Part

1

The VPN Overview

Chapter 1

VPN-in-Brief

1.1 VPN Overview

This is the information age. We no longer have to commute physically from one place to another to complete a set of tasks or to gather pieces of information. Everything can be done virtually with a mouse click on an online host. In a way, everything we do in our daily lives is related in one way or another to information access. This has made information sharing almost mandatory and indispensable. These days, a customer can retrieve and compare products or services information promptly online, anytime, anywhere. For competitive reasons, organizations that provide this information have to make the information readily available online. In other words, the concept of a shared infrastructure is undisputedly important. A shared infrastructure is none other than a public network. At present, the biggest public network is the Internet, which has over 100,000 routes and is still growing rapidly.

As more and more companies link up their corporate network to the Internet, we are faced with an inevitable issue—information security. Sharing information on a public network also implies giving access and visibility to everyone who wants to retrieve these data. What if the person who has the accessibility and visibility to the information decides to create havoc? Some of the general threat types that are posed by malicious hackers include eavesdropping, denial of service, unauthorized access, data manipulation, masquerade, session replay, and session hijacking.

How do we ensure the safe passage of data across a shared infrastructure? The answer is to deploy a secured virtual private network (VPN). VPNs are networks deployed on a public network infrastructure that utilize the same security, management, and quality of service policies that are applied in a private network. VPNs provide an alternative to building a private network for site-to-site communication over a public network or the Internet. Because they operate across a shared infrastructure rather than a private network, companies can cost effectively extend the corporate WAN to telecommuters, mobile users, and remote offices as well as to new constituencies, such as customers, suppliers, and business partners.

Traditional private WANs connect customer sites via dedicated point-to-point links. This means that multiple independent circuits have to terminate at the corporate network egress, making the deployment nonscalable and difficult to maintain. VPNs extend the classic WAN by replacing the physical point-to-point links with logical point-to-point links sharing a common infrastructure, allowing all the traffic to be aggregated into a single physical connection. This scenario results in potential bandwidth and cost savings at the network egress. Because customers no longer need to maintain a private network, and because a VPN itself is cheaper to own and offers significant cost savings over private WANs, operation costs are reduced.

VPNs provide an alternative WAN infrastructure that can replace or augment commercial private networks that use leased-line or frame relay/ATM networks. There are two ways business customers can implement and manage their VPNs. They can either roll out their own VPNs and manage them internally, or outsource the VPN management to their service providers for a total VPN package that is tailored to their particular business needs.

Last but not least, from the service providers' perspective, VPNs are a fundamental building block in delivering new value-added services that benefit their business customers as well as themselves. In this instance, the service providers deploy the VPNs for their customers, and the customers need only subscribe to the service providers for the VPN services.

1.2 VPN Types and Solutions

In this section we address three types of VPNs: remote access, site-to-site, and firewall-based (a site-to-site variation). The variation between remote access and site-to-site VPNs will become more ambiguous as new devices such as hardware VPN clients, become more prevalent. These appear as a single device accessing the network, albeit there may be a network with several devices behind it. In all cases, the VPN comprises two endpoints that may be represented by routers, firewalls, client workstations, or servers.

1.2.1 Remote access

Remote access VPNs or virtual private dialup network (VPDN) are deployed for individual remote users, commonly referred to as mobile users and telecommuters. In the past, corporations supported these remote users via dialup networks, which required the remote users to make toll calls to access the corporate network directly. This was not a cost-effective solution, especially when an international user made a call back.

With the introduction of remote access VPNs, a mobile user can make a local call to their Internet service provider (ISP) to access the corporate network with their PC via the Internet wherever they may be. Remote access VPNs are an extension of the traditional dialup networks. In this case, the software on the PC provides a secure connection, often known as a tunnel, back to the corporation. Since the users need only make local calls, the operation cost is reduced. Remote access VPNs and their corresponding technologies are beyond the scope of this book.

1.2.2 Site-to-site

Site-to-site VPNs are deployed for interconnecting corporate sites. In other words, the network of one location (site) is connected to the network of another location (site) via a VPN. In the past, a leased line or frame relay connection was required to connect the sites; however, these days, most corporations have Internet access. With Internet access, leased lines and frame relay circuits can be replaced with site-to-site VPNs. Site-to-site VPNs are an extension of legacy WAN networks.

