CHAPTER 5

MPLS in the DCN

First Published: January 3, 2008
Last Updated: January 3, 2008

Finding Support Information for Platforms and Cisco IOS and Catalyst OS Software Images

Use Cisco Feature Navigator to find information about platform support and Cisco IOS software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn. An account on Cisco.com is not required.

Introduction

This section describes an architecture for converged networks that provides a framework data communications network (DCN) for a growing range of new technologies being deployed, such as Metro Ethernet, L2VPN, and L3VPN on a single foundation. A DCN is the out-of-band operations support network (OSN) that service providers use for connectivity between their Operations Support System (OSS) applications and network elements (transport, switching, routing, and so on). The OSS applications perform network surveillance, provisioning, service restoral, collection of billing data, and other applications.

This DCN architecture uses the Multiprotocol Label Switching (MPLS) technology to provide many distinct advantages to the service provider in deploying a more robust, foolproof, and secure DCN over the traditional IP packet forwarding as shown in Figure 5-1.

Figure 5-1 MPLS/IP in the DCN

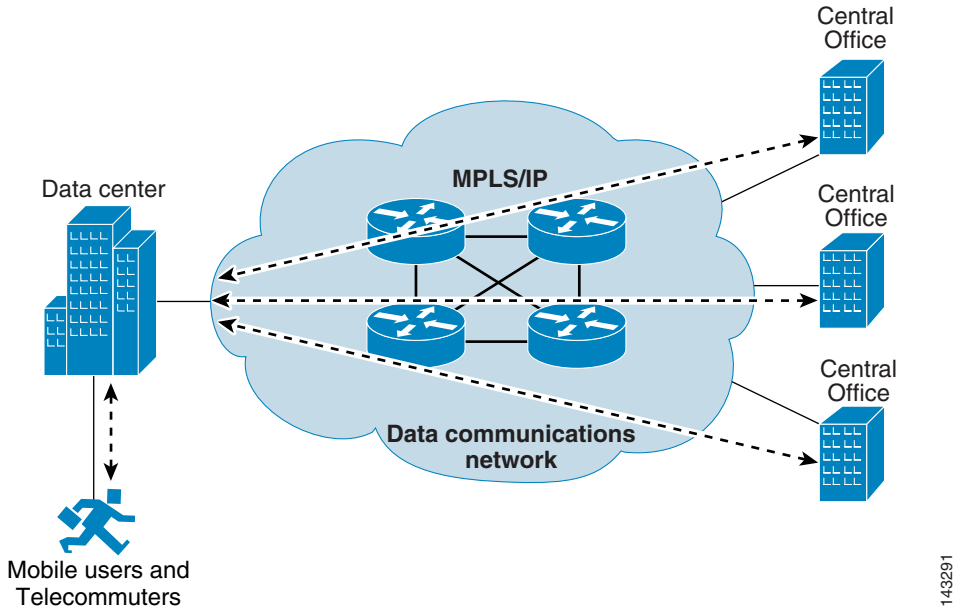

The DCN architecture described in this section uses MPLS technology. MPLS provides many advantages to the service provider over the traditional IP packet forwarding technologies, such as deploying a DCN that is more robust and secure. The DCN architecture and related software features are described in the following sections:

- MPLS in the DCN: Overview, page 5-2
- Deploying MPLS VPNs on a DCN, page 5-16
- Configuration Examples for MPLS VPNs on a DCN, page 5-18

MPLS in the DCN: Overview

MPLS is a high-performance, enhanced packet forwarding technology that improves the performance and traffic management capabilities of Layer 2 (data link layer) and Layer 3 (network layer) of the Open System Interconnection (OSI) model. MPLS provides improved switching with more flexibility and controlled routing. Many service providers use MPLS in the main network to provide assured bandwidth and advanced Service Level Agreements (SLAs) services. Deploying MPLS cuts network costs and provides more ways for the service provider to improve network efficiency.

An MPLS virtual private network (VPN) can provide network services that enable connectivity among multiple sites on a shared infrastructure with the same access or security mechanisms that a separate private network would offer. The network is made virtually private for traffic separation using MPLS VPNs.

Service providers can use MPLS to provide VPN services from the central office (CO) to the remote OSS in the data center. This functionality has been traditionally done using IPsec tunnels often requiring a large number of VPN concentrators or firewall products.

Figure 5-2 shows an example of MPLS VPN architecture over a DCN infrastructure.

Figure 5-2 MPLS VPN Architecture over a DCN Infrastructure

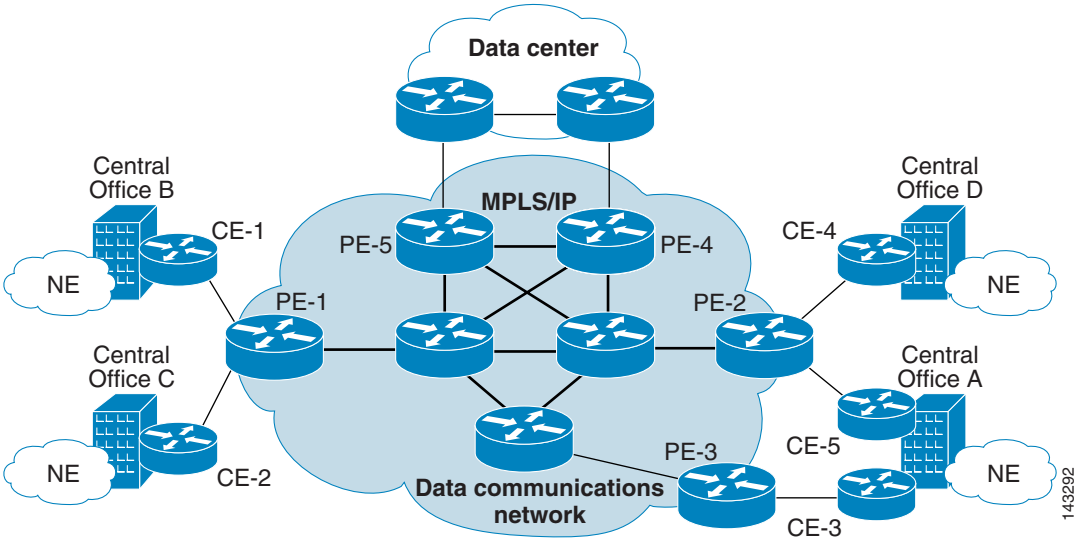

This section provides the following information about MPLS in the DCN:

- Scenarios for Service Providers Deploying MPLS VPNS in the DCN, page 5-3
- Benefits of MPLS VPNs on a DCN, page 5-12
- Supported Platforms, page 5-13
- Design Details, page 5-14

Scenarios for Service Providers Deploying MPLS VPNS in the DCN

A small number of service providers have started using the MPLS VPN technology within the DCN during the last few years. Traditionally, service providers have kept their internal enterprise networks and their DCNs separate. In the early days of the DCN, service providers placed the OSS directly on the DCN, as shown Figure 5-3. The OSSs were connected to the DCN and also often connected to the enterprise network. In addition, network operations center (NOC) computers and NOC technicians were directly connected to the DCN.

