



Interested in learning more about security?

SANS Institute InfoSec Reading Room

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

Comparing BGP/MPLS and IPSec VPNs

Network managers have many options for site-to-site connectivity. Traditional leased lines, Frame Relay, and ATM based connectivity solutions are giving way to newer VPN technology. Two types of modern VPNs, BGP/MPLS and IPSec are becoming increasingly attractive to network managers. This paper gives an overview of MPLS and then discusses the mechanisms used to provide VPNs based upon BGP/MPLS and IPSec. The paper assesses the security provided by both solutions and suggests guidelines for network managers to assist in...

Copyright SANS Institute
Author Retains Full Rights

AD

Build your business'
breach action plan.

START NOW

 **LifeLock**
BUSINESS SOLUTIONS

No one can prevent all identity theft. © 2016 LifeLock, Inc. All rights reserved. LifeLock and the LockMan logo are registered trademarks of LifeLock, Inc.

Comparing BGP/MPLS and IPsec VPNs

Abstract

Network managers have many options for site-to-site connectivity. Traditional leased lines, Frame Relay, and ATM based connectivity solutions are giving way to newer VPN technology. Two types of modern VPNs, BGP/MPLS and IPsec are becoming increasingly attractive to network managers. This paper gives an overview of MPLS and then discusses the mechanisms used to provide VPNs based upon BGP/MPLS and IPsec. The paper assesses the security provided by both solutions and suggests guidelines for network managers to assist in evaluating these two options.

Introduction

Recent advances in Virtual Private Network (VPN) technology offer great benefits to network managers, especially in terms of cost savings. VPN technology can offer the ability to connect remote offices at a fraction of what it costs for a traditional leased line. Traditional service provider VPNs, such as ATM or Frame Relay, are giving way to more cost effective alternatives. VPN networks based upon IP Security (IPsec) or a combination of Border Gateway Protocol (BGP) and Multiprotocol Layer Switch (MPLS) are viable alternatives that are growing in popularity.

IP Security (IPsec) VPNs are proliferating throughout the enterprise and medium-small business space due to the ability to not only connect remote users, but disparate offices over an existing IP infrastructure. In the case of IPsec, the existing IP infrastructure is the Internet and an individual company can own the VPN equipment. In addition, service providers, such as Virtela, are beginning to offer IPsec based VPN services to clients. (Jones, 1)

BGP/MPLS VPNs, based upon RFC 2547, are an alternative to IPsec VPNs. MPLS VPNs are network based, meaning that they are not the CPE-based VPNs that are more prevalent in the US today. Rather, MPLS VPNs travel over existing service provider IP infrastructures and can only be only utilized as a purchased service.

IPsec VPNs based upon CPE have been popular in the US for the past few years. However, that is not the case throughout the world. "Amid U.S. enterprises, traditional CPE-based VPN solutions are favored. VPN technology overall has enjoyed wide acceptance abroad where networked-based solutions are taking hold more quickly." (Jones, 2)

Despite international differences, it is clear that the use of VPN technology will continue to grow, and an increasing amount of providers, in both the United States and abroad, will offer BGP/MPLS based VPN service. This gives network managers a choice between MPLS or IPsec

VPN technology. This choice should be made after a careful comparison between the two options.

MPLS - The Basics

An understanding of the basics behind MPLS is required for understanding MPLS based VPNs. MPLS evolved for a number of reasons. The first was that it provided a more scaleable method of allowing IP traffic to travel over an ATM network.

Second, MPLS enhances routing functionality. On a traditional service provider IP Network, traffic is routed via an interior routing protocol such as OSPF. The sum result of relying on traditional Layer 3 routing protocols is that all traffic routes based upon best path. This means that the best paths through service provider cores are used heavily while other, redundant paths, are underutilized. MPLS can be used to move network traffic over these underutilized paths allowing a greater return on redundant architecture investment for the service provider.

Finally, MPLS allows service providers to offer customers enhanced services. Utilizing constraint based routing and MPLS's quality of service features, service providers can provide several layers of service to their customers based upon many factors. Of course, MPLS also allows a service provider to offer value added services such as VPNs.

MPLS is based upon routers, or switches, performing label switching to provide a Label Switched Path (LSP) through a network. Essentially, when an IP packet enters an interface of an MPLS ingress router, that router assigns the packet to a Forwarding Equivalency Class (FEC). FEC assignment can be based on a variety of Access Control List (ACL) matches such as source address, destination address, BGP next hop, application type, and Differentiated Service flag. (Cosine, 2) The router then assigns the packet to an LSP and "pushes" a label onto the packet header before sending it out the appropriate interface. Since FEC assignment is based upon a number of different factors, rather than just destination, packet forwarding can be policy based providing an advantage over traditional IP networks.

