Miercom

The leading edge in networking information

White Paper

Cisco MPLS based VPNs: Equivalent to the security of Frame Relay and ATM

March 30, 2001

Abstract: The purpose of this white paper is to present discussion and findings that conclude that Cisco MPLS-based VPNs are as secure as their layer 2 counterparts such as Frame-Relay and ATM. This document details a series of tests were carried out on a Cisco router test bed validating that MPLS based VPNs (MPLS-VPN) provide the same security as Frame-Relay or ATM.

ATM and Frame-Relay have a reputation in the industry as being secure foundations for enterprise connectivity. Essential items that make ATM and Frame-Relay a secure network were considered and tested on an MPLS-VPN.

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

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.

As described in greater detail through out this paper a test bed of 22 Cisco routers was used, including- two Cisco 12000 series Internet routers, two 7505s, four 7206 VXRs, five 3640s, five 2611s, and four 1750s running IOS version (12.0) and (12.1) to implement the necessary functions to provide a stable and secure MPLS core.

410 Hightstown Road
Princeton Junction, NJ 08550
609-490-0200; fax 609-490-0610
info@mier.com www.mier.com

Introduction

Today, business customers accept the level of security that Frame-Relay and ATM offer as layer 2 VPNs, however they might have concerns about the level of security that an MPLS based VPN offers. The goal of this paper is to answer those questions and provide proof with test results that an MPLS based VPN solution is as secure as a comparable layer 2 VPN. A basic understanding of MPLS and MPLS-VPN principles is assumed for this paper.

Virtual Private Networks

A virtual private network (VPN) can be defined loosely as a network in which customer connectivity amongst multiple sites is deployed on a shared infrastructure, with the same access or security policies as a private network. As a alternative solution to expensive leased-lines or circuit-switched infrastructures, the growth rate of virtual private networks in the business world has been expanding.

Currently most of these VPN infrastructures are built on Frame-Relay or ATM networks connecting customer sites via Virtual Circuits (VCs.) The hub and spoke topologies, common of VPNs, today are being replaced by an any-to-any mesh that increases the complexity and number of VCs needed. This increase in VCs and the complexity that goes with them is driving the need for a more scalable VPN solution.

VPN topology today

Today VPNs are implemented using the overlay model, where the service provider provides an enterprise customer with the ability to inter-connect many sites utilizing a private WAN IP network. Each site requiring connectivity will receive a router that needs to be peered through an appropriate interior gateway protocol (IGP) to at least one head end router. The backbone here is owned by the service provider and shared between multiple enterprise customers. So the network is not really a private network but a Virtual Private Network.

Currently the enterprise IP network is overlaid on top of the Service Provider backbone (figure 1); the enterprise network is the higher layer network (layer 3) while the backbone network is the lower layer (layer 2). Both networks exist, but independently of each other. The enterprise establishes router-to-router communication using some IGP and the service provider views the routing information as merely more data.

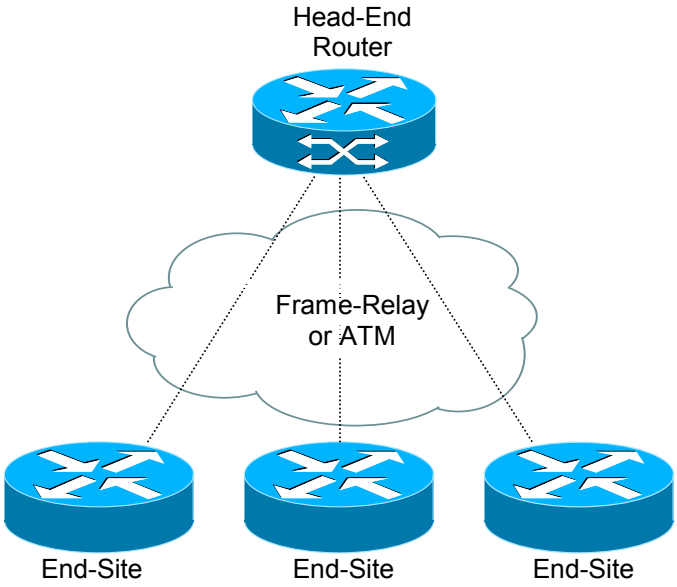

Figure 1: Overlay VPN

For an enterprise to be able to route optimally in this model, it is necessary for the network to be fully meshed (figure 2). This means that every site must have a link to every other site increasing the number of VCs to a total of n*(n-1)/2 where n = number of sites. That increase in the number of VCs required also greatly increases the complexity of the network and the routing protocol. This added complexity makes adding additional sites painful for both the enterprise and the service provider. Traffic engineering is also made more difficult in this model as knowledge of site-to-site traffic is necessary to properly provision the VCs. Plainly stated this model does not scale well for large more meshed topologies.