Figure 5-3 Classic DCN and Enterprise Deployments

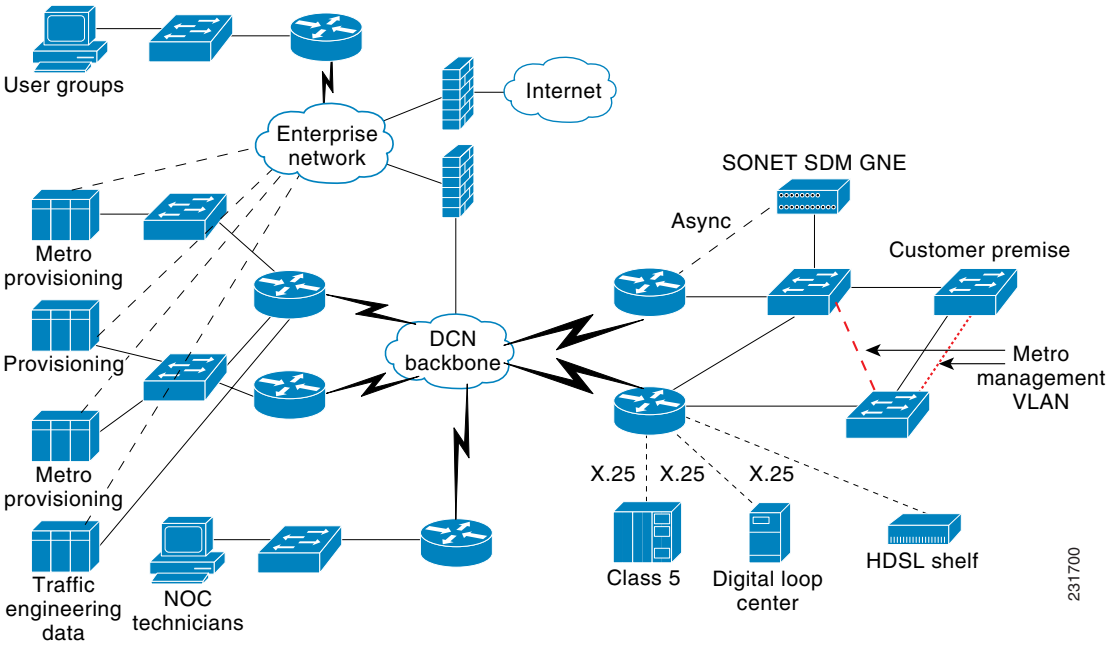

Over time, some service providers have chosen to move their OSSs and the NOC technicians to the enterprise portion of the network. The service providers moved the OSSs to the enterprise network because the OSSs (see Figure 5-4) were dual homed. So the OSSs were connected to the enterprise and the DCN, but the service providers considered this to be a security risk. Moving the OSSs to the enterprise networks eliminated the dual homing. The idea was to authenticate user traffic entering the DCN with a firewall. Also, user groups and NOC technicians were required to authenticate when entering the DCN.

The two-network strategy of both an enterprise and a DCN allows the service provider to control access to the DCN. The firewall is the gatekeeper.

Figure 5-4 Migrating OSSs to an Enterprise Network

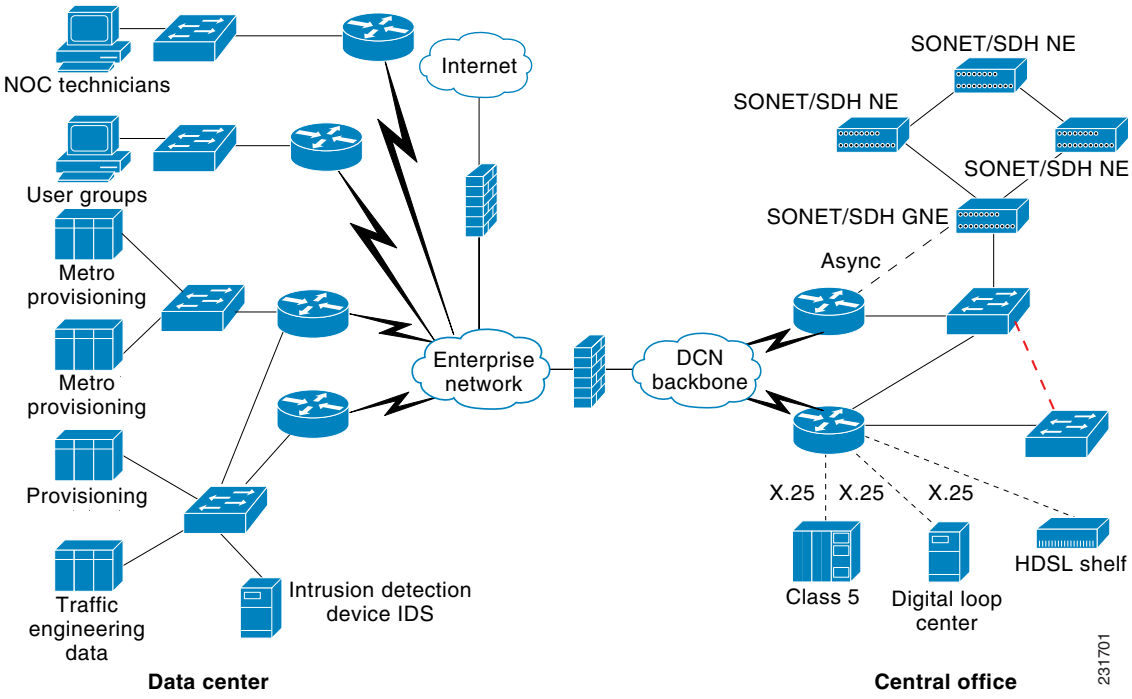

This section describes the business problems that service providers are solving with the MPLS technology, in the following sections:

- Shared Core Network, page 5-5
- Multiple DCNs, page 5-6
- Untrusted Management Traffic, page 5-9
- Enterprise User Traffic VPN, page 5-10
- Service Provider Network for Customer Data with a Management VPN for a DCN, page 5-11

Shared Core Network

Service providers implement MPLS in the DCN in ways different than is typical for an MPLS-based network. For example, in one of the first implementations for MPLS VPNs, service providers combined the enterprise network and DCN. One approach was to build a common core for the DCN and enterprise networks, but keep the distribution and access layers separate for both networks. In this shared core scenario, the core links and core routers are shared between the networks. The enterprise and DCN traffic is placed in separate VPN routing and forwarding (VRF) tables. As shown in Figure 5-5, the enterprise traffic is in its own VRF represented by blue (solid) lines. The DCN traffic is in a separate VRF represented by red (broken) lines. The service provider in this example uses dedicated provider edge (PE) aggregation routers and access routers for the DCN, and uses dedicated PE aggregation routers and access routers for enterprise traffic.

