

# Security of the MPLS Architecture



## Scope and Introduction

Many enterprises are thinking of replacing traditional Layer 2 VPNs such as ATM or Frame Relay (FR) with MPLS-based services. As Multiprotocol Label Switching (MPLS) is becoming a more widespread technology for providing virtual private network (VPN) services, MPLS architecture security is of increasing concern to service providers (SPs) and VPN customers. This paper gives an overview of MPLS architecture security for both SPs and MPLS users, and compares it with traditional Layer 2 services from a security perspective. This paper also recommends how to secure an MPLS infrastructure. The focus is specifically on the MPLS/Border Gateway Protocol (BGP) VPN architecture.

The Miercom group has also undertaken research in this field and conducted practical testing of MPLS architecture security [Miercom].

MPLS is being used to achieve the following results: to engineer the core network more easily and efficiently (traditional MPLS and MPLS traffic engineering), to provide VPN services (MPLS-VPN), and to facilitate quality of service (QoS) across a network core (MPLS-DBP). In this paper, the main emphasis is on security of the VPN provisioning aspect of MPLS, although most of it applies to other aspects of MPLS.

This paper assumes that the MPLS core network is provided in a secure manner. Thus, it does not address basic security concerns such as securing the network elements against unauthorized access, misconfigurations of the core, internal (within the core) attacks, and so on. If a customer does not wish to assume the SP network is secure, it becomes necessary to run IP Security (IPSec) over the MPLS infrastructure (Section 6).

Analysis of the security features of routing protocols is covered only to the extent that it influences MPLS. This paper does not cover IPSec technology, except to highlight the combination of MPLS with IPSec.

Part A covers an analysis of the security that MPLS provides, compared to similar Layer 2 infrastructures. It targets the frequently asked question whether MPLS-based VPN services offer at least the same degree of security as ATM or Frame Relay-based VPNs. Section 2 outlines the security requirements for such networks, and Section 3 analyzes MPLS BGP/VPN with respect to these requirements.

Part B offers guidelines to secure an MPLS infrastructure. It discusses securing routing toward an MPLS core and interconnections between VPNs and Internet access. For additional security such as encryption, IPSec over an MPLS infrastructure is discussed, as well as remote access via IPSec into a specific VPN. The last section outlines topics that the MPLS architecture does not cover.

This paper is targeted at technical staff of SPs and enterprises. Knowledge of the basic MPLS architecture is required to understand this paper.

## Part A: Analysis of the Security of the MPLS Architecture

This part answers the frequently asked question, whether MPLS provides the same level of security as traditional Layer 2 VPNs such as ATM and Frame Relay. Section 2 contains the requirements typically put forward by users of ATM or Frame Relay services, and Section 3 examines whether MPLS complies with these requirements.

### Security Requirements of MPLS Networks

Both SPs offering MPLS services and customers using them have specific demands for the security of this special VPN solution. Mostly they compare MPLS-based solutions with traditional Layer 2-based VPN solutions such as Frame Relay and ATM, because these are widely deployed and accepted. This section outlines the security requirements typical in MPLS architectures. The next section discusses if and how MPLS addresses these requirements, for both the MPLS core and the connected VPNs.

#### Address Space and Routing Separation

Between two nonintersecting VPNs of an MPLS VPN service, it is assumed that the address space between different VPNs is entirely independent. This means, for example, that two nonintersecting VPNs must be able to use the 10/8 network without any interference. From a routing perspective, this means that each end system in a VPN has a unique address, and all routes to this address point to the same end system. Specifically:

- Any VPN must be able to use the same address space as any other VPN.
- Any VPN must be able to use the same address space as the MPLS core.
- Routing between any two VPNs must be independent.
- Routing between any VPN and the core must be independent.

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 given VPN reach a host with the same address in another VPN or the core.

#### Hiding of the MPLS Core Structure

The internal structure of the MPLS core network (provider edge (PE) and provider (P) elements) should not be visible to outside networks (Internet or any connected VPN). Although a breach of this requirement does not lead to a security problem, many SPs feel this is advantageous if the internal addressing and network structure remains hidden to the outside world. A strong argument is that denial-of-service attacks against a core router, for example, are much easier to carry out if an attacker knows the address. Where addresses are not known, they can be guessed, but with this limited visibility, attacks become more difficult.

Ideally, the MPLS core should be as invisible to the outside world as a comparable Layer 2 (such as Frame Relay or ATM) infrastructure.



## Resistance to Attacks

There are two basic types of attacks: denial-of-service (DoS) attacks, where resources become unavailable to authorized users, and intrusions, where the underlying goal is to gain unauthorized access to resources. Table 1 shows the two basic types of attack.

Table 1 Types of Attacks

	Has Access	Has No Access
Authorized User	Normal	Denial of service
Unauthorized User	Intrusion	Normal

- For attacks that give unauthorized access to resources (intrusions), there are two basic ways to protect the network: first, to harden protocols that could be abused (such as Telnet to a router), and second, to make the network as inaccessible as possible. The latter is achieved by a combination of packet filtering or use of firewalls and address hiding, as discussed above.
- DoS attacks are easier to execute, because in the simplest case, a known IP address might be enough to attack a machine. The only way to be certain the network is invincible to this kind of attack is to make sure that machines are not reachable, again by packet filtering and address hiding.

MPLS networks must provide at least the same level of protection against both forms as current Layer 2 networks do. Note that this paper concentrates on protecting the core network against attacks from the “outside,” or the Internet and connected VPNs. This paper does not consider protection against attacks from the “inside,” for example, if an attacker has logical or physical access to the core network, because any network can be attacked with access from the inside.

### Impossibility of Label Spoofing

In a pure IP network, it is easy to spoof IP addresses, a key issue in Internet security. Because MPLS works internally with labels instead of IP addresses, the question arises whether these labels can be spoofed as easily as IP addresses.

