

## CHAPTER 3

# TECHNOLOGY OVERVIEW: MAKING THE TECHNOLOGY CASE FOR MPLS AND TECHNOLOGY DETAILS

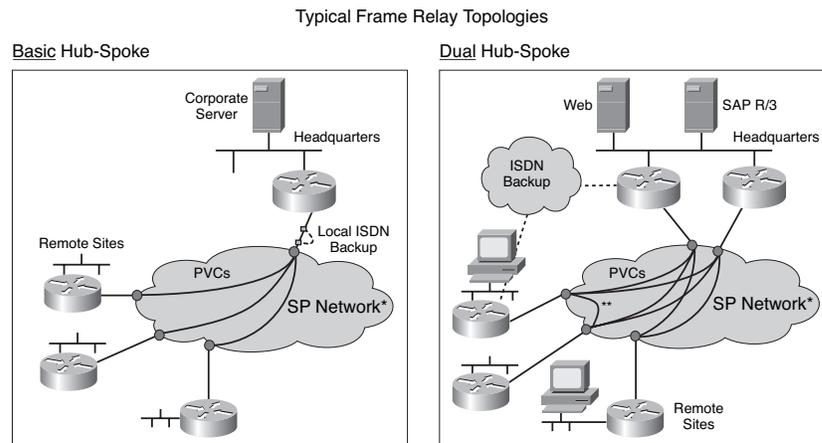
This chapter highlights all the available technologies for creating the services described in the previous chapters. It provides pros and cons for each option and builds a case for multiprotocol label switching (MPLS) as a baseline technology for service creation. It also discusses the MPLS technology details. From a service provider perspective, it is pivotal that MPLS as a technology has been adopted by service providers as a key architectural component for next-generation networks (NGNs) because it is an enabler for services based on IP. For enterprise organizations, the virtualization capabilities inherent in MPLS facilitate LAN/WAN segmentation rather than the implementation of static circuits and mechanisms that can be costly in the end.

## Available Technologies and Options

Layer 2 technologies, such as Frame Relay and ATM, have long been deployed to provide a VPN-like service. The attributes of both technologies are quite similar, as follows:

- A virtual circuit has bi-directionality.
- A virtual circuit is established via signaling.
- A fixed hierarchy exists of a virtual path or virtual circuit.
- The virtual circuit is connection oriented and not tied to an IP control plane.
- A single route exists between the source and destination.
- A full-mesh of VCs is required to have any-to-any connectivity.

A typical topology for Layer 2 implementations has been hub and spoke, in which all VCs terminate at a central location—for example, at the enterprise headquarters. Hub and spoke topologies are depicted in Figure 3-1.



**Figure 3-1** Typical Frame Relay Topologies

The attributes of a Layer 2 technology, such as Frame Relay, include the following:

- Secure, closed user group connectivity exists amongst corporate sites.
- Statistical performance guarantees throughput via permanent virtual circuit (PVC) constructs with a committed information rate (CIR) and excess information rate (EIR).
- Approximately 80% of the traffic over a Frame Relay network is IP.

As an *unbundled* service, Frame Relay is Layer 2–centric where the target market consists of enterprise customers who implement their own corporate virtual private networks (VPNs). The enterprise purchases a PVC from a service provider; consequently the enterprise is responsible for designing the VPN topology and managing the customer edge router (CE) IP routing, quality of service (QoS) policies, and application prioritization. For a service provider, Layer 2 virtual circuits are easy to sell, manage, and bill.

Another type of service using Frame Relay technology (there is a similar service in ATM service) is a *bundled* Frame Relay managed router service, which has a look and feel similar to that of an IP VPN. The target market is customers who want to outsource a VPN (Layer 2-based) to a service provider. The enterprise customer subscribes to Layer 2-based VPN services and is not involved in the PVC complexity discussions. The service provider must manage the PVC complexity, the corresponding topology, and the CE and address customer routing, application prioritization, and service level agreement management issues.

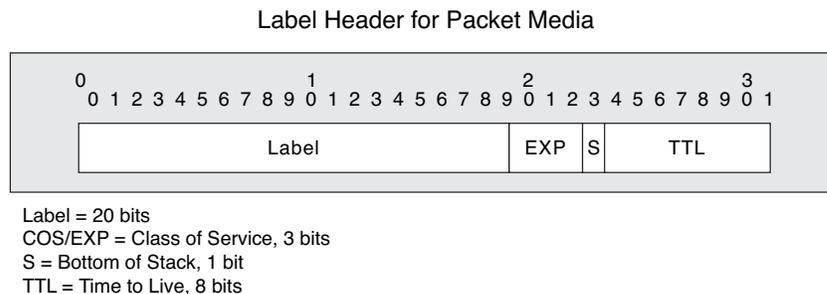
What are the possible limitations of a Layer 2 technology, such as Frame Relay, as customers request value-added services, such as a service provider–hosted IP telephony? The service provider must provision a full mesh of PVCs among all sites—for example, a VPN with 50 sites would require 1225 PVCs. Due to the requirement to prioritize Voice over IP (VoIP), the service provider must deploy separate voice and data PVCs. With shared service provider–hosted PBXs and offnet gateways, the service provider must provision PVCs from each customer site to the service provider data center. As a result, enterprise customers often do not accept a bill for the cost-prohibitive PVC mesh and the service provider consequently bears the cost itself. So, scalable value-added service architecture is needed, and MPLS technology possesses attributes that contribute to a scalable architecture for managed VPNs with value-added service elements.

## Why MPLS? (High-Level Detail)

Multiprotocol label switching architecture, as discussed in IETF RFC 3031, combines the benefits of the hardware packet switching approach of ATM and the Layer 3 approach of IP. The MPLS architecture separates the control information for packets required for packet transfer itself; that is, it separates the control and data planes. In traditional IP routing, a packet is assigned in each router to a particular flow corresponding to a class of routing or a forward equivalence class (FEC). In contrast, in MPLS this assignment is performed once at the entry, or *ingress*, to the MPLS network. In an MPLS network, the FEC is identified by the network exit destination, or *egress*, and by the ingress label-switched router (LSR).

The FEC consists of a simple group of IP destinations for which a transfer can be managed in the same manner and which is assigned a fixed-length identifier called a *label*. The path corresponding to each FEC between the ingress and egress LSRs is called a *label-switched path (LSP)*. An FEC, therefore, determines how packets are mapped to an LSP. This means that a packet entering at the egress LSR of an MPLS domain is assigned to an FEC following the analysis of the IP header.

A label is assigned to the FEC imposition operation either by tagging an existing field or as a complement in the packet header. The label is pivotal to the establishment of the LSP through all the routers or switches in the MPLS domain. Each LSR analyzes the incoming packet label. Then after consulting a label table that permits it to recognize the LSP, the LSR switches the packet to the next LSR after changing the value of the label. The label is removed at the egress LSR or a disposition operation is performed. By definition, an LSP is *unidirectional*—that is, two LSPs are required to support bi-directional traffic. We can compare the MPLS behavior as a Layer 2 switching approach to ATM and a Layer 3 routing approach to IP. Figure 3-2 depicts the actual MPLS label.

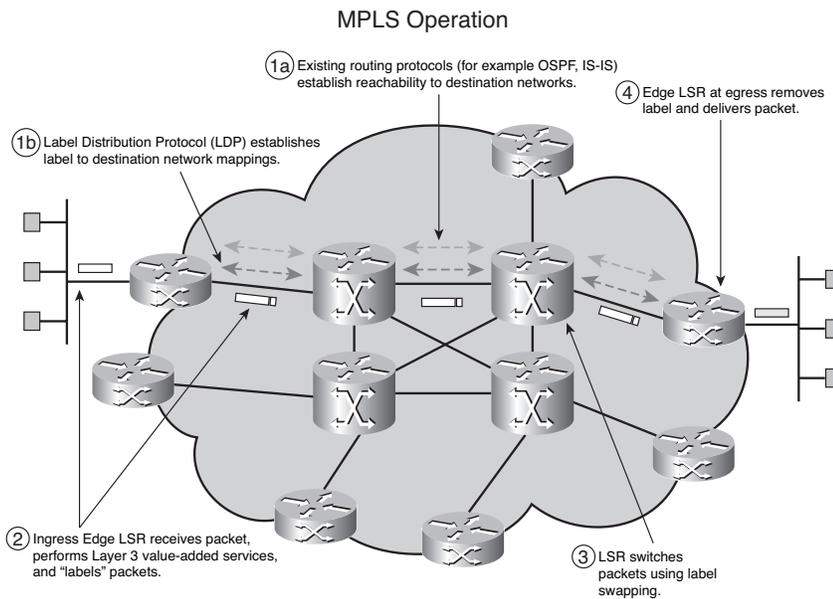


- Can be used over Ethernet, 802.3, or PPP links
- Uses two new Ethertypes/PPP PIDs
- Contains everything needed at forwarding time
- One word per label

**Figure 3-2** *Label Header for Packet Media*

The tag frame encapsulation uses a *shim header*. The shim is a header that sits between a transport header and the Layer 3 header in the packet. The label format is generic because it can be used on link layers, such as Ethernet 802.3, PPP, Frame Relay, ATM PVC, and so on. The label value consists of four octets, although several labels can be assigned to a packet, because of a concept called

*label stacking.* The label can be tagged in AM in the virtual circuit identifier (VCI) and virtual path identifier (VPI) fields of the cell headers. For Ethernet, PPP, FDDI, and other technologies, an interposed header (shim) is located between the link and network headers and is used to transport the label. The LSR performs the control and transfer functions, and the transfer element uses fixed-length labels. These labels are memorized in a table with a path indication for outgoing packets. The control element consists of network-layer routing protocols and one or more label allocation mechanisms. Figure 3-3 summarizes the fundamental MPLS operation.