Figure 5-5 Shared Core Utilizing MPLS VPN

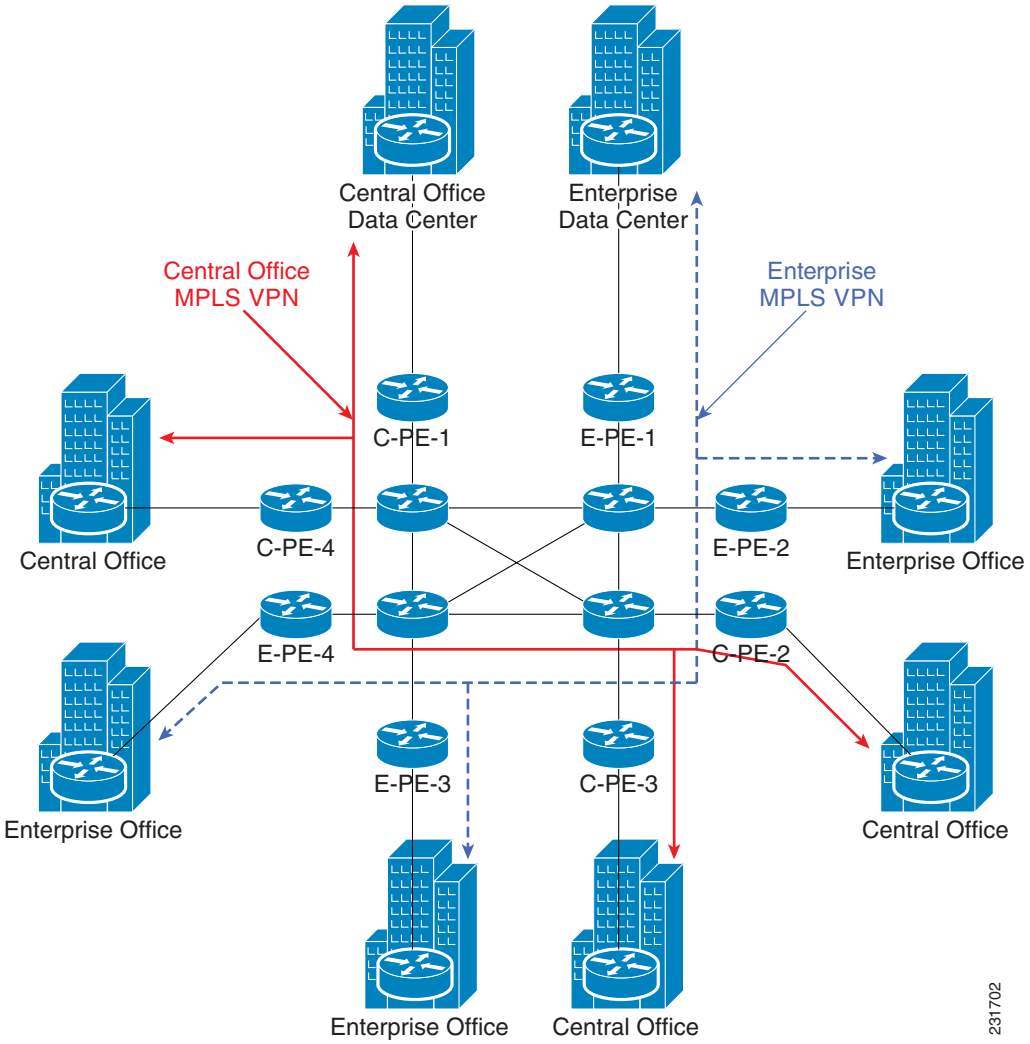

Multiple DCNs

Some service providers have multiple DCNs in their network. Sometimes, service providers have deployed separate DCNs for different functions. For example, some service providers use one DCN to monitor and provision transmission network elements. A second DCN is used for collection of billing data from Class 5 switches. A third DCN is used for the management and provisioning of Class 5 switches.

Figure 5-6 shows an example of multiple DCNs by application.

Figure 5-6 Multiple DCNs by Application

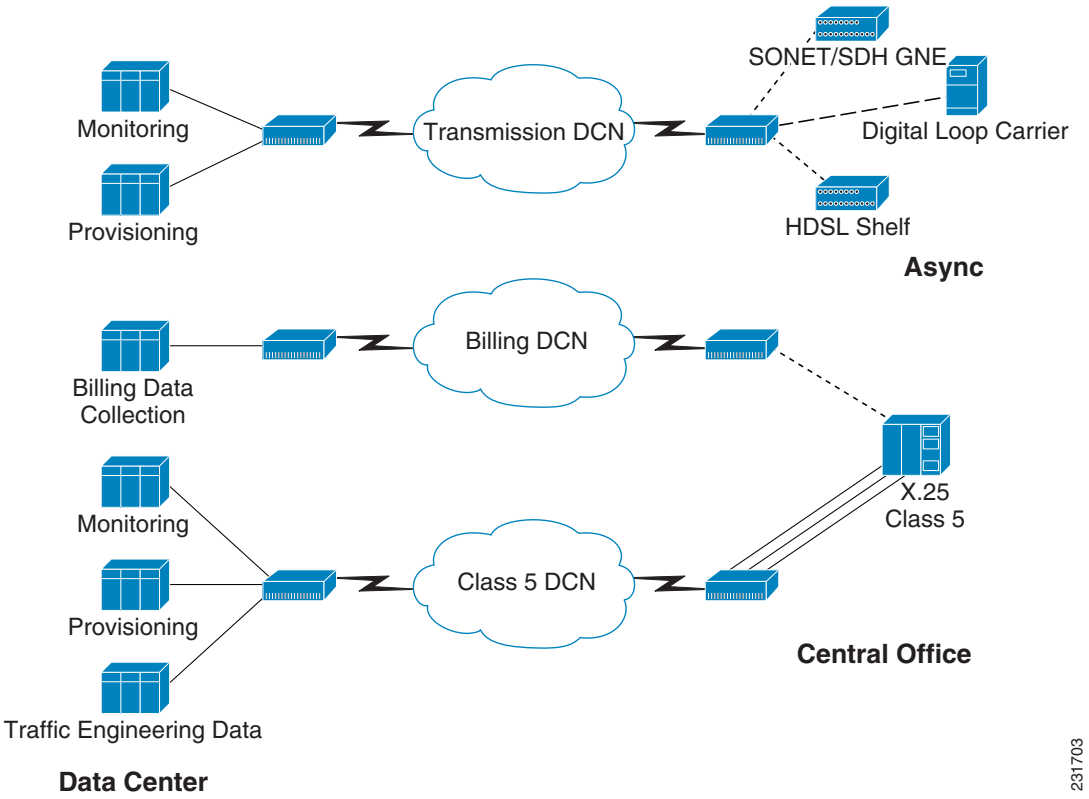

Some manufactures of network elements have required that their traffic be kept separate from other vendors, as shown in Figure 5-7.

Figure 5-7 Multiple DCNs by Vendor

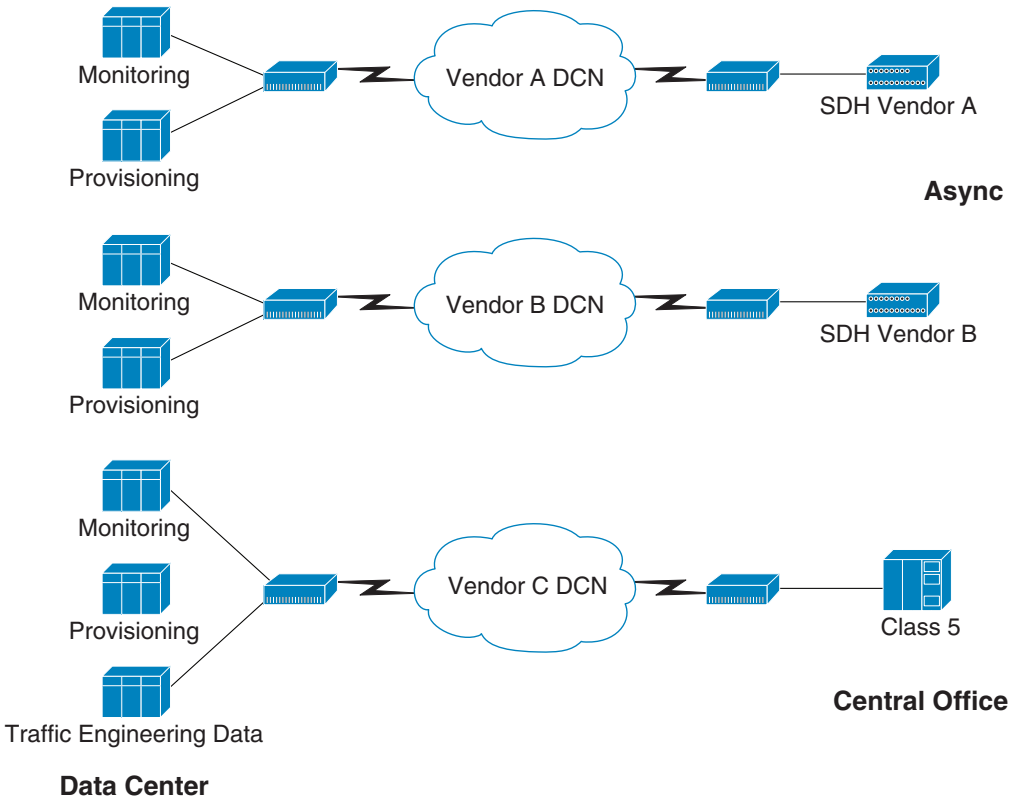

Today, service providers want to consolidate the multiple DCNs into one DCN. In some cases, older DCN equipment no longer being supported by the manufacturer is the compelling reason to consolidate the networks. For example, vendors are getting out of the X.25 equipment market. Also, service providers no longer want to operate and maintain multiple DCNs. Maintaining multiple DCNs requires multiple support staffs, multiple redundant circuits, and multiple support contracts. The MPLS VPN solution is one method for building a common IP core and consolidating multiple DCNs into one DCN. The service provider can create discrete VRFs to isolate vendor traffic or keep applications separate.

