

# Designing MPLS Extensions for Customer Edge Routers

## Support of Multi-VRFs in a Single CE Extending Limited MPLS Functionality to CE Routers

### Executive Summary

Virtual Private Networks (VPNs) have become increasingly important as more and more businesses are connecting to a Service Provider's network. Keeping data private as it travels across a Service Provider's network is the ultimate concern for both the Service provider providing the VPN service as well as the companies sending the data.

While deploying a single VPN service model would simplify network operations, this approach cannot satisfy diverse customer requirements because each subscriber (company) is unique. To satisfy a broad range of customer requirements, service providers must offer subscribers a portfolio that contains a number of different VPN service delivery models. A number of VPN models have been proposed.

#### Traditional VPNs

- Frame Relay (Layer 2)
- ATM (Layer 2)

#### CPE-based VPNs

- L2TP and PPTP (Layer 2)
- IPSec (Layer 3)

#### Provider Provisioned VPNS

- MPLS-based Layer 2 VPNs
- MPLS-VPNs based on RFC2547bis (Layer 3)

Service providers are already generating a tremendous amount of interest in MPLS-VPNs as a mechanism to simplify WAN operations for a diverse set of customers. As a result of this surge of interest in service provider-based MPLS-VPNs, a new feature has been developed to extend the MPLS-VPNs to the branch office. This new feature is called Multi-VRF CE.

This paper focuses on first establishing a basic understanding of MPLS-VPNs and then developing a detailed understanding of Multi-VRF CE, including the requirements needed and a configuration example.

## MPLS-VPN Overview

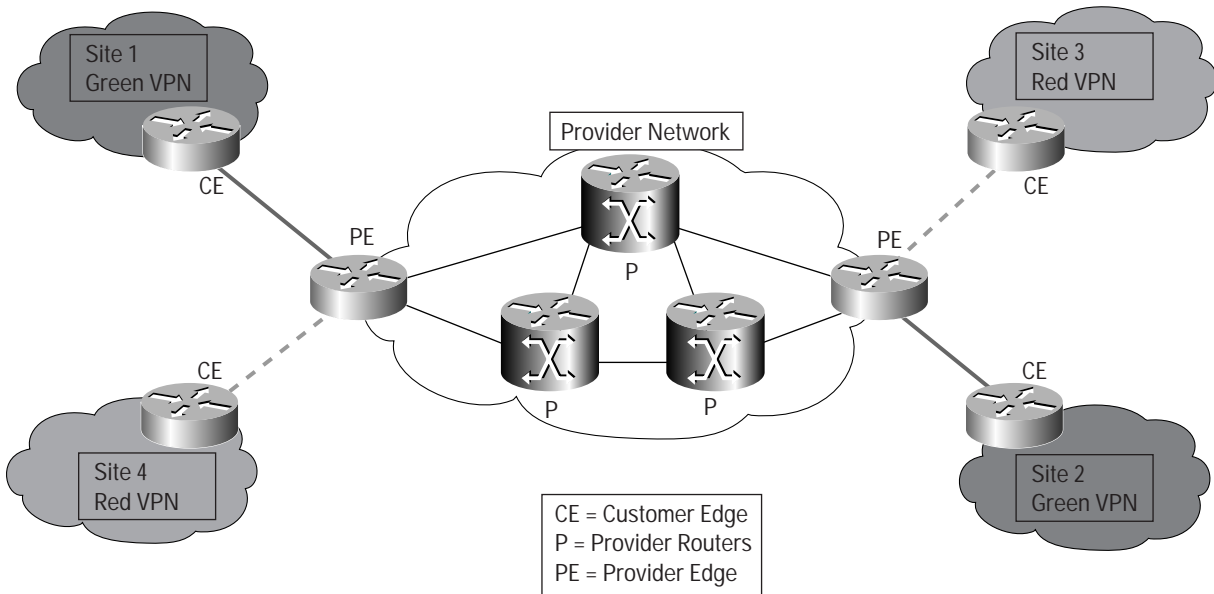
Before discussing Multi-VRF CE, a basic understanding of MPLS-VPNs is necessary as Multi-VRF CE extends the functionality of the current MPLS-VPN model out to the branch office.

