White Paper

Secure Extension of Community of Interests Across Wide Area Networks

Authors

Mark “Mitch” Mitchiner
Solutions Architect
U.S. Federal Area
mmitchin@cisco.com

Craig Hill
Distinguished Systems Engineer
U.S. Federal Area
crhill@cisco.com

Abstract

This paper examines how recent network-based virtualization technology can be used to simplify community of interest (COI) deployment and operations within Department of Defense (DoD), Intelligence Community (IC), and secure enterprise networks.

The primary innovations addressed in this paper are Multiprotocol Label Switching (MPLS) over multipoint GRE (mGRE), combined with Group Encrypted Transport (GET) Virtual Private Network (VPN) technology while utilizing Next Generation Encryption ([NGE], also known as Suite B). These technologies, when combined as an architectural framework, address some of the major scaling, deployment, and operational challenges common in secure Wide Area Networks (WANs) today when Layer 3 network virtualization is required.

This paper compares the use of MPLS VPN over the WAN with network virtualization technologies typically deployed today. It also highlights the advantages of Cisco GET® VPN over multipoint IP tunnel-based overlay networks and how it simplifies operations and deployment. Finally, this paper describes and compares NGE to legacy cipher solutions, offering cryptographic algorithms designed to meet large-scale security requirements for decades to come.

Problem Statement

Introduction

Over the past decade, corporate networks have evolved to become mission-critical lifelines for businesses and governments. As the demand for IT increases, so does the pressure for IT organizations to “do more with less,” dictating how networks and data centers are being designed today. Cost reduction initiatives such as consolidation, maximizing the use of existing hardware, as well as “green” initiatives are driving how IT managers and architects design their networks.

Given this focus on consolidation, one area that is gaining increased interest is the area of virtualization. Although virtualization is typically thought of in the data center, specifically on servers with the use of virtual machines, network virtualization is also an important element to the overall consolidation of network devices. Network virtualization simplifies network operations by enabling customers to securely share a common network infrastructure between groups of users, applications, and devices. Using these network segmentation techniques over a common infrastructure provides the same look and feel of dedicated hardware to the end user, allowing virtualized domains, thus reducing the cost and simplifying management by reducing the number of network devices needed.

The purpose of this paper is to discuss next-generation network virtualization techniques as they apply to enterprise WAN design and how these virtualization solutions can be extended end to end in a secure manner.

Requirements to Support Community of Interest (COI) Separation

Overview of Network Virtualization Concepts

Like many current networks being implemented, the need for advanced network virtualization implementations is being seen as vital for maintaining secure separation of network resources over a common infrastructure. A common reference to a network virtualized domain is a COI or closed user group. A COI requires access to shared network resources or services, while maintaining separation from users outside of their group. The concept that a COI must maintain separation from other groups within the enterprise, campus, or data center is becoming common in today's network architecture, where the segmentation of data is needed throughout the infrastructure. However, extending COIs across WANs in a scalable and flexible manner provides a new set of challenges for network architects. (See Figure 1.)

Figure 1. Network Virtualization Overview

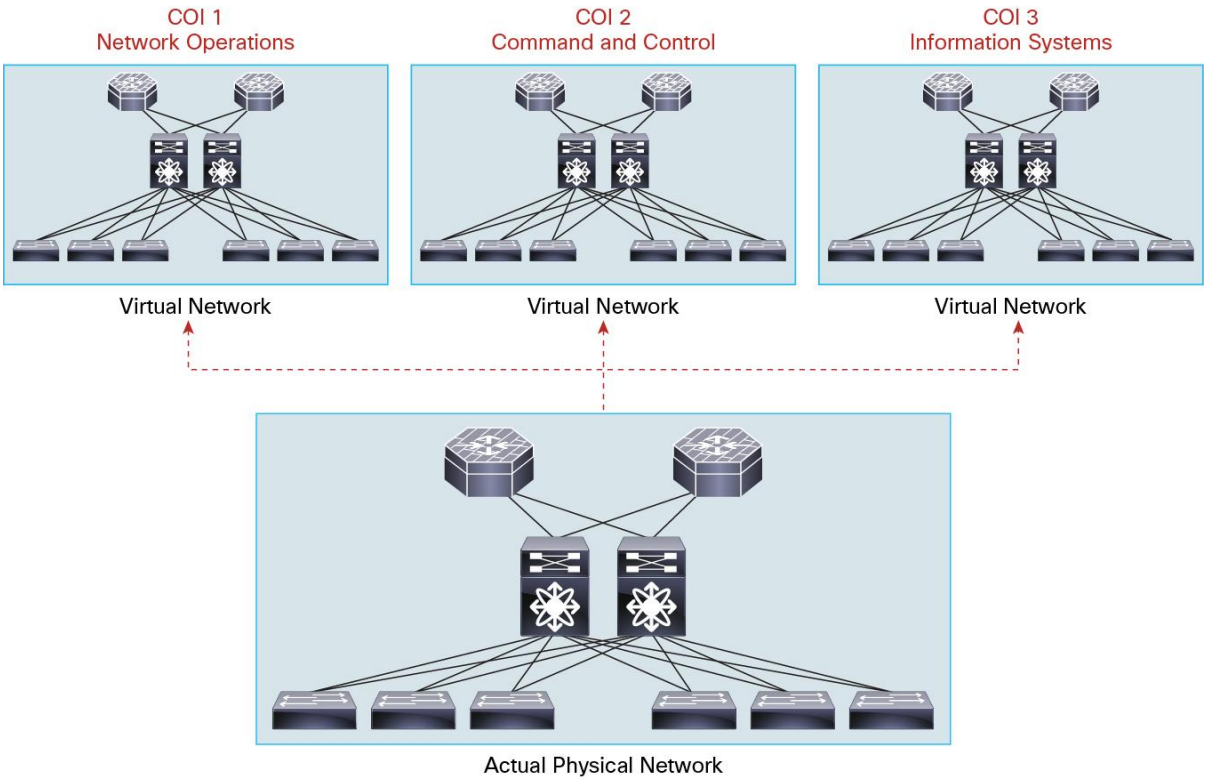

There are fundamental building blocks for designing this separation into the network framework. The first element is device separation and the concept of virtualizing the hardware for forwarding. The two common methods here are separation done for Layer 2 (L2) forwarding (that is, Ethernet and MAC layer), using the well-known concept of VLANs. A second alternative involves segmentation at the IP forwarding layer, where both the routing and forwarding tables are virtualized within a single device. This is most commonly referred to as a Virtual Routing and Forwarding instance, or VRF. For the purposes of this paper, a VRF correlates directly to a COI.

The second important building block is the practice of interconnecting these virtualized devices over the network, either campus, WAN, or data center. Virtualization extension is much more complicated and may include signaling protocols to exchange information and use multiplexing capabilities. A common mode for extending L2 or L3 virtualization is the use of Multiprotocol Label Switching (MPLS) for forwarding. MPLS in the forwarding plane allows the separation of L2 and/or L3 traffic over a common network infrastructure. Several examples of L2 and L3 MPLS-enabled services that can use MPLS forwarding in the backbone are Virtual Private LAN Service (VPLS), Ethernet Pseudowire, and BGP/MPLS IP VPN (RFC 4364) separation. In each solution, MPLS labels are enabled end to end in order to maintain the separation of L2 and L3 VPN COIs.