In Figure 5-8, vendor A traffic is placed in the red VPN, and vendor B traffic is placed in the blue VPN, so there can be overlapping address space and vendor traffic will be kept separate.

Figure 5-8 MPLS VPN Implementation for Vendor Traffic Separation

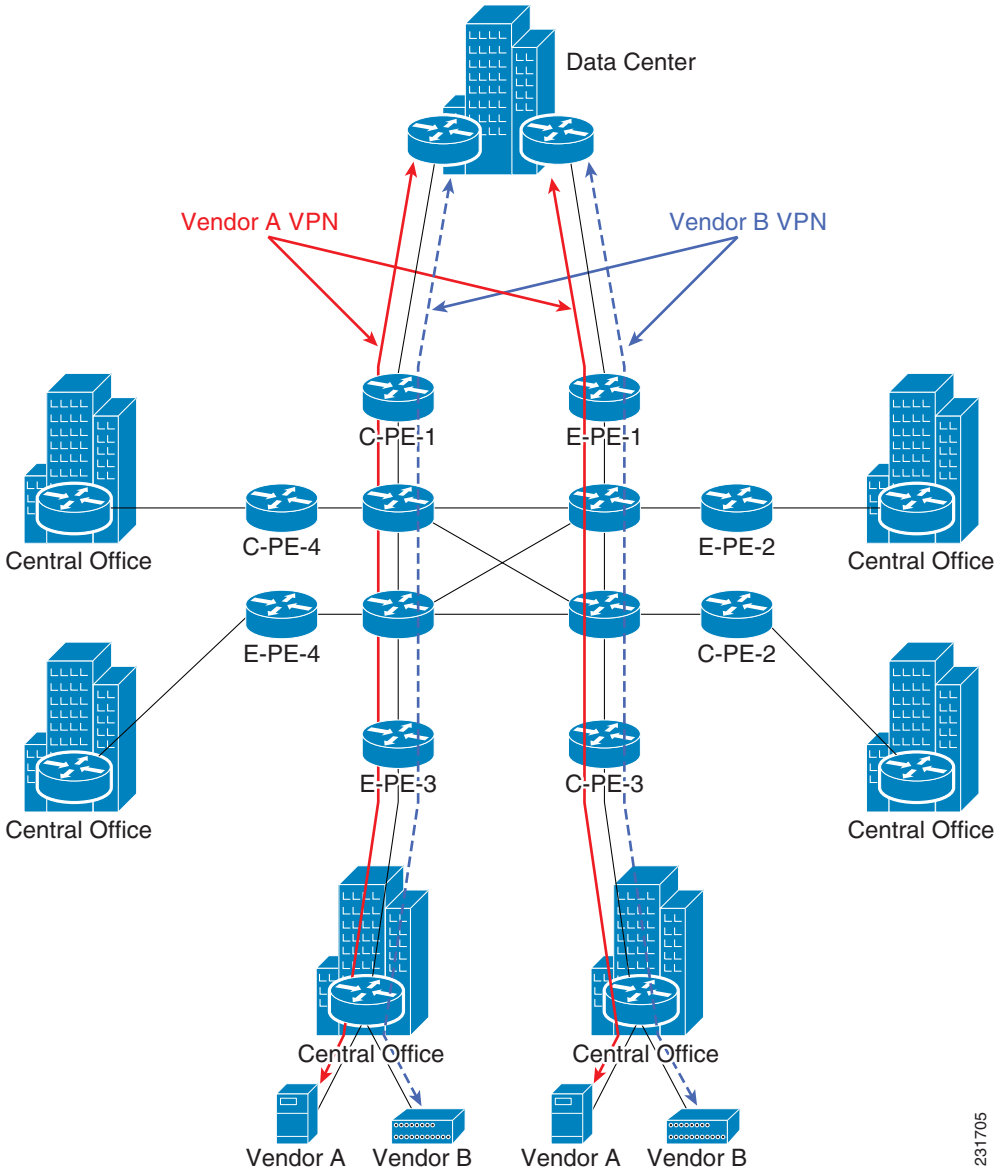

Untrusted Management Traffic

In next generation services such as Metro Ethernet, service providers are deploying on customer premises equipment (CPE) that is managed by the service provider. In other words, the service provider deploys an Ethernet switch on a CPE and manages the switch over a management VLAN connection from the CO, as shown in Figure 5-9. The service provider is concerned about an outsider using the CPE located at the customer premises to break into the network. So the service provider treats management data in the management VLAN as suspect and places the management data in a VPN. Also, the VPN prevents someone from plugging into the customer premises switch and breaking into the classic DCN. If a Metro Ethernet user does manage to break into the management VLAN, the user can only access the management VPN for the Metro Ethernet equipment. The intruder cannot access the classic DCN used for monitoring and provisioning traditional CO equipment such as Class 5 telephone switches, digital loop carrier systems, and other transmission gear.

Figure 5-9 MPLS VPN in the DCN

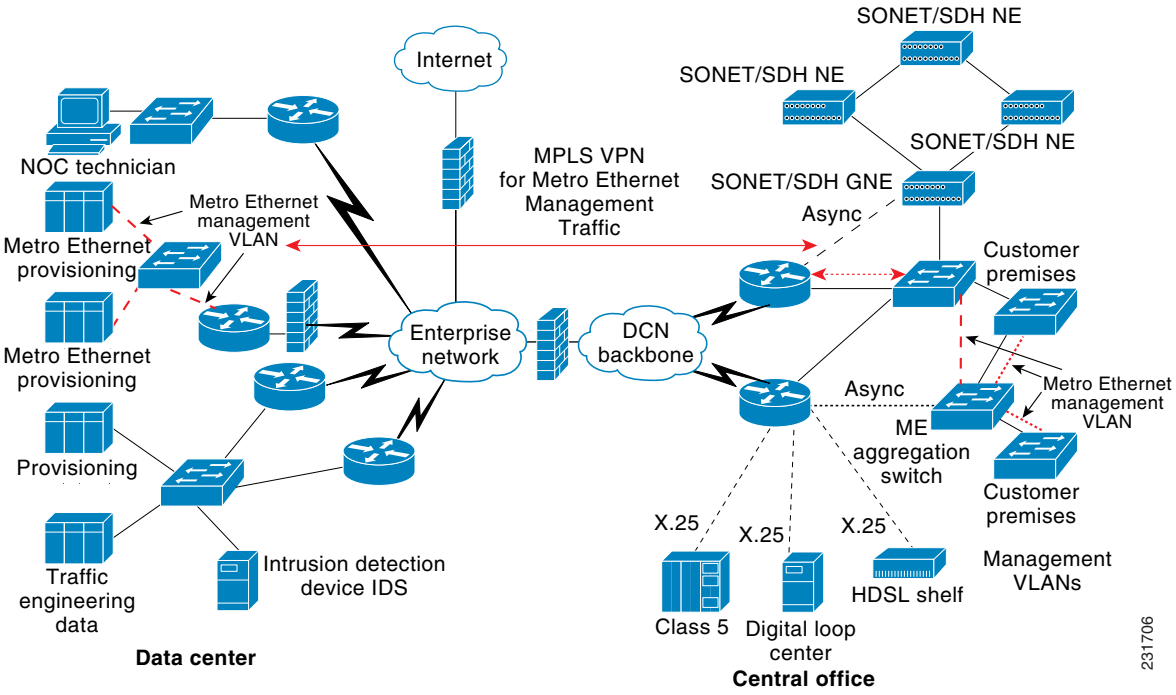

Enterprise User Traffic VPN

Some service providers have built enterprise networks out to their CO to provide connectivity to the technicians and other user groups located in the CO. These service providers have two networks deployed to the CO building. The first network is for the personnel in the building, and the second network is for CO equipment. An alternative is to use the MPLS VPN solution shown in Figure 5-10.