Site-to-site VPNs can be further viewed as intranet VPNs or extranet VPNs. Intranet VPNs refer to connections between sites that all belong to the same organization. User access between these sites is less restraining than for extranet VPNs. Extranet VPNs refer to connections between an organization and its business partners. User access between these sites should be tightly controlled by both entities at their respective sites.

1.2.3 Firewall-based

A firewall-based VPN is intrinsically a site-to-site implementation. Firewall-based VPN solutions are not a technical but a security issue. They are deployed when a corporation requires more advanced perimeter security measures for its VPNs. Corporations can enhance their existing firewalls to support firewall-based VPNs.

1.3 VPN Terminology

This section lists some of the common VPN terminology that is used in the subsequent sections:

Provider network (P-Network): the service provider infrastructure that is used to provide VPN services.

Customer network (C-Network): the part of the network that is still under customer control.

Customer site: a contiguous part of the C-Network that can comprise many physical locations.

Provider (P) device: the device in the P-Network with no customer connectivity and without any “knowledge” of the VPN. This device is usually a router and is commonly referred as the P router.

Provider edge (PE) device: the device in the P-Network to which the CE devices are connected. This device is usually a router and is often referred as the PE router.

Customer edge (CE) device: the device in the C-network that links into the P-network; also known as customer premises equipment (CPE). This device is usually a router and is normally referred as the CE router.

Virtual circuit (VC): logical point-to-point link that is established across a shared layer-2 infrastructure.

1.4 VPN Models

VPN services can be offered based on two major paradigms:

1. The overlay VPNs, whereby the service provider furnishes virtual point-to-point links between customer sites.
2. The peer-to-peer VPNs, whereby the service provider participates in customer routing.

In the following sections, we discuss these two different models in details.

1.4.1 Overlay model

The overlay VPN is deployed via private trunks across a service provider's shared infrastructure. These VPNs can be implemented at layer-1 using leased/dialup lines, at layer-2 using X.25/frame relay/ATM Virtual Circuits, or at layer-3 using IP (GRE) tunneling.

In the overlay VPN model, the service provider network is a connection of point-to-point links or virtual circuits (VCs). Routing within the customer network is transparent to the service provider network, and routing protocols run directly between customer routers. The service provider has no knowledge of the customer routes and is simply responsible for providing point-to-point transport of data between the customer sites.

Figure 1.1 illustrates the deployment of an overlay VPN. The scenario adopts a hub-and-spoke topology whereby the Paris site is the hub, and both the London and Zurich sites are the spokes. The London site is linked up to the Paris site via a point-to-point VC #1. Likewise, the Zurich site is linked up to the Paris site via a point-to-point VC #2. In this instance, the layer-3 routing adjacencies are established between the CE routers at the various customer sites, and the service provider is not aware of this routing information at all. As illustrated in Figure 1.2, from the perspective of the CE routers, the service provider infrastructure appears as point-to-point links between Paris–London and Paris–Zurich.

The overlay VPN model has two further constraints. One is the high level of difficulty in sizing the intersite circuit capacities. The other is the requirement