An alternative to MPLS forwarding in the backbone is use of native IP. That is, MPLS VPN services (and control plane) may be utilized for extending the device virtualization where L3 VPNs are required; however, IP encapsulation is used to forward the MPLS VPN traffic across the backbone between VPN provider edge (PE) devices. BGP VPNs over IP can take advantage of existing MPLS VPN control plane standards and virtualization techniques, while using the simplicity of any IP infrastructure as the transport.

Common MPLS Deployment Models

A common confusion point around the topic of MPLS is in the understanding of how MPLS fits into the overall architecture, specifically in the area of service provider (SP)-managed service offerings versus customers enabling MPLS in their privately owned backbone. For example, an SP will offer an “any-to-any” IP VPN managed service using MPLS as the underlying technology. However, MPLS is completely transparent to the customer.

As seen in Figure 2, the customer’s customer edge (CE) router peers with the SP’s router (PE) and, through a negotiated routing protocol (or static route), advertises the customer’s IP routes to the SP, and the SP’s MPLS network propagates these customers’ routes to all participating locations. The primary element here is that the CE does not participate in the SP’s MPLS backbone network or exchange any labels with the SP PE devices.

Figure 2. Service Provider-Managed IP VPN Service

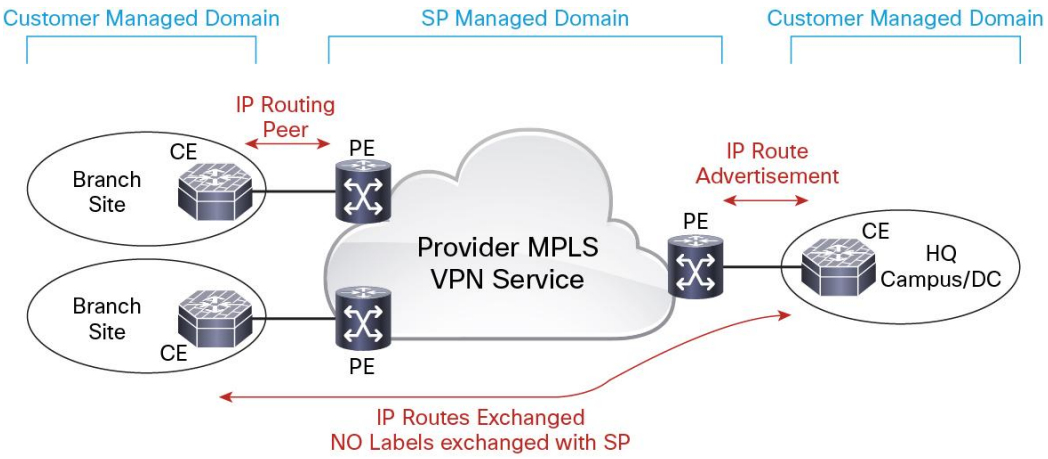

In the alternative model (shown in Figure 3), the customer deploys, owns, and manages the MPLS backbone, which is referred to as the “self-deployed” model. This model allows the customer to have total control of the services offered, service-level agreements (SLAs), as well as how rapidly a service can be stood up and rolled out to the end customer. The customers that use a “self-deployed” MPLS backbone normally need control and have the proper in-house expertise to design, manage, and maintain the infrastructure.

Figure 3. Self-Deployed MPLS

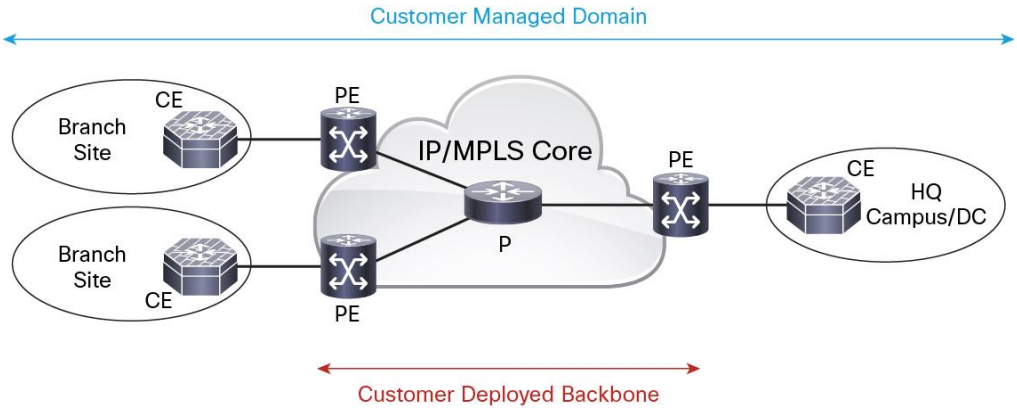

The “SP managed” versus “self-deployed” MPLS models are two common forms of transport solutions in the enterprise space.

A third model that is beginning to gain traction, specifically for customers that have the requirement to deploy network virtualization, is the combination of self-deployed MPLS VPN capabilities (Figure 3) while using the simplicity of an “SP managed” IP VPN transport (Figure 2.) This model introduces a new set of requirements on the CE router, as MPLS VPNs will need to be extended over the IP service transport, a model we refer to as “over the top.” This also introduces the concept of a customer PE (c-PE), as the CE router takes on full PE functionality (that is, RFC 4364) over IP. The following section will explore these “over-the-top” requirements, functions, and solutions when MPLS VPN requires the transport to be over IP.

Proposed Solution

MPLS VPN over IP

As discussed earlier, IP and IP VPN transport offerings from SPs are becoming much more prevalent and cost effective for enterprise and federal customers, which in turn require customers wanting a L3 virtualization solution to install their own form of MPLS VPN “over the top” of IP solutions. Although several IP encapsulation solutions exist, the most common form of IP encapsulation has been generic route encapsulation, or GRE. (See Figure 4.)

Figure 4. GRE Datagram Format

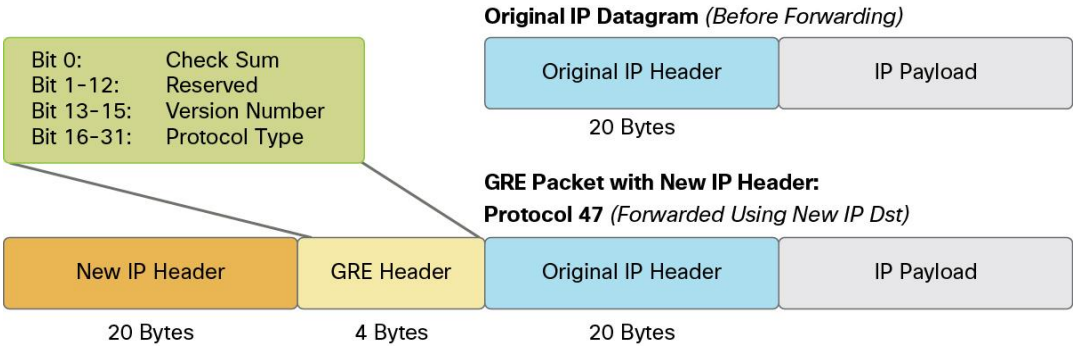

GRE Tunnel Overview

GRE has had many uses and capabilities over the years, specifically transporting non-IP protocols such as IPX, AppleTalk, DECnet, and even IP, most notably when using Internet Protocol Security (IPsec) and the customer requires IP multicast and/or a routing protocol to run over the GRE tunnel. GRE tunneling has two primary topologies as it relates to transporting IP, as shown in Figure 5.