Assuming the address and routing separation as discussed above, a potential attacker might try to gain access to other VPNs by inserting packets with a label that he doesn't “own.” This could be done from the outside, for example, another customer edge (CE) router or from the Internet, or from within the MPLS core. This paper does not discuss the latter case (from within the core), because the assumption is that the core network is provided in a secure manner (see also Section 8). If a network requires protection against an insecure core, it is necessary to run IPSec on top of the MPLS infrastructure (discussed in Section 6).


It must be impossible to send packets with wrong labels from a CE router (the “outside”) through a PE into the MPLS cloud, because this would make packet spoofing possible.

## Analysis of MPLS Security

In this section the MPLS architecture is analyzed with respect to the security requirements listed above.

### Address Space and Routing Separation

Figure 1 Format of a VPN IPv4 Address



MPLS allows distinct VPNs to use the same address space, which can also be private address space [RFC1918]. This is achieved by adding a 64-bit route distinguisher (RD) to each IPv4 route, making VPN-unique addresses also unique in the MPLS core. This “extended” address is also called a “VPN-IPv4 address” and is shown in Figure 1. Thus, customers of an MPLS service do not need to change current addressing in their networks.

There is only one exception, which is the IP addresses of the PE routers the CE routers are peering with, in the case of using routing protocols between CE and PE routers (for static routing this is not an issue). Routing protocols on the CE routers need to have configured the address of the peer router in the core, to be able to “talk” to the PE router. This address must be unique from the perspective of the CE router and thus belongs logically to the address space of the VPN. In an environment where the SP also manages the CE routers as CPE, this setup can be made invisible to the customer.

Routing separation between the VPNs can also be achieved. Every PE router maintains a separate Virtual Routing and Forwarding instance (VRF) for each connected VPN. Each VRF on the PE router is populated with routes from one VPN, through statically configured routes or through routing protocols that run between the PE and the CE router. Because every VPN results in a separate VRF, there will be no interferences between the VPNs on the PE router.

Across the MPLS core to the other PE routers, this separation is maintained by adding unique VPN identifiers in multiprotocol BGP (MP BGP), such as the route distinguisher. VPN routes are exclusively exchanged by MP-BGP across the core, and this BGP information is not redistributed to the core network; it is redistributed only to the other PE routers, where the information is kept again in VPN-specific VRFs. Thus, routing across an MPLS network is separate per VPN.

Given the addressing and routing separation across an MPLS core network, we can assume that MPLS offers, in this respect, the same security as comparable Layer 2 VPNs such as ATM or Frame Relay. It is not possible to intrude into other VPNs through the MPLS cloud, unless this has been configured specifically.



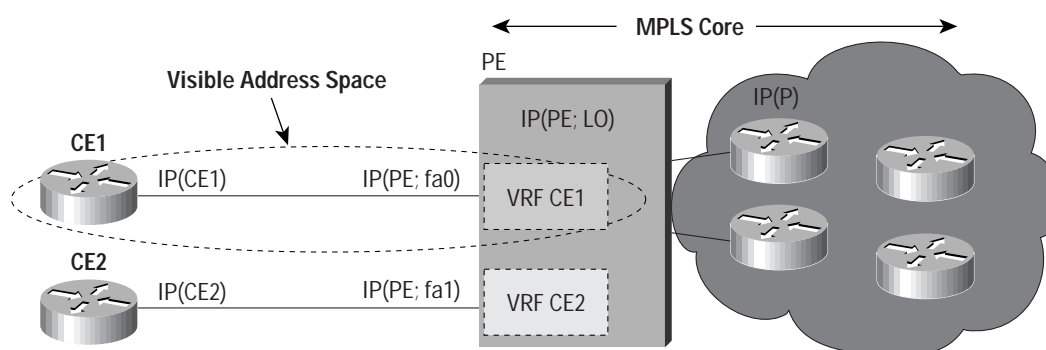
## Hiding of the MPLS Core Structure

For reasons of security, SPs and end customers do not normally want their network topologies revealed to the outside. This makes attacks more difficult. If an attacker does not know the target, he/she can only guess the IP addresses to attack or try to find out about addressing through a form of intelligence. Because most DoS attacks do not provide direct feedback to the attacker, a network attack is difficult.

With a known IP address, a potential attacker can launch a DoS attack against that device. So the ideal is to not reveal any information of the internal network to the outside. This applies equally to the customer networks as to the MPLS core. In practice, numerous additional security measures have to be taken, primarily extensive packet filtering.

Figure 2 shows the visible address space of a given VPN. No P routers or other VPNs are visible to VPN1. The link between the CE and PE routers, which includes the interface address of the PE router, belongs to the VPN address space. All other addresses on the PE router, such as loopback interfaces, are not part of the VPN address space.

Figure 2 Hiding of the Core Infrastructure



MPLS does not reveal unnecessary information to the outside, not even to customer VPNs. Core addressing can be conducted with private addresses [RFC1918] or public addresses. Because the interface to the VPNs—and potentially the Internet—is BGP, there is no need to reveal any internal information. The only information required in the case of a routing protocol between PE and CE is the address of the PE router (IP PE in Figure 2). 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.

Customer VPNs will have to advertise their routes as a minimum to the MPLS core, to ensure reachability across the MPLS cloud. Although this could be seen as too “open,” the following must be noted: First, the information known to the MPLS core is not about specific hosts, but networks (routes); this offers some degree of abstraction. Second, in a VPN-only MPLS network (such as one with no shared Internet access), this is equal to existing Layer 2 models in which the customer must trust an SP to some degree. Also, in a FR or ATM network, routing information about the VPNs can be seen on the core network.

In a VPN service with shared Internet access, an SP will typically announce the routes of customers who wish to use the Internet to upstream or peer providers. This can be done via a Network Address Translation (NAT) function to further obscure the addressing information of the customers’ networks. In this case, the customer does not reveal more information to the general Internet than with a general Internet service. Core information will still not be revealed at all, except for the peering address(es) of the PE router(s) that hold(s) the peering with the Internet.

In summary, in a pure MPLS-VPN service, where no Internet access is provided, the information hiding is as good as on a comparable FR or ATM network; no addressing information is revealed to third parties or the Internet. If a customer chooses to access the Internet via the MPLS core, the customer must reveal the same addressing structure as for a normal Internet service. NAT can be used for further address hiding.