Similar to an ATM VPI/VCI or a Frame Relay DLCI, MPLS labels only have local significance. Intervening Label Switch Routers (LSRs) "swap" the incoming label with a label defined in their own MPLS forwarding database. When the MPLS egress, or final, router is reached, the label is permanently removed, or "popped", prior to the egress router forwarding the regular IP packet.

An example of this process is shown in Figure 1. An IP packet arrives at the ingress router with a destination of 192.4.2.10. The ingress router pushes a label of 10 onto the packet and forwards it out a second interface. The intermediary LSR does a lookup on the label of 10 and swaps the 10 with a label of 2. The egress router terminates the LSP by popping the label of 2 and forwarding the regular IP packet onward.




Figure 1. Basic MPLS Operation

When a label is pushed, a MPLS header is added to each packet. MPLS headers contain the following information (Figure 2):

- A 20 bit label field that contains the actual label information.
- A 3 bit Class of Service field that can have an impact on the packet as it travels the network.
- A 1 bit stack field that supports a hierarchical label stack.
- An 8 bit Time to Live field which can support conventional IP TTL functionality (Semeria, 13)




Figure 2. MPLS Header Information

LSPs can be setup throughout a network utilizing a variety of methods. LSPs can be manually configured on each LSR that participates within the LSP. Constraint based routing protocols such as Constrained Shortest Path First (CSPF) can also be used to define LSPs. Label Distribution Protocol (LDP), Resource Reservation Protocol (RSVP), and Constraint Based Routing-Label Distribution Protocol (CR-LDP) are signaling methods used to setup the LSP through a network.

BGP/MPLS VPN Mechanics

RFC 2547 defines BGP/MPLS VPNs. In defining BGP/MPLS VPNs, the authors write:

MPLS is used for forwarding packets over the backbone, and BGP is used for distributing routes over the backbone. The primary goal of this method is to support the outsourcing of IP backbone services for enterprise networks. It does so in a manner which is simple for the enterprise, while still scalable and flexible for the Service Provider, and while allowing the Service Provider to add value. These techniques can also be used to provide a VPN which itself provides IP service to customers. (Rosen and Rekhter, 1)

The RFC defines three functional types of routers that comprise a VPN. Customer edge (CE) routers sit at the customer site and are typically owned by the customer. However, some service providers provide equipment for CE routers. CE routers are connected to provider edge (PE) routers. Owned by service providers, these are usually located at POPs and are the entry points into the provider's network. Finally, provider (P) routers are defined as transit routers within the provider's core.

Routing information is passed between a CE and its directly connected PE through an IGP, BGP, or through default routes defined on each router. Each PE router keeps one or more per-site forwarding tables known as VPN Routing and Forwarding Tables (VRFs). Within every PE, each VRF serves a particular interface, or set of interfaces that belong to each individual VPN. In other words, each interface that leads to a customer within a PE has its own VRF. The only time this is not the case is when a customer has two or more sites connected to the same PE. In this case, the VRF is used for all interfaces facing that customer. Each PE is configured with its own unique VRF meaning that routers do not share VRF information directly. (Figure 3)




Figure 3. VRFs are unique for each VPN connected to a particular PE.

One of the requirements for a service provider VPN is that it support overlapping address space. Since many enterprises utilize RFC 1918 address space within their networks, it is likely that at least two customers utilize the same IP addresses. This can cause confusion since routing protocols assume that each IP address is unique. BGP/MPLS VPNs utilize the VPN-IPv4 address family combined with multiprotocol extensions to BGP to solve this issue. Since these mechanisms function on the service provider's backbone, they are hidden from customers.

A VPN-IPv4 address is a 12 byte address that begins with 8 byte Route Distinguisher (RD) and ends with a 4 byte IPv4 address. RFC 2547 states, "If two VPNs use the same IPv4 address prefix, the PEs translate these into unique VPN-IPv4 address. This ensures that if the same

address is used in two different VPNs, it is possible to install two complete different routes to that address, one for each VPN.” (Rosen and Rekhter, 8)

RDs were developed so that service providers could administer their own address space without conflicting with a different provider. The 8 byte RD is made up of a 2 byte Type field and a 6 byte Value field. The Value field contains two subfields, Administrator and Assigned Number. The Type field determines the lengths of the two subfields. There are currently two types defined. In Type 0, the Administrator subfield contains 2 bytes and the Assigned Number 4 bytes. The Administrator subfield is made up of the service provider's Autonomous System Number (ASN). The Assigned Number subfield is a number assigned by the service provider that is globally unique within that provider only.

