



CHAPTER 2

MPLS VPN Concepts

This chapter provides a conceptual information useful for understanding MPLS. It contains the following sections:

- [MPLS VPNs, page 2-1](#)
- [MPLS VPN Security, page 2-8](#)

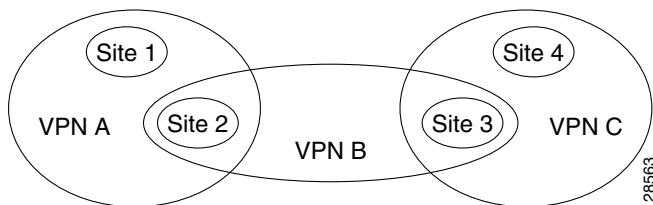
MPLS VPNs

At its simplest, a virtual private network (VPN) is a collection of sites that share the same routing table. A VPN is also a network in which customer connectivity to multiple sites is deployed on a shared infrastructure with the same administrative policies as a private network. The path between two systems in a VPN, and the characteristics of that path, might also be determined (wholly or partially) by policy. Whether a system in a particular VPN is allowed to communicate with systems not in the same VPN is also a matter of policy.

In an MPLS VPN, the VPN generally consists of a set of sites that are interconnected by means of an MPLS provider core network, but it is also possible to apply different policies to different systems that are located at the same site. Policies can also be applied to systems that dial in; the chosen policies would be based on the dial-in authentication processes.

A given set of systems can be in one or more VPNs. A VPN can consist of sites (or systems) that are all from the same enterprise (intranet), or from different enterprises (extranet); it might consist of sites (or systems) that all attach to the same service provider backbone, or to different service provider backbones.

Figure 2-1 **VPNs Sharing Sites**



MPLS-based VPNs are created in Layer 3 and are based on the peer model, which makes them more scalable and easier to build and manage than conventional VPNs. In addition, value-added services, such as application and data hosting, network commerce, and telephony services, can easily be targeted and deployed to a particular MPLS VPN because the service provider backbone recognizes each MPLS VPN as a secure, connectionless IP network.

The MPLS VPN model is a true peer VPN model that enforces traffic separations by assigning unique VPN route forwarding tables (VRFs) to each customer's VPN. Thus, users in a specific VPN cannot see traffic outside their VPN. Traffic separation occurs without tunneling or encryption because it is built directly into the network. (For more information on VRFs, see [VPN Routing and Forwarding Tables, page 2-3](#).)

The service provider's backbone is comprised of the PE and its provider routers. MPLS VPN provides the ability that the routing information about a particular VPN be present *only* in those PE routers that attach to that VPN.

Characteristics of MPLS VPNs

MPLS VPNs have the following characteristics:

- Multiprotocol Border Gateway Protocol (MP-BGP) extensions are used to encode customer IPv4 address prefixes into unique VPN-IPv4 Network Layer Reachability Information (NLRI) values. NLRI refers to a destination address in MP-BGP, so NLRI is considered "one routing unit." In the context of IPv4 MP-BGP, NLRI refers to a network prefix/prefix length pair that is carried in the BGP4 routing updates.
- Extended MP-BGP community attributes are used to control the distribution of customer routes.
- Each customer route is associated with an MPLS label, which is assigned by the provider edge router that originates the route. The label is then employed to direct data packets to the correct egress customer edge router. When a data packet is forwarded across the provider backbone, two labels are used. The first label directs the packet to the appropriate egress PE; the second label indicates how that egress PE should forward the packet.
- Cisco MPLS CoS and QoS mechanisms provide service differentiation among customer data packets.
- The link between the PE and CE routers uses standard IP forwarding.

The PE associates each CE with a per-site forwarding table that contains only the set of routes available to that CE.

Principal Technologies

There are four principal technologies that make it possible to build MPLS-based VPNs:

- Multiprotocol Border Gateway Protocol (MP-BGP) between PEs carries CE routing information.
- Route filtering based on the VPN route target extended MP-BGP community attribute.
- MPLS forwarding carries packets between PEs (across the service provider backbone).
- Each PE has multiple VPN routing and forwarding instances (VRFs).

Intranets and Extranets

If all the sites in a VPN are owned by the same enterprise, the VPN is a corporate *intranet*. If the various sites in a VPN are owned by different enterprises, the VPN is an *extranet*. A site can be in more than one VPN. Both intranets and extranets are regarded as VPNs.

While the basic unit of connection is the site, the MPLS VPN architecture allows a finer degree of granularity in the control of connectivity. For example, at a given site, it might be desirable to allow only certain specified systems to connect to certain other sites. That is, certain systems at a site might be members of an intranet and members of one or more extranets, while other systems at the same site might be restricted to being members of the intranet only.