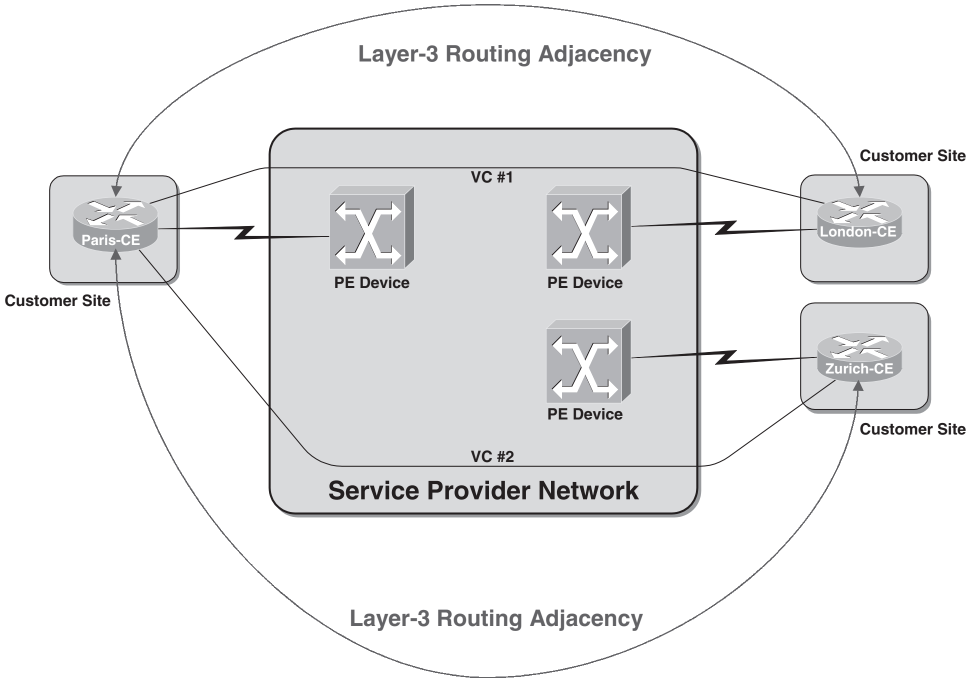

Figure 1.1: The overlay VPN model.

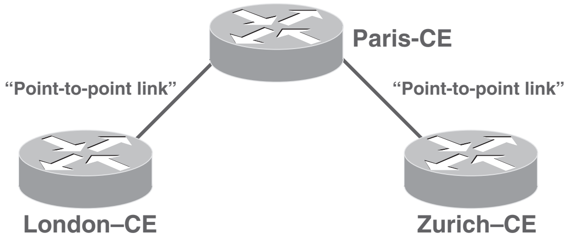

Figure 1.2: Perception of the SP infrastructure from the CE routers.

of a fully meshed deployment of point-to-point links or VCs over the service provider's backbone to attain optimal routing.

1.4.1.1 Layer-1 implementation. Figure 1.3 illustrates the overlay VPN layer-1 implementation, which adopts the traditional time division multiplexing (TDM) solution. In this scenario, the service provider assigns bit pipes and establishes the physical-layer (Layer-1) connectivity between customer sites via ISDN, DS0, T1, E1, SONET, or SDH, and the customer is accountable for implementation of all higher layers, such as PPP, HDLC, and IP.

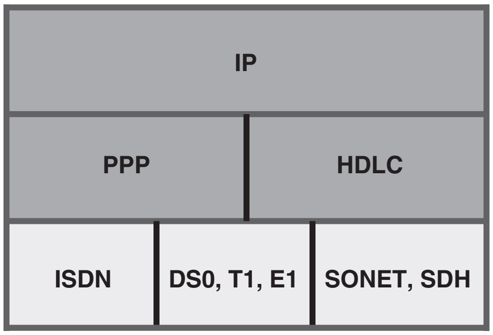

Figure 1.3: Overlay VPN layer-1 implementation.

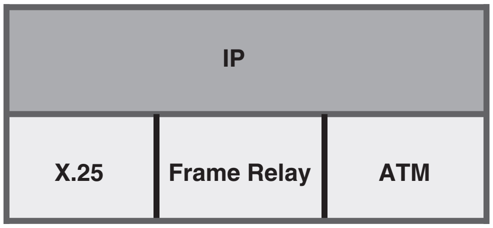

Figure 1.4: Overlay VPN layer-2 implementation.

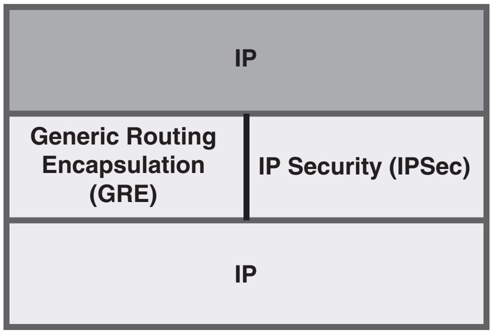

Figure 1.5: Overlay VPN layer-3 implementation.

1.4.1.2 Layer-2 implementation. Figure 1.4 illustrates the overlay VPN layer-2 implementation, which adopts the traditional switched WAN solution. In this scenario, the service provider is responsible for establishing layer-2 VCs between customer sites via X.25, Frame Relay, or ATM, and the customer is accountable for the IP layer and above.

1.4.1.3 Layer-3 implementation. Figure 1.5 illustrates the overlay VPN layer-3 implementation, whereby the VPN is implemented with point-to-point IP-over-IP tunnels. This is commonly referred as IP tunneling whereby a destination can be reached transparently without the source having to know the topology specifics. Therefore, virtual networks can be created by tying otherwise unconnected devices or hosts together through a tunnel. Tunnels also enable the use of private network addressing across a service provider's backbone without the need for network address translation (NAT). Tunnels are established with generic routing encapsulation (GRE) or IP security (IPSec). The GRE implementation is simpler and quicker but less secure, whereas the IPSec deployment is more complex and resource-intensive (CPU cycle) but it provides more robust security.

