MPLS - VPN Services and Security

Ravi Sinha

MPLS - VPN Services and Security

Ravi Sinha

May 29 2003

GSEC Practical Version 1.4b

Option 1

Contents

VPN Services and Security in MPLS Architecture.

1. Abstract.

This paper will focus on the issues related to providing VPN services in a MPLS environment. The paper will discuss traditional routing and ATM in a service provider network. It will discuss the MPLS architecture and security issues and its benefits in providing a scalable connection oriented networking solution. This will provide the foundation for the discussion on providing scalable VPN services in a MPLS environment.

2. Traditional IP Routing and ATM in a Service Provider Network.

2.1 Traditional Routing.

The traditional IP network consists of a series of routers interconnected by physical media that communicate via standard routing protocols. The need for robust communication was one of the goals of IP networking in its early days. Delivery of packets with specific delays or bandwidth requirements was not an issue. Even though IP has the concept of type of service it has rarely been utilized. IP has been a very powerful networking technology. Conforming to open standards and its flexibility enables it to transfer a wide range of data types.

The growth of the Internet has put IP to the forefront of the communication world. The Internet is generally segmented into many autonomous system domains and uses typically a interior gateway protocol such as OSPF to route packets inside the AS and a exterior gateway protocol such as BGP is used to communicate between routes of separate AS. Traditional routing being connectionless has some benefits in terms of scalability and network resiliency. In a service provider network this connectionless nature also has some limitations.

Within the network OSPF establishes links using the open shortest path first algorithm. Some of the difficulties encountered are a likelihood of congestion on some of the links and a limited ability to distribute traffic over all available links. One other issue to consider is that traffic is sent across links on a hop-by-hop basis. Routing decisions are made at every node. This can create congestion on the network because routers base their forwarding decision on the destination address on the packet header and the least path cost to that destination. This causes all traffic to that destination to take the least cost path leaving the other links underutilized. With traditional routing service providers can provide only a best effort networking. All traffic is essentially treated equal and packets may be discarded on congestion. This is acceptable for application like e-mail and others with no specific requirements for latency or bandwidth.

2.2 The Need for Switching.

As service provider networks grew large the problems of the forwarding component of traditional routing and the difficulty of predicting performance in a large meshed network grew. Many service providers have enhanced their IP services by incorporating ATM and Frame-Relay, a connection oriented service into their networks.

ATM and Frame-Relay networks use a different forwarding method based on a label-swapping algorithm. This forwarding is done in hardware and yields a much better performance than traditional routing. Frame-Relay and ATM are connection-oriented technologies. A connection is established between the two end points before the traffic is sent. Since the traffic is sent along a pre-determined path the network becomes more predictable. When data is sent through an ATM network the end-to-end connection takes into account the state of the network, the latency and bandwidth requirements of the application and the preferred routes. These are critical to ATM's connection oriented performance.

Many large networks now have a switching fabric at the core. IP routing still dominates at the network edge. The need to overlay and integrate the best of traditional layer 2 and layer 3 technologies raise some interesting issues. When an IP network is overlaid on top of an ATM switched network all the routers seem to be connected at the network layer. This requires every router to have an adjacency with every other router. The adjacency must be established using ATM Virtual Circuits (VC). As the network expands, the VC's requirements grew exponentially limiting the scalability of the network. This is generally referred to as the "n-squared" problem. One other issue on the ATM network is the 5-byte overhead on every 53-byte ATM packet and the difficulty in performing very high speed segmentation and reassemble (SAR). Current IP networks are still far from meeting all the requirements of service providers and their customers.

2.3 The March towards MPLS.

As the demand for newer application like video conferencing, layer 3 VPN and VoIP increases so do the need for low latency networks. Many customers now require SLA to meet layer 2, layer 3 and layer 4 requirements to support newer applications. To meet these requirements vendors must need to add new management capabilities and predictability to their networks. To meet these commitments and requirements what is needed is a connection oriented link layer (COLL) that is aware of the end-to-end state of the network. It should be able to route traffic based on application requirements and load balance traffic across all available links. It should be able to overcome the limitations of scalability of a combined routing and switching network. That is available with newer technologies like MPLS, which is a natural evolution for networks that provide predictable IP services.

