CHAPTER 11

MPLS VPN Technologies

Chapter 2 provided some brief discussion of Virtual Private Network (VPN) architecture with respect to connectivity options for teleworkers. Remote-access VPNs and IPsec VPNs were both discussed along with some key differences between the two. Among the items discussed was the fact that a remote-access VPN is an on-demand connection, whereas an IPsec VPN is an always-on connection. Each has its particular place in the bigger picture of the Intelligent Information Network (IIN).

The Service-Oriented Network Architecture (SONA) framework encourages the offering of applications and services to all network users so that they may have the same network experience regardless of how they access the network. The Multiprotocol Label Switching (MPLS) VPN is another piece of the SONA framework that allows those applications and services to be offered to remote branch offices and small office/home office (SOHO) sites. With MPLS VPNs, two key pieces of the framework fall into place: the teleworker and, now, the branch office sites. For SOHO sites, any of the three VPN options is viable depending on the implementation.

"Do I Know This Already?" Quiz

The purpose of the "Do I Know This Already?" quiz is to help you decide whether you really need to read the entire chapter. If you already intend to read the entire chapter, you do not necessarily need to answer these questions now.

The 12-question quiz, derived from the major sections in the "Foundation Topics" portion of the chapter, helps you to determine how to spend your limited study time.

Table 11-1 outlines the major topics discussed in this chapter and the "Do I Know This Already?" quiz questions that correspond to those topics.

Table 11-1 *"Do I Know This Already?" Foundation Topics Section-to-Question Mapping*

Foundation Topics Section	Questions Covered in This Section	Score
MPLS VPN Architecture	1–2	
Traditional VPNs	3	
Peer-to-Peer VPNs	4	
MPLS VPNs	5–12	
Total Score		

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark this question wrong for purposes of self-assessment. Giving yourself credit for an answer that you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which type of VPN does not require any participation by the service provider in the routing functionality?
   a. Overlay VPN
   b. Peer-to-peer VPN
   c. Overlay-to-overlay VPN
   d. MPLS VPN
2. Which of the following is implemented with routers, ACLs, and dedicated routers per customer?
   a. Overlay VPNs
   b. Peer-to-peer VPNs
   c. Overlay-to-overlay VPNs
   d. MPLS VPNs
3. In a Layer 2 overlay VPN model, how is redundancy achieved?
   a. It is automatic due to routing protocol convergence.
   b. By provisioning additional circuits between critical sites.
   c. Only through the hub router.
   d. Redundancy is the responsibility of the provider.

4. Which is a characteristic of a peer-to-peer VPN?
   a. Dedicated PE router per customer
   b. Shared PE routers
   c. MPLS VPNs
   d. Lack of dynamic routing
5. Which of the following comprise all or part of the LSP?
   a. C network
   b. CE router
   c. P router
   d. PHP
6. Which of the following is prepended to a customer route?
   a. VPNv4 address
   b. RD
   c. RT
   d. LDP
7. Which of the following is appended to a customer route to indicate VPN membership?
   a. VPNv4 address
   b. RD
   c. RT
   d. LDP
8. Which protocol runs in the P network with the express purpose of propagating customer routes between PE routers?
   a. BGP
   b. OSPF
   c. MPBGP
   d. MPOSPF
9. Where would an import RT most likely be used?
   a. Ingress PE
   b. Egress PE
   c. P router
   d. CE router

10. Customer routes from a VRF are exported as VPNv4 routes into what?

    a. LDP
    b. Egress PE
    c. MPBGP
    d. CE router

11. PE routers use a label stack consisting of how many labels in a typical MPLS VPN?

    a. 1
    b. 2
    c. 3
    d. 4

12. When the final P router in an LSP removes the top label in the stack, this is known as?

    a. Label unstacking
    b. Penultimate hop popping
    c. VRF export
    d. VPN label

The answers to the "Do I Know This Already?" quiz are found in Appendix A, "Answers to the 'Do I Know This Already?' Quizzes and Q&A Sections." The suggested choices for your next step are as follows:

- **7 or fewer overall score**—Read the entire chapter. This includes the "Foundation Topics," "Foundation Summary," and "Q&A" sections.
- **8 or 9 overall score**—Begin with the "Foundation Summary" section, and then go to the "Q&A" section.
- **10 or more overall score**—If you want more review on these topics, skip to the "Foundation Summary" section, and then go to the "Q&A" section. Otherwise, move to the next chapter.

Foundation Topics

MPLS VPN Architecture

To properly understand MPLS VPNs as a solution, it is important to understand the problem. MPLS VPNs are a Layer 3 WAN solution to an age-old Layer 2 WAN problem—that is, the quest to provide any-to-any connectivity among sites in a cost-efficient manner. In the past, WAN architects struggled with topological design principals that amounted to choosing the least of all evils. A full mesh topology was too expensive but most robust. A hub-and-spoke topology was least expensive but least robust. A failure at the hub site would have a severe network impact. Partial mesh topologies created a balance of pain created by leveraging cost against connectivity.

MPLS is the answer to the problem. With MPLS, it is possible to have a fully meshed network, but beyond that, it is a Layer 3–capable, fully meshed network. The possibilities for architecting a WAN solution are greatly expanded with little or no incremental cost over traditional Layer 2 circuits.

The idea of a VPN brings to mind the concepts of security and privacy. These things have always been an enterprise solution that had to be implemented by knowledgeable individuals within a particular company or by an outside consultant brought in for just such a deployment. The term VPN still brings to mind, for most people, the IPsec and remote-access VPNs discussed in Chapter 2.

All-in-all, the term VPN has become rather wide reaching. Figure 11-1 illustrates this fact in detailing what VPN has come to mean in a wider sense.