GRE tunnels provide a specific conduit across the shared WAN and encapsulate traffic with new packet headers to ensure delivery to specific destinations. Since traffic can enter a tunnel only at an endpoint, the network is private. GRE tunnels do not provided true confidentiality but can carry encrypted traffic. For better protection, GRE can be used in conjunction with IPSec. When we deploy a layer-3 VPN over a public network, the VPN tunnels are built across a distrusted shared infrastructure. IPSec provides robust security services such as confidentiality, integrity, and authentication to ensure that sensitive information is securely transported across these tunnels and not circumvented accidentally or intentionally.

IPSec is an Internet Engineering Task Force (IETF) standard that enables encrypted communication between users and devices. It can be implemented transparently and seamlessly into the network infrastructure. End users need not have any knowledge that packets are being intercepted and transformed by IPSec. Because it operates at the network layer (layer-3), IPSec is ideally positioned to enforce corporate network security. Since IPSec encryption works only on IP unicast frames, it can be used in conjunction with GRE to alleviate this deficit. This is because GRE is capable of handling the transportation of multiprotocol and IP multicast traffic between two sites, which have only IP unicast connectivity.

1.4.2 Peer-to-peer model

The peer-to-peer model adopts a simple routing scheme for the customer. Both provider and customer network use the same network protocol and all the customer routes are carried within the core network (service provider network). The PE routers exchange routing information with the CE routers, and layer-3 routing adjacencies are established between the CE and PE routers at each site. Because peer-to-peer routing has been implemented, routing between sites is now optimal. Fully meshed deployment of point-to-point links or VCs over the service provider backbone is no longer applicable to attain optimal routing. Since there is no overlay mesh to contend with, the addition of new sites is easier, and circuit capacity sizing is not an issue. Because the service provider now participates in customer routing, provider-assigned or public address space needs to be deployed at the customer's network, so private addressing is no longer an option.

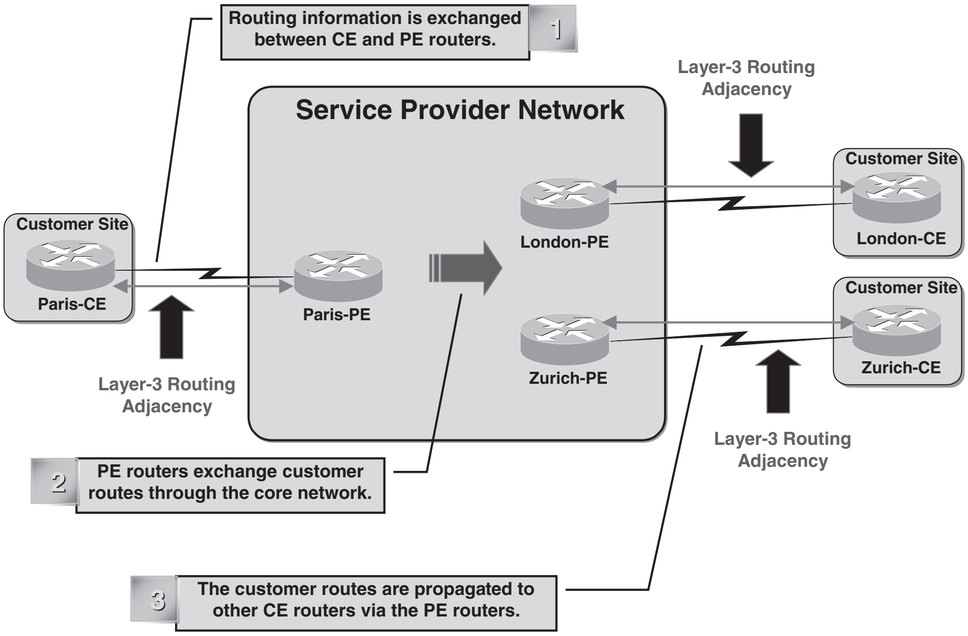
Routing information is exchanged between CE and PE routers.
1
Layer-3 Routing Adjacency
Service Provider Network
Customer Site
London-CE
Customer Site
Paris-CE
London-PE
Paris-PE
Customer Site
Zurich-CE
Zurich-PE
Layer-3 Routing Adjacency
Layer-3 Routing Adjacency
2
PE routers exchange customer routes through the core network.
3
The customer routes are propagated to other CE routers via the PE routers.