A CE router can be in multiple VPNs, although it can only be in a single site. When a CE router is in multiple VPNs, one of these VPNs is considered its primary VPN. In general, a CE router's primary VPN is the intranet that includes the CE router's site. A PE router might attach to CE routers in any number of different sites, whether those CE routers are in the same or in different VPNs. A CE router might, for robustness, attach to multiple PE routers. A PE router attaches to a particular VPN if it is a router adjacent to a CE router that is in that VPN.

VPN Routing and Forwarding Tables

The VPN routing and forwarding table (VRF) is a key element in the MPLS VPN technology. VRFs exist on PEs only (except in the case of a Multi-VRF CE). A VRF is a routing table instance, and more than one VRF can exist on a PE. A VPN can contain one or more VRFs on a PE. The VRF contains routes that should be available to a particular set of sites. VRFs use Cisco Express Forwarding (CEF) technology, therefore the VPN must be CEF-enabled.

A VRF is associated with the following elements:

- IP routing table
- Derived forwarding table, based on the Cisco Express Forwarding (CEF) technology
- A set of interfaces that use the derived forwarding table
- A set of routing protocols and routing peers that inject information into the VRF

Each PE maintains one or more VRFs. Prime Fulfillment software looks up a particular packet's IP destination address in the appropriate VRF only if that packet arrived directly through an interface that is associated with that VRF. The so-called "color" MPLS label tells the destination PE to check the VRF for the appropriate VPN so that it can deliver the packet to the correct CE and finally to the local host machine.

A VRF is named based on the VPN or VPNs it services, and on the role of the CE in the topology. The schemes for the VRF names are as follows:

- The VRF name for a hub: `ip vrf vx:[VPN_name]`
- The *x* parameter is a number assigned to make the VRF name unique.

For example, if we consider a VPN called Blue, then a VRF for a hub CE would be called:

```
ip vrf V1:blue
```

A VRF for a spoke CE in the Blue VPN would be called:

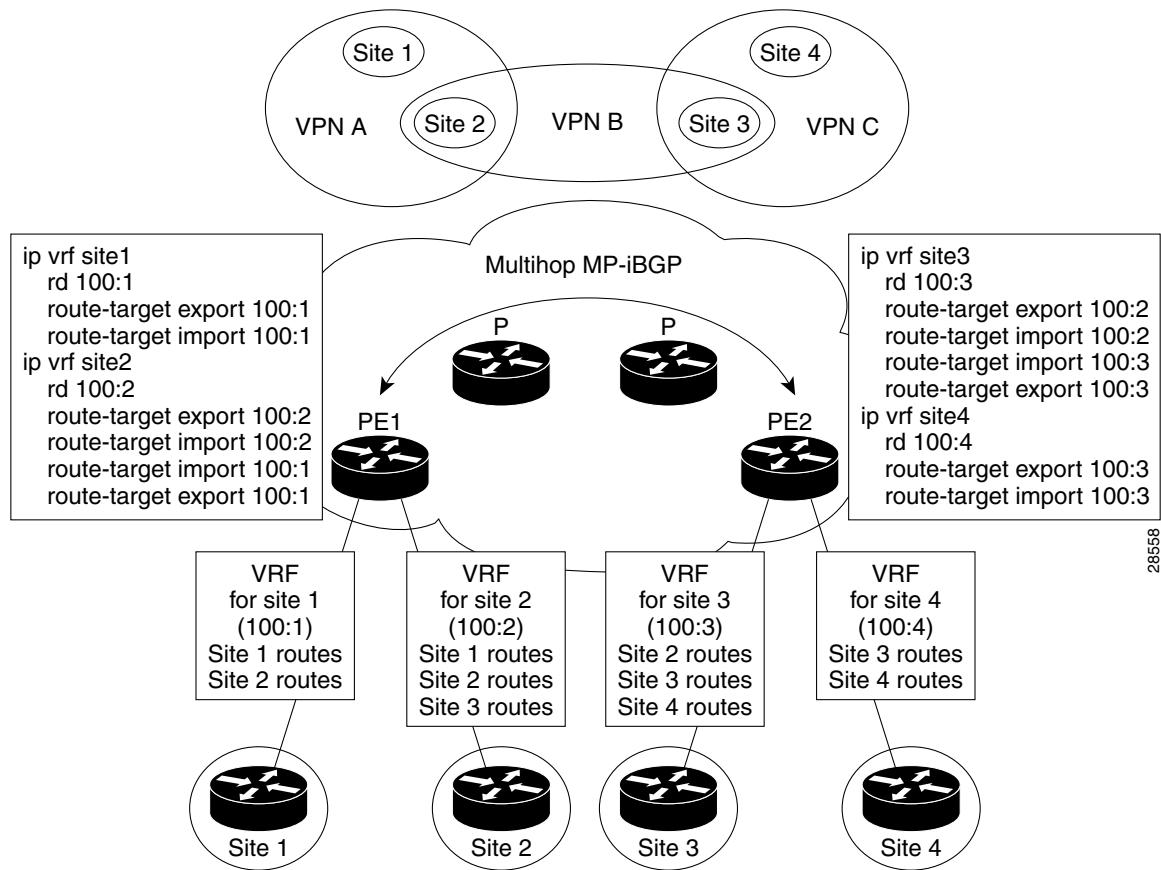
```
ip vrf V1:blue-s
```

A VRF for an extranet VPN topology in the Green VPN would be called:

```
ip vrf V1:green-etc
```

Thus, you can read the VPN name and the topology type directly from the name of the VRF.

[Figure 2-2](#) shows a network in which two of the four sites are members of two VPNs, and illustrates which routes are included in the VRFs for each site.

Figure 2-2 VRFs for Sites in Multiple VPNs

VRF Implementation

When implementing VPNs and VRFs, we recommend you keep the following considerations in mind:

- A local VRF interface on a PE is not considered a directly-connected interface in a traditional sense. When you configure, for example, a Fast Ethernet interface on a PE to participate in a particular VRF/VPN, the interface no longer shows up as a directly-connected interface when you issue a **show ip route** command. To see that interface in a routing table, you must issue a **show ip route vrf *vrf_name*** command.
- The global routing table and the per-VRF routing table are independent entities. Cisco IOS commands apply to IP routing in a global routing table context. For example, **show ip route**, and other EXEC-level **show** commands—and utilities such as **ping**, **traceroute**, and **telnet**—all invoke the services of the Cisco IOS routines that deal with the global IP routing table.

- You can issue a standard Telnet command from a CE router to connect to a PE router. However, from that PE, you must issue the following command to connect from the PE to the CE:
`telnet CE_RouterName /vrf vrf_name`
 Similarly, you can utilize the **Traceroute** and **Ping** commands in a VRF context.
- The MPLS VPN backbone relies on the appropriate Interior Gateway Protocol (IGP) that is configured for MPLS, for example, EIGRP, or OSPF. When you issue a **show ip route** command on a PE, you see the IGP-derived routes connecting the PEs together. Contrast that with the **show ip route vrf VRF_name** command, which displays routes connecting customer sites in a particular VPN.

VRF Instance

The configuration commands to create a VRF instance are as follows:

	Command	Description
Step 1	<code>Router# configure terminal</code> <code>Router(config)#</code>	Enter global configuration mode.
Step 2	<code>Router(config)# ip vrf vrf_name</code>	For example, <code>ip vrf CustomerA</code> initiates a VPN routing table and an associated CEF table named CustomerA. The command enters VRF configuration submode to configure the variables associated with the VRF.
Step 3	<code>Router(config-vrf)# rd RD_value</code>	Enter the eight-byte route descriptor (RD) or IP address. The PE prepends the RD to the IPv4 routes prior to redistributing the route into the MPLS VPN backbone.
Step 4	<code>Router(config-vrf)# route-target import export both community</code>	Enter the route-target information for the VRF.

Independent VRF Object Management

Prime Fulfillment allows you to specify VPN and VRF information in an independent VRF object, which is subsequently deployed to a PE device and then associated with an MPLS VPN link via an MPLS VPN service request. For details on using this feature, see to the *Cisco Prime Fulfillment User Guide 6.1*.

Route Distinguishers and Route Targets

MPLS-based VPNs employ BGP to communicate between PEs to facilitate customer routes. This is made possible through extensions to BGP that carry addresses other than IPv4 addresses. A notable extension is called the *route distinguisher* (RD).

The purpose of the route distinguisher (RD) is to make the prefix value unique across the backbone. Prefixes should use the same RD if they are associated with the same set of route targets (RTs) and anything else that is used to choose routing policy. The community of interest association is based on the route target (RT) extended community attributes distributed with the Network Layer Reachability Information (NLRI).The RD value must be a globally unique value to avoid conflict with other prefixes.

The MPLS label is part of a BGP routing update. The routing update also carries the addressing and reachability information. When the RD is unique across the MPLS VPN network, proper connectivity is established even if different customers use non-unique IP addresses.