MPLS-VPNs define a mechanism that allows service providers to use their IP backbone to provide VPN services to their customers. This model can also be termed BGP/MPLS-VPNs because BGP is used to distribute VPN routing information across the provider's backbone and because MPLS is used to forward VPN traffic from one VPN site to another. An MPLS-based VPN network has three major components:

1. VPN route target communities—A VPN route target community is a list of all other members of a VPN community. VPN route targets need to be configured for each VPN community member.
2. Multiprotocol BGP (MP-BGP) peering of VPN community PE routers—MP-BGP propagates VRF reachability information to all members of a VPN community. MP-BGP peering needs to be configured in all PE routers within a VPN community.
3. MPLS forwarding—MPLS transports all traffic between all VPN community members across a VPN service-provider network.

In the MPLS-VPN model a VPN is defined as a collection of sites sharing a common routing table. A customer site is connected to the service provider network by one or more interfaces, where the service provider associates each interface with a VPN routing table. A VPN routing table is called a *VPN routing and forwarding (VRF) table*. Figure 1 illustrates the fundamental building blocks of an MPLS-VPN.

Figure 1 MPLS-VPN Network Components





## Customer Edge Devices

A customer edge (CE) device provides a customer access to the service provider network over a data link to one or more provider edge (PE) routers. The CE device is an IP router that establishes an adjacency with its directly connected PE routers. After the adjacency is established, the CE router advertises the site's local VPN routes to the PE router and learns remote VPN routes from the PE router. Any router in Cisco's portfolio can act as a CE router as the CE router only exchanges routing information to the PE router. Typically in a branch office, the Cisco 2600 series serves as the CE router.

## Provider Edge Routers

PE routers exchange routing information with CE routers using static routing, RIPv2, OSPF, or EIGRP. While a PE router maintains VPN routing information, it is only required to maintain VPN routes for those VPNs to which it is directly attached. This design eliminates the need for PE routers to maintain all of the service provider's VPN routes.

Each PE router maintains a VRF for each of its directly connected sites. Multiple interfaces on a PE router can be associated with a single VRF if these sites all participate in the same VPN. Each VPN is mapped to a specific VRF, i.e., an interface on the PE router and a site is not associated with a VRF. The PE router is able to maintain multiple forwarding tables that support the per-VPN segregation of routing information. After learning local VPN routes from CE routers, a PE router exchanges VPN routing information with other PE routers using IBGP. Only the routes pertinent to the PE router's VRFs are exchanged.

The following is a list of router platforms supported at the provider edge.

- Cisco 3640 series
- Cisco 3660 series
- Cisco 3700 Series
- Cisco 7200 series
- Cisco 7500 series
- Cisco 10000 series
- Cisco 10720 series
- Cisco 12000 series

## Provider Routers

A provider (P) router is any router in the provider's network that does not attach to CE devices. P routers function as MPLS transit LSRs when forwarding VPN data traffic between PE routers. Since traffic is forwarded across the MPLS backbone using a two layer label stack, P routers are only required to maintain routes to the provider's PE routers; they are not required to maintain a specific VPN routing information for each customer site.

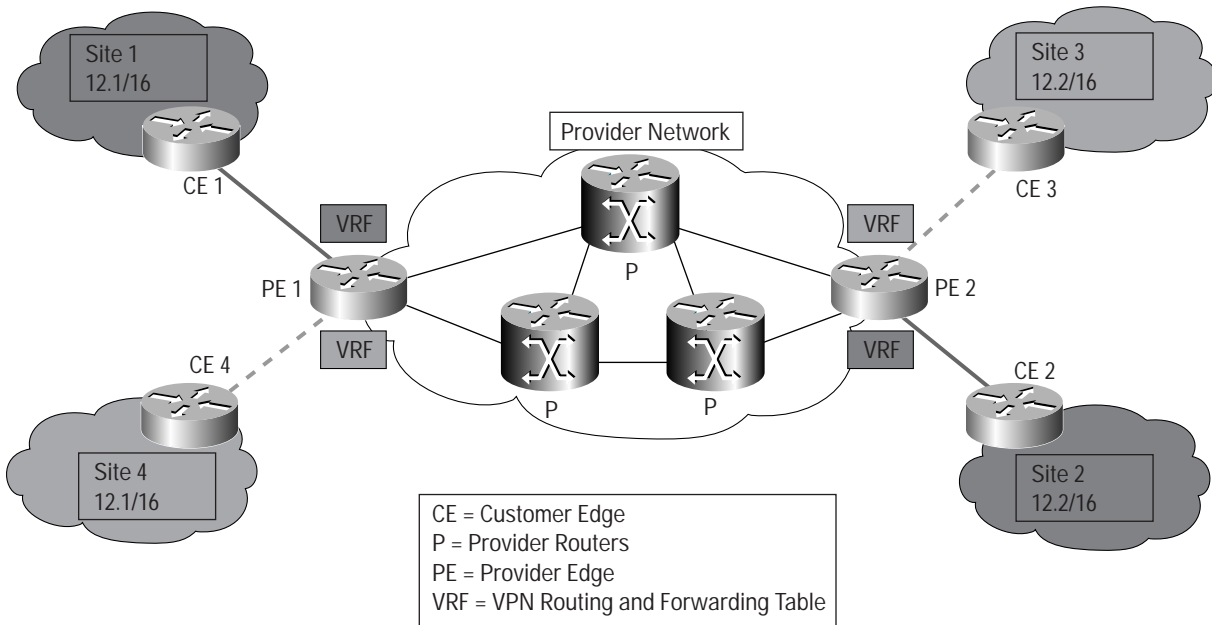
The following is a list of router platforms supported at the provider core.