Figure 5. Point-to-Point and Point-to-Multipoint GRE Tunnels

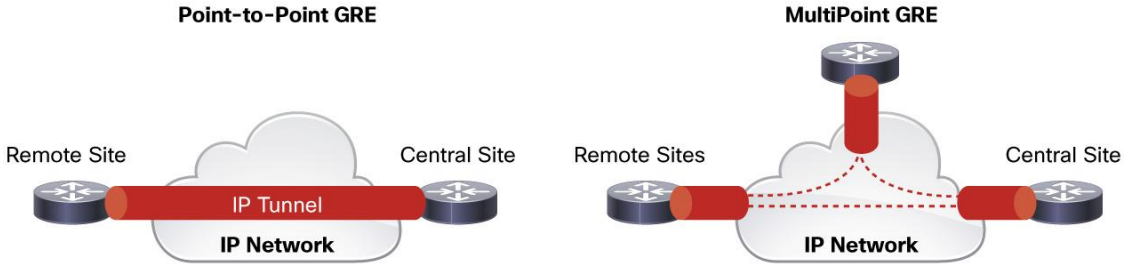

The figure on the left shows a typical “tunneling” deployment for a GRE overlay network, where two routers form the endpoints of a point-to-point tunnel. The point-to-point GRE model can be thought of as a “stateful” solution, where the endpoints of the tunnel are manually configured and specified by source and destination IP address.

To expand on the point-to-point GRE tunnel topology, Figure 6 highlights a common model, and as the number of sites in the network grows in a full-mesh environment, the number of point-to-point tunnels increases, normally in an “N - 1” fashion, where “N” is the number of locations, so in a network with 50 locations, each router will need to support 49 GRE tunnels.

Figure 6. Point-to-Point GRE Tunnels

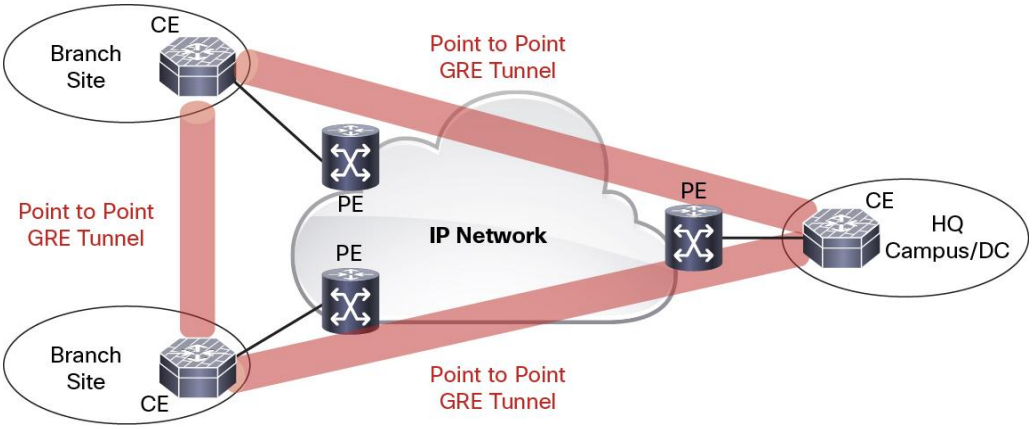

GRE also has the ability to function in a multipoint fashion (mGRE), which is shown on the right in Figure 5. In this model, a single tunnel interface on each endpoint can connect to multiple endpoints within the network domain, thus eliminating the N - 1 requirement, which is the primary advantage that mGRE offers.

Referring to Figure 7, mGRE only requires a single GRE interface/tunnel per router, so as the number of routers in the mesh increases, the number of GRE tunnel interfaces on the router remains constant (one). mGRE, because it does not require a stateful configuration, does require some form of dynamic mechanism to discover the destination endpoints (such as NHRP or BGP), but mGRE has much better scale and management simplicity than point-to-point GRE. This greatly reduces the amount of configuration required by the operator as locations are added to the network.

Figure 7. Point-to-Multipoint GRE Tunnels

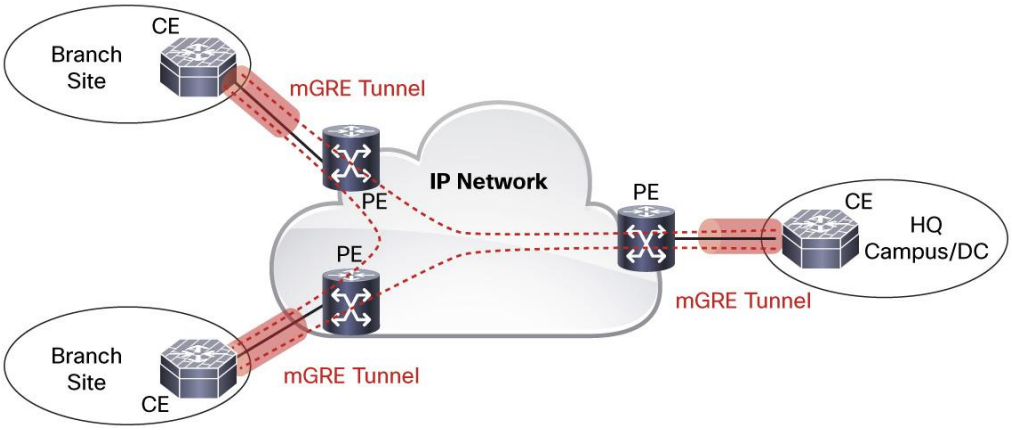

MPLS VPN over Point-to-Point GRE

As mentioned earlier, GRE is highly flexible, with the ability to transport a wide variety of network protocols in a multitude of topologies. In addition to supporting a variation of protocols, GRE also has the ability to carry MPLS labels, which opens up an entirely new set of capabilities for supporting network virtualization over an IP infrastructure (see RFC 4023). That is, MPLS over GRE can be utilized to offer L3 VPN services over any IP transport.

Typical L3 VPN over MPLS forwarding uses a label stacking concept, where the “inner” label identifies the VPN identifier (L2 or L3), while the “outer” label (typically referred to as the tunnel label) is used for forwarding the labeled packet to the destination PE through the various MPLS backbone routers (that is, P routers). It is this outer label where MPLS over GRE changes the forwarding paradigm.

As shown In Figure 8, when L3 VPN over GRE is used, the “inner” VPN label remains untouched, but the “outer” label is replaced with a GRE tunnel header plus the destination IP address of the remote PE router. This model, as outlined in RFC 4797, allows the standard RFC 4364 BGP VPN control plane to be used, while transporting the VPN data traffic over any IP transport offering. This versatility to use an IP transport is very appealing for enterprise and federal customers that desire MPLS VPN segmentation, but only have access to an IP transport.