Figure 11-1 *VPN Taxonomy*

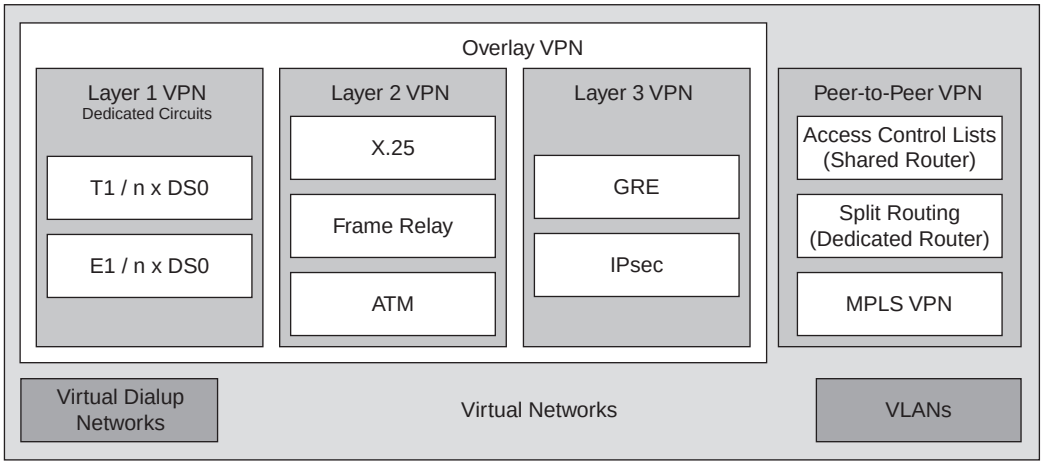

In essence, Figure 11-1 shows an evolutionary path of the VPN and how it has come to encompass a very different set of technologies depending on how it is to be deployed.

Virtual local-area networks (VLAN) allow the isolation of traffic on a per-subnet basis across a common physical infrastructure.

Virtual private dialup networks (VPDN) allow the use of dialup infrastructure via private implementation or as a service offered by a service provider.

VPNs allow the use of a shared infrastructure offered by a service provider to implement private networks. The degree of security is, of course, subject to negotiation. Many service provider offerings now include a "firewall in the cloud" offering to filter traffic to and from an Internet connection or other network. Also available are managed voice, content caching, and content filtering services. It all depends on the negotiated package.

From a typical VPN implementation standpoint, there are essentially two models:

- **Overlay VPNs**—Include older technologies such as X.25, Frame Relay, and Asynchronous Transfer Mode (ATM) for Layer 2 overlay VPNs as well as generic routing encapsulation (GRE) tunnels and IPsec for Layer 3 overlay VPNs
- **Peer-to-peer VPNs**—Implemented with shared service provider router infrastructure using access control lists (ACL) and providing separate routers per customer

Traditional VPNs

Traditional VPNs, or overlay VPNs, are essentially what has been considered a WAN solution for the past few decades and then some. These are based on a Layer 2 overlay model in which a service provider sells permanent virtual circuits (PVC) and/or switched virtual circuits (SVC). The drawbacks of the Layer 2 overlay have been discussed in quite a bit of detail up to this point.

Like most other networking technologies, VPN connections have evolved from Layer 1 up. The concept of Overlay VPNs began years ago in the form of dedicated circuits primarily used for Time-Division Multiplex (TDM) traffic. This evolution continued upward to reach Layers 2 and 3 in their respective forms.

Layer 1 Overlay

Layer 1 overlay VPN implementations were also sold by service providers in the form of Layer 1 circuits. These included such technologies as Integrated Services Digital Network (ISDN). Not to be excluded are the circuits that formed the backbone of the access technology offerings, the digital service (DS) hierarchy, DS0, DS1, and so on. A single DS0 offers 64 kbps of bandwidth

but when time-division multiplexing (TDM) implementations grouped 24 DS0s together, a DS1 was the result, offering 1.544 Mbps of bandwidth or what is more commonly referred to as a T1 line. In Europe and other locales around the globe, service providers would group 30 DS0s into a bundle, use an additional DS0 for framing operations, and use yet another DS0 for signaling. This 32 DS0 implementation, known as E1, offers 2.048 Mbps of bandwidth.

Other higher-speed technologies such as Synchronous Optical Network (SONET) and Synchronous Digital Hierarchy (SDH) were brought about by the ever-present need for more speed.

Service providers delivered the Layer 1 and the customer was responsible for applying a Layer 2 and any other features that might be appropriate. Today's market calls for much more on the part of the service provider.

Layer 2 Overlay

Layer 2 VPN overlay, as mentioned, is more along the lines of what most network administrators and IT staff think of as a traditional WAN service. This includes X.25, Frame Relay, ATM, High-Level Data Link Control (HDLC), Synchronous Data Link Control (SDLC), and Switched Multimegabit Data Service (SMDS), to name a few. At this point, the service provider is delivering Layer 1 and Layer 2, leaving the higher-level services at the discretion of the customer. Again, today's market demands yet more from the service provider as protection of applications and services traffic becomes more significant across the WAN. The momentum behind this is driven by the ideas expressed in the SONA framework and the desire to deliver a single experience for all users, regardless of location or access method. Figure 11-2 illustrates a classic example of a Layer 2 overlay VPN.

In Figure 11-2, a headquarters site is connected via Layer 2 virtual circuits (VCs) in a hub-and-spoke topology. The Layer 3 connectivity is unknown to the provider's network and routing updates must be sent across the VCs to each site. All traffic between the remote sites traverses the hub router at the headquarters site. Should the router at the headquarters site experience a failure, there will be considerable impact on the other remote sites.

In such scenarios, enterprise network administrators implement such backup features as dial-backup to facilitate data flow between sites in the event of a primary WAN link failure.