For the RD, every CE that has the same overall role should use a VRF with the same name, same RD, and same RT values. The RDs and RTs are *only* for route exchange between the PEs running BGP. That is, for the PEs to do MPLS VPN work, they have to exchange routing information with more fields than usual for IPv4 routes; that extra information includes (but is not limited to) the RDs and RTs.

The route distinguisher values are chosen by the Prime Fulfillment software.

- CEs with hub connectivity use `bgp_AS:value`.
- CEs with spoke connectivity use `bgp_AS:value + 1`

Each spoke uses its own RD value for proper hub and spoke connectivity between CEs; therefore, the Prime Fulfillment software implements a new RD for each spoke that is provisioned.

Prime Fulfillment chooses route target values by default, but you can override the automatically assigned RT values if necessary when you first define a Route Target in the Prime Fulfillment software.

Route Target Communities

The mechanism by which MPLS VPN controls distribution of VPN routing information is through the VPN route-target extended MP-BGP communities. An extended MP-BGP community is an eight octet structure value. MPLS VPN uses route-target communities as follows:

- When a VPN route is injected into MP-BGP, the route is associated with a list of VPN route-target communities. Typically, this is set through an export list of community values associated with the VRF from which the route was learned.
- An import list of route-target communities is associated with each VRF. This list defines the values that should be matched against to decide whether a route is eligible to be imported into this VRF.

For example, if the import list for a particular VRF is {A, B, C}, then any VPN route that carries community value A, B, or C is imported into the VRF.

Route Targets

A VPN can be organized into subsets called Route Targets. A Route Target describes how the CEs in a VPN communicate with each other. Thus, Route Targets describe the logical topology of the VPN.

Prime Fulfillment can be employed to form a variety of VPN topologies between CEs by building hub and spoke or full mesh Route Targets. Route Targets are building blocks that allow you to form complex VPN topologies and CE connectivity.

The most common types of VPNs are *hub-and-spoke* and *full mesh*.

- A hub-and-spoke Route Target is one in which one or a few CEs act as hubs, and all spoke CEs talk only to or through the hubs, never directly to each other.
- A full mesh Route Target is one in which every CE connects to every other CE.

These two basic types of VPNs—full mesh and hub and spoke—can be represented with a single Route Target.

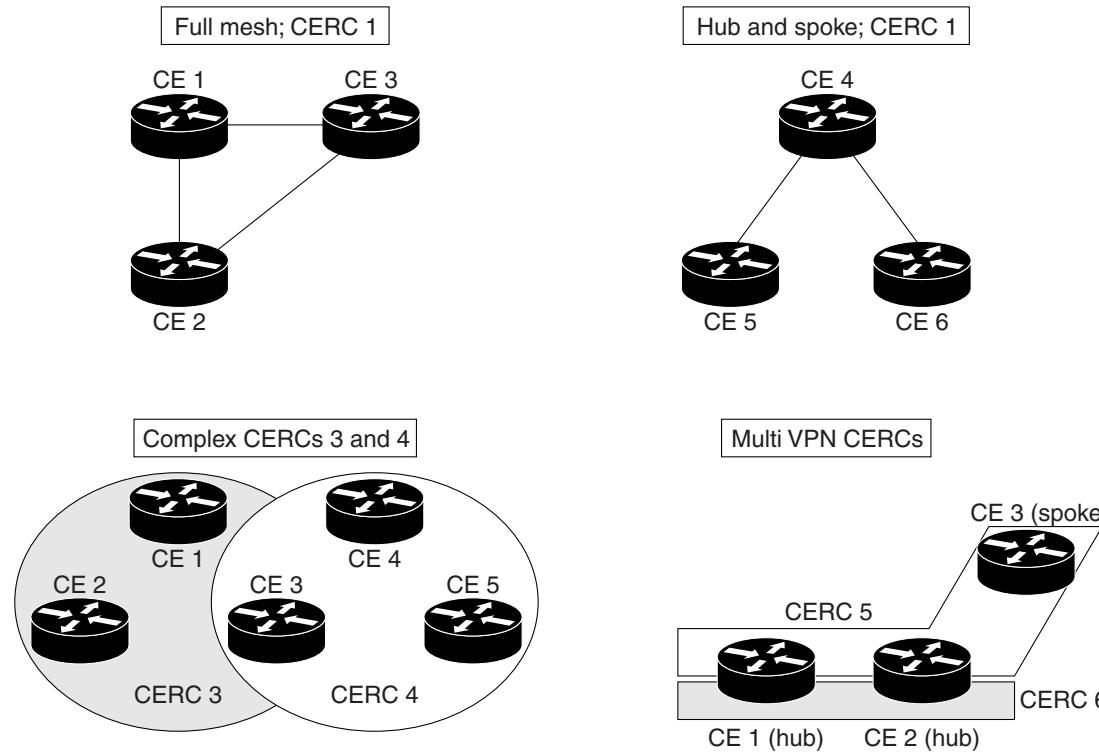
Whenever you create a VPN, the Prime Fulfillment software creates one default Route Target for you. This means that until you need advanced customer layout methods, you will not need to define new Route Targets. Up to that point, you can think of a Route Target as standing for the VPN itself—they are one and the same. If, for any reason, you need to override the software's choice of route target values, you can do so only at the time you create a Route Target in the Prime Fulfillment software.