Figure 1.6: The peer-to-peer VPN model.

Figure 1.6 illustrates the deployment of a peer-to-peer VPN. In this scenario, customer routing information is exchanged between Paris-CE and Paris-PE. The customer routes are then broadcast through the core network to London-PE and Zurich-PE, which in turn, propagate these routes to their respective CE routers. The shared router and dedicated router approaches are derivatives of the peer-to-peer model. We discuss these two approaches in Sections 1.4.2.1 and 1.4.2.2.

1.4.2.1 Shared PE router model. In this model, a common PE router that carries customer routes is deployed. Individual customers routes are separated with packet filters on PE-CE interfaces. The packet filters are managed so that information goes to the proper site and different customers are separated. The complexity of these packet filters results in high maintenance costs and a significant impact on performance.

Figure 1.7 illustrates the Shared PE router model. In this scenario, there are three separate VPNs: VPN-101, VPN-201, and VPN-301. These VPNs are deployed over four different customer sites. VPN-101 is deployed for Paris as well as Lyon; VPN-201 is deployed for Brussels; and VPN-301 is deployed for Munich. In Figure 1.7, we can see that London-PE carries all the customer routes for VPN-101, VPN-201, and VPN-301. Isolation between VPNs is achieved with packet filters (access lists) on the PE-CE interfaces: Serial0/0, Serial0/1, Serial0/2, and Serial0/3.

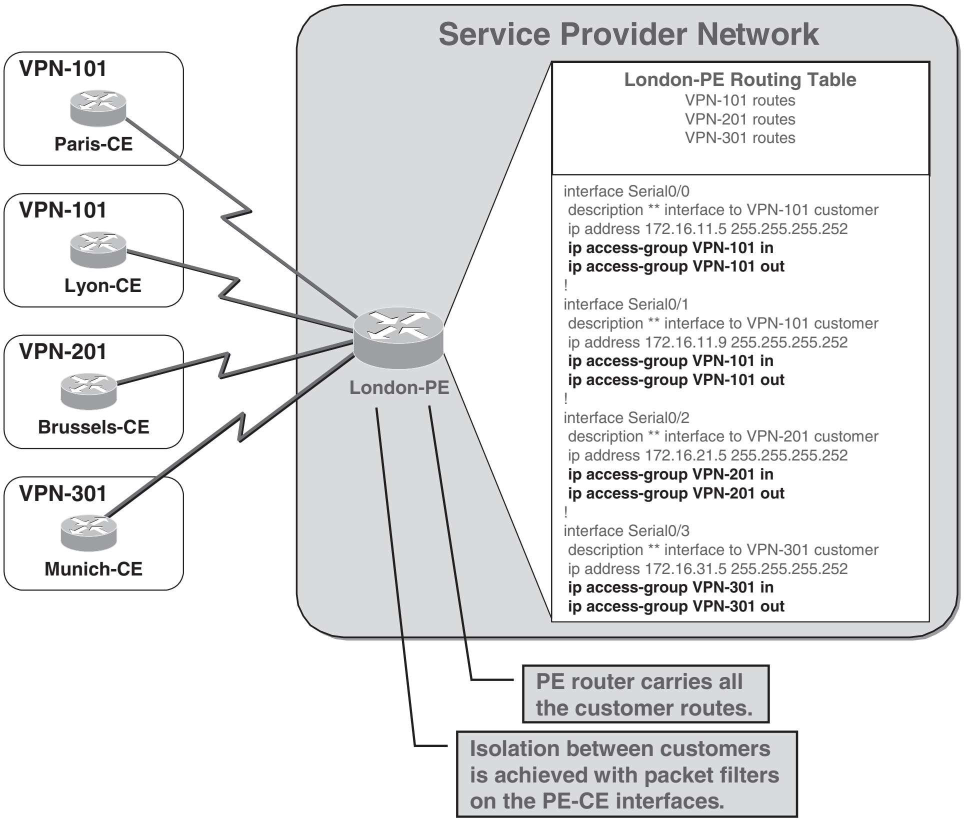

```
interface Serial0/0
 description ** interface to VPN-101 customer
 ip address 172.16.11.5 255.255.255.252
 ip access-group VPN-101 in
 ip access-group VPN-101 out
!
interface Serial0/1
 description ** interface to VPN-101 customer
 ip address 172.16.11.9 255.255.255.252
 ip access-group VPN-101 in
 ip access-group VPN-101 out
!
interface Serial0/2
 description ** interface to VPN-201 customer
 ip address 172.16.21.5 255.255.255.252
 ip access-group VPN-201 in
 ip access-group VPN-201 out
!
interface Serial0/3
 description ** interface to VPN-301 customer
 ip address 172.16.31.5 255.255.255.252
 ip access-group VPN-301 in
 ip access-group VPN-301 out
```

Figure 1.7: Shared PE router model.