If an MPLS network has no interconnections to the Internet, this is equal to FR or ATM networks. With an Internet access from the MPLS cloud, the SP has to reveal at least one IP address (of the peering PE router) to the next provider, and thus the outside world.

#### Resistance to Attacks

Section 3.1 shows that it is not possible to directly intrude into other VPNs. The only other possibility is to attack the MPLS core, and try to attack other VPNs from there. The MPLS core can be attacked in two basic ways:

- By attacking the PE routers directly
- By attacking the signaling mechanisms of MPLS (mostly routing)

To attack an element of an MPLS network, it is first necessary to know its address. As discussed in Section 3.2, it is possible to hide the addressing structure of the MPLS core to the outside world. Thus, an attacker does not know the IP address of any router in the core that he/she wants to attack. The attacker could now guess addresses and send packets to these addresses. However, because of the address separation of MPLS, each incoming packet will be treated as belonging to the address space of the customer. Thus it is impossible to reach an internal router, even through IP address guessing. This rule has only one exception, which is the peer interface of the PE router.

The routing between the VPN and the MPLS core can be configured two ways:

1. Static—In this case the PE routers are configured with static routes to the networks behind each CE, and the CEs are configured to statically point to the PE router for any network in other parts of the VPN (mostly a default route). There are now two subcases: The static route can point to the IP address of the PE router, or to an interface of the CE router (for example, serial0).
2. Dynamic—Here a routing protocol (for example, Routing Information Protocol [RIP], Open Shortest Path First [OSPF], BGP) is used to exchange the routing information between the CE and the PE at each peering point.

In the case of a static route from the CE router to the PE router, which points to an interface, the CE router does not need to know any IP address of the core network, not even of the PE router. This has the disadvantage of a more extensive (static) configuration, but from a security point of view is preferable to the other cases.

In all other cases, each CE router needs to know at least the router ID (RID; peer IP address) of the PE router in the MPLS core, and thus has a potential destination for an attack. One could imagine various attacks on various services running on a router. In practice, access to the PE router over the CE/PE interface can be limited to the required routing protocol by using ACLs (access control lists). This limits the point of attack to one routing protocol, for example 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 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 Message Digest 5 (MD5) authentication for routing protocols—This is available for BGP [RFC2385], OSPF [RFC2154], and RIP2 [RFC2082], for example. It prevents packets from being spoofed from parts of the customer network other than the CE router. Note that this requires that the SP and customer agree on a shared secret between all CE and PE routers. The problem here is that it is necessary to do this for all VPN customers—it is not sufficient to do this for the customer with the highest security requirements.
- Configure, where available, parameters of the routing protocol, in order to further secure this communication—In BGP, for example, it is possible to configure dampening, which limits the number of routing interactions. Also, a maximum number of routes accepted per VRF should be configured where possible.

It should be noted that although in the static case the CE router does not know any IP address of the PE router, it is still attached to the PE router via some method; therefore, it could guess the address of the PE router and try to attack it with this address.

In summary, it is not possible to intrude from one VPN into other VPNs, or the core. However, it is theoretically possible to exploit the routing protocol to execute a DoS attack against the PE router. This in turn might have a negative impact on other VPNs. Therefore, PE routers must be extremely well secured, especially on their interfaces to the CE routers. ACLs must 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 on all PE/CE peerings. It is easily possible to track the source of such a potential DoS attack.

### 3.4 Label Spoofing

Within the MPLS, network packets are not forwarded based on the IP destination address, but based on labels that are prepended 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 first section, the assumption was made that the core network is secured by the SP. (If this assumption cannot be made, IPSec must be run over the MPLS cloud.) Thus in this section the emphasis is on whether it is possible to insert packets with (wrong) labels into the MPLS network from the outside, that is, from a VPN (CE router) or from the Internet.

Principally, the interface between any CE router and its peering PE router is an IP interface (that is, 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 prepended to the packet. This is the case for all PE routers, toward CE routers as well as the upstream SP. All interfaces into the MPLS cloud require only 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 packets that arrive on a CE interface with a label will be dropped. Thus it is not possible to insert fake labels, because no labels at all are accepted.

There remains the possibility to spoof the IP address of a packet that is being sent to the MPLS core. However, because there is strict addressing separation within the PE router, and each VPN has its own VRF, this can harm only the VPN that the spoofed packet originated from; in other words, a VPN customer can attack himself/herself. MPLS does not add any security risk here.

### Comparison with ATM/FR VPNs

ATM and FR VPN services often enjoy a very high reputation in terms of security. Although ATM and FR VPNs can also be provided in a secure manner, it has been reported that these technologies can also have severe security vulnerabilities [DataComm]. Also, in ATM/FR the security depends on the configuration of the network being secure, and errors can also lead to security problems.

## Part B: Options for Securing an MPLS Core

This part targets the SP: It tries to outline how MPLS-based VPN services can be secured, and what has to be addressed to implement network-based services such as Internet access, remote access to a VPN, or firewalling. It also explains what MPLS does not provide.

### Securing the MPLS Core

This section is targeted toward the SP, to give guidelines about secure configuration of an MPLS core network. Many general security mechanisms, such as securing routers, are not discussed here. More information can be found at the *Site Security Handbook* [RFC2196], and *Recommended Internet Service Provider Security Services and Procedures* [RFC3013]. The following is a list of recommendations and considerations on configuring an MPLS network securely.

- **Trusted devices**—The PE and P devices, as well as remote-access servers and authentication, authorization, and accounting (AAA) servers, have to be treated as trusted systems. This requires strong security management, starting with physical building security and including issues such as access control, secure configuration management, and storage. Ample literature is available on how to secure network elements, so this topic is not treated here in more detail. CE routers are typically not under full control of the SP and, therefore, have been treated as untrusted.
- **CE/PE interface**—The interface between the PE and CE routers is crucial for a secure MPLS network. The PE router should be configured as close as possible. From a security point of view, the best option is to configure the interface to the CE router unnumbered, and route statically.