**Figure 3-3** *MPLS Operation*

As previously mentioned, one of the key advantages of the MPLS architecture is the separation into two planes—the data plane that contains the information required to transfer a packet and a control/signaling plane that allocates the transfer information. The data plane is used for the transport of packets (or label swapping algorithm), and the control plane is analogous to routing information (for example, the location to which to send the packet). This capability is

programmed into hardware by the control plane. This separation permits applications to be developed and deployed in a scalable and flexible manner. Examples of applications that are facilitated by MPLS technology include the following:

- **MPLS QoS**—This implements a quality-of-service mechanism that enables the creation of LSPs with guaranteed bandwidth.
- **BGP VPNs**—Border Gateway Protocol (BGP) is used to exchange FEC-label binding. Further, a service provider can use BGP in its network with IP routing protocols or static routing between the service provider and the customer to create a Layer 3 VPN service.
- **Traffic engineering**—Traffic engineering enables one to control traffic routing via constraint-based routing. Constraint-based routing enables a demand-driven, resource-reservation aware, routing paradigm to co-exist with current topology-driven hop-by-hop Internet interior gateway protocols.
- **Multicast routing**—Protocol Independent Multicast (PIM) is the control protocol used to create FEC tables; extensions of version 2 of the PIM protocol are used to exchange FEC-label binding.
- **Pseudowires**—These can be used to evolve legacy networks and services, such as Frame Relay, ATM, PPP, High-Level Data Link Control (HDLC), and Ethernet. Traffic is accepted into the network via a variety of access technologies, labeled at the edge, and transported over a common MPLS core. At the network egress, the label is removed and delivered in a manner similar to the ingress implementation.
- **Generalized MPLS (GMPLS)**—The goal of GMPLS is to integrate control of the routing layer with that of the optical transmission layer, thus facilitating the implementation of traffic engineering across the network. Optical cross-connect platforms do not examine traffic passing through them—in contrast to routers, for example. GMPLS deployment links capacity provisioning in the optical layer for an automated execution of resource reservation (for example, bandwidth brokering and provisioning).

As an overview, MPLS uses label swapping rather than conventional IP routing. IP routing defaults might appear, such as in traffic engineering, where the establishment of optimal routes and the analysis of available bandwidth on the various links are necessary to optimize the use of network resources. Although conventional IP routing can examine the optimal route by applying metrics, it cannot analyze the available bandwidth on the individual links.

The term *traffic engineering* refers to the specific actions performed to ensure that the express demand remains within the available capacity of network resources. These actions include routing, dimensioning procedures, and demand estimation. Current routing on IP networks is based on computing the shortest paths, where the “length” of the link is determined by an administrative assigned weight. If the traffic matrix (defining expected demand between all network end points) is known, then by appropriately setting the value of these weights, you can ensure that traffic flows are routed optimally. For example, you can ensure that available capacity is used to its maximal effect.

MPLS offers additional possibilities for routing traffic over links with sufficient capacity. LSPs completely specify the path for broadly defined traffic aggregates (defined by source and destination addresses, for instance). These attributes can be constructed in real time as required or by network management procedures by using an estimated traffic matrix. Route selection can be performed end-to-end by the edge routers or on a hop-by-hop basis.

## MPLS and Quality of Service

For QoS, the integrated services model (InServ) specifies two classes of services—controlled load (CL) and guaranteed service (GS)—and uses a signaling protocol known as Resource Reservation Protocol (RSVP). Briefly, the quality of CL end-to-end connections (IETF RFC 2211) is intended to be equivalent to that provided by the traditional best effort service in a lightly loaded network. Here is an example: A large percentage of packets is successfully transmitted to the recipient and latency is no greater than the minimum delay for packets circulating in a lightly loaded network. To ensure compliance with these conditions, applications addressing CL requests (via RSVP) supply the network with an estimate of the traffic they are likely to generate via the parameters of a

“leaky-bucket.” This so-called *traffic specification (Tspec)* is used by each network node on the flow path to carry out admission control. The following are possible mechanisms for implementing CL:

- **Priority queuing**—It uses two queues, a high priority queue subject to CL traffic admission control and a best-effort queue.
- **Weighted fair queuing (WFQ)**—It enables you to regulate the way link capacity is shared between various traffic flows. All flows have access to the full connection bandwidth, but when several flows have packets in the queue, the service rate of each flow is proportional to its assigned “weight.” By selecting the appropriate weights, you can therefore reserve capacity for CL more efficiently.
  - **Class-based queuing (CBQ)**—This is an alternative algorithm that also permits rate control for various classes of traffic.
  - **Random early detection (RED)**—This protects CL traffic to some extent from any unresponsive best-effort flows.

RED is an active queue management mechanism that tends to ensure a fairer distribution of bandwidth between contending flows.

Additionally, low latency queuing (LLQ), which is in fact Class Based Weighted Fair Queuing with a Priority Queue (known as PQCBWF), is a critical mechanism that supports both data class of service and VoIP.

- **Weighted random early detection (WRED)**—This combines the capabilities of the RED algorithm with IP precedence. This combination provides for preferential traffic handling for higher-priority packets. It can selectively discard lower-priority traffic when the interface starts to get congested and can provide differentiated performance characteristics for different classes of service. WRED is also RSVP aware and can provide an integrated services controlled-load QoS.

The guaranteed service (IETF RFC 2212) permits applications with strict requirements for both assigned bandwidth and packet delay. It ensures that all packets are delivered within a given time and not lost as a result of queue overflow. This service is first invoked by the sender, who specifies the Tspec and QoS requirements. Resource reservation is performed in the reverse direction with the receiver specifying the desired level of service (Rspec). As for CL, Tspec corresponds to the parameters of the leaky-bucket.

The InServ model did not achieve the success anticipated because its implementation is much more complex than the best-effort model. The fact that all routers must be RSVP-capable and able to store the details of every reserved CS and GS flow, although feasible on small networks, makes it unwieldy when applied to large backbones. Additionally, the guarantees defined in the two service classes tend to be either too strict (GS) or too vague (CL) for most applications.

The differentiated services model (DiffServ) relies on a broad differentiation between a small number of service classes. DiffServ support over MPLS is documented in IETF RFC 3270. Packets are identified as belonging to one class or another via the content of the differentiated services (DS) field in the IP header. Packets are generally classified and marked at the network edge depending on the type of service contract or service level agreement (SLA) between the customer and the service provider. The different classes of packet then receive different per-hop behaviors (PHBs) in the network core nodes. Service differentiation, therefore, implies differential tariffs depending on the QoS offered to flows and packets belonging to different classes. The DiffServ architecture consists of a set of functional elements embodied in the network nodes, as follows:

- The allocation of buffering and bandwidth to packet aggregates corresponding to each PHB
- Packet classification (FEC)
- Traffic conditioning, metering, and shaping

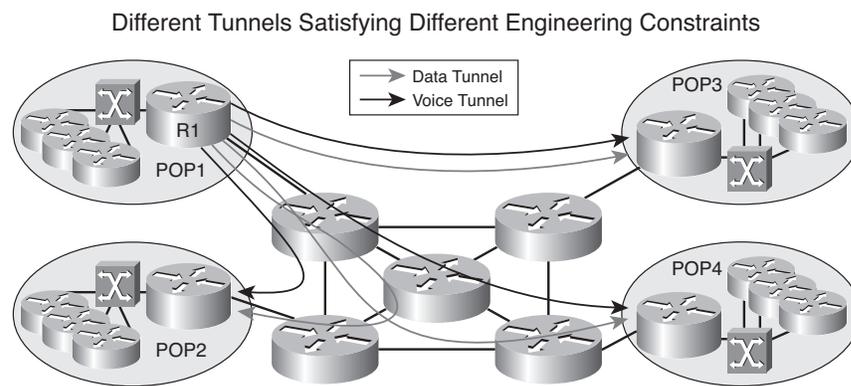
The DiffServ architecture avoids the requirement to maintain per-flow or per-user state within the network core, as is the case of InServ. The DS field (IETF RFC 2474) replaces existing definitions in the type of service (TOS) byte in IPv4 and the traffic class byte in IPv6. Six bits of the DS field are used in the form of the DS code point (DSCP) to identify the PHB to be received by a packet to each node.

Packets must first be classified according to the content of certain header fields that determine the aggregates defined in the user's SLA. Each aggregate is checked for conformity against SLA traffic parameters, and the contents of the DSC field are suitably marked to indicate the appropriate level of priority and PHB. The flow produced by certain aggregates can be reshaped to make these conform to the SLA.

In addition to best effort, considered to be the default PHB, two other PHBs have been defined by the IETF: expedited forwarding (EF) (IETF RFC 2598) and assured forwarding (AF) (IETF RFC 2597). These attributes are further discussed in Chapter 9, “Quality of Service.” Service implementations using DiffServ include a virtual leased line for Vo IP via EF PHB and a so-called Olympic service using the AF PHB group where the four AF classes are used to create four service qualities referred to as platinum, gold, silver, and bronze.

### Differentiating Service with Traffic Engineering

Deploying different tunnels satisfying a variety of engineering constraints can be done via DiffServ traffic engineering (DS-TE). Figure 3-4 depicts the implementation of DiffServ traffic engineering.



**Figure 3-4** *Different Tunnels Satisfying Different Engineering Constraints*

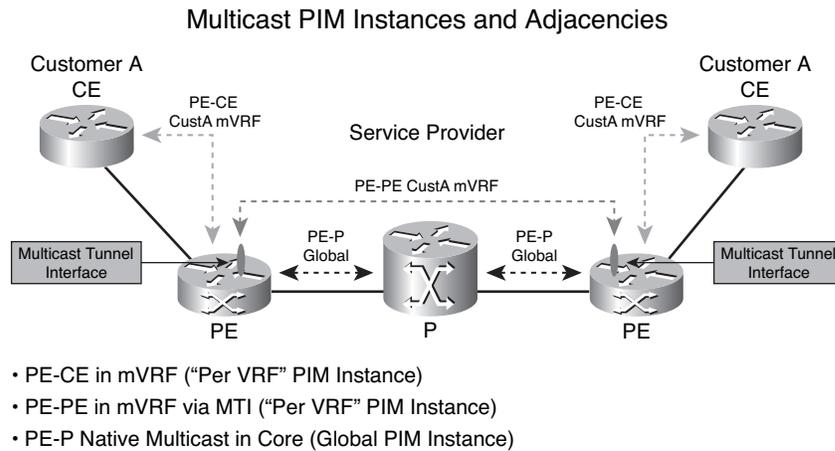
For example, with DS-TE in Figure 3-4:

- R1 can build a voice tunnel and a data tunnel to every POP.
- If R1 sends a data packet in a data tunnel (with EXP = Data), it gets the correct QoS for data.
- If R1 sends a voice packet in a voice tunnel (with EXP = Voice), it gets the correct QoS for voice.

Class of service–based traffic engineering tunnel selection (CBTS) provides a mechanism for dynamically using different tunnels—that is, dynamically steering packets to the designated DS-TE tunnel depending on the destination or class of service (CoS). Therefore, CBTS involves minimum configuration and automatic routing and rerouting when required. CBTS complements DS-TE to achieve dynamic QoS-based routing over an MPLS core where each CoS is transported over a tunnel engineered for its specific requirements; finally, CBTS achieves strict QoS with “right-provisioning” using the mechanism available with this technology, instead of wasteful “over-provisioning.”

