Group Encrypted Transport VPN Security Analysis

Group Encrypted Transport (GET) VPN is a new method for encrypting any-to-any WANs, available on Cisco routers. This document provides an overview of GET VPN, including the trust model and security considerations, and gives guidelines on when to deploy GET VPN.

Overview of GET VPN

There is an increasing need for enterprises to encrypt private WANs that are built using service provider networks (e.g. BGP over MPLS). This is often prompted by regulatory requirements such as Health Insurance Portability and Accountability Act (HIPAA), Sarbanes-Oxley Act, and Payment Card Industry (PCI) security standards.

Typically, IP Security (IPsec) VPN gateways "surround" the service provider links, encrypting and decrypting communications at the border of the WAN. This technology has evolved over the years and scales well for hub-and-spoke configurations, where potentially thousands of remote sites can be aggregated into a central site gateway. With the increasing use of voice over IP (VoIP) and video applications that are sensitive to latency and delays, enterprise networks are tending toward meshed configurations, where remote sites are directly connected to each other rather than through a central site.

The traditional approach of surrounding the WAN links with pairs of Internet Key Exchange (IKE)/IPsec connections between VPN gateways quickly breaks down when meshed networks are deployed, as figure 1 shows. For example, a meshed network with 100 sites would require around 5000 (n × [n – 1] / 2) VPN connections, and each site would need VPN gateways that can handle such numbers.

Figure 1. Challenges in Scaling Full-Mesh Networks

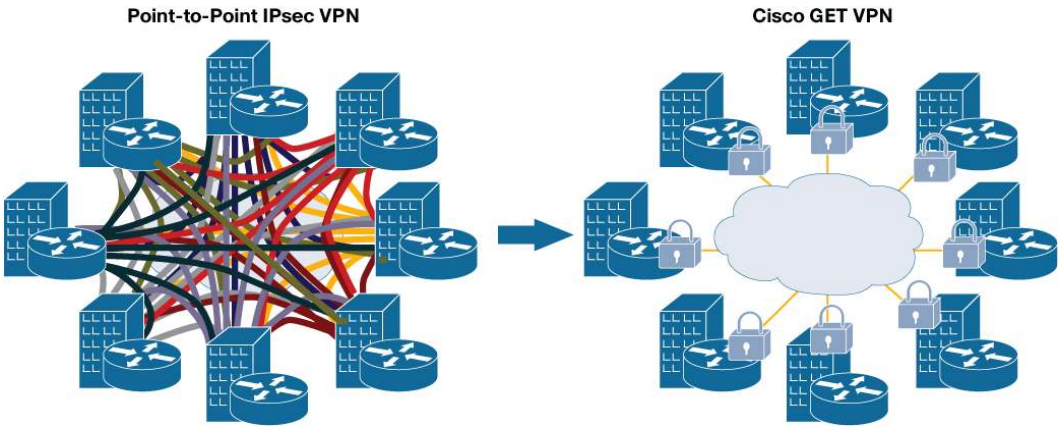

Cisco® Group Encrypted Transport (GET) VPN provides an innovative, scalable solution to protect enterprise traffic as it passes through the meshed private WAN. Using the underlying network intelligence, Cisco GET VPN enables encrypted IP Unicast and Multicast packets to be routed

© 2008 Cisco Systems, Inc. All rights reserved. This document is Cisco Public Information.

directly to remote sites based on routing protocol decisions and to be re-routed around failed paths, providing enhanced availability.

Cisco GET VPN avoids scaling issues by replacing pair-wise IKE connections with a dynamic group key management system. The dynamic group security in GET VPN has three steps, shown in Figure 2.

