



IPsec VPN WAN Design Overview

This design guide defines the comprehensive functional components that are required to build a site-to-site virtual private network (VPN) system in the context of enterprise wide area network (WAN) connectivity. This design overview defines, at a high level, the available design choices for building an IPsec VPN WAN, and describes the factors that influence the choice. Individual design guides provide more detailed design and implementation descriptions for each of the major design types.

This design overview is part of an ongoing series that addresses VPN solutions using the latest VPN technologies from Cisco, and based on practical design principles that have been tested to scale.

Contents

Introduction	4
Target Audience	6
Scope of Work	6
Design Guide Structure	7
IP Security Overview	8
Introduction to IPsec	8
Tunneling Protocols	9
IPsec Protocols	9
Encapsulating Security Protocol	9
Authentication Header (AH)	10
Using ESP and AH Together	10
IPsec Modes	11
Tunnel Mode	11
Transport Mode	11
Internet Key Exchange	12
Security Association	12
IKE Phase One	13



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

Copyright © 2006 Cisco Systems, Inc. All rights reserved.

- IKE Phase Two 15
- Fragmentation Issues 16
 - Setting MTU on Client and Server Network Interface Cards 17
 - Path MTU Discovery 17
 - Interface MTU 17
 - Look Ahead Fragmentation 17
 - TCP Maximum Segment Size 17
- Why Customers Deploy IPsec VPNs 18
 - Business Drivers 18
 - Bandwidth 19
 - Cost Reduction 19
 - Security 19
 - Deployment Flexibility 19
 - Resiliency 19
 - Customer Requirements 20
 - Encryption 20
 - IKE Authentication 20
 - Quality of Service 21
 - Interface Level 21
 - Connection or Session Level 21
 - IP Multicast 22
 - Non-IP Protocols 22
 - Routing 23
 - Dynamically Addressed Remotes 23
 - High Availability 23
 - Headend Failure 23
 - Site Failure 24
 - Branch Office Failure 24
 - Stateful versus Stateless Failover 24
 - Integrated Security 24
 - Dynamic Meshing 25
 - Scalability 25
 - Provisioning and Management 25
 - Understanding the Technologies 25
 - Touchless Provisioning 26
 - Ongoing Management 26
 - Service Provider 26
 - Design Selection 27
 - IPsec Direct Encapsulation Design 27

Design Overview	27
Advantages	28
Disadvantages	28
Most Common Uses	28
Point-to-Point GRE over IPsec Design	29
Headend Architecture—Single Tier Headend versus Dual Tier Headend	29
Design Overview	29
Advantages	30
Disadvantages	31
Most Common Uses	31
Dynamic Multipoint VPN—Hub-and-Spoke Topology Design	31
Headend Architecture—Single Tier Headend versus Dual Tier Headend	32
Design Overview	33
Advantages	34
Disadvantages	34
Most Common Uses	34
Dynamic Multipoint VPN—Spoke-to-Spoke Topology Design	35
Design Overview	35
Advantages	36
Disadvantages	36
Most Common Uses	37
Virtual Tunnel Interface Design	37
Design Overview	37
Advantages	39
Disadvantages	39
Most Common Uses	39
Design Comparison	40
Major Feature Support	40
Platform Support	40
Selecting a Design	41
Scaling a Design	42
Critical Scalability Criteria	42
Number of Branch Offices	42
Connection Speeds	43
IPsec Throughput	43
Routing Peers	44
Quality of Service	45
High Availability	45
IP Multicast	45
Internet Access Strategy	46


Integrated Services	47
Appendix A—Evaluating Design Scalability	48
Test Methodology	48
Traffic Mix	48
Finding Limits	49
Conservative Results	50
Cisco Platforms Evaluated	50
Appendix B—References and Recommended Reading	51
Appendix C—Acronyms	52

Introduction

This document serves as a design guide for those intending to deploy a site-to-site VPN based on IP Security (IPsec). The designs presented in this document focus on Cisco IOS VPN router platforms.


The primary topology described in this document is a hub-and-spoke design, where the primary enterprise resources are located in a large central site, with a number of smaller sites or branch offices connected directly to the central site over a VPN. A high-level diagram of this topology is shown in [Figure 1](#).

Figure 1 Hub-and-Spoke VPN Topology



The introduction of dynamic multipoint VPN (DMVPN) makes a design with hub-and-spoke connections feasible, as well as the ability to create temporary connections between spoke sites using IPsec encryption. This topology is shown in [Figure 2](#).

Figure 2 DMVPN Spoke-to-Spoke VPN Topology



This design guide begins with an overview of various VPN solutions, followed by critical selection criteria as well as a guide to scaling a solution. Finally, a platform overview is presented.

Target Audience

This design guide is targeted at systems engineers to provide guidelines and best practices for customer deployments.

Scope of Work

The following design topologies are currently within the scope of this design guide:

- IPsec Direct Encapsulation
- Point-to-Point (p2p) Generic Route Encapsulation (GRE) over IPsec
- Dynamic Multipoint VPN (DMVPN)
- Virtual Tunnel Interface (VTI)

The following major features and services are currently within the scope of this design guide:

- Dead Peer Detection (DPD)
- Reverse Route Injection (RRI)
- Internet Key Exchange (IKE) authentication using digital signatures or certificates
- Cisco VPN routers running Cisco IOS
- EIGRP and OSPF as dynamic Interior Gateway Protocol (IGP) routing protocols across the VPN

- Quality of service (QoS) and Voice and Video Enabled IPsec VPN (V3PN)
- Hot Standby Routing Protocol (HSRP) and Stateful Switchover (SSO) as appropriate for high availability
- IP multicast services over the VPN

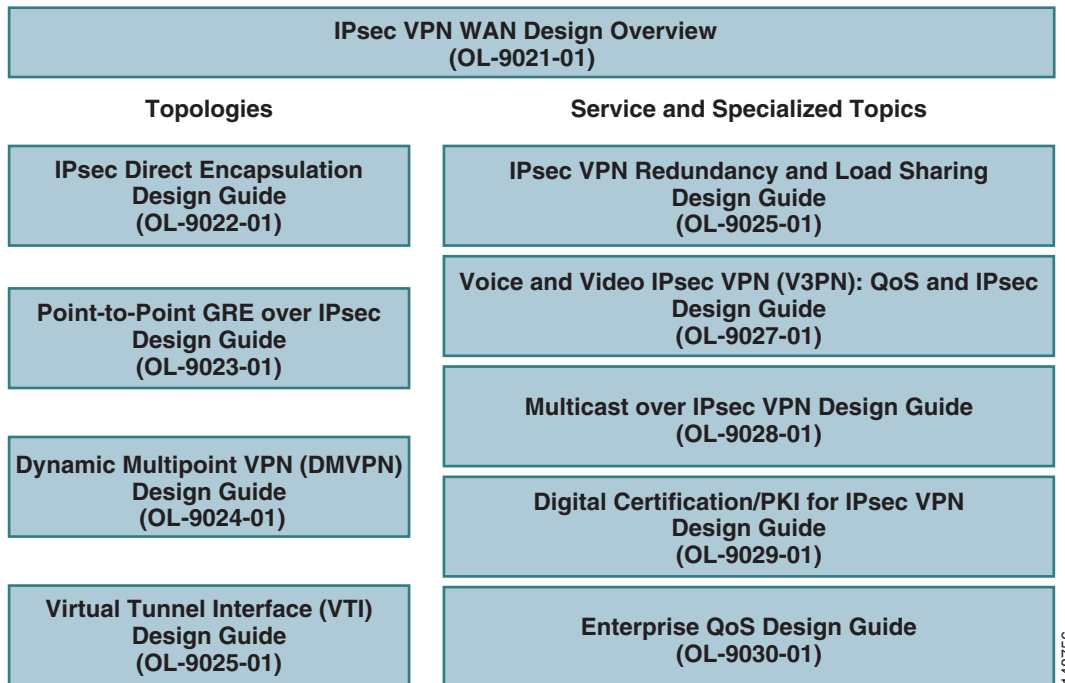
The following features and services are currently outside the scope of this design overview and the design guides it provides:

- Easy VPN authentication and design topology
- Cisco non-IOS platforms including PIX Series and VPN3000 Series
- Remote access applications (client-based)
- Layer 2 tunneling protocols such as Layer 2 Tunneling Protocol (L2TPv3), Point-to-Point Tunneling Protocol (PPTP), and WebVPN (SSL/TLS VPNs)
- MPLS-based VPNs
- Network Management

Design Guide Structure

This design overview is part of a series of design guides, each based on different technologies for the IPsec VPN WAN architecture. (See [Figure 3](#).) Each technology uses IPsec as the underlying transport mechanism for each VPN.

Figure 3 IPsec VPN WAN Design Guides



The operation of IPsec is outlined in this guide, as well as the criteria for selecting a specific IPsec VPN WAN technology.

148756

IP Security Overview

The purpose of this overview is to introduce IP Security (IPsec) and its application in VPNs. For a more in-depth understanding of IPsec, see the Cisco SAFE documentation at the following URL: <http://www.cisco.com/go/safe>.

Introduction to IPsec

The IPsec standard provides a method to manage authentication and data protection between multiple crypto peers engaging in secure data transfer. IPsec includes the Internet Security Association and Key Management Protocol (ISAKMP)/Oakley and two IPsec IP protocols: Encapsulating Security Protocol (ESP) and Authentication Header (AH).

IPsec uses symmetrical encryption algorithms for data protection. Symmetrical encryption algorithms are more efficient and easier to implement in hardware. These algorithms need a secure method of key exchange to ensure data protection. Internet Key Exchange (IKE) ISAKMP/Oakley protocols provide this capability.

This solution requires a standards-based way to secure data from eavesdropping and modification. IPsec provides such a method. IPsec provides a choice of transform sets so that a user can choose the strength of their data protection. IPsec also has several Hashed Message Authentication Codes (HMAC) from which to choose, each giving different levels of protection for attacks such as man-in-the-middle, packet replay (anti-replay), and data integrity attacks.

Tunneling Protocols

Tunneling protocols vary in the features they support, the problems they are designed to solve, and the amount of security they provide to the data being transported. The designs presented in this architecture focus on the use of IPsec as a tunneling protocol alone, and IPsec used in conjunction with Generic Route Encapsulation (GRE) and Virtual Tunnel Interfaces (VTI).

When used alone, IPsec provides a private, resilient network for IP unicast only, where support is not required for IP multicast, dynamic IGP routing protocols, or non IP protocols. When support for one or more of these features is required, IPsec should be used in conjunction with either GRE or VTI.

The p2p GRE over IPsec design allows for all three features described in the preceding paragraph, while a DMVPN design or a VTI design fulfills only the IP multicast and dynamic IGP routing protocol requirements.

Other possible tunneling protocols include the following:

- Secure Sockets Layer/Transport Layer Security (SSL/TLS)
- VPN (WebVPN)
- Point-to-Point Tunneling Protocol (PPTP)
- Layer Two Tunneling Protocol (L2TP)

These protocols are based on user- or client-to-gateway VPN connections, commonly called remote access solutions, and are not implemented in this solution.

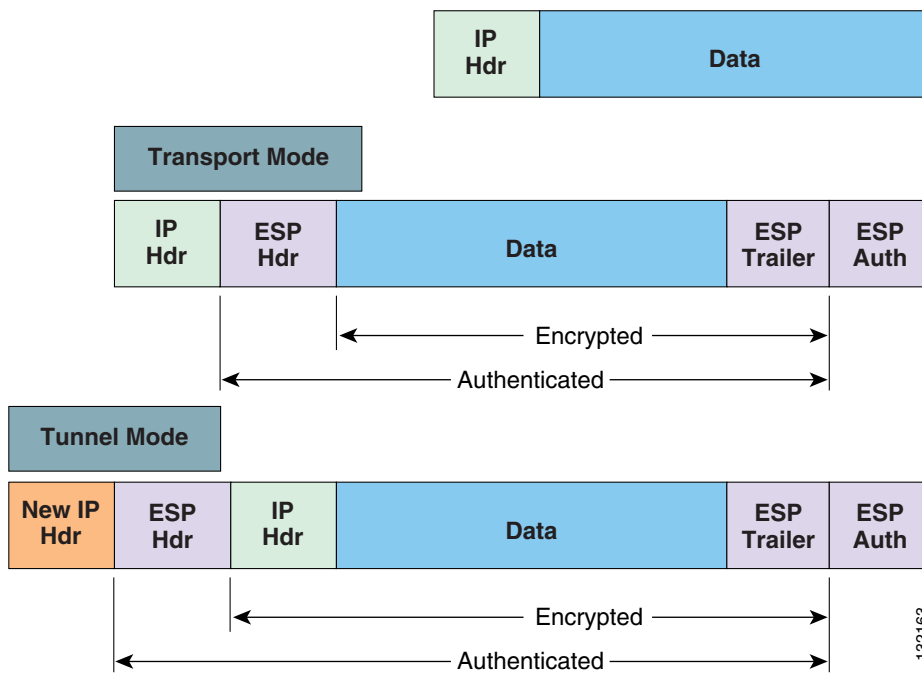
IPsec Protocols

The following sections describe the two IP protocols used in the IPsec standard: ESP and AH.

Encapsulating Security Protocol

The ESP header (IP protocol 50) forms the core of the IPsec protocol. This protocol, in conjunction with an agreed-upon set of security parameters or transform set, protects data by rendering it indecipherable. This protocol encrypts the data portion of the packet only and uses other protections (HMAC) for other protections (data integrity, anti-replay, man-in-the-middle). Optionally, it can also provide for authentication of the protected data. [Figure 4](#) illustrates how ESP encapsulates an IP packet.