Figure 11-2 *Layer 2 Overlay VPN*

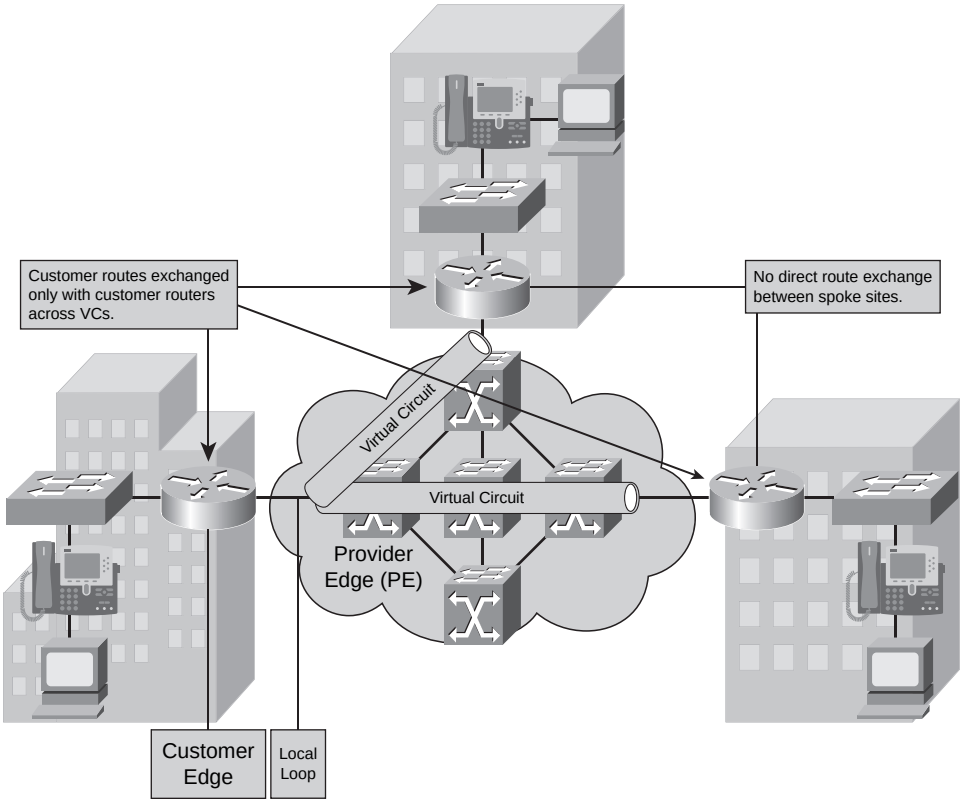

Layer 3 Overlay

Traditional WAN connectivity would entail the configuration of Layer 3 options manually to send routing information via WAN circuits. For example, the use of the **broadcast** keyword when configuring **frame-relay map** statements when mapping a next-hop IP address to a local data-link connection identifier (DLCI) would complete a necessary Layer 2 to Layer 3 address mapping, allowing routing updates to be transmitted across the link.

Even with such a configuration in place, there is no real Layer 3 capability to adapt to changes brought about by routing protocol updates. Each circuit is still a point-to-point connection in every sense of the concept. While Layer 3 protocols may flow across the links, the links are not Layer 3 aware. Customer routes flow directly between customer routers across the WAN connection.

Peer-to-Peer VPNs

The introduction of a peer-to-peer VPN causes the service provider to take a more active role in the routing operations of its customer base. This means that the service provider will be

maintaining customer routing information stored in a separate routing instance within its network. The customer edge (CE) router exchanges routing information not with the far-end CE router, but with the local, provider edge (PE) router. These routes are conveyed across the provider network to other CE routers.

This connection to and sharing of routing information with the service provider facilitates the concept of a peer-to-peer VPN. This evolutionary step forward allows the WAN to be Layer 3 aware rather than simply a Layer 3 transport. Figure 11-3 illustrates this concept.

Figure 11-3 *Peer-to-Peer VPN*

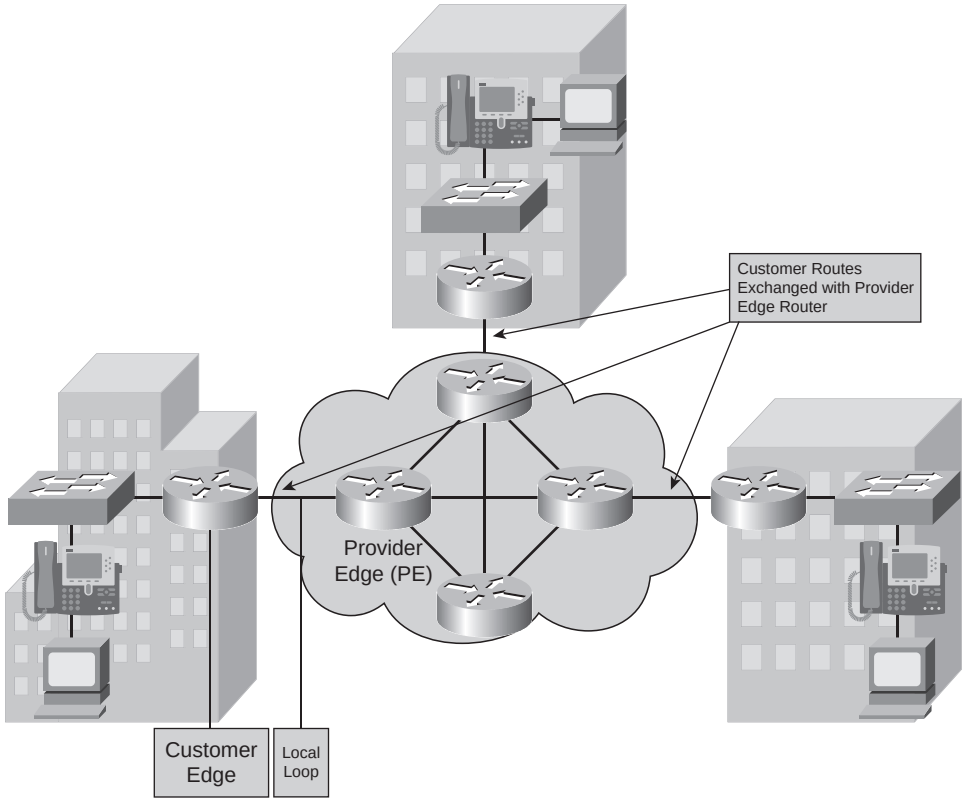

With a peer-to-peer network, the provider is handing off a Layer 1, Layer 2, and Layer 3 connection. Typically, the Layer 2 is still Frame Relay simply because most network administrators are comfortable with it. However, the next-hop addresses are those of the PE router. Most providers allow the customer to choose the routing protocol that is used across the local loop. Once the routes hit the PE, they are redistributed into the provider's Border Gateway Protocol (BGP) table.

Even though the local loop has not changed, the essence of the network has changed. The provider is now part of the customer routing infrastructure. A full mesh topology is accomplished through a single link to the provider network. The added benefits of a full mesh network come to bear. The network is more resilient because it is simply an extension of the existing customer routing infrastructure.

VPN Benefits

As access technologies advance, options become more numerous. The choices made for connectivity will be driven primarily by the needs of the business constructing the network architecture. The needs of a large enterprise network will be somewhat different from those of a small business.

Overlay VPNs are well known and have gone down in price to a large degree. They are easily implemented from both provider and customer points of view. They are now seen as a less-complex solution because the provider does not participate in the customer's routing infrastructure. This means that route redistribution need not be a concern when passing information between sites.