## **Multicast**

For multicast VPN (MVPN) implementation, the VPN multicast flow is encapsulated inside an IP multicast GRE packet at the provider edge (PE) replicated inside the MPLS cloud. This encapsulation and replication are performed via regular IP multicast methods toward the far PE, which unwraps the GRE packet to obtain the customer multicast packet. The multicast destination of the GRE packet is unique per multicast domain (that is MPLS VPN). Two kinds of multicast trees can be created in the core: default-mdt and data-mdt. The default-mdt is the basic vehicle that allows the VPN routing and forwarding (VRFs) in the PEs to establish PIM neighbor relationships and pass multicast data between the PEs. All the multicast-enabled PEs of a VRF are members of the default-mdt. The “all” requirement means that PEs that are not interested in particular (S,G) flow still get it. The data-mdt is a traffic-triggered multicast tree created separately from the default-mdt that consists only of the PEs that want to get a particular customer (S,G). Figure 3-5 summarizes the multicast VPN implementation.



*Figure 3-5 Multicast PIM Instances and Adjacencies*

We have provided an overview of the MPLS operation with traffic engineering, quality of service, and multicast descriptions for use in an MPLS-based network. The next section discusses the benefits of MPLS as a technology foundation for service development and deployment.

## Benefits

This section focuses on MPLS technology as a service building block and foundation for enterprise virtualization implementation.

MPLS offers the following benefits for service providers and enterprises:

- Flexible classification of packets and the optimization of network resources.
- Label distribution through various protocols such as BGP, LDP, RSVP, and PIM.
- The coexistence of different distribution protocols in the same LSR.

- The redundancy of numbering and global label allocation, as labels that have only a local significance.
- The introduction of modular value-added applications such as traffic engineering, quality of service, multicast, and VPN.
- Facilitation of the evolution of legacy services via Any Transport over Multiprotocol Label Switching (AToM) and even the introduction of Layer 2 VPNs as the cost of retaining Frame Relay and ATM infrastructures becomes prohibitive.
- Unification of optical and routing control planes in GMPLS to evolve SDH and Sonet services. Also, GMPLS is used to generalize the MPLS control plane over many types of transports, including packet-type networks.

MPLS, therefore, provides the predictability of routing performance required to support differentiated services and the capability to offer tight SLAs associated with these differentiated service constructs. MPLS facilitates the integration of multiple services over a common switching platform, therein contributing to the reduction of operating expense. MPLS traffic engineering can reduce the management burden for IP-based services via the creation of backup paths and by facilitating the deployment of VoIP VPNs.

Path diversity can result in unpredictability in end-to-end delay because the number of links and routers by successive packets can be varied. With path diversity, each router must perform a full routing table lookup to determine the next-hop router along the path. This process is time-consuming and produces difficulties in attaining end-to-end delay within acceptable bounds for voice and video applications.

MPLS addresses the problem in several ways. Label-switched networks fundamentally implement a simpler procedure to determine the exit path for any incoming packet (as previously discussed). In addition, traffic can be fixed to certain paths (constrained routing) via traffic engineering, which allows the service provider to exert more control over traffic congestion. For resiliency, the service provider can create backup paths such that in the event of a link or node failure, the alternative path can be activated to reduce service failure. Therefore, MPLS opens up new possibilities for traffic engineering. The definition of LSPs

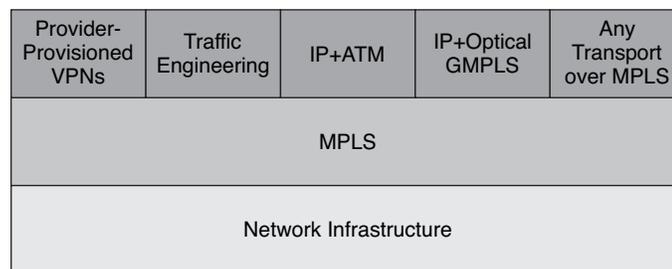
and their FECs allows specific traffic flows to follow paths that deviate from the shortest path designated by classical IP routing protocols.

Implementing the DiffServ architecture with MPLS can provide traffic CoS capabilities over a packet-based network, therein providing the capability to deploy voice and multimedia applications marked with a service priority. Service providers can also deploy MVPNs to support applications using streaming, such as IPTV, Windows Media Player, Real Player, Quick-Time Video Conferencing, and Netmeeting.

Service providers are deploying Layer 2 VPNS to reduce TDM switching and transmission costs as AToM technology emulates Layer 2 services, such as Frame Relay, ATM, PPP, HDLC, and Ethernet. Further, Fast Reroute is used to provide network resilience in place of SDH. Finally, GMPLS can be deployed by organizations with mixed networks and services that require control of multiple technologies, including the optical domain with rapid bandwidth allocation as a key driver to GMPLS implementation. Some current issues with GMPLS include a lack of standards for interdomain routing, integration across nonGMPLS networks, and end-to-end instantiation.

In summary, MPLS technology offers service providers the capability to develop and deploy value-added services and to implement these services in an evolutionary manner. The service architecture is depicted in Figure 3-6.

MPLS as a Foundation for Value-Added Services



**Figure 3-6** *MPLS as a Foundation for Value-Added Services*

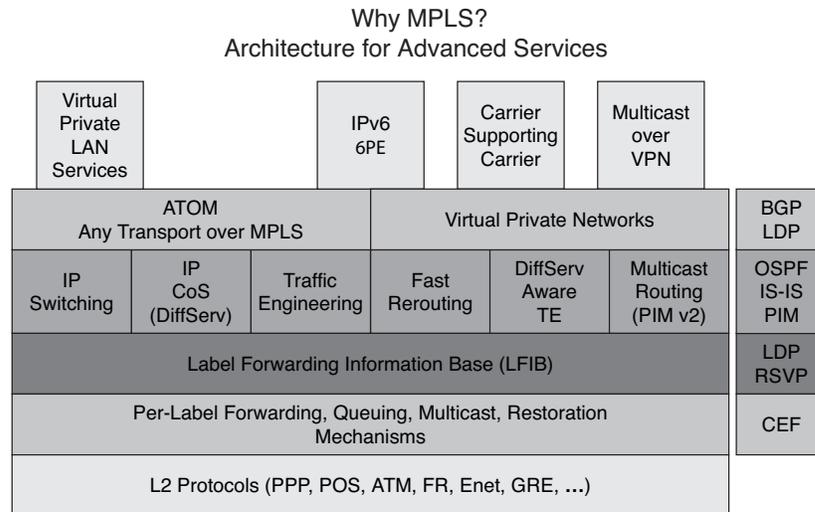
## MPLS Technology Details

This section examines how MPLS facilitates the development of service types, such as Layer 3 VPNs and traffic engineering. Figure 3-7 depicts the MPLS advanced service architectural components that include Layer 3; traffic engineering; differentiated services; Layer 2 VPNs; Virtual Private LAN Service (VPLS); IPv6, multicast GMPLS; and the key control protocols, such as Label Distribution Protocol (LDP), BGP, RSVP, and so on, that activate these service functions. As mentioned previously in this chapter, one of the key benefits of the MPLS architecture is the separation into two planes—one containing information required to transfer a packet (the data plane) and the other allocating the transfer information (the control plane). This separation permits applications to be developed and deployed in a scalable and flexible manner.

Several applications that are facilitated by the implementation of MPLS include:

- **MPLS QoS**—Implements quality of service mechanisms, such as differentiated service, which enables the creation of LSPs with guaranteed bandwidth.
- **Layer 3 VPN**—Uses BGP in the service provider’s network with IP routing protocols or static routing between the service provider and the customer. The BGP protocol is used to exchange the FEC-label binding.
- **Traffic engineering**—Uses extensions of IS-IS or OSPF to distribute attributes in the network. Control processes the FEC-binding through RSVP. Traffic engineering enables you to control traffic routing and thus optimize network utilization.
- **Multicast routing via PIM**—The protocol used to create FEC tables; extensions of version 2 of the PIM protocol are used to exchange FEC-label binding.
- **Layer 2 VPN**—Can be created via a Layer 2 circuit over MPLS, commonly referred to as *Any Transport over MPLS*. Layer 2 VPNs, therefore, use Layer 2 transport as a building block to construct a Layer 2 VPN service that includes auto configuration, management, QoS, and so on.

The sections that follow focus on the technology details for a base of these services, such as Layer 3 VPNs, traffic engineering, differentiated services, and Layer 2 VPNs. Multicast, IPv6, and GMPLS are discussed in later chapters.



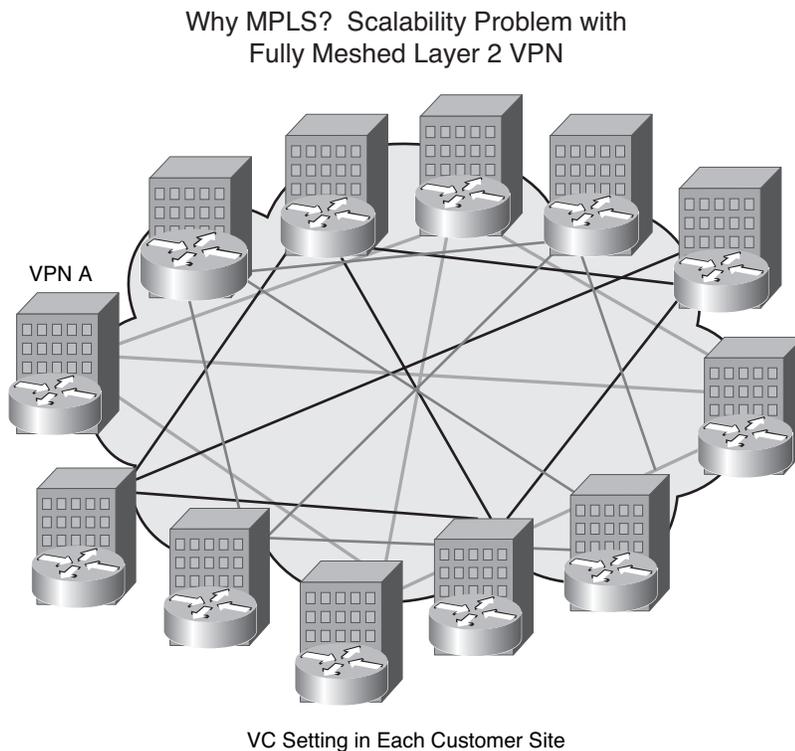
*Figure 3-7 Architecture for Advanced Services*

## Layer 3 VPNs

A virtual private network can be defined as a network shared between organizations, each one with its own individual policy concerning addressing, routing, and security. A VPN thus offers significant savings to organizations because the network investment and operating costs are shared between all users. As long as the service provider ensures that traffic belonging to the various companies is isolated and the preceding policies are respected, the VPN can be considered *transparent*. Because the VPN is managed by an operator external to the company, service provision is subject to a contract in which the operator agrees to respect the terms of an SLA. This specifies, for example, the degree of network availability (number of outages, average time to repair, and so on), the minimum transmission rate between sites, packet loss, jitter, and the maximum latency between sites.