Packet filters (ACLs) should be configured to permit only one specific routing protocol to the peering interface of the PE router, and only from the CE router. All other traffic to the router and the internal SP network should be denied. This scenario prevents attack on the PE and P routers, because the PE router will drop all packets to the corresponding address range. The only exception is the peer interface on the PE router for routing purposes. This needs to be secured separately.

If private address space [RFC1918] is used for the PE and P routers, the same rules with regard to packet filtering apply: All packets must be filtered to this range. However, because addresses of this range should not be routed over the Internet, attacks to adjacent networks are limited.

- **Routing authentication**—Routing is the signaling mechanism between the CEs and the PEs. To introduce bogus information into the core, routing protocols are the most obvious point for an attack. Thus it is essential that routing information is as secure as possible, and that it comes really from the router it is expected from, and not from a hacker's router. Toward this goal, all routing protocols should be configured with the corresponding authentication option toward the CEs and toward any Internet connection (specifically: BGP [RFC2385], OSPF [RFC2154], and RIP2 [RFC2082]). All peering relationships in the network need to be secured this way: CE/PE (with BGP MD5 authentication), PE/P (with Label Distribution Protocol [LDP] MD5 authentication) and P/P. This setup prevents attackers from spoofing a peer router and introducing bogus routing information. Note specifically here the importance of secure management: Configuration files often contain shared secrets in cleartext (for example, for routing protocol authentication).
- **Separation of CE/PE links**—If several CEs share a common Layer 2 infrastructure to access the same PE router (for example, an Ethernet virtual LAN [VLAN]), a CE router can spoof packets as belonging to another VPN that also has a connection to this PE router. Securing the routing protocol as described above is not sufficient, because this does not affect normal packets. To avoid this problem, it is recommended to implement separate physical connections between CEs and PEs. The use of a switch between various CE routers and a PE router is also possible, but it is strongly recommended to put each CE/PE pair into a separate VLAN, to provide traffic separation. Note that although switches with VLANs increase security, they are not unbreakable. A switch in this environment must thus be treated as a trusted device, and configured with maximum security.



- LDP authentication—The LDP can also be secured with MD5 authentication across the MPLS cloud. This scenario prevents hackers from introducing bogus routers, which would participate in the LDP.

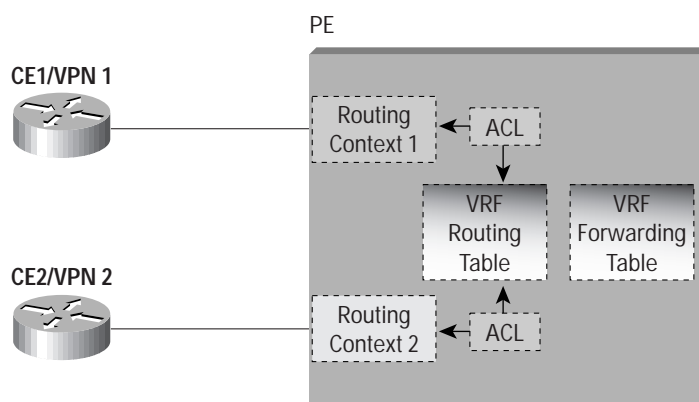
Note: The security of the overall solution depends on the security of its weakest link. This could be the weakest single interconnection between a PE and a CE, an insecure access server, or an insecure Trivial File Transfer Protocol (TFTP) server.

## Interconnections between VPNs and Internet Access

### Connectivity between VPNs

MPLS provides VPN services with address and routing separation between VPNs. In many environments, however, destinations outside the VPN must also be reachable. This could be for Internet access or for merging two VPNs, for example, in the case of two companies merging. MPLS not only provides full VPN separation, but also allows merging VPNs or access to the Internet.

Figure 3 Connectivity between VPNs



To achieve this access, the PE routers maintain various tables: A *routing context* is specific to a CE router, and contains only routes from this particular VPN. From there, routes are propagated into the *VRF* (virtual routing and forwarding instance) *routing table*, from which a *VRF forwarding table* is calculated. For separated VPNs, the VRF routing table contains only routes from one routing context. To merge VPNs, different routing contexts (from different VPNs) are put into one single VRF routing table. This way, two or several VPNs can be merged to a single VPN. Note that in this case all merged VPNs must have mutually exclusive addressing spaces; in other words, the overall address space must be unique for all included VPNs. It is possible to control with ACLs which routes get redistributed into VRF tables.

For a VPN to have Internet connectivity, the same procedure is used: Routes from the Internet VRF routing table (the default routing table) are propagated into the VRF routing table of the VPN that requires Internet access. Alternatively to propagating all Internet routes, a default route can be propagated. In this case, the address space between the VPN and the Internet must be distinct. In other words, the VPN must use either publicly registered or private address space [RFC1918], because all other addresses can occur in the Internet.

From a security point of view, the merged VPNs behave like one logical VPN, and the security mechanisms described above apply now between the merged VPN and other VPNs: The merged VPN must have unique address space internally, but further VPNs may use the same address space without interference. Packets from and to the merged VPNs cannot be routed to other VPNs, so all the separation functions of MPLS apply also for merged VPNs with respect to other VPNs.

If two VPNs are merged in this way, hosts from either part can reach the other part as if the two VPNs were a common VPN. This means that with the standard MPLS features, there is no separation or firewalling/packet filtering between the merged VPNs. Also, if a VPN receives Internet routes through MPLS/BGP VPN mechanisms, firewalling or packet filtering has to be engineered in addition to the MPLS features.

## 5.2 Firewalling Options

Two scenarios are examined in this section: securing VPNs against each other while maintaining inter-VPN connectivity, and securing Internet access.

### Scenario 1: Firewalls between VPNs

One reason for merging two previously independent VPNs is two companies merging or interoperating over the network. In most of these cases, the companies want to maintain a logical separation from *other* companies, even if connectivity between the companies is required. Typically, firewalls are placed in such circumstances. As in traditional networks, the interconnection points between the two VPNs have to be secured with firewalls. However, whereas in traditional networks the border router is normally under the control of the company, in the MPLS/BGP VPN environment, the “peering point” between the VPNs is a PE router under the control of the SP.