1.4.2.2 Dedicated PE router model. In this model, each customer has a dedicated PE router that carries only its own routes. Customer segregation is achieved through lack of routing information on the PE router. The P router contains all customer routes and filters routing updates between different PE routers using Border Gateway Protocol (BGP) Communities. Because each customer has a dedicated PE router, this approach is expensive to deploy, and hence it is not a cost-effective solution.

Figure 1.8 illustrates a dedicated PE router model. In this scenario, there are two separate VPNs: VPN-101 and VPN-201. These VPNs are deployed over four different customer sites. VPN-101 is deployed for two locations in Paris and one location in Brussels, and VPN-201 is deployed for one location in London. In Figure 1.8, we can see that the P router in the service provider network contains all the customer routes for VPN-101 and VPN-201. It filters routing updates between Paris-PE, London-PE, and Brussels-PE using BGP Communities. VPN-101 routes are tagged with a community value of 101:1, and VPN-201 routes are tagged with a community value of 201:1. In other words, the P

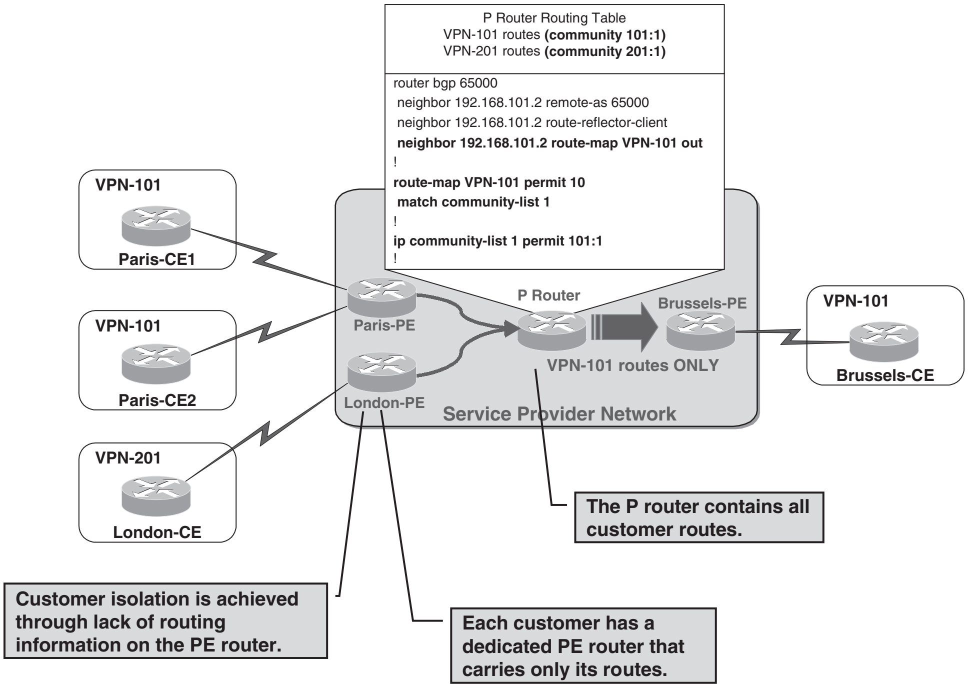
P Router Routing Table
VPN-101 routes (community 101:1)
VPN-201 routes (community 201:1)
VPN-101
Paris-CE1
VPN-101
Paris-CE2
VPN-201
London-CE
Paris-PE
London-PE
P Router
Brussels-PE
VPN-101 routes ONLY
Service Provider Network
VPN-101
Brussels-CE
The P router contains all customer routes.
Customer isolation is achieved through lack of routing information on the PE router.
Each customer has a dedicated PE router that carries only its routes.

```
router bgp 65000
 neighbor 192.168.101.2 remote-as 65000
 neighbor 192.168.101.2 route-reflector-client
 neighbor 192.168.101.2 route-map VPN-101 out
!
route-map VPN-101 permit 10
 match community-list 1
!
ip community-list 1 permit 101:1
!
```

Figure 1.8: Dedicated PE router model.

router will propagate routes with a community value of 101:1 to VPN-101 and routes with a community value of 201:1 to VPN-201. Each customer or VPN has a dedicated PE router that carries only its routes. Customer isolation is simply achieved through the lack of routing information on the PE routers.