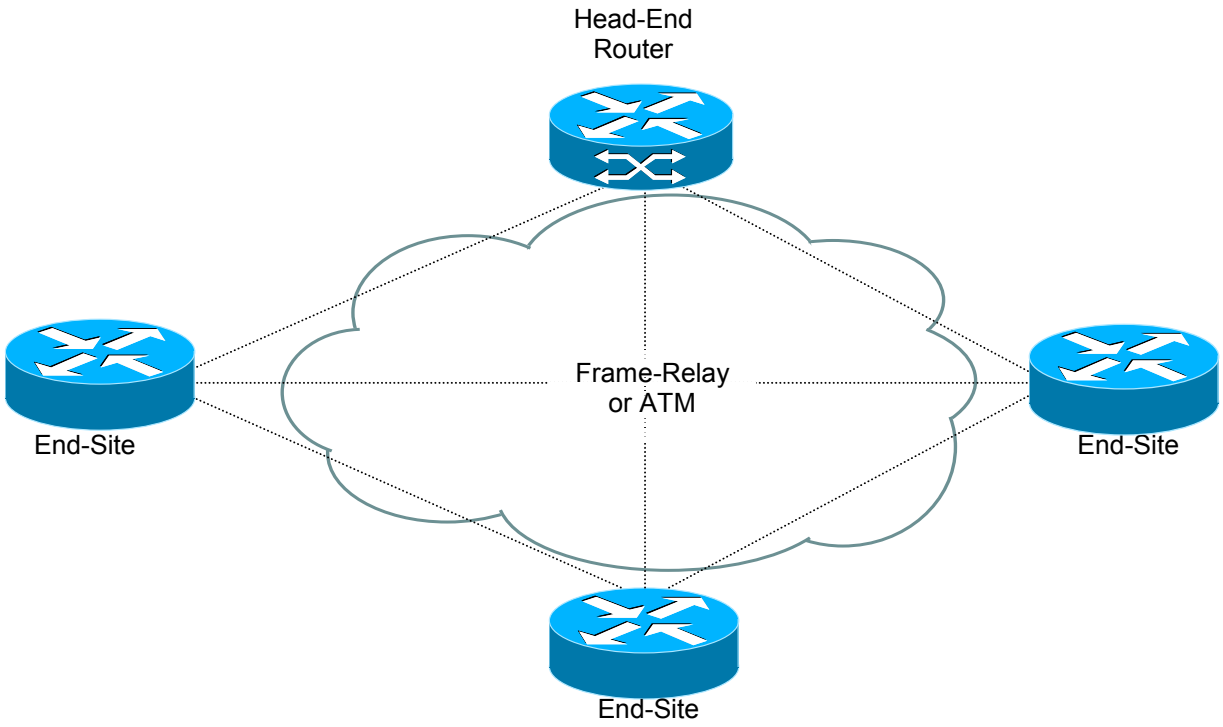

Figure 2: Fully Meshed VPN

Peer Model

Utilizing the peer model, both the service provider and the customer use the same network protocol. In this model the Provider Edge (PE) device is a router that directly exchanges routing information with the CPE router. This provides the ability to simplify the routing from the customer's perspective, as they no longer have to peer with every other end-site instead, only with one PE-router. Routing is now optimal between customer's sites, as the provider routers now know the customer's network topology. Also the addition of a new site is significantly simpler due to the service provider not having to provision a whole new set of VCs.

Two implementation options existed for the peer model prior to MPLS based VPNs, the shared router approach and the dedicated router approach. The shared router approach is where several VPN customers share the same PE-router. This approach has to be concerned with access control, making sure that there is no crossover between different customer's traffic. While the dedicated router utilizes a separate PE router for each VPN customer, causing scalability concerns for the provider. Neither approach allows for the use of private IP addresses (RFC 1918), as each customer would have to have unique addressing.

A major drawback of both of these peer models is their inability to provide traffic isolation. Once the customers are connected to the provider network they need to use unique addressing as all routes are placed in the global routing table. Unlike layer 2

based VPNs it is necessary to look at the layer 3 header to make the forwarding decision. In the early models forwarding over the backbone was done by IP routing.

MPLS-VPN

In this VPN model, MPLS is used for forwarding packets over the backbone, and BGP is used for distributing routes over the backbone. The method is simple for the customer and scalable and flexible for the Service Provider. This method also allows the Service Provider the ability to provide Internet access to these customers as well.

An MPLS-VPN is a "true peer VPN" model that performs traffic separation at Layer 3, through the use of separate IP VPN forwarding tables. MPLS-VPN enforces traffic separation between customers by assigning a unique VRF to each customer's VPN. This compares to the security of a Frame-Relay or ATM network, because users in a specific VPN cannot see traffic outside their VPN.

This is due to the fact that forwarding within the Service Provider backbone is based on labels. These label switched paths (LSPs), setup by MPLS, begin and terminate at the PE routers while the CE routers perform normal routing. It is the job of the incoming interface on the PE to determine which forwarding table to use when handling a packet because each incoming interface on a PE router is associated with a particular VPN. That shows that a packet can enter a VPN only through an interface that is associated with that VPN.