3. MPLS.

3.1 Introducing MPLS.

Multi Protocol Label Switching (MPLS) is designed to meet the mandatory characteristics of a large-scale carrier class network. It uses existing layer 3 routing protocols as well as the widely available layer 2 transport mechanisms and protocols. The IETF set up the MPLS working group in 1997 to develop a common standardized approach. The goal of the MPLS working group was to standardize protocols that used Label Swapping forwarding techniques. The use of label swapping has powerful advantages. It separates the routing problem from the forwarding problem. Routing is a global networking problem and requires the cooperation of all participating routers. Forwarding is a local problem. The router/switch decides entirely on it's own about which path to take. MPLS has one more advantage. It reintroduces the Connection State into the IP dataflows.

MPLS integrates the best of layer 2 and layer 3 technologies. The key component within a MPLS network is the label switching router (LSR), which is capable of understanding and participating in IP routing, and layer 2 switching. MPLS has provided significant new capabilities in four areas that have ensured its popularity: QoS support, Traffic Engineering, Virtual Private Network, and Multiprotocol Support.

3.2 QoS Support.

QoS is the ability to assure delivery of important data flows. Network Mangers require QoS for many reasons. Some of the requirements are

a) Guarantee a fixed amount of bandwidth for various applications.
b) Control latency.
c) Provide quantifiable SLA.
d) Ability to configure various levels of QoS for multiple customers.

In a connection less IP based internetwork it is very difficult to provide any true Qos commitments. A Differentiated Service (DS) or Integrated Services (IS) with RSVP is limited in terms of flexibility and scalability and may prove inadequate in a heavily loaded network. A connection-oriented service has powerful traffic management and QoS capabilities. MPLS imposes a connection-oriented framework and provides the foundation for reliable QoS traffic contracts.

3.3 Traffic Engineering.

Traffic Engineering is the ability to dynamically plan resource commitments on the basis of known demands, define routes dynamically and optimize network utilization. Traffic engineering seeks to control traffic flows and network resources so that predefined objectives can be met. With basic IP routing there is a primitive form of automated traffic engineering. Dynamic routing reacts in a very simple manner to congestion and does not provide a way to support QoS. When MPLS is applied, the layer 2 circuits are replaced by Label Switched Path (LSP). A set of protocols and tools are designed to measure traffic

within the LSP and provide feedback so that traffic can be adjusted. OSPF extensions for traffic engineering have been designed with MPLS LSP in mind.

3.4 VPN Support.

MPLS provides an effective mechanism for supporting VPN's. MPLS technology provides the ability to separate traffic belonging to different VPN's. In addition the establishment of LSP tunnel satisfies the formation of a virtual topology. One more advantage of MPLS as a VPN tunnel technology is MPLS traffic engineering can dedicate resources to a LSP. The security of a VPN tunnel using MPLS is equivalent to that provided by ATM/Frame-Relay PVC.

3.5 Multiprotocol Support.

MPLS can be used on many networking technologies. MPLS enabled routers can coexist with ordinary routers. MPLS is designed to work in ATM, Frame-Relay networks. MPLS enabled ATM, Frame-Relay switches can also work with ordinary switches.

3.6 MPLS Operation.

The MPLS network consists of a set of nodes capable of switching and routing on the basis of label appended to each packet. A MPLS domain consists of a contiguous or connected set of MPLS enabled nodes. These nodes are called Label Switched Router (LSR). The labels define the flow of packets between the two endpoints. A specific path through the network of LSRs for each distinct flow called a Forwarding Equivalence Class (FEC) is defined. MPLS is a connection-oriented technology. With each FEC is associated a traffic characterization that defines a QoS requirements for that flow. Because the LSR forwards the packet based on its label value this ensures that the forwarding process is simpler than with an IP router. Figure 1 depicts the operation of MPLS enabled router. The following are the key elements.