Type 1 RDs reverse the byte counts for the two Value subfields. The Administrator subfield is 4 bytes and the Assigned Number subfield is 2 bytes long. In a type 1 RD, the Administrator subfield holds an IPv4 address. This address is from the address space assigned to the service provider and is most commonly the loopback address of the PE router that originates the route. (Semeria, 13)

There are two important traffic flows that occur in a BGP/MPLS VPN. The first is a control flow that is used for route distribution and LSP definition. The second data flow is used to forward customer traffic.

There are two control mechanisms within BGP/MPLS VPNs. The first is responsible for the exchange of routing information between different PE routers that make up a VPN. The second is for the establishment of LSPs across a service provider's backbone (Semeria, 8).

PEs learn customer routes from CE routers. This can be done through an IGP, BGP, or through static configuration on the PE. PEs install these routes into the appropriate VRF along with an assigned MPLS label. The assigned MPLS label is used to signify which CE interface the routes point to. The routes are advertised via RFC 2283 Multiprotocol-IBGP to other PE routers. Multiprotocol BGP allows BGP to support address families other than IPv4 such as IPv6 or VPN-IPv4 address.

Route distribution is filtered based upon route targets using BGP extended community attributes. When a VPN route is injected into BGP, VPN route target extended community attributes are attached to it. These are taken from the BGP export list associated with the VRF from which the route is learned. Each PE also contains an import list of route target communities associated with each VRF. The import list acts as an ACL defining which routes can be imported into a VRF. For example, if an import list for a particular VRF defined within a PE includes target communities A and C, but not B, then any VPN route that carries communities A and C is imported into that VRF. However, VPN routes that carry VPN community B will not be imported. This assures that VPN routes belonging to a particular customer are only installed in VRFs associated with that customer (Cisco, 9).

LSP establishment for VPN tunnels may be accomplished through LDP or RSVP. A provider would use LDP if it wants to establish best effort routing between PEs using a particular IGP.

However, if a service provider wants to assign bandwidth requirements, other constraints, or offer advanced services RSVP must be used to signal LSP path.

Packet flow within a BGP/MPLS VPN is shown in Figure 4. The router, CE1, is sending data destined for a network behind CE2. It forwards its packets to its default gateway, PE1. PE1 performs a route lookup in the VRF that is defined for Site 1. It finds that PE2 advertised the route to Site 2 with a label of 5. It also finds the BGP next hop for the route, which is the loopback address of PE2. Finally, PE1 looks up the route to PE2 and finds that it should use an LSP where it pushes a label of 100 before forwarding the data out a particular interface. In other words, the ingress router, PE1 pushes two labels onto the packet. The inner label, 5, is used by PE 2 to match the packet with a particular CE. The outer label, 100, is used to label switch the packet through the network.

The packet arrives at router P1. P1 swaps the outer label of 100 with 2 before forwarding it. P2 is the penultimate router for the LSP, so it pops the outer label before forwarding the packet to PE2. PE2 uses the inner label, 5, to identify the CE that is the next hop. It pops the inner label creating an IPv4 packet that it then forwards to CE2.




Figure 4. BGP/MPLS VPN Packet Flow

Evaluating BGP/MPLS VPNs

BGP/MPLS VPNs offer an alternative for secure site-to-site communication. Cisco reports, “The final report [Miercom] shows that it was not possible to attack the network in any way, nor other VPNs on the same network. In addition, neither the address space nor the routing separation between VPNs could be overcome. The end result of the report is that MPLS is at least as secure as Frame Relay and ATM networks. (Cisco, 17).”

BGP/MPLS VPNs have built within them several mechanisms to provide security. Address space separation may be a concern, especially considering the possibility of overlapping address space. However, the use of VPN-IPv4 addresses allows for independent VPNs to remain separate despite any addressing overlap.

Routing separation is achieved on PE routers through separate VRFs for each customer connected to that PE. Furthermore, routes are kept separate in BGP routing updates through the use of unique identifiers in BGP route target extended attributes. When properly configured for address space and routing separation, BGP/MPLS VPNs offer “the same security as comparable

Layer 2 VPNs such as ATM or Frame Relay. It is not possible to intrude into other VPNs through the MPLS cloud, unless this has been configured specifically (Cisco, 4).”