Figure 2. GET VPN Overview

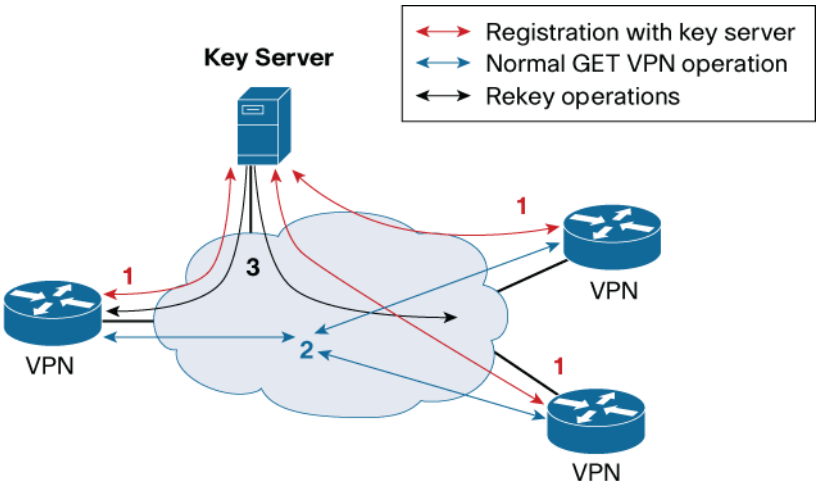

Step 1. VPN gateway joins a group by contacting a centralized key server using a registration protocol. After a mutual authentication process, the key server validates that the VPN gateways are authorized devices in the group VPN, and issues GET VPN IPsec policy and current keys to the VPN gateways.

Step 2. The VPN gateways use the group policy and keys issued by the key server to participate in the GET VPN group—they do not need to directly authenticate each other. The VPN gateways are thus able to encrypt and decrypt IP packets being circulated within the group.

Step 3. The key server distributes new group keys to the VPN gateways as needed, using a rekey protocol.

Using these steps, a GET VPN group can efficiently and securely protect traffic between an authorized mesh of VPN gateways.

Cisco GET VPN Trust Model

Understanding the trust model of a security solution is necessary before considering it for deployment. Traditional IPsec VPNs use a simple bilateral trust model. This means that each pair of VPN gateways directly authenticate each other, and set up IPsec sessions that are private to the pair. This is necessary when the information passed between the two VPN gateways must be kept secret between the two VPN gateways and also delivered only to that peer VPN gateway. For example, consider the extranet scenario shown in Figure 3, where Enterprise 1 has set up IPsec VPN connections to other partner or supplier enterprises.

Figure 3. Bilateral Trust Model (Traditional IPsec VPN)

© 2008 Cisco Systems, Inc. All rights reserved. This document is Cisco Public Information.

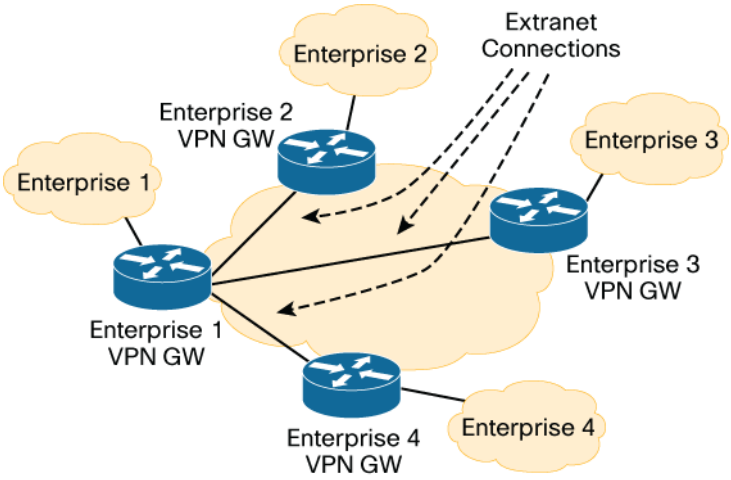

In this scenario, Enterprise 1 does not want the other enterprises to communicate using this VPN—they must only communicate to Enterprise 1; packets delivered within a single VPN link are intended for only one peer VPN gateway. When IPsec VPNs have these requirements, traditional IKE/IPsec VPNs should continue to be used to protect data. Cisco VPN solutions (IPsec/Generic Routing Encapsulation [GRE], Easy VPN, Dynamic Multipoint VPN [DMVPN]) adopt the bilateral trust security model to protect traffic between two peers.