Peer-to-peer VPNs provide optimal routing solutions and full mesh topological redundancy for WAN-connected sites. There is no real additional planning or design for the implementation on the part of the customer. The provider will have already traffic engineered the network based on services offered and service level agreements (SLA) negotiated. Provisioning of additional sites is as simple as placing a router and dropping a local loop into place. The configuration does not require the creation of multiple VCs to provide the full mesh capabilities.

VPN Drawbacks

The cost and administrative overhead associated with a large enterprise full mesh Layer 2 topology is daunting on any scale. To reduce the number of VCs required, redundancy is sacrificed. Each site requires manual provisioning of a VC to get the required connectivity and traffic flow. Overlay VPNs also incur encapsulation overhead when IPsec or GRE tunneling is involved.

The chief benefit in the peer-to-peer VPN model is also, at times, its greatest drawback. The provider is involved in the customer routing process. Routing information is redistributed at CE and PE routers to be passed into or out of each respective network. Route filters should be placed on each interface to protect both parties from route floods sometimes caused by convergence events. The customer now must place additional trust in the capabilities of the service provider to properly configure and maintain their routing infrastructure. This can be problematic at times.

At critical sites with redundant routers and connections to the service provider network, care should be taken to ensure diversity in the connection so that both circuits do not land on the same PE router. The goal is to eliminate any single point of failure. Figure 11-4 illustrates this concept.

Figure 11-4 *Redundant Connections*

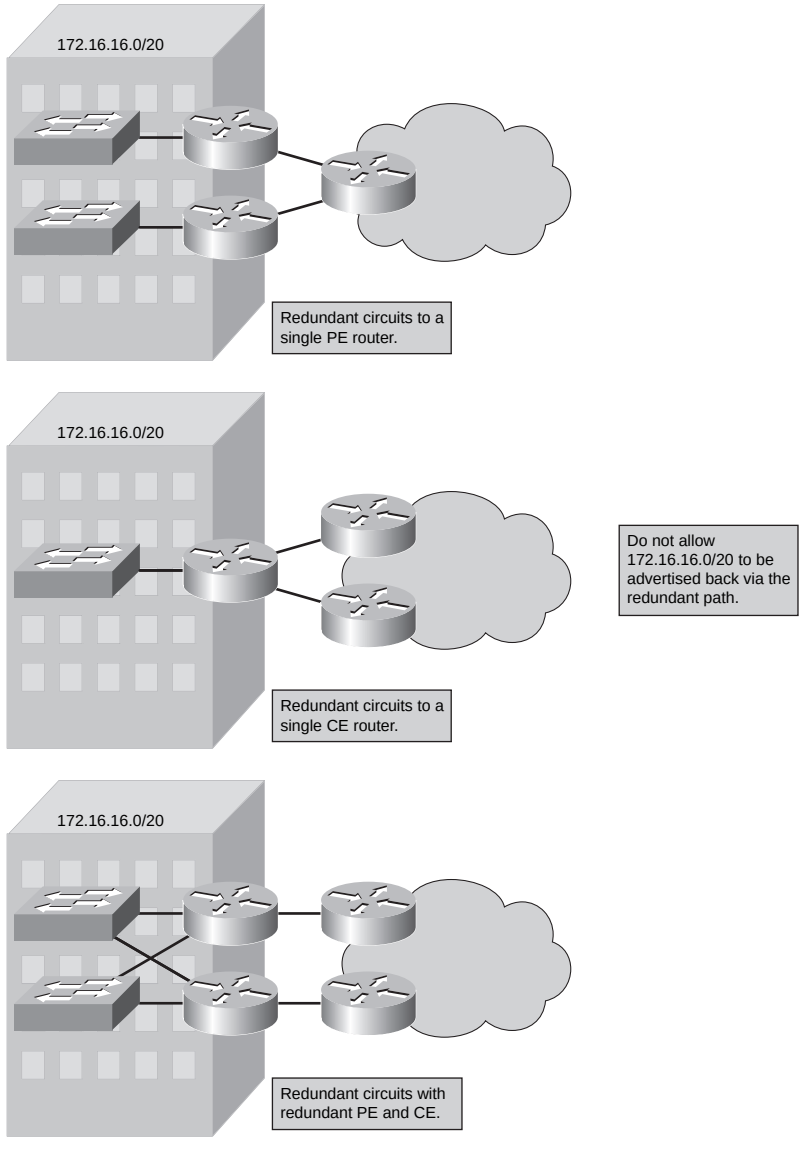

As Figure 11-4 points out, it is also necessary to ensure that routes advertised via one circuit are not redistributed out to the PE and then right back in via the redundant circuit to the CE. This will

cause a significant routing loop. Split horizon will not stop it, because the update is not received via the interface through which it was initially sent. Suddenly the routers have an erroneously valid path to the 172.16.16.0/20 subnet via a PE router. A simple inbound route filter blocking 172.16.16.0/20 on both CE routers or, more preferably, an outbound route filter on both PE routers will remedy the situation.

Another potential drawback is that PE routers will most likely be a shared resource. That is, there may be many other customers sharing the resources of a single PE. There are quite a few providers, however, that will negotiate a dedicated PE per customer connection. Peer-to-peer VPNs are very much a case of getting that which is paid for.

Along with resource allocation, the provider must be able to effectively deal with the fact that most, if not all, of its customers will be using RFC 1918 addressing. This makes the job of maintaining individual customer routing information that much more important.

With that in mind, customers can be sure that there is significant use of route filters throughout the provider network and that some degree of service degradation may occur due to such filtering, especially if done incorrectly.

MPLS VPNs

The MPLS VPN takes the best aspects of overlay VPNs and the best aspects of peer-to-peer VPNs and assembles them into a single product offering. MPLS VPNs are essentially peer-to-peer VPN implementations. Each customer's routing information is kept securely separate from every other customer's routing information through the use of a route distinguisher (RD) that is unique to a particular customer. The use of the RD allows the provider to give each customer a logically separate PE router, though not always physically separate. PE routers will remain a shared resource unless otherwise negotiated.

The customer routing information is maintained by a specific routing protocol instance tied to its RD. The routing table assembled by this routing protocol instance is known as a virtual routing and forwarding (VRF) table. In essence, it is simply an extension of the customer's routing table, because it includes all of the customer's advertised prefixes.

The following sections focus on terminology associated with MPLS VPNs, architectural needs of both the provider and customer networks, and some discussion on how a technology such as MPLS can maintain routing information for individual customers in a shared routing infrastructure environment.