Furthermore, because these mechanisms are internal to the service provider, the structure of the service provider core is invisible to customers. CE routers know their next hop for PE routers, but not the address of any P routers. Thus, there is no intelligence about the core that customers can gain. This protects against potential attackers using information about the core for a denial of service, spoof, or session hijack attack.

It may be possible to spoof the customer address space. However, since each VRF is tied to the interface, or interfaces, that point to the customer, the attacker must be on the customer premises to spoof an IP address within the VPN. If a VPN customer takes precautions to secure its internal network within a site it is unlikely that this can occur. In this regard, this provides the same level of security as IPSec. The existence of a given link between PE and CE router and its association with a particular VRF within the PE replaces the need to authenticate the VPN end points since they must physically be connected by the customer in the proper location.

It may also be possible to spoof MPLS labels. However, there are two mitigating factors. First, the interface between the CE and the PE is an IP interface. No LSP should travel that interface, and a properly configured PE should drop any MPLS traffic coming from a CE. Second, labels only have local significance to each P router. That means that an attacker would not only have to gain physical access to a service provider network, but also be in the right location on the network to target a specific VPN with a spoofed label. This is not necessarily impossible to do, but appropriate security procedures at service providers should prevent this.

One of the largest concerns considering the security of MPLS VPNs is the complexity in configuring service provider networks with multiple VRFs and target communities. Interactive week quotes Randy Bush of AT&T Labs:

The technique uses Border Gateway Protocol, which sets up routing tables in large networks. The tables are code sequences that tell routers how to forward packets from one machine to another. Many service providers already have trouble managing these tables to ensure good connectivity, because a change in one table affects many others. Adding thousands of such tables – as proposed by Cisco and Juniper for individual VPN sessions – would make the risk of managing this software close to impossible, Bush said.

“It severely complicates the core,” Bush said. “It’s a serious issue of managing one BGP table – you want me to have how many thousand of them? But it is a great scam to sell more routers to hold all those BGP tables!” (Smetannikov, 1)

It should be noted that Juniper Networks offers software to assist in managing BGP\MPLS VPN setup and service providers such as AT&T are developing their own specialized tools (Smetannikov, 2).

Internet access is vital for modern companies. Any VPN must provide a mechanism for users to access the Internet. There are three different ways to accomplish this using MPLS VPNs. First, a second circuit could be used to access the Internet. This is analogous to how most Frame or ATM customers currently access the Internet. Second, entries into the VRF for a particular VPN could point to the Internet and an LSP could end at any one of a service provider's Internet connected routers. This requires each site to maintain a firewall on premises, or service providers to place a firewall for each customer before the PE router facing the Internet (Cisco, 11). Third, some service provider equipment vendors are developing virtual firewalls that can run within each PE router. Cosine is one of those providers. Their equipment allows service providers to run instances of Checkpoint or Gauntlet within their PE equipment.

MPLS VPNs offer an economic advantage over traditional layer 2 VPN technologies or service based IPsec VPNs. Cosine estimates the savings to service providers as a minimum of \$400 per site. (Cosine, 3) In a competitive marketplace, this is likely to be passed onto customers.

Finally, the largest benefit to MPLS VPNs is their scalability from an enterprise perspective. Layer 2 and IPsec VPNs traditionally are point-to-point. Thus, a hub and spoke architecture is the most logical to design from a cost perspective. However, providers can easily configure MPLS VPNs for a full mesh topology. Not only does this make it easier to troubleshoot inter-site routing issues, but also it can also greatly ease the deployment of latency sensitive applications such as voice or video. BGP/MPLS VPNs could also be combined with the advanced QoS features of MPLS for quality guarantees. This allows network managers to take advantage of advanced MPLS features.

The advantages of BGP/MPLS VPNs were best expressed by Cosine:

The economic value proposition of BGP/MPLS VPNs will benefit small and medium businesses and large enterprises since both will eagerly pursue connectivity solutions that allow them to connect remote office sites together in a manner not previously affordable. High-speed, low-cost broadband technologies coupled with BGP/MPLS VPNs solidify the economic argument for these services. Additionally, the network-based delivery model allows for hassle-free provisioning and maintenance of fully meshed site-to-site VPNs. BGP/MPLS VPNs allow enterprises to move away from the rigid hub-and-spoke hierarchy of Frame Relay/ATM to an any-to-any connectivity model that leverages the flexibility of the IP network. (2)