- Cisco 3600 series
- Cisco 7200 series
- Cisco 7500 series
- Cisco 8540 series
- Cisco 8800 series
- Cisco 12000 series

## Sample MPLS-VPN Network Topology

Figure 2 illustrates a sample network topology where a single service provider delivers an MPLS-VPN service to different enterprise customers. In this network there are two PE routers connected to four different customer sites

Figure 2 Sample MPLS-VPN Network Topology



The inter-site connectivity can be described by the following policies.

- Any host in Site 1 can communicate with any host in Site 2.
- Any host in Site 2 can communicate with any host in Site 1.
- Any host in Site 3 can communicate with any host in Site 4.
- Any host in Site 4 can communicate with any host in Site 3.

To make sure these policies are followed, two processes must be used.

1. Exchange of routing information between the CE and PE routers at the edge of the provider's backbone and between the PE routers across the provider's backbone.
2. Establishment of Label Switch Paths (LSPs) across the provider's backbone between PE routers.

### PE/CE Exchange of Routing Information

In this example, PE 1 is configured to associate VRF Green with the interface or sub-interface over which it learns routes from CE1. When CE 1 advertises the route for prefix 12.1/16 to PE 1, PE 1 installs a local route to 12.1/16 in VRF Green.

PE 1 advertises the route for 12.1/16 to PE 2 using IBGP. Before advertising the route, PE 1 selects an MPLS label (for this example, 426) to advertise with the route and assigns its loopback address as the BGP next hop for the route.

MPLS-VPN supports overlapping address spaces by the use of route distinguishers (RDs) and the VPN-IPv4 address family. It also constrains the distribution of routing information among PE routers by the use of route filtering based on BGP extended community attributes (route targets).

