CHAPTER 3

MPLS-Based VPN MAN Reference Topology

MAN Topology

The MAN topology described in Figure 3-1 shows a network that serves three different organizations in a single corporation.

Figure 3-1 Typical Self-Managed MPLS MAN Deployment

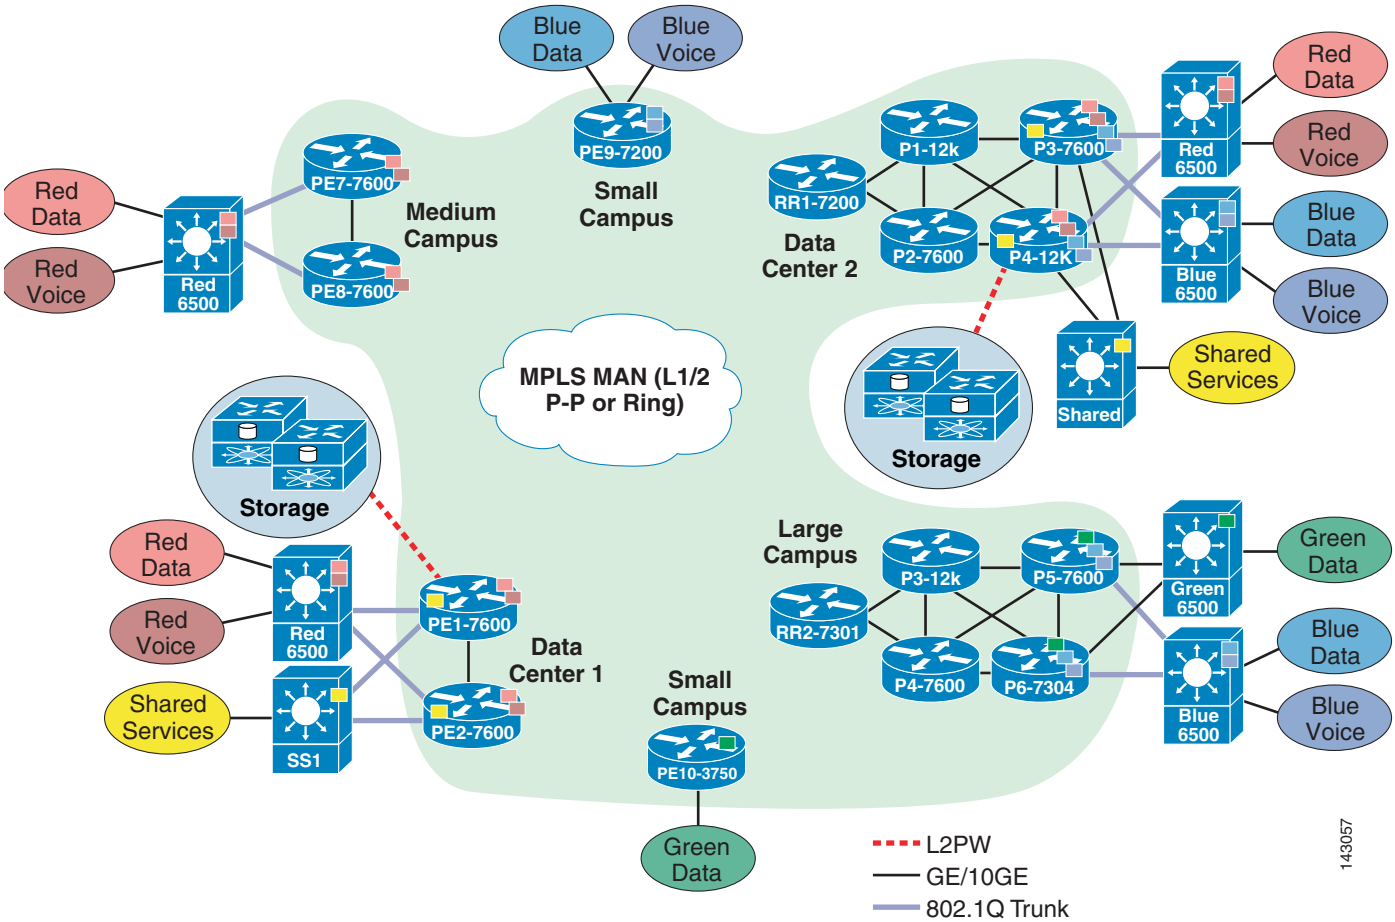

This corporate network spans five different locations: large campus (LC), medium campus (MC), two small campuses (SC1, SC2), and two data centers (DC1, DC2).