Figure 8. GRE Tunnel Format with MPLS

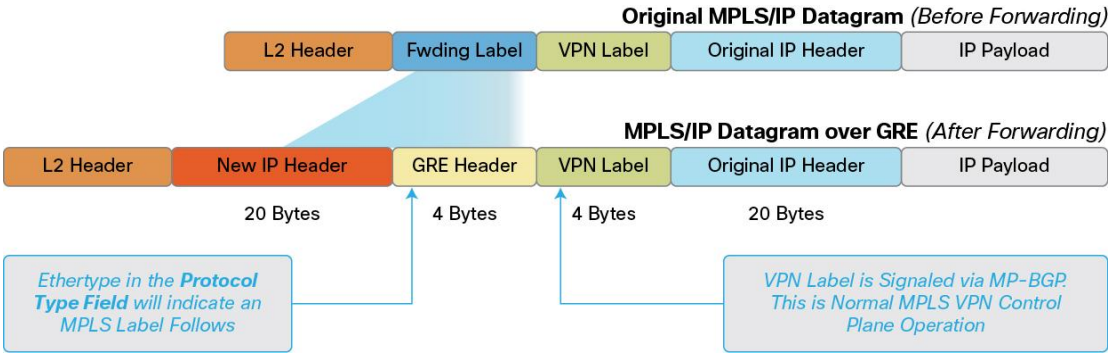

MPLS VPN over Multipoint GRE

Overview

As previously discussed, point-to-point GRE is flexible in that it offers the ability to encapsulate network-layer packets, including MPLS, inside IP tunneling packets. However, scale becomes an issue with the point-to-point tunnel model as additional sites are added to the network.

MPLS VPN over mGRE uses the flexibility that point-to-point MPLS VPN over IP solutions offer, while simplifying management, configuration, control plane protocols, and overall network operations of any MPLS VPN solution over IP. Because mGRE is a point-to-multipoint model, fully meshed GRE tunnels are not required to interconnect MPLS VPN PE devices. Thus, mGRE solves the cumbersome configuration issue that exists when attempting to configure hundreds of sites, requiring a full mesh of connectivity, with point-to-point GRE tunnels.

Although MPLS requires that all core routers support MPLS label forwarding, the MPLS VPN over mGRE feature overcomes this requirement by still requiring the MPLS VPN standard control plane, while eliminating the need that all forwarding routers support MPLS label forwarding. This allows the MPLS forwarding to use an IP encapsulation (GRE in this case) between endpoints across any form of existing IP transport networks, including L3 VPN providers, and public/private IP transports, including the Internet.

When the MPLS VPN over mGRE feature is enabled, it allows the operator to deploy a “self-deployed” MPLS VPN service on the c-PE router connected using the SP service offering. This allows the operator to provision L3 VPN services without using Label-Switched Path (LSP) or a Label Distribution Protocol (LDP). MPLS VPN over mGRE uses IPv4-based mGRE tunnels to encapsulate VPNv4/v6 labeled packets between c-PE devices, over the IP transport. It should also be noted that because the mGRE solution has no established destination in its configuration, some form of “signaling” is required to discover the interested endpoints. For the MPLS VPN over mGRE solution, BGP builds this IP tunnel endpoint database between each c-PE, thus allowing IP encapsulation between MP-BGP endpoints.

In addition, MPLS VPN over mGRE also allows both IP transport and legacy MPLS forwarding to coexist in the same PE. The ingress PE/c-PE router determines which encapsulation technology to use when a packet is sent to the remote PE/c-PE router, based on which PE signaled Network Layer Reachability Information (NLRI) using MP-BGP. This feature is extremely powerful, as it offers a smooth transition from MPLS VPNs over an MPLS core as well as utilizing IP encapsulation.

It is obvious the MPLS VPN over mGRE solution offers clear advantages for network designers wanting to use MPLS VPNs over an IP transport, specifically when a fully meshed traffic pattern is required between c-PE or PE routers. MPLS VPN over mGRE drastically eliminates the amount of configuration required on each device and simplifies the ability to support MPLS VPNs over IP.

L3 VPN over mGRE: Control Plane

L3 VPN over mGRE maximizes the use of this “multipoint” GRE concept, allowing BGP to discover the remote nodes and establish a “tunnelless” GRE connection through its normal BGP peering process. This allows the COIs to extend to multiple locations without the operator having to manually configure a mesh of GRE tunnels each time a COI is requested.

On the terminating mGRE node, through the use of a “tunnel profile,” a single multipoint tunnel interface is created dynamically by the router, and only a single source IP address is required for each tunnel endpoint. The tunnel destination is derived dynamically through the use of standards-based RFC 4364 MP-BGP (MP-iBGP) as the control plane mechanism. That is, MP-iBGP neighbors are established “over the top” of the SP VPN cloud, and iBGP is used to advertise VPNv4 routes, exchange VPN labels, and learn tunnel endpoints. VPNv4 labels and VPN payload are carried across the network through mGRE tunnel encapsulation. Figure 9 represents the MP-BGP control plane for this solution.

Figure 9. MPLS VPN over mGRE: Control Plane

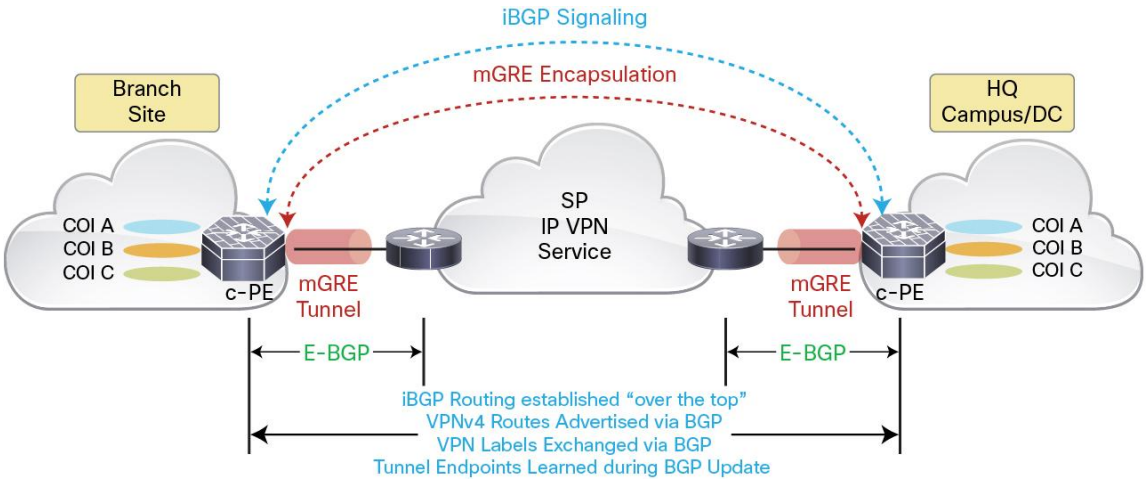

As shown in Figure 9, the use of MP-BGP control plane greatly enhances the scalability of the solution, simplifying the process of adding new VRFs, eliminating the need for a Label Distribution Protocol (LDP or RSVP-TE), and using IP encapsulation for forwarding.

L3 VPN over mGRE: Forwarding Plane

The MPLS VPN over mGRE solution simplifies operations in that it only requires a single IP address from each site for transport over the SP network (this is typically a loopback address or the interface facing the SP network). Moreover, this solution is fully capable of supporting quality of service (QoS) because any ToS/EXP markings from packets transiting the PE router will be “reflected” to the outer header of the GRE header, allowing the IP WAN transport the ability to properly prioritize MPLS VPN traffic. In addition to marking, any ingress/egress QoS policies can also be applied such as policing, shaping, and/or scheduling prior to sending to the WAN. Figure 10 details the label stack and forwarding operation for MPLS VPN over mGRE as a packet is sent between sites.