In contrast, Cisco GET VPN uses a group trust model, and is suitable for deployment where a set of VPN gateways belong to the same domain (i.e., the same enterprise or coalition) and are authorized to share the same information between them.

For example, in the network shown in Figure 4, the VPN gateways act as corporate gateways to a private WAN such as an MPLS VPN. All the gateways (labeled “GW1” through “GW9”) are within the same enterprise, and are identically authorized to pass enterprise traffic between them. Together they protect enterprise traffic that passes through the private WAN, and traffic can be forwarded from any member to any other member in the WAN.

The bilateral trust model can add a substantial amount of unnecessary overhead to this system, in terms of VPN gateway management of pair-wise encryption state, VPN gateway processing, and IP packet latency waiting for tunnels to be set up.

Figure 4. GET VPN Group Trust Model

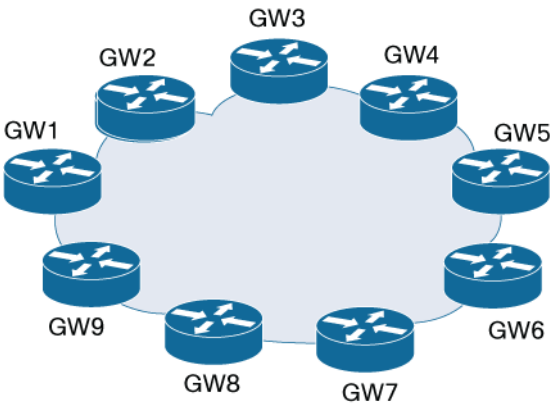

© 2008 Cisco Systems, Inc. All rights reserved. This document is Cisco Public Information.

When to Deploy GET VPN

The group trust model used by GET VPN is appropriate only when the VPN gateways are truly acting as a group. Usually, the gateways are acting as a group if the following characteristics are true:

- It is not necessary for each VPN gateway to directly authenticate each of its peer VPN gateways. It is acceptable to depend on a centralized entity (i.e., a key server), and for that key server to decide which VPN gateways are authorized VPN peers.
- All VPN gateways are trusted to decrypt any packets encrypted and forwarded by other GET VPN gateways. In addition, it is not necessary for a VPN gateway to determine exactly which packets were sent by which peer VPN gateway.
- Network Address Translation (NAT) is not used to obscure the IP addresses of VPN gateways or key servers.

These characteristics are usually found in MPLS networks, as well as other private WANs such as Frame Relay. If any of the above statements are not true for the network (e.g., direct links to the Internet), it is not suitable to deploy GET VPN.

GET VPN Protocols

GET VPN uses several protocols, many of which are Internet standards. In this section, each protocol is described, including its security features.

VPN Gateway Registration to a Key Server

GET VPN uses the Group Domain of Interpretation (GDOI) group key management protocol (RFC 3547) developed by the IETF. The GDOI registration protocol provides strong bilateral authentication through the use of preshared keys or digital certificates. Using the registration protocol, as shown in Step 1, each VPN gateway registers with the key server, requesting policy and keys to join a group. The GDOI key server ensures that the VPN gateway is authorized to join the group, before issuing group policy and keys for the group. Rekey policy and keys are distributed to the VPN gateways along with the IPsec policy and keys. All messages are protected with confidentiality, integrity, and replay protection.

Data Security Between VPN Gateways

GET VPN uses the Encapsulating Security Payload (ESP) protocol (RFC 2406) to provide confidentiality, integrity, and replay protection for packets flowing between VPN gateways. This is the same protocol used by traditional IPsec VPNs. The security policy determining how the data packets are protected (e.g., Security Parameters Index [SPI] value, cipher choice, key lifetime) is configured on the GDOI key server and transmitted during the VPN gateway registration process. The keys used for a particular IPsec Security Association are also downloaded during the registration process.

GET VPN uses ESP in tunnel mode,[1] which protects the entire data packet, as well as the IP header received by the VPN gateway. Tunnel mode processing adds a new IP header to the packet after ESP encapsulation. GET VPN uses a method of tunnel mode called “tunnel mode with address preservation” that copies the original source and destination from the inner IP header to the outer IP header (as shown in Figure 5).