Traffic separation occurs without tunneling or encryption because it is built directly into the network itself. MPLS-VPN uses Multi-protocol BGP extensions to encode customer IPv4 address prefixes into unique VPN-IPv4 NLRIs. Through the use of the Extended BGP community attribute the PE routers are able to control the distribution of these routes. These PE routers also assign a label with each VPN customer route and share these labels with other PEs, assuring that data packets are directed to the correct egress CE.

When a data packet is forwarded two labels are used. The top label directs the traffic to the correct PE router while the second label indicates how the PE should handle that packet. MPLS then takes over by forwarding the packet across the backbone using dynamic IP paths or traffic engineered paths.

To simplify things further, standard IP forwarding is used between the PE and CE routers. The PE has a per-site VRF forwarding table that contains only the set of routes available to that CE router. The CE router is a routing peer of the PE to which it is directly connected but is not a routing peer of CE routers at other sites. Routers at different sites don't directly exchange routing information with one another. This allows for very large VPNs to be easily supported while simplifying the routing configuration at each individual site.

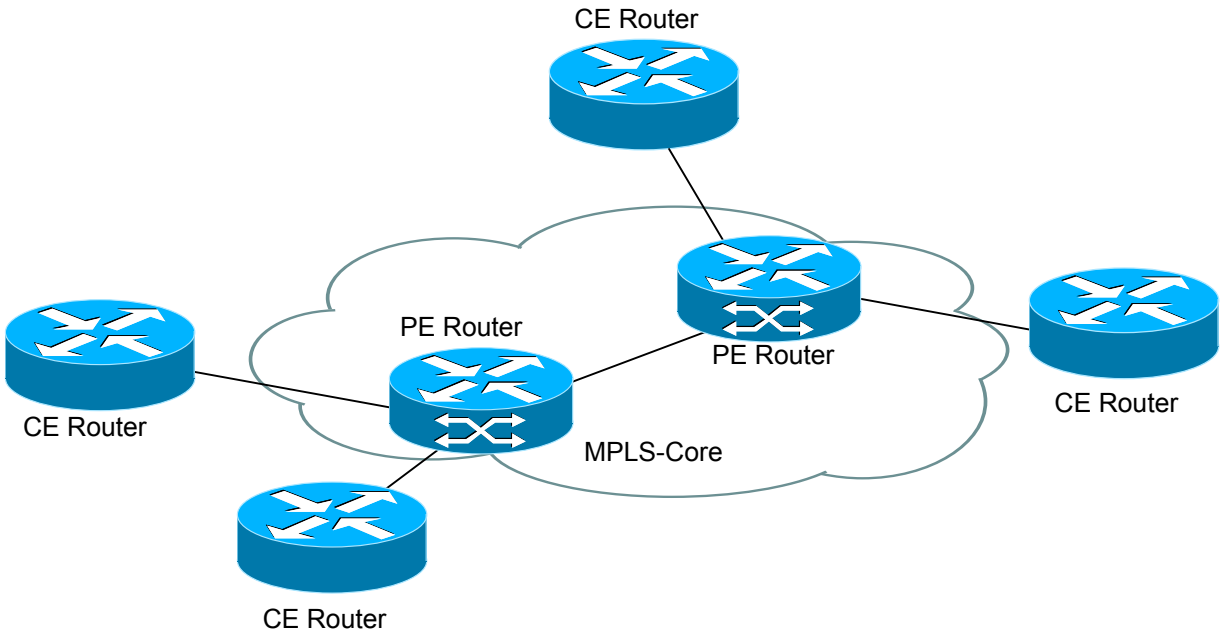

Figure 3: MPLS-VPN

Requirements of a Secure Network

When comparing MPLS-VPN based solutions to traditional layer 2 based VPN solutions such as Frame-Relay and ATM, several key security requirements need to be addressed.

- It is necessary to have addressing and routing separation.
- The internal structure of the backbone network must be hidden from the outside. Just as a Frame-Relay or ATM network core is hidden, so must an MPLS-VPN core.
- The network must have resistance to attacks, both Denial-of-Service (DoS) and intrusion attacks.

Addressing separation implies that between two non-intersecting VPNs the address spaces between them are entirely independent. For example two VPNs can use the exact same address space and not interfere with each other. From the routing perspective this means that each end system in a VPN has a unique address, so no two sites in the same VPN share the same address space. ATM and Frame-Relay have no problem implementing these features, as they never look at the layer 3 information. The forwarding decision is made on layer 2 based criteria such as DLCIs and VPI/VCI pairs.

Hiding the internal structure of the backbone states that there should be little or no visibility into the core from outside networks. As there is no layer 3 connectivity between the customer equipment and the Frame-Relay or ATM switch the only visibility

into the internal network is the VC information needed to bring up the connection. Ideally the MPLS core should be as invisible as a comparable Frame-Relay or ATM core.