Figure 10. MPLS VPN over mGRE: Data Plane

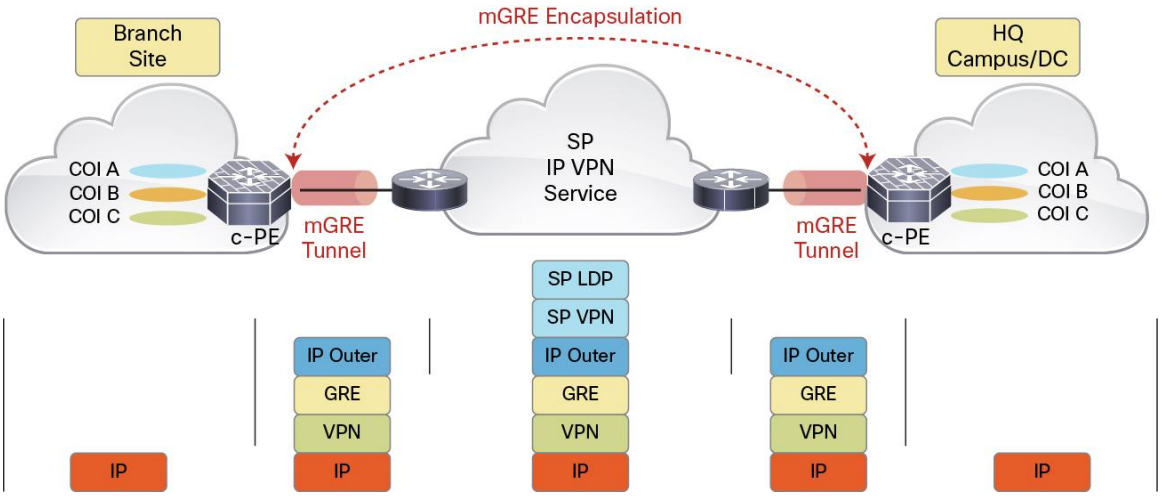

As show in Figure 10, the MP-BGP control plane process has taken place, so as an IP packet is sent from the branch site, the c-PE router (running MPLS VPN over mGRE) attaches the appropriate VPN label for the destination (learned during the control place process) that matches the COI/VRF. The c-PE router then encapsulates the labeled packet in a GRE format, with the HQ campus/DC c-PE router as the destination, before sending the packet to the SP network. The SP then prepends its own VPN and LDP labels for transport (assuming the SP is running MPLS VPN) and forwards the packet with the destination IP address of the branch site. The HQ c-PE router then receives the GRE-encapsulated labeled packet (indicated by the protocol ID 0x8847/8848), deencapsulates the GRE tunnel header, does a lookup on the L3 VPN label, and forwards the packet out the appropriate egress interface associated with that VPN label.

Deployment Models: Extending COIs Across the WAN

There are two models for deploying MPLS VPN over mGRE. The first model, shown in Figure 11, applies to SP-managed IP VPN service, where customers connect to an SP over IP. Although the SP may utilize an MPLS core, the CE/c-PE to PE protocol is IP, and routes are distributed either statically or using a routing protocol, such as eBGP or OSPF. When the MPLS VPN over mGRE feature is configured, the operator can deploy L3 VPN services (RFC 4364) “over the top” of the SP transport network, allowing COIs to be distributed over a WAN. This is a common “over-the-top” model where the core is not owned or managed by the customer.

Figure 11. MPLS over mGRE: SP-Managed Model

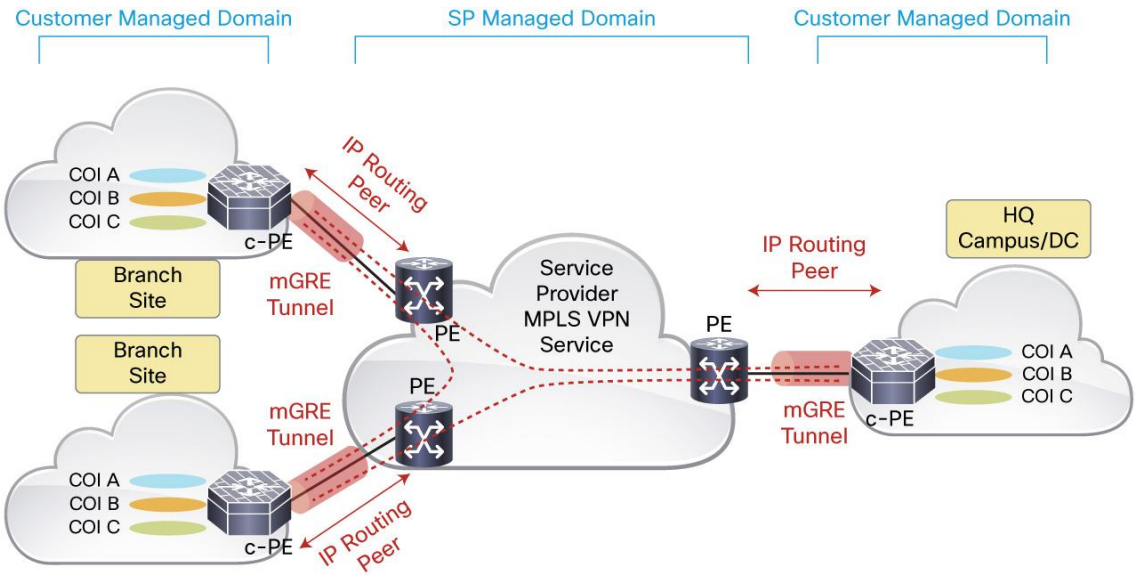

The second model, shown in Figure 12, typically applies to what we have called a “self-deployed” MPLS model, where the customer manages the infrastructure end to end. In this case, the CE, PE, and P routers and entire WAN infrastructure are managed by the IT organization/agency. This offers the IT group complete control of provisioning, SLAs, and security/control as well as rapid time to deployment of service “turnup.”

Figure 12. MPLS over mGRE: Self-Deployed Model

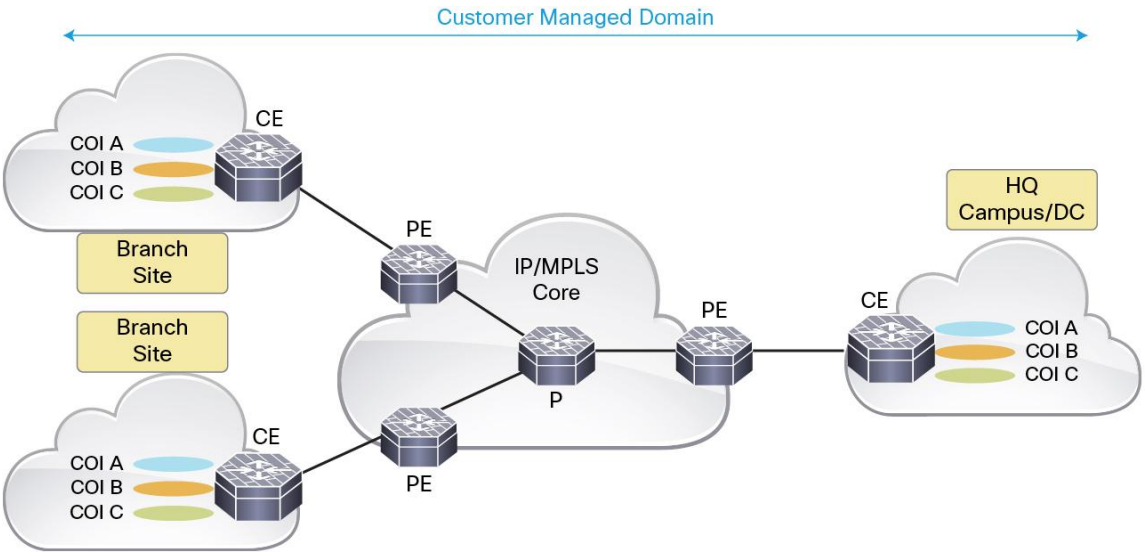

Each model allows one to provision VPN services without using LSP or LDP. The system also utilizes IPv4-based mGRE tunnels to encapsulate both VPN-labeled IPv4 packets, as well as IPv6, allowing a smooth transition to IPv6 services.