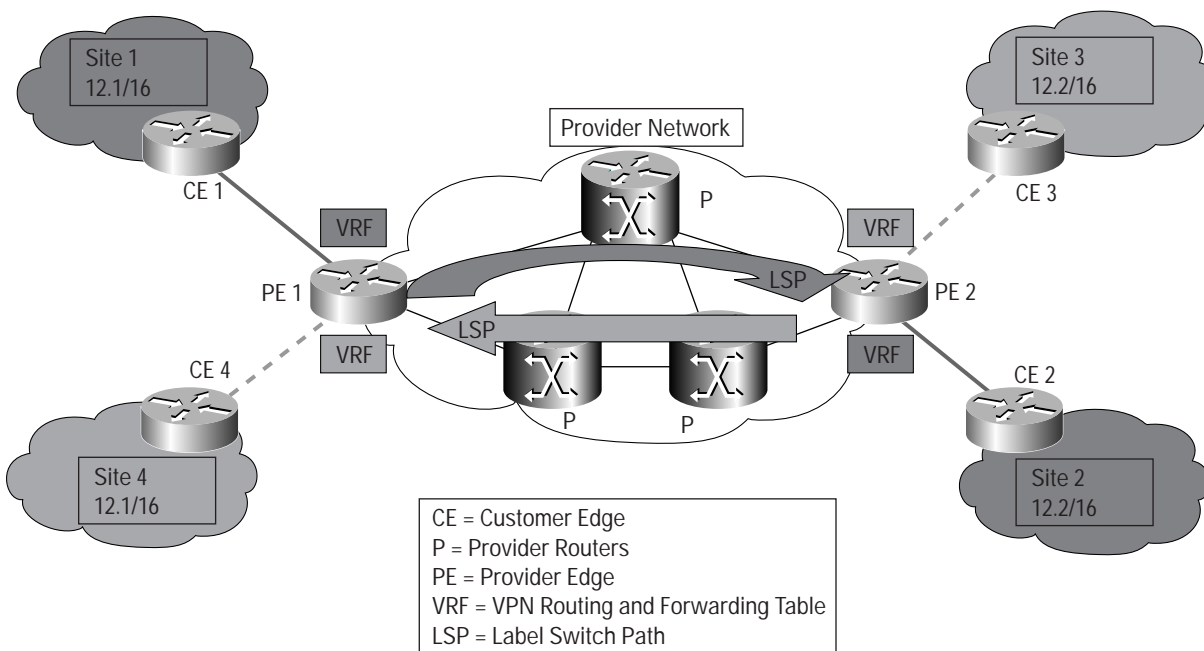


When PE 2 receives PE 1's route advertisement, it determines if it should install the route to prefix 12.1/16 into VRF Green by performing route filtering based on the BGP extended community attributes carried with the route. If PE 2 decides to install the route in VRF Green, it then advertises the route to prefix 12.1/16 to CE 2.

### LSP Establishment

In order to use MPLS to forward VPN traffic across the provider's backbone, LSPs must be established between the PE router that learns the route and the PE router that advertises the route (Figure 3).

Figure 3 LSPs between Site 1 and Site 2



LSPs can be established and maintained across the service provider's network using one of the following techniques.


- Label Distribution Protocol (LDP) for assigning labels associated with the PE loopback
- BGP for assigning VPN specific labels
- Resource Reservation Protocol (RSVP) for traffic engineering tunnels

Note that there can be a single LSP or several parallel LSPs (perhaps with different QoS capabilities) established between PE routers. Also, note that it is possible to use RSVP to assign labels for PE loopbacks, although this is not recommended. LDP provides more flexibility and is less manually intensive to configure.

## Traffic Flow

Figure 4 shows the flow of VPN traffic across the service provider's backbone from one customer site to another customer site. Assume that Host 1.2.3.4 at Site 2 wants to communicate with Server 2.1.3.8.

Figure 4 Traffic Flow from Site 2 to Site 1



Host 1.2.3.4 forwards all data packets for Server 2.1.3.8 to its default gateway. When a packet arrives at CE 2, it performs a longest-match route lookup and forwards the IPv4 packet to PE 2.

PE 2 receives the packet, performs a route lookup in VRF Green. User traffic is forwarded from PE 2 to PE 1 using MPLS with a label stack containing two labels. For this data flow, PE 2 is the ingress LSR for the LSP and PE 1 is the egress LSR for the LSP. Before transmitting a packet, PE 2 pushes the label, 426 in this example, onto the label stack making it the bottom (or inner) label. This label is originally installed in VRF Green when PE 2 receives PE 1's IBGP advertisement for the route 12.1/24. Next, PE 2 pushes the label stack making it the top (or outer) label. When the packet arrives from CE2, PE2 inserts a VPN label for that customer (inner label), does a lookup in the proper VPN FIB (LFIB), and then inserts a label for forwarding to PE1 (outer label).

After creating the label stack, PE 2 forwards the MPLS packet on the outgoing interface to the first P router along the LSP from PE 2 to PE 1. P routers switch packets across the core of the provider's backbone network based on the top (outer) label. The penultimate router to PE 1 pops the top label (exposing the bottom or inner label) and forward the packet to PE 1.

When PE 1 receives the packet, it pops the label creating a native IPv4 packet. PE 1 uses the bottom label (426) to identify the directly attached CE that is the next hop to 12.1/16. Finally, PE 1 forwards the native IPv4 packet to CE 1, which forwards the packet to Server 2.1.3.8 at Site 1.

For additional information on MPLS, <http://www.cisco.com/go/mpls> contains references to MPLS and MPLS-VPNs.



## Multi-VRF CE Overview

MPLS-VPNs provide security and privacy as traffic travels through the provider network. The CE router has no mechanism to guarantee private networks across the traditional LAN network. Traditionally to provide privacy, either a switch needed to be deployed and each client be placed in a separate VLAN or a separate CE router is needed per each client's organization or IP address grouping attaching to a PE. Figure 5 and 6 shows the traditional solutions for LAN security within an MPLS-VPN network design.