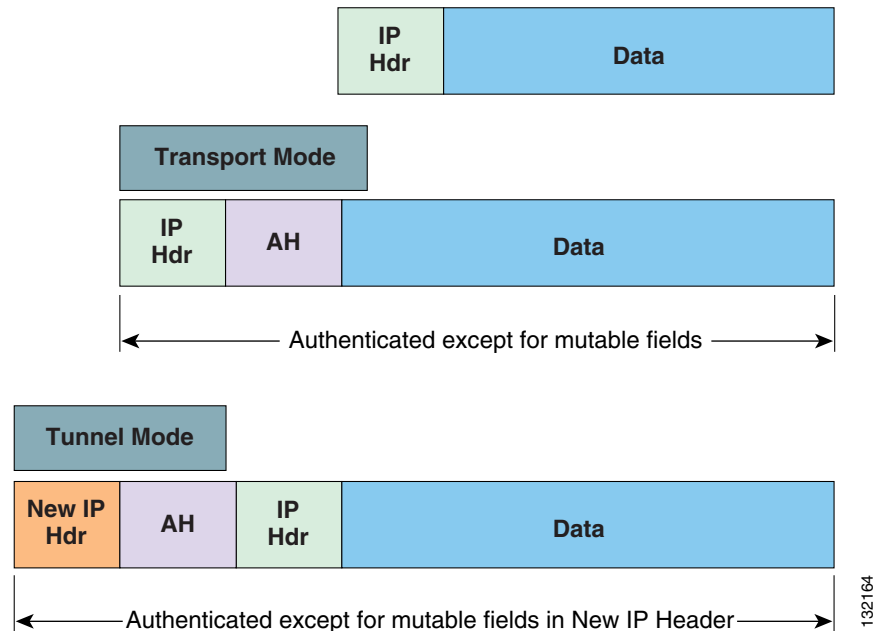
Figure 4 Encapsulating Security Protocol (ESP)



Authentication Header (AH)

The AH protocol (IP protocol 51) forms the other part of IPsec. The AH does not encrypt data in the usual sense, by hiding the data, but it adds a tamper-evident seal to the data. It also protects the non-mutable fields in the IP header carrying the data, which includes the address fields of the IP header. The AH protocol should not be used alone when there is a requirement for data confidentiality. [Figure 5](#) illustrates how AH encapsulates an IP packet.

Figure 5 Authentication Header (AH)



Using ESP and AH Together

It is possible to use ESP and AH together on the same IPsec Security Association (SA). ESP includes the same authentication as AH, as well as providing data encryption and protection. Only the use of ESP alone is shown in the architecture described in this guide.

IPsec Modes

IPsec has the following two modes of forwarding data across a network:

- Tunnel mode
- Transport mode

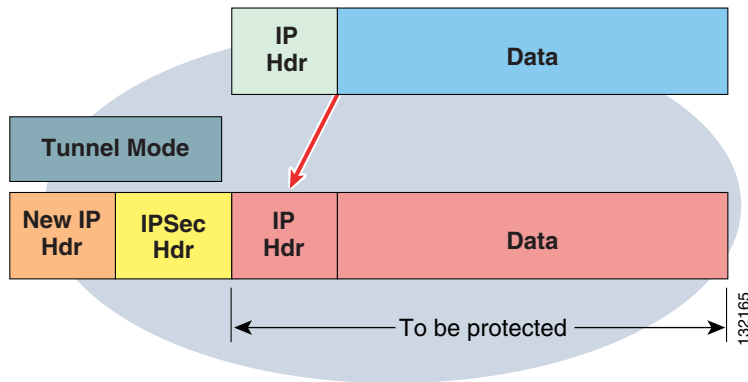
Each differs in its application as well as in the amount of overhead added to the passenger packet. These modes are described in more detail in the next two sections.

Tunnel Mode

Tunnel mode works by encapsulating and protecting an entire IP packet. Because tunnel mode encapsulates or hides the IP header of the pre-encrypted packet, a new IP header is added so that the packet can be successfully forwarded. The encrypting devices themselves own the IP addresses used in this new header. These addresses can be specified in the configuration in Cisco IOS routers. Tunnel mode can be employed with either or both IPsec protocols (ESP and AH). Tunnel mode results in additional packet expansion of approximately 20 bytes because of the new IP header. Tunnel mode is widely considered more secure and flexible than transport mode. IPsec tunnel mode encrypts the source and destination IP addresses of the original packet, and hides that information from the unprotected network. This helps prevent social engineering attacks.

Figure 6 illustrates the expansion of the IP packet.

Figure 6 IPsec Tunnel Mode




Transport Mode

IPsec transport mode works by inserting the ESP or AH header between the IP header and the next protocol or the transport layer of the packet. Both IP addresses of the two network nodes whose traffic is being protected by IPsec are visible in the IP header of the post-encrypted packet. This mode of IPsec can be susceptible to traffic analysis attacks. However, because no additional IP header is added, it results in less packet expansion. Transport mode can be deployed with either or both ESP and AH. Transport mode can be used with p2p GRE over IPsec, because this design hides the addresses of the end stations by adding their own IP header. If the source IP or destination IP address is an RFC 1918 compliant address, the packet cannot be transmitted over the public Internet, and these addresses cannot transit a Network Address Translation (NAT) or Port Address Translation (PAT) device without invalidating the HMAC of the crypto packet.

Figure 7 illustrates the expansion of the IP packet.

Figure 7 IPsec Transport Mode



Internet Key Exchange

To implement a VPN solution with encryption, periodic changing of session encryption keys is necessary. Failure to change these keys makes the VPN susceptible to brute force decryption attacks. IPsec solves the problem with the IKE protocol, which makes use of two other protocols to authenticate a crypto peer and to generate keys. IKE uses a mathematical algorithm called a Diffie-Hellman exchange to generate symmetrical session keys to be used by two crypto peers. IKE also manages the negotiation of other security parameters such as the data to be protected, the strength of the keys, the hash methods used, and whether the packets are protected from anti-replay. ISAKMP normally uses UDP port 500 as both the source and destination port.

Security Association

A Security Association (SA) is an agreement between two peers engaging in a crypto exchange. This agreement includes the type and strength of the encryption algorithm used to protect the data. The SA includes the method and strength of the data authentication and the method of creating new keys for that data protection. Crypto peers are formed as described in the following sections.

Each SA possesses a lifetime value for which an SA is considered valid. The lifetime value is measured in the both time (seconds) and volume (byte count) and is negotiated at SA creation. These two lifetime values are compared, and agreement is reached on the lower of the two. Under normal circumstances, the lifetime value expires via time before the volume limit. Thus, if an interesting packet matches the SA within the final 120 seconds of the lifetime value of an active SA, the crypto re-key process is typically invoked. The crypto re-key process establishes another active SA before the existing SA is deleted. The result is a smooth transition with minimum packet loss to the new SA.

ISAKMP Security Association

An ISAKMP SA is a single bi-directional secure negotiation channel used by both crypto peers to communicate important security parameters to each other, such as the security parameters for the IPsec SA (data tunnel).

In Cisco IOS, the ISAKMP SA policy has a default lifetime value of 86,400 seconds with no volume limit.

IPsec Security Associations (Data Tunnel)

An IPsec SA is a uni-directional communication channel between one crypto peer to another. The actual customer data traverses *only* an IPsec SA, and never over the ISAKMP SA. Each side of the IPsec tunnel has a pair of IPsec SAs per connection; one to the remote, one from the remote. This IPsec SA pair information is stored locally in the SA database.

In Cisco IOS, the IPsec SA policy has a default lifetime value of 3600 seconds with a 4,608,000 Kbytes volume limit.

IKE Phase One

IKE Phase One is the initial negotiation of a bi-directional ISAKMP SA between two crypto peers, often referred to as main mode. IKE Phase One begins with an authentication in which each crypto peer verifies their identity with each other. When authenticated, the crypto peers agree upon the encryption algorithm, hash method, and other parameters described in the following sections to build the ISAKMP SA. The conversation between the two crypto peers can be subject to eavesdropping with minimal risk

of the keys being recovered. The ISAKMP SA is used by the IKE process to negotiate the security parameters for the IPsec SAs. The ISAKMP SA information is stored locally in the SA database of each crypto peer. Table 1 illustrates the various security parameters defined in the following sections.

Table 1 ISAKMP SA Security Parameters

Default in Cisco IOS	Authentication	Encryption	HMAC	Diffie-Hellman	Lifetimes	NAT-T
ISAKMP SA parameters	RSA signatures (PKI) (default)	DES (default)	SHA-1 (default)	Group 1 (default)	86,400 seconds No volume limit (default)	Enabled (default)
	PSK	3DES	MDS	Group 2	User definable	Disabled
	RSA nonce	AES 128		Group 5		
		AES 192				
		AES 256				

Authentication Methods

IKE Phase One has three possible authentication methods: Pre-Shared Keys (PSK), Public Key Infrastructure (PKI) using X.509 Digital Certificates, and RSA encrypted nonces. For the purpose of this architecture, only PSK and PKI with X.509 Digital Certificates are described, but the design is feasible with any of these authentication methods.

Pre-Shared Keys

PSKs are an administrative pre-defined key string in each crypto peer used to identify each other. Using the PSK, the two crypto peers are able to negotiate and establish an ISAKMP SA. A PSK usually contains a host IP address or subnet and mask that is considered valid for that particular PSK. A *wildcard PSK* is special kind of PSK whose network and mask can be any IP address.

Public Key Infrastructure using X.509 Digital Certificates

An alternative to implementing PSK is the use of Public Key Infrastructure (PKI) with X.509 Digital Certificates. Digital Certificates make use of a trusted third party, known as a certificate authority (CA), to digitally sign the public key portion of the encrypted nonce.

Included with the certificate is a name, serial number, validity period, and other information that an IPsec device can use to determine the validity of the certificate. Certificates can also be revoked, which denies the IPsec device the ability to successfully authenticate.

Configuration and management of Digital Certificates is covered in detail in *Digital Certification/PKI for IPsec VPN Design Guide* at the following URL: <http://www.cisco.com/go/srnd>.

Encryption Algorithms

Crypto uses various encryption algorithms. At the core of the encryption algorithm is a shared secret key to authenticate each peer. When authenticated, clear text data is fed into the algorithm in fixed-length blocks and is converted to cipher text. The cipher text is transmitted to the crypto peer using ESP. The peer receives the ESP packet, extracts the cipher text, runs it through the decryption algorithm, and outputs clear text identical to that input on the encrypting peer.

Cisco IOS supports DES, 3DES, AES 128, AES 192, and AES 256 encryption algorithms, with DES designated as the default.

Hashed Message Authentication Codes

The fundamental hash algorithms used by main mode are the cryptographically secure Message Digest 5 (MD5) and Secure Hash Algorithm 1 (SHA-1) hash functions. Hashing algorithms have evolved into Hashed Message Authentication Codes (HMAC), which combine the proven security of hashing algorithms with additional cryptographic functions. The hash produced is encrypted with the private key of the sender, resulting in a keyed checksum as output.

Both MD5 and SHA-1 are supported within Cisco IOS, with SHA-1 designated as the default.

Diffie-Hellman Key Agreement

The Diffie-Hellman key agreement is a public key encryption method that provides a way for two crypto peers to establish a shared secret key that only they know, while are communicating over an insecure channel.

With the Diffie-Hellman key agreement, each peer generates a public and private key pair. The private key generated by each peer is kept secret and never shared. The public key is calculated from the private key by each peer and is exchanged over the insecure channel. Each peer combines the public key of the other with its own private key, and computes the same shared secret number. The shared secret number is then converted into a shared secret key. The shared secret key is never exchanged over the insecure channel.

Diffie-Hellman Groups 1, 2, and 5 are supported within Cisco IOS. Group 1 is the default value, with a key length of 768 bits. Group 2 has a key length of 1024 bits and Group 5 has a key length of 1536 bits.

NAT Transparency (NAT Traversal)

IPsec NAT Transparency (NAT-T) introduces support for crypto peers to travel through NAT or PAT points in the network by encapsulating crypto packets in a UDP wrapper, which allows packets to traverse NAT devices. NAT-T was first introduced in Cisco IOS 12.2(13)T, and is enabled by default as a global command. NAT-T is auto-negotiated between the two crypto peers during ISAKMP negotiation with a destination UDP port of 4500. The source uses the next available higher port. When UDP port 4500 is used, the destination port moves to UDP port 4501, 4502, and so on, until an ISAKMP session is established. NAT-T is defined in RFC 3947.

IKE Phase Two

In IKE Phase Two, the IPsec SAs are negotiated by the IKE process using the ISAKMP bi-directional SA, often referred to as quick mode. The IPsec SAs are uni-directional in nature, causing a separate key exchange for data flowing in each direction. One of the advantages of this strategy is to double the amount of work required by an eavesdropper to successfully recover both sides of a conversation. During the quick mode negotiation process, the crypto peers agree upon the transform sets, hash methods, and other parameters. [Table 2](#) illustrates the various security parameters.