MPLS VPN Terminology

Much of the terminology of MPLS VPNs has been discussed at one point or another in previous chapters. It is prudent to touch on it once more at this point to ensure that all of the terms associated with the technology are in the forefront of the mind while taking in the information in the remainder of the chapter.

- **C network**—The customer-controlled internal network.
- **CE router**—The customer edge router (also known as customer premises equipment, or CPE), which connects to a PE router.
- **Label-switched path (LSP)**—The pathway established for use by a label-switched packet through a P network in transit to a particular destination.
- **P network**—The service provider–controlled internal network comprised of core routers providing transport across the provider backbone but carrying no customer routing information.
- **P router**—A service provider MPLS core or backbone router with no customer-facing interfaces and carrying no VPN routes.
- **PE router**—A provider edge MPLS router containing customer-facing interface(s) and connecting to CE router(s) for the purpose of customer routing information exchange.
- **Penultimate hop pop (PHP)**—The final P router in the P network pops the label prior to the packet's arrival at the egress PE router.
- **PoP**—Service provider point of presence.
- **Route distinguisher (RD)**—A 64-bit identifier prepended to an IPv4 address to make it a globally unique VPNv4 address.
- **Route target (RT)**—An attribute appended to a VPNv4 BGP route to indicate VPN membership.
- **Virtual routing and forwarding (VRF) table**—A customer-specific routing table instance.

CE Router Architecture

Over the course of the discussions of the technologies involved in this chapter, the CE router will play an important role. Regardless of what designation is applied to it, the CE router is a router. It runs an IGP (available protocols include BGP, OSPF, EIGRP, RIP, or static routing) and exchanges routes with a neighboring router discovered through whatever routing protocol process the chosen protocol uses.

The CE router is not MPLS aware and does not participate in the MPLS architecture in any way other than the sending and receiving of customer routing information. The provider's MPLS P routers are similarly invisible to CE routers. The MPLS architecture simply appears to be an extension of an intra-company BGP routing implementation between WAN sites with little or no visibility beyond the customer-facing PE router interface.

All redistribution and MPLS-related manipulation will be done on the PE router and will remain transparent to the CE routers at each site. A CE router will be no different, architecturally or functionally, from any other router in the C network.

PE Router Architecture

The architecture of the PE routers in a provider's network is similar to that of a typical PoP in a dedicated peer-to-peer model. The major difference is that the architecture is compressed into a single device. The PE routers are usually relatively high-end routers such as the Cisco 7200VXR router.

Each customer is assigned its own RD and VRF table dedicated to maintaining routing information within the provider infrastructure. Routing across the provider backbone is performed by yet another routing process meant to bring some sense of simplification back into the picture in the form of a global IP routing table. The PE is managed as a single router but runs multiple instances of a routing protocol to maintain customer-specific routes and redistribute them into the global IP routing table. Figure 11-5 illustrates the concept of the PE router architecture.

Figure 11-5 *PE Router Architecture*

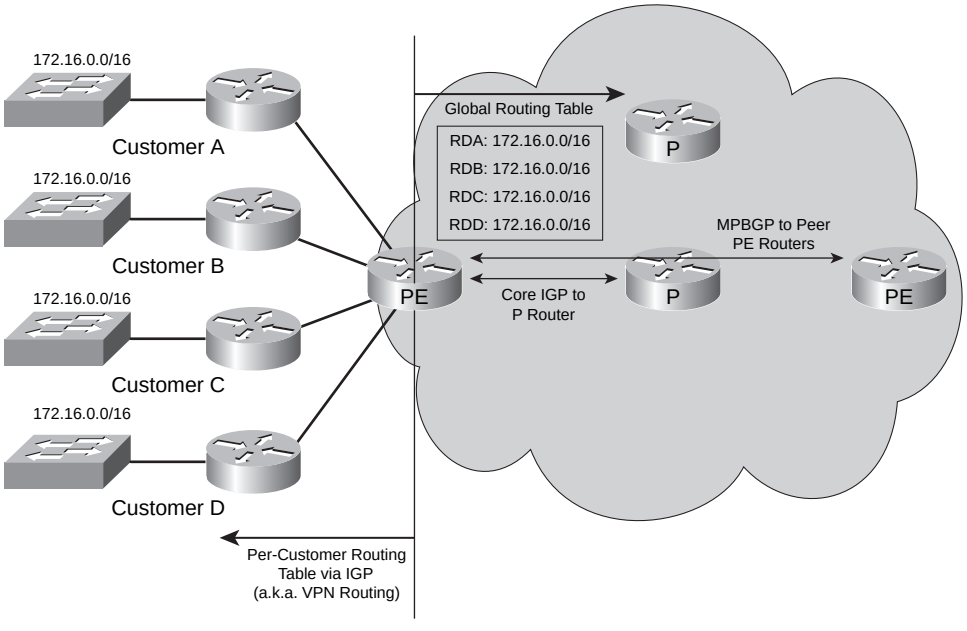

As Figure 11-5 shows, the VRF provides isolation between customer routes. The information from these routing tables still must be exchanged between various PE routers. Therefore, a routing protocol is needed that will allow the transport of all customer routes across the P network while allowing the continued independence of each customer's address space.

The decision was made that a single routing protocol be run between PE routers that will exchange customer routes without the involvement of the P routers. The PE routers that connect to a given customer network will be peered to each other and routes will be exchanged. With this model, the number of routing protocols between PE routers need not increase in proportion to the number of customers served.

This also has the added benefit of keeping the customer routes off of the P routers as they are unicast from peer to peer.

The number of prefixes advertised by each customer, when added to those P network routes already in existence, can combine to create an excessively large routing table overall. BGP is the only protocol with the scalability to handle these types of operations while giving the most flexibility in manipulation of routing and traffic flow in general. BGP neighbor relationships are configured between PE routers directly so that prefixes can be exchanged for a given customer. The global IP routing table in the P network need not actually carry any of the actual customer routes.

P Router Architecture

P routers make up the backbone of the P network. They do not carry VPN routes and do not participate in MPLS routing. They do provide transport for traffic between PEs but that is essentially where their job stops. They run a routing protocol such as IS-IS, OSPF, or BGP across the provider backbone and carry only P network routing information in their routing tables. They interface with PE routers to facilitate the transport of BGP peering information across to remote PE routers.

BGP is typically the protocol of choice for P networks due to its scalability and functionality, not for any MPLS-related need or requirement.