Resistance to attacks includes both Denial-of-Service (DoS), where resources become unavailable to authorized users, and intrusions attacks or gaining unauthorized access. As most DoS attacks are based on layer 3 attributes, Frame-Relay and ATM aren't particularly vulnerable to this type of attack. If an attack were committed it would be internal to the VPN, as the network would simply pass these attacking packets through without looking above the layer 2 DLCI or VPI/VCI.

Validating MPLS-VPN as a Secure Network

MPLS-VPNs were explained briefly in a previous section. What the next sections concentrate on is how MPLS based VPNs compare to Frame-Relay and ATM. Frame-Relay and ATM are well known in the industry and have the reputation as being secure. In order to consider MPLS-VPNs to be as secure as layer 2 based VPNs, the security characteristics described earlier must be met or exceeded.

We will go over the testing that was performed in order to prove that MPLS based VPNs are as secure as comparable Frame-Relay or ATM based VPNs. The format of the next sections will be as follows:

- The security characteristic being tested will be defined
- How layer 2 based VPNs handle this characteristic
- How MPLS-VPNs handle this characteristic
- How we tested this characteristic with MPLS-VPNs
- The results of those tests

Note that this paper concentrates on protecting the core from outside attacks. Protection against inside attacks is not considered here as any network can be attacked with access from the inside.

Address Space and Routing Separation

Business customers today need the flexibility of maintaining their own addressing plans and the freedom to use either public or private address space. Both ATM and Frame-Relay, as layer 2 based solutions, provide this flexibility. Neither technology examines the layer 3 portion of the packet, which contains the addressing, but rather makes the forwarding decision on DLCI (Frame-Relay) or VPI/VCI (ATM) information.

MPLS on the other hand, does look at the layer 3 portion of the packet but still is able to allow multiple VPNs to use the same address space. Also MPLS-VPNs allows the use of public or private addressing. This is possible by adding a 64-bit route distinguisher (RD) to each IPv4 route. This new route called a "VPN-IPv4 address" ensures that VPN-unique addresses are also unique in the MPLS core. The only exception here is the IP addressing of the PE to CE links, they will need to be unique if using dynamic routing protocols.

Routing separation between business customers is also a necessity. Again because layer 2 based VPNs never look at the layer 3 header they don't route, instead they switch by examining the layer 2 information (DLCI, VPI/VCI). MPLS provides route separation by having each PE router maintain a separate routing table for each connected VPN. This routing table called a Virtual Routing and Forwarding instance (VRF) contains the routes from one VPN that were learned statically or through a dynamic routing protocol. These VRFs are separate from each other as well as from the global routing table.

This separation is maintained across the MPLS core to the other PE routers by utilizing multi-protocol BGP (MP-BGP.) By adding unique VPN identifiers such as the route distinguisher, multi-protocol BGP has provided the ability to uniquely identify VPN routes through the core of the network. MP-BGP is the only way that VPN routes are exchanged across the core. These BGP routes are not re-distributed into the core network but only to the other PE routers, in fact the core network routers do not need to run BGP. Instead the PE routers exchange the information with each other and then place the information in VPN specific VRFs. Using these features, routing across an MPLS network is separate per VPN.

For our test we built our addressing scheme so that we had the ability to test this functionality first hand. Basically we had a scenario that involved three different VPNs; two of which shared the exact same address space utilizing private addressing while the third VPN used the public space. Connectivity was verified by using ICMP and telnet to make sure that traffic stayed within the VPN boundaries.

To prove that MPLS based VPNs provided both addressing and routing separation we examined the routing table of every CE, PE, and P router. On the CE routers we verified that the routes that existed belonged solely to the VPN that the CE was a member of and there were no routes to other VPNs or the core. We repeated this step on the PE routers by verifying that each VRF routing table contained the same information. Once we reached the P routers we verified that there were no VRF routing tables and the only routes that appeared were to other routers in the providers network.

Next we wanted to verify traffic being initiated from inside the VPN stayed inside that VPN. Our next test employed a traffic injection tool to verify that when you have two VPNs off the same PE router that the traffic stayed isolated. In other words when you have two CE's with the same address space in different VPNs, only the CE in the same VPN as the source received traffic.