[1] Note: IPsec also defines a transport mode for ESP that does not add a new IP header to the packet. Transport mode may be safely used in some IPsec applications, but fragmentation and reliability issues render it unsuitable for use with GET VPN.

© 2008 Cisco Systems, Inc. All rights reserved. This document is Cisco Public Information.

Figure 5. ESP Tunnel Mode with Address Preservation Encapsulation

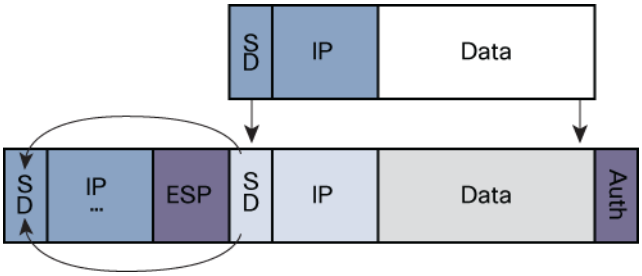

Tunnel mode with address preservation differs from normal tunnel mode, where the source and destination on the packet are taken as the source and destination VPN gateways. With normal tunnel mode, the packet can only be routed between the two VPN gateways through the private WAN service provider network and cannot be delivered via any alternate route to the destination, even if such an alternate route exists. This is a problem if routing to or through the destination VPN gateway fails because the IPsec protected traffic may be dropped. Address preservation allows packets to traverse through asymmetric paths in a private WAN network, which allows load-sharing between links and fast failover. Other parts of the inner IP header are copied to the outer IP header in both normal tunnel mode and tunnel mode with address preservation. Depending on configuration, these include the Type of Service field, Identification field, and the Don't Fragment This Datagram bit.

Figure 6. Routing Through a Private WAN

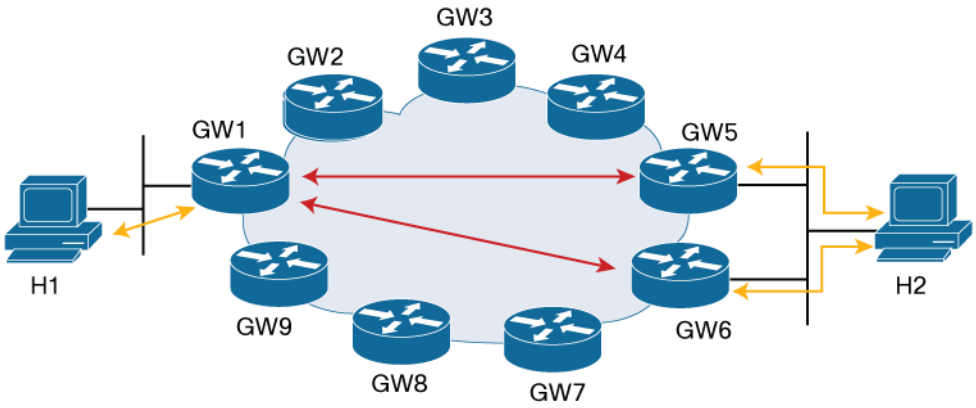

Figure 6 gives an example of routing through a service provider private WAN. When no tunneling is used, packets from H1 to H2 are sent across the private WAN with the source and destination (S, D) fields in the IP header set to H1 and H2. This allows the private WAN routing protocol to route the packets through either GW5 or GW6. However, if a traditional IPsec tunnel mode is added, the packets may be delivered across the private WAN between GW1 and GM5. In some cases (e.g., routing to GW5 becomes unstable), the packets will continue to be tunneled from GM1 to GM5 even though they could be delivered through GW6, resulting in "black holes" where failover does not kick in, and requiring considerable administrative effort to detect and fix. Tunneling mechanisms intended to react to this kind of packet "black-hole" situation (e.g., dead peer detection) are not effective in all black-hole scenarios.

© 2008 Cisco Systems, Inc. All rights reserved. This document is Cisco Public Information.