Technically, the interconnection by announcing the routes of both VPNs to the other VPN as described above happens in one router. This way of interconnecting alone does not provide firewall capabilities. To position a firewall between two VPNs, the firewall must be provisioned as a separate entity in addition to the PE router. The PE router manages the two VPNs completely separate, as described above. This setup provides the required security between the two VPNs.

To interconnect the two VPNs via a firewall, an additional interface that leads to the firewall must be provisioned for each VPN. This way, packets from VPN A to VPN B would come from a router in VPN A, and they would be routed to the interconnecting PE router. The PE router has a route to VPN B, which points to the interface to which the firewall is connected. The packets traverse the firewall, and enter the PE router through another interface, which belongs to VPN B. This way, it is also possible to use NAT on the firewall, with the effect that the merged VPNs do not have to have mutually exclusive address space.

The note on “Separation of CE/PE links” in Section 4 also applies here: Switches do not necessarily provide traffic separation. Thus if switches are used, it is strongly recommended not to put the interfaces of the firewall onto the same switch, but to use separate switches. If this is not possible, different VLANs must be used for the two sides of the firewall. Because hubs do not provide any traffic separation, their use is strongly discouraged.

There can be more than one interconnection point between VPNs. All interconnection points can be engineered this way.

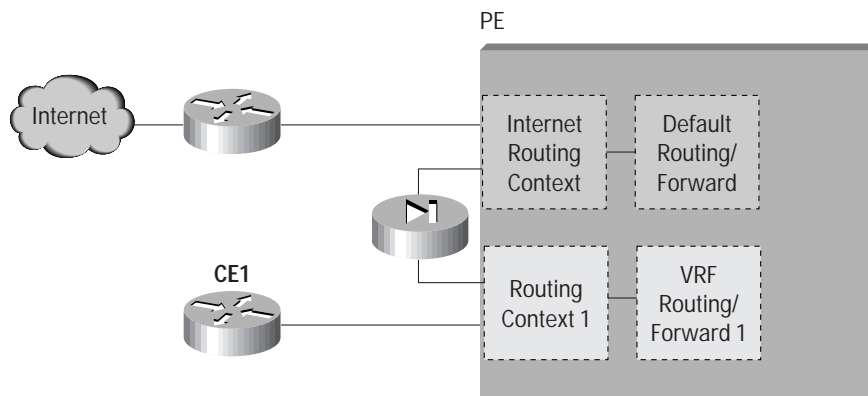
### Scenario 2: Firewalls to the Internet

The provisioning of a firewall for Internet access is similar: The PE router that connects to one or more other SPs will have to traverse a firewall before sending or receiving packets to/from the Internet. If one firewall can be applied to all VPNs equally (shared firewall), the setup consists of a PE router, which connects to a firewall before going to other SPs. The PE router maintains in the default VRF routing table the Internet routes or a default route to the Internet. The Internet routes (or the default) are propagated to the VRF routing tables of VPNs that require Internet connectivity. The routes from the VPNs are propagated to the default VRF routing table of the PE router, which announces them to the Internet over the firewall. Instead of dynamic routing, static routes can also be used. The addressing space of VPNs using Internet connectivity must be publicly registered address space.

Figure 4 shows one possible way to secure Internet access with a firewall of choice. The Internet routing table is treated as another VPN, and the connectivity to other VPNs is passing through an external firewall, providing all the features of this firewall, including NAT if required.



Figure 4 Example for a Firewall Installation to the Internet



This option is relatively easy to engineer, but has the disadvantage that all VPNs use the same firewall and are thus bound to one security policy at this point. To engineer an Internet firewall for each VPN separately, the above setup needs to be multiplied.

For separate firewalls per VPN, the PE router that connects to the Internet needs one interface per VPN, leading to one firewall per VPN. Beyond the firewalls, the connections can come together again in one router, which connects then to other providers. The advantage of this option is the capability to have a NAT function per customer, so that internally each VPN can use random address space, and on the firewalls this is mapped to publicly registered space.

### Scenario 3: A Firewall per CE Router

Big networks tend to become unmanageable in terms of security, unless there is some form of separation between parts of the network. In a country-wide network that is internally completely open, a security incident such as a break-in in one office might require all hosts of the entire network to be reinstalled, to ensure that the attacker has not left some Trojan horses somewhere. An increasing number of companies are securing their internal networks additionally by, for example, separating offices with firewalls. This is, in general, a good security practice. Given that Cisco routers can function as a firewall, the additional costs are normally manageable, because often only a software upgrade is required.

In the case where a company separates its network, implemented as a VPN on an MPLS service, putting a firewall on every CE router can make the overall network easier to manage and more secure. In this case, everything outside an office (connected with a CE router to the MPLS network) is treated as untrusted, and traffic from the same company is checked as well as traffic from other companies, which might be merged over the MPLS structure.

### Combining IPSec and MPLS

From a customer perspective, it is impossible to control the whole network; the SP must be trusted to some extent. If the MPLS core is not properly configured with the necessary security measures, the connected VPNs will be exposed to some forms of attack. IPSec offers additional security over an MPLS network.

IPSec can be run on the CE routers, or on devices further away from the core. If the CE router is under control of the customer, this could be an obvious choice. If the SP controls the CEs as part of the service, the customer has to decide whether to trust the SP to configure IPSec for him/her on the CE routers, or whether to maintain control over the IPSec in additional equipment outside the SP's scope.