Figure 1: MPLS Operation

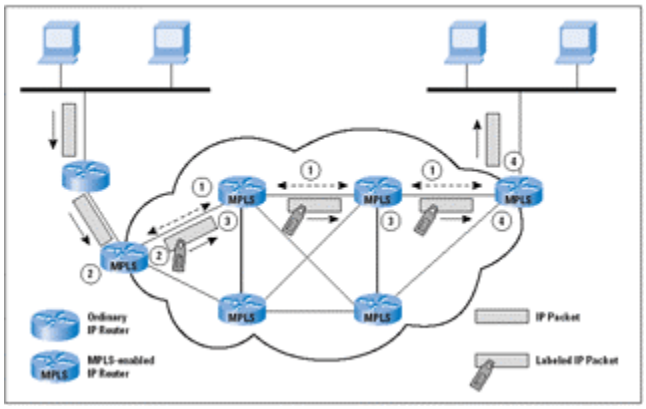

Figure 1 MPLS Operation [1]

Prior to the delivery of packets a path through the network called a Label Switched Path (LSP) must be defined for the packets in the given FEC. The QoS

parameters must be established. The QoS parameters determine how many resources to commit to that path and what is the queuing and discard policy at each LSR for the FEC. To accomplish the above an interior gateway protocol like OSPF is required for reachability and routing information. Labels are assigned to each packet in the FEC. These labels have only local significance. A protocol like Label Distribution Protocol (LDP) or an enhanced version of RSVP is used to determine routes and establish labels values. These can also be setup manually by an operator.

A packet enters the MPLS domain through an ingress edge LSR. It is here that the packet is processed to determine its need for network layer services, defining its QoS. The LSR assigns this to a particular FEC and therefore a LSP, and then forwards the packet.

Each LSR that receives a labeled packet removes the incoming label and attaches the appropriate outgoing label to the packet and forwards the packet to the next LSR in the LSP.

The egress edge LSR strips the label reads the IP packet header and forwards the packet on to its final destination.

One of the most important features of MPLS is label stacking. A labeled packet may carry many labels organized as a last-in-first-out. Processing is based on the top label. At any LSR labels can be added to the stack or removed from the stack. This allows aggregation of LSPs into a single LSP for the portion of the route creating a tunnel.

The FEC for a packet can be determined by one or more parameters, such as source or destination IP address, source or destination ports, IP protocol ID, differentiated service codepoints or IPv6 flow labels. A per-hop behavior (PHB) can be defined at a LSR for a FEC. The PHB defines the queuing priority of the packets for this FEC and the discard policy. Packets send to the same endpoints may belong to different FEC and will be labeled differently and experience different PHB at each LSR and may follow different paths through the network. The essence of MPLS functionality is that traffic is grouped into FECs. The traffic in a FEC transits a MPLS domain along a LSP. Individual packets in a FEC are uniquely identified as being a part of a given FEC by means of a locally significant label.

Route selection refers to the selection of a LSP for a particular FEC. MPLS supports hop-by-hop routing and explicit routing. With hop-by-hop routing each LSR independently chooses the next hop for each FEC. This option makes use of an ordinary routing protocol such as OSPF. This has some advantages but because of the limited use of performance metrics, hop-by-hop routing does not readily support traffic engineering or policy related to QoS and security. With explicit routing a single LSR specifies some or all the LSRs in the LSP for a FEC.

Explicit routing provide all the benefits of MPLS, including the ability to do traffic engineering and policy routing. Dynamic explicit routing provides the best scope for traffic engineering. In this mode the LSR setting up the LSP would need information about the topology as well as QoS related information for the MPLS domain. The enhanced version of OSPF for MPLS has some newer metrics that would be useful in constraint based routing including maximum link data rates, current capacity reservation, packet loss rate and link propagation delay. Route selection consists of defining a LSP for a FEC. A separate function is the actual setting up of the LSP and for this each LSR on the LSP must

a) Assign a label to the LSP to be used to recognize incoming packets that belong to the corresponding FEC.
b) Inform potential upstream nodes of the label assigned by this LSR to this FEC.
c) Learn the next hop for this LSP and the label that the down stream node has assigned to this FEC.

4. VPN in a MPLS environment.

4.1 ABC's of VPN.

The basic idea of Virtual Private Network (VPN) is quite simple. A Company may have many offices at different locations. Each of these locations has its own local network. Internetworking these separate networks over a shared network creates a VPN. VPN's today are built over tunnels and security. The tunnel transports traffic between locations and the encryption provides the confidentiality. Three types of VPN are considered here. They are the access VPN, the Intranet VPN and the Extranet VPN. Each type of VPN meets different business needs.