Figure 5 Traditional Branch Office Network Topology in an MPLS-VPN Network Utilizing a Switch to Segment LAN Traffic


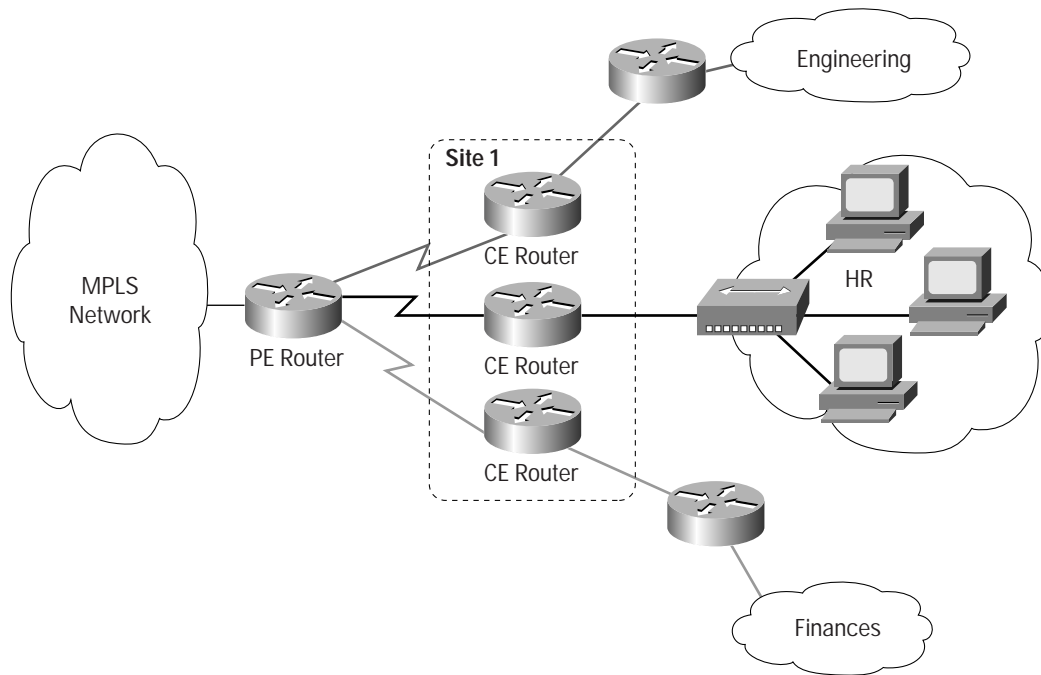


Figure 6 Traditional Branch Office Network Topology in an MPLS-VPN Network Utilizing Separate Routers to Segment LAN Traffic



These solutions are both costly to the customer as additional equipment is needed and requires more network management and provisioning of each client site.

Multi-VRF CE is a new feature, introduced in Cisco IOS release 12.2(4)T, that addresses these issues. Multi-VRF CE extends limited PE functionality to a CE router in an MPLS-VPN model. A CE router now has the ability to maintain separate VRF tables in order to extend the privacy and security of an MPLS-VPN down to a branch office rather than just at the PE router node.

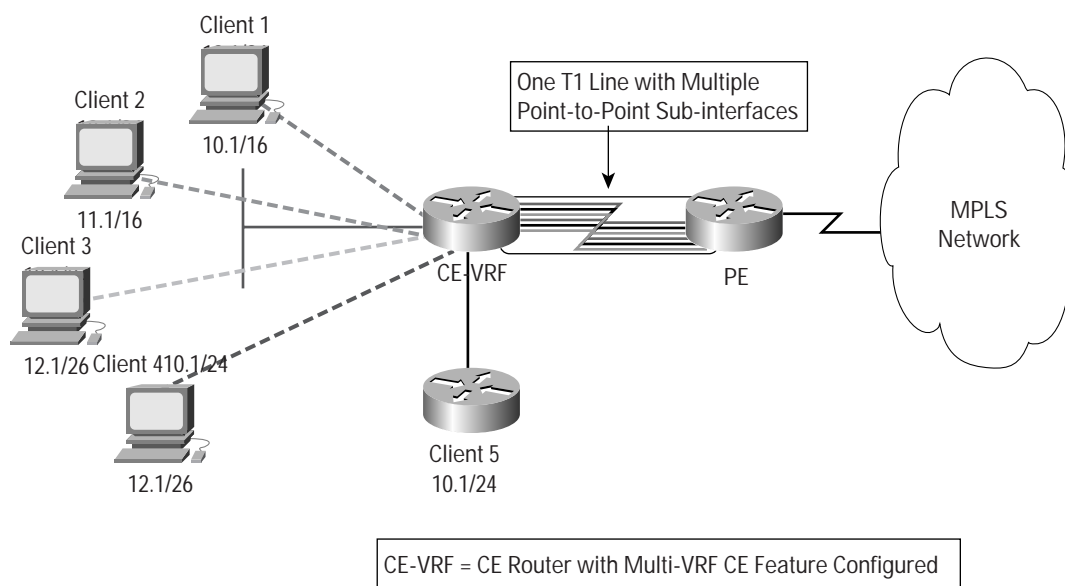
CE routers use VRF interfaces to form a VLAN-like configuration on the customer side. Each VRF on the CE router is mapped to a VRF on the PE router. With Multi-VRF CE, the CE router can only configure VRF interfaces and support VRF routing tables. Multi-VRF CE extends SOME of the PE functionality to the CE router—there is no label exchange, there is no LDP adjacency, there is no labeled packet flow between PE and CE. The only PE-like functionality that is supported is the ability to have multiple VRFs on the CE router so that different routing decisions can be made. The packets are sent toward the PE as IP packets.