Table 2 IPsec SA Security Parameters

Default in Cisco IOS	Encryption	HMAC	PFS	Lifetimes	IPsec Mode
IPsec SA parameters	DES (default)	SHA-1 (default)	Disabled (default)	3600 seconds 4,608,000 Kbytes (default)	Tunnel mode (default)
	3DES	MD5	Group 1	User definable	Transport mode
	AES 128		Group 2		
	AES 192		Group 5		
	AES 256				

Encryption Algorithms

As in main mode, quick mode uses an encryption algorithm to establish the IPsec SAs. The encryption algorithm negotiated by the quick mode process can be the same or different from that in the main mode process. Cisco IOS supports DES, 3DES, AES 128, AES 192, and AES 256 encryption algorithms, with DES designated as the default.

Hashed Message Authentication Codes

As in main mode, quick mode uses an HMAC to establish the IPsec SAs. The HMAC negotiated by the quick mode process can be the same or different from that in the main mode process. Both MD5 and SHA-1 are supported within Cisco IOS, with SHA-1 designated as the default.

Perfect Forward Secrecy

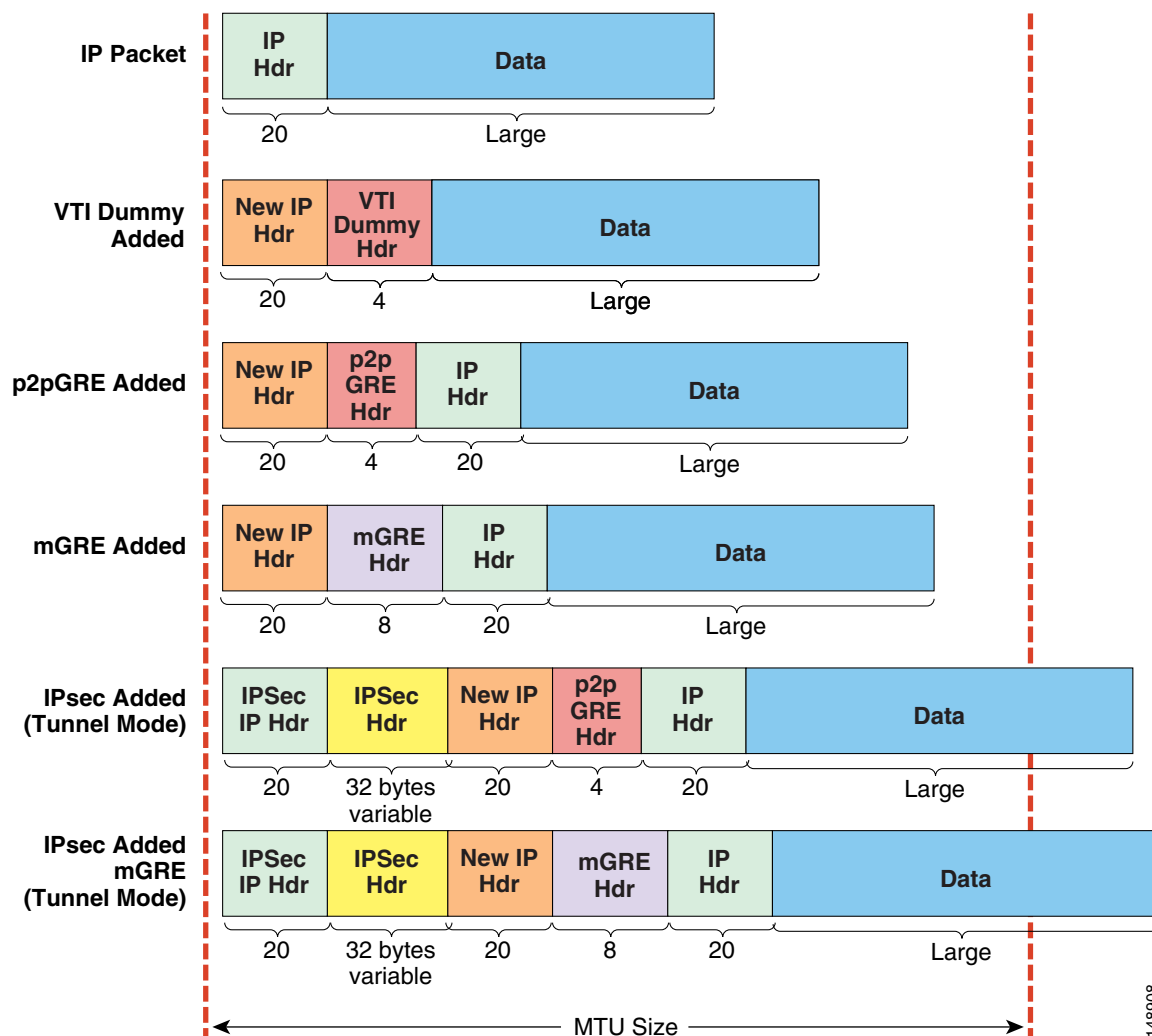
If perfect forward secrecy (PFS) is specified in the IPsec policy, a new Diffie-Hellman exchange is performed with each quick mode negotiation, providing keying material that has greater entropy (key material life) and thereby greater resistance to cryptographic attacks. Each Diffie-Hellman exchange requires large exponentiations, thereby increasing CPU use and exacting a performance cost.

PFS (Diffie-Hellman) Groups 1, 2, and 3 are supported within Cisco IOS. PFS is disabled by default. Group 1 has a key length of 768 bits, Group 2 has a key length of 1024 bits, and Group 5 has a key length of 1536 bits.

Fragmentation Issues

The various IPsec VPN designs use encapsulation of the original IP datagram using one of the following: IPsec Direct Encapsulation design, Point-to-Point GRE over IPsec design, DMVPN (mGRE) design, or VTI design. These encapsulations add to the original packet size. [Figure 8](#) illustrates the various packet expansions.

Figure 8 Various Packet Expansions



148908

When a packet is expanded beyond an interface maximum transmission unit (MTU), the initiating router must fragment the packet before transmission, which means the receiving crypto router must re-assemble the fragments before decryption. The reassembly process is usually performed at the process level, which seriously impacts router performance. Therefore, fragmentation should be avoided if at all possible. There are several options for preventing fragmentation, some of which are configured within Cisco IOS, and some of which require changes to the VPN clients or end stations.

Setting MTU on Client and Server Network Interface Cards

The best way to avoid fragmentation issues in a VPN environment is to manually set the MTU on all client and server Network Interface Cards (NIC) to a smaller value than the Ethernet standard of 1500 bytes. The “tried and true” value to use is 1300 bytes. However, because the average enterprise network has potentially thousands of client workstations, it is not always possible to accomplish this task because of the sheer scale of devices. The second-best strategy is to set the MTU on the NIC on the application servers to 1300 bytes, because these devices are usually in a secure location accessible to network administrators, and because they are often handling the largest packets.

Path MTU Discovery

A feature of IP called Path MTU Discovery (PMTUD) can eliminate the possibility of fragmentation if it is supported by the end stations. This feature can determine the smallest MTU between two end stations to ensure the sender does not transmit a packet that results in fragmentation.

With PMTUD enabled, all packets are sent with the do not fragment (DF) bit set. If a packet encounters a link with a lower MTU than the packet size, an ICMP error message is generated with a 3 in the type field (destination unreachable), a 4 in the code field (fragmentation needed and DF set), and the next hop MTU size in the unused field of the ICMP header. After receiving the ICMP error message, the original sender lowers the MTU of the subsequent packets transmitted.

Interface MTU

Unfortunately, the effectiveness of PMTUD is negated if some device in the transmission path, such as a network or personal firewall, blocks or filters the ICMP messages used. If this is the case, the next fragmentation-avoidance technique is to set the MTU on VPN interfaces to a lower size. Again, the recommended value is 1300 bytes. In designs with GRE tunnels implemented, the configuration is applied to the tunnel interface. However, in Cisco IOS, the tunnel interface default MTU value of 1514 bytes cannot be changed, so changing the IP MTU is the only option. The IP MTU can be changed by using the **ip mtu 1300** command.

Look Ahead Fragmentation

A feature called Look Ahead Fragmentation (sometimes abbreviated LAF and sometimes called *Pre-Fragmentation*) is supported by current versions of Cisco IOS. With Look Ahead Fragmentation enabled, the crypto router looks at the MTU of the outbound crypto interface, evaluates the crypto headers to be added to a packet, and performs fragmentation at the IP level before sending the fragments to the encryption process. The receiving crypto peer decrypts the fragments independently and forwards them to the receiving host for re-assembly of the original packet. In a Cisco IOS router running 12.1(11)E, 12.2(13)T or later, LAF is enabled by default on physical interfaces.

TCP Maximum Segment Size


The TCP maximum segment size (MSS) value influences the resulting size of TCP packets. The majority of data packets on a network are TCP. Other than video, suitable UDP applications (such as DNS and NTP) exhibit an average packet size of less than 300 bytes. Use the router to influence or set the TCP MSS for TCP flows so as to reduce the data packet size. The effect is to reduce the impact of serialization delay, where no Layer 2 Fragmentation and Interleaving (LFI/FRF.12) technique exists.

Before the implementation of PMTUD, the maximum IP packet size for off-net (hosts not on a directly connected interface) was 1300 bytes. The TCP MSS is the number of bytes following the IP and TCP header, so the default MSS size was 1260 bytes. The IP and TCP header are each 20 bytes, so 1300 minus 40 equals 1260.

The MSS option can appear only in a TCP SYN packet, and each end announces its own MSS. Although not required, it is frequently the same in both directions. The recommended behavior changed with the introduction of PMTUD, which allows greater data throughput by transmitting more payload in each packet. This is fine for data, but given the relatively low speed links of broadband connections and lack of LFI, it introduces serialization delay for voice packets. By default, PMTUD is disabled for TCP sessions originated by the router. The Cisco IOS interface configuration command **ip tcp adjust-mss 1260** overrides the value of the MSS option for TCP SYN packets received through that interface, allowing the router to override the host-provided MSS value and substitute one that is optimal.

Figure 9 illustrates an MSS in a packet.

Figure 9 MSS Packet Breakdown



Why Customers Deploy IPsec VPNs

This section describes the motivations and business drivers for customers who are deploying IPsec VPNs as part of their WAN strategy.

Business Drivers

Up to 40 percent of typical enterprise employees work in branch offices, away from the central sites providing mission-critical applications and services required for business operations. As these services are extended to branch office employees, requirements increase for bandwidth, security, and high availability.

Because of the flexibility, security, and cost effectiveness, many customers are deploying IPsec VPNs in their corporate WAN strategy. The most common business drivers are described in the following sections.

Bandwidth

Traditional WANs, such as Frame Relay and ATM, have typically provided 128 Kbps, 256 Kbps, and 512 Kbps connection speeds. As services and advanced applications increase at the branch office, so do bandwidth requirements. Customers are faced with either doubling or tripling their existing WAN circuits, which is often cost prohibitive, or seeking out higher bandwidth alternatives.

Cost Reduction

Often the cost of a relatively high-bandwidth IP connection, such as an ISP connection, IP VPN provider, or broadband DSL/cable access, is lower than existing or upgraded WAN circuits. As a result, many customers are either migrating their primary WAN connectivity to these services, or deploying such WAN alternatives as a secondary high-speed WAN circuit to augment their existing private WAN.

The prevalence of high speed T1 ISP services, as well as broadband cable and DSL access services, are putting tremendous pressure on costs of traditional WAN services.

Security

Regulations such as the Health Insurance Portability and Accountability Act (HIPAA), the Sarbanes-Oxley Act (S-Ox), and the Basel II Agreement (in EMEA) recommend or mandate the need for companies to implement all reasonable safeguards to protect personal, customer, and corporate information.

IPsec VPNs inherently provide a high degree of data privacy through establishment of trust points between communicating devices, and data encryption with the Triple Data Encryption Standard (3DES) or Advanced Encryption Standard (AES) standard.

Customers are using IPsec to encrypt their WAN communications, whether using a private WAN, IP VPN, or the Internet for connectivity.

Deployment Flexibility

Because IPsec VPNs can be quickly established wherever an Internet access connection is available, they offer a great degree of flexibility in connecting branch offices, even in locations which do not offer private WAN or IP VPN business services.

Resiliency

As applications such as Voice over IP (VoIP) and mission-critical business applications are deployed to branch offices, resiliency and high availability become primary concerns. Enterprise customers are faced with duplicating their private WAN circuits to provide a level of redundancy, which can be cost prohibitive.

IPsec VPNs over a high-speed ISP connection or broadband cable/DSL access can provide a very cost-effective secondary WAN connection for branch offices. Many customers continue to route their most critical traffic across their private WAN circuits, and route higher-bandwidth, less critical traffic across IPsec VPNs as a secondary connection path. If a failure occurs of their primary WAN circuit, the IPsec VPN can also function as an established backup path.

Customer Requirements

When relying on an IPsec VPN for their primary or secondary WAN strategy, enterprise customers expect the same functionality, performance, and reliability as with their private WANs. This includes QoS, IP multicast support, and high scalability. This section covers many common enterprise customer requirements.

Encryption

To ensure confidentiality of data transported over the VPN, encryption algorithms are used to encrypt the payload of IP packets. The following are three common encryption standards in use:

- Data Encryption Standard (DES)
- Triple Data Encryption Standard (Triple DES or 3DES)
- Advanced Encryption Standard (AES)

DES is considered least secure, Triple DES is newer and considered more secure, and AES is the newest standard and is considered very secure. Various laws and restrictions govern domestic and international use and export of encryption technology.

All Cisco VPN router platforms, including 870, 1800 ISR, 2800 ISR, 3800 ISR, 7200VXR (with SA-VAM2+), and 7600 or Catalyst 6500 (with VPN SPA) support all three encryption standards with hardware-acceleration.