Access VPN provides remote access connection for telecommuters and small offices through a dial or broadband access technology. VPN users access a local service provider POP and PPP or IPSec tunnels are routed to the corporate gateway through the provider network or the Internet. Access VPN modes can either be compulsory or voluntary. Voluntary mode gives the user the ability to setup the VPN at their discretion. The client software initiates the tunnel and the encryption so that the encapsulated traffic is invisible to the service provider network. In the compulsory mode the tunnel is imposed in the POP so it is invisible to the user.

Intranet VPN's extended the resources and applications to the branch offices. Extranet VPN's extend resources and applications to the companies suppliers and partners. These are typically full time connections. These links carry critical data and are typically encrypted.

The major building blocks of VPN's consist of Tunneling, Security, QoS, Management and Provisioning.

The purpose of tunneling is to create a pathway across the shared network and to encapsulate traffic with the new packet headers for delivery. Tunnels do not provide confidentiality, which is provided by encryption. Tunnels can be either layer 2 or Layer 3. Layer 2 tunnels are based on PPP, which include PPTP, L2TP and PPoE. Layer 3 tunnels include Generic Route Encapsulation, which is a dependable way to encapsulate non-IP traffic for transport over IP networks. IPSec is an IETF standard which incorporate both tunneling and encryption.

VPN security requires user authorization, authentication and data encryption. The initial authentication is used to verify user/router and permit certain actions and deny others. Often VPN tunneling provides sufficient protection, but certain traffic may require encryption. IPSec is an IETF standard that provides 56/128/256-bit encryption.

Enterprise manager or service providers are usually responsible for the management of the VPN services. They enforce security and QoS throughout the network. The also manage the authentication, authorization and accounting of the systems.

4.2 VPN and MPLS.

Enterprises building Virtual Private Networks (VPN) can treat the public Internet as their own private WAN for connecting remote offices. Early VPN implementations using PVC and tunneling techniques were successful. However as the connectivity requirements grow wider, scalability and management challenges are presenting themselves. One VPN solution creates a mesh of point-to-point tunnels among routers at the edge of the service provider network. The provider router must maintain complete routing information about all customer networks. Every edge router must exchange information with every other router. The volume of routing information increases as new sites are added. This creates scalability problems. One other mode of VPN services is building a layer 2 point-to-point network using a Frame-Relay or ATM PVC. The need to establish and manage a full mesh of virtual circuits is difficult because each new connection must add to a growing population of sites. QoS is another VPN challenge. Managing multiple service levels across multiple PVC's can be an administrative nightmare. In the tunneling scenario IPSec or GRE do not support QoS on their own. Fortunately MPLS technology has arrived that enables the construction of full mesh VPN that supports multiple service levels and scales infinitely. MPLS VPN technologies incorporate all the capabilities of existing VPN networks while also delivering QoS capabilities. While both IPSec and MPLS can be used to build VPNs both were designed for completely different tasks.

IPSec is a highly secure infrastructure for transporting sensitive information over the public Internet. It provides data privacy through a flexible suite of encryption and tunneling mechanisms that protect data payload. In an

IPSec based VPN no modification to the application is required. IPSec functions at the network layer. Untrusted public networks are the most likely candidates for IPSec tunneling and encryption technology.

MPLS is best deployed at the core of the service provider network. While IPSec functions best at the outer regions of the network. At the core of the network is where QoS, Traffic engineering and bandwidth utilization can be fully controlled, enabling service providers SLA's. MPLS can also run over IP, ATM, and Frame-Relay networks. It also provides for load balancing. MPLS like IPSec also provide for end-to-end security for transmission. It separates traffic using Route Distinguishers (RD). IPSec also requires site-to-site peering to operate while VPN's built using MPLS do not require pre-defined relationships. Packets are labeled for specific VPNs within the network and only those ports that are part of the specified VPN receive the traffic. For mobile users and telecommuters MPLS is not an option. MPLS is a network-based solution and does not go all the way out to the computing endpoints. MPLS stops at the edge of the service provider network. IPSec is the only practical application to enable secure remote access. While service providers can deploy either an IPSec or MPLS based VPN architecture to deliver new value added services a greater benefit can be realized if they converge. Service providers may choose IPSec for traffic that needs strong authentication and confidentiality and choose MPLS for its broader connectivity, QoS support and traffic engineering capabilities.