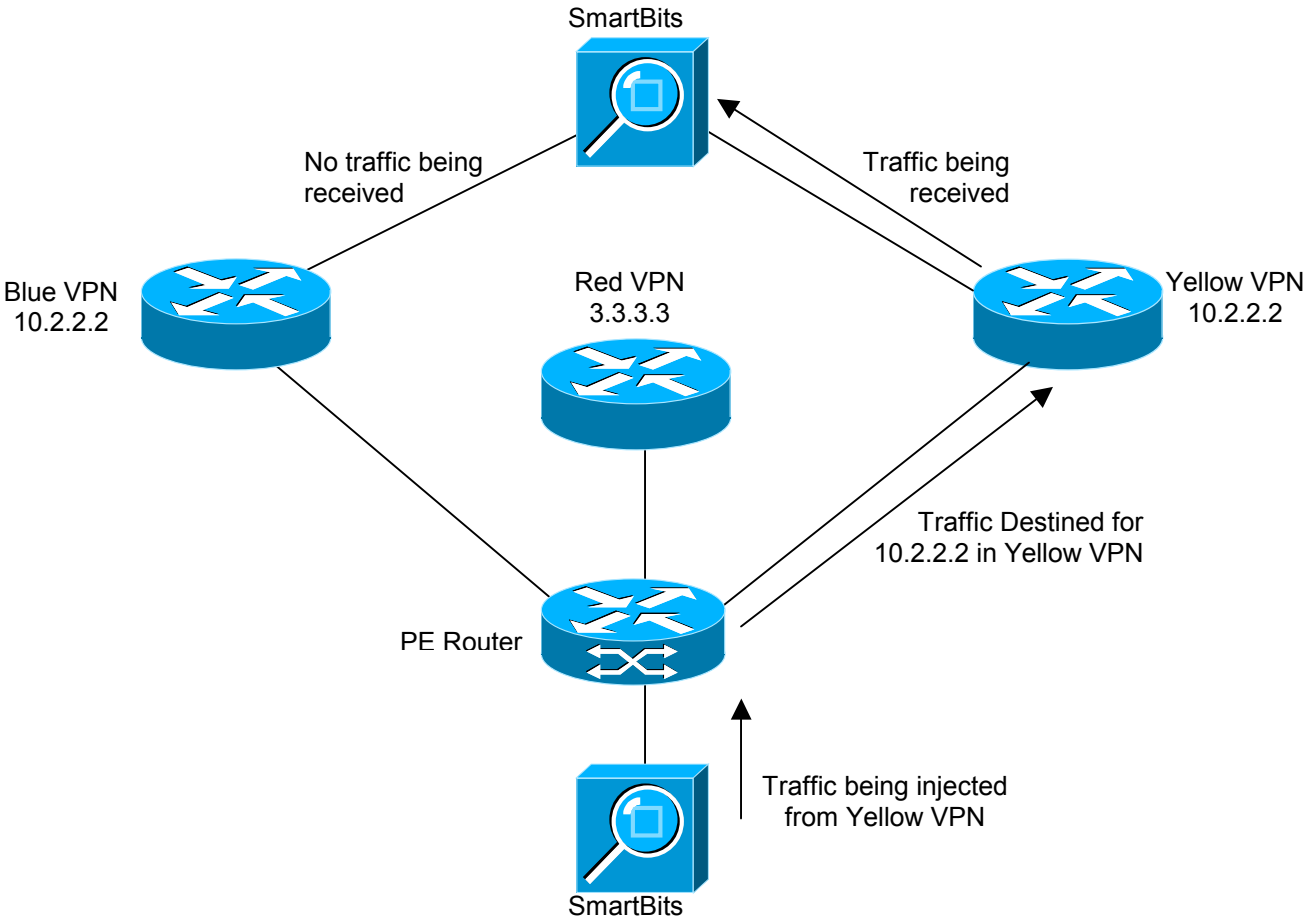

Figure 4: Traffic Isolation Test

In these tests MPLS based VPNs were shown to have the same addressing and routing separation capabilities as comparable layer-2 VPNs such as Frame-Relay or ATM. The only way possible to intrude into other VPNs through the MPLS core is if this has been specifically configured (extranet configuration.)

Hiding the Service Provider Core Network

Service providers and customers do not want their network topology revealed to the outside world. Without knowledge of the network topology an attacker can only guess the IP addresses to attack, thus making the network much more difficult to attack. However with a known IP address an attacker can launch a DoS attack against that device without a high degree of difficulty. So ideally we do not want to reveal any information of the internal network to the outside.

Currently layer 2 based VPNs such as Frame-Relay and ATM handle this quite well. The only information that is shared between the service provider network and the

customer network is information about the customer's VCs. This limits the view of the provider's topology from the customer's perspective. The customer is aware of the core due to the information he received regarding the VCs, but has no other knowledge.

The same ideals apply to customer networks as to the MPLS core. MPLS doesn't reveal additional unnecessary information even to customer VPNs. Since the interface to the VPNs is BGP there is no need to reveal any information about the core. The only information required in the case of a routing protocol between PE and CE is the address of the PE router. If this is not desired, static routing can be configured between the PE and CE. With this measure the MPLS core can be kept completely hidden and be addressed using public or even private address.

The way we tested this functionality was again by sending ICMP packets and performing telnet tests. We had an advantage in this test because we knew the addressing used in the core, however even then we were not able to have any reach ability into the service provider's core network. These tests proved that there was no access from the CE routers to the PE and P routers[1].