IPSec should be used on top of an MPLS infrastructure if the customer does not want to trust the SP's network to be fully secure, or if he/she wants his/her security to be independent of potential problems in the core network. Specifically, IPSec should be used if one or several of the following requirements exist:

- *Encryption* of parts or all traffic over the MPLS core—If an attacker is able to sniff traffic on the core, with IPSec he/she will be able to see only the site from which the traffic came, and the site to which it goes. Note that for application-specific encryption such as secure Web transactions, other forms of encryption such as transport layer security (TLS) might be more appropriate. Also, for example, Telnet can be replaced by Secure Shell Protocol (SSH).
- *Authentication* of the endpoints—The crypto-endpoints (probably the CE routers) can authenticate each other, so that an attacker cannot introduce another router in the VPN, even if the SP's MPLS core is not fully secured. This setup can provide authentication for all protocols between the routers, specifically routing protocols, but also general traffic. Note that MPLS on its own can provide authentication throughout the network, but on a hop-by-hop basis. The added value of IPSec is that there is a direct authentication between the remote routers, so that even misconfigurations in the MPLS core cannot endanger the security of the customer's network.
- *Integrity* of the traffic—Packets cannot be changed on their way through the core without the change being noticed. Bogus packets cannot be introduced from the core.
- *Replay detection*—If IPSec Authentication Header (AH) is used, an attacker cannot save a packet flow and replay it later. This can be crucial in application environments where simple messages such as “close connection” exist, which could be saved by an attacker and later be used for a DoS attack against this service.

IPSec running on CE routers or behind can be seen as an overlay network over the MPLS cloud: The MPLS cloud is not aware of the IPSec layer, and vice versa, and the two layers do not influence each other. Thus to apply IPSec over MPLS, all the general rules to build IPSec networks apply. This paper does not discuss IPSec design in detail, but refers to specific IPSec literature. The general topology options that are available with IPSec include the following:

- *Point to point*—In this case, each pair of IPSec-aware routers has the other end of the IPSec connection statically configured with IP address and further parameters. This approach is easy to engineer, but difficult to deploy in larger environments, because a large number of one-to-one connections need to be statically configured. This option is advisable only if the number of IPSec routers is very small; otherwise the configuration effort becomes too complex. This setup is technically achieved by configuring *static cryptomaps*.
- *Hub and spoke*—Many networks such as banks' networks consist naturally of a design with few (typically one to two) central sites, and many remote sites. In such an environment, it is often easier to configure only the remote offices statically to the one or two central sites. The central office can use in this case a *dynamic cryptomap* to accept any connection from remote offices, so that it is not necessary to add static configurations per remote office into the central routers.
- *Full mesh*—A full mesh of a large number of IPSec routers would be extremely complex to deploy with static configurations, because the total number of peering relationships becomes quickly too large to manage. For this case a specific protocol is available, called *Tunnel Endpoint Discovery (TED)*. With this technology, it is not necessary to have any static IPSec peering configuration in the routers, making deployment very easy. In this mode, IPSec tunnels will be automatically established between routers.

MPLS networks on their own provide a high level of security, comparable to existing ATM or Frame Relay networks. They do not, however, provide encryption. IPSec can further increase the security of the customer's network by not putting any trust on the SP's infrastructure and handling all security relevant functions outside the core network. MPLS and IPSec together provide a very high level of security for VPNs.



## Remote Access to an MPLS VPN

In the typical Intranet VPN scenario, an MPLS VPN interconnects static devices, which are attached to a defined port, and have a defined addressing. In the case of remote access, neither the location nor the address of the connecting device is known deductively. Remote access can be provided in various ways; for example, over dialup, using IPSec, or both.

From a technical point of view, remote access to an MPLS VPN can be provided in two ways: internally to a company network, and thus invisible to the MPLS service, or as part of the shared infrastructure of the MPLS core. Commercially, there is merit in sharing the remote-access solution among customers. Sharing remote-access solutions, however, requires a clear understanding of the security of the setup. This section focuses on how to securely map a user into a VPN and to assign him/her an address (and other parameters) of this VPN.

The main security considerations for remote access to a shared MPLS network follow:

- Mapping of user into a VPN
- Location of the AAA server (shared or per VPN)
- Security of the connection of a remote-access server to the MPLS cloud

The mapping of the user to a VPN is the first step, and it is crucial for the security of the overall MPLS VPN infrastructure. If a user can pretend to belong to another VPN, all other security measures discussed previously are void. The assignment of users to an organization and thus a VPN can be done in several ways. The user can add the organization manually at the time of connection setup (user@domain), or the organizational information can be deduced from a certificate.

The AAA [RFC2903] server interfaces the remote-access server. It holds all the user-related information. This server could be shared, or each company could have its own AAA server. From a security point of view, both options are similar, assuming that the customer trusts the SP.

The mapping of user to VPN must be made before the AAA server is consulted, because the AAA server might be part of a company VPN and thus the access server has to know in advance which AAA server to ask. The communication between the AAA server and the remote-access server is implemented typically using the TACACS+ or Remote Access Dial-In User Service (RADIUS) protocols. Both provide security so that the AAA server cannot be spoofed.

If the remote-access server is a PE device (that is, if it participates in the label distribution and attaches labels to incoming packets), according to the previous steps, then the security of this device can be treated as the security of a PE router in the cases above. If the remote-access server is not a PE, the interconnection between the server and the PE device becomes crucial to security. The access server has to be a trusted device, because the PE device has to rely on correct VPN assignment of traffic by the access server. The PE also must be a trusted device, as described in the previous section. The interconnection between the two must be engineered such that no third party can possibly interfere in the communications between access server and PE. For example, if the server and PE are interconnected by an Ethernet connection, it must be a dedicated Ethernet connection. The PE and access server must either be colocated in a physically secured SP environment, or the communications between the two must be further secured with IPSec. In general, the access server must be part of the trusted system, which also includes the PE and P routers.

The type of remote access (standard dialup, IPSec, and so on) is generally not of concern to the overall MPLS solution. It should be noted, however, that unsecured access imposes an increased risk for intrusions, because it is, for example, possible to steal TCP sessions over a shared infrastructure. Thus an already-established connection between a remote client and the access server could be used by a hacker to gain unauthorized access. It is highly recommended to use IPSec for remote access in security-aware environments.

In a correct configuration, remote users cannot access a VPN other than their own. Thus adding a secured remote-access solution does not break the overall security of MPLS architecture. However, a DoS attack against the remote-access server could be possible. This scenario is not discussed here, because it is highly device dependent.