## Operational Model

Figure 7 illustrates one method in which Multi-VRF CE can be used at the CE router. The connection from the PE router to the provider network uses the same path that was discussed in the MPLS-VPN overview section.

Figure 7 One Method for Deploying Multi-VRF CE



The CE router using Multi-VRF CE can segment its LAN traffic by placing each client or organization with its own IP address space either on separate Ethernet interfaces such as Client 5 or through one Fast Ethernet interface segmented into multiple sub-interfaces. Each sub-interface contains its own IP address space to separate each different client.

When receiving an outbound customer data packet from a directly attached interface, the CE router then performs a route lookup in the VRF that is associated with that site. The specific VRF is determined by the interface or sub-interface over which the data packet is received. Support for multiple forwarding tables makes it easy for the CE router to provide the per-VPN segregation of routing information before it is sent to the PE router. The use of a T1 line with multiple point-to-point sub-interfaces allows traffic from the CE router to the PE router to be segmented into each separate VRF.

Using Figure 7, the data path is as follows from the Clients to the PE router with Multi-VRF CE configured on the CE router.

1. CE-VRF learns Client 1's VPN Red routes from a sub-interface of the Fast Ethernet interface directly attached to CE-VRF. CE-VRF then installs these routes into VRF Red.
2. PE 1 learns Client 1's VPN Red routes from the CE-VRF and installs them into VRF Red.
3. Local VPN Blue routes from Client 2 are not associated with VPN Red and are not imported into VRF Red.

In this model, the CE router associates a specific VRF by the clients connected to its interfaces and exchanges that information with the PE router. Routes are installed in the VRF on the Multi-VRF CE. There also needs to be a routing protocol or a static route that propagates routes from a specific VRF on the Multi-VRF router to the same VRF on the PE router.

Note: Platform restrictions with respect to processor and memory apply i.e., the number of VRFs supported by the CE router is dependent on the platform, processing power and available memory. As a design practice, you must factor in the processing and memory requirements for routing processes, management, packet forwarding etc.

#### Benefits of Multi-VRF CE

1. Without the use of cryptographic techniques (IPSec), security on customer's LAN is equivalent to that supported by existing Layer 2 (ATM or Frame Relay) connections with out the use of an additional switch.
2. Only one CE router is needed thus simplifying provisioning and network management rather than a multiple CE router solution.
3. CE router has VRF functionality to provide VPN routing information. Less routing updates to manage.
4. Overlapping Customer address spaces  
VPN customers often manage their own networks and use private address spaces. If customers do not use globally unique IP addresses, the same 32-bit IPv4 address can be used to identify different systems in different VPNs. The result can be routing difficulties because BGP assumes that each IPv4 address it carries is globally unique. To solve this problem, MPLS-VPNs supports a mechanism that converts nonunique IP address into globally unique addresses by combining the use of VPN-IPv4 address family with the deployment of Multiprotocol BGP Extensions (MP-BGP).
5. No need for NAT to allow support of overlapping IP address space. However, NAT may still be required in order to send traffic to the Internet.
6. Extends PE routers. A Multi-VRF router could use 5 different OSPF processes to connect to 5 different customers in the same site, and then use BGP to propagate the routes to the PE router.

#### Multi-VRF CE Requirements

##### Supported Platforms

The following is a list of router platforms supported at the provider edge.

- Cisco 1700 Series as of 12.2(8)YN and higher
- Cisco 2600 series
- Cisco 3600 series
- Cisco 3700 Series as of 12.2(8)T and higher
- Cisco 7200 series
- Cisco 7500 series

##### Minimum IOS Required

Multi-VRF CE is introduced in Cisco IOS release 12.2(4)T.

Note: A PLUS feature set is required for this feature.