Figure 5-10 MPLS VPN for Enterprise Traffic

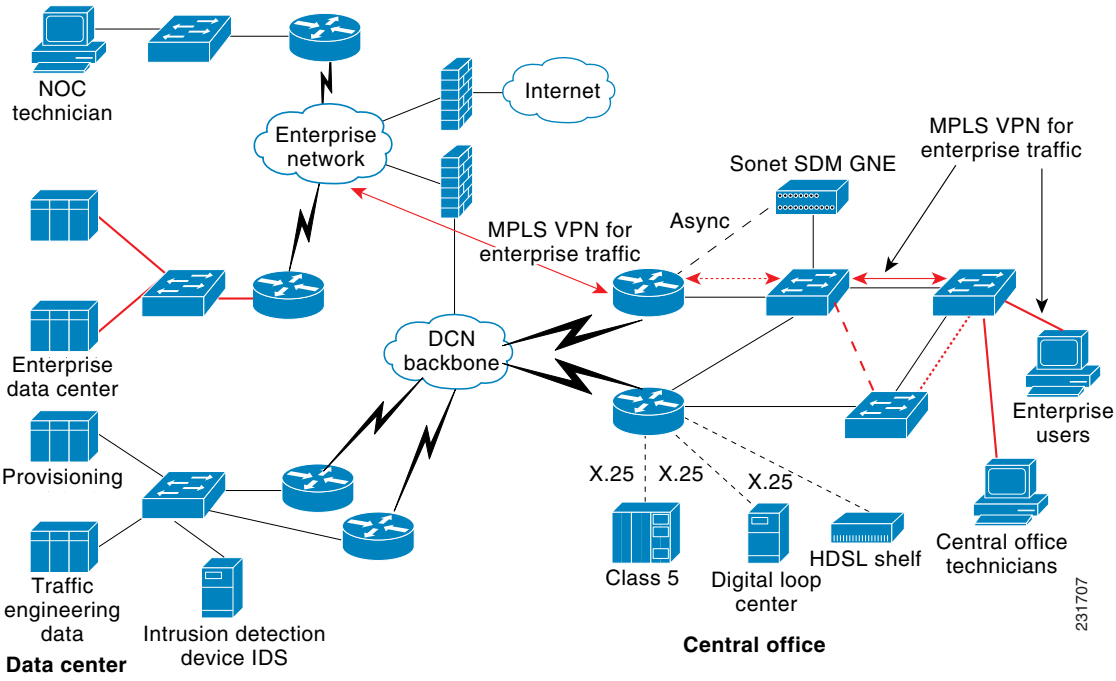

Service Provider Network for Customer Data with a Management VPN for a DCN

Almost all service providers have built separate DCNs. Cisco is beginning to see service providers create a management VPN on the network that provides customer services. The customers that have implemented this option have implemented an alternate access method in the event the user network becomes unstable. The alternate access may be using an existing DCN and connecting the asynchronous console access to the network elements. A second method may be to use ISDN backup access as a secondary WAN access to the DCN access router. So if the service provider lost the management VPN, the DCN access router could use the ISDN dialup to dial back into the NOC. The service provider will need to determine the number of ISDN dialup connections that the service provider would be required to support in case of a large network outage.

The concept of a service provider network connecting customer sites over a VPN and a separate management VPN is shown in Figure 5-11. Both the Customer A and Customer B sites are connected with a VPN. The customers are aggregated together on a PE dedicated to the customers. The CO access DCN routers are connected to PE aggregation routers dedicated to the DCN. Notice in Figure 5-11 that the DCN access routers have an ISDN backup connection to the data center with the OSSs.

Figure 5-11 Service Provider Network for Customer Data with a Management VPN for a DCN

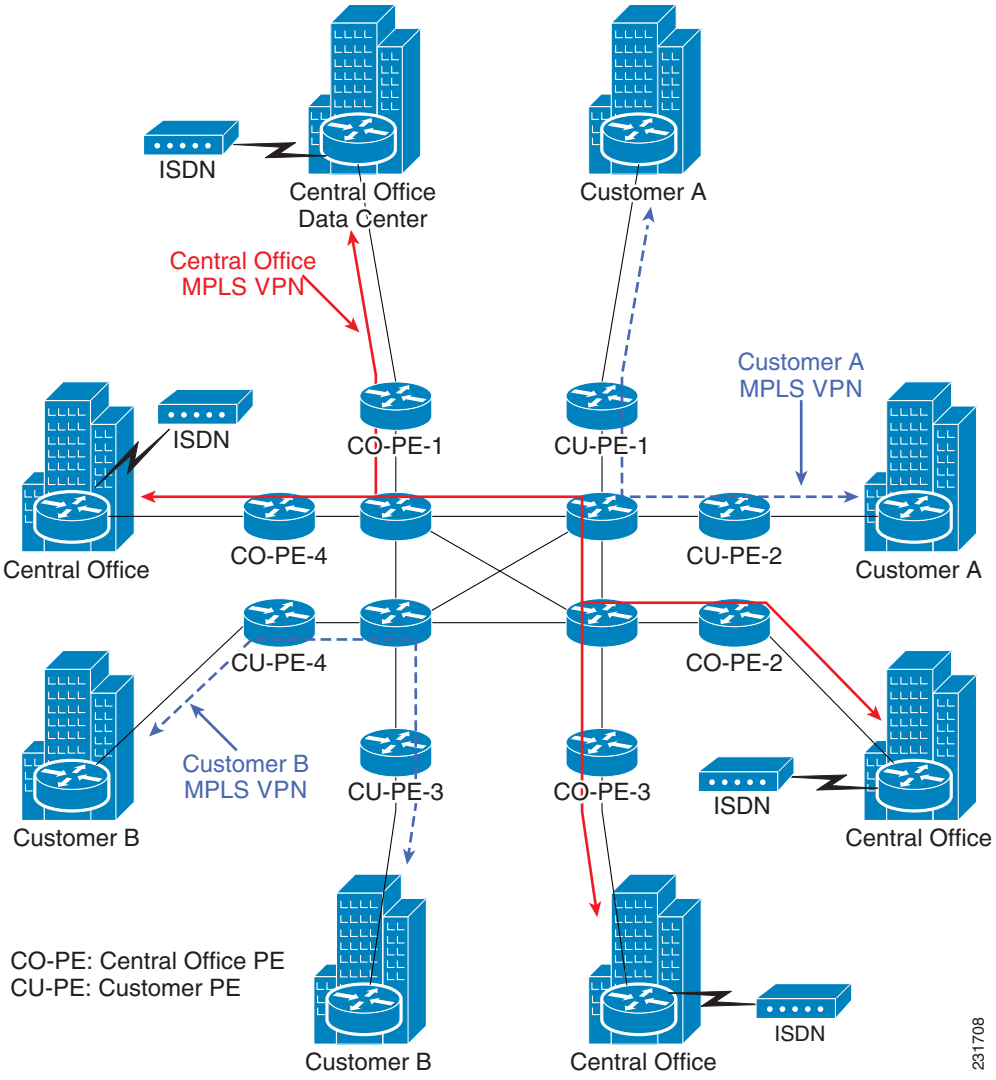

Benefits of MPLS VPNs on a DCN

Deploying MPLS in the DCN offers the following advantages to service providers:

- Outstanding scalability for the DCN. MPLS VPNs easily scale while providing the same level of security and access as Layer 2 technologies. Adding, moving, or integrating COs is simplified. MPLS technology allows the service provider to quickly consolidate multiple DCNs into one DCN and maintain traffic separation.
- Easier network management. The service provider backbone DCN does not need to be reconfigured to implement a new CO; only the PE router where the CO connects to the main network needs to be added to the VRF table. Cisco IP Solution Center (ISC) Version 4.0 can be used to comprehensively manage the DCN on which the MPLS VPN is implemented.
- Quality of service (QoS) can be applied to the MPLS VPN for guaranteed bandwidth. MPLS QoS features enable the efficient use of existing network elements in the DCN to meet growing bandwidth demands. Multiple class of service (CoS) classes can be assigned to traffic on the VPNs.