Each organization requires that their traffic be kept separate from the traffic of any other organization. Within each organization, voice and data traffic should also be separated. To meet these requirements, traffic is segmented into five different VPNs:

- Red-data to serve data traffic in the Red organization
- Red-voice to serve voice traffic in the Red organization
- Blue-data to serve data traffic in the Blue organization
- Blue-voice to serve voice traffic in the Blue organization
- Green-data to serve the Green organization

The Blue organization resides in a large campus (LC) and a small campus 1 (SC1). The Red organization resides in a medium campus (MC). Data center 2 (DC2) serves both the Blue and Red organizations. Data center 1(DC1) serves only the Red organization. The Blue and Green organizations are using EIGRP and the Red organization is using OSPF to relay VPN subnet information.

Because voice and data traffic for each VPN user need to be segmented end-to-end, it is essential that the segmentation starts on the first Layer 3 device in the path. Note that segmentation from the access layer to a distribution layer is maintained using VLANs. Multi-VRF is deployed on the distribution layer devices (DLs), which are the first Layer 3 devices to maintain and carry the segmentation throughout the network. Users in each VLAN are directly mapped to associated VRFs on the Multi-VRF devices (DL switches) and the traffic is kept separate as it traverses through each campus and data center.

Because EIGRP and OSPF are widely adapted routing protocols in enterprise networks, testing is done with EIGRP as a core IGP and then repeated using OSPF. Note that this does not affect the VPN site routing protocols in use. The Blue and Green organizations continue to use EIGRP and the Red organization continues to use OSPF for distributing each VLAN subnet into MPLS PE regardless of the IGP used in the core.

Frame-mode unsolicited downstream on demand LDP mode is used to allocate and distribute MPLS labels. The traffic in the core is switched based on the pre-calculated incoming and outgoing labels.

Notice that dual-homing at the MPLS edge as well as redundant links in the core and distribution layer to MPLS edge exist. For example, DL1, which connects red-data and red-voice VPNs, is dual-homed to both PE1 and PE2. The distribution layer to the MPLS edge, as well as the load balancing in the core, is done by IGP Cisco Express Forwarding. To make sure the traffic is load balanced between ingress and egress edge points, MP-iBGP multipath load balancing is enabled on all the PEs. To ensure multiple paths are stored in the forwarding tables, unique RDs are used on each PE for the same VPN.

The network also has backdoor links between sites DC2 and MC for the Red organization and between sites LC and DC2 for the Blue organization. The backdoor links are not to be used for load balancing with the MPLS VPN network, but are backup links to the MPLS VPN network. To prevent routing loops for the Red organization, OSPF sham link is configured on PE3, PE4, PE7, and PE8. To prevent routing loops for the Blue organization, EIGRP SoO is configured on all the Blue site-connected PEs.

- Core AS is 1 when EIGRP is used as the IGP.
- Core area is 0 when OSPF is used as the IGP.
- MP-iBGP peers are in AS 1.
- EIGRP edge is using AS 10 for blue-data, 11 for blue-voice, and 12 for green-data.
- The Red organization uses OSPF Area 0.

VPN Information

Table 3-1 shows VPN information.

Table 3-1 VPN Information

Organization	VPN Name	RD	RT
Blue	blue-data	10:1055	10:105
	blue-voice	10:1066	10:106
Green	green-data	10:107	10:107
Red	red-data	10:1031	10:103
	red-voice	10:1042	10:104

Inventory of Devices

Table 3-2 lists an inventory of devices.