Hardware acceleration of encryption is important for the following two reasons:

- Throughput is greatly improved compared to software-only encryption
- Latency/jitter-sensitive applications, such as VoIP, require hardware acceleration

Testing with hardware acceleration has shown that performance is not significantly affected by choice of encryption method.

IKE Authentication

To ensure the authentication of the IPsec peers, Internet Key Exchange (IKE) is used. Several types of IKE authentication are possible, including the following two most common types:

- Pre-Shared Keys (PSK)
- Public Key Infrastructure (PKI) using X.509 Digital Certificates

For implementations with a small number of branch offices, the choice might be PSK. However, as the number of branches increases, managing individual PSKs is more challenging, and X.509 Digital Certificates are likely a better option.

With the Cisco IOS-CA functionality, a Cisco IOS router can act as a CA, instead of customers having to purchase a more expensive third-party CA server or managed CA service. For more information on using Cisco IOS-CA see the *Digital Certification/PKI for IPsec VPN Design Guide* at the following URL: <http://www.cisco.com/go/srnd>.

Quality of Service

If IPsec VPN designs are proposed as a replacement or supplement to traditional WAN services, customers expect the same level of QoS functionality to be provided. IPsec VPNs and QoS have been integrated in Cisco IOS with the implementation of Voice and Video IPsec Enabled VPN (V3PN). However, there are at least two levels to consider for QoS.

Interface Level


QoS can be very effective in mitigating traffic congestion at an interface level. For example, if a branch office is connected with a T1 access connection, QoS can ensure that the traffic does not exceed the T1 rate, and also prioritize more important or sensitive traffic such as VoIP. Most Cisco platforms readily support QoS, including Class-Based Weighted Fair Queuing (CBWFQ) and traffic shaping at an interface level, otherwise known as Low Latency Queuing (LLQ).

Connection or Session Level


Private WANs, such as Frame Relay, can implement QoS for each connection, such as a permanent virtual circuit (PVC), between a sender and receiver. A service policy can be configured at the PVC level to traffic shape the connection to a matched speed so that a very high-speed headend cannot overrun a lower-speed remote. Figure 10 illustrates this concept.

Figure 10 Traffic Shaping Comparison


Traditional WAN QoS



IPsec VPN without per-tunnel QoS



IPsec VPN with per-tunnel QoS



In the case of IPsec VPNs, a logical tunnel connection exists between sender and receiver that does not have a direct bandwidth specification synonymous with a FR PVC. In addition, IPsec tunnels are not configurable with a QoS service policy, so the only option is to assign QoS at the interface level. However, because IPsec headend routers usually have much higher connection speeds than branch routers, it is possible for the headend to overrun the branch router. This is shown in the middle scenario in Figure 10 above.

There are several options to address this issue:

- Use a WAN transport that provides a sub-interface (such as FR PVC) where QoS can be applied
- Use a service provider that offers QoS services at the provider edge facing the branch office
- Use a branch access that has several times the downlink bandwidth relative to uplink bandwidth, such as 384 Kbps and 2 Mbps cable or DSL
- Implement QoS per VPN tunnel

Although progress is being made, there are still challenges with providing the equivalent QoS guarantees at an IPsec VPN tunnel level. The p2p GRE over IPsec design can be implemented by assigning a QoS service policy (including generic traffic shaping) per p2p GRE tunnel interface. Alternatively, VTI can be implemented that dynamically clones a specified QoS service policy profile per connection. However, the performance of applying a QoS service policy with generic traffic shaping to the tunnel interface is less desirable. When applied, all packets are process switched, which causes high CPU.

Neither of the design options for QoS per VPN tunnel is currently very scalable.

For more information on integration of QoS and IPsec for supporting latency/jitter-sensitive applications, see the *Voice and Video Enabled IPsec VPN (V3PN) Design Guide*. For more generic QoS information, see the *Enterprise Quality of Service SRND*. Both guides are available at the following URL: <http://www.cisco.com/go/srnd>.

IP Multicast

IP multicast is an increasing requirement for many enterprise customers, and presents some challenges for IPsec VPNs.

First, IPsec does not inherently support transport of IP multicast packets, so an encapsulation of this traffic is required with p2p GRE or multipoint GRE (mGRE). More recently, VTI has been implemented in Cisco IOS, which extends IPsec to transport IP multicast without explicitly having to encapsulate the transit packet in a GRE header.

The second challenge involves the problem of IP multicast packet replication and fan-out at the enterprise WAN edge. Each branch office must receive a copy of the IP multicast packet, so the WAN edge router must replicate IP multicast packets for each connection to branch offices. Cisco IOS has been optimized to replicate IP multicast packets very quickly.

However, when IPsec VPNs are deployed, each sender or receiver establishes a trustpoint between them and typically a unique encryption key. Each replicated IP multicast packet is first encapsulated in either a p2p GRE or mGRE header and then encrypted by IPsec with the unique encryption key for each destination. Encrypting such IP multicast fan-outs can be extremely resource-intensive on encrypting routers and VPN acceleration hardware, and can lead to design scalability issues.

For more information on supporting IP multicast applications over an IPsec VPN, see the *IP Multicast over IPsec VPN Design Guide* at the following URL: <http://www.cisco.com/go/srnd>.

Non-IP Protocols

Some customers have requirements for transporting non-IP protocols, such as IPX, over the IPsec VPN. IPsec does not inherently support transport of non-IP packets, so an encapsulation of this traffic is required with p2p GRE over IPsec *only*.

Routing

Many enterprise WANs use IGP dynamic routing protocols such as EIGRP and OSPF to provide routing and maintain link state and path resiliency. All IGP routing protocols use either broadcast or IP multicast as a method of transmitting routing table information.

IPsec does not inherently support transport of broadcast or IP multicast packets, so an encapsulation of this traffic is required with p2p GRE, mGRE, or VTI.

Dynamically Addressed Remotes

Traditional WAN services, such as Frame Relay or ATM, typically rely on static addresses. Increasingly, IP transport options such as high-speed ISP connections, and especially broadband cable and DSL, are being used for primary or alternate WAN connectivity for branch offices.

With some of these access types, static addressing might not exist, or might be more expensive; therefore, it is an increasing customer requirement for IPsec VPN designs to operate with branch offices receiving their addresses via DHCP or PPPoE from the access provider.

Dynamically addressed remotes present some challenges, because there is no fixed address to which to configure a tunnel destination on the headend router. Generally, headend routers must be configured with dynamic crypto maps, and if a tunneling or encapsulation method (such as p2p GRE, mGRE, or VTI) is also being used, headends must be configured to accept dynamic tunneling connections as well.

High Availability

Although high availability is a broad topic that requires several entire design guides to address appropriately, there are several key considerations to understand in the context of an IPsec VPN design. This section explores several forms of high availability and their relationship to IPsec VPNs.

For more information on designing IPsec VPNs for high availability and resiliency, see the *IPsec VPN Redundancy and Load Sharing Design Guide* at the following URL: <http://www.cisco.com/go/srnd>.

Headend Failure

The following are the primary headend points of failure in an IPsec VPN design:

- Headend aggregation routers
- VPN hardware acceleration modules (typically an add-on blade or card)
- WAN routers (can be standalone or integrated into VPN routers)
- WAN connections to the service provider(s)

Any headend device failures affect a large part of the connectivity; therefore, redundancy should be strongly considered. Most Cisco IPsec VPN designs support redundancy for each device type, including VPN card-to-card failover and router-to-router failover.

Another possible failure at the headend is the high-speed connection pipe (WAN) to the service provider. Again, because an outage on this connection can affect a large part of the functionality, multiple connections to service providers should be considered.

Site Failure

Failure of an entire headquarters location is increasingly part of the design requirements of large enterprise customers. IPsec VPN designs can be engineered such that multiple headend aggregation sites that are geographically dispersed provide a high level of resiliency in the event of a total site failure.

Branch Office Failure

The following are primary branch office points of failure in an IPsec VPN design:

- Branch office routers
- VPN hardware acceleration modules (can be on-board or add-on card)
- WAN connections to the service provider(s)

Any branch office device failures affect local functionality; therefore, redundancy should be strongly considered and weighed with costs of an outage for that branch office. Again, most Cisco IPsec VPN designs support redundancy for each device type, including VPN card-to-card failover and router-to-router failover.

Another possible failure at the branch office is the connection pipe to the service provider. Again, because an outage on this connection can affect local functionality, multiple connections to service providers should be considered.

Stateful versus Stateless Failover

Whether the requirement is for stateless or stateful, failover should be considered in each failure mode. Stateless failover maintains no state information of traffic protocol sessions that might be in progress at the time of failure, while stateful failover maintains the traffic state, and continues nearly immediately with the next packet in the session.

Failover to a secondary tunnel via routing protocol convergence is typically stateless, and might require TCP/IP sessions in progress to retransmit or restart. IPsec Direct Encapsulation designs can also support a stateful failover, with the IPsec session state information being exchanged between two headend routers.

Stateful failover can typically minimize loss of functionality for 1–3 seconds, before switching over to the standby headend router. Stateless failover typically requires 20–60 seconds before the routing protocol in use can converge and resume traffic over the alternate path.

Integrated Security

Because IPsec VPNs can be deployed across essentially any IP transport, including traditional WAN (such as FR, ATM), IP VPN, and Internet, integrated security services might be a customer requirement.

For example, if the Internet is being used as a transport, it might be desirable to have integrated firewall, Intrusion Prevention System (IPS), and denial of service (DoS) prevention systems integrated with the IPsec VPN design.

Integration of security tends to be high at branch offices, and is one of the primary advantages of the Cisco Integrated Services Router (ISR). At headend locations, security functions have historically been distributed or dedicated devices, but increasingly integrated security functions are given as customer requirements.

Typically, security functions such as firewalls are relatively intensive computing operations, so the impact on headend or branch office routers should be considered if the same router provides both VPN services and other security services.

Dynamic Meshing

The typical design of a traditional private WAN, as well as commonly deployed IPsec VPN designs, are hub-and-spoke topologies. Branch offices have connections to one or more VPN headend aggregation hubs.

Some enterprise customers might have requirements for direct communication between branch offices, and if these requirements are significant enough, customers might request additional meshing of their IPsec VPN design topology.

If the desired direct branch-to-branch connections are for a few large branch offices in the topology and are fairly well known, additional IPsec VPN connections can be pre-established in the design between these sites.

If the customer has requirements for branch offices to dynamically establish connection paths to other branch offices, Cisco VPN routers and Cisco IOS can provide this functionality. Implications for the WAN topology need to be considered if implementing a meshed topology.

Scalability

The flexibility of IPsec VPNs leads customer expectations for larger-scale aggregation points than they expect from private WANs. Where traditional routed WANs were typically designed to aggregate approximately 200 PVCs, it is common for customers to expect to aggregate 500, 1000, or even 5000 IPsec VPN tunnel connections to one or more hub locations.

Many factors affect scalability of an IPsec VPN design, including access connection speeds, routing peer limits, IPsec encryption engine throughput, and IPsec tunnel termination. How to scale large aggregations while maintaining performance and high availability is challenging, and requires careful planning and design.

See [Scaling a Design, page 41](#) for a more thorough description of scalability considerations for IPsec VPNs.

Provisioning and Management

Because of the flexibility of IPsec VPNs, many enterprise customers expect to deploy these networks to connect relatively large numbers of branch offices. This can present challenges for provisioning and management.

Understanding the Technologies

Most enterprise WAN network staff are versed in routing and private WAN technologies. Many are becoming experienced with QoS as well. However, IPsec has historically been a security and remote access technology, and has most likely been managed by the enterprise security network staff, which might not be familiar with WAN technologies and routing.

Similarly, private WANs are generally considered a trusted medium, while IPsec VPNs are often deployed over untrusted media such as the Internet. WAN network staff might not have much experience with security issues, and InfoSec staff might need to be involved in the design process.

This can present a challenge, in that enterprise customer network staff must understand (or be educated in) a number of technologies to confidentially implement an IPsec VPN as a WAN strategy. Simplifying the planning, design, and deployment needs to be a primary objective.

Touchless Provisioning

Enterprise customers like to configure their VPN headend aggregation routers to allow touchless provisioning of new branch offices. Ideally, customers want to configure a headend router once, and then not have to add new configuration lines to provision new branch offices, which disrupts operation of the VPN headend routers.

To provide touchless provisioning, headend routers need to be configured with dynamic crypto maps or profiles and dynamic tunneling or encapsulation (such as mGRE or VTI) configurations, so that new branch connections can be established without modifying the headend router.

Cisco provides several *touchless* design options, as described in [Design Selection, page 25](#).

Ongoing Management

Enterprise customers might have their own management tools and systems, or might desire a specific management system for the IPsec VPN deployment. If desired, Cisco offers a number of management tools for this purpose, including the following:

- VPN Management System (VMS)
- IP Solutions Center (ISC)

Further descriptions on network management products are outside the scope of this design overview.

Service Provider

VPNs inherently rely on one or more service providers to provide access services to the headend and branch offices to deploy the network. Choosing a service provider is a critical element of deploying an IPsec VPN.

Factors to be considered include cost, services available, reliability, and the expected geographical coverage of the customer VPN. At a minimum, the enterprise should have an SLA with the service provider that outlines the critical service elements of their VPN. These factors include availability, bandwidth, and latency.