With MPLS VPNs, advanced IP CoS mechanisms such as Weighted Fair Queuing (WFQ) and Weighted Random Early Detection (WRED) can be applied to VRFs. As described in the "Service Provider Network for Customer Data with a Management VPN for a DCN" section, service providers can have implemented a separate DCN for each vendor so that the vendor traffic is isolated from other vendors. The MPLS VPN solution allows the customer to create one network but meet the vendor's DCN requirement for SLAs and traffic isolation.

- VRF monitoring using a VRF-aware IP SLAs. IP SLAs uses the Service Assurance Agent (SAA) to measure the hop-by-hop response time for the path through the MPLS network, and can be used to monitor a VRF because the IP SLAs is VRF-aware.
- Path selection using MPLS Traffic Engineering/Fast Reroute (TE/FRR) in the DCN. The MPLS TE feature allows network operations to specify routing paths through the network. This feature can also dynamically look for a route to carry a specified bandwidth traffic capacity through the network. FRR functions when a route fails. Convergence times of less than 50 msec can be achieved using FRR. Such a fast recovery prevents applications from timing out and also prevents loss of data.
- MPLS TE paths are called TE tunnels, and accomplish the following:
  - Automatically ensure that the required bandwidth is reserved for the label switched path (LSP) through the network. The bandwidth requirement is provisioned and then checked by the tunnel using Resource Reservation Protocol (RSVP) to verify adequate capacity along the LSP.
  - Keep their own topology map and detect a link or router fault.
  - Provide operating efficiencies and flexibility.

  There is flexibility in terms of routing protocols used at the CO without added overhead. MPLS VPNs greatly simplify service deployment compared to traditional IP VPNs when the number of COs or the number of routes inside the COs increase. MPLS VPNs can provide a simplified managed network without the need to provision a new IPsec tunnel every time a traffic flow is provisioned in the network, and also provide simpler deployment at the CO. MPLS VPNs also do not require translation for the private IP addresses used in the COs.
- Financial benefits: MPLS VPNs provide a robust and scalable platform for IP convergence, enabling service providers to take advantage of lower total cost of ownership, improved operation effectiveness, and improved business performance. Service providers are able to consolidate multiple DCNs into one DCN, which lowers the cost of maintaining and managing the network.

  MPLS VPNs simplify the management of the network, thus reducing cost. The network at the COs also becomes more robust and requires fewer upgrades. Advanced MPLS VPN failure recovery mechanisms also reduce network downtime.
- Load distribution: MPLS can be used to distribute traffic load more effectively throughout the DCNs.
- Fault detection and correction: MPLS provides rapid mechanisms to detect and correct faults. MPLS technology has evolved to provide subsecond convergence and improved network traffic reroutes in times of network outage because of human and network errors.

Supported Platforms

A good reference source for platforms that support MPLS can be found in the *MPLS VPN and Multi-Virtual Route Forwarding Support for Cisco ISR* application note at the following URL: http://www.cisco.com/en/US/products/ps6557/products_white_paper0900aecd8051fbdc.shtml

Design Details

Find design suggestions for an MPLS DCN in the following sections:

Traffic Separation

MPLS VPNs in the DCN allow service providers to differentiate traffic from separate network elements performing different functions in a converged network. In Figure 5-12, traffic is separated using VRFs shown in blue and red in the DCN.

Figure 5-12 Traffic Separation Using Blue and Red VRFs on the DCN

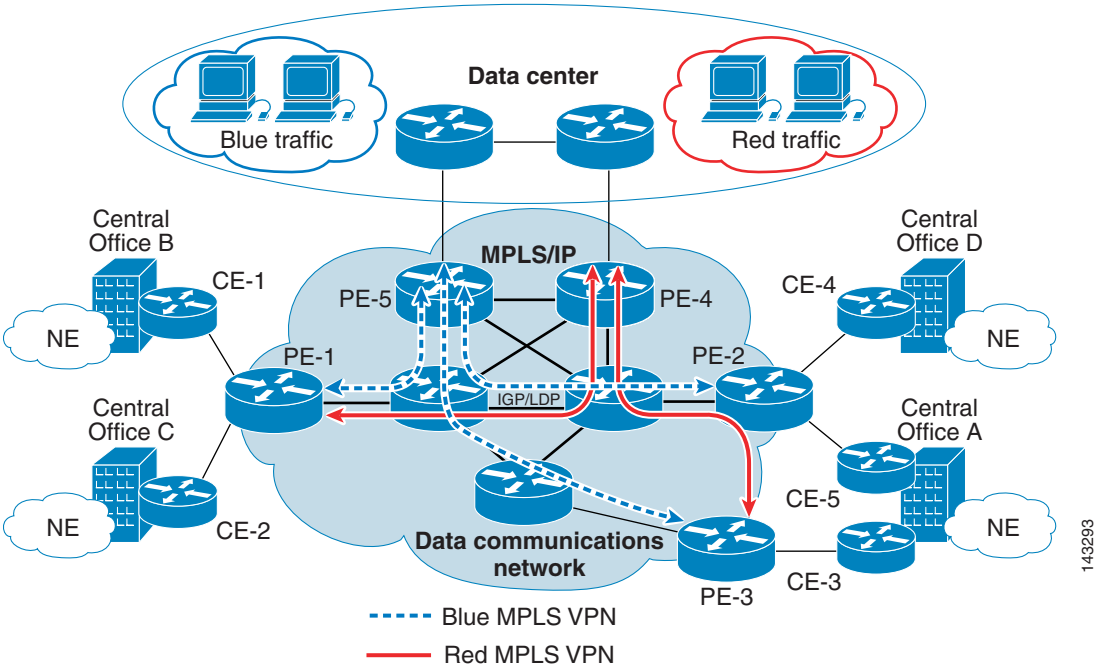

MPLS technology forwards packets while Border Gateway Protocol (BGP) takes care of route distribution over the network. MPLS VPN enforces traffic separation by assigning a unique VRF table to each traffic type.

The routes in the VRF are called the VPN-IPv4 routes and are kept in a table separate from the global routes. In the global routing table, PE routers store Interior Gateway Protocol (IGP) routes and associated labels distributed via Label Distribution Protocol (LDP)/Tag Distribution Protocol (TDP). In the VRFs, PE routers store VPN routes and associated labels distribution through multiprotocol internal BGP (MP-iBGP), which can run over any interior gateway protocol such as Open Shortest Path First (OSPF), Intermediate System-to-Intermediate System (IS-IS), Routing Information Protocol (RIP), and so on.

An MP-iBGP update is sent to all PE neighbors. PEs receive MP-iBGP routes and translate the VPN-IPv4 addresses to determine the VRF. The elements in a specific VPN thus do not see traffic outside the VPN to which they belong.

In Figure 5-13, the incoming interface on the PE determines the forwarding table to use to label-switch the packet.