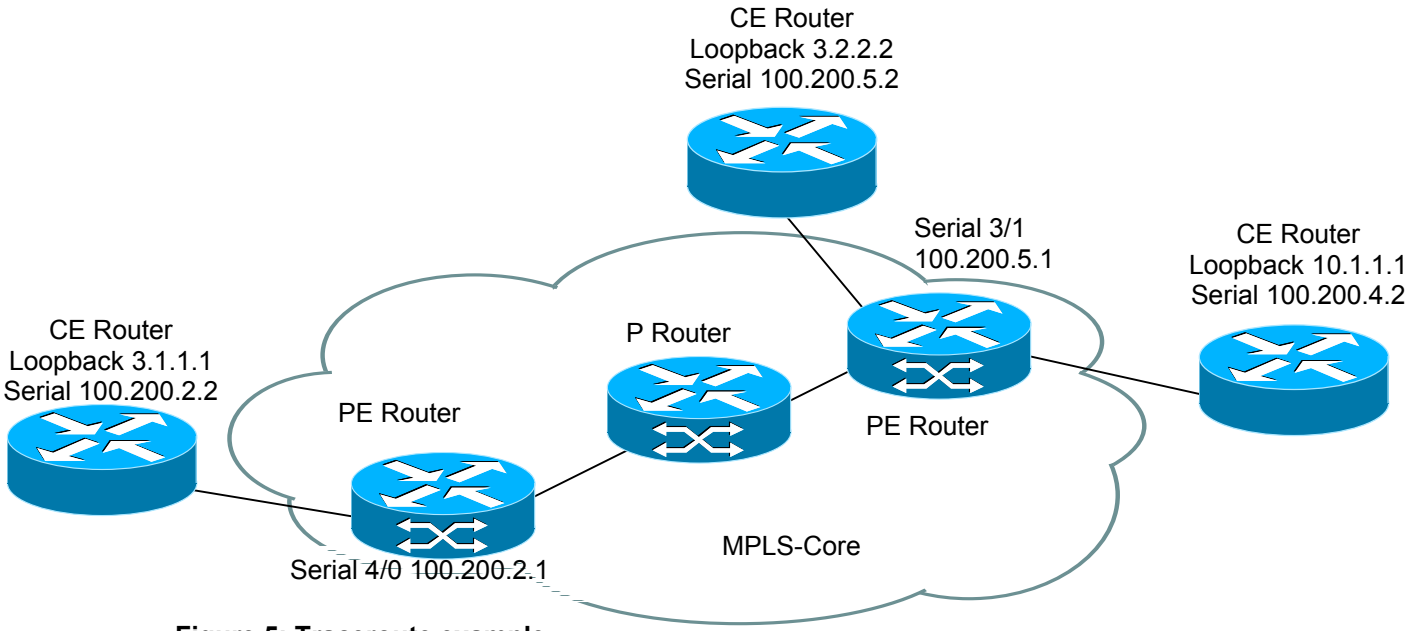

Figure 5: Traceroute example

[1] Due to the fact that we had IP addresses on the interfaces between the PE and CE routers we had to place access-lists on the PE and CE router to deny telnet traffic. This is due to the fact that the IP address on the PE router belongs to the VPN not the global routing table. As part of a good security practice we have configured access lists to deny telnet traffic, this should be a part of any router security policy

Also while using the traceroute utility the MPLS cloud does not show up as a hop in the output. See this example of a traceroute through the network in figure 5.

```
CE-router>trace 3.1.1.1

Type escape sequence to abort.
Tracing the route to 3.1.1.1

  1 100.200.5.1 0 msec 0 msec 0 msec
  2 100.200.2.2 4 msec 4 msec 4 msec
```

Resistance to Attacks

Whether it is a layer 2 VPN or an MPLS-VPN the service provider's network should be resistant to attacks. An attack from inside a VPN should be contained to that VPN, not having any service effect on the other VPNs. And an attacker should not be able to gain access into other VPNs or into the service provider's network.

Traditional layer 2 VPNs such as Frame-Relay and ATM are able to handle this by some of the methods already addressed in previous sections. As the provider only tells the customer about VCs there is a lack of places to attack. Also due to the fact they operate at layer 2 and not layer 3 there are few types of attacks that can be launched against them.

We have proved earlier that it is impossible to gain access into other VPNs and thus it is impossible to attack other VPNs from within a VPN. However in MPLS there is the possibility to attack the MPLS core and to attack other VPNs from there. There are two basic ways the MPLS core can be attacked: first by attempting to attack the PE routers directly, and second by attempting to attack the signaling mechanisms of MPLS.

In order to attack the PE routers directly it is necessary to know its address. As discussed in hiding the MPLS Core it is possible to hide the addressing structure of the MPLS core from the outside world except for when running a dynamic routing protocol. In that case the router will know at least the router ID of the PE router in the core.

If an attacker does not know the IP address of any router in core the attacker now has to guess addresses and send packets to these addresses. However, due to the address separation of MPLS each incoming packet will be treated as belonging to the address space of the customer. Thus it is highly difficult to reach an internal router, even through IP address guessing.

