint Discovery
- IPsec MIB and Third Party Applications

Tunnel Endpoint Discovery

Tunnel Endpoint Discovery (TED) enhances the IPsec feature. Defining a dynamic crypto map allows you to be able to dynamically determine an IPsec peer; however, only the receiving router has this ability. With TED, the initiating router can dynamically determine an IPsec peer for secure IPsec communications.

TED allows IPsec to scale to large networks by reducing multiple encryptions, reducing the setup time, and allowing for simple configurations on participating peer routers. Each node has a simple configuration that defines the local network that the router protects and the required IPsec transforms.

TED mechanisms best function in partially or fully meshed networks, which require spoke-to-spoke connectivity on an infrequent basis.

IPSec MIB and Third Party Applications

The IPSec Management Information Base (MIB) feature allows users to configure and monitor their IPSec MIB tunnel tables and their trap notifications using Simple Network Management Protocol (SNMP). Utilizing a MIB can increase the performance of your network. It automates the gathering and organization of network management data, which would otherwise add significant CPU overhead to the gateway.

This feature allows users to specify the desired size of a tunnel history table or a tunnel failure table. The history table archives attribute and statistic information about the tunnel; the failure table archives tunnel failure reasons along with the time failure occurred. A failure history table can be used as a simple method to distinguish between a normal and an abnormal tunnel termination. That is, if a tunnel entry in the tunnel history table has no associated failure record, the tunnel must have terminated normally. However, a tunnel history table does not accompany every failure table because every failure does not correspond to a tunnel. Thus, supported setup failures are recorded in the failure table, but an associated history table is not recorded because a tunnel was never set up.

This feature also allows a router to send IPSec trap notifications, which are MIB related, to a random or specified host. A trap notification may be sent when a particular event, such as an error, occurs.

The primary benefit of IPSec MIB is that trap notifications can be sent only once and are discarded as soon as they are sent, thereby reducing traffic and creating lower overhead on your network. Third party MIB applications are available to monitor and control the management information base. One such example is HP Openview, which is a component of several Cisco network management products.