Until the introduction of MPLS architecture, private networks were deployed by using one of two basic techniques: the *overlay* and the *peer-to-peer* models. The overlay model typically uses the virtual circuits of a Frame Relay or ATM service, which means that sites can be interconnected by stacking the IP layer above a Layer 2 connectivity service. The overlay model has advantages such as permitting the duplication of addresses and the isolation of the control and security planes. The overlay model, however, also has drawbacks, including the difficulty in optimizing the size of the virtual circuits between sites, the requirement for meshed circuits that optimize routing, and the obligation to manage Layer 3 adjacencies for all circuits. Conversely, implementing IPSec over the Internet or via general routing encapsulation (GRE) tunnels are examples of Layer 3 overlay over IP for private or public network constructs.

An example of the overlay model is shown in Figure 3-8.



**Figure 3-8** Scalability Problem with Fully Meshed Layer 2 VPN

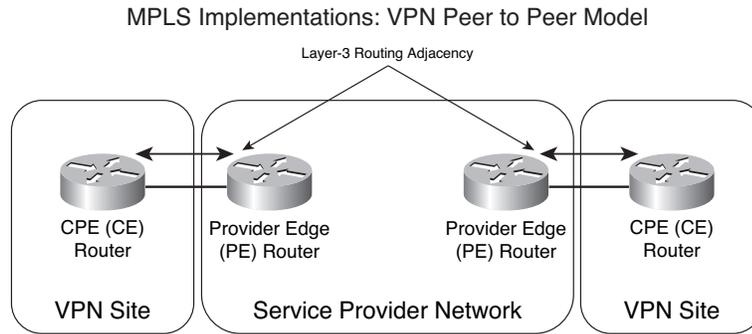
## Peer-to-Peer Model

In the peer-to-peer model, certain limitations of the overlay model are overcome by replacing the use of multiple virtual circuits with a direct exchange of routing information between the service provider and the customer's equipment. The main advantage of this model is the simplification of routing as it appears from the CE installation, thanks to the elimination of multiple virtual circuits. Moreover, the size of the circuits is no longer problematic, and intersite routing is optimal from the moment the service starts up. The main disadvantages of the model are the requirement for the service provider's IGP protocol to manage all the customer VPN routes and the fact that duplication of addressing between clients is impossible.

The VPN service defined over the MPLS architecture allows a group of customers to share common routing information. Thus, a site can belong to one or more VPNs. An MPLS VPN operates at Layer 3 and is also referred to as a *BGP-VPN* because multiprotocol BGP is used to transport the VPN constructs, as is discussed later in this section. The MPLS VPN architecture is based on a VPN router from the customer site (CE) and a provider edge router (PE). The service provider's backbone—specifically the provider (P) routers—have no knowledge of the routing information specific to the various customer VPNs. The PE performs the most important function in the MPLS VPN architecture; the VPN's intelligence is located on the PE, but only for VPNs directly attached to it. The PE manages two or more separate tables for storing routing information.

- **Global table**—Contains all the service provider's internal routes as well as the interface addresses of routers not linked to the VPN (P routers) and the PEs. The global table can contain external IPv4 routes that are useful for providing Internet access for example.
- **VPN routing and forwarding instances (VRF) table**—Includes the customer VPN routes associated with one or more directly connected sites (CE routers). The notion of a VRF is similar to a virtual router.

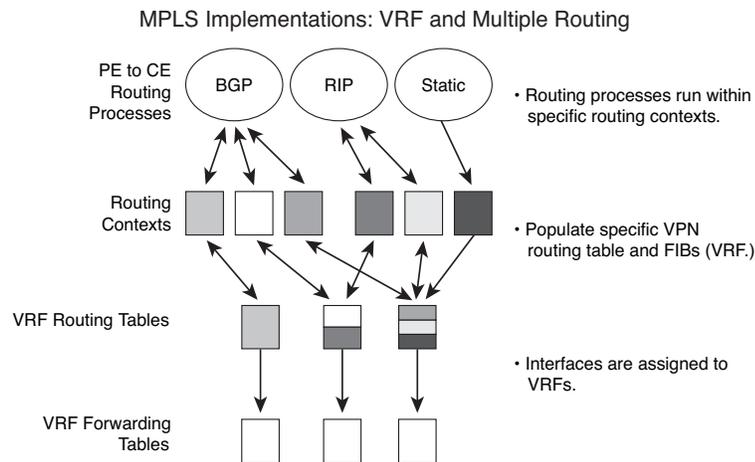
MPLS-VPN is an example of a peer-to-peer model and is depicted in Figure 3-9.



**Figure 3-9** *MPLS Implementations: VPN Peer-to-Peer Model*

### VRF and its Function

A VRF table can be associated with all types of interfaces, logical or physical, and these interfaces can share the same routing information. Whenever a route is defined for a VPN site, the corresponding VRF is informed thanks to the routing context associated with the incoming interface. A routing context can be thought of as the capacity to manage several instances of a given routing protocol, but with a total separation of routing information between the contexts, as summarized in Figure 3-10.

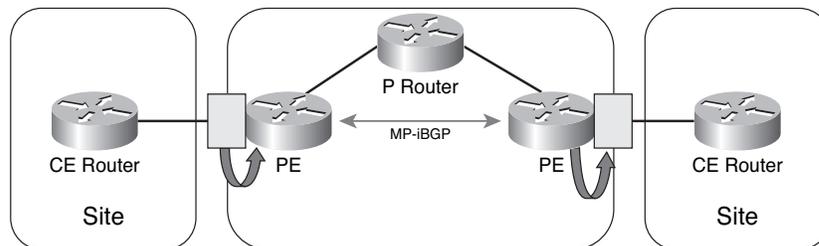


**Figure 3-10** *MPLS Implementations: VRF and Multiple Routing*

To allow the duplication of addressing between VPN clients, a single identifier is required: the *route distinguisher (RD)*. The RD is added to the beginning of an IPv4 route before the route is distributed in BGP and is used for exchanging VPN routes between the PE routers. The combination of the RD and the IPv4 prefix constitutes the VPNv4 prefix. The exchange of routing information for MPLS-VPN or Layer 3 VPN is carried out using the dynamic routing protocols (BGP-4, OSPF, RIPv2, and EIGRP), one the PE-CE links, (or by static routing), and by using multiprotocol BGP between the PE routers. The multiprotocol extension of BGP is necessary because BGP does not carry simple IPv4 prefixes in the MPLS-VPN architecture. In fact, with the creation of the VPNv4 prefix through the addition of the RD to the IPv4 prefix, BGP should be able to transport prefixes that are no longer IPv4. After the route is memorized in the VRF, it is redistributed through the backbone as a VPNv4 prefix via multiprotocol BGP to the other PE routers, as shown in Figure 3-11.

#### MPLS Implementations: VRF Route Distribution

- PE routers distribute local VPN information across the MPLS/VPN backbone through the use of MP-iBGP and redistribution from VRF receiving PE imports routes into attached VRFs.



**Figure 3-11** MPLS Implementations: VRF Route Distribution

Clearly, a mechanism must exist to permit the receiving PE to distribute the information on the input routes to all VPN sites concerned. The PE router must be able to obtain the suitable routes from the appropriate VRF table and then inform the affected VPN sites. To do this, the sending PE router inserts an extended community attribute called a *route target (RT)* in the BGP message. The route target is used by the receiving PE to identify which of its various VRF tables

should receive the route. This function is entirely managed by configuration. Each VRF of the receiving PE must be configured with the acceptable RT values that allow it to import the appropriate routes. After the routes are imported into VRFs, the PE router can transmit them to the affected VPN sites, providing the routing information that ensures connectivity between the VPN sites. Figure 3-12 depicts the RT operation.

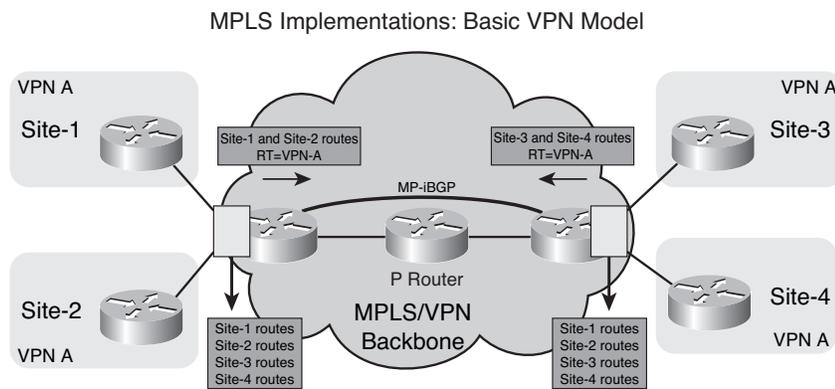


Figure 3-12 MPLS Implementations: Basic VPN Model

### MPLS Label Stack Role

To permit label switching in the MPLS backbone, an MPLS label used by the receiving PE as an index in the transfer table is associated with each VPN route communicated via multiprotocol BGP. An additional label is then used to switch the packet to the source PE. To make the MPLS network transparent for the transmitted and received data on a VPN link, a stack or hierarchy of labels allows the transfer of information between the two PEs, while a second level (announced via multiprotocol BGP) informs the exit or egress PE on which VPN interface to send the packet. In this way, a two-level stack of labels is used for end-to-end transfer, as shown in Figure 3-13.