Figure 5-13 PE Forwards Traffic into Blue and Red VRFs

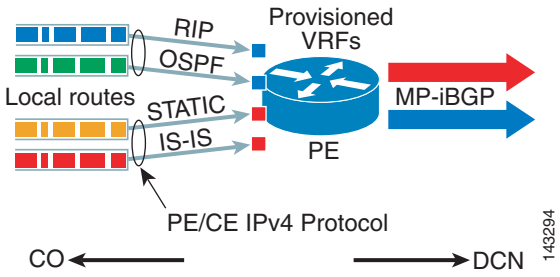

Traffic forwarding within the network is based on labels. LSPs start and end at the PE routers. Because each interface on a PE router is associated with a particular VPN, a packet can enter a VPN only through that interface. Standard IP forwarding can be used between the PE and customer edge (CE) routers. The CEs can use their own routing mechanism.

MPLS VPNs use two labels to forward the data traffic. The top label forwards the traffic to the correct PE router and the label underneath indicates how the other PE should handle that incoming packet. Thus, the VPN-based traffic separation and distribution occurs without IPsec tunneling or any kind of encryption.

Routing Information in a VRF

VPN routing information is propagated through the use of VPN route target communities implemented using the BGP extended communities. VPNv4 routes are exported to and imported from VRFs in the following ways:

- Exported IPv4 routes are brought from the VRF, translated into VPN-IPv4 routes, and inserted into the MP-BGP table. When a VPN route learned from a CE router is injected into BGP, the list of route target community values is thus set from an export list of route targets associated with the VRF from which the route was learned.
- Imported VPN-IPv4 routes are brought from the MP-BGP table, translated into IPv4 routes, and inserted into the VRF. An import list of route target extended communities is thus associated with each VRF.

PEs need to know which route is intended for which VRF. Figure 5-14 shows a proposed VPN architecture in which the data center PEs import all VRF routes from the COs. Routes are differentiated using red and blue **route-target** commands.

Figure 5-14 Proposed VPN Architecture-Data Center PEs Import All VRF Routes from COs

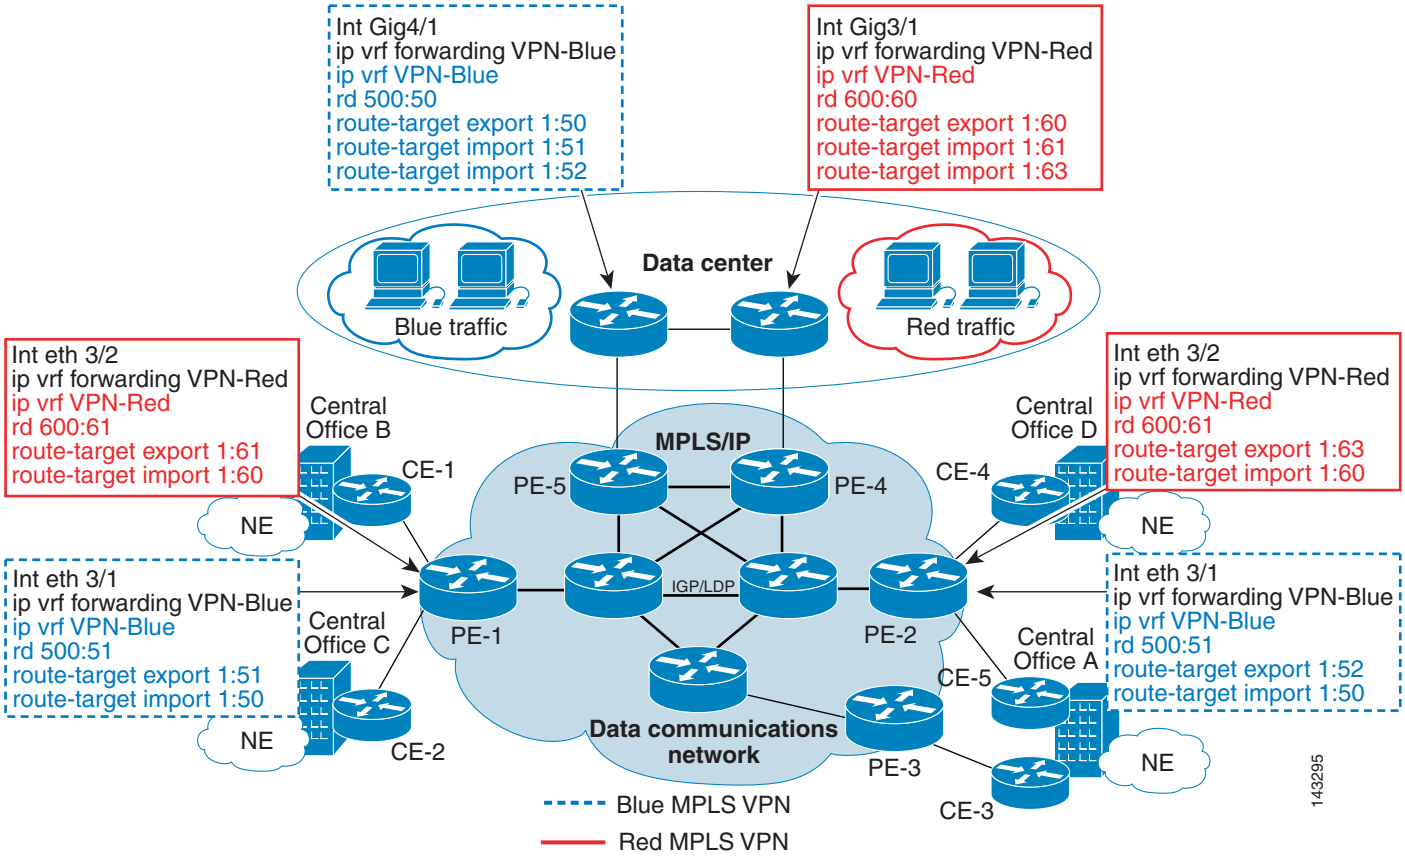

Deploying MPLS VPNs on a DCN

Key MPLS infrastructure elements to keep in mind when deploying MPLS on the service provider network are described in the following sections:

- Core Architecture, page 5-16
- Route Reflectors in an MPLS Network, page 5-17
- IPsec-Aware MPLS VPN, page 5-18

Core Architecture

The first step in building an MPLS VPN architecture is the design and development of the core MPLS network. The routers in the core are called Provider (P) routers and the routers on the edge are called the Provider Edge (PE) routers. The following are key elements to consider when designing the core:

- The initial reachability information exchanged between PE and P routers is achieved with loopback addresses. Cisco recommends that these addresses be assigned to identify the device in terms of its location in the network cloud. Loopback interfaces are virtual interfaces that are never shut down unless the device is completely isolated or powered off. Loopback interfaces can be set up to allow label switching to use these addresses as the endpoints of a tunnel.

Cisco recommends that loopback interfaces be defined on PE routers for a specific VPN. VPN reachability can be checked after the PE router is installed at the CO. The advantage with this approach is that the VPN reachability on the PE of the CO in question can be verified before the CO is activated on that MPLS VPN.

- IGP routing such as OSPF, RIP, and IS-IS needs to be set up among all the P and PE routers in the core network.
- Labels should be distributed using LDP. The global routing table created by IGP is used to distribute label information by P and PE routers in the MPLS cloud.
- BGP should be enabled for VRF information. BGP is used to locate the hop closest to a destination. BGP establishes the peer relation to its neighbor using TCP port 179. With iBGP, the neighbor need not be directly connected to the BGP speaker; IGP is used to achieve this.

Route Reflectors in an MPLS Network

BGP route reflectors (RRs) are not essential for a DCN MPLS network to function. BGP routing requires all the iBGP speakers to be fully meshed. This requirement is not scalable when there are a large number of iBGP hops in the core. The RR feature is a good solution to this iBGP mesh issue. Figure 5-15 shows how RRs reduce network complexity.

