



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Multiprotocol Label Switching Virtual Private Networks and the enterprise – Do they fit in the security model?

Abstract

Multiprotocol label switching virtual private networks have gained press as a new service provider method to provide a secure path in the public Internet space. The question arises if this technology is the latest marketing ploy of the service providers or if this is a valid option for the enterprise within the security framework? If this is a valid option for the enterprise security framework is it a WAN technology only or does it have a place within the MAN or LAN environments? The first part in looking at the issue is to first understand if multiprotocol label switching virtual private network technology is a virtual private network and to determine how it works. A determination of what security is provided or not provided; also looking at what flaws it may have as a virtual private network. If the technology is found to be acceptable as a virtual private network, a look at how different enterprise customers can take advantage of this technology.

Virtual Private Network definition

Virtual private network has come to take on several meanings forcing an understanding of the way the term is used. The term is often associated with IPSEC, Layer 2 Tunneling Protocol, and Point-to-Point Tunneling Protocol in which a secure channel is achieved of the public or non-private network via encryption mechanisms. In this same definition, hardware encryption also became a valid method of setting up Virtual private networks in a non-secure environment. The definition was later expanded to include any mechanism, which brings the ability to provide segregated channel for the data being transferred in the public or non-private network.

The acid test of this as a virtual private network for this paper comes when looking at the following [1]:

Can two customers use the same IP space?

Can a customer use the same IP space as the service provider?

Is routing between Virtual private networks independent?
Is routing between the Core and Virtual private networks independent?

Multiprotocol label Switching Network Overview

When the IETF introduced multiprotocol label switching it was intended to provide quality of service and traffic engineering, not as a choice in security. As several service providers began to roll out multiprotocol label switching backbones, the realization came that it may have the ability to provide a security option.

To get to a multiprotocol label switching virtual private network, the first step is to set up a multiprotocol label switching network. The service provider sets up a standard routing infrastructure. This includes an IGP to provide transport and control for the service provider and an iBGP upon which the customer traffic is to be routed. After establishing the network the multiprotocol label switching overlay is added.

The multiprotocol label switching network consists of setting up all routes in the multiprotocol label switching mesh to be Label Switch Routers. This is simply turning on multiprotocol label switching and providing for the traffic to use either RSVP or LDP as the mechanism to set up the Label Switch Path through the network over the underlying IGP network. Label Switch Paths act as unidirectional paths between two end points. The path does not use traditional routing to make decisions on how to transmit traffic. The traffic, when entering the multiprotocol label switching mesh, has a 64-bit route distinguisher placed onto it in the same method that an 802.1Q network places a 4-byte tag onto a packet over a tagged port. [2] At each Label Switch Router, the Label Switch Path route distinguisher is read, this acts as an index to a route and traffic is forwarded to the next router in the Label Switch Path path. This is done instead of longest path match of a traditional routed network. This allows traffic to be forwarded independent of IP address allowing the customer to choose to use public IP space or RFC 1918 space. In longest path match, the IP address is read and matched against the routing table.

Multiprotocol label Switching Virtual Private Network

The multiprotocol label switching virtual private network has two major options: Layer 3 virtual private network or Layer 2 virtual private network. Each uses the

multiprotocol label switching mesh that was created above, but the method of getting traffic to and from the service provider network to the customer varies greatly. In both cases it assumes the service provider has provided a secure infrastructure using industry best practices.

Multiprotocol label Switching Layer 3 Virtual Private Network Overview

The layer 3 version typically uses 2547 as the mechanism to provide the virtual private network service and can also be referred to as multiprotocol label switching BGP/virtual private network. While other IETF drafts are currently underway, 2547 is still the most predominant and will be the focus of this discussion.

The layer 3 version requires the use of virtual routing and forwarding instances. At each point in which a customer connects to a service provider a virtual routing and forwarding instance must be maintained on the service provider router for the customer. The customer views the virtual routing instance as an extension of their network, which is controlled by the service provider. This causes the service providers to either manage or heavily dictate the layer 3 interaction with the customer. Almost always the service providers require that it is a multiprotocol BGP connection through their network, as to not have to manage the IGP interaction with the customers (most vendors implemented 2547 allow IGP connections and static routes, but most service providers avoid this type of connection). Customers are still allow to uplink in whatever manner the service provider permits on any Internet connection; static route, BGP, or IGP. The multiprotocol BGP extensions in the service provider network keep the traffic being routed separate for each virtual routing and forwarding instances. [7]

The customer still connects to the service provider router via a standard mechanism such as T(X), OC(X), or other support network type at each location. As the service provider must take part in the BGP routing of the customer, the first IP address of the service provide virtual routing and forwarding instance is known at each connected site. This is the only portion of the service provider network that is known to the customer. The routers in the service provider core are hidden, as there is no other layer 3 interaction with the customer. This allows the customer to choose all of the IP addresses for their network except the connection to and from the service providers.