In this example, routes from Paris-CE1 and Paris-CE2, which are both in VPN-101, are announced to a dedicated PE router—Paris-PE. Paris-PE in turn advertises the routes to the P router. The P router uses a community list to match routes with a community value of 101:1 and propagates the routes that match this community value to another dedicated PE router, Brussels-PE, which serves the customer site that is in VPN-101 at Brussels.

1.4.3 MPLS VPN model

MPLS VPN is a true peer-to-peer model that combines the best of both worlds. It unites the customer security and segregation features implemented in the overlay model with the simplified customer routing deployed in the traditional peer-to-peer model. The MPLS VPN architecture is very similar to the dedicated PE router model, except the dedicated per customer routers are implemented as virtual routing tables within the PE router. In other words, customer segregation is achieved through the concept of virtual routing and forwarding (VRF) whereby the PE router is subdivided into virtual routers serving differ-

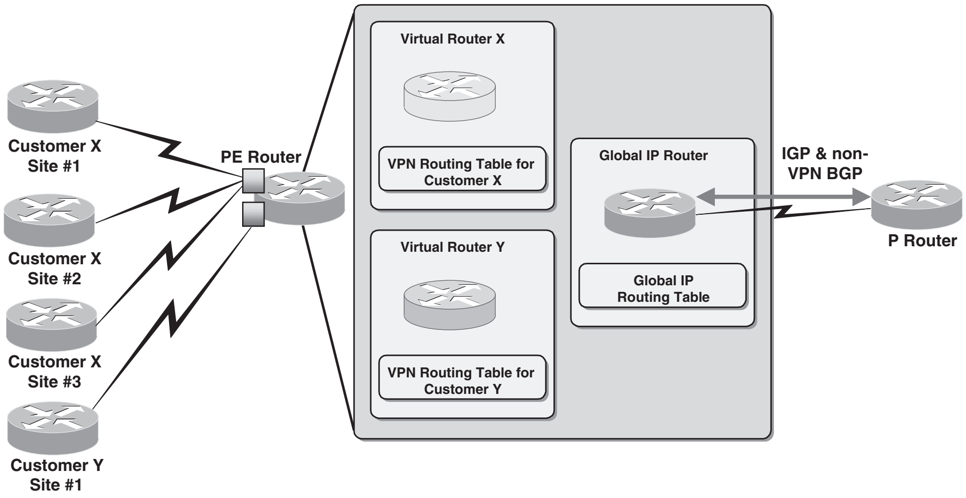

Figure 1.9: MPLS VPN model.

ent VPNs (or customer sites). This establishes overlapping addresses in different customer sites since each customer is now assigned an independent routing table.

The PE routers hold the routing information only for directly connected VPNs. As a result, the size of the PE routing table is significantly reduced. The amount of routing information is proportional to the number of VPNs attached to the PE router. As such, the PE routing table will still grow when the number of directly connected VPNs increases. In addition, the PE routers participate in customer routing, ensuring optimal routing between sites and easy provisioning. Full routing within the service provider backbone is no longer required because multi protocol label switching (MPLS) (see Chapter 7 for more in-depth MPLS concepts), and not traditional IP routing, is used to forward packets. The details of the MPLS VPN architecture are discussed in Chapter 8.

Figure 1.9 illustrates the MPLS VPN model. There are two separate customers—Customer X and Customer Y. Customer X spans three different sites, while Customer Y has only one site. Customer isolation is achieved with a dedicated per customer virtual router. In Figure 1.9, we can see that routing within the PE router is split into two separate planes, one for VPN routing and the other for global IP routing.

In the VPN routing plane, the PE router is subdivided into virtual router X which serves customer X, and virtual router Y which serves customer Y. Each of these virtual routers participates in customer routing and keeps an independent VPN routing table for its respective customers. Meanwhile, the PE router also has a global IP router in the global IP routing plane that take cares of the IGP and non-VPN BGP routing between the various PE and P routers.