Situations that require an enterprise to use multiple service providers to cover their branch locations add a level of complexity, and it can be potentially problematic to obtain the desired level of end-to-end VPN service. For this reason, Cisco recommends seeking an SLA with a single service provider that can guarantee a level of end-to-end service for the enterprise locations.

Also, some Internet service providers for DSL and cable services implement policing of traffic for residential class service. This means that protocols such as IPsec might be blocked unless the customer subscribes to business class service.

Design Selection

This section gives a high-level overview of several different IPsec VPN design topologies that can be currently deployed. Advantages and disadvantages are presented for each. The subsequent section then maps customer requirements to select the most appropriate design recommendation.

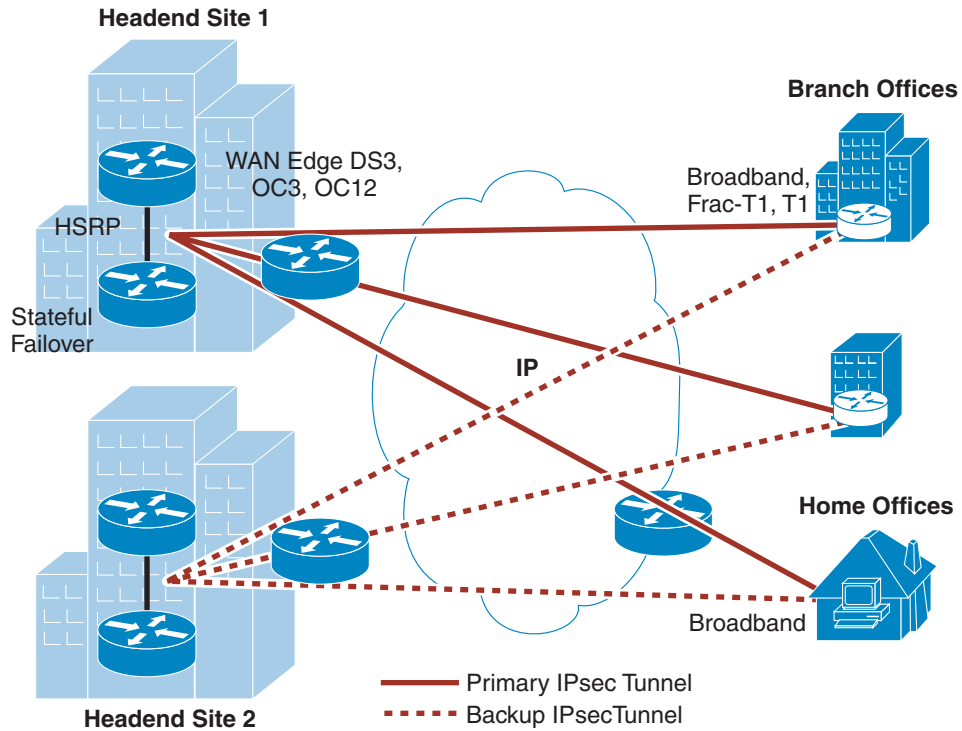
IPsec Direct Encapsulation Design

IPsec itself provides a tunnel mode of operation that enables it to be used as a standalone connection method. This option is the most fundamental IPsec VPN design model. IPsec Direct Encapsulation designs cannot transport IGP dynamic routing protocols or IPmc traffic.

Design Overview

Figure 11 illustrates the IPsec Direct Encapsulation design.

Figure 11 IPsec Direct Encapsulation Design



Headend			Branch		
Routing Control Plane	Route Redistr.	RRI	Static Route		
IPsec Control Plane	Dynamic Crypto Map	DPD	Static Crypto Map	Peer List	DPD

148182

Each remote site initiates an IPsec tunnel to a pre-defined headend. Remotes can have static or dynamic IP addresses, while headends must have static IP addresses.

Resiliency can be provided by IPsec Stateful Failover at the headend locations. Branch routers can be configured with a list of headends. If a connection cannot be established with the first headend, subsequent headends are tried until successful connection is made.

Dynamic crypto maps can be configured on the headend routers, such that new incoming tunnel connections can be established without having to manually provision each new tunnel on the headend router.

Typically, Dead Peer Detection (DPD) is enabled to detect loss of a peer connection, and Reverse Route Injection (RRI) is configured on the headends to inject routes for the branch router subnets into the routing table. No routing protocol is exchanged between the headend and branch routers.

Advantages

IPsec Direct Encapsulation designs offer the following advantages:

- Configurations are fairly straightforward
- All Cisco IOS router platforms support this design, including the 870, ISR 1800/2800/3800, 7200VXR, and 7600. The PIX, ASA, and VPN3K platforms also provide support.
- Dynamic crypto maps can be configured on the headend routers such that new incoming tunnel connections can be established without having to manually provision each new tunnel on the headend router.
- There is interoperability with non-Cisco peer devices that are RFC compliant.
- IGP routing peers are not a scalability limitation because no dynamic IGP routing protocol runs over the tunnel.

Disadvantages

IPsec Direct Encapsulation designs offer the following disadvantages:

- No support for IP multicast or non-IP protocols (multiprotocols).
- No support for dynamic IGP routing protocols over the VPN tunnel.
- If the primary tunnel is lost, no secondary tunnel is pre-established, so the new tunnel must be established to the alternate headend before traffic can continue.
- Distribution of IPsec tunnels to headend routers can be non-deterministic, because loss of a connection results in remote routers initiating a tunnel to subsequent headend peers in the peer list. For example, remotes do not automatically switch back to their primary headend after a failure recovery.
- It is not possible to implement a QoS service policy per VPN tunnel.
- When QoS service policies are configured with IPsec designs, interaction between IPsec and QoS can cause IPsec anti-replay packet drops.

Most Common Uses

IPsec Direct Encapsulation designs are commonly used when there is no requirement for dynamic IGP routing or IP multicast, and branch offices have very few or a single subnet, such as teleworkers and small office/home office (SOHO) deployments. IPsec Direct Encapsulation designs are also commonly used in remote access applications, where a single device (such as a laptop) is initiating a tunnel via a client.

For more information on using IPsec Direct Encapsulation designs for enterprise WAN connectivity, see the *IPsec Direct Encapsulation Design Guide* at the following URL: <http://www.cisco.com/go/srmd>.

Point-to-Point GRE over IPsec Design

IPsec can be deployed in conjunction with p2p GRE (an IPsec encrypted point-to-point GRE tunnel) to provide additional functionality. With the addition of p2p GRE to IPsec, dynamic IGP routing protocols and IP multicast traffic can be transported over the VPN tunnel.

Headend Architecture—Single Tier Headend versus Dual Tier Headend

When implementing a p2p GRE over IPsec design, the following two headend architectures can be implemented at the central site:

- Single Tier
- Dual Tier

Single Tier Headend Architecture

In a Single Tier Headend Architecture, both the p2p GRE and crypto functionally co-exist on the same router CPU. Headend routers service multiple p2p GRE over IPsec tunnels for a prescribed number of branch office locations. In addition to terminating the VPN tunnels at the central site, headend routers can advertise branch routes using IP routing protocols (EIGRP, OSPF, and so on).

Dual Tier Headend Architecture

In a Dual Tier Headend Architecture, the p2p GRE and crypto functionally do not co-exist on the same router CPU. There are p2p GRE headend routers, as well as crypto headend routers, that together service multiple p2p GRE over IPsec tunnels for a prescribed number of branch office locations. In addition to terminating the VPN tunnels at the central site, the p2p GRE headend routers can advertise branch routes using IP routing protocols (EIGRP, OSPF, and so on).

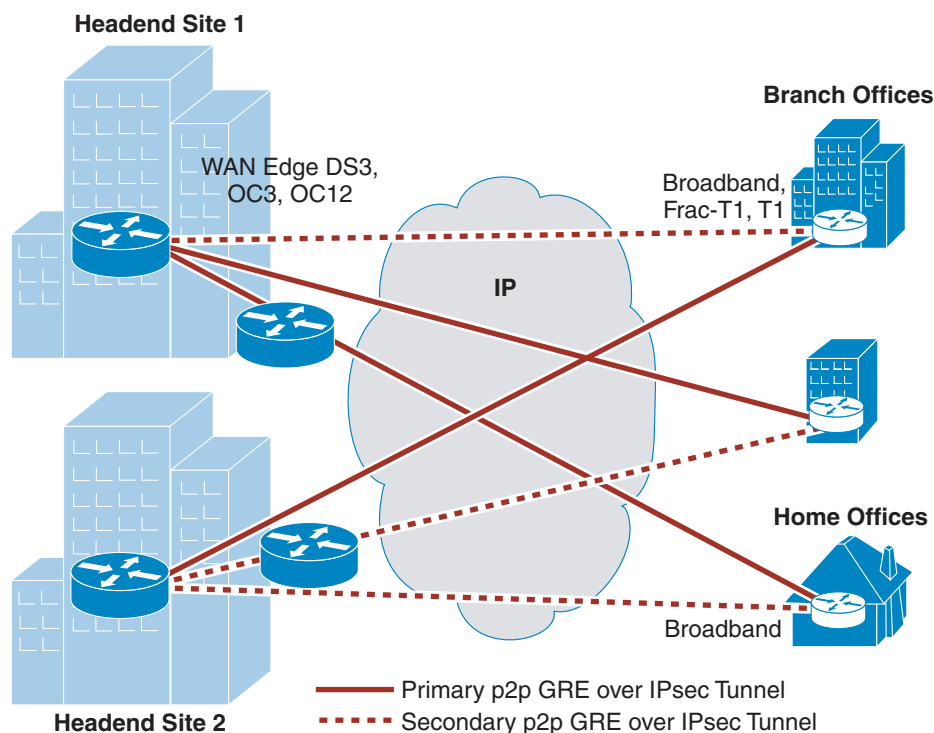
Performance and Value

When considering performance and value, the two should be considered together. Performance is based on the number of packets a platform can forward in a given time frame, or packets per second (pps). Value is the price for a specific platform based on the pps rate. For more information when choosing a Single Tier Headend Architecture versus a Dual Tier Headend Architecture, see the *Point-to-Point GRE over IPsec Design Guide* at the following URL: <http://www.cisco.com/go/srnd>.

Design Overview

Figure 12 illustrates the p2p GRE over IPsec design.

Figure 12 p2p GRE over IPsec Design



	Headend	Branch
Routing Control Plane	Dynamic Routing	Dynamic Routing
GRE Control Plane	Point-to-Point GRE	Point-to-Point GRE
IPsec Control Plane	Dynamic or Static Crypto Map DPD	Static Crypto Map DPD or Tunnel Protection

148878

Each remote site is connected with a p2p GRE over IPsec tunnel to a pre-defined headend. Remotes can have static or dynamic IP addresses, while headends must have static IP addresses.

Resiliency can be provided by configuring p2p GRE over IPsec tunnels to multiple headend routers at one or more geographic hub locations. An IGP dynamic routing protocol is exchanged over the p2p GRE over IPsec tunnels, and primary/secondary tunnels are differentiated by configuring slightly different routing metrics.

DPD can be enabled to detect loss of a peer connection.

Advantages

p2p GRE over IPsec designs offer the following advantages:

- IP multicast and non-IP protocols are supported.
- Dynamic IGP routing protocols over the VPN tunnel are supported.
- Supported on all Cisco IOS router platforms.

- QoS service policies can be configured per p2p GRE over IPsec tunnel (scalability might be an issue).
- Distribution of IPsec tunnels to headend routers is deterministic, with routing metrics and convergence choosing the best path.
- All primary and secondary/backup p2p GRE over IPsec tunnels are pre-established, such that a new tunnel does not have to be established in the event of a failure scenario.

Disadvantages

p2p GRE over IPsec designs have the following disadvantages:

- Configuration of each p2p GRE tunnel interface is static and can lead to lengthy headend configurations.
- Provisioning of new branch offices typically requires a configuration change/addition to the headend router(s).
- IGP routing peers tend to limit scalability more than IPsec Direct Encapsulation designs.
- Per-tunnel QoS service policies are limited in scalability because of generic traffic shaping currently being process switched.
- When QoS service policies are configured with IPsec designs, interaction between IPsec and QoS can cause IPsec anti-replay packet drops.

Most Common Uses

p2p GRE over IPsec designs are commonly used when there are requirements for routing and/or IPmc. They are also used in circumstances where branch offices have multiple subnets that make it desirable to exchange IGP dynamic routing protocols.

For more information on using p2p GRE over IPsec designs for enterprise WAN connectivity, see the *Point-to-Point GRE over IPsec Design Guide* at the following URL: <http://www.cisco.com/go/srnd>.

Dynamic Multipoint VPN—Hub-and-Spoke Topology Design

DMVPNs combine IPsec, mGRE, and Next Hop Resolution Protocol (NHRP). DMVPN has the following two modes of operation:

- DMVPN hub-and-spoke topology design—Functions very similarly to a p2p GRE over IPsec design in that all tunnels are established between headend and remote routers only.
- DMVPN spoke-to-spoke topology design—In addition to the DMVPN hub-and-spoke topology functionality, it is also possible for spokes to establish dynamic tunnels to other spokes for direct communications.

This section covers the DMVPN hub-and-spoke topology design mode of operation. [Dynamic Multipoint VPN—Spoke-to-Spoke Topology Design, page 34](#) covers the DMVPN spoke-to-spoke topology design mode of operation.