Putting It All Together—Forwarding Plane

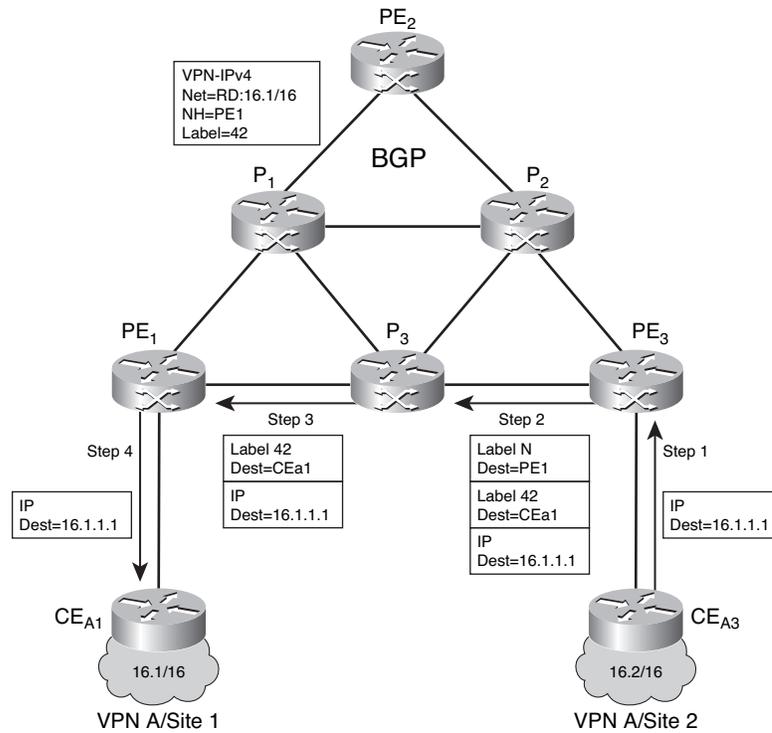


Figure 3-13 Putting it All Together: Forwarding Plane

### Topologies

We have discussed building a basic VPN or an intranet. Via the manipulation of RTs, so-called *extranets* can also be deployed; additionally, hub and spoke topologies can be supported. The next section discusses advanced services, such as Carriers Carrier (CsC) and Inter-Provider Autonomous System (Inter-AS). CsC and Inter-AS are described in detail in the book *MPLS and VPN Architectures*, Vol 2, J. Guichard, et al., Cisco Press.

Finally, you could use a subset of the MPLS-VPN architecture. For example, you could use virtual routing forwarding instances to support multiple (overlapping and independent) routing tables (and forwarding tables) per customer, which is referred to as *Multi-Lite VRF*. The CE supports traffic separation between customer networks. In addition, no MPLS functionality exists on the CE and no label exchange exists between the CE and PE.

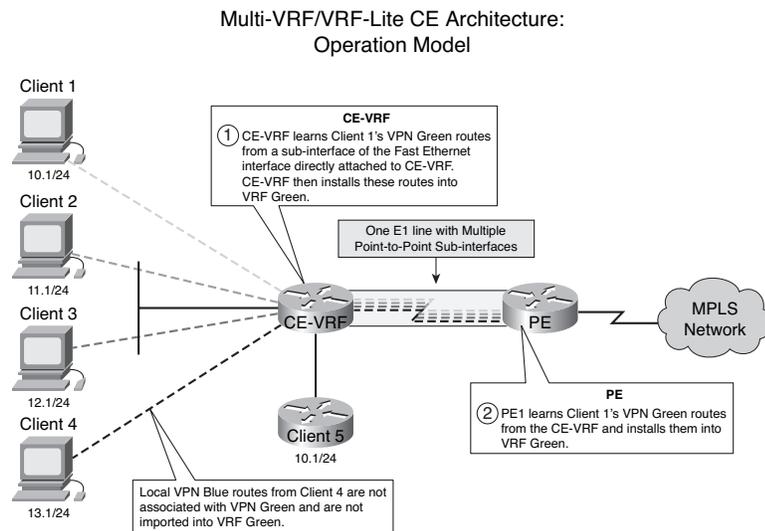
A customer could implement Multi-Lite CE in an enhanced branch office capability where CE routers use VRF interfaces. VLAN-like configuration on the customer side and the CE router can only configure VRF interfaces and support VRF routing tables. An alternative to Multi-Lite VRF is to use separate CE routers per each client's organization.

---

**NOTE** When deploying Multi-Lite VRF in a multiclass configuration that has different class treatments per VRF, certain complexities are introduced that require careful rule sets to preserve traffic characteristics for each class or QoS set.

---

Figure 3-14 shows an example of Multi-VRF deployment.



**Figure 3-14** Multi-VRF/VRF-Lite CE Architecture: Operational Model

We have discussed the Layer 3–centric examples of MPLS used to build VPNs or BGP-based VPNs. We have also highlighted VRF’s attributes and the types of topologies that can be supported.

In the next section, we explore advanced MPLS VPN implementations, such as Carrier Supporting Carrier and Inter-AS constructs, specifically used across multiple operator domains.

## Carrier Supporting Carrier and Inter-Provider Autonomous Systems

MPLS-VPN architecture can be further extended to implement advanced services, such as CsC or Inter-AS. For example, VPN sites might be geographically dispersed, requiring connectivity to different MPLS VPN service providers. That is, the transit between VPN sites might pass through multiple providers’ MPLS backbones implying an exchange of VPN routing information between providers and the provider backbones might or might not provide VPN service directly. Figures 3-15 and 3-16 summarize the Inter-AS service problem and available options.

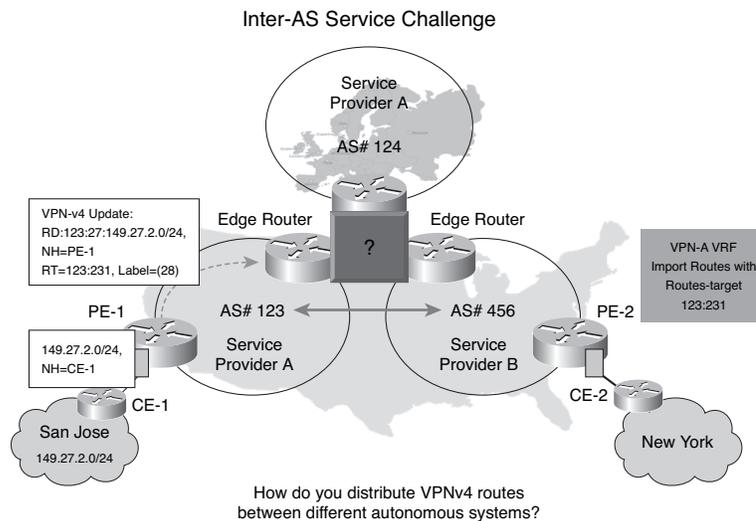
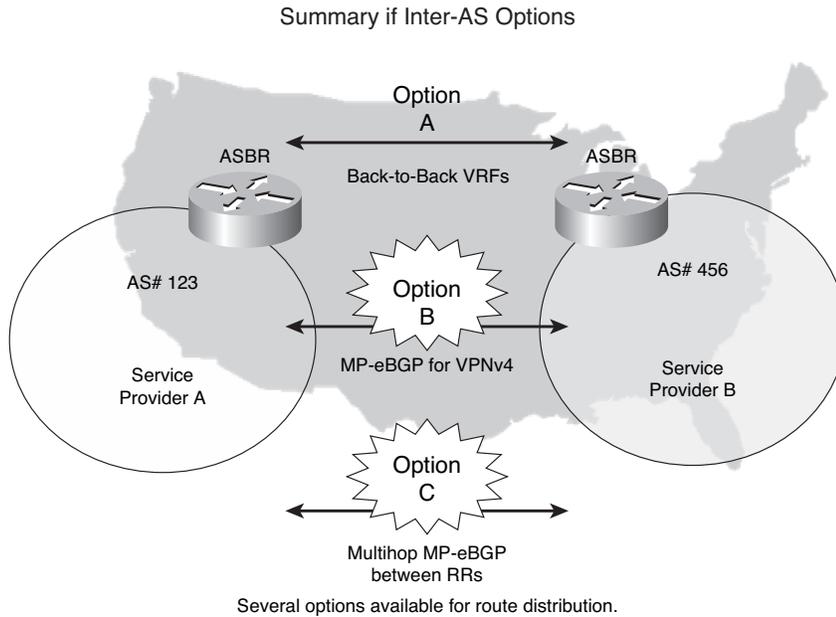


Figure 3-15 Inter-AS Service Challenge



**Figure 3-16** Summary of Inter-AS Options

In the next section, we examine traffic engineering implementations as a service building block for MPLS-based networks.

## Traffic Engineering

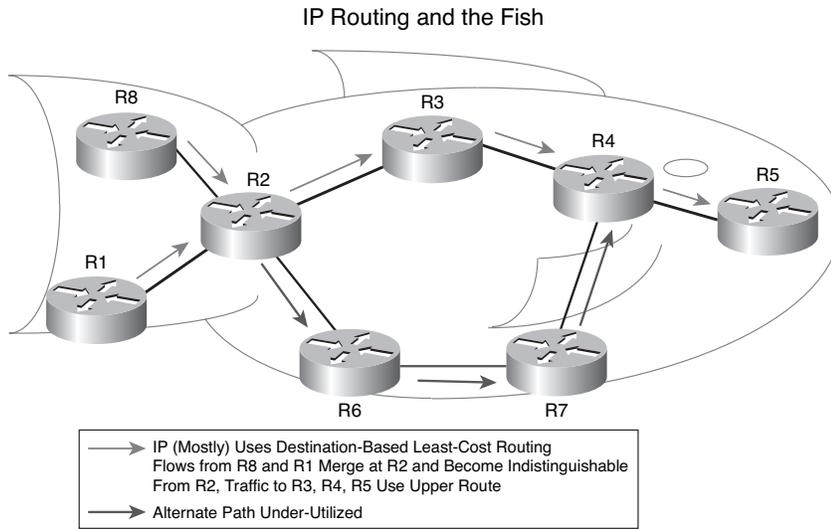
Traffic engineering is the process of routing data traffic to balance the traffic load on the various links, routers, and switches in the network and is most applicable in networks where multiple parallel or alternate paths are available. Fundamentally, traffic engineering involves provisioning the network to ensure sufficient capacity exists to handle the forecast demand from the different service classes while meeting their respective QoS objectives. Current routing on IP

networks is based on computing the shortest path where the “length” of a link is determined by an administratively assigned weight. Reasons to deploy traffic engineering include the following:

- Congestion in the network due to changing traffic patterns
- Election news, online trading, or major sports events
- Better utilization of available bandwidth
- Route on the path that is not the shortest
- Route around failed links/nodes; fast rerouting around failures, transparently to users like SONET Automatic Protection Switching (APS)
- Building of new services—virtual leased-line services
- VoIP Toll-Bypass applications, point-to-point bandwidth guarantees
- Capacity planning traffic engineering improves aggregate availability of the network

Additional reasons to consider traffic engineering are that IP networks route based only on destination (route) and ATM/FR networks switch based on both source and destination (PVC and so on). Some large IP networks were built on ATM or FR to take advantage of source and destination routing, and overlay networks inherently hinder scaling (see “The Fish Problem” in Figure 3-17). MPLS-TE allows you to do source and destination routing while removing the major scaling limitation of overlay networks. Finally, MPLS-TE has since evolved to do things other than bandwidth optimization, which is discussed in detail in Chapter 8, “Traffic Engineering.”