## What MPLS Does Not Provide

### Protection against Misconfigurations of the Core and Attacks “within” the Core

The security mechanisms discussed here assume correct configuration of the involved network elements on the MPLS core network (PE and P routers). Deliberate or inadvertent misconfigurations from SP staff may result in undesired behavior, including severe security leaks.

Note that this paragraph refers specifically to the core network; that is, the PE and P elements. Misconfiguration of any of the customer-side elements such as the CE router *is* covered by the security mechanisms above, meaning that a potential attacker must have access to either PE or P routers to gain advantage from misconfigurations. If an attacker has access to core elements or is able to insert additional equipment into the core, he/she will be able to attack both the core network and the connected VPNs. Thus the following is important:

- To avoid the risk of misconfigurations, it is important that the equipment is easy to configure, and that SP staff has the appropriate training and experience when configuring the network.
- To avoid the risk of “internal” attacks, the MPLS core network must be properly secured. This security includes network-element security, management security, physical security of the SP infrastructure, access control to SP installations, and other standard SP security mechanisms.

MPLS can provide a secure service only if the core network is provided in a secure fashion. This paper assumes that it is.

### Data Encryption, Integrity, and Origin Authentication

MPLS itself does not provide encryption, integrity, or authentication services. If these features are required, IPSec should be used over the MPLS infrastructure as described in Section 6.

### Customer Network Security

MPLS can be secured so that it is comparable with other VPN services. However, the security of the core network is only one factor for the overall security of a customer’s network. Threats in today’s networks come not only from the “outside” connection, but also from the “inside” and from other entry points (modems, for example). To reach a good security level for a customer network in an MPLS infrastructure, MPLS security is necessary but not sufficient. (See also [RFC2196] for more information on how to secure a network.)

## Summary and Conclusions

This section summarizes the findings of the previous sections, to give an overview of the various deployment options. It should be noted that security has many subtle degrees, which cannot all be discussed here. Thus these guidelines are only approximate. For each deployment option, the potential security risks are outlined.

All options below are based on an MPLS core network with VPN services. The basic assumption for all the scenarios is that the MPLS core is configured and managed in a secure fashion.

### Summary of Configuration Options

#### Option 1: Dynamic versus Static Routing between CEs and PEs

Routing protocols between CEs and PEs must be secured with the appropriate authentication mechanisms [RFC2082, RFC2154, RFC2385] to ensure that only the CE router can send routing updates. Furthermore, the routing protocols must be further secured from the SP side in order to not be vulnerable to routing attacks (malicious or inadvertent). For example, in BGP, it is possible to configure dampening parameters, where only a limited number of routing updates are accepted in a period of time.



Even with these precautions, an attacker cannot be prevented from finding a way to flood the router with bogus routing messages. The existence of a peer IP address might be enough to prevent this type of attack. Flooding the PE router from a CE can not break security as far as the MPLS mechanisms are concerned, but might have an effect on router performance, which might in turn influence performance of other VPNs.

However, practically, it must be considered that in this scenario the potential attack can come only from a legal user of the VPN service. This type of attack is easily traceable to one port—and thus one VPN customer. It can easily be stopped by shutting down that particular interface.

If there is no routing protocol running between CEs and PEs, static routing is required. If this routing is configured on an unnumbered link, just pointing to an interface rather than to a peer IP address, the CE does not need to know any addressing information of the MPLS core. Sending any type of message to the PE router will not have an effect, because it will be treated within the VRF. An attacker could still guess the address space of the router (for example, the loopback address), but it can be protected with ACLs that do not permit any packet, because no communication is required with this address.

In this scenario, the security is very high and fully comparable to similar Layer 2 services (FR, ATM).

#### Option 2: Internet Service

As long as MPLS/BGP VPNs are not connected to the Internet or other VPNs, MPLS provides a high level of security. In the case of Internet access through the MPLS network, all the rules of accessing the Internet in general apply. Most important, a firewall should be placed between the VPN and the Internet. The various options are described above. If configured correctly, Internet access over MPLS can be offered in a secure manner.

The same applies to various VPNs that are merged on the MPLS network. MPLS itself does not provide firewalling mechanisms, but an MPLS core can be engineered such that firewalls secure VPNs but allow connectivity.

#### Option 3: Running IPSec over the MPLS Cloud

If the security of the SP MPLS network is considered insufficient, there is the additional option to run IPSec on the CE routers or behind or over the MPLS cloud, with encryption (Encapsulating Security Protocol [ESP]) and authentication (AH). In this case, even attacks within the MPLS cloud cannot break the security of the overall VPN. Traffic on the MPLS network can be traced back to only two routers; the content is not legible. Changes on existing packets as well as fake or spoofed packets will be detected by the IPSec AH mechanisms.

In the case of dynamic routing, the potential routing attacks as described above can still be carried out, so DoS from a neighbouring VPN might be possible. Static routing provides more security than dynamic routing, but in this case static routing can be easily configured also from the SP side, because the SP sees only packets between the CE routers. Thus on each PE router, one static route to the corresponding CE router is sufficient. Note, however, that security of a given VPN depends on the security of the overall MPLS service.

#### Option 4: Including the CE Router in the SP Management

All discussions so far have assumed that the interface between the customer and the SP is between the CE and PE routers. However, in reality, many existing service offerings include the CE router as SP-managed customer premises equipment (CPE). This setup has numerous consequences for security:

- The core network can now be completely hidden to the customer networks, because customers have no direct connection to the MPLS core. This setup improves security.
- In addition to the PE, the CE can now also be configured with strict ACLs that also control access to the PE router.

- The routing protocol between the CE and the PE is now under the control of the SP. Routing must also run between the customer's network and the CE, so there is still some potential for routing attacks. But because the CE can be secured with ACLs and additional routing security, the overall setup will be more secure.
- If IPsec is required, running IPsec from the CE routers means in this scenario giving control of encryption to the SP. Some organizations might prefer to keep control of the encryption, in which case the IPsec should be done on another device before the CE router.