Headend Architecture—Single Tier Headend versus Dual Tier Headend

When implementing a DMVPN hub-and-spoke topology design, the following two headend architectures can be implemented at the central site:

- Single Tier
- Dual Tier

Single Tier Headend Architecture

In a Single Tier Headend Architecture, both the mGRE and crypto functionally co-exist on the same router CPU. Headend routers service multiple DMVPN tunnels for a prescribed number of branch office locations. In addition to terminating the VPN tunnels at the central site, headend routers can advertise branch routes using IP routing protocols (EIGRP, OSPF, and so on).

Dual Tier Headend Architecture

In a Dual Tier Headend Architecture, the mGRE and crypto functionally do not co-exist on the same router CPU. There are mGRE headend routers, as well as crypto headend routers, that together service multiple DMVPN tunnels for a prescribed number of branch office locations. In addition to terminating the VPN tunnels at the central site, the mGRE headend routers can advertise branch routes using IP routing protocols (EIGRP, OSPF, and so on).


Performance and Value

When considering performance and value, the two should be considered together. Performance is based on the pps a platform can forward in a given time frame. Value is the price for a specific platform based on the pps rate. For more information when choosing a Single Tier Headend Architecture versus a Dual Tier Headend Architecture, see the *Dynamic Multipoint VPN (DMVPN) Design Guide* at the following URL: <http://www.cisco.com/go/srnd>.

Design Overview

Figure 13 illustrates the DMVPN hub-and-spoke topology design.

Figure 13 DMVPN Hub-and-Spoke Topology Design



Headend			Branch	
Routing Control Plane	Dynamic Routing	NHRP	Dynamic Routing	NHRP
GRE Control Plane	Multipoint GRE		Point-to-Point GRE	
IPsec Control Plane	Dynamic Crypto Map or Tunnel Protection	DPD	Static Crypto Map or Tunnel Protection	DPD

148911

Each remote site is connected with a p2p GRE tunnel interface to a pre-defined headend. The headend router(s) use mGRE interfaces to dynamically accept new tunnel connections.

Resiliency can be provided by configuring DMVPN tunnels mapped to mGRE interfaces on multiple headend routers at one or more geographic hub locations.

Remotes can have static or dynamic IP addresses, while headends must have static IP addresses. An IGP dynamic routing protocol is exchanged over the DMVPN tunnels, and primary/secondary tunnels are differentiated by configuring slightly different routing metrics.

IPsec Tunnel Protection is generally used to map the crypto attributes to the tunnel that is originated by the remote router. DPD can be enabled to detect loss of a peer connection.

NHRP is configured on both the headend and branch office routers, and is a requirement for using mGRE interfaces.

Advantages

DMVPN hub-and-spoke topology designs offer the following advantages:

- IP multicast is supported.
- Dynamic IGP routing protocols over the VPN tunnel are supported.
- Supported on all Cisco IOS router platforms (some limitations on high-end router platforms).
- Distribution of IPsec tunnels to headend routers is deterministic, with routing metrics and convergence choosing the best path.
- All primary and secondary/backup DMVPN tunnels are pre-established, such that a new tunnel does not have to be established in the event of a failure scenario.
- Configuration of both IPsec and mGRE is dynamic, which simplifies and shortens configurations on the headend only. Provisioning of new branch offices can be done without a configuration change/addition to the headend router(s).

Disadvantages

DMVPN hub-and-spoke topology designs have the following disadvantages:

- No support for non-IP protocols.
- IGP routing peers tend to limit the design scalability.
- No interoperability with non-Cisco IOS routers.
- There is some added complexity with DMVPN in having to configure NHRP, as well as added complexity in troubleshooting.
- Not possible to implement a QoS service policy per VPN tunnel.
- When QoS service policies are configured with IPsec designs, interaction between IPsec and QoS can cause IPsec anti-replay packet drops.
- There is no direct acceleration support for the extra 4-byte mGRE tunnel key on the high-end router platforms, such as the Cisco 7600 (or Cisco Catalyst 6500) with a VPNSM or VPN SPA blade. However, there is a work-around that enables some acceleration capability (see [Scaling a Design, page 41](#) for more information).

Most Common Uses

DMVPN hub-and-spoke topology designs are commonly used when there are requirements for IGP routing and/or IP multicast. They are also used in circumstances where branch offices have multiple subnets that make it desirable to exchange IGP dynamic routing protocols.

In addition, because of the easier configuration and touchless provisioning of new branch offices, a DMVPN hub-and-spoke topology design is preferred for headend configurations over a p2p GRE tunnel approach. Even in cases where static p2p GRE configurations are used on the branch routers, it is advantageous to use mGRE on the headends.

Finally, DMVPN hub-and-spoke topology designs are deployed as a first step towards a DMVPN spoke-to-spoke topology design (see next sections).

For more information on using DMVPN hub-and-spoke topology designs for enterprise WAN connectivity, see the *Dynamic Multipoint VPN (DMVPN) WAN Design Guide* at the following URL: <http://www.cisco.com/go/srnd>.

Dynamic Multipoint VPN—Spoke-to-Spoke Topology Design

DMVPNs can also be configured in a spoke-to-spoke topology design, in which individual branch office routers can dynamically initiate DMVPN tunnel connections between each other, bypassing the headend router.


A DMVPN hub-and-spoke topology design always exists as a prerequisite to a DMVPN spoke-to-spoke topology design, meaning that branch offices always have an DMVPN tunnel connection to their assigned headend router(s).

With a DMVPN spoke-to-spoke topology design, each branch office router is additionally configured with one or more mGRE interfaces.

Design Overview

Figure 14 illustrates the DMVPN spoke-to-spoke topology design.

Figure 14 DMVPN Spoke-to-Spoke Topology Design



	Headend		Branch	
Routing Control Plane	Dynamic Routing	NHRP	Dynamic Routing	NHRP
GRE Control Plane	Multipoint GRE		Multipoint GRE	
IPsec Control Plane	Tunnel Protection	DPD	Tunnel Protection	DPD

148912

Each remote site is connected with a DMVPN tunnel to a pre-defined headend. The headend router(s) use mGRE interfaces to dynamically accept new tunnel connections.

Resiliency can be provided by configuring DMVPN tunnels mapped to mGRE interfaces on multiple headend routers at one or more geographic hub locations.

Remotes can have static or dynamic IP addresses, while headends must have static IP addresses.

An IGP dynamic routing protocol is exchanged over the DMVPN hub-and-spoke tunnels *only*, and primary/secondary tunnels are differentiated by configuring slightly different routing metrics. No routing information is exchanged between spokes.

IPsec Tunnel Protection is generally used to map the crypto attributes to the tunnels. DPD can be enabled to detect loss of a peer connection.

NHRP is configured on both the headend and branch office routers, and is a requirement for using mGRE interfaces.

With the addition of mGRE on the remote routers, spoke-to-spoke tunnels can be established between remote peers. Traffic between spokes always starts via the spoke-hub-spoke path. NHRP facilitates a new tunnel being established directly between the two spoke routers. After the new path is established, traffic flows from spoke-to-spoke.

Advantages

DMVPN spoke-to-spoke topology designs offer the following advantages:

- IP multicast is supported on the hub-and-spoke tunnels only (not between spokes directly).
- Dynamic IGP routing protocols over the hub-and-spoke VPN tunnel are supported (but not exchanged between spokes).
- Supported on all Cisco IOS router platforms (some limitations on high-end router platforms).
- Distribution of IPsec tunnels to headend routers is deterministic, with routing metrics and convergence choosing the best path.
- All primary and secondary/backup DMVPN tunnels are pre-established, such that a new tunnel does not have to be established in the event of a failure scenario.
- Configuration of both IPsec and mGRE is dynamic, which simplifies and shortens configurations. Provisioning of new branch offices can be done without a configuration change/addition to the headend router(s).

Disadvantages

DMVPN spoke-to-spoke topology designs have the following disadvantages:

- No support for non-IP protocols.
- There is no QoS between spoke routers for spoke-to-spoke tunnels, making it possible for a destination spoke router to become overwhelmed with traffic. As a result, latency/jitter/drop-sensitive applications such as VoIP and Video over IP are “best effort” in spoke-to-spoke topologies.
- IGP routing peers tend to limit the design scalability.
- No interoperability with non-Cisco IOS routers.
- There is some added complexity with DMVPN in having to configure NHRP, as well as added complexity in troubleshooting.

- Not possible to implement a QoS service policy per VPN tunnel.
- When QoS service policies are configured with IPsec designs, interaction between IPsec anti-replay and QoS can cause packet drops.
- There is no direct acceleration support for the extra 4-byte mGRE tunnel key on the high-end router platforms, such as the Cisco 7600 (or Cisco Catalyst 6500) with a VPNSM or VPN SPA blade. However, there is a work-around that enables some acceleration capability (see [Scaling a Design, page 41](#) for more information).

Most Common Uses

DMVPN spoke-to-spoke topology designs are commonly used when there are requirements for dynamic meshing when customers anticipate having significant traffic requirements between branch offices.

However, because of the current limitations in QoS, expectations need to be set appropriately if latency/jitter/drop-sensitive applications will be transported over DMVPN spoke-to-spoke topology designs. Such designs are currently “best effort”.

For more information on using DMVPN spoke-to-spoke topology designs for enterprise WAN connectivity, see the *Dynamic Multipoint VPN (DMVPN) Design Guide* at the following URL: <http://www.cisco.com/go/srnd>.

Virtual Tunnel Interface Design


Virtual Tunnel Interface (VTI) design is the newest of the IPsec VPN design options available in Cisco IOS. VTI designs have a number of distinct advantages over other IPsec design options, including the ability to transport IGP dynamic routing protocols and IPmc traffic without the addition of p2p GRE or mGRE headers.

In addition, VTI tunnels are assigned an interface, such that tunnel level features can be enabled on each tunnel, for example a QoS service policy. This makes it possible to have per VPN tunnel/destination QoS.

Design Overview

[Figure 15](#) illustrates the VTI design.

Figure 15 VTI Design



Headend			Branch	
Routing Control Plane	Dynamic Routing		Dynamic Routing	
IPsec Control Plane	Dynamic Virtual Tunnel Interface	DPD	Static Virtual Tunnel Interface	DPD

148913

The topology is very similar to a p2p GRE over IPsec design, except when used in conjunction with the Virtual Template Service. VTI also provides some of the same advantages of a DMVPN hub-and-spoke topology design.

Each remote site is connected with a VTI tunnel to a pre-defined headend. The headend router(s) use dynamic VTI interfaces to dynamically accept new tunnel connections, much like an mGRE interface, except each new tunnel clones the defined template and can also clone other attributes of a tunnel, such as a QoS service policy.

Resiliency can be provided by configuring VTI tunnels mapped to VTI interfaces on multiple headend routers at one or more geographic hub locations.

Remotes can have static or dynamic IP addresses, while headends must have static IP addresses. A dynamic IGP routing protocol is exchanged over the VTI tunnels, and primary/secondary tunnels are differentiated by configuring slightly different routing metrics.

IPsec Tunnel Protection is generally used to map the crypto attributes to the tunnel that is originated by the remote router. DPD can be enabled to detect loss of a peer connection.

Advantages

VTI designs offer the following advantages:

- IP multicast is supported.
- Dynamic IGP routing protocols over the VPN tunnel are supported.
- Supported on most Cisco IOS router platforms starting with IOS 12.3(14)T or 12.3(4).
- Distribution of IPsec tunnels to headend routers is deterministic, with routing metrics and convergence choosing the best path.
- All primary and secondary/backup VTI tunnels are pre-established, such that a new tunnel does not have to be established in the event of a failure scenario.
- Configuration of VTI is dynamic, which simplifies and shortens headend configurations. Provisioning of new branch offices can be done without a configuration change/addition to the headend router(s).
- Possible to implement a QoS service policy per VTI tunnel (scalability could be a concern).
- When QoS service policies are applied to VTIs, IPsec anti-replay drops induced by the encrypting router are eliminated because of QoS being performed before encryption. However, note that the ISP can reorder packets and cause anti-replay drops.
- Backwards compatible with an IPsec Direct Encapsulation design on other Cisco IOS routers because it offers a migration path from an IPsec Direct Encapsulation design headend to a VTI design dynamic headend.

Disadvantages

VTI designs have the following disadvantages:

- IP Unnumbered is required.
- No support for non-IP protocols.
- Might have limited interoperability with non-Cisco peer devices.
- IGP routing peers tend to limit the design scalability.
- Per-tunnel QoS service policies are limited in scalability because of generic traffic shaping currently being process switched.
- There is no support for VTI on the high-end router platforms, such as the Cisco 7600 (or Cisco Catalyst 6500) with a VPNSM or VPN SPA blade. VTI is scheduled for a future Cisco IOS release.

Most Common Uses

VTI designs are relatively new, but can be used for the same customer requirements where a p2p GRE over IPsec design or DMVPN hub-and-spoke topology design would be recommended; for example, where there are requirements to transport dynamic IGP routing protocols and/or IP multicast traffic over the VPN.

In cases where a DMVPN hub-and-spoke topology design is being considered because of the easier configuration and touchless provisioning of new branch offices, VTI designs offer further simplification of the configuration and design by not having to configure mGRE or NHRP.

In addition, VTI designs offer the capability to configure QoS service policies at the tunnel level, which makes it the only VPN design topology that can easily support QoS per VPN tunnel/destination, equivalent to traditional WAN connections with generic traffic shaping in both the hub-and-spoke and spoke-to-hub directions.

For more information on using VTI designs for enterprise WAN connectivity, see the *Virtual Tunnel Interface (VTI) Design Guide* at the following URL: <http://www.cisco.com/go/srnd>.

Design Comparison

This section offers a comparison of the various design options described in this guide.