If an attacker does know the IP address of the router he wants to attack then you could imagine the possible attacks on various services running on the router. In practice access to the PE router over the CE-PE interface can be limited to the required routing protocol by using access control lists. This limits the point of attack to one routing protocol, for example RIP or BGP. A potential attack could be to send an extensive number of routes, or to flood the PE router with routing updates. Both could lead to a DoS, however, not to unauthorized access. To restrict this risk it is necessary to configure the routing protocol on the PE router as securely as possible. This can be done in various ways:

- By Access Control List (ACL), allow the routing protocol only from the CE router, not from anywhere else. Furthermore, no access other than that should be allowed to the PE router in the inbound ACL on each CE interface.
- Where available, configure MD-5 authentication for routing protocols. This is available for BGP, OSPF, and RIPv2, which are the only supported dynamic routing protocols for MPLS-VPN. It avoids packets that could be spoofed from parts of the customer network other than the CE router.
- To configure available parameters of the routing protocols, such as BGP where it is possible to configure route dampening, which limits the number of routing interactions.
- VRFs can limit the maximum number of routes that are accepted into that VRF routing table.

It has to be mentioned that although in the static case the CE router doesn't know any of the IP addresses of the PE router, it is still attached to the PE router via some method, and could thus guess the address of the PE router and try to attack it with this address.

We divided this section up into two separate tests attacking the PE router and then attacking the MPLS signaling methods. The MPLS signaling methods are covered in the next section, MPLS Label Spoofing, this section focuses on the tests that we performed to attack the PE.

We performed in depth testing of DoS attacks by utilizing a tool that injected high numbers of RIP (50,000) and OSPF (82,000) routes into a PE router. We proceeded to separate the test into three scenarios tested with each of the routing protocols: First inject routes into the PE without using BGP or VRF route filtering, second inject routes into the PE while using BGP route filtering, and finally inject routes with both BGP and VRF router filtering.

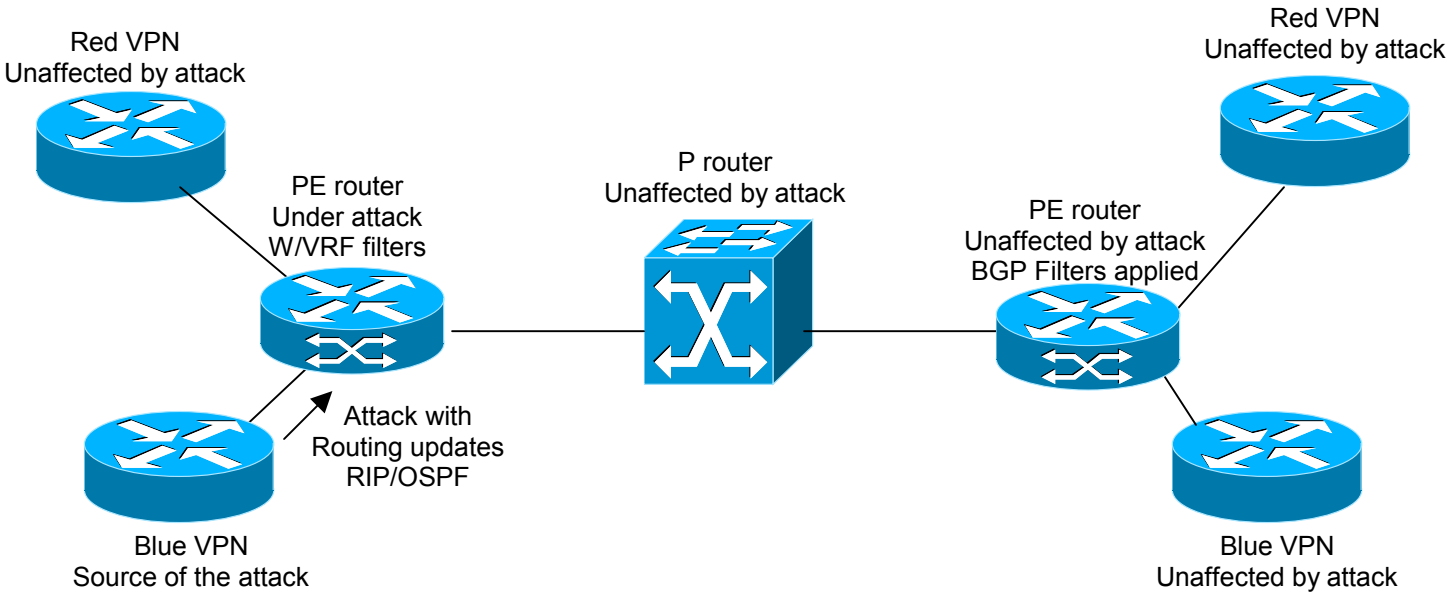

Figure 6: DoS Attack

Our results showed that when BGP and VRF route filtering are configured we were not able to disrupt traffic for any other VPN when attacking from inside a VPN. We observed that a traffic stream from one VPN was not disrupted by the attack on the PE from a different VPN. This provides not only substantial security against DoS attacks but also by being able to limit the number of routes a customer can inject provides a scenario where service providers can implement different service levels.