## Conclusions

MPLS provides full address and routing separation as in traditional Layer 2 VPN services. It hides addressing structures of the core and other VPNs, and it is in today's understanding not possible from the outside to intrude into the core or other VPNs abusing the MPLS mechanisms. It is also not possible to intrude into the MPLS core if it is properly secured. However, there is a significant difference between MPLS-based VPNs and, for example, FR- or ATM-based VPNs: The control structure of the core is on Layer 3 in the case of MPLS. This fact has caused significant scepticism in the industry toward MPLS, because this setup might open the architecture to DoS attacks from other VPNs or the Internet (if connected). Table 2 compares ATM/FR with MPLS.

Table 2 Comparison between ATM/FR and MPLS

	ATM/FR	MPLS
Address Space Separation	Yes	Yes
Routing Separation	Yes	Yes
Resistance to Attacks	Yes	Yes
Resistance to Label Spoofing	Yes	Yes

As shown in this paper, it is possible to secure an MPLS infrastructure to the same level of security as a comparable ATM or FR service. It is also possible to offer Internet connectivity to MPLS VPNs in a secure manner, and to interconnect different VPNs via firewalls. Although ATM and FR services have a strong reputation with regard to security, it has been shown that security problems can also exist in these networks [DataComm].

With regard to attacks from within the MPLS core, all VPN classes (MPLS, FR, ATM) have the same problem: If an attacker can install a sniffer, he/she can read information in all VPNs, and if the attacker has access to the core devices, he/she can execute a large number of attacks, from packet spoofing to introducing a new peer router. Numerous precaution measures that an SP can use to tighten security of the core are outlined above, but the security of the MPLS architecture depends on the security of the SP. If the SP is not trusted, the only way to fully secure a VPN against attacks from the "inside" of the VPN service is to run IPsec on top, from the CE devices or beyond.

This paper discusses many aspects of MPLS security. It should be noted explicitly that the overall security of MPLS architecture depends on all components, and is determined by the security of the weakest part of the solution. For example, a perfectly secured static MPLS network with secured Internet access and secure management is still open to many attacks if there is a weak remote-access solution in place.

The Miercom group has tested the security of the Cisco MPLS/BGP VPN solution by probing a network in various ways to prove the points made in this paper. The final report [Miercom] shows that it was not possible to attack the network in any way, nor other VPNs on the same network. In addition, neither the address space nor the routing separation between VPNs could be overcome. The end result of the report is that MPLS is at least as secure as Frame Relay and ATM networks.



## References

- DataComm. *Frame Relay and ATM: Are they really secure?* Data Communications Report, Vol. 15, No. 4: February 2000. (<http://www.yankeegroup.com>)
- Miercom. *Cisco MPLS-Based VPNs: Equivalent to the Security of Frame Relay and ATM.* Miercom Report: March 2001. (<http://www.mier.com/reports/cisco/MPLS-VPNs.pdf>)
- Rekhter, Y. *Address Allocation for Private Internets.* RFC1918: February 1996. (<http://search.ietf.org/rfc/rfc1918.txt>)
- Baker, F. and R. Atkinson. *RIP-2 MD5 Authentication.* RFC2082: January 1997. (<http://search.ietf.org/rfc/rfc2082.txt>)
- Murphy, S., M. Badger, and B. Wellington. *OSPF with Digital Signatures.* RFC2154: June 1997. (<http://search.ietf.org/rfc/rfc2154.txt>)
- Fraser, B. *Site Security Handbook.* RFC2196: September 1997. (<http://search.ietf.org/rfc/rfc2196.txt>)
- Heffernan, A. *Protection of BGP Sessions via the TCP MD5 Signature Option.* RFC2385: August 1998. (<http://search.ietf.org/rfc/rfc2385.txt>)
- Rosen, E. and Y. Rekhter. *BGP/MPLS VPNs.* RFC2547: March 1999. (<http://search.ietf.org/rfc/rfc2547.txt>)
- Baker, F., B. Lindell, and M. Talwar. *RSVP Cryptographic Authentication.* RFC2747: January 2000. (<http://search.ietf.org/rfc/rfc2747.txt>)
- Ferguson, P. and D. Senie. *Network Ingress Filtering: Defeating Denial of Service Attacks which Employ IP Source Address Spoofing.* RFC2827: May 2000. (<http://search.ietf.org/rfc/rfc2827.txt>)
- Shirey, R. *Internet Security Glossary.* RFC2828: May 2000. (<http://search.ietf.org/rfc/rfc2828.txt>)
- de Laat, C., G. Gross, L. Gommans, J. Vollbrecht, and D. Spence. *Generic AAA Architecture.* RFC2903: August 2000. (<http://search.ietf.org/rfc/rfc2903.txt>)
- Killalea, T. *Recommended Internet Service Provider Security Services and Procedures.* RFC3013: November 2000. (<http://search.ietf.org/rfc/rfc3013.txt>)



Corporate Headquarters  
Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
www.cisco.com  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 526-4100

European Headquarters  
Cisco Systems Europe  
11, Rue Camille Desmoulins  
92782 Issy-les-Moulineaux  
Cedex 9  
France  
www-europe.cisco.com  
Tel: 33 1 58 04 60 00  
Fax: 33 1 58 04 61 00

Americas Headquarters  
Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
www.cisco.com  
Tel: 408 526-7660  
Fax: 408 527-0883

Asia Pacific Headquarters  
Cisco Systems Australia, Pty., Ltd  
Level 9, 80 Pacific Highway  
P.O. Box 469  
North Sydney  
NSW 2060 Australia  
www.cisco.com  
Tel: +61 2 8448 7100  
Fax: +61 2 9957 4350

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the

**Cisco Web site at [www.cisco.com/go/offices](http://www.cisco.com/go/offices)**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia  
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia  
Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru  
Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa  
Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2001, Cisco Systems, Inc. All rights reserved. Printed in the USA. Cisco, Cisco IOS, Cisco Systems, and the Cisco Systems logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0106R)

08/01 LW2498