Major Feature Support

Table 3 summarizes some of the major features and requirements and indicates whether each of the IPsec VPN design options supports the requirement.

Table 3 Major Feature Support

	Dynamic Routing	Tunnel Keep-Alive	HA	IPmc	Interface QoS	Per-Tunnel QoS	Branch Dynamic IP Address	Dynamic Headend Config	Dynamic Meshing
IPsec Direct Encapsulation	Partial (DPD/RRI)	N/A	Stateful failover	No	Yes	No	Yes	Yes	No
p2p GRE Over IPsec	Yes	Yes	RP	Yes	Yes	Yes	Yes	p2p GRE, No IPsec, Yes	No
DMVPN Hub-and-Spoke Topology	Yes	No	RP	Yes	Yes	No	Yes	Yes	No
DMVPN Spoke-to-Spoke Topology	Yes	No	RP	Yes	Yes	No	Yes	Yes	Yes
VTI	Yes	No	RP	Yes	Yes	Yes	Yes	Yes	Yes

Platform Support

Table 4 summarizes the current Cisco IOS router platforms and indicates whether each of the IPsec VPN design options is supported on the platform today.

Table 4 Platform Support

	Cisco 870	Cisco ISR 1800	Cisco ISR 2800	Cisco ISR 3800	Cisco 7200VXR	Cisco 7301	Cisco 7600	Cisco Catalyst 6500
IPsec Direct Encapsulation	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Point-to-Point GRE Over IPsec	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
DMVPN Hub-and-Spoke Topology	Yes	Yes	Yes	Yes	Yes	Yes	Work-around	Work-around
DMVPN Spoke-to-Spoke Topology	Yes	Yes	Yes	Yes	Yes	Yes	No	No
VTI	Yes	Yes	Yes	Yes	Yes	Yes	No	No

Selecting a Design

One of the strengths of Cisco that should be presented to customers is the richness of design options that can be provided.

The following are key questions that should help a customer choose an appropriate design option:

- IP unicast only or IP unicast and IP multicast combined?

If the customer traffic is IP unicast only, and the branch offices are relatively small, an IPsec Direct Encapsulation design might meet the requirements.

More often, the customer has requirements for IP multicast or dynamic IGP routing protocols, in which case an encapsulation method such as p2p GRE, mGRE, or VTI is required.

If the customer has requirements for IP multicast, dynamic IGP routing protocols, and non-IP protocol support, the only encapsulation method supported is p2p GRE.

- Large or small number of branch offices?

If the customer deployment is for a relatively small number of branch offices (100 or so), the p2p GRE over IPsec design might meet the requirements.

Often, the customer has requirements for several hundred, one thousand, or several thousand branch offices, in which case a dynamic touchless headend configuration is advantageous, such as the DMVPN hub-and-spoke topology or VTI design.

- What level of high availability is required?

If the customer requirement is for 1–2 second stateful failover, the IPsec Direct Encapsulation design might be the only design option.

If the customer has expectations of 20–60 second failover based on IGP routing convergence, the designs supporting IGP routing protocols should be considered, such as p2p GRE over IPsec, DMVPN hub-and-spoke topology, or VTI.

- Dynamic meshing required?

If the customer has significant requirements for branch offices to establish direct connections to other branch offices because of significant branch-to-branch traffic requirements, the DMVPN spoke-to-spoke topology design is the most likely option.

Scalability is another factor that can influence design selection, which is covered in the next section.

Scaling a Design

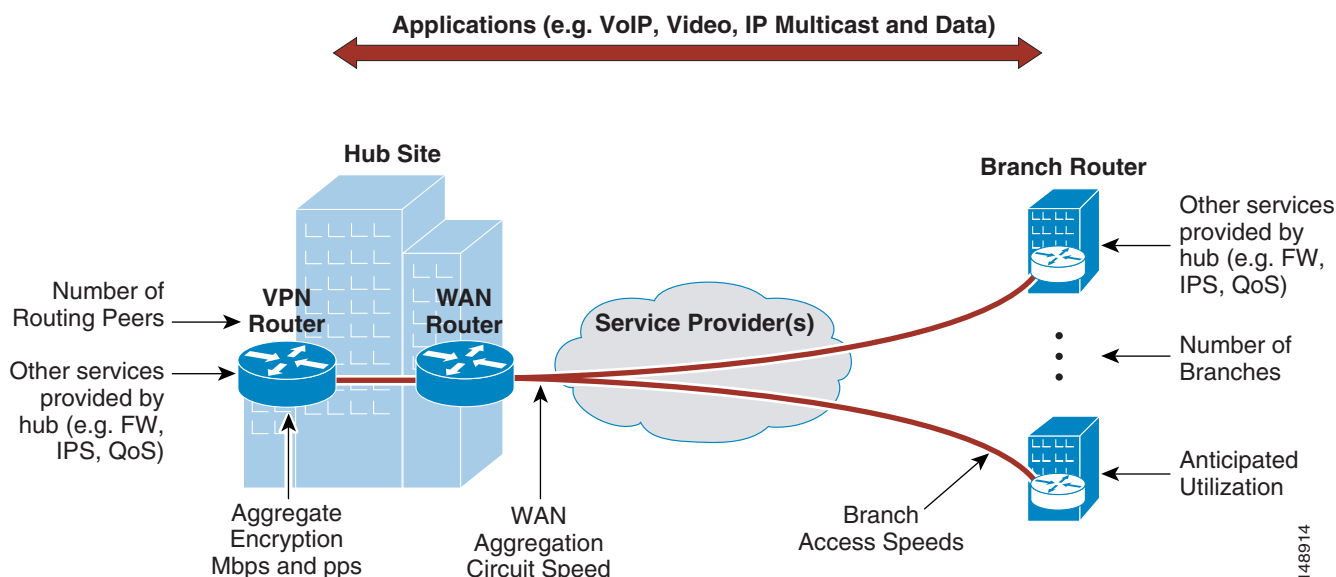
This section describes the critical factors that affect the scalability of an IPsec VPN design. Ideally, design selection and design scalability should be somewhat independent decisions, but in reality both must be considered together for a practical design proposal.

Also see each of the design sections for IPsec Direct Encapsulation, p2p GRE over IPsec, DMVPN, and VTI for specific scalability and performance information.

Critical Scalability Criteria

Scalability of an IPsec VPN design depends on a number of factors, and is not trivial. [Figure 16](#) illustrates some of these factors.

Figure 16 Critical Scalability Criteria



The main factors to consider include the number of branch offices, connection speeds at the hub(s) and branches, IGP routing peers, high availability requirements, and applications that will be transported over the IPsec VPN. These critical factors are covered in the following sections.

Number of Branch Offices

The number of branch offices that will be aggregated over the IPsec VPN is a primary factor in determining scalability of a design.

In addition to being a primary factor of the design, the number of branch offices also affects the routing plan, high availability design, and ultimately the overall throughput that must be aggregated by the VPN headend router(s).

Consider the number of branch offices that exist today, and also future expansion needs.

Connection Speeds

Because IPsec VPN connections do not (normally) have a bandwidth associated with them, the overall physical interface connection speeds of both the headend and branch routers largely determine the maximum speeds at which the IPsec VPN must operate. Some typical connection speeds are shown in [Table 5](#) and [Table 6](#).

Table 5 *Typical Headend Connection Speeds*

Connection Type	Speed
T3/DS3	44.76 Mbps
Fast Ethernet	100 Mbps
OC3	155 Mbps
OC12	622 Mbps

Table 6 *Typical Branch Connection Speeds*

Connection Type	Speed
T1	1.544 Mbps
2 x T1	3.088 Mbps
T3/DS3	44.76 Mbps
Broadband cable/DSL	384 Kbps uplink / 2 Mbps downlink

IPsec Throughput

IPsec throughput depends on several factors, including connection speeds, capacity of the crypto engine, and CPU limits of the router.

Bi-directional Traffic Handling

The first important factor to understand about IPsec throughput is that an IPsec crypto engine in a Cisco IOS router, whether software-only or hardware-accelerated, is a uni-directional device that must process bi-directional packets. Outbound packets must be encrypted by the IPsec crypto engine, while inbound packets must be decrypted by the same device. Thus, for each interface having packets encrypted, it is necessary to consider the bi-directional speed of the interface. For example, a T1 connection speed is 1.544 Mbps, but the IPsec throughput required is 3.088 Mbps.

Packets per Second (pps) is More Accurate

The second most important factor to understand is that pps rate matters more than throughput bandwidth (bps) for the connection speeds being terminated or aggregated. In general, routers and crypto engines have upper boundaries for processing a given number of pps. Size of packets used for testing and

throughput evaluations can understate or overstate true performance. For example, if a router/VPN module combination can handle 20 Kpps, then 100-byte packets lead to 16 Mbps throughput, while 1400-byte packets at the same packet rate lead to 224 Mbps.

Because of such a wide variance in throughput, pps is generally a better parameter to consider for scalability than bps.

Number of Tunnels May be a Factor

Each time a crypto engine encrypts or decrypts a packet, it performs mathematical computations on the IP packet payload using the unique crypto key for the trustpoint, agreed upon by the sender and receiver. If more than one IPsec tunnel is terminated on a router, the router has multiple trust points and therefore multiple crypto keys. When packets are to be sent or received to a different tunnel than the last packet sent or received, the crypto engine must swap keys to use the right key matched with the trustpoint. This key swapping can degrade the performance of a crypto engine, depending on its architecture, and increase the router CPU utilization.

For some Cisco platforms, such as the 7200VXR with SA-VAM2+, as the number of tunnels increases, throughput of the IPsec crypto engine decreases. For other Cisco platforms, such as the 7600 with VPN SPA, performance is relatively linear, with relatively the same throughput for a single tunnel as for 1000 or even 5000.

Throughput Summary

There are many factors (as described here) that affect throughput, but it is helpful to at least understand the general IPsec throughput that can be achieved with various platforms. [Table 7](#) provides a list of throughput per platform.

Table 7 IPsec Throughput Capabilities by Platform

Platform	Performance and Scalability Testing	Comparable Connection Speed
Cisco 1800 ISR On-board VPN Module	4–6 Mbps 2–2.5 Kpps	2 x T1
Cisco 2800 ISR On-board VPN Module	6–23 Mbps 2.5–11 Kpps	4 x T1
Cisco 3800 ISR On-board VPN Module	36–48 Mbps 18–24 Kpps	Fractional DS3
Cisco 7200VXR NPE-G1 Dual SA-VAM2+	80–100 Mbps 40–50 Kpps	DS3
Cisco 7600 Sup720 VPN SPA	1.1 Gbps–1.2 Gbps 480–600 Kpps	OC12

The performance and scalability testing column shows what has been measured under a relatively IP unicast aggressive traffic mix (described later) and conservative performance parameters, and can be considered relatively conservative design recommendations. The comparable connection speed column shows the approximate connection speed that can be driven by the platform.

Routing Peers

For routed IPsec VPN designs (such as p2p GRE over IPsec, DMVPN, and VTI), the number of IGP routing peers that must be maintained by the headend aggregation router(s) is a primary determining factor in the scalability of the design.

Specific performance depends on the aggregation router platform and its CPU capacity, which IPsec design is implemented, the number of peer branch routers exchanging routing, and whether there are secondary (and backup) IPsec tunnels to alternate headends for high availability. Table 8 illustrates very general ranges of IGP routing peers that have been verified in the Cisco test lab for various designs.

Table 8 IGP Routing Peers by Platform

Platform	Performance and Scalability Testing
Cisco 7200VXR NPE-G1 Dual SA-VAM2+	500–700
Cisco 7600 Sup720 VPN SPA	1000



Note

Whether a platform and design can really handle the number of peers suggested here depends on all the other factors that are covered in this section.

Quality of Service

The effect of QoS on scalability largely depends on the scope in which QoS is deployed in. If deployed at an interface level, scalability and performance testing has shown approximately 10–15 percent impact on the Cisco IOS router when QoS and in particular generic traffic shaping (GTS) is engaged. This is primarily because GTS is process switched, not CEF switched.

If there is a requirement for per-VPN tunnel QoS, this means a number of traffic shapers are engaging simultaneously on the Cisco IOS router during periods of tunnel congestion. Because of process switching, if a significant number of traffic shapers are engaged on a headend platform, scalability is significantly limited.

For more information on scalability limitations with per-VPN tunnel QoS, see the *Virtual Tunnel Interface (VTI) Design Guide* at the following URL: <http://www.cisco.com/go/srnd>.

High Availability

The nature of the high availability design requirements affects the scalability of the deployment. If multiple headends are deployed at a single or different geographic sites, with each branch router having an IPsec tunnel to primary or secondary (and possibly backup) headends, the number of total IPsec tunnels that needs to be aggregated is doubled (or tripled, or even quadrupled).

For example, if 300 branches are to be aggregated to a pair of headend routers, each with a primary and secondary tunnel, there are a total of 600 tunnels. If a second backup data center is included with a second pair of headend routers, there are now potentially 1200 total tunnels from those 300 branches.

High availability convergence times after a failure is also a critical factor in design scalability. For example, if a routed IPsec VPN design has a requirement for convergence within 20 seconds, the frequency of routing hellos must be more aggressive than a design with a convergence requirement of 40 seconds. The more aggressive the convergence time, the higher the burden on the headend aggregation router CPU for processing the number of hellos sent or received to all the peer (branch office) routers.


IP Multicast

If IP multicast is a design requirement, this can definitely limit the scalability of an IPsec VPN design. As stated earlier, p2p GRE over IPsec, DMVPN, or VTI designs must be implemented to transport IP multicast traffic.