4.3 Security.

As MPLS is becoming a more widespread technology for providing VPN services MPLS architecture security is of increasing concern to both service providers and VPN customers. This section will look at the security that MPLS provides. It assumes that the MPLS core is provided in a secure manner. It does not address basic security concerns such as the misconfiguration of the core, internal attacks and securing the network elements against unauthorized access. The four issues discussed are security requirements typical in MPLS architectures.

a) Address space and routing separation.

From a security perspective, the basic requirement is to avoid the situation in which packets destined to a host a.b.c.d within a VPN reach a host in another VPN or the core. Between any two non-intersecting VPN in the MPLS it is assumed that the address space between different VPNs is entirely independent. From a routing perspective this means that any VPN must be able to use the same address space as any other VPN or the MPLS core and that the routing between any two VPNs or the VPN and the MPLS core must be independent. MPLS adds a 64-bit route distinguisher (RD) to each IPv4 route making VPN unique addressing also unique to the MPLS core. There is one exception, which is the IP address of the provider edge (PE) the customer edge (CE) router are peering with. This address must be unique from the perspective of the CE router. Every PE router maintains a Virtual Routing and Forwarding (VRF) instance for

each connected VPN. This ensures routing separation between the VPNs. Between the core to the other PE router the separation is maintained by unique VPN identifiers in multiprotocol BGP. This BGP information is not redistributed to the core. It is redistributed to the other PE routers. It is not possible to intrude into other VPNs through the MPLS cloud unless it has been specifically configured.

b) Hiding of the MPLS core structure.
The internal structure of the MPLS core (provider edge (PE) and provider (P) elements) should not be visible to outside networks. A denial-of-service-attack against the core is much easier to carry out if the attacker knows the addresses. MPLS does not reveal any unnecessary information to the outside not even to the VPN customer. The only information needed in the case of the routing protocol between the PE and the CE is the address of the PE router. If this is not desired static routing can be configured. With this measure the MPLS core can be completely hidden. Customer VPNs will have to advertise their network routes (not hosts) at a minimum to the MPLS core. In a VPN service with a shared Internet access a service provider will announce routes to customers wishing to use the Internet or upstream peer providers. This can be done via Network Address Translation (NAT) function to further obscure the addressing information of customer's networks.

c) Resistance to attacks.
There are two basic types of attack: A DoS attack where resources become unavailable. They are easier to execute. The only way to be sure that the network is invincible to this kind of attack is to make sure that the machines are not reachable, by packet filtering and address hiding. For attacks that give unauthorized access to resources, there are two basic ways to protect the network. First is to harden the protocols that could be abused and second to make the network as inaccessible as possible. It is not possible to intrude from one VPN to another or the core. It is theoretically possible to exploit the routing protocol to execute a DoS attack against a PE router. PE routers must be extremely well secured especially their interface to the CE router. ACLs and MD5 authentication must be used on all PE/CE peering.

d) Label spoofing.
In an MPLS network packets are forwarded on the basis of labels and not on IP destination address. It is possible to insert wrong labels into a MPLS network from the outside i.e. from the VPN CE or from the Internet. The CE router is unaware of the MPLS core and thinks it is sending an IP packet to a simple router. It is in the PE router that a label is chosen and added to the packet. For security reasons the PE router should never accept a packet with a label from the CE router. There is a possibility of spoofing the IP address of the packet being sent to the MPLS core, but

because of strict address separation the spoofed packet can only harm the VPN that the spoofed packet originated from.

From a customers perspective it is impossible to control the whole network. If the MPLS core is not properly configured, the VPNs will be exposed to some form of attacks. IPSec offers additional security over a MPLS network. IPSec can be run on the CE routers or devices further away from the core. IPSec should be used if one of the following requirements exist.

1) Encryption of parts or all traffic over MPLS core.
2) Authentication of the endpoints.
3) Integrity of the traffic.
4) Replay detection.

MPLS and IPSec together can provide a very high level of security for VPNs.

5. Conclusions.

MPLS provides benefits that service providers need urgently in their networks, such as predictability, scalability and manageability. MPLS will require modifications to existing equipment; it will not require an extensive overall. MPLS defines an evolutionary networking methodology that combines the principles of layer 2 and layer 3 technologies while preserving the service provider investment in routing technology at the edge and the switching technology at the core. MPLS improves the scalability of routing and forwarding and provides traffic engineering capabilities for better network provisioning. The MPLS infrastructure provides at least the same level of security as a comparable ATM or Frame-Relay service. It would be an excellent choice for providing VPN services.