Route Distinguishers

On PE routers, there is obviously a need to deal with the fact that most, if not all, customers will be using RFC 1918 addressing and that all that common space will be allocated in varying manners. So, there is a need to be able to keep individual customer routes separate and distinct so that each network is reachable. One customer's 10.1.1.0/24 subnet will likely co-exist with another customer's 10.1.1.0/24 subnet, for example. These will obviously have differing outbound interfaces.

An RD allows these prefixes to be kept unique. The RD is a 64-bit identifier that is tacked on to the front of the IPv4 address. These VPNv4 addresses are advertised between BGP peers on PE routers. The BGP implementation known as Multiprotocol BGP (MPBGP) supports address families other than IPv4 addressing. This creates a 96-bit entity known as a VPNv4 address. Figure 11-6 illustrates the mechanics involved.

Figure 11-6 *PE Peers*

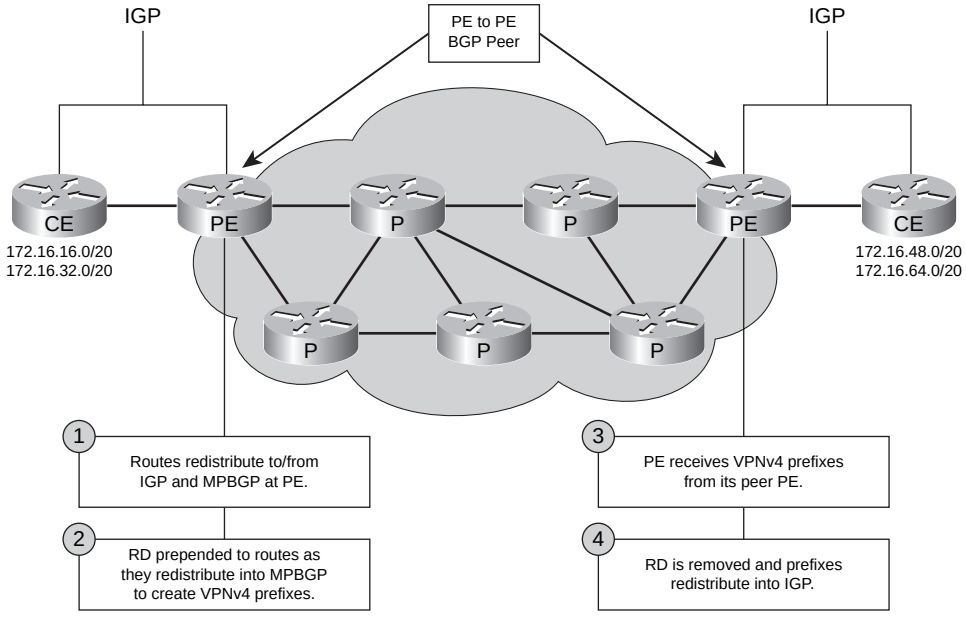

An IGP running across the local loop serves to move customer routing information between the PE and CE routers. This routing information is redistributed into MPBGP where the prefixes are converted to VPNv4 addresses. The PE routers are peered directly to each other via an Interior BGP (IBGP) peering so that they exchange routes directly with one another. Once the neighbor PE receives VPNv4 information from its peer, the RD is removed so that routes can be redistributed back into the customer IGP and sent to the CE router for propagation through the enterprise.

RD values have no real specific meaning. They are only meant to allow the routing architecture to deal with overlapping address space. So long as each is unique within the P network, there should be no risk of route overlap. Because there has to be a unique mapping between the RD and the VRF, the RD can be viewed as the VRF identifier in Cisco implementations.

Usually, each customer has a single RD assigned to its prefixes. There are times, however, when customers will want to protect interdepartmental routing information or business-to-business

connectivity via an MPLS VPN. A single RD per customer would preclude some scenarios and create a need for a more versatile form of management. On the surface, this would seem to require the use of multiple RDs and redistribution of the desired routes between the VRFs. This is indeed the case.

Consider a deployment of an enterprise Voice over IP (VoIP) solution managed by the provider, similar to that shown in Figure 11-7.

Figure 11-7 *VoIP Service Example*

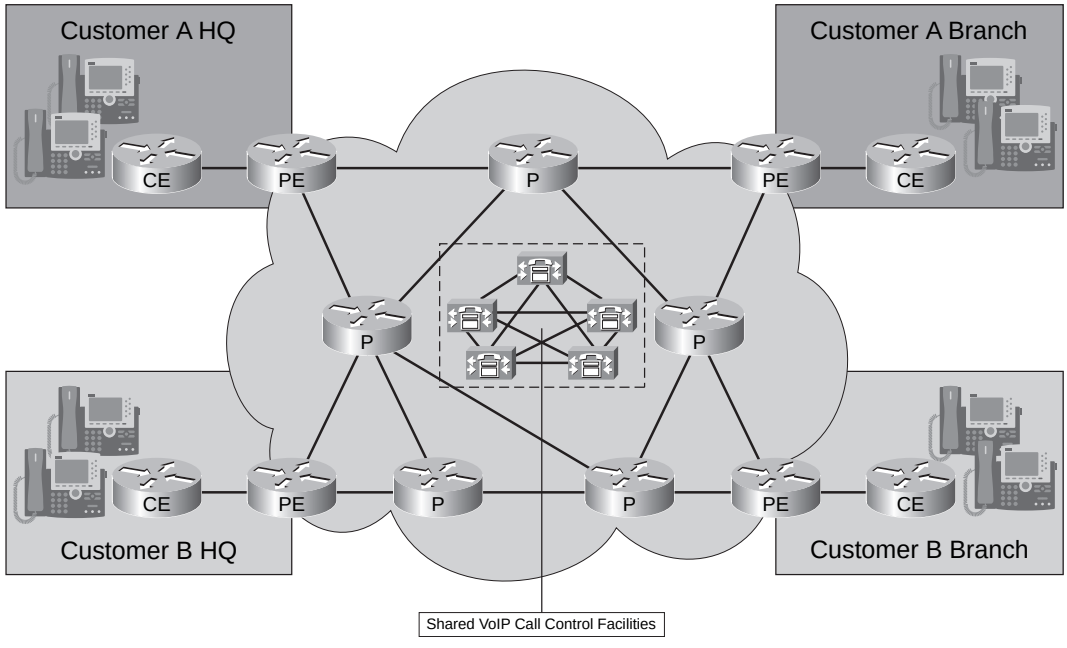

The provider would be responsible for all call control for both customer-internal calling between sites and PSTN calling. The provider would also have particular designs for calling between customers across the network. These calls are no different from typical Public Switched Telephone Network (PSTN) calls to each customer, but the traffic need never leave the provider's network if both are MPLS VPN customers.