Table 3-2 Inventory of Devices

	Reference	Actual	Loop0
DC1	PE1	7600-DC1-PE1	125.1.125.5/32
	PE2	7600-DC1-PE2	125.1.125.6/32
	SS1	7600-DC1-SS1	125.1.125.17/32
DC2	P1	12k-DC2-P1	125.1.125.1/32
	P2	7600-DC2-P2	125.1.125.2/32
	RR1	7200-DC2-RR1	125.1.125.15/32
	PE4	12k-DC2-PE4	125.1.125.8/32
	PE3	7600-DC2-PE3	125.1.125.7/32
LC	P3	12k-LC-P3	125.1.125.3/32
	PE5	7600-LC-PE5	125.1.125.9/32
	PE6	7304-LC-PE6	125.1.125.10/32
	RR2	7200-LC-RR2	125.1.125.16/32
	P4	7600-LC-P4	125.1.125.4/32
	SS2	7600-LC-SS2	125.1.125.18/32
MC	PE8	7600-MC-PE8	125.1.125.12/32
	PE7	7600-MC-PE7	125.1.125.11/32
SC1	PE9	7200-SC1-PE9	125.1.125.13/32
SC2	PE10	3750-SC2-PE10	125.1.125.14/32

Building a MAN MPLS VPN Network

Layer 1 and some Layer 2 as well as IP addresses are enabled on the interfaces. Interfaces are in up and up spoofing mode.

To build a MAN MPLS VPN network, complete the following steps:

Step 1 Build the MPLS core:

a. Enable EIGRP (or OSPF) on the core routers, RR, and PE core-facing interfaces.

b. Enable Cisco Express Forwarding:

```
Router(config)# ip cef
```

c. Select the LDP router id. Enable LDP on the core routers, RR, and PE core-facing interfaces:

```
Router(config)#mpls ldp router-id loopback0 force
Router(config)#interface interface #
Router(config)#mpls ip
Router(config)#mpls label protocol ldp
```

Step 2 Build the MPLS Layer 3 VPNs:

a. Enable MP-iBGP on PEs and establish BGP peering among all the PE routers and RRs.

Although route reflectors are not necessary, they are used for scalability purpose. BGPv4 requires that all the iBGP devices be fully meshed. With route reflector in use, PEs would only have to peer with a route reflector instead of peering with each other. This reduces having to fully mesh PEs by n(n-1), n being total numbers of PEs. As the RRs do not need to be in the data path, they could be local on any router that does not require a powerful switching capacity.

When route reflectors are in use, and all the outgoing updates have the same policy, it further helps fine tune the network by using peer-groups on route reflectors. This reduces the number of outgoing updates (per client) that a route reflector has to generate. Here two route reflectors are used for high availability purpose. All the PEs peer with RR instead of with each other.

PEs peer with route reflectors to exchange VPNv4 routing information:

```
7200-DC2-RR1(config)#router bgp 1
7200-DC2-RR1(config-router)#
```

Route reflectors peer with PEs to reflect VPNv4 routing information learned from other PEs:

```
Peer-Group Setup:
7200-DC2-RR1(config)#router bgp 1
7200-DC2-RR1(config-router)#neighbor CampusPE peer-group
       7200-DC2-RR1(config)#neighbor CampusPE remote-as 1
7200-DC2-RR1(config)#neighbor CampusPE update-source
Loopback0

7200-DC2-RR1(config)#neighbor <PE loopback#> peer-group CampusPE
```

VPNv4 BGP peering between PEs and RRs on a route reflector:

```
7200-DC2-RR1(config-router)#  address-family vpnv4
7200-DC2-RR1(config-router-af)# neighbor CampusPE activate
7200-DC2-RR1(config-router-af)# neighbor CampusPE route-reflector-client

7200-DC2-RR1(config-router-af)# neighbor CampusPE send-community extended

7200-DC2-RR1(config-router-af)# neighbor <PE loopback#> peer-group CampusPE
```

VPNv4 BGP peering between PEs and RRs on a PE.