To build very complex topologies, it is necessary to break down the required connectivity between CEs into groups, where each group is either fully meshed, or has a hub and spoke pattern. (Note that a CE can be in more than one group at a time, so long as each group has one of the two basic patterns.) Each subgroup in the VPN needs its own Route Target. Any CE that is only in one group just joins the corresponding Route Target (as a spoke if necessary). If a CE is in more than one group, then you can use the Advanced Setup choice during provisioning to add the CE to all the relevant groups in one service request. Given this information, the provisioning software does the rest, assigning route target values and VRF tables to arrange exactly the connectivity the customer requires. You can use the Topology tool to double-check the Route Target memberships and resultant VPN connectedness.

Prime Fulfillment supports multiple CEs per site and multiple sites connected to the same PE. Each Route Target has unique route targets (RT), route distinguisher (RD) and VRF naming. After provisioning a Route Target, it is a good idea to run the audit reports to verify the Route Target deployment and view the topologies created by the service requests. The product supports linking two or more Route Targets in the same VPN.

[Figure 2-3](#) shows several examples of the topologies that Prime Fulfillment Route Targets can employ.

Figure 2-3 Examples of Route Target Topologies



Hub and Spoke Considerations

In hub-and-spoke MPLS VPN environments, the spoke routers have to have unique Route Distinguishers (RDs). In order to use the hub site as a transit point for connectivity in such an environment, the spoke sites export their routes to the hub. Spokes can talk to hubs, but spokes never have routes to other spokes.

Due to the current MPLS VPN implementation, you must apply a different RD for each spoke VRF. The MP-BGP selection process applies to all the routes that have to be imported into the same VRF plus all routes that have the same RD of such a VRF. Once the selection process is done, only the best routes are imported. In this case this can result in a best route which is not imported. Thus, customers must have different RDs per spoke-VRF.

Full Mesh Considerations

Each Route Target has two distinct RTs, a hub RT and a spoke RT. When building a full mesh topology, always use the hub RT. Thus, when a need arises to add a spoke site for the current full mesh topology, you can easily add the spoke site without reconfiguring any of the hub sites. The existing spoke RT can be used for this purpose. This is a strategy to prevent having to do significant reprovisioning of a full mesh topology to a hub-and-spoke topology.

MPLS VPN Security

This section discusses the security requirements for MPLS VPN architectures. This section concentrates on protecting the core network against attacks from the “outside,” that is, the Internet and connected VPNs.



Note Protection against attacks from the “inside,” that is, when an attacker has logical or physical access to the core network is not discussed here, since any network can be attacked with access from the inside.

Address Space and Routing Separation

Between two non-intersecting 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 non-intersecting VPNs must be able to both 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.

Address Space Separation

From a security point of view, the basic requirement is to avoid that 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.

MPLS allows distinct VPNs to use the same address space, which can also be private address space. 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*. Thus customers of an MPLS service do not need to change current addressing in their networks.

In the case of using routing protocols between CE and PE routers (for static routing this is not an issue), there is one exception—the IP addresses of the PE routers the CE routers are peering with. To be able to communicate to the PE router, routing protocols on the CE routers must configure the address of the peer router in the core. This address must be unique from the CE router’s perspective. In an environment where the service provider manages also the CE routers as CPE (customer premises equipment), this can be made invisible to the customer.

Routing Separation

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. Since every VPN results in a separate VRF, there are no interferences between the VPNs on the PE router.

Across the MPLS core to the other PE routers, this routing separation is maintained by adding unique VPN identifiers in multi-protocol BGP, such as the route distinguisher (RD). VPN routes are exclusively exchanged by MP-BGP across the core, and this BGP information is not redistributed to the core network, but 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 addressing and routing separation across an MPLS core network, 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 core, unless this has been configured specifically.