IPsec VPN Mechanics

Today's IPsec VPNs are currently defined by RFC 2401 which made obsolete RFC 1825. The authors explain:

IPsec provides security services at the IP layer by enabling a system to select required security protocols, determine the algorithm(s) to use for the service(s), and put in place any cryptographic keys required to provide the requested services. IPsec can be used to protect one or more "paths" between a pair of

hosts, between a pair of security gateways, or between a security gateway and a host. (Kent and Atkinson, 4)

Essentially, IPSec is a suite of protocols designed to provide for a secure IP based pathway between two devices. Since IPSec is device dependant, rather than network dependant, it offers greater flexibility than BGP/MPLS VPNs in terms of deployment. Service providers are not required to perform extra configuration on their routers. IPSec VPNs can be configured between hosts or between sites using security gateways.

IPSec utilizes two protocols that provide for secure transport over IP. Authentication Headers (AH) can be added to an IP packet to provide for message integrity. AH adds a keyed-hash to the header of an IP packet to ensure message integrity and authenticity. Encapsulating Security Payload (ESP) headers are added to an IP datagram to provide for encryption and traffic flow confidentiality.

AH and ESP are appended to an IP packet after the layer 3 header. Figure 5 provides a diagram demonstrating ESP placement in an IP packet.

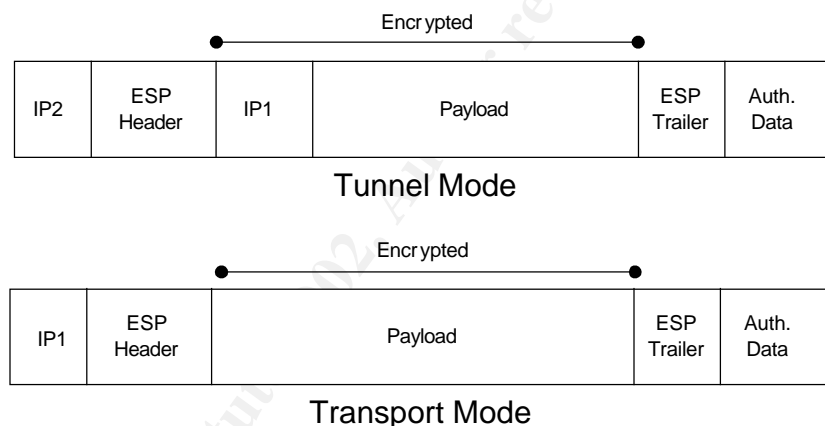


Figure 5. ESP Placement in an IP Packet Using Tunnel or Transport Mode

IPSec provides for two different modes to transport encrypted data: transport and tunnel modes. Transport mode is used solely between two hosts. Transport mode can provide separate tunnels for each protocol session between two hosts. It inserts a header immediately after the IP header and any options, and before any higher layer protocols (Kent and Atkinson, 8). Tunnel mode must be used for gateway-to-gateway sessions as well as host-to-gateway sessions. Tunnel mode provides a single encrypted tunnel over which all traffic must travel. It works by encryption the original IP header for a packet and creating a new header with a destination address of the remote IPSec gateway.

IPSec does not specify which encryption and hashing algorithms must be used. Common encryption methods include DES and 3DES. Popular hashing algorithms used are SHA-1 and MD5. Diffie-Hellman is the preferred method of exchanging encryption keys. However, since IPSec does not specifically require certain encryption and hashing algorithms, it provides a

method for devices to negotiate a common scheme. Two devices, or end points, decide on which algorithms to use and share session keys to create a Security Association (SA). SA information is installed into a Security Policy Database (SPD). VPN gateways use the SPD to lookup information about how to handle packets for a particular VPN tunnel. Before sending a packet, an IPSec gateway adds a Security Policy Identified (SPI) into the header of the packet. The receiving gateway uses destination information and the SPI to lookup information within its own SPD on how to decrypt and/or authenticate the packet.

IPSec involves the exchange of encrypted data. In order to do this properly, encryption keys must be exchanged. Internet Key Exchange (IKE) is the preferred method of doing so. In addition, IKE is the protocol used to establish the SA since IPSec already assumes the SA is in place. Alternatively, each device can be manually keyed. However, manual keying has a severe limitation. “You would have to manually key each device. This solution may be acceptable for small environments with few users, but it does not scale well. If you plan to have more than just a handful of systems passing IPSec traffic, then Internet Key Exchange is a beautiful thing – another TCO triumph! Also, recall that when using manual-keyed IPSec, no replay protection is provided (Smith, 3).”