Enable PEs to exchange VPNv4 routing information to the RRs.

```
7600-DC1-PE1(config)#router bgp 1
7600-DC1-PE1(config-router)#no synchronization
7600-DC1-PE1(config-router)#bgp log-neighbor-changes
7600-DC1-PE1(config-router)#neighbor <RR1 loopback ip#> remote-as 1

7600-DC1-PE1(config-router)#neighbor <RR1 loopback ip#> update-source Loopback0

7600-DC1-PE1(config-router)# address-family vpnv4
7600-DC1-PE1(config-router-af)# neighbor <RR1 loopback ip#> activate

7600-DC1-PE1(config-router-af)# neighbor 125.1.125.15 send-community extended
```

If the network does not have RRs, set up VPNv4 peering with other PEs in the network by using the PEs loopback IP addresses. It is important to set up BGP peering before VPN site routing information is redistributed into BGP for easier management and troubleshooting of the network.

b. Create a VPN.

MPLS allows support for multiple VPNs in a scalable way. VLANs terminating on the distribution layer can be individually mapped into a VRF with its own routing instance. The VRF-name is a unique and case-sensitive value. It is used to identify the VRF.

```
7600-DC1-PE1(config)#ip vrf vrf name
7600-DC1-PE1(config)#ip vrf red-data
```

c. Create an RD under its associated VRF.

Use a unique RD per VRF.

```
7600-DC1-PE1(config-vrf)#rd route-distinguisher unique value
7600-DC1-PE1(config-vrf)#rd 10:1031
```

Note You can assign only one RD to a VRF. If an RD needs to be changed for any reason after VRFs are operational, make sure to save the entire VRF-related configuration. Changing the RD requires removing the current RD, which removes the associated VRF. VPN site routes as well as MP-iBGP redistribution also need to be reconfigured.

d. Create an RT under its associated VRF.

```
7600-DC1-PE1(config-vrf)#route-target {import | export | both} route- target
7600-DC1-PE1(config-vrf)#route-target export 10:103
7600-DC1-PE1(config-vrf)#route-target import 10:103
```

e. Bind a VRF to an interface.

To specify the interfaces belonging to a VPN (VRF), use the following command:

```
Router(config)#interface interface #
Router(config-if)#ip vrf forwarding <vrf-name>
7600-DC1-PE1(config)#interface GigabitEthernet1/7.1
7600-DC1-PE1(config-subif)#  ip vrf forwarding red-data
```

Note You can only bind one VRF to an interface. After an interface is bound to a VRF, its IP address is not be part of the global routing table. You need to examine the associated VRF instance.

f. Redistribute VPN site routing information into MP-iBGP.

Reachability information for the adjacent VPN sites at the ingress PE needs to be sent to the egress PE so that it can update its adjacent VPN sites. To establish the connectivity between two VPN sites, redistribute routing into MP-iBGP at the ingress PE and back to the remote site IGP at the egress PE.

For OSPF:

```
!
7600-DC1-PE1(config)#router bgp 1
7600-DC1-PE1(config-router)#address-family ipv4 vrf red-data

7600-DC1-PE1(config-router-af)# redistribute ospf 2 vrf red-data match internal
external 1 external 2

7600-DC1-PE1(config-router-af)# maximum-paths ibgp unequal-cost 6

7600-DC1-PE1(config-router-af)# no auto-summary
7600-DC1-PE1(config-router-af)# no synchronization
7600-DC1-PE1(config-router-af)# exit-address-family
!
```

For EIGRP:

```
!
7600-DC2-PE3(config)#router bgp 1
7600-DC2-PE3(config-router)#address-family ipv4 vrf blue-voice

7600-DC2-PE3(config-router-af)#redistribute eigrp 11
7600-DC2-PE3(config-router-af)#maximum-paths ibgp unequal-cost 8

7600-DC2-PE3(config-router-af)#no auto-summary
7600-DC2-PE3(config-router-af)#no synchronization
7600-DC2-PE3(config-router-af)#exit-address-family
!
```

g. Redistribute remote site routing information learned via MP-iBGP into local site IGP.

Configuration varies based on the routing protocol used in VPN sites.