Hiding the MPLS Core Structure

The internal structure of the MPLS core network (PE and Provider router devices) should not be visible to outside networks (either the Internet or any connected VPN). While a breach of this requirement does not lead to a security problem itself, it is generally advantageous when the internal addressing and network structure remains hidden to the outside world. 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.

Denial-of-service attacks against a core router, for example, are much easier to carry out if an attacker knows the IP address. Where addresses are not known, they can be guessed, but when the MPLS core structure is hidden, attacks are more difficult to make. Ideally, the MPLS core should be as invisible to the outside world as a comparable Layer 2 infrastructure (for example, Frame Relay or ATM).

In practice, a number of additional security measures have to be taken, most of all *extensive packet filtering*. MPLS does not reveal unnecessary information to the outside, not even to customer VPNs. The addressing in the core can be done with either private addresses or public addresses. Since the interface to the VPNs, and potentially to 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 a PE and CE is the address of the PE router. If this is not desired, you can configure static routing between the PE and CE. With this measure, the MPLS core can be kept completely hidden.

To ensure reachability across the MPLS cloud, customer VPNs will have to advertise their routes as a minimum to the MPLS core. While this could be seen as too open, the information known to the MPLS core is not about specific hosts, but networks (routes); this offers some degree of abstraction. Also, in a VPN-only MPLS network (that is, no shared Internet access), this is equal to existing Layer 2 models, where the customer has to trust the service provider to some degree. Also in a Frame Relay or ATM network, routing information about the VPNs can be seen on the core network.

In a VPN service with shared Internet access, the service provider typically announces the routes of customers that want to use the Internet to his upstream or peer providers.

In summary, in a pure MPLS VPN service, where no Internet access is provided, the level of information hiding is as good as on a comparable Frame Relay or ATM network—no addressing information is revealed to third parties or the Internet. If a customer chooses to access the Internet by way of the MPLS core, he will have to reveal the same addressing structure as for a normal Internet service.

If an MPLS network has no interconnections to the Internet, this is equal to Frame Relay or ATM networks. With Internet access from the MPLS cloud, the service provider 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

It is not possible to directly intrude into other VPNs. However, it is possible to attack the MPLS core, and try to attack other VPNs from there. There are two basic ways the MPLS core can be attacked:

- Attacking the PE routers directly.
- Attacking the signaling mechanisms of MPLS (mostly routing)

There are two basic types of attacks: *denial-of-service (DoS) attacks*, where resources become unavailable to authorized users, and *intrusion attacks*, where the goal is to gain unauthorized access to resources.

For intrusion attacks, give unauthorized access to resources, there are two basic ways to protect the network:

- Harden protocols that could be abused (for example, Telnet to a router)
- Make the network as inaccessible as possible. This is achieved by a combination of filtering packets and hiding the IP addresses in the MPLS core.

Denial-of-service attacks are easier to execute, since in the simplest case, a known IP address might be enough to attack a machine. The only way to be certain that you are not be vulnerable to this kind of attack is to make sure that machines are not reachable, again by packet filtering and pinging IP addresses.

MPLS networks must provide at least the same level of protection against both forms of attack as current Layer 2 networks provide.

To attack an element of an MPLS network it is first necessary to know this element, that is, its IP address. It is possible to hide the addressing structure of the MPLS core to the outside world, as discussed in the previous section. Thus, an attacker does not know the IP address of any router in the core that he wants to attack. The attacker could guess addresses and send packets to these addresses. However, due to the address separation of MPLS, each incoming packet is treated as belonging to the address space of the customer. It is therefore impossible to reach an internal router, even through guessing the IP addresses. There is only one exception to this rule—the peer interface of the PE router.

Securing the Routing Protocol

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 (usually a default route).

The static route can point to the IP address of the PE router, or to an interface of the CE router (for example, serial0).

Although in the static case the CE router does not know any IP addresses of the PE router, it is still attached to the PE router by way of some method, and could guess the address of the PE router and try to attack it with this address.

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, it is preferable to the other cases.

2. **Dynamic.** A routing protocol (for example, RIP, OSPF, or BGP) is used to exchange the routing information between the CE and the PE at each peering point.

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.

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 (ACLs). This limits the point of attack to one routing protocol, for example BGP. A potential attack could send an extensive number of routes, or flood the PE router with routing updates. Both of these attacks could lead to a denial-of-service attack, however, not to an intrusion attack.