The next issue is, if the customer wishes to have separation within the multiprotocol BGP mesh created by the service provider for the customer. If the customer has 2 sites that have both Human Resources and Engineering connected and they would like to segregate the traffic across the service provider backbone, two virtual route and forwarding instances would be needed. The

service provider would just treat this as two customers and connect to the sites via two connections, 2 DLCI's, 2 VCI/VPIs, or 2 VLAN depending upon if frame relay, ATM, or 802.1Q VLAN were available. The customer would have two separate multiprotocol BGP instances on the service provider router and would be required to use internal security, routing, and switching policies and procedures to separate the 2 virtual private networks.

In addition to the address shown for the multiprotocol label switching layer 3 virtual private network, the customer may also chose to have a traditional Internet connection provided by the ISP. This has the same potential issues as any Internet connection plus the customer having to be savvy enough to not connect the non-virtual private network traffic into the virtual private network without going through proper security measures. As the customer may run BGP with both the virtual private network and Internet link, proper safeguards that may include; correct filters, firewalls, and/or proxies should be appropriately established. The proper security for this type of connection is beyond the scope of this paper.

While the reader may note that a virtual routing and forwarding instance for each customer may get cumbersome and draw routing resources for the services provider, that discussion will not be covered in this paper.

The reader should also note that only IP connections are allowed via the layer 3 method as multiprotocol BGP is used to connect to the provider.

Multiprotocol label Switching Layer 2 Virtual Private Network Overview

In the Layer 2 option the service provider looks like a layer two connection to the customer and the customer provides the layer three management of his or her own network. The service provider has no interaction with the layer three routing management as they did in the RFC 2547 multiprotocol label switching Layer 3 VLAN method.

The layer two method makes use of the same multiprotocol label switching mesh that the layer 3 method made use of and transfers traffic through the service provider core identically as layer 3, but uses a different mechanism to get the traffic between the customer and the service provider edge.

The service provider again drops off a T(X), OC(X), or other supported media type. This circuit could be provisioned as direct T(X), OC(X), untagged Ethernet, Frame Relay DLCI, ATM VPI/VCI, or tagged Ethernet or other supported media type. The two end points are not required to be the same media type, but some service providers will require it to simplify their own network and not have to translate a VLAN to a DLCI mapping. The customer sees a simple layer 2

connection while the service provider uses a layer three network to accomplish the connection. In this case, the customer is able to run any protocol across the link that they would run on their local LAN. There are exceptions for the service provider allowing large enough MTU and other assorted problems when going across the network. [5] Again, this is beyond the scope of this paper.

If a third site is added, 1 or 2 additional label switch path(s) will also need to be added to allow all sites to connect to each other. This is the same as ATM or Frame relay where the customer must decide if they would rather run a hub and spoke connection or a fully meshed network. For this example, a fully meshed network will be discussed. If site 1 and site 2 are connected via label switch path A, then site 3 can connect to site 1 via label switch path B and to site 2 via label switch path C.

This would require either multiple circuits to each of the site or using Frame Relay DLCIs, ATM VPI/VCI, or VLANs to allow multiple connections across a single connection. Assuming VLANs are the method of delivery, site 1 would have a VLAN X that the service provider would map to label switch path A to connect it to site 2 and a second VLAN Y that the service provider would map to label switch path B to connect it to site 3. The customer would decide if they wished to provide routing between the sites or if they wished to tie the VLANs together and make a wide area multi-site layer 2 connection. In most cases routing would be used, but in limited cases some networks may wish to connect multiple sites as one layer two connection. This is significantly more complicated than just connecting another site to the multiprotocol BGP mesh as done in the multiprotocol label switching Layer 3 connection. This is being addressed in several different multiprotocol label switching drafts, most notably, the Martini draft. [10]

If separation is required, as in the example of a Human Resources and an Engineering team split between two sites again. A second connection or user two DLCIs, VPI/VCI, VLANs or other support mechanism would be added and mapped to a different label switch path between the two sites. The customer would again have to keep appropriate traffic segregated between the two virtual private networks locally via proper internal security, routing, and switching policies and procedures.

