



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.

What Is an MPLS VPN Anyway?

This paper is about a new technology, MPLS VPN, that is being offered by service providers to compete with Frame Relay and ATM networks. When a company wants to connect its geographically different sites they don't have to purchase a Frame Relay circuit, or purchase an ATM circuit, or lease a dedicated telco line. They can now go to their Internet service provider and purchase a MPLS VPN to connect their geographically different sites. This paper will give a basic understanding of how a MPLS VPN works.

Copyright SANS Institute
Author Retains Full Rights

AD

Let Us Hack You.
Before Hackers Do!
It's Here — The Cenzic Website HealthCheck

FREE

CENZIC

Request one now

The advertisement features a dark red background with a stylized image of a person's face on the left. The text is in yellow and white. A yellow starburst contains the word "FREE". The Cenzic logo is on the right, and a button with a play icon says "Request one now".

What is a MPLS VPN anyway?
Kelly DeGeest

This paper is about a new technology, MPLS VPN, that is being offered by service providers to compete with Frame Relay and ATM networks. When a company wants to connect its geographically different sites they don't have to purchase a Frame Relay circuit, or purchase an ATM circuit, or lease a dedicated telco line. They can now go to their Internet service provider and purchase a MPLS VPN to connect their geographically different sites. This paper will give a basic understanding of how a MPLS VPN works.

First lets start with a little background to explain why the need for MPLS. As for the need for a VPN, there is plenty of reading material in the SANS reading room on VPN technology. Andrew Egorov has written an excellent piece on VPN deployment that explains the need for a VPN and some problems encountered deploying them. The title is Implementing Virtual Private Networks – Observations from the field. It can be found at the following URL; http://www.sans.org/infosecFAQ/encryption/implement_VPN.htm .

You see the Internet has gotten very big with Internet backbone routers having to hold 100,000+ BGP routes(1). Border Gateway Protocol version 4 (BGP4) is the defacto routing protocol of the Internet; it is an exterior routing protocol used to exchange routing information between Autonomous Systems (AS). An autonomous system is basically a network of routers that are under the control of a single network administration. The Internet backbone is made up of different AS that exchange routing information. If you go to the following URL you can look at the major players United States Internet backbones; <http://www.nthelp.com/maps.htm> .

The people at the Internet Engineering Task Force (IETF) decided something had to be done to speed up the process of routing packets on the Internet or the whole thing was going to come to a screeching halt. They developed a new protocol that is called MPLS; it was actually based on Cisco's proprietary tag switching protocol. MPLS stands for Multi-Protocol Label Switching. It was a protocol that was developed to help speed up the process of routing packets on the Internet. The MPLS architecture is defined in RFC 3031.

In traditional routing as an IP packet travels from one router to the next, every router makes it's own decision on where the packet should go. Each router reads the packet network layer header, and then runs a routing algorithm against the destination address to determine the next hop. Every router then chooses its own next hop for the packet based on the packet's header and the routing algorithm. Routers will assign each packet into a set of "Forwarding Equivalence Classes (FECs)"(2). They will then map each FEC to a next hop. As far as the router is concerned there is no difference between packets that get mapped into the same FEC when its making a forwarding decision for each packet, different packets

which get mapped into the same FEC are indistinguishable. Every packet in the FEC will go to the next hop assigned to that FEC. As the packet moves from hop to hop across the network each router reexamines the packet network layer header and assigns it to a FEC and sends it out the corresponding interface until it reaches its destination.

With MPLS every packet only has its network layer header examined once, when it enters the MPLS network. After the initial FEC assignment a 32 bit fixed length label is inserted into the packet that contains the assigned FEC then is sent to the next hop router with the label attached. The label is of local significance only. When MPLS routers, which are called label switch routers, are provisioned they will set up a table of label to FEC mappings. Each FEC is assigned a next hop. A label distribution protocol is used to exchange label information between label switch routers that have a direct connection to each other. The protocol usually rides on top of the routing protocol in use by the use of extensions that have been developed for MPLS. As the packet goes from hop to hop across the MPLS network the network layer header no longer has to be examined by every router. Instead, the label is used to determine the next hop and which new label to use. The old label is replaced with the new label, and the packet is forwarded to its next hop. With MPLS forwarding, once a packet is assigned to a FEC, subsequent routers do no further network layer header analysis; the labels drive all forwarding decisions.