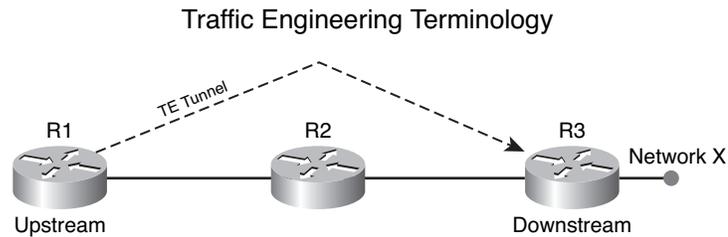
The challenge with destination leased cost routing is that alternate links are often underutilized, as shown in Figure 3-17.



**Figure 3-17** IP Routing and the Fish

To demonstrate how traffic engineering addresses the problem of underutilized links, we will take an example in Figure 3-18 by first defining the traffic engineer terminology:

- **Head-End**—A router on which a TE tunnel is configured (R1)
- **Tail-End**—The router on which the TE tunnel terminates (R3)
- **Mid-point**—A router through which the TE tunnel passes (R2)
- **LSP**—The label-switched path taken by the TE tunnel; here it's R1-R2-R3
- **Downstream router**—A router closer to the tunnel tail
- **Upstream router**—A router farther from the tunnel tail (so R2 is upstream to R3's downstream, and R1 is upstream from R2's downstream)



**Figure 3-18** Traffic Engineering Terminology

Continuing the traffic engineering building block, information distribution is done via a link state protocol, such as IS-IS or OSPF. The link state protocol is required only for traffic engineering, not for the implementation of Layer 3 VPNs. A link state protocol is required to ensure that information gets flooded and to build a topology of the entire network.

Information that is flooded includes link, bandwidth, and attributes. After available bandwidth information is flooded, a router can calculate a path from head to tail. The TE head-end performs a constrained SPF (CSPF) calculation to find the best path. CSPF is just like regular IGP SPF, except that it takes required bandwidth into account and looks for the best path from a head to a single tail, not to all devices.

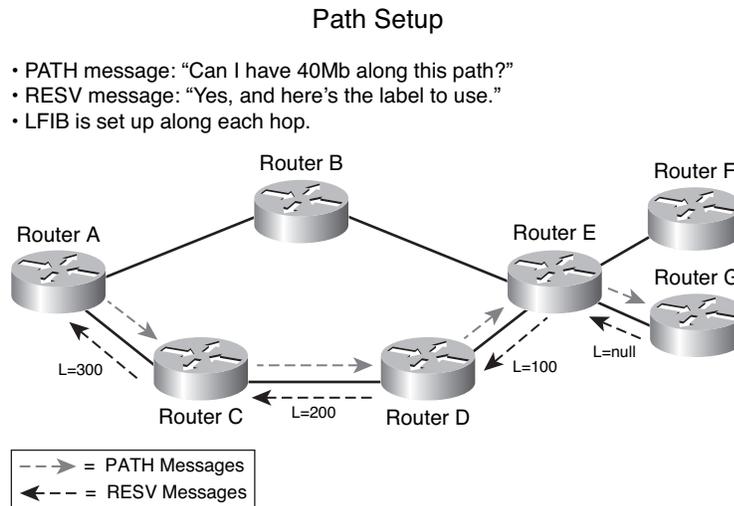
Note that control capabilities offered by existing Internet Gateway Protocols (IGPs) are adequate for traffic engineering. This makes actualizing effective policies to address network performance problems difficult. IGPs that are based on shortest path algorithms contribute to congestion problems in autonomous systems within the Internet. SPF algorithms generally optimize based on a simple additive metric. These protocols are topology driven so bandwidth availability and traffic characteristics are not factors in routing decisions. (Refer to IETF RFC 2702, “Requirements for Traffic Engineering over MPLS.”)

In practice, there has been zero impact from CSPF CPU utilization on even the largest networks. After the path is calculated, you need to signal it across the network.

To reserve any bandwidth so that other LSPs cannot overload the path and to establish an LSP for loop-free forwarding along an arbitrary path, a path setup is done via PATH messages from head to tail and is similar to “call setup.” A PATH MESSAGE carries a LABEL\_REQUEST, whereas RESV messages are done from tail to head and are analogous to “call ACK.” RESV messages transport the LABEL.

Other RSVP message types exist for LSP teardown and error signaling. The principles behind path setup are that you can use MPLS-TE to forward traffic down a path other than that determined by your IGP cost and that you can determine these arbitrary paths per tunnel head-end.

Figure 3-19 describes the path setup operation.



**Figure 3-19** Path Setup

After having established the TE tunnel, the next step in deploying MPLS-TE is to direct traffic down the TE tunnel. Directing traffic down a TE tunnel can be done by one of the following four methods:

- **Autoroute**—The TE tunnel is treated as a directly connected link to the tail IGP adjacency and is *not* run over the tunnel. Unlike an ATM/FR VC, autoroute is limited to single area/level only.

- **Forwarding adjacency**—With autoroute, the LSP is not advertised into the IGP, and this is the correct behavior if you are adding TE to an IP network. However, it might not be appropriate if you are migrating from ATM/FR to TE. Sometimes advertising the LSP into the IGP as a link is necessary to preserve the routing outside the ATM/FR cloud.
- **Static routes**
- **Policy routing**

With autoroute and static route, MPLS-TE provides for unequal cost load balancing. Static routes inherit unequal cost load sharing when recursing through a tunnel. IP routing has equal-cost load balancing but not unequal cost. Unequal cost load balancing is difficult to implement while guaranteeing a loop-free topology. Therefore, because MPLS does not forward based on IP header, permanent routing loops do not occur. Further, 16 hash buckets are available for the next hop, and these are shared in rough proportion to the configured tunnel bandwidth or load-share value. Autoroute, forward adjacency, and static and policy routing are further explained in Chapter 8. To summarize, MPLS-TE operational components include the following:

- Resource/policy information distribution
- Constraint-based path computation
- RSVP for tunnel signaling
- Link admission control
- LSP establishment
- TE tunnel control and maintenance
- Assignment of traffic to tunnels

MPLS-TE can be used to direct traffic down a path other than that determined by your IGP cost. Fast Reroute (FRR) builds a path to be used in case of a failure in the network and minimizes packet loss by avoiding transient routing loops. To deploy FRR, you must pre-establish backup paths such that when a failure occurs, the protected traffic is switched onto backup paths after local repair and the tunnel head-ends are signaled to recover. Several FRR modes, such as link node and path protection, exist. In link protection, the backup tunnel tail-head is one hop away from the point of local repair (PLR). In node protection, the backup tunnel tail-end

is two hops away from the PLR. Figures 3-20 and 3-21 depict link, node, and path protection mechanisms.

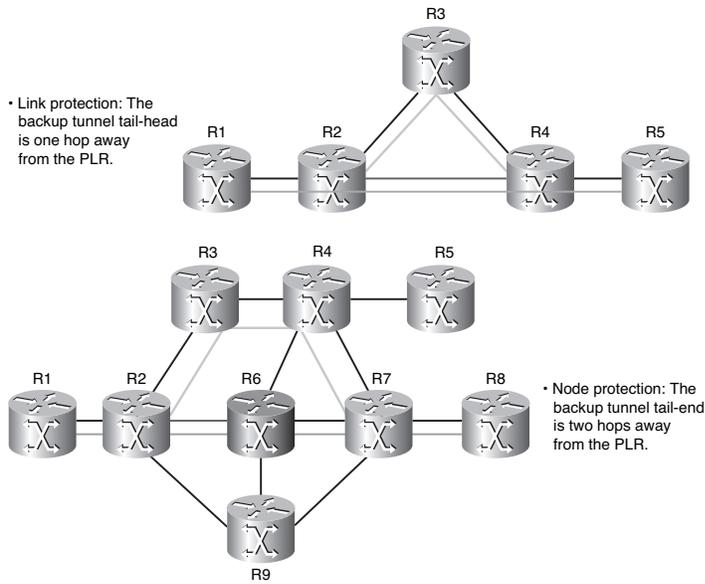


Figure 3-20 FRR Link and Node Protection

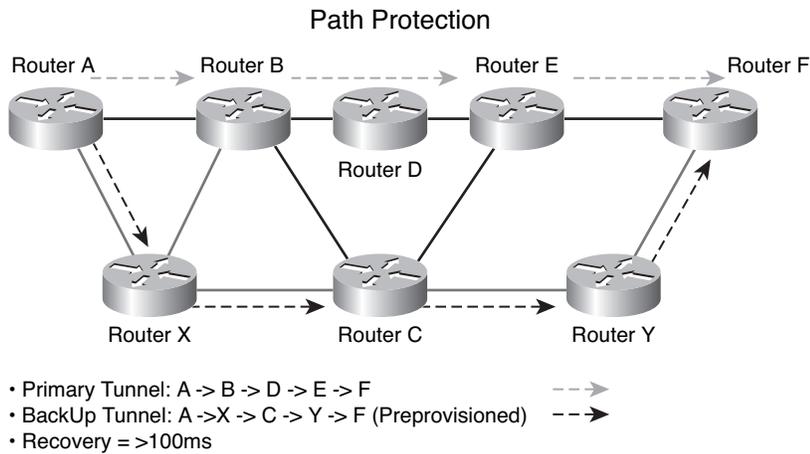


Figure 3-21 Path Protection

One application for MPLS-TE is to implement a virtual leased line (VLL) with bandwidth guarantees. This can be done via MPLS-TE or differentiated service-traffic engineering (DiffServ-TE) with QoS. Diff-Serv is covered in the next section of this chapter. Figure 3-22 shows an example of VLL deployment via MPLS-TE.

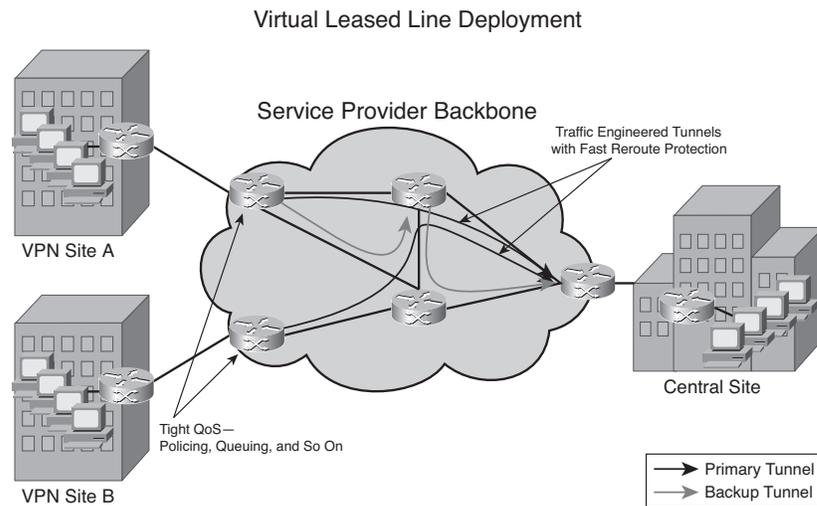


Figure 3-22 Virtual Leased Line Deployment

The next section discusses class of service implementation based on the differentiated service architecture or DiffServ. Details of DiffServ are described in Chapter 9. The next section highlights the architecture and provides linkage to service development.