IKE has a two step process for establishing a SA between two nodes. In Phase 1, two nodes establish a secure, authenticated channel with which to communicate. They do this through a 6-way handshake in which they share temporary IKE SA proposals, Diffie-Hellman keys and nonces, and verify each other’s identity. In Phase 2, permanent security associations are negotiated and a secure tunnel for traffic is established. The Phase 2 3-way handshake takes place over the encrypted tunnel formed in Phase 1. In other words, Phase 1 sets up an encrypted tunnel for Phase 2 information to be passed. Phase 2 then sets up the VPN.

Figure 6 provides an example of a tunnel mode VPN. Host A sends traffic ultimately destined for Host B to its default gateway, G1. G1 is an IPSec gateway and firewall. It performs a rule lookup and determines that traffic to Host B’s network must be encrypted. It then performs a lookup in its SPD and does not find an already established SA for Host B’s gateway, G2. G1 must then establish an SA with G2 before traffic can be passed. G1 initiates Phase 1 IKE. After Phase 1, both gateways use Phase 2 to create data transport tunnel SA. By comparing IKE “proposals” they determine they can commonly communicate using 3DES/SHA1 for encryption and hash algorithms. Each side inserts this information into its SPD. Encryption keys are exchanged and G1 encrypts the packet using 3DES and adds a SHA-1 hash. G2 receives the packet, looks up the SPI contained in the header in its own SPD, and finally checks the hash and decrypts the packet. G2 then sends the packet to Host B. The VPN is invisible to both hosts and to the Internet




Figure 6. Tunnel Mode IPsec VPN

Finally, it should be noted that like an LSP, IPsec SAs are unidirectional. In the example above, G2 must initiate the creation of an SA to allow Host B to respond to Host A.

Evaluating IPsec VPNs

IPsec VPNs offer a high degree of security for information transmitted over the Internet. In discussing the threats to traffic transmitted over the Internet, Cisco suggests:

Confidentiality, integrity, and authentication are key services used to protect against the threats outlined in the previous section. Obviously, if data is encrypted while in transit, it is impossible for a perpetrator to observe or modify. The other threats, identity spoofing, and denial-of-service, can be prevented with strong network-layer authentication. If devices can positively identify the source of data, then it is much harder to impersonate a friendly device and to anonymously implement a denial-of-service attack. (4)

Privacy is maintained in IPsec VPNs through the use of encryption. SAs are renegotiated periodically to provide for even greater protection. Although 3DES is recommended over DES due to the fact that DES can be decrypted, renegotiating session keys limits the amount of exposure.

IPsec requires each side to authenticate with the other. This is accomplished using pre-shared keys, PKI, or digital signatures (Cisco, 9). Authentication when combined with hashing and encryption provides all the methods needed for IPsec to provide for confidentiality and integrity of information.

However, IPsec is not without its drawbacks. First, IPsec is not inherently scalable. This becomes evident if one attempts to create a large site-to-site VPN or an enterprise sized remote access VPN. Each host in the remote access VPN must be configured independently. Each site in a site-to-site VPN must also be configured independently. This creates a large administrative burden on enterprises implementing VPNs internally. Furthermore, they can be complicated to configure. Errors in configuring IPsec VPNs can have serious consequences. Eric Hines notes, "While VPN technologies are powerful, they are not foolproof. More to the point, they are not

hack-proof. They can be extremely difficult to configure properly, and even a small error can create a serious hole in a firewall.” (4)

More importantly, although IPSec theoretically provides for ample security, implementations of IPSec fall short. Hines again notes:

Instead of targeting the encryption itself, attackers may prefer to set their sites on the improper implementation of such a technology from vendor to vendor. This should not therefore be construed as a criticism of IPSec itself, rather as an attack on the way vendors have architected their VPNs. With no standards set in place, each vendor is faced with their own interpretation and ideology of what a VPN appliance should be. In the experience of the author, this has been nothing but a history of failed attempts.

Another difficult that the VPN industry has faced (and continues to face) is an influx of vendors attempting to develop products with a degree of functionality that does not belong in a security device. This misguided tendency is exemplified by vendors who have added such functions as archaic versions of Sendmail to turn a VPN into a combination VPN and Mail Server. (3)

Hines continues by noting two examples where specific vendor implementations led to inherent security flaws in their IPSec VPN gateways. So, while they do provide a high degree of security, IPSEC VPNs are not necessarily foolproof.