When a packet first enters into the MPLS network on an interface of Router A, known as the edge label switch router, Router A examines the network layer header determines the FEC that the packet belongs to. Then it checks the label to FEC mapping table to see which label to use. It then puts Label X into the packet and sends it out the interface that corresponds to the next hop for the assigned FEC. Router B receives the packet from Router A and reads Label X. Router B looks in his table and sees that when it receives a Label X from Router A it's new label for the packet will be Label Y. It removes Label X, adds Label Y and sends it out the interface to the next hop that corresponds to the FEC for Label Y. This continues until the packet reaches its destination. Then the label is stripped from the packet and sent out the interface that the destination is on.

This method of packet forwarding has many advantages over traditional network layer forwarding. Since a packet is assigned to a FEC when it enters the network, the edge label switch router can use any information about the packet in determining which FEC to use, even if the information is not contained in the network layer header. Packets with the same destination arriving on different ports of the router can be assigned to different FECs. Conventional forwarding, on the other hand, can only consider information that travels with the packet in the packet header. A packet that enters the network at a particular router can be labeled differently than the same packet entering the network at a different router, and as a result forwarding decisions that depend on the ingress router

can be easily made. This cannot be done with traditional forwarding, since the identity of a packet's ingress router does not travel with the packet.

The methods used to determine how a packet is assigned to a FEC can become even more complicated, without any additional effect on the rest of the routers in the MPLS network that merely forward labeled packets. There are times when you may want to have a packet follow a particular route which is chosen when the packet enters the network. This may be done as a matter of policy, or to support traffic engineering requirements. In traditional forwarding this is accomplished by using source routing, where the path of routers is contained inside the packet. In MPLS, labels can be used to represent the route, so that the identity of the explicit route need not be carried within the packet. MPLS can stack labels on the packet to set the path of the packet. Also many routers can analyze a packet's network layer header not only to choose the packet's next hop, but also to determine what precedence or class of service the packet has. They may then use this information to assign different quality of services to each packet. MPLS allows for the precedence or class of service to be fully or partially inferred from the label. This way the label actually represents the combination of a FEC and a precedence or class of service. Now that we have a basic understanding of what MPLS is let's move on to how the MPLS VPN works

With the ability to determine the path of the packet through the network, Service Providers could offer a Virtual Private Network across their backbones that could compete with Frame Relay and ATM networks. They make it work with the MPLS network.

The service provider will have a customer edge router connect to an interface on the service providers edge label switch router. Each geographically different site that will belong to the VPN will connect a customer edge router into a service provider edge label switch router. The customer edge router will be a routing peer of the service provider's edge label switch router and can exchange routing information. Individual customer sites will not be routing peers with each other and they don't even have to know about each other. Because of this the customer does not have to manage the VPN backbone. The service provider will handle all the routing that happens between the customer's sites. The customer will not have access to the service providers edge label switch router and the service provider will not have access to the customer's edge router. The customer will be responsible for maintaining his own sites' edge routers.

The service providers edge label switch router will maintain a number of different forwarding tables. An edge label switch router can have multiple customers connecting to it. It will map each customer's VPN to its own individual forwarding table. The forwarding table will only contain routes to the rest of the customer's sites that belong to the VPN for the customer. Each forwarding table for each VPN is known as a VPN Routing and Forwarding table. In this way there can be no communications between customers that do not have any VPN in common.