The final type of connection is a traditional Internet connection. This has the same potential issues as any Internet connection plus the customer having to be savvy enough to not connect the non-virtual private network traffic into the virtual private network without going through proper security measures. As the customer may run DLCIs, VPI/VCI, or VLANs into the same router or switch proper safe guards and correct filters, firewalls, and/or proxies should be appropriately established to keep the traffic separate. The proper security for this type of connection is beyond the scope of this paper.

The question remains: Are multiprotocol label switching virtual private networks really a virtual private network?

The acid test for virtual private network for this paper required looking at the following questions:

Can two customers use the same IP space?

Can a customer use the same IP space as the service provider?

Is routing between Virtual private networks independent?

Is routing between the Core and virtual private network independent?

Two customers can use the same IP address space in multiprotocol label switching Layer 2 virtual private networks. In Layer 3 virtual private networks it is a qualified yes, as two customers cannot use the same space to connect to the provider.

Customer can use the same IP space as the provider in multiprotocol label switching Layer 2 virtual private networks. This again is a qualified yes as for multiprotocol label switching Layer 3 virtual private networks due to the provider edge to customer issue.

Routing between both Layer 2 and Layer 3 virtual private network is independent. In the Layer 3 case separate virtual routing and forwarding instances are created for each virtual private network, giving each virtual private network its own routing table. In the layer 2, case, each customer controls their own routing with the service provider only giving a virtual path between sites.

Routing between the core and the Virtual private networks is independent in both cases. In the layer case 3, the core is routed via the standard routing table (or separate virtual routing and forward instance is the service provider wishes). In the layer 2 case, the core is routed by the service provides and the customer does the routing for the virtual private networks.

The multiprotocol label switching Layer 2 connection is a virtual private network by the definition used. The multiprotocol label switching Layer 3 connection is a qualified virtual private network due to the service provider and customer sharing IP space on the connection between them.

What has not been provided or what else has been provided by multiprotocol label switching virtual private networks?

At this point, the only thing that it has been shown to provide is a segregated path between termination points. [3]

IPSEC, L2TP, PPTP virtual private network provide a means for verification of Integrity, Encryption, and Origin Authentication of data, these are not provided by multiprotocol label switching Virtual private networks. [12] In addition it does not make any claim of security on the provider or customer networks.

Multiprotocol label switching Virtual private networks also provide the customer the responsibility to verify or assume the service provider's network is secure. If this is in doubt, multiprotocol label switching does not provide any recourse to make the path more secure than what the provider has delivered. Given that it does not provide the integrity, data encryption and origin authentication, the two major issue of note are label spoofing for both layer 2 and 3 virtual private networks and exposure of the IP addresses for customer to provider connection.

Label spoofing is a concern at any point the label switch paths are accepted to the service provider backbone. As the service provider actually applies the label, it should not accept labels from any customer interface eliminating this concern. In a case where 2 service providers allow a label switch path to run between them, both must only allow label switch paths on selected connect points (proper DLCIs, VPI/VCIs, VLANs, or other accepted media) and not from any other source.

In the layer 3 case, IP addresses have to be exposed that connect the service provider and customer. This makes the addresses susceptible to standard IP attacks such as DoS. The threat of these attacks can be minimized using RFC 1918 address, having proper ACLs in place, using authentication mechanisms such as MD5 if the connecting protocol allows, limiting the number of routes a virtual routing and forwarding instance can have, and configuring BGP mechanisms such as route damping. [7]

How do different Enterprise customers use multiprotocol label switching Virtual private networks?

When looking at multiprotocol label switching Virtual private networks, the Enterprise customer must first look at what type of customer they are. Do they have WAN or MAN connections over the public or any non-private networks? Are they a small, medium, or large Enterprise? Do they function as a combination enterprise backbone and service provider? How is there customer base and services distributes? What level of security do they need/want?

If an Enterprise does have customers over a non-private WAN or MAN, they need to look at the data being transferred. If the data or path does not require any security level, this is not a place for any virtual private network. If the data or path does require security what level is needed? Is it full encryption or only traffic segregation to a desired termination point? In either case a multiprotocol label switching virtual private network may have a place. If the data or path could survive an ATM or Frame relay virtual private network then multiprotocol label switching virtual private network is a valid substitute. If more than a segregated path is needed, a multiprotocol label switching virtual private network is not the right solution. An IPSEC style virtual private network would be needed.

The size of the enterprise brings in all of the scaling issues. One of the major advantages of using a multiprotocol label switching virtual private network is the scaling and flexibility of the virtual private network. If a new router is added it is configured to run the proper multiprotocol label switching configuration and is now ready to add customers (depending on the way Label Switch Paths are configured, this may not be trivial). New Label Switch Paths can be added by simply configuring the new Label Switch Path name on the terminating routers at the provider edge. If the endpoint changes, the configuration of the moving end is removed and placed on the new termination point while the fixed end has the terminating information updated if any changes are needed (this is again dependent on base configuration of the multiprotocol label switching network).