##### Memory Requirements

###### Flash

- The feature set chosen determines the amount of flash needed.
- Check CCO (<http://www.cisco.com>) for the minimum flash requirements per each IOS release

###### DRAM

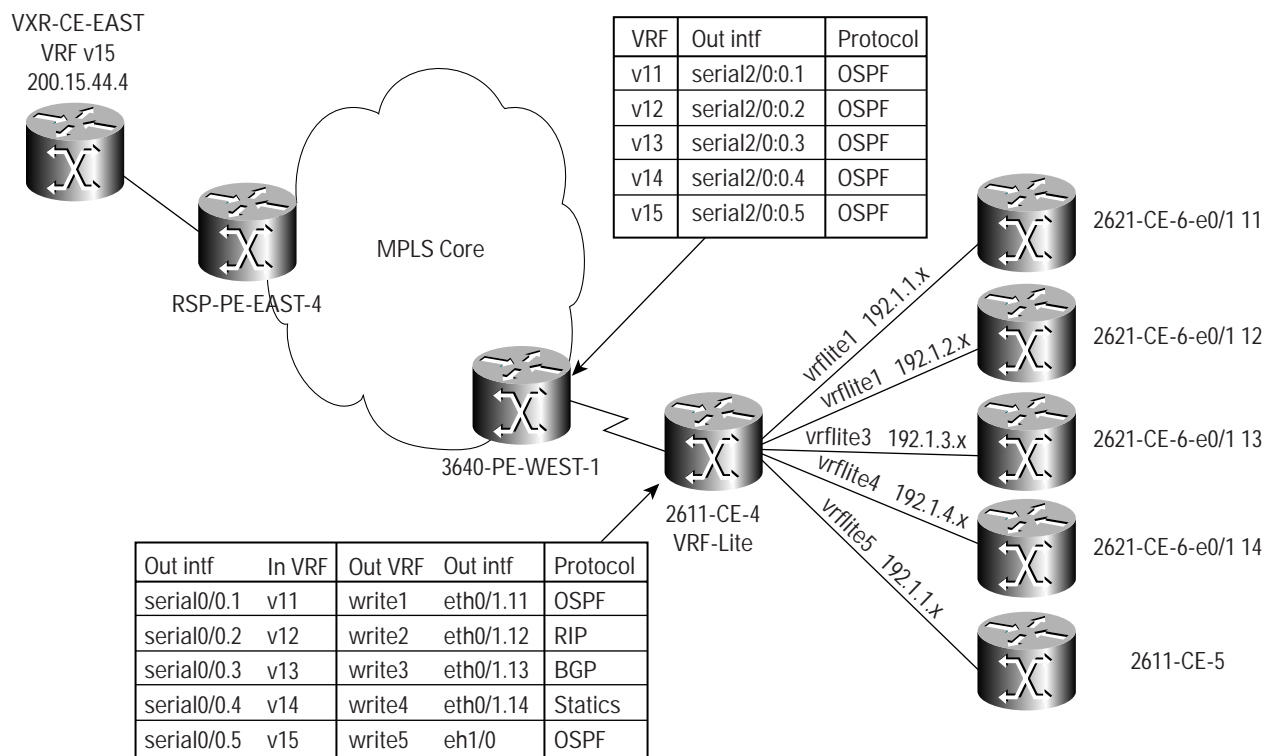
Maximum DRAM memory for each platform is recommended.



## Case Study: Multi-VRF CE

Assume a single service provider has an IP backbone to deliver MPLS-VPN Services to different enterprises. There are 2 PE routers in the network and one 2611 configured for Multi-VRF CE

Figure 8 Case Study: Multi-VRF CE Network Topology



In this test set up, the 2611-CE4 represents a typical branch office with several hosts connected to it. For example, the host 2621-CE-6-e0/1.11 could represent the HR organization; whereas, host 2621-CE-6-e0/1.12 could represent Finances. Both hosts connect to the 2611-CE-4 via the Ethernet interface but both hosts' data requires that it remain private.

VXR-CE-East represents a typical corporate office where multiple branch offices connect via the MPLS network.

RSP-PE-East-4 and 3640-PE-West-1 represent PE routers and perform all PE functionality that has been discussed in the MPLS-VPN Overview Section.

The following policies describe the desired inter-site connectivity for this case study.