Enterprises using a private WAN service provider network typically expect the private WAN service provider to route their packets by original source and destination packets. That is, part of the service provided by the private WAN is to separate enterprise packets from other enterprises in the private WAN. This results in a network topology where there is no need to “hide” the enterprise addresses.

Tunnel mode with address preservation is the only method of tunnel mode that allows IP Multicast packets to be routed natively. GET VPN applies this mode to both Multicast and Unicast packets to optimize packet delivery.

Replay protection is an important function provided by ESP: it stops an attacker from capturing and resending data packets. The traditional ESP replay protection method uses sequence numbers; however, this does not scale for GET VPN. Instead, GET VPN includes a time-based replay protection method where any ESP packet not recently sent by a VPN gateway is not accepted. GET VPN administrators can configure the age of packets, which if exceeded, results in the packets being rejected.

Key Server Rekey

A GDOI key server needs to periodically distribute the new group state. It does so by sending a “rekey” message to the VPN gateways. The VPN gateways validate the rekey message against the policy received at registration time. The rekey protocol is labeled “3” in Figure 2. The rekey message is often sent as an IP Multicast message, which is the recommended method for GET VPN. In addition, GET VPN supports the distribution of rekey messages as Unicast packets, suitable for private WANs over service provider networks that do not support IP Multicast. Replacement rekey policy and keys are supplied during rekey.

The GDOI rekey message is authenticated using a digital signature. It is also encrypted for confidentiality, and includes replay protection. Accidental or malicious deletions of GDOI rekey message will not cause VPN gateways to be left without Security Associations. A GDOI key server can be configured to retransmit GDOI rekey messages to increase the likelihood that all VPN gateways will receive a particular GDOI message. However, when Security Associations are near the end of their life and no replacement Security Associations have been received, a VPN gateway will initiate the registration process with the key server to obtain replacement Security Associations. Thus, while rekey messages ease the distribution of replacement group policy, non-receipt of rekey messages will not affect connectivity of the group.

Key Server Synchronization

A single GDOI key server in a GET VPN can be a single point of failure in a network. GET VPN supports a set of cooperative key servers that jointly accept registrations from VPN gateways and distribute the GDOI rekey message. The GDOI key servers use a proprietary Cooperative Key Server Protocol to communicate between themselves and exchange current group policy and keys. If any key server becomes unreachable, the remaining cooperative key servers continue to distribute group policy and keys to the VPN gateways. This helps ensure that the group will continue to function normally as long as any one of the key servers is reachable.

The GDOI key synchronization protocol provides mutual authentication of the key servers. Group state is only exchanged with authorized cooperative key servers specified in the configuration. All packets are protected by confidentiality, integrity, and replay protection.

GET VPN Reliability

© 2008 Cisco Systems, Inc. All rights reserved. This document is Cisco Public Information.

GET VPN is designed for use cases where a large number of VPN gateways need to communicate in a mesh. Failure of GET VPN key management would affect the entire VPN. System reliability is therefore a mandatory component of the solution and is implemented in GET VPN components as follows:

- Support for multiple cooperative key servers that maintain and distribute the same group policy. This allows the system to operate even when a subset of key servers is nonoperational or unreachable.
- A VPN gateway has the ability to register with a set of key servers. This increases its chances of obtaining group state even when a subset of key servers is nonoperational or unreachable.
- Both the key server and VPN gateway take responsibility for ensuring that the VPN gateway has fresh keys before the current keys expire.
- In traditional IPsec, tunneled ESP packets directed to a single VPN gateway can be dropped if that VPN gateway becomes unreachable. GET VPN, however, uses the address preservation mode of ESP, which helps ensure that enterprise packets will not be dropped but rerouted within the service provider network.

Open and Closed Failure Modes

Whether or not a VPN gateway protects enterprise traffic in the presence of a VPN failure is called its “failure mode.” Typically, VPN gateways are configured to protect enterprise data at all costs, including dropping enterprise data that cannot be delivered as an encrypted packet. However, some enterprises’ risk assessments conclude that dropped packets are a larger risk to the enterprise livelihood than the risk of exposure in the service provider network private WAN. GET VPN supports both a fail-closed mode and a fail-open mode.