IP multicast requires the crypto headend router to replicate each IP multicast packet for each VPN tunnel that is joined to the IP multicast stream. The IP multicast fan-out speed creates challenges to supporting significant numbers of IP multicast listeners joined to a single IP multicast stream without overwhelming the input queue of the crypto engine in a given platform.

For example, scalability and performance testing done with the Cisco VPNSM (in both Cisco Catalyst 6500 and Cisco 7600) has shown that it is susceptible to packet drops whenever the number of IP multicast listeners joined to a single IP multicast stream is more than 200. Figure 17 illustrates this issue.

Figure 17 IP Multicast Buffer Ring Issues



The newer Cisco VPN SPA has shown improvement with the ability to handle up to 1000 simultaneous listeners joined to a single multicast stream.

For more information on supporting IP multicast applications over an IPsec VPN, see the *IP Multicast over IPsec VPN Design Guide* at the following URL: <http://www.cisco.com/go/srnd>.

Internet Access Strategy

How Internet-bound traffic from branch offices is routed can also affect the IPsec VPN design and scalability. This consideration is primarily applicable when the Internet is being used as the transport connectivity to branch offices. Two common alternatives are described in the next two sections, along with the implications to the design scalability.

Backhaul

Enterprise customers commonly choose to *backhaul* all traffic to the headend site, regardless of whether the traffic is destined for the enterprise corporate network or for the Internet. It might seem inefficient to encrypt and transport Internet traffic over the Internet via an IPsec tunnel, only to be decrypted and then routed out the corporate Internet gateway; possibly through the same gateway those packets had just arrived.

However, there are some legitimate reasons for doing so, including the following:

- Keeping the branch routing simplistic for security reasons, instead of a more complex configuration to send some traffic over the corporate tunnel and other traffic straight to the Internet.
- Internet traffic monitoring, such as URL filtering, Websense, e-mail scanning, and anti-virus scanning, can be centralized at the corporate headquarters location(s), instead of trying to manage distributed monitoring functions.

If all traffic, including Internet-destined, is backhauled over the IPsec VPN connection, this traffic must be factored into the overall traffic bandwidth and IPsec throughput requirements for the design.

Split Tunneling

Split tunneling is the process by which packets being transmitted from a branch office can either be protected by IPsec and sent to or received from the headend aggregation router, or unprotected by IPsec and sent to or received from the Internet.

If split tunneling is implemented, the following extra factors must be considered in the design:

- The bandwidth and IPsec throughput saved by routing Internet traffic directly instead of over the VPN to headquarters.
- Most likely NAT or PAT will need to be running as a service on the branch office router, which can have performance and scalability implications.
- Security services, such as firewall and IPS, might need to be running as services on the branch office router, which will have performance and scalability implications. Whether or not these mandatory security services are running on the branch router or as standalone devices must be considered.
- Is the Internet traffic from the employee or from a spouse and child or guest to be split from the IPsec-protected traffic? In which subnet(s) or VLAN(s) are those hosts located in the branch?

Integrated Services

Integrated services running on the same routers as the IPsec VPN must be considered in the overall scalability, because depending on the platform, additional services can impact router performance. Examples of integrated services are covered in the next sections.

Security

Integrated security services can include firewall, IPS, DoS prevention, Network Admission Control (NAC), and others. Security services tend to be computationally intensive while performing packet inspections and analysis.

Integrated security services might not have a direct impact on router performance, depending on the function and what level of acceleration is provided in hardware on the specific platform. If the security services are performed in the main router CPU, enabling the service will most likely affect performance.

Integrated security services are most commonly deployed in branch office routers, but increasingly are being deployed in headend aggregation routers as well, especially in scenarios where an untrusted transport service, such as the Internet, is being used for WAN or VPN transport.

VoIP

Integrated voice services include voice station cards that terminate handsets, VoIP conferencing and transcoding services such as Digital Signal Processors (DSPs), a Public-Switched Telephone Network (PSTN) gateway, Survivable Remote Site Telephony (SRST), and others.

Integrated voice services are common in branch office routers, especially with the Integrated Services Router (ISR) series, which is specifically designed with integrated services in mind. Integrated voice services are generally not deployed in headend aggregation routers.

Other Integrated Service Types

Other types of integrated services that can impact router performance include the following:

- NAT or PAT
- DHCP server
- Content caching

See the individual IPsec design guides for IPsec Direct Encapsulation, p2p GRE over IPsec, DMVPN, and VTI for more information on the specific impacts of these services.

Appendix A—Evaluating Design Scalability

This appendix describes how Cisco conducts performance and scalability testing.

Test Methodology

The Cisco test methodology is to build networks using best practices and then to apply traffic loads that approximate as closely as possible real customer networks. There are several proof-of-concept test beds as well as a large-scale test bed containing nearly 1500 routers, capable (using VRFs) of emulating up to 5000 branch office routers with the ability to drive up to 1.5 Gbps of customer traffic.

The objectives of testing are to find real world performance and scalability guidelines for Cisco platforms functioning in an end-to-end system. The objective is not to exhaustively test all features and functionality, and also not to find the absolute maximum performance points of platforms. Instead the objective is to provide conservative benchmarks for use in sizing customer design proposals.

Traffic Mix

Cisco gets a lot of questions from customers regarding testing with certain packet sizes or with the popular IMIX testing profile. Cisco does not test with any particular packet sizes and does not use IMIX. As stated earlier, the testing methodology is to apply real customer traffic profiles, and as such server endpoints (such as Sun Netras or Penguins) are used to create end-to-end protocol flows over the network under test. This allows for real world behaviors of protocols, such as TCP/IP flow control, to be characterized properly.

Traffic flows are established primarily using the NetIQ Chariot and Ixia testing tools. The mix of traffic is primarily made up of the following types of traffic:

- DNS
- FTP
- POP3
- HTTP
- TN3270
- g.729 VoIP

The mix of traffic is approximately 35 percent UDP and 65 percent TCP in bps, which correlates to ~80 percent of the packets being UDP, with the remaining 20 percent being TCP. Of the UDP packets, the majority are g.729 VoIP RTP streams to simulate a converged enterprise network. VoIP and QoS deployment guidelines stipulate that bandwidth consumed by VoIP should be no more than 33 percent of total bps bandwidth.

The most noticeable impact of this traffic profile, compared to a profile such as IMIX, is that results are conservative for a data-only network. IMIX does not take VoIP into account. Inclusion of VoIP causes an increase in the number of small packets in the traffic mix, driving the overall pps rate up, which in turn drives the router CPU higher.

Finding Limits

Finding the performance limits is a difficult process, because it is especially difficult to determine a particular bps or pps rate, given that the testing methodology involves real end-to-end protocol flows. Essentially what occurs is an iterative process of increasing the traffic flows while keeping the defined ratios of protocols, until the desired limits are found.

The limits themselves depend on the nature of the testing being conducted, but, in general, the parameters shown in [Table 9](#) are monitored.

Table 9 *Finding Limits*

Parameter	Acceptable Value
CPU	Performance reported at 50%, 65%, and 80% points
Latency	<50ms one way
Jitter	<10ms one way
Packet drops	<0.5%
Process switching	Monitored
Anti-replay drops	<1%
Failover	Must converge and continue operation in a reasonable time frame

The targeted CPU utilization for a router deployment is always the subject of debate. Cisco attempts to provide a representative range of CPU utilizations on the higher side of a normal network deployment. These utilizations are 50 percent, 65 percent, and 80 percent. Although the high number (80 percent) is not recommended during normal operation, this number is provided to enable a network engineer to see what traffic levels can be handled by a platform in a failover scenario.

Conservative Results

Cisco performance and scalability results tend to be conservative. They are not meant to contradict published data sheet performance numbers for products. Instead they are intended to provide conservative design performance guidelines that can be used for deployments with a level of confidence.

Any testing has its own unique set of limitations. Keep in mind that Cisco test results are specific results for a specific set of conditions and do not reflect the conditions that can occur in each customer network.

Cisco Platforms Evaluated

Table 10 illustrates various platforms that were evaluated in the Cisco IPsec VPN test lab.

Table 10 *Current Cisco VPN Router Platforms Evaluated*

Application	Cisco VPN Router	Processing Engine	VPN Acceleration Options
Headend aggregation	Catalyst 6500	Sup2	VPNSM
	7600	Sup 720	VPN SPA
	7200VXR	NPE-G1	SA-VAM2+
	7301	NPE-G1 (Equivalent)	SA-VAM2+
Large branch office	3845 ISR	On-Board	On-Board, AIM-VPN/HP/II+ (optional)
	3825 ISR	On-Board	On-Board, AIM-VPN/EP/II+ (optional)
Medium branch office	2851 ISR	On-Board	On-Board, AIM-VPN/EP/II+ (optional)
	2821 ISR	On-Board	On-Board, AIM-VPN/EP/II+ (optional)
	2811 ISR	On-Board	On-Board, AIM-VPN/EP/II+ (optional)
Small office	2801 ISR	On-Board	On-Board, AIM-VPN/EP/II+ (optional)
	1841 ISR	On-Board	On-Board, AIM-VPN/BP/II+ (optional)
	1811W	On-Board	On-Board
	871W	On-Board	On-Board

Table 11 illustrates various platforms previously evaluated in the Cisco test lab.

Table 11 *Legacy Cisco VPN Router Platforms Evaluated*

Application	Cisco VPN Router	Processing Engine	VPN Acceleration Options
Headend aggregation	7200VXR	NPE-400	SA-VAM
	7200VXR	NPE-300	SA-VAM
Large branch office	3745	On-Board	AIM-VPN/HPII
	3725	On-Board	AIM-VPN/EPII
Medium branch office	2691	On-Board	AIM-VPN/EPII
	2651XM	On-Board	AIM-VPN/BPII
	1760	On-Board	MOD1700-VPN
Small office	1711	On-Board	On-Board
	831	On-Board	On-Board

Appendix B—References and Recommended Reading

This section provides the following references and additional information related to the subjects covered in this design guide:

- Documents (available at <http://www.cisco.com/go/srnd>):
 - IPsec VPN WAN Design Overview
 - IPsec Direct Encapsulation Design Guide
 - p2p GRE over IPsec Design Guide
 - Dynamic Multipoint VPN (DMVPN) Design Guide
 - Virtual Tunnel Interface (VTI) Design Guide
 - Voice and Video Enabled IPsec VPN (V3PN) Design Guide
 - Multicast over IPsec VPN Design Guide
 - IPsec Redundancy and Load Sharing Design Guide
 - Digital Certification/PKI for IPsec Design Guide
- Request For Comment (RFC) papers
 - Security Architecture for the Internet Protocol—RFC 2401
 - IP Authentication Header—RFC 2402
 - The Use of HMAC-MD5-96 within ESP and AH—RFC 2403
 - The Use of HMAC-SHA-1-96 within ESP and AH—RFC 2404
 - The ESP DES-CBC Cipher Algorithm With Explicit IV—RFC 2405
 - IP Encapsulating Security Payload (ESP)—RFC 2406
 - The Internet IP Security Domain of Interpretation for ISAKMP—RFC 2407
 - Internet and Key Management Protocol (ISAKMP)—RFC 2408
 - The Internet Key Exchange (IKE)—RFC 2409
 - The NULL Encryption Algorithm and Its Use With IPsec—RFC 2410
 - IP Security Document Roadmap—RFC 2411
 - The OAKLEY Key Determination Protocol—RFC 2412

Appendix C—Acronyms

Term	Definition
3DES	Triple Data Encryption Standard
AES	Advanced Encryption Standard
AIM	Advanced Integration Module
ATM	Asynchronous Transfer Mode
CA	Certificate Authority
CBWFQ	Class Based Weighted Fair Queuing

Term	Definition
CPE	Customer Premises Equipment
DES	Data Encryption Standard
DHCP	Dynamic Host Configuration Protocol
DMVPN	Dynamic Multipoint Virtual Private Network
DPD	Dead Peer Detection
DSL	Digital Subscriber Line
EIGRP	Enhanced Interior Gateway Routing Protocol
FR	Frame Relay
FRTS	Frame Relay Traffic Shaping
FTP	File Transfer Protocol
GRE	Generic Route Encapsulation
HSRP	Hot Standby Router Protocol
IKE	Internet Key Exchange
IOS	Internetwork Operating System
IPmc	IP Multicast
IPsec	IP Security
ISP	Internet Service Provider
LFI	Link Fragmentation and Interleaving
mGRE	Multipoint Generic Route Encapsulation
MPLS	Multi-Protocol Label Switching
NAT	Network Address Translation
NHRP	Next Hop Resolution Protocol
OSPF	Open Shortest Path First
p2p GRE	Point-to-Point GRE
PAT	Port Address Translation
PVC	Permanent Virtual Circuit
QoS	Quality of Service
RADIUS	Remote Authentication Dial In User System
RTP	Real-Time Protocol
SLA	Service Level Agreement
SPA	Shared Port Adapter
SRST	Survivable Remote Site Telephony
TCP	Transmission Control Protocol
ToS	Type of Service
UDP	User Datagram Protocol
VoIP	Voice over IP
V3PN	Voice and Video Enabled IPsec VPN

Term	Definition
VAM	VPN Acceleration Module
VPN	Virtual Private Network
VPNSM	VPN Service Module
VPN SPA	VPN Shared Port Adapter
VTI	Virtual Tunnel Interface
WAN	Wide Area Network