Because some or all customers would share a common call-control facility, certain routing changes would be necessary to ensure that all customers can reach this common point inside the provider network. A single RD would preclude this capability. In some cases, the provider would institute a specific voice RD for reachability to the shared call-control and PSTN gateway devices. Firewalls, ACLs, and more would be necessary to ensure security of all signaling and media traffic so that no unauthorized traffic would be able to traverse the alternate RDs.

In such an example, both Customer A and B sites would be participating in their own customer-specific VPN as well as the shared voice VPN. To mitigate the possibility for unauthorized access or activity, the Customer A and B branch sites may route in hub-and-spoke fashion via the HQ site to place and receive voice calls. This would mean that the branch sites would participate only in the customer-specific VPNs, leaving the HQ sites as the sole point of contact with any shared infrastructure.

Route Targets

To indicate that a site participates in multiple VPNs, a method is needed in which a set of VPN identifiers can be attached to a route to indicate that membership. An RD is adequate for a single VPN. Route targets (RT) were introduced to facilitate a more complex VPN topology. An RT is an additional attribute that is attached to a VPNv4 BGP route to indicate VPN membership.

The RT is appended at the time the IPv4 route is converted to a VPNv4 route by the PE router. RTs attached to routes are called *export RTs* and are configured separately for each VRF in a PE router. Export RTs identify the VPNs to which the sites associated with a particular VRF belong.

Import RTs are those RTs that specify the routes associated with a particular VRF. When VPNv4 routes are propagated to neighboring PE routers, routes meant to be imported into a particular VRF need to be selected. This is accomplished based on import RTs. Each VRF in a PE router can have multiple import RTs identifying the set of VPNs from which the VRF is accepting routes. In cases of overlapping VPN topologies, RTs are used to identify VPN membership and allow for more complex scenarios.

With this implementation, as the CE router advertises routes to the PE router, the inbound routes are prepended with the RD to create VPNv4 addresses, and then the RTs are appended based on VPN membership. These routes are exported into the appropriate VRFs for propagation to the remote PEs. Routes will be imported by remote PEs based on import RT values and redistributed to the remote CE routers.

End-to-End Routing Update Flow

Now that all of the pertinent pieces of the MPLS VPN puzzle have been introduced, a final walk through the routing update flow is in order. Figure 11-8 provides a visual aid for the flow of the discussion.

Figure 11-8 *End-to-End Routing Updates*

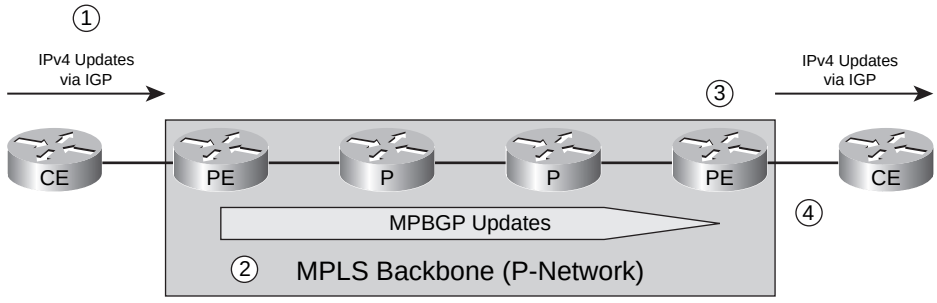

In Figure 11-8, there are four designated steps in the routing update process:

Step 1 PE routers receive IPv4 routing updates from the CE router via a configured common IGP. These routes are installed in the appropriate VRF table.

Step 2 Customer routes from the VRF are exported as VPNv4 routes into the MPBGP instance and propagated to other PE routers. To become VPNv4 routes, RDs must be prepended to the route entries. To be exported, export RTs are appended to specify VPN membership.

Step 3 The PE routers receiving MPBGP updates import the incoming VPNv4 routes into the appropriate VRFs according to the values specified by the import RTs attached to the routes and the individual VRF tables.

Step 4 The VPNv4 routes installed in the VRF table(s) are redistributed into the IGP instance running between PE and CE and then propagated to the CE and into the C network.

From the CE standpoint on both sides of the P network, the P network simply looks like any other routing instance. The CE routers have no visibility to the MPLS network or its structure. Once routing updates are successfully flowing, end-user traffic can begin to flow as well.

MPLS VPN Packet Forwarding

PE routers use a two-label stack to label the VPN packets for forwarding across the P network. The label stack is imposed by the ingress PE router.

The top label in the stack will be used by LDP for P network traversal along an LSP that will get it to the egress PE router. The S-bit in the top label will be set to 0.

The second label will be assigned by the egress PE router. Remember, the label values are downstream-assigned. The purpose of the second label is to tell the router how to forward the

incoming VPN packet. This label could point to a particular outbound interface or to a VRF table. If the label points to an outbound interface, a label lookup is performed on the VPN packet itself. If a VRF table pointer is specified, a label lookup is performed to find the target VRF instance. An IP routing lookup is then performed within that VRF instance. The S-bit in the second label will be set to 1. The S-bit is the "end-of-stack" pointer. When set to 0, there will be further labels in the stack. The bottom label in the stack will have the S-bit set to 1, indicating its position as the last label.

Either method is acceptable. The second label in the stack points to an outbound interface when the CE router is the next hop in the VPN route. The second label points to a VRF table for aggregate VPN routes, VPN routes to the null interface, and directly connected VPN interfaces.

The P routers perform label switching based only on the top label. They never see the second label because they do not analyze the structure any further than the first label.

The egress PE performs a label switch on the second label because the first one has been popped. It will then forward the packet according to the parameters of the packet, which point it to a VRF or an outbound interface.

MPLS VPN PHP

It seems rather inefficient for the egress PE to deal with both labels. The use of PHP allows the final P router in the LSP to pop the label, thereby relieving the egress PE router of the need to do so. This allows the egress PE router to simply perform its function using only the VPN label in the stack. Once that label is removed, an IP routing lookup can take place and the packet can be forwarded.

Foundation Summary

MPLS VPNs are somewhat of a departure from traditional WAN technologies. However, the benefits of being able to deploy a fully Layer 3–aware WAN topology with built-in redundancy is very alluring. The possibilities for service and application offerings by both providers and enterprise customers are exceedingly diverse.