Cisco Group Encrypted Transport

Introduction

As powerful as MPLS VPN technology has become for large-scale L3 VPN solutions, one area in which a native MPLS network is limited is the requirement for encryption. MPLS does not natively support encryption, so a primary use case for MPLS over IP is the ability to use industry standard IPsec. After the MPLS packet is encapsulated into IP, it has the ability to be encrypted with IPsec, and this is a typical deployment when the encryption requirement exists. In the case of MPLS VPN over a multipoint GRE tunnel, more innovation is required to use the multipoint nature of mGRE, while also meeting the encryption requirement.

Cisco Group Encrypted Transport (GET) is a next-generation WAN VPN solution, providing a new category of virtual private networks, and a paradigm shift from point-to-point IPsec VPN tunnels for transport security. GET provides end-to-end security for network traffic over a WAN in a tunnelless fashion. It maintains network intelligence such as full mesh topology, utilizing the core network's ability to route traffic using the natural routing path to all destinations across the WAN. GET is ideally suited to encrypt traffic over an IP/MPLS-based core network, because it preserves the original source and destination addresses in the header of the encrypted traffic. GET is a highly scalable technology, eliminating the need for complex peer-to-peer security associations and tunnel configuration. GET also maintains QoS services across the WAN and scalable key management through the use of group encryption keys. Finally, GET supports a universal encryption model, natively allowing the encryption of both IP multicast and unicast traffic.

Cisco GET VPN is standards based, utilizing Group Domain of Interpretation (GDOI, RFC 3547) as the keying protocol and IPsec (RFC 4301) for encryption. The use of GET VPN is described in depth in the Design and Implementation Guide (GET VPN_DIG.pdf) and the Configuration Guide (GET VPN_CG.pdf), so they will not be detailed here. However, some primary concepts must be introduced in order to understand how GET VPN can be utilized to encrypt COI data traffic as it traverses the WAN.

GET VPN Technology Overview

GET VPN uses a tunnelless model of encryption. GET utilizes a group IPsec Security Association (SA) that allows any CE device to encrypt and decrypt traffic from any other CE device on an MPLS/IP WAN. Within the GET VPN architecture, there are two components: Key Servers (KSs) and Group Members (GMs). Referring to our MPLS over mGRE design, the group members are the CE/c-PE devices at the WAN edge. The GET security model is based on the concept of “trusted” group members. The GM registers with the key server to get the IPsec SA that is necessary to encrypt data traffic and communicate with the group.

GET is highly flexible, and GET-based networks can be utilized across a number of different WAN environments, including IP or MPLS cores. Figure 13 shows our MPLS over mGRE design as it applies to the GET VPN architecture.

Figure 13. GET VPN Architecture

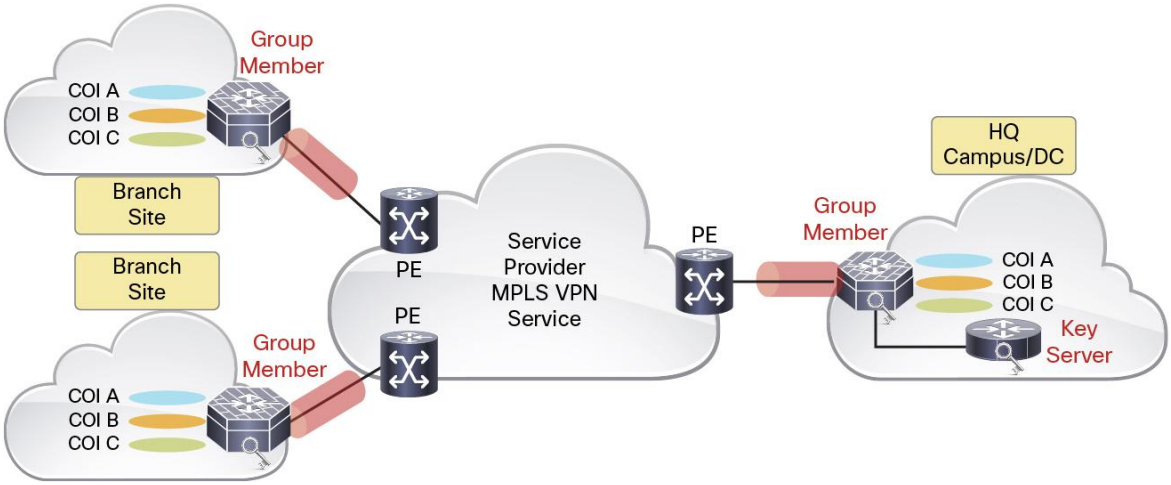

GET Key Server and Group Member

The key server is responsible for maintaining security policies, authenticating the GMs, and providing and maintaining the session key for encrypting traffic. All necessary cryptographic policies are defined on the key server. The KS authenticates the individual group members at the time of registration. The group member registers with the key server to get the IPsec SAs that are necessary to communicate with the group. A GM can register at any time and receive the most current policy and keys.

There are two types of keys that the GM will receive from the KS: the Key Encryption Key (KEK) and the Traffic Encryption Key (TEK). The KEK is used to secure rekey messages between the key server and the group members. The KS sends out rekey messages if an impending IPsec SA expiration occurs or if the policy has changed on the key server. The TEK becomes part of the IPsec SA with which the group members within the same group encrypt the data. Figure 14 highlights the Group Policy, the KEK, and the TEK being distributed from the Key Server to each of the GMs in the GET VPN domain.