## DiffServ

DiffServ architecture relies on a broad differentiation between a small number of service classes. Packets are identified as belonging to one class or another through the content of the DiffServ field in the IP header. Packets are classified and marked at the network edge depending on the kind of traffic contract

or SLA between the customer and the service provider. The different classes of packets then receive different per-hop behaviors (PHBs) in the nodes of the network core. Service differentiation also implies different tariffs depending on the QoS offering to flows and packets belonging to different classes. The DiffServ architecture consists of a set of functional elements embodied in the network nodes.

- Allocation of buffering and bandwidth aggregates corresponding to each PHB
- Packet classification (FEC)
- Traffic conditioning, metering, marking, and shaping

The sophisticated operations of packet classification are implemented at the edge of the network or in the customer equipment. The architecture avoids the requirement to maintain per-flow or per-user state within the network core.

The implementation, configuration, operation, and administration of the PHB groups supported of a DiffServ domain are dependent on sufficient resources being available. You must ensure that the amount of resources available is adequate, given the traffic conditioning parameters for contracted SLAs.

The DiffServ field (IETF RFC 2474) replaces the existing definitions of the TOS byte in IPv4 and the traffic class byte in IPv6. Six bits of the DiffServ field are used in the form of the differentiated services code point (DSCP) identifies the PHB to be received by a packet at each node.

In addition to traditional best effort, considered to be the default PHB, two other PHBs have been defined by the IETF. These are expedited forwarding (EF) (IETF RFC 2598) and assured forwarding (AF) (IETF RFC 2597). The EF PHB is designed to provide an end-to-end service with low packet loss, delay, and jitter and a guaranteed bit rate. It can be used to create a virtual leased line as described previously under MPLS-TE. The AF PHB group permits a service provider to offer differentiated levels of performance to different traffic aggregates received from customers. For example, AF packets can be divided into four subclasses with a separate resource allocation for each class.

Using DiffServ in MPLS (IETF RFC 3270) the following two types of LSPs exist:

- **EXP-Inferred-PSC LSP (E-LSP)**—Can transport different service classes and the differentiated handling of packets being carried out at the level of the LSR on the basis of the EXP field where up to eight PHBs per LSPs can be deployed.
- **Label-Only-Inferred-PSC LSP (L-LSP)**—Handles only one type of DiffServ aggregate, the label defining the LSP corresponding to a DiffServ class. The information on the DiffServ class is provided when the LSP is set up using LDP or RSVP protocol. An LSR can merge L-LSPs only if they belong to the same DiffServ class. The EXP field can be used for the discard priorities of DiffServ classes.

The advantages of E-LSP are that multiple PHBs (8) per LSP are supported, thus reducing the number of labels required; implementation is a bit easier in that you just need to configure LSRs to map EXP values to PHBs. The disadvantage is that, although 8 PHBs can be supported, DiffServ actually supports up to 64 PHBs and cannot be implemented when the shim header is not used—for example, in ATM or FR.

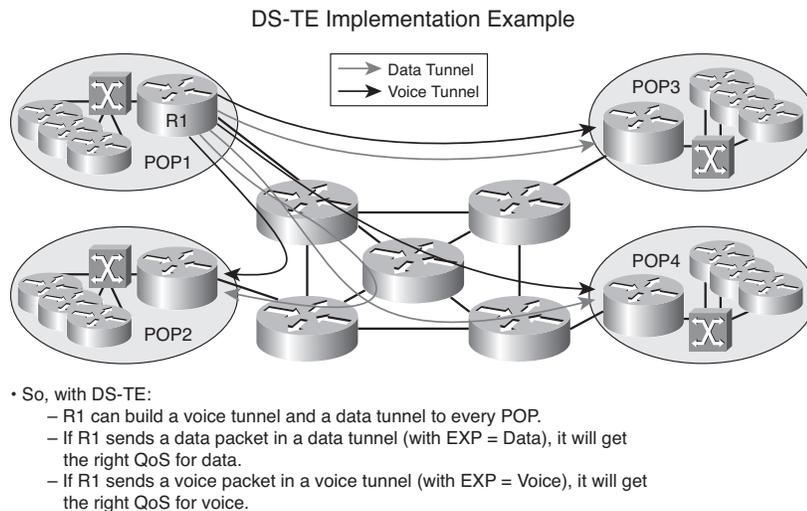
The advantages of L-LSP are that it can support an arbitrarily large number of PHBs in excess of 64 and can use multiple paths for different PHBs via traffic engineering. The disadvantages of L-LSP are that it consumes more labels and is more difficult to configure. For example, you need to configure LDP to signal PHBs during label establishment.

The majority of current deployments consist of less than 8 PHBs in the core. Actual deployment models on Router-LSRs consist of a single E-LSP for all CoS or 1 E-LSP per CoS (for example, 1 E-LSP for voice + 1 E-LSP for data). A strict priority queue exists for EF and a high weight exists for premium/AF (for example, 90%) with WRED for in and out-of contract. Additionally, a low weight exists for best effort (for example, 10%) with RED. The L-LSPs used today on ATM-LSRs might be required in the future on Router-LSRs, if and when more than 8 PHBs are needed.

MPLS DiffServ and MPLS TE can coexist because MPLS TE is designed as a tool to improve backbone efficiency independently of QoS. MPLS TE computes routes for aggregates across all classes and performs admission control over the “global” bandwidth pool for all classes (that is, MPLS TE does not consider the bandwidth allocated to each queue). MPLS TE and MPLS DiffServ can run simultaneously and provide their own benefits. For example, TE distributes aggregate load and DiffServ provides differentiation. This construct is referred to as differentiated services-traffic engineering (DS-TE). DS-TE makes MPLS TE aware of DiffServ in that DS-TE establishes separate tunnels for different classes and takes into account the “bandwidth” available to each class (for example, to queue).

DS-TE also considers separate engineering constraints for each class. Here are two examples: I want to limit voice traffic to 70% of link max., but I do not mind having up to 100% of best effort traffic, or I want an overbook ratio of 1 for voice but 3 for best effort.

DS-TE can take into account different metrics (for example, delay) and ensures that a specific QoS level of each DiffServ class is achieved. DS-TE provides a mechanism where different tunnels can satisfy various engineering constraints as summarized in Figure 3-23.



**Figure 3-23** DS-TE Implementation Example

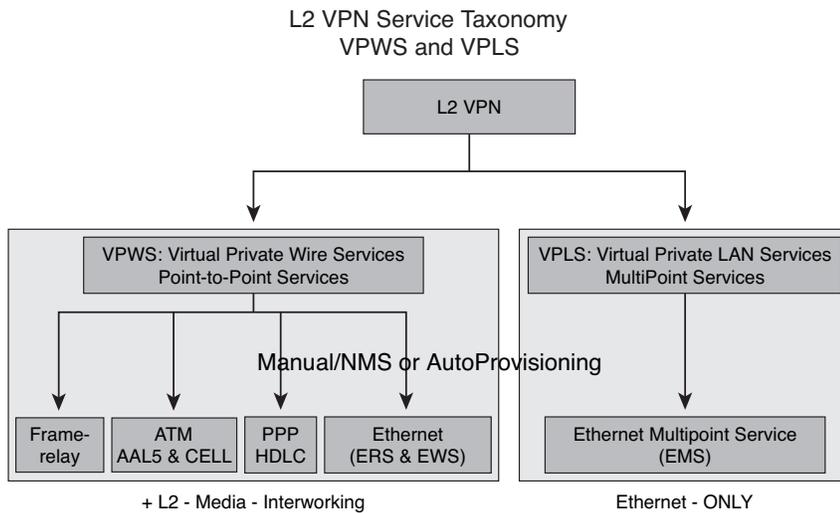
MPLS DiffServ and DS-TE are discussed further in Chapter 9.

## Layer 2 VPNs

There is a broad taxonomy for Layer 2 transport consisting of the following components:

- **L2 Transport**—Provides point-to-point Layer 2 connectivity.
- **L2VPNs**—Use Layer 2 transport as a building block to build a Layer 2 VPN service that includes autoconfiguration, management, QoS, and so on. A concept of pseudowires to emulate a Layer 2 service is a key attribute for a Layer 2 VPN over MPLS.
- **Virtual private wire service (VPWS)**—Has a characteristic of a fixed relationship between an attachment-virtual circuit and an emulated virtual circuit. VPWS-based services are point-to-point (for example, Frame-Relay/ATM/Ethernet services over IP/MPLS).
- **Virtual private LAN switching service (VPLS)**—It's fundamentally an end-to-end service. It is "virtual" because multiple instances of this service share the same physical infrastructure; it is "private" because each instance of the service is independent and isolated from the others. It is "LAN service" because it tries to provide a multipoint connectivity among the participant endpoints that looks like a LAN. A dynamic relationship (learned) exists between an attachment-virtual circuit and emulated virtual circuits that is determined by customer MAC address. An example of a VPLS-based service is an Ethernet multipoint service.
- **IP LAN services (IPLS)**—A service similar to VPLS, in that all LAN interfaces are implemented in promiscuous mode and frames are forwarded based on their MAC destination addresses. However, the maintenance of the MAC forwarding tables is done via signaling rather than via the MAC address learning procedures of IEEE 802.1D. Further, the Address Resolution Protocol (ARP) messages are proxied, rather than carried transparently. You could use routers and a single MAC address rather than the more complex bridging of customer LANs. IPLS is currently an IETF draft.

Figure 3-24 summarizes the Layer 2 taxonomy.