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:

- Use VRFs. There are mechanisms within the context of a VRF for a service provider to monitor and control the number of routes that a customer can have in the VPN. When such thresholds are breached, for example 80 percent of the allowed number of routes syslog messages can be generated indicating to the service provider that the VRF is reaching the allowed limit.
- Use ACLs. 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 PE interface.

ACLs must be configured to limit access only to the port(s) of the routing protocol, and only from the CE router.

- Where available, configure MD-5 authentication for routing protocols.

This is available for BGP, OSPF, and RIP2. It avoids the possibility that packets could be spoofed from other parts of the customer network than the CE router. This requires that the service provider 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 only for the customer with the highest security requirements.



Note

Prime Fulfillment does not provide for the provisioning of MD5 authentication on PE-CE links using routing protocols. The VPN customer and the service provider must manually configure this.

MD5 authentication in routing protocols should be used on all PE-CE peers. It is easy to track the source of such a potential denial-of-service attack.

- Configure, where available, the parameters of the routing protocol 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.

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 denial-of-service attack against the PE router. This in turn might have negative impact on other VPNs. For this reason, PE routers must be extremely well secured, especially on their interfaces to the CE routers.

Label Spoofing

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 does not own. This is called *label spoofing*. This kind of attack can be done from the outside, that is, another CE router or from the Internet, or from within the MPLS core. The latter case (from within the core) is not discussed since the assumption is that the core network is provided in a secure manner.

Within the MPLS network, packets are not forwarded based on the IP destination address, but based on the 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 possible to spoof the label of an MPLS packet.

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 is only aware of the destination router. The intelligence exists in the PE device, where based on the configuration, the PE chooses a label and prepends it to the packet. This is the case for all PE routers, toward CE routers, and to the upstream service provider. All interfaces into the MPLS cloud require IP packets without labels.

For security reasons, a PE router should never accept a packet with a label from a CE router. Cisco routers implementation is such that packets that arrive on a CE interface with a label are dropped. Thus, it is not possible to insert fake labels because no labels are accepted. Additional security can be implemented by using MD5 authentication between peer routers in the core if the service provider is using LDP to distribute labels.

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 does not add any security risk here.

Securing the MPLS Core

The following is a list of recommendations and considerations on configuring an MPLS network securely.



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 TFTP server.

Trusted Devices

The PE and P devices, and remote access servers and AAA servers must 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. There is ample literature available on how to secure network elements, so these topics are not discussed here in more detail.

CE routers are typically not under full control of the service provider and must be treated as “untrusted.”

PE-CE Interface

The interface between 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 (Access Control Lists) 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 service provider network should be denied. This avoids the possibility that the PE and P routers can be attacked, since all packets to the corresponding address range are dropped by the PE router. The only exception is the peer interface on the PE router for routing purposes. This PE peer interface must be secured separately.

If private address space is used for the PE and P routers, the same rules with regard to packet filtering apply—it is required to filter all packets to this range. However, since addresses of this range should not be routed over the Internet, it limits attacks to adjacent networks.

Routing Authentication

All routing protocols should be configured with the corresponding authentication option toward the CEs and toward any Internet connection. Specifically: BGP, OSPF, and RIP2. All peering relationships in the network need to be secured this way:

- CE-PE link: use BGP MD-5 authentication
- PE-P link: use LDP MD5 authentication
- P-P

This prevents attackers from spoofing a peer router and introducing bogus routing information. Secure management is particularly important regarding configuration files, which often contain shared secrets in clear text (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 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 is not sufficient, since this does not affect normal packets.

To avoid this problem, we recommend that you 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. 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 Label Distribution Protocol (LDP) can also be secured with MD-5 authentication across the MPLS cloud. This prevents hackers from introducing bogus routers, which would participate in the LDP.

Connectivity Between VPNs

MPLS provides VPN services with address and routing separation between VPNs. In many environments, however, the devices in the VPN must be able to reach destinations outside the VPN. 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 and accessing the Internet.

To achieve this, the PE routers maintain various tables: A *routing context table* 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. In this way, two or several VPNs can be merged to a single VPN. In this case, it is necessary that all merged VPNs have mutually exclusive addressing spaces; in other words, the overall address space must be unique for all included VPNs.

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. The VPN must use private address space since 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 can use the same address space without interference. Packets from and to the merged VPNs cannot be routed to other VPNs. 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. With the standard MPLS features, there is no separation or firewalling or 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.

MP-BGP Security Features

Security in Prime Fulfillment MPLS-based networks is delivered through a combination of MP-BGP and IP address resolution. In addition, service providers can ensure that VPNs are isolated from each other.