Service provider offerings such as firewall-in-the-cloud and managed voice service are just the beginning of what is possible with a creative architect.

A great deal of information has been covered in a short span in this chapter. The information that follows serves to summarize the key points discussed herein. Table 11-2 revisits the roles of routers in MPLS VPN architectures.

Table 11-2 *MPLS VPN Router Roles*

Router	Location	Purpose	Description
C router	C network, internal	Maintains C network routes and forwards traffic	A router internal to the customer-controlled network
CE router	C network, edge	Exchanges C network routes with a PE router	A customer-controlled router that interfaces and exchanges routing information with a PE router
P router	P network, internal	Maintains P network routes and forwards traffic	A router internal to the provider-controlled network, usually an LSR
PE router	P network, edge	Exchanges VPN routes with CE router	A provider-controlled router that interfaces and exchanges routing information with a CE router

Various protocols are present in MPLS VPN architectures. Table 11-3 provides a snapshot review of them as they pertain to the MPLS technologies.

Table 11-3 *MPLS VPN Related Protocols*

Protocol	Where	Description
Customer IGP	C network and CE-PE router connection	The customer internal routing protocol used to maintain routing information throughout the enterprise
Provider IGP	P network	The provider internal routing protocol used to maintain routing information, usually BGP, IS-IS, and/or OSPF
MPBGP	PE-to-PE peering	Multiprotocol BGP maintaining peer connections between PE routers for the express purpose of propagating C network routing information

Q&A

The questions and scenarios in this book are designed to be challenging and to make sure that you know the answer. Rather than allowing you to derive the answers from clues hidden inside the questions themselves, the questions challenge your understanding and recall of the subject.

Hopefully, mastering these questions will help you limit the number of exam questions on which you narrow your choices to two options, and then guess.

You can find the answers to these questions in Appendix A. For more practice with exam-like question formats, use the exam engine on the CD-ROM.

1. Consider a traditional Layer 2 overlay VPN. List some technologies and possible topologies that are available for such implementations.
2. What is the primary benefit of a peer-to-peer VPN over a Layer 2 overlay VPN?
3. When using redundant connections at a single site, what are some pitfalls that should be avoided?
4. Consider Figure 11-9. The routing entry for 192.168.1.0/24 needs to make its way to the routing table of Router B. Trace its path from left to right, explaining the process.

Figure 11-9 *MPLS Routing Information Flow*

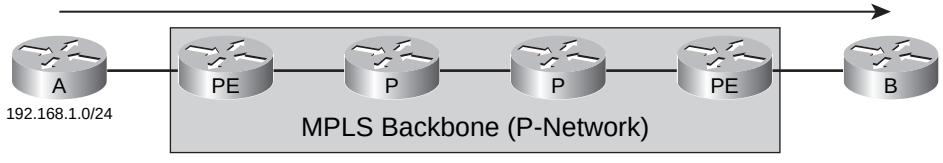

5. Consider Figure 11-10. Now that the 192.168.1.0/24 network is known in Router B, the host at 192.168.5.3 would like to ping the host at 192.168.1.5. Trace the path of the first ICMP echo-request packet from 192.168.5.3 to 192.168.1.5 from CE to CE. Assume that any and all address resolution activities have been successfully completed and that full routing convergence has been reached.

Figure 11-10 *End-to-End Traffic Flow Over MPLS*

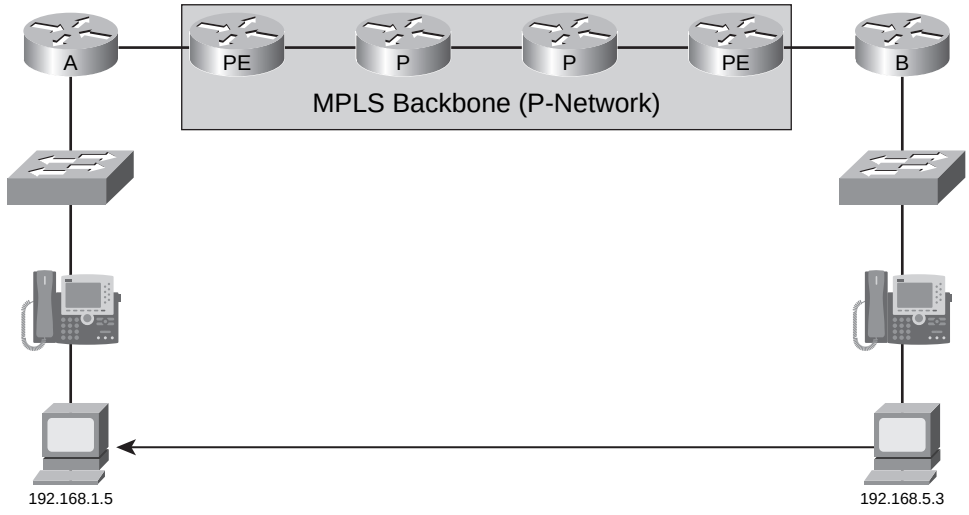

This part of the book covers the following ISCW exam topics:

Implement a site-to-site IPSec VPN.

- Describe the components and operations of IPSec VPNs and GRE Tunnels.
- Configure a site-to-site IPSec VPN/GRE Tunnel with SDM (i.e., preshared key).
- Verify IPSec/GRE Tunnel configurations (i.e., IOS CLI configurations).
- Describe, configure, and verify VPN backup interfaces.
- Describe and configure Cisco Easy VPN solutions using SDM.

Part III: IPsec VPNs

Chapter 12 IPsec Overview

Chapter 13 Site-to-Site VPN Operations

Chapter 14 GRE Tunneling over IPsec

Chapter 15 IPsec High Availability Options

Chapter 16 Configuring Cisco Easy VPN

Chapter 17 Implementing the Cisco VPN Client

Exam Topic List

This chapter covers the following topics that you need to master for the CCNP ISCW exam:

- **IPsec**—Internet Protocol Security (IPsec) is a suite of protocols that can provide data confidentiality, data integrity, and data origin authentication to IP packets.
- **Internet Key Exchange (IKE)**—A framework used to exchange security parameters and authentication keys between IPsec endpoints.
- **Encryption Algorithms**—Mathematical algorithms (and the associated keys) used to make data unreadable to everyone except those who have the proper keying material.
- **Public Key Infrastructure**—A hierarchical framework for managing the security attributes for devices that engage in secure communications across a network.