Overall it is not possible to intrude from one VPN into other VPNs, or the core making MPLS-VPNs as secure as their Frame-Relay and ATM counterparts. It is theoretically possible to exploit the routing protocol to execute a DoS attack against the PE router that might have negative impact on other VPNs. However if the PE routers do the proper filtering this threat is negated. PE routers must be extremely well secured, especially on their interfaces to the CE routers. Access control lists should be configured to limit access only to the port(s) of the routing protocol, and only from the CE router. MD5 authentication in routing protocols should be used with all PE-CE peers.

MPLS Label Spoofing

At the core of the MPLS network packets are not forwarded based on the IP destination address, but rather based on labels that are pre-pended by the PE routers. Similar to IP spoofing attacks, where an attacker replaces the source or destination IP address of a packet, it is also theoretically possible to spoof the label of an MPLS packet. In the earlier sections the assumption was made that the service provider secures the core network. Thus in this section we emphasize whether it is possible to insert packets with

(wrong) labels into the MPLS network from the outside, i.e., from a VPN (CE router) or from the Internet.

In the Frame-Relay and ATM world this would be equivalent to inserting DLCIs or VPI/VCI pairs. However if those DLCIs or VPI/VCIs are not configured on the specific port the traffic is dropped.

In MPLS the interface between a CE router and its peering PE router is an IP interface, i.e., an interface without labels. The CE router is unaware of the MPLS core, and thinks it is sending IP packets to a simple router. The "intelligence" is done in the PE device, where based on the configuration, the label is chosen and pre-pended to the packet. This is the case for all PE routers, towards CE routers as well as the upstream service provider. All interfaces into the MPLS cloud only require IP packets, without labels. For security reasons a PE router should never accept a packet with a label from a CE router. In Cisco routers the implementation is such that labeled packets that arrive on any interface where label switching is not enabled will be dropped. Thus it is not possible to insert fake labels, since no labels will be accepted.

There remains the possibility to spoof the IP address of a packet that is being sent to the MPLS core. However, since there is strict addressing separation within the PE router, and each VPN has its own VRF, this can only do harm to the VPN the spoofed packet originated from, in other words, a VPN customer can attack himself. MPLS doesn't add any security risk here as the service provider's network is not threatened and service to other VPNs will not be impacted. It is the responsibility of the customer to properly secure their CE routers against this.

This section of the test was done using the industry standard SmartBits traffic injection tool from Spirent. This box has the capabilities to inject MPLS labeled packets. We proceeded to inject traffic with labels on it into PE routers and watched the traffic drop, even if it was a valid label. We repeated the test on P routers as well.

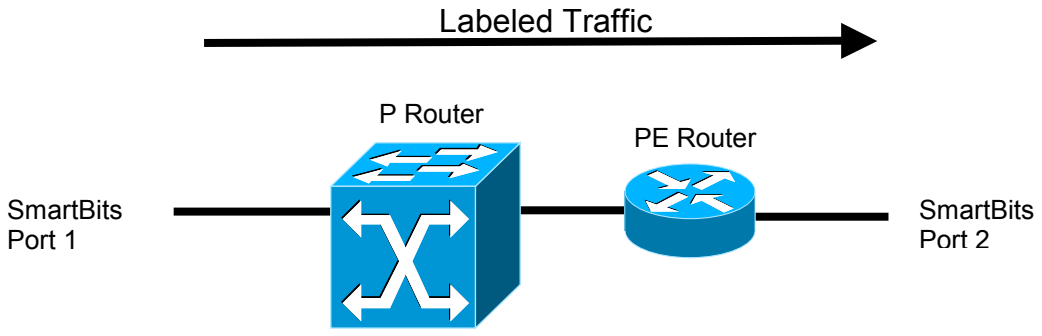

Figure 7: Spoofed MPLS traffic

As it is impossible to insert a 'spoofed' label into an MPLS network and thus gain access to another VPN or the MPLS core. In this capability a MPLS based VPN provides the same security as a Frame-Relay or ATM based VPN.

Summary

Our tests results have demonstrated that MPLS based VPN networks have met or exceeded all of the security characteristics of a comparable layer 2 based VPN such as Frame-Relay or ATM. Business customers, due to several features, consider layer 2 based VPNs secure. MPLS-VPNs perform the same features:

- Address space and routing separation are achieved through the use of a per VPN routing table and MPLS switching in the core.
- While at the same time not revealing the service provider core structure so it appears to be as invisible as their Frame-Relay and ATM counterparts.
- And even though the MPLS based VPN solution exchanges layer 3 routing information with the CE routers there are mechanisms in place to limit the impact of DoS attacks to the VPN where they originated. Which equates to the way these attacks are handled in Frame-Relay and ATM based network.

Therefore it is our conclusion that Cisco MPLS-VPNs can offer the same level of security as Frame Relay or ATM to both business customers and service providers.