BGP/MPLS and IPSec VPNs Compared

Network managers should consider both BGP/MPLS and IPSec VPNs when looking to implement a VPN solution for site-to-site connectivity. The following factors should be used to evaluate the two solutions:

- **Data Confidentiality** – IPSec VPNs provide data confidentiality through robust encryption algorithms. BGP/MPLS VPNs seek to ensure data confidentiality by defining a single path between physical sites on a service provider network. This prevents attackers from accessing transmitted data unless they place sniffers on the service provider network. Though BGP/MPLS minimizes the chance that data may be intercepted, IPSec provides for better data confidentiality through encryption. A third option is to use IPSec over BGP/MPLS VPNs. This option would certainly provide a very high degree of data confidentiality.
- **Data Integrity** – IPSec uses hashing algorithms to ensure data integrity. There are no inherent methods to provide data integrity within BGP/MPLS VPNs. However, the odds of data being altered by a man-in-the-middle attack is low due to the separation of address space and routing information provided by BGP/MPLS VPNs.
- **Data Availability** – IPSec relies on the Internet for transport. Although an attacker could not read the data, an attacker could DoS an IPSec VPN by entering false routes into

Internet routing tables. BGP/MPLS VPNs rely on LSPs for transport, and since LSPs have local significance only, spoofing is very difficult to accomplish. BGP is used to transport routing information within the VPN, however, the use of route targets and the BGP community attributes makes this extremely unlikely. Miercom tested Cisco's implementation of BGP/MPLS technology and was unable to accomplish this feat. (Cisco, 17) Thus, BGP/MPLS can be said to provide for better data availability in this regard. Both implementations allow for redundant gateways and paths between sites, though this is easier to achieve with a BGP/MPLS VPN.

- Remote Access – Though many service provider equipment vendors provide support for remote access solutions, they are not inherent to BGP/MPLS VPNs. Furthermore, they either require remote access users to be on the same service provider network, or service providers must be agree to implement hierarchical BGP/MPLS VPNs. IPsec provides native remote access VPN support and is ultimately superior in this regard.
- Internet Access – Most IPsec VPNs utilize the Internet for transport. Therefore, most IPsec VPN architectures allow VPN access to sites they are connecting. This is more difficult to accomplish within BGP/MPLS VPN architecture, though it is possible. Of the three options to connect BGP/MPLS VPNs to the Internet, a separate Internet connection is the most likely to be chosen by network managers who wish to manage their own Internet perimeter security solutions. Both options provide risk, especially considering configuration errors.
- Scalability – IPsec VPNs are difficult to scale. Due to configuration requirements, IPsec usually leads to point-to-point connections and a hub-and-spoke architecture. BGP/MPLS VPNs are configured by service providers and can easily be used to provide a fully meshed network architecture. Furthermore, BGP/MPLS VPNs allow network managers to take advantage of advanced MPLS features such as Quality of Service guarantees. Therefore, BGP/MPLS VPNs can be considered more scaleable than IPsec in an enterprise environment.

Both BGP/MPLS and IPsec VPNs have their advantages. BGP/MPLS VPNs are more scaleable and most likely provide better availability. However, IPsec VPNs provide for better data confidentiality and integrity. Both types of VPN are difficult to configure and poor implementation is a concern for either solution. However, both solutions are viable for site-to-site connectivity and network managers should carefully weigh the pros and cons of each when deciding on alternatives to traditional leased lines.

Cited References

Cisco Systems. "Product Documentation: MPLS Virtual Private Networks". December 8, 1999.

URL:

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120t/120t5/vpn.htm>

Cisco Systems. "White Paper: Security of the MPLS Architecture". August 14, 2001. URL: http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/mxinf_ds.htm

Cisco Systems. "White Paper: IPSec". November 21, 2000. URL: http://www.cisco.com/warp/public/cc/so/neso/sqso/eqso/ipsec_wp.htm

Cosine Communications. "White Paper: Integrated Provider Edge (PE) Services Switch: Maximizing the Value of Your BGP/MPLS VPN Deployment". November, 2001. URL: http://www.cosinecom.com/library/mpls_vpn_wp.html

Hines, Eric. "Virtual Private Networks: A Broken Dream?". Security Focus, September 10, 2001. URL: <http://www.securityfocus.com/infocus/1461>

Jones, Jennifer. "VPN Connection Options Expand". InfoWorld, November 2, 2001. URL: <http://staging.infoworld.com/articles/fe/xml/01/11/05/011105feedge.xml>

Kent & Atkinson. "Security Architecture for the Internet Protocol". RFC 2401, November 1998. URL: <http://www.ietf.org/rfc/rfc2401.txt>