**Figure 3-24** L2 VPN Service Taxonomy: VPWS and VPLS

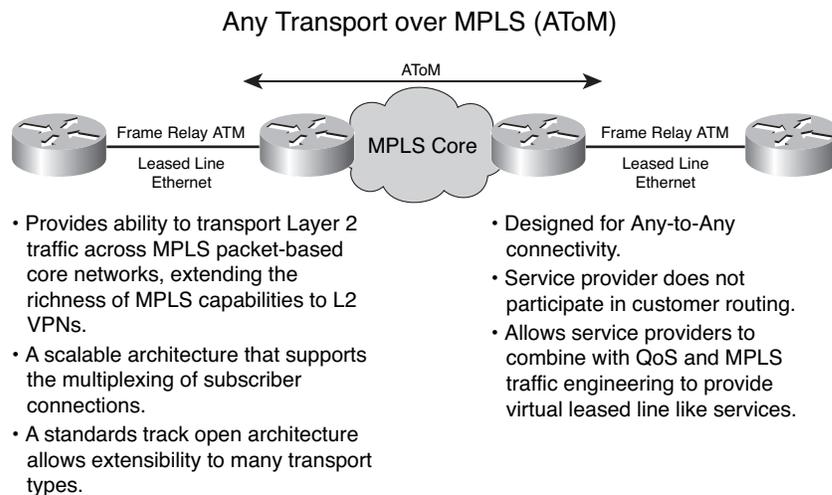
A Layer 2 transport over MPLS is referred to as *Any Transport over MPLS (AToM)* by Cisco. Any transport over MPLS is required to support several services, such as Layer 2 transport over packet-based infrastructure. ATM and Frame Relay service is popular, and the provisioning of these services is easily understood. Currently, VPNs are built using either IPsec tunnels or PVCs with ATM or Frame Relay.

Layer 3 VPNs are available to offer branch office connectivity; however, this connectivity is limited to IP-based services and other protocols must be encapsulated in IP. First, encapsulating everything in IP might not always be possible. Second, service requirements, such as ATM cell transport, IGP trunking, and non-IP transport, are also needed. This trunking of Layer 2 frames can be done with Any Transport over MPLS in MPLS IP networks.

While deploying IP core, the trunking of Layer 2 can be considered because there might be existing revenue generating services for service providers or service providers might want to offer services similar to point-to-point virtual leased lines to their customers. Service providers are used to build Layer 2

services. These services are attractive in terms of revenue because the provider is not required to participate in any customer routing information. Because MPLS can transport both Layer 2 and Layer 3, it offers a viable alternative and convergence point for diverse infrastructures. Moreover, specific services, such as transparent LAN connectivity, extension of broadcast domain, virtual leased line, and voice trunking, can be easily built when AToM functionality is combined with MPLS QoS and traffic engineering.

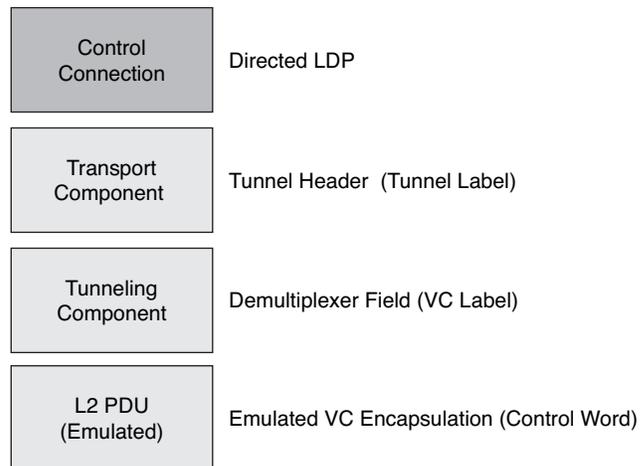
An AToM overview is depicted in Figure 3-25.



**Figure 3-25** Any Transport over MPLS (AToM)

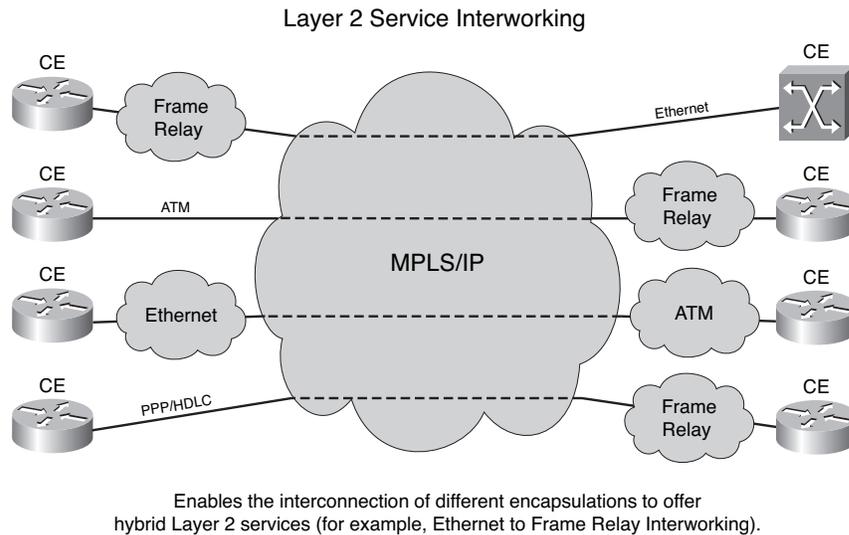
Layer 2 transport options include Frame Relay, ATM AAL5 and ATM Cell Relay, Ethernet, 802.1q (VLAN), Packet Over Sonet (POS), TDM, and Cisco HDLC and PPP. The architectural elements of AToM consist of the following: a control connection performed by directed LDP, a transport component called a *tunnel header* or *tunnel label*, a tunnelling component that consists of a demultiplexer field or virtual circuit label, and a Layer 2 protocol data unit (PDU) that provides an emulated virtual circuit encapsulation via the control word attribute. Figure 3-26 summarizes the AToM architectural elements.

### Layer-2 Transport across MPLS



**Figure 3-26** Layer 2 Transport Across MPLS

- AToM is used for the point-to-point transport of Layer 2 PDUs across an MPLS-enabled core via directed LDP sessions to negotiate virtual circuit labels between participating peers. AToM can use a control word to preserve relevant information in transported PDUs (for instance, Frame Relay BECN, FECN, or DE). AToM can also interwork with native service management protocols, such as ILMI/LMI, to indicate the local circuit status to remote peers. Layer 2 service interworking enables the interconnection of different encapsulations to offer hybrid Layer 2 services (for example, Ethernet to Frame Relay Interworking) over an IP/MPLS core and can facilitate the convergence of existing services. The Layer 2 service interworking construct is shown in Figure 3-27.



**Figure 3-27** *Layer 2 Service Interworking*

Layer 2 and Layer 3 VPNs are summarized as follows:

Layer 3 VPNs are concerned primarily with looking at the Layer 3 information and making forwarding decisions. MPLS VPNs require a CE-to-provider edge routing process plus PE-to-PE signaling via multiprotocol BGP. With Layer 2 VPNs, the information used to make forwarding decisions is based on Layer 2 information—for instance, via a MAC address, via a VLAN ID, on DLCI information, or on an input interface for lease line connectivity. Figure 3-28 provides a comparison between Layer 3 and Layer 2 VPNs. Figure 3-29 depicts the L2VPN constructs as has been discussed.

Layer 3 and Layer 2 VPN Comparison

Layer 3 VPNs	Layer 2 VPNs
<ul style="list-style-type: none"> <li>• Provider devices forward customer packets based on Layer 3 information (for example, IP).</li> </ul>	<ul style="list-style-type: none"> <li>• Provider devices forward customer packets based on Layer 2 information.</li> </ul>
<ul style="list-style-type: none"> <li>• <b>SP Involvement in Routing</b></li> </ul>	<ul style="list-style-type: none"> <li>• <b>Tunnels, Circuits, LSPs, MAC Address</b></li> </ul>
<ul style="list-style-type: none"> <li>• MPLS/BGP VPNs (RFC 2547), GRE, Virtual Router Approaches</li> </ul>	<ul style="list-style-type: none"> <li>• Pseudo-Wire Concept</li> </ul>

Figure 3-28 Layer 3 and Layer 2 VPN Comparison

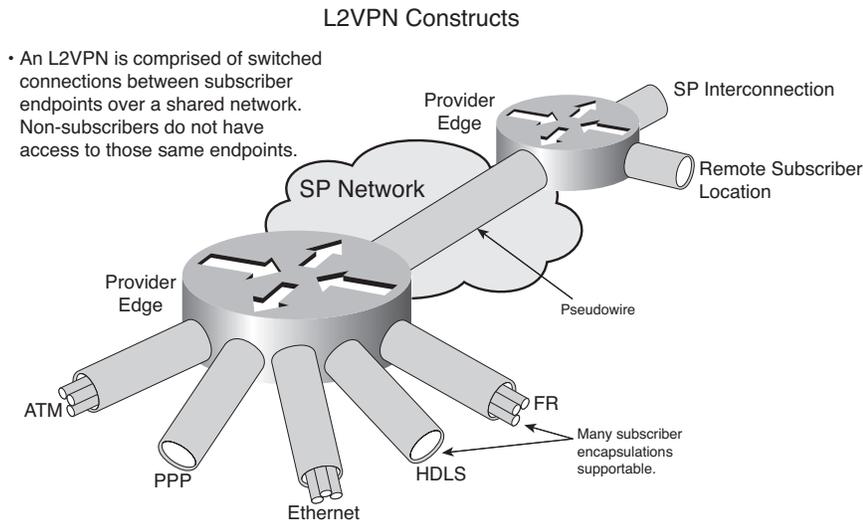


Figure 3-29 L2VPN Constructs

Layer 2 VPNS are further discussed in Chapter 4.

## Summary

MPLS technology is fundamentally a service enabler for Layer 3 VPNs and provides for support of CoS and QoS guarantees along with traffic engineering, DiffServ-TE, and fast reroute that are required to manage tight SLAs for such services as voice, video, and data. Multicast VPNs can support enhanced services for push applications, such as streaming, IPTV, videoconferencing, and e-learning. As IP commences to dominate the majority of public network traffic and the requirement for bandwidth increases the need to reduce management, service providers definitely require configuration and provisioning in the network.

With GMPLS, the IP routing tables of an optical LSR enable you to activate a lambda of dense wave divisional multiplexing (DWDM) immediately, according to the needs of the network. It therefore becomes possible to establish connections in a dynamic fashion for rapid provisioning through the SDH/Sonet, optical, or packet network layers.

GMPLS could be used to support bandwidth-intensive applications, such as GRiD, which is emerging in the industry amongst auto, financial, and pharmaceutical vertical segments. Thus, MPLS technology is flexible and can be used to develop and implement services serving current and future market needs.

We have technically deconstructed Layer 3 VPNs as these apply to MPLS and have described the functionality of traffic engineering and differentiated services essential for the deployment of services that require tight SLAs, such as voice and videoconferencing.

We have further provided an overview of Layer 2 VPNs that can be implemented over an MPLS core. Layer 2 service interworking as part of an evolving convergence strategy can facilitate the migration of legacy Layer 2-based services over an MPLS infrastructure. Finally, we have provided a building block of the MPLS service architecture essential to highlighting its value-added proposition, particularly toward developing service provider NGNs for implementing LAN/WAN virtualization within enterprise organizations.