For OSPF:

```
!
7600-DC1-PE1(config)#router ospf 1 vrf red-data
 7600-DC1-PE1(config-router)#log-adjacency-changes
        7600-DC1-PE1(config-router)#redistribute bgp 1 subnets
7600-DC1-PE1(config-router)#network <site's subnet> area 0
!
```

Note OSPF process 1 is used for VRF red-data. The routes learned from remote VPN site via BGP are distributed into OSPF process 1 to update adjacent VPN segments. A different process number would be used to support additional VPN.

For EIGRP:

```
       !
7600-DC2-PE3(config)#router eigrp 10
7600-DC2-PE3(config-router)# address-family ipv4 vrf blue-voice

7600-DC2-PE3(config-router-af)#redistribute bgp 1 metric 1000000 100 255 1 1500

 7600-DC2-PE3(config-router-af)#network <VPN site Subnet or network #>

 7600-DC2-PE3(config-router-af)#maximum-paths 8
```

```
7600-DC2-PE3(config-router-af)#no auto-summary
7600-DC2-PE3(config-router-af)#autonomous-system 11
 7600-DC2-PE3(config-router-af)#exit-address-family
      !
```

h. For OSPF VPN sites with backdoor links, configure a sham-link on a pair of ingress/egress PEs:

- Configure an additional /32 loopback interface and bind the associated VRF to this loopback interface:

```
interface Loopback1
ip vrf forwarding red-data
ip address 125.1.125.103 255.255.255.255
```

- Advertise the loopback interface address through BGP and not through OSPF process:

```
7600-DC2-PE3(config)#router bgp 1
7600-DC2-PE3(config-router)# address-family ipv4 vrf red-data
7600-DC2-PE3(config-router-af)#redistribute connected metric 1
```

- Associate the sham-link with an existing OSPF area and configure under the associated VRF OSPF process between a pair of ingress egress PEs:

```
7600-DC2-PE3(config)#router ospf 1 vrf red-data
7600-DC2-PE3(config-router)#area 0 sham-link 125.1.125.103 125.1.125.107
7600-DC2-PE3(config-router)# area 0 sham-link 125.1.125.103 125.1.125.108
```

Note The sham-link is set up between PE3 and PE5 and PE3 and PE6 as the backdoor link exists between these two sites.

i. For EIGRP sites with backdoor links, configure SoO on PE and CE interfaces.

Step 3 Enable segmentation on multi-VRF devices.

Note This is done on a distribution layer device. Access layer VLANs are terminated and mapped into the associated VRF on a distribution layer device (DL1). End-to-end segmentation across multiple campuses is achieved by maintaining this segmentation using dedicated interfaces for each VPN subnet to connect to the MPLS ingress PE.

a. Create VRFs on the distribution layer device.

```
!
ip vrf red-data
 rd 10:103
!
ip vrf red-voice
 rd 10:104
!
```

b. Bind a VRF to a pair of ingress-egress interfaces.

Ingress interfaces connecting to VLAN:

```
!
interface GigabitEthernet5/7
 description To RT - port 103/1
 no ip address
interface GigabitEthernet5/7.1
 encapsulation dot1Q 505
 ip vrf forwarding red-data
 ip address 125.1.1.65 255.255.255.224
```

```
!
interface GigabitEthernet5/7.2
 encapsulation dot1Q 506
 ip vrf forwarding red-voice
 ip address 125.1.10.9 255.255.255.252
!
```

Egress interfaces connecting to PE:

```
!
interface GigabitEthernet5/7.1
 encapsulation dot1Q 505
 ip vrf forwarding red-data
 ip address 125.1.1.65 255.255.255.224
!
interface GigabitEthernet5/7.2

encapsulation dot1Q 506
 ip vrf forwarding red-voice
 ip address 125.1.10.9 255.255.255.252
!
```

c. Redistribute VPN site routing information.

Use separate routing processes per VPN to exchange routing information with PEs:

```
!
router ospf 1 vrf red-data
 log-adjacency-changes
 capability vrf-lite
 redistribute static subnets
 network <> area <>
 maximum-paths 6
!
router ospf 2 vrf red-voice
 log-adjacency-changes
 capability vrf-lite
 redistribute static subnets
 network <> area <>
 maximum-paths 6
!
```