Figure 14. GET VPN Security Elements

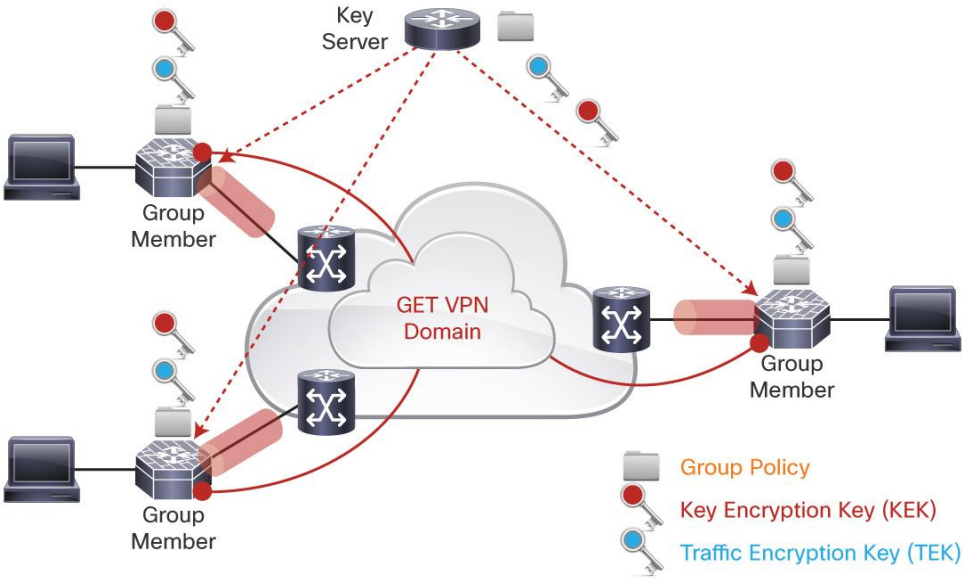

Group Member Authentication

Although group members may authenticate to the key server at registration time using preshared keys, the use of Public Key Infrastructure (PKI) with digital certificates is recommended because it is considered more secure. PKI uses its infrastructure to overcome the key management distribution and scaling issues encountered when preshared keys are used. The PKI infrastructure acts as a certificate authority (CA) where router certificates are issued and maintained.

Key Distribution

The GDOI protocol is used for group key and group SA management. GDOI is defined at the Internet Security Association Key Management Protocol (ISAKMP) Domain of Interpretation (DOI) for group key management. GET utilizes GDOI for authenticating the GMs and KSs. Digital certificates are recommended as the ISAKMP authentication mechanism.

Although a single KS is shown in Figure 14, multiple key servers are supported by the GET VPN architecture to make sure of redundancy, high availability, and fast recovery if the primary key server fails. It is recommended that redundant key servers be utilized in order to support these features and eliminate any single points of failure.

Address Preservation

GET VPN allows IPsec-protected data to carry the original IP header information in the outer header, a technique known as IPsec Tunnel Mode with address preservation. This technique is depicted in Figure 15.

Figure 15. GET VPN Address Preservation

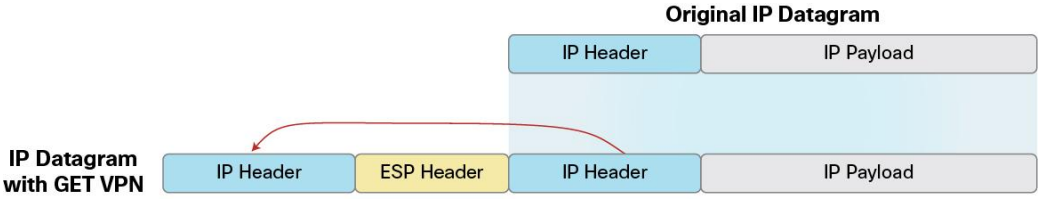

Address preservation allows GET VPN to use the routing functionality present within the core network. For our purposes, the source and destination addresses of the mGRE tunnel will be preserved. Figure 16 shows the frame format for GET VPN, as applied to an MPLS over mGRE datagram.

Figure 16. GET VPN Address Preservation with MPLSomGRE

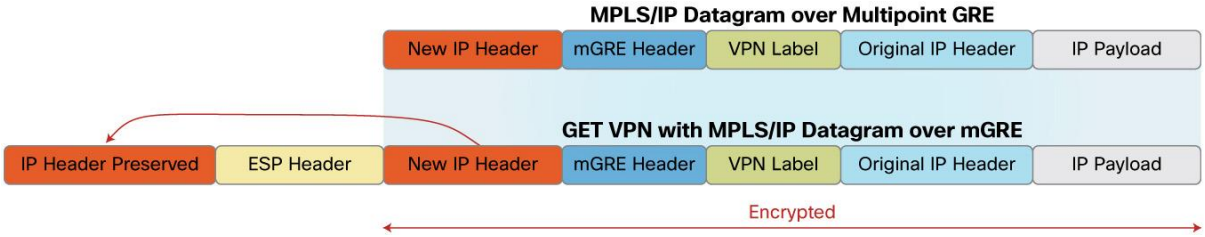

Fail-Close Mode

It is recommended that GET VPN be deployed in Fail-Close mode. With this feature enabled, all traffic passing through the GM will be dropped until the GM has successfully registered with the KS. This makes sure that data that is intended to be encrypted will not be sent out of the router on the protected interface unencrypted or in the clear. With the exception of control plane traffic, all traffic sent across the WAN will exit the router as encrypted data.

MPLS VPN over mGRE with GET VPN

GET VPN in combination with MPLS VPN over mGRE offers a scalable, secure network virtualization solution when IP encapsulation is required. In the MPLS VPN over mGRE deployment with GET, all L3 VPN traffic is encrypted as it egresses the c-PE router, maintaining a secure transport of COI traffic in a multipoint IP environment.

MPLS VPN over mGRE with GET VPN preserves the original source and destination addresses of the mGRE tunnel interface, while encrypting the data payload (MPLS VPNv4 content) in an any-to-any fashion over the WAN. Figure 17 depicts the MPLS VPN over mGRE with GET VPN architecture. Traffic is sent from a particular COI from a branch site to the HQ site, being encrypted/decrypted on each c-PE device. This is the case for both unicast and multicast traffic, without any effects on the MPLS data plane or control plane.

Figure 17. MPLS VPN over mGRE with GET VPN

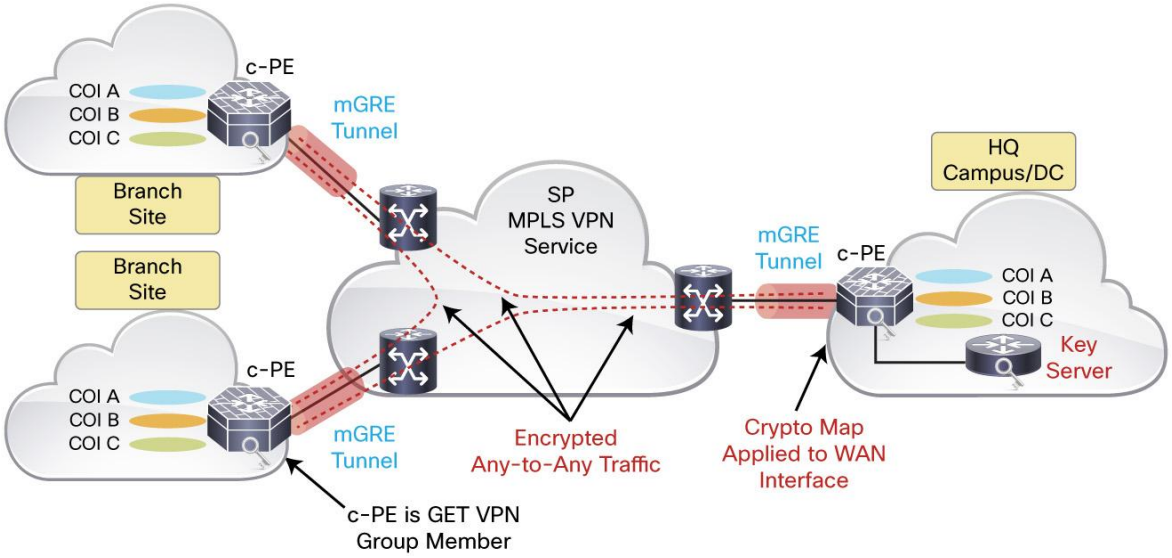

Next-Generation Encryption

Introduction

Cryptography allows for secure communications through protocols and algorithms that allow authentication, data confidentiality, and data integrity. Because modern cryptography relies heavily on mathematics and computing, the science of cryptography is ever changing. To keep up with advances in these fields, it is necessary to adopt newer, stronger algorithms and larger key sizes. Next-Generation Encryption (NGE) is a set of secure ciphers and cryptographic algorithms designed to meet the security and scalability requirements of both governments and businesses for the next two decades. NGE provides new algorithms for encryption, authentication, digital signatures, and key exchange.