Fail-Closed Mode

The safest security policy is a consistent security policy where no unencrypted data leaves the enterprise. The fail-closed mode is recommended for GET VPN. GET VPN supports this on Cisco IOS® Software-based devices through the use of an interface access control list (ACL). Figure 7 shows a simple ACL that creates a fail-closed mode. This ACL allows only GDOI packets (on UDP port 848) and ESP packets to be sent or received through Ethernet0/0 facing the service provider network. (This access list is likely to also include exceptions for routing packets and other control plane traffic, but these have been omitted for clarity.)

Figure 7. Fail-Closed Model ACL Example

```
access-list 101 permit udp any eq 848 any eq 848
access-list 101 permit esp any any
access-list 101 deny   ip  any any

interface Ethernet0/0
 ip access-group 101 in
 ip access-group 101 out
```

© 2008 Cisco Systems, Inc. All rights reserved. This document is Cisco Public Information.

Fail-Open Mode

If a VPN gateway does not have an interface ACL on the service provider private WAN interface, it is using a fail-open policy. When a VPN gateway is first configured or brought online, it will not yet have registered with a GDOI key server and will thus not have any knowledge of what data GET VPN administrators intend to be protected. At this stage, the VPN gateway is in fail-open mode.

When the VPN gateway completes registration with the GDOI key server, it will have installed entries in its IPsec security policy database that exactly reflect the desire of the GET VPN administrators. The VPN gateway is now in fail-closed mode and will remain there unless the GET VPN data is completely deactivated (e.g., the crypto map containing the GET VPN policy is removed from the service provider interface).

Cryptographic Recommendations

Table 1 provides general guidance regarding cryptographic algorithms for the GDOI registration, IPsec ESP, GDOI rekey, and cooperative key server protocols.

Table 1. Cryptographic Algorithms

Protection	Recommended	Acceptable	Not Recommended
Peer Authentication	• Digital certificates	• Pre-Shared Keys (Peer-Specific Pair-Wise Keys)	• Pre-Shared Keys (Group Pre-Shared Keys)
Encryption Cipher	• AES-CBC (128-bit keys)	• 3DES-CBC • AES-CBC (192-bit keys) • AES-CBC (256-bit keys)	• DES-CBC
Integrity Algorithm	• HMAC-SHA	• HMAC-MD5	
IPsec Security Association Lifetime	• Less than the computed maximum lifetime (see the IPsec Security Association Lifetime Computation section), with a maximum of 24 hours	• 1 hour to 24 hours	• Less than 1 hour • More than the computed maximum lifetime (see the IPsec Security Association Lifetime Computation section)
Rekey Security Association Lifetime	• 24 hours	• 8 hours to 24 hours	• Less than 8 hours

Encryption Cipher

The Advanced Encryption Standard (AES) cipher is the recommended cipher for use with GET VPN. AES uses large key sizes, allowing AES in the CBC mode with 128-bit keys to provide the necessary security for small and large GET VPN deployments. Use of larger key sizes (e.g., 256-bit keys) is allowed if policy requires it, but will result in increasing the size of key management messages transporting those keys. In addition, use of hardware acceleration is recommended on platforms in which it is available.

Encryption Cipher with Address Preservation

GET VPN's use of ESP tunnel mode with address preservation might suggest that a small of amount of encrypted information is "leaked" (i.e., the original source and destination addresses are copied to the outside IP header). However, with the recommended ciphers (3DES-CBC and AES-CBC), there is no indirect leakage via known plaintext attacks that occur with GET VPN that cannot also occur when used with traditional VPNs. Modern ciphers such as Triple Data Encryption Standard (3DES) and AES resist known plaintext attacks. They also resist chosen-plaintext attacks, in which an adversary within the enterprise chooses some plaintext that gets encrypted by an unknown key. Even if such an attacker obtains a gigabyte of plaintext encrypted under a given

© 2008 Cisco Systems, Inc. All rights reserved. This document is Cisco Public Information.

key, the attacker gains only a negligible amount of information about any other plaintext encrypted by the same key. More specifically, there is very strong security as long as a 3DES-CBC key protects fewer than a theoretical maximum number of bytes (as discussed in the next section).