Rosen and Rekhter. "BGP/MPLS VPNs". RFC 2547, March 1999. URL: <http://www.ietf.org/rfc/rfc2547.txt>

Semeria, Chuck. "Mutiprotocol Label Switching: Enhancing Routing in the New Public Network". Juniper Networks, September 1999. URL: <http://www.juniper.net/techcenter/techpapers/200001.html>

Semeria, Chuck. "RFC 2547bis: BGP/MPLS VPN Fundamentals". Juniper Networks, March 2001. URL: <http://www.juniper.net/techcenter/techpapers/200012.html>

Smetannikov, Max. "Beware MPLS VPN Tech Challenges". Interactive Week, August 20, 2001. URL: <http://www.zdnet.com/filters/printerfriendly/0,6061,2805408-2,00.html>

Smith, Christopher. "IPSec's Role in Network Security: Past, Present, and Future". SANS, September 2001. URL: http://www.sans.org/infosecFAQ/encryption/ipsecs_role.htm

Additional References

Semeria, Chuck. "RFC 2547bis: BGP/MPLS VPN Hierarchical and Recursive Applications". Juniper Networks, July 2001. URL: <http://www.juniper.net/techcenter/techpapers/200014.html>

Davie, Bruce and Yakov, Rekhter. MPLS Technology and Applications. Academic Press, 2000.



Upcoming SANS Training

[Click Here for a full list of all Upcoming SANS Events by Location](#)

SANS Security East 2018	New Orleans, LAUS	Jan 08, 2018 - Jan 13, 2018	Live Event
SANS Amsterdam January 2018	Amsterdam, NL	Jan 15, 2018 - Jan 20, 2018	Live Event
Northern VA Winter - Reston 2018	Reston, VAUS	Jan 15, 2018 - Jan 20, 2018	Live Event
SEC599: Defeat Advanced Adversaries	San Francisco, CAUS	Jan 15, 2018 - Jan 20, 2018	Live Event
SANS Dubai 2018	Dubai, AE	Jan 27, 2018 - Feb 01, 2018	Live Event
SANS Las Vegas 2018	Las Vegas, NVUS	Jan 28, 2018 - Feb 02, 2018	Live Event
Cyber Threat Intelligence Summit & Training 2018	Bethesda, MDUS	Jan 29, 2018 - Feb 05, 2018	Live Event
SANS Miami 2018	Miami, FLUS	Jan 29, 2018 - Feb 03, 2018	Live Event
SANS Scottsdale 2018	Scottsdale, AZUS	Feb 05, 2018 - Feb 10, 2018	Live Event
SANS London February 2018	London, GB	Feb 05, 2018 - Feb 10, 2018	Live Event
SANS Southern California- Anaheim 2018	Anaheim, CAUS	Feb 12, 2018 - Feb 17, 2018	Live Event
SANS Secure India 2018	Bangalore, IN	Feb 12, 2018 - Feb 17, 2018	Live Event
SANS Dallas 2018	Dallas, TXUS	Feb 19, 2018 - Feb 24, 2018	Live Event
SANS Brussels February 2018	Brussels, BE	Feb 19, 2018 - Feb 24, 2018	Live Event
SANS Secure Japan 2018	Tokyo, JP	Feb 19, 2018 - Mar 03, 2018	Live Event
Cloud Security Summit & Training 2018	San Diego, CAUS	Feb 19, 2018 - Feb 26, 2018	Live Event
SANS New York City Winter 2018	New York, NYUS	Feb 26, 2018 - Mar 03, 2018	Live Event
CyberThreat Summit 2018	London, GB	Feb 27, 2018 - Feb 28, 2018	Live Event
SANS London March 2018	London, GB	Mar 05, 2018 - Mar 10, 2018	Live Event
SANS San Francisco Spring 2018	San Francisco, CAUS	Mar 12, 2018 - Mar 17, 2018	Live Event
SANS Secure Osaka 2018	Osaka, JP	Mar 12, 2018 - Mar 17, 2018	Live Event
SANS Secure Singapore 2018	Singapore, SG	Mar 12, 2018 - Mar 24, 2018	Live Event
SANS Paris March 2018	Paris, FR	Mar 12, 2018 - Mar 17, 2018	Live Event
SANS SEC460: Enterprise Threat Beta	OnlineCAUS	Jan 08, 2018 - Jan 13, 2018	Live Event
SANS OnDemand	Books & MP3s OnlyUS	Anytime	Self Paced