- All sub-interfaces off the 2621-CE-6 can communicate with VXR-CE-East but not with each other.
- 2611-CE-5 can communicate with VXR-CE-East but not with any host off 2621-CE-6
- All Traffic off 2611-CE-4 is segmented into 5 separate VRFs (labeled vrflite1-5)

Frame Relay over T1 point-to-point sub-interfaces is used to connect 3640-PE-WEST1 and 2611-CE-4. Ethernet is used to connect 2621-CE-6 to 2611-CE-4. 802.1Q sub-interfaces comprise 4 of the 5 Ethernet connections from 2621-CE-6 to 2611-CE-4.

Duplicate IP address spaces were given to two hosts [2621-CE-6-e0/1.11 and 2611-CE-5] to show the benefit of the VRF feature with respect to IP addresses. These two connections use OSPF as the routing protocol to exchange updates with 2611-C-4. All other hosts off 2611-CE-4 use a combination of OSPF, EBGp, RIPv2 and static routes as show in the charts within Figure 8. These routes are redistributed into OSPF at 2611-CE-4. The T1 interface between 3640-PE-WEST-1 and 2611-CE-4 is segmented into 5 point-to-point Frame Relay sub-interfaces which are mapped directly to each separate VRF attached to the Ethernet hosts off 2611-CE-4. OSPF is the routing protocol for each of the Frame Relay links. BGP is redistributed into OSPF on the 3640-PE-WEST-1 to allow routes to propagate from the separate Ethernet hosts. This route redistribution allows a certain VRF off of 2611-CE-4 to have a connection to a remote host across the MPLS core on the same VRF.

## Configurations

The configurations for 3640-PE-WEST-1, 2611-CE-4, 2621-CE-6, and 2611-CE-5 are listed below. The configurations for RSP-PE-EAST-1 and VXR-CE-EAST are not shown, as they are not pertinent to the Multi-VRF CE configuration. You can find sample configurations for these routers at <http://www.cisco.com/go/mpls>

### 3640-PE-WEST-1

```
3640-PE-WEST-1#sh run
hostname 3640-PE-WEST-1
ip subnet-zero
!
ip vrf v11
  rd 11:1
  route-target export 11:1
  route-target import 11:1
!
ip vrf v12
  rd 12:1
  route-target export 12:1
  route-target import 12:1
!
ip vrf v13
  rd 13:1
  route-target export 13:1
  route-target import 13:1
!
ip vrf v14
  rd 14:1
  route-target export 14:1
  route-target import 14:1
!
ip vrf v15
  rd 15:1
  route-target export 15:1
  route-target import 15:1
!
ip cef
!
controller T1 2/0
  framing esf
  linecode b8zs
  channel-group 0 timeslots 1-24
!
interface Loopback0
  description Router ID
  ip address 10.13.1.65 255.255.255.255
```



```
!  
interface FastEthernet1/0  
  description FE to GSR-P-CENTRAL-A - 4.16  
  ip address 10.13.4.18 255.255.255.252  
  duplex auto  
  speed auto  
!  
interface Serial2/0:0  
  description T1 connection to CE - VRF_Lite  
  no ip address  
  encapsulation frame-relay  
!  
interface Serial2/0:0.1 point-to-point  
  description PE to VRF_Lite CE connection 1  
  ip vrf forwarding v11  
  ip address 220.1.65.5 255.255.255.252  
  frame-relay interface-dlci 21  
router bgp 1  
  no synchronization  
  no bgp default ipv4-unicast  
  bgp log-neighbor-changes  
  neighbor 10.13.1.48 remote-as 1  
  neighbor 10.13.1.48 update-source Loopback0  
  neighbor 10.13.1.48 activate  
  neighbor 10.13.1.61 remote-as 1  
  neighbor 10.13.1.61 update-source Loopback0  
  neighbor 10.13.1.61 activate  
  no auto-summary  
  !  
  address-family ipv4 vrf v15  
  redistribute ospf 15  
  default-metric 10  
  no auto-summary  
  no synchronization  
  exit-address-family  
  !  
  address-family ipv4 vrf v14  
  redistribute ospf 14 match internal external 1 external 2  
  default-metric 10  
  no auto-summary  
  no synchronization  
  exit-address-family  
  !  
  address-family ipv4 vrf v13  
  redistribute ospf 13 match internal external 1 external 2  
  default-metric 10  
  no auto-summary  
  no synchronization  
  exit-address-family  
  !  
  address-family ipv4 vrf v12  
  redistribute ospf 12 match internal external 1 external 2  
  default-metric 10  
  no auto-summary  
  no synchronization  
  exit-address-family  
  !  
  address-family ipv4 vrf v11  
  redistribute ospf 11  
  default-metric 10
```

```

no auto-summary
no synchronization
exit-address-family
!
address-family vpnv4
neighbor 10.13.1.