Multiprotocol BGP is a routing information distribution protocol that, through employing multiprotocol extensions and community attributes, defines who can talk to whom. VPN membership depends upon logical ports entering the VPN, where MP-BGP assigns a unique Route Distinguisher (RD) value (see [Route Distinguishers and Route Targets, page 2-5](#)).

RDs are unknown to end users, making it impossible to enter the network on another access port and spoof a flow. Only preassigned ports are allowed to participate in the VPN. In an MPLS VPN, MP-BGP distributes forwarding information base (FIB) tables about VPNs to members of the same VPN only, providing native security by way of logical VPN traffic separation. Furthermore, iBGP PE routing peers

can perform TCP segment protection using the MD5 Signature Option when establishing iBGP peering relationships, further reducing the likelihood of introducing spoofed TCP segments into the iBGP connection stream among PE routers (for information on the MD5 Signature Option, see RFC 2385).

The service provider, not the customer, associates a specific VPN with each interface when provisioning the VPN. Users can only participate in an intranet or extranet if they reside on the correct physical or logical port and have the proper RD. This setup makes a Cisco MPLS VPN virtually impossible to enter.

Within the core, a standard Interior Gateway Protocol (IGP) such as OSPF or IS-IS distributes routing information. Provider edge routers set up paths among one another using LDP to communicate label-binding information. Label binding information for external (customer) routes is distributed among PE routers using MP-BGP multiprotocol extensions instead of LDP, because they easily attach to VPN IP information already being distributed.

The MP-BGP community attribute constrains the scope of reachability information. MP-BGP maps FIB tables to provider edge routers belonging to only a particular VPN, instead of updating all edge routers in the service provider network.

Security Through IP Address Resolution

MPLS VPN networks are easier to integrate with IP-based customer networks. Subscribers can seamlessly interconnect with a provider service without changing their intranet applications because MPLS-based networks have built-in application awareness. Customers can even transparently use their existing IP address space because each VPN has a unique identifier.

MPLS VPNs remain unaware of one another. Traffic is separated among VPNs using a logically distinct forwarding table and RD for each VPN. Based on the incoming interface, the PE selects a specific forwarding table, which lists only valid destinations in the VPN. To create extranets, a provider explicitly configures reachability among VPNs.

The forwarding table for a PE contains only address entries for members of the same VPN. The PE rejects requests for addresses not listed in its forwarding table. By implementing a logically separate forwarding table for each VPN, each VPN itself becomes a private, connectionless network built on a shared infrastructure.

IP limits the size of an address to 32 bits in the packet header. The VPN IP address adds 64 bits in front of the header, creating an extended address in routing tables that classical IP cannot forward. The extra 64 bits are defined by the Route Distinguisher and the resultant route becomes a unique 96-bit prefix. MPLS solves this problem by forwarding traffic based on labels, so one can use MPLS to bind VPN IP routes to label-switched paths. PEs are concerned with reading labels, not packet headers. MPLS manages forwarding through the provider's MPLS core. Since labels only exist for valid destinations, this is how MPLS delivers both security and scalability.

When a virtual circuit is provided using the overlay model, the egress interface for any particular data packet is a function solely of the packet's ingress interface; the IP destination address of the packet does not determine its path in the backbone network. Thus, unauthorized communication into or out of a VPN is prevented.

In MPLS VPNs, a packet received by the backbone is first associated with a particular VPN by stipulating that all packets received on a certain interface (or subinterface) belong to a certain VPN. Then its IP address is looked up in the forwarding table associated with that VPN. The routes in that forwarding table are specific to the VPN of the received packet.

In this way, the ingress interface determines a set of possible egress interfaces, and the packet's IP destination address is used to choose from among that set. This prevents unauthorized communication into and out of a VPN.

Ensuring VPN Isolation

To maintain proper isolation of one VPN from another, it is important that the provider routers not accept a labeled packet from any adjacent PE unless the following conditions are met:

- The label at the top of the label stack was actually distributed by the provider router to the PE device.
- The provider router can determine that use of that label will cause the packet to exit the backbone before any labels lower in the stack and the IP header will be inspected.

These restrictions are necessary to prevent packets from entering a VPN where they do not belong.

The VRF tables in a PE are used only for packets arriving from a CE that is directly attached to the PE device. They are not used for routing packets arriving from other routers that belong to the service provider backbone. As a result, there might be multiple different routes to the same system, where the route followed by a given packet is determined by the site from which the packet enters the backbone. So one might have one route to a given IP network for packets from the extranet (where the route leads to a firewall), and a different route to the same network for packets from the intranet.