6. References.

[1] Stallings, William. “MPLS.” The Internet Protocol Journal. September 2001(2001): 2-14.

[2] Meredith, Gail. “Going Connectionless.” Packet Magazine. Fourth Quarter 2000(2000): 78-82.

[3] Meredith, Gail. “Express Delivery.” Packet Magazine. Fourth Quarter 2000(2000): 83-87.

[4] Editorial Staff. “Inner Workings of MPLS.” Packet Magazine. First Quarter 2000(2000): 83-87.

[5] Editorial Staff. “MPLS and IPSec .” Packet Magazine. First Quarter 2001(2001): 37-43.

[6] Dickson, Kevin. “ABCs of VPNs.” Packet Magazine. Fourth Quarter 1999(1999): 50-53,94.

[7] Telstra, GeoffHuston. “Quality of Service-Fact or Fiction.” The Internet Protocol Journal. March 2000(2000): 27-34.

[8] Cisco Systems. “Internet Connectivity Options.” White Paper. URL: http://www.cisco.com/en/US/tech/tk436/tk428/technologies_white_paper09186a00801281f1.shtml 04/29/2003).

[9] NetPlane Systems. “Layer 3 Switching Using MPLS.” White Paper. URL: http://www.netplane.com/pdf/Layer%20switching%20whitepaper%20July27_FINAL.pdf (05/12/2003).

[10] Cisco Systems. “Security of the MPLS Architecture.” White Paper. URL: http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/mxinf_ds.htm (04/29/2003).

[11] Vishwanathan, Arun., et al. “Evolution of Multiprotocol Label Switching.” URL: http://www.comsoc.org/ci/private/1998/may/Feldman.html (05/08/2003).

[12] Rosen, Eric., et al. “Multiprotocol Label Switching Architecture.” RFC 3031 IETF. January 2001 URL: http://www.ietf.org/rfc/rfc3031.txt (04/16/2003).

[13] Rosen, Eric., et al. “MPLS Label Stack Encoding .” RFC 3032 IETF. January 2001. URL: http://www.ietf.org/rfc/rfc3032.txt (04/16/2003).

Last Updated: September 25th, 2020

Upcoming SANS Training

Click here to view a list of all SANS Courses

SANS October Singapore 2020	**Singapore, SG**	**Oct 12, 2020 - Oct 24, 2020**	**Live Event**
SANS Community CTF	**,**	**Oct 15, 2020 - Oct 16, 2020**	**Self Paced**
SANS SEC504 Rennes 2020 (In French)	**Rennes, FR**	**Oct 19, 2020 - Oct 24, 2020**	**Live Event**
SANS SEC560 Lille 2020 (In French)	**Lille, FR**	**Oct 26, 2020 - Oct 31, 2020**	**Live Event**
SANS Tel Aviv November 2020	**Tel Aviv, IL**	**Nov 01, 2020 - Nov 06, 2020**	**Live Event**
SANS Sydney 2020	**Sydney, AU**	**Nov 02, 2020 - Nov 14, 2020**	**Live Event**
SANS Secure Thailand	**Bangkok, TH**	**Nov 09, 2020 - Nov 14, 2020**	**Live Event**
APAC ICS Summit & Training 2020	**Singapore, SG**	**Nov 13, 2020 - Nov 21, 2020**	**Live Event**
SANS FOR508 Rome 2020 (in Italian)	**Rome, IT**	**Nov 16, 2020 - Nov 21, 2020**	**Live Event**
SANS Community CTF	**,**	**Nov 19, 2020 - Nov 20, 2020**	**Self Paced**
SANS Local: Oslo November 2020	**Oslo, NO**	**Nov 23, 2020 - Nov 28, 2020**	**Live Event**
SANS Wellington 2020	**Wellington, NZ**	**Nov 30, 2020 - Dec 12, 2020**	**Live Event**
SANS OnDemand	**OnlineUS**	**Anytime**	**Self Paced**
SANS SelfStudy	**Books & MP3s OnlyUS**	**Anytime**	**Self Paced**