The edge label switch router can map different sites to the same forwarding table only if the different sites belong to the same VPN. The forwarding tables get populated with the BGP routing protocol. The customer has a MPLS VPN with Site 1, Site 2, and Site 3 connected to service provider Router 1, Router 2, Router 3 respectively. Router 1, Router 2, and Router 3 will exchange routing information for their respective sites with the use of the BGP routing protocol. The service provider edge label switch router will also contain a default forwarding table that will be populated by the service providers normal routing protocol and will not contain any MPLS VPN routes. After all this router can still be providing Internet access for other customers.

There is a possibility that different companies are using the same IP address space. They may be using a RFC 1918 private IP address space and doing network address translation for their Internet access. In fact this has become very common in today's networks. This is not a problem for MPLS VPN, because each VPN uses its own forwarding table you can have overlapping IP address space between VPNs and not have any routing problems. When the different service provider edge label switch routers exchange their routing information they maintain the separate routes for the same IP address space with the use of the BGP Multiprotocol extension. The extension makes use of a new VPN-IPv4 address. The address is 12 bytes with 8 bytes for the Route Distinguisher portion of the address and 4 bytes for the actual IP address. When multiple MPLS VPN use the same IP address space the edge label switch router will translate the address into the new unique VPN-IPv4 address. This way the routers will populate the multiple forwarding tables with different routes with the same address space for each MPLS VPN. The Route Distinguisher portion of the VPN-IPv4 address is controlled by the service provider and structured so there will be no conflict between Route Distinguishers from different service providers.

If every service provider's backbone routers had to maintain routing information for every VPN that the service provider was supporting, severe scalability problems would arise. Because of the label technology employed in the backbone the routing information only needs to be held by the edge label switch router that the VPN attaches to. This makes MPLS VPNs very scaleable, much more so than Frame Relay or ATM networks. The service provider only has to manage its own backbone and not multiple VPN backbones.

The customer has a lot of flexibility with how they want their MPLS VPN set up. They can have multiple entry points into the service provider's edge label switch router. The customer might want multiple MPLS VPN set up as Extranets between business partners and some MPLS VPN for their own geographically different offices to be part of their Intranet. Then the customer can control which network traffic goes to which site because they control their own edge router. The MPLS VPN can also be used with VLAN technology. The service provider edge label switch router can analyze the VLAN tag of the packet from the customer edge router and assign it to the correct MPLS VPN for each VLAN.

MPLS VPN security is accomplished by using a data plane and control plane approach for security. The data plane protects against a packet from within a MPLS VPN from traveling outside of its VPN boundaries and from packets from outside a MPLS VPN traveling into the boundaries of a MPLS VPN. The service provider will ensure that routers will drop packets that do not belong to MPLS VPN by examining the label of the packet. Control plane security ensures that non-trusted peers can not inject routes into the MPLS VPN. This is accomplished by the use of the MD5 authentication feature of BGP. Control plane security will also ensure that physical security of the routers is maintained to eliminate unauthorized access.

Miercom conducted an independent test of MPLS VPN security with Cisco equipment in March of 2001. The testing took the following considerations for security into account:

- Address and routing separation equivalent to layer 2 models. (3)
- A service provider core network that is not visible to the outside world. (3)
- A network that is resistant to attacks. (3)

To quote from the report:

“The test results show that MPLS-VPNs provide the previous features at or above the level of a layer 2 VPN such as Frame-Relay or ATM.” (3)

To test the requirement Miercom set up three MPLS VPNs, two VPNs used the same RFC 1918 private address space and the third used a public address space. They used Telnet and ICMP to ensure connectivity and that traffic remained inside its own VPN. They also examined all routing tables of every device, the customer edge router, the edge label switch router, and label switch router, to ensure they maintained address and routing separation. To test that packets do not leak between VPNs they used a packet injection tool to inject packets into a VPN and monitor the other VPNs for leakage. They tested for the core network being hidden by doing ICMP and Telnet tests. Even though Miercom knew the addresses of the service provider core, the tests proved they could not reach inside the service providers core network and had no access to the edge label switch router or the label switch router. The also proved that by using access lists and MD5 authentication that the service provider core network was not susceptible to DoS attacks with false routing information. They also tested for attacks against one VPN would not be propagated to other VPNs. They then tried to inject spoofed MPLS labels into the network through the edge label switch router. These packets were dropped because edge label switch routers will not accept labeled packets on interfaces from outside the MPLS network.