If the enterprise acts as a combination enterprise backbone and service provider, they will face the standard enterprise issues plus many of the same issues service providers have encountered. This typically arises from mergers or spin-offs of companies. When a company acquires another company, two enterprises now must act as a service provider to the other. Just as service providers are using multiprotocol label switching Virtual private networks as method to connect two disparate sites for the customers who needs a method of providing access between two LANs, two enterprises can gain the same benefit. The same issues can come during the split of an enterprise in two units that need to have their traffic separated, when it used to be allowed on the same backbone without segregation. In this case two multiprotocol label switching virtual private networks can be run to use the same infrastructure to provide to virtual backbones.

Where does it fit in the security framework?

The above discussion shows multiprotocol label switching Virtual private networks simply provides a segregated virtual private network path to transfer

data over. This allows the Technology to be used in any location where current ATM or Frame relay style Virtual private networks would normally fit. This can be over the WAN, MAN, or LAN with the same caveats provided about multiprotocol label switching Virtual private networks that are provided about Frame Relay or ATM Virtual private networks.

Conclusion

Several services providers are touting multiprotocol label switching virtual private network to many enterprise customers as security method to be used to connect sights across the wide area. While Layer 2 multiprotocol label switching virtual private networks meet all of the requirements to be a virtual private network, the Layer 3 multiprotocol label switching virtual private networks meets it in a qualified manner. Both methods allow for segregated traffic to be passed across a non-private network meeting the necessary requirements where ATM or frame relay Virtual private networks would have traditionally been used. This is accomplished as multiprotocol label switching virtual private network have the same issues as frame relay and ATM virtual private networks as they do not provide a means for verification of integrity, encryption, and origin authentication of data. This does not preclude using multiprotocol label switching virtual private networks with protocols such as IPSEC that do provide the missing mechanisms. multiprotocol label switching virtual private networks also have several current drafts that may include the mechanisms to provide for more than just segregation of traffic in the near future. In the meantime, multiprotocol label switching virtual private networks literally have become another choice in addition to frame relay and ATM virtual private networks for the enterprise customer and multiprotocol label switching virtual private networks allow the same additional security layer that frame relay and ATM virtual private networks currently provide.

List of References

- [1] Tolly, Kevin “Will the real VPN please stand up?” March 9, 2003
http://www.tolly.com/NEWS/KT_NWW/20020204VPNs.asp (taken offline)
- [2] “MPLS VPNs: The real deal” Network Computing workshop 2001
<http://www.nc-india.com/workshop/stories37075.html> (taken offline)
- [3] Mier, Edwin “Tester’s Choice: MPLS takes on security role” May 5, 2001
<http://www.nwfusion.com/research/2001/0521feat2.html>
- [4] “A Compromise Between IPSEC and Multiprotocol Label Switching Virtual Private Networks” Cisco White Paper 2000
- [5] Petrosky, Mary “The promised LAN” April 8, 2002,
www.nwfusion.com/research/2002/0408feat2.html
- [6] “Security of the MPLS Architecture” Cisco White Paper September 20, 2002
http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/mxinf_ds.htm
- [7] Behringer, Michael “Analysis of the Security of the MPLS Architecture” October, 2002
<http://www.ietf.org/internet-drafts/draft-behringer-mpls-security-03.txt>
- [8] Greenfield, David “MPLS in Brief” February 06, 2002
<http://www.commweb.com/article/NMG20020206S0016>
- [9] Senevirathne, Tissa “Secure MPLS – Encryption and Authentication of MPLS Payloads” July, 2002
<http://www.ietf.org/internet-drafts/draft-tsenevir-smpls-02.txt>
- [10] MPLS Resource Center “MPLS VPNs” 2002
<http://www.mplsvpn.net/support/mplsvpn.html>
- [11] Young, Clifford; Merali, Mehmood; and James, Bradley “Looking to MPLS, LDAP, and policy for better security” September 2000
<http://www.serverworldmagazine.com/sunserver/2000/09/mpls.shtml>
- [12] “IPSEC and MPLS: Which VPN is right for you?” Network Magazine India 2001
- [13] Wu, Tim and Walden, Andy “MPLS VPNs: Layer 2 or Layer? Understanding the choice” March 2003
http://www.riverstonenet.com/technology/mpls_vpn.shtml