IPsec Security Association Lifetime Computation

There are theoretical limits to the amount of data encrypted with a single cipher key. This is called the volume threshold of the key. If an attacker captures the entire volume, he could theoretically analyze and begin to learn about the keying material used to encrypt the date. The volume threshold governs when the rekey needs to happen. The following volume thresholds are acceptable ciphers for use with GET VPN:

- 3DES keys should be used to encrypt no more than 2^{35} bytes.
- 128-bit AES keys should be used to encrypt no more than 2^{68} bytes.

While both of those volume thresholds appear to include a large amount of data, it is necessary to consider the aggregate usage of the key as it protects data within the GET VPN. That is, the volume threshold of the key depends on the total number of VPN gateways encrypting data with the key, and the rate of encryption of those VPN gateways. It is not possible for the GET VPN group to keep track of a volume threshold based on the actual amount of data encrypted because no VPN gateway or key server sees all of the data packets encrypted with a single key. Instead, the use of the key must be managed using a time-based value (e.g., lifetime). It is possible to construct a maximum lifetime from a volume threshold using the following procedure.

Step 1. Calculate the theoretical maximum possible throughput of the VPN gateways.

For example, consider a GET VPN where 500 VPN gateways are connected via a T1 interface (1.544 Mbps) and two VPN gateways are connected via an OC3 interface (155 Mbps). The total aggregate data for the GET VPN is computed as follows:

```
Aggregate Data = 500 * 1.544 Mbits/sec + 2 * 155 Mbits/sec
 = 772 Mbits/sec + 310 Mbits/sec
 = 1082 Mbits/sec
 = 135.25 Mbytes/sec
 = 0.13525 Gbytes/sec
```

Step 2. Determine the volume threshold for the cipher in use.

Assume the example GET VPN is using AES-CBC 128-bit keys, which have a volume threshold of 2^{68} bytes.

Step 3. Calculate the time-based lifetime by dividing the volume threshold by the aggregate data.

In this example, the lifetime is computed by dividing the AES-CBC 128-bit key volume threshold determined in Step 2 by the value computed in Step 1.1

```
Lifetime = Volume Threshold / Aggregate Data
= 2^68 bytes / 0.13525 Gbytes/sec
= 2^38 Gbytes / 0.13525 Gbytes/sec
= 2,032,368,997,738 secs
= 64,446 years
```

As shown by the example, an AES-CBC 128-bit key is often strong enough to be used by many VPN gateways simultaneously—the lifetime of the key is not short. Except in extremely large

© 2008 Cisco Systems, Inc. All rights reserved. This document is Cisco Public Information.

networks, the largest configurable lifetime of 24 hours will be acceptable for an IPsec Security Association using AES-CBC. However, note that the lifetime using the 3DES-CBC cipher in this same network is very short. Using the 3DES-CBC Volume Threshold specified above results in the following computation.

```
Lifetime = Volume Threshold / Aggregate Data
  = 2^35 bytes / 0.13525 Gbytes/sec
  = 2^5 Gbytes / 0.13525 Gbytes/sec
  = 237 secs
```

As can be seen, it is critical to choose the correct lifetime value when deploying an IPsec Security Association using 3DES-CBC. Because of this, it is recommended that AES-CBC be used rather than 3DES-CBC.

Conclusion

The GET VPN solution provides a scalable, reliable, and secure VPN solution when a group trust model is appropriate. A group trust model is generally appropriate when all VPN gateways are working together to provide confidentiality and integrity of enterprise data as it passes over a private WAN such as an MPLS VPN. The protocols used by GET VPN are standards-based, and the cryptographic properties of using a group key are well understood.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco StadiumVision, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn is a service mark; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0804R)

Printed in USA C11-471053-00 04/08

© 2008 Cisco Systems, Inc. All rights reserved. This document is Cisco Public Information.