In summary MPLS VPN offerings are starting to get the attention of enterprise customers as an alternative to Frame Relay and ATM for connecting their

geographically different sites. It has been tested to show that it is as secure or more secure than Frame Relay and ATM. It has the ability to scale into very large networks and provide a quality of service. It allows for the customer to have control over which type of data he wants to use on his VPN. I think as time goes by that MPLS VPNs will become the VPN of choice for enterprise customers when deciding how to connect their geographically different sites.

These are many service providers offering MPLS VPNs services and also all major network equipment companies build MPLS capable routers. ComputerWorld has a recent article showing that companies are choosing MPLS VPNs over Frame Relay networks.

http://www.idg.net/english/crd_frame_533574.html . NetworkWorldFusion has had articles about MPLS offering by service providers.

<http://www.nwfusion.com/newsletters/vpn/1025vpn2.html> . ZDNet had an article telling how carriers are embracing MPLS VPN:

<http://www.zdnet.com/eweek/news/0223/23vpn.html> Major network companies advertise their ability to run MPLS on their equipment:

<http://www.nortelnetworks.com/corporate/technology/mpls/index.html>

<http://www.juniper.net/techcenter/techpapers/200012.html>

http://www.cisco.com/warp/public/cc/pd/rt/10000/prodlit/c10mp_ds.htm

http://www.riverstonenet.com/technology/mpls_ethernet.shtml

References and research material:

(1) <http://www.telstra.net/gih/papers/ipj/4-1-bgp.pdf>

(2) <http://www.ietf.org/rfc/rfc3031.txt>

(3) <http://www.mier.com/reports/cisco/MPLS-VPNs.pdf>

http://www.sans.org/infosecFAQ/encryption/implement_VPN.htm

<http://www.nthelp.com/maps.htm>

http://www.cisco.com/warp/public/cc/so/neso/vpn/vpnsp/mpls_wi.htm

<http://www.mentortech.com/learn/welcher/papers/mplsvpn.html>

<http://www.faqs.org/rfcs/rfc2917.html>

<http://www.ietf.org/internet-drafts/draft-rosen-rfc2547bis-03.txt>

<http://www.ietf.org/internet-drafts/draft-behringer-mpls-security-00.txt>

BGP4 Inter-Domain Routing in the Internet, Author: John W. Stewart III

Publisher: Addison-Wesley





Upcoming SANS Training

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

SANS London 2009	London, United Kingdom	Nov 28, 2009 - Dec 06, 2009	Live Event
SANS WhatWorks in Incident Detection Summit 2009	Washington, DC	Dec 09, 2009 - Dec 10, 2009	Live Event
SANS CDI East 2009	Washington, DC	Dec 11, 2009 - Dec 18, 2009	Live Event
SANS WhatWorks in Data Leakage Prevention and Encryption Summit 2010	New Orleans, LA	Jan 07, 2010 - Jan 12, 2010	Live Event
SANS Security East 2010	New Orleans, LA	Jan 10, 2010 - Jan 18, 2010	Live Event
SANS AppSec 2010 and WhatWorks in AppSec Summit	San Francisco, CA	Jan 29, 2010 - Feb 05, 2010	Live Event
SANS Phoenix 2010	Phoenix, AZ	Feb 14, 2010 - Feb 20, 2010	Live Event
SANS Tokyo 2010 Spring	Tokyo, Japan	Feb 15, 2010 - Feb 20, 2010	Live Event
SANS Geneva CISSP at HEG 2009 Autumn	OnlineSwitzerland	Nov 23, 2009 - Nov 28, 2009	Live Event
SANS OnDemand	Books & MP3s Only	Anytime	Self Paced