Figure 5-15 Route Reflectors Reduce Network Complexity

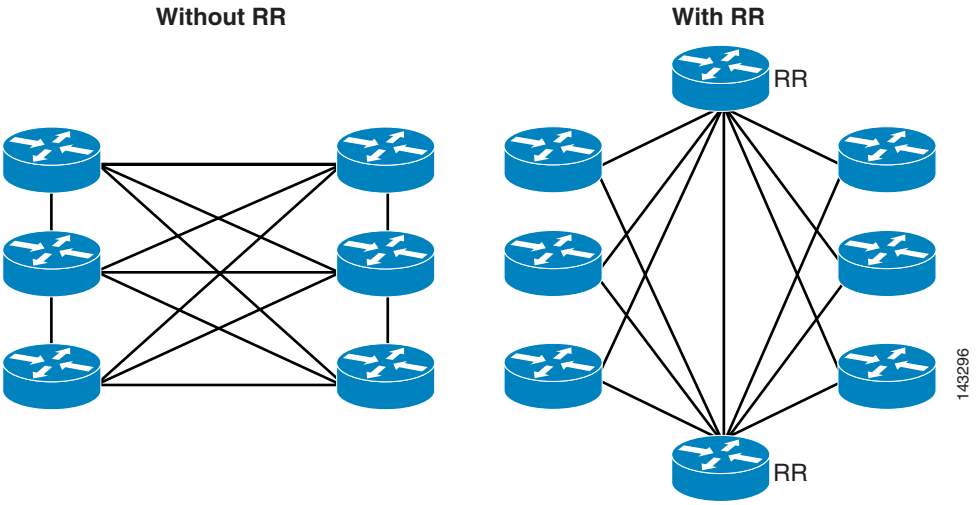

The choice of RRs in a DCN should include the following factors:

- RR should be used when the core is handling a large amount of edge devices establishing peering relationships.
- More than one RR should be deployed for redundancy when deploying in a DCN.

IPsec-Aware MPLS VPN

Figure 5-16 shows how IPsec tunnels can exist between the CO and the DCN PE router. The two peers can secure different kinds of data streams where each IPsec tunnel uses a separate set of traffic associations. For example, some data streams might be console traffic generated from an optical network, while other data streams might be billing data mapped to an IPsec tunnel.

Figure 5-16 IPsec-Aware MPLS VPN

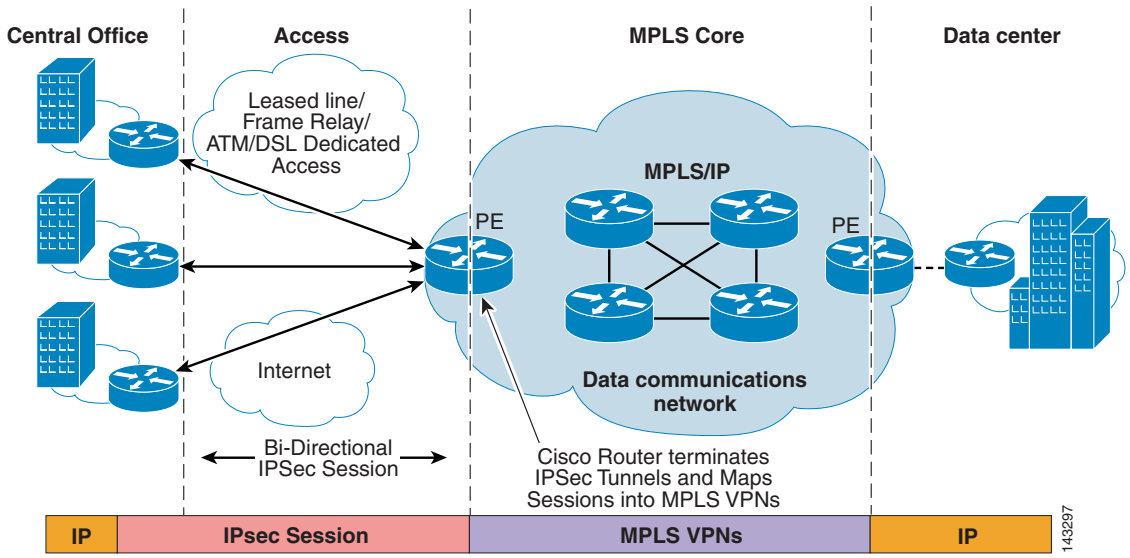

IPsec sessions can be terminated on the provider edge of the MPLS/IP backbone, and each of these tunnels can be mapped into their respective MPLS VPNs. The mapping between the IPsec and the MPLS VPN can be done based on the deployment model and the policies that need to be applied.

Configuration Examples for MPLS VPNs on a DCN

This section provides the following configuration examples for the blue and red VRFs, as described in the "Traffic Separation" section on page 5-14:

- Data Center PE (PE-5)-Blue Routes Imported: Example, page 5-19
- Data Center PE (PE-4)-Red Routes Imported: Example, page 5-19
- Central Center PE (PE-1)-Blue and Red Routes Exported: Example, page 5-19
- Central Center PE (PE-2)-Blue and Red Routes Exported: Example, page 5-20

See Figure 5-12 for an example of the routes.

Data Center PE (PE-5)-Blue Routes Imported: Example

The following configuration defines the VRF:

```
ip vrf VPN-Blue
rd 500:50
route-target export 1:50
route-target import 1:51
route-target import 1:52
```

The following configuration applies the VRF to the interface facing the DCN:

```
interface GigabitEthernet4/1
description link TO PE-5
ip vrf forwarding VPN-Blue
ip address 10.1.2.2 255.255.255.0
no ip directed-broadcast
ip pim sparse-mode
load-interval 30
negotiation auto
tag-switching ip
```

Data Center PE (PE-4)-Red Routes Imported: Example

The following configuration defines the VRF and the import and export communities:

```
ip vrf VPN-Red
rd 600:60
route-target export 1:60
route-target import 1:61
route-target import 1:62
```

The following configuration applies the VRF to the interface facing the DCN:

```
interface GigabitEthernet3/1
description link TO PE-4
ip vrf forwarding VPN-Red
ip address 10.1.2.1 255.255.255.0
no ip directed-broadcast
ip pim sparse-mode
load-interval 30
negotiation auto
tag-switching ip
```

Central Center PE (PE-1)-Blue and Red Routes Exported: Example

The following configuration defines the VRF:

```
ip vrf VPN-Red
rd 600:61
route-target export 1:61
route-target import 1:60

ip vrf VPN-Blue
rd 500:51
route-target export 1:51
route-target import 1:50
```

The following configuration applies the VRF to the interface facing the DCN:

```
interface ethernet 3/1
description link TO CE-2
ip vrf forwarding VPN-Blue
ip address 10.1.100.2 255.255.255.0
no ip directed-broadcast
ip pim sparse-mode
load-interval 30
negotiation auto
tag-switching ip


interface ethernet 3/2
description link TO CE-1
ip vrf forwarding VPN-Red
ip address 10.1.100.5 255.255.255.0
no ip directed-broadcast
ip pim sparse-mode
load-interval 30
negotiation auto
tag-switching ip
```

Central Center PE (PE-2)-Blue and Red Routes Exported: Example

The following configuration defines the VRF:

```
ip vrf VPN-Red
rd 600:61
route-target export 1:63
route-target import 1:60

ip vrf VPN-Blue
rd 500:51
route-target export 1:52
route-target  import 1:50
```

The following configuration applies the VRF to the interface facing the DCN:

```
interface ethernet 3/1
description link TO CE-5
ip vrf forwarding VPN-Blue
ip address 10.1.100.2 255.255.255.0
no ip directed-broadcast
ip pim sparse-mode
load-interval 30
negotiation auto
tag-switching ip

interface ethernet 3/2
description TO link CE-4
ip vrf forwarding VPN-Red
ip address 10.1.100.5 255.255.255.0
no ip directed-broadcast
ip pim sparse-mode
load-interval 30
negotiation auto
tag-switching ip
```