TABLE 1.1 Upside and Downside of Overlay VPN versus Peer-to-Peer VPN.

Overlay VPN Model		Peer-to-Peer VPN Model	
Upside	Downside	Upside	Downside
■ Allows replicate IP addressing ■ Full isolation between customers ■ Secure VPN service	■ Difficult to size intersite circuit capacity ■ Fully meshed circuit requirement for optimal routing ■ Layer-3 CE routing adjacencies between sites	■ Routing between sites is optimal ■ Circuit capacity sizing between sites is not an issue ■ Simpler routing configuration for customers (no overlay mesh)	■ All VPN routes are carried in the service provider IGP ■ Replicate IP addressing is no longer an option ■ Complex filters or dedicated devices

TABLE 1.2 Benefits of MPLS VPN

MPLS VPN Model	
Combined benefits of overlay and peer-to-peer VPN models	
■ Routing between sites is optimal	■ PE routers hold only pertinent VPN routers
■ Allows replicate IP addressing	■ Full isolation between customers
■ Secure VPN service	■ No complex filters or dedicated routers

1.5 Comparison Between Various VPN Implementations

This section compares some of the benefits and drawbacks between overlay VPN and peer-to-peer VPN. Table 1.1 gives a quick appraisal of the upside and downside of the overlay VPN model versus the peer-to-peer VPN model.

Table 1.2 illustrates the benefits of the MPLS VPN model, which are combined from the overlay and peer-to-peer VPN models for the best of both worlds.

1.6 IPSec versus MPLS VPNs

Both IPSec and MPLS VPNs are implemented at the network layer of the OSI model. In other words, they are layer-3 VPNs that use the IP protocol as the network layer protocol. However, the deployment of these different VPN architectures is rather controversial because IPSec VPN is based on the overlay model (see Section 1.4.1), while MPLS VPN is based on the peer-to-peer model (see Section 1.4.3) which offers the combined benefits of traditional overlay and peer-to-peer VPN models (see Table 1.2). The controversy between the IPSec and MPLS VPNs is narrowed down to one important factor: who owns the VPN.

In other words, the implementation of these VPN architectures depends on whether the VPN is managed by the customer or the service provider.

Customers who manage their own VPNs usually adopt the overlay approach. Overlay VPNs are connection oriented. They are based on creating point-to-point connections and not networks. In IPSec VPNs, IPSec tunnels are created to provide point-to-point connectivity at the customer's business site. This also implies that IP connectivity is required to establish these point-to-point tunnels. The best way to interconnect the sites is via the Internet through a service provider who is responsible only for providing Internet connectivity between the customer sites. In this case, the customers will have to build and secure their own VPNs across the Internet by deploying IPSec. These site-to-site enterprise-managed VPNs are also known as Internet VPNs. Because site-to-site peering is required when implementing IPSec, scalability issues will arise when the number of customer sites grows.

MPLS is typically offered as a site-to-site VPN service from a service provider. The service provider builds and manages a private IP-based network and offers multiple customers IP connectivity between their sites across this network. This highly scalable connectionless architecture allows individual customers to view the MPLS service as though they had an IP VPN connecting their sites. This setup provides customers the same benefits of a layer-2 private network (see Section 1.4.1.2), but with scalability and the easy management features of an IP (layer-3) network, by eliminating the need for tunnels or VCs. Since MPLS VPN runs across a private IP-based network rather than the Internet, the service provider has the capabilities to offer differentiated levels of services and service level agreements (SLAs) to its customers. A typical MPLS VPN deployment is capable of supporting tens of thousands of VPN groups over the same network. However, because MPLS VPN is based on a service provider's private network, the reach of the service is constrained to the locations at which the service provider operates.

1.7 Summary

There is a growing demand for VPNs. The different types of VPNs such as remote access, site-to-site, and firewall-based VPNs have been examined. The layer-1, layer-2 and layer-3 implementations of overlay VPNs have been explained, along with the shared router and dedicated router approaches for peer-to-peer VPNs. We also briefly examined how the MPLS VPN model operates and compared the benefits and drawbacks of overlay VPNs versus peer-to-peer VPNs and listed the benefits of MPLS VPNs, which are the best of both worlds. The last section of the chapter addressed the controversy on the deployment of IPSec and MPLS VPN architectures. The remainder of this book discusses implementing IPSec VPNs in Enterprises and deploying MPLS VPNs for Service Providers.