NGE Cryptographic Algorithms

Cisco NGE technology is a complete set of algorithms, replacing legacy algorithms for encryption, authentication, integrity, and key exchange. Advanced Encryption Standard (AES) Cipher Block Chaining (CBC) is replaced by AES Galois Counter Mode (GCM) for authenticated encryption at high data rates. Public key algorithms for encryption and digital signatures such as RSA and Diffie-Hellman (DH) key exchange are replaced by Elliptic Curve Cryptography (ECC). For hash functions, SHA-2 replaces the legacy MD5 and SHA-1 algorithms. A detailed overview of NGE algorithms can be found here: Next-Generation Encryption.

NGE, Suite B, and Classified Information

Cisco's NGE technology is compatible with the U.S. National Security Agency's (NSA) Suite B set of ciphers. Indeed, Cisco led the design and standardization of the AES-GCM algorithm that is used for encryption in Suite B. The algorithms utilized in NGE are identical to the algorithms used in Suite B. These include the cryptographic algorithms in AES-GCM, as well as algorithms for hashing, digital signatures, and key exchange. Suite B is the cryptographic basis for protecting both unclassified data as well as classified information, including Secret and Top Secret information. Companies in the private sector can increase the security of transmitting sensitive data over nontrusted networks by utilizing these algorithms.

Cisco has introduced Suite B cryptography in its VPN products, and Suite B algorithms can be applied to IPsec VPNs (defined in RFC 6379). Cisco's NGE algorithms can be employed within our GET VPN framework for encryption of both unclassified as well as classified information as it is transmitted over the WAN.

MPLS VPN over mGRE with GET VPN Utilizing NGE Algorithms

As stated, NGE is compatible with the GET VPN security standard. The GET VPN support for the Suite B features allows these algorithms to be used with GDOI and GET VPN in various ways, including SHA-2 for hash functions and AES-GCM for encryption within IPsec. SHA-2 is utilized within the KEK rekey policy, while AES-GCM is used in TEK IPsec policies as IPsec Security Association (SA) encryption and integrity algorithms.

From a data forwarding perspective, plaintext traffic enters the c-PE router from a particular COI, which is then encapsulated in GRE. The GET VPN process then encrypts the data using IPsec AES-GCM, which is then forwarded out the egress interface toward the WAN. The process is reversed on the receiving side, with traffic exiting the c-PE router as plaintext COI traffic. Figure 18 depicts this process.

Figure 18. MPLS VPN over mGRE with GET VPN and Suite B Encryption

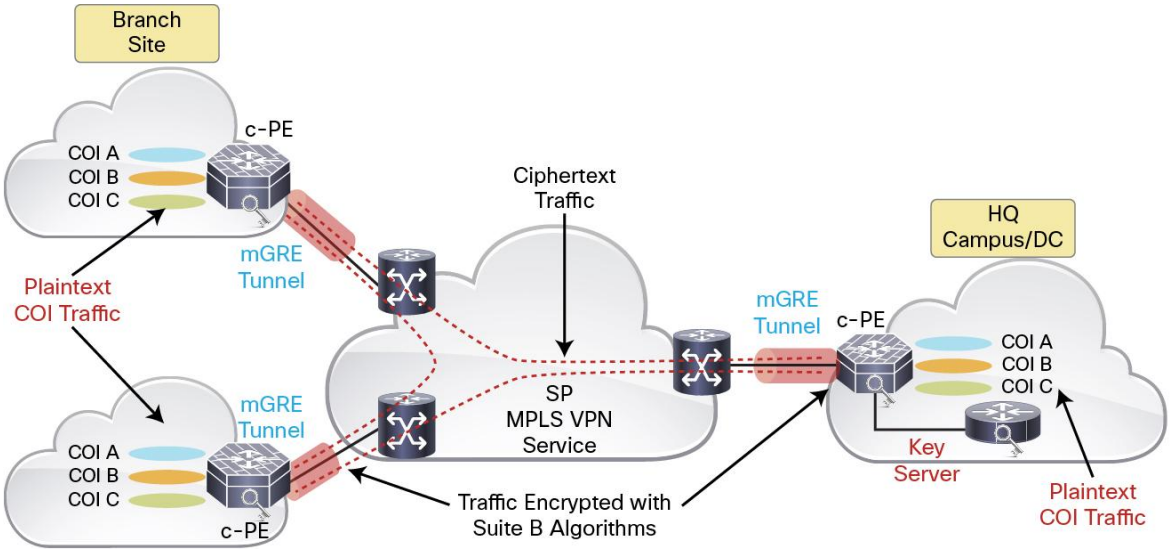

Conclusion

Efficiencies within IT organizations have never been more critical. Budget constraints and cost reductions are on the rise; at the same time, there are competing demands for IT and its role in the business process. The network fabric is mission critical for this transformation, and virtualization of the network infrastructure is a primary enabler of both efficiency and cost reductions. Network virtualization simplifies network operations by enabling customers to securely share a common network infrastructure between groups of users, applications, and devices. Network virtualization allows the participation of communities of interests in a secure manner, while maintaining separation from users outside of the group, over a common infrastructure. The challenge has been the ability to extend COIs across a WAN in a manner that is both scalable and flexible, without decreasing the overall security posture of the network.

Using MPLS VPN over IP innovations described in this paper directly simplifies the scaling, deployment, and operational challenges of current MPLS VPN over IP solutions. This technology, combined with GET VPN utilizing Suite B (NGE), allows the extension of COIs across the WAN in a highly secure manner, suitable for deployment within the DoD, intelligence community, and secure enterprise networks.

References

- RFC 2408: “Internet Security Association and Key Management Protocol (ISAKMP)”
- RFC 2784: “Generic Routing Encapsulation (GRE)”
- RFC 4023: “Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)”
- RFC 4106: “The Use of Galois/Counter Mode (GCM) in IPsec Encapsulating Security Payload(ESP)”
- RFC 4271: “A Border Gateway Protocol 4 (BGP-4)”
- RFC 4301: “Security Architecture for the Internet Protocol”
- RFC 4364: “BGP/MPLS IP Virtual Private Networks (VPNs)”
- RFC 4760: “Multiprotocol Extensions for BGP-4”
- RFC 4797: “Use of Provider Edge to Provider Edge (PE-PE) Generic Routing Encapsulation (GRE) or IP in BGP/MPLS IP Virtual Private Networks”
- RFC 5036: “LDP Specification”
- RFC 6407: “The Group Domain of Interpretation”
- RFC 6379: “Suite B Cryptographic Suites for IPsec”
- NSA Suite B Cryptography: http://www.nsa.gov/ia/programs/suiteb_cryptography/

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.

Printed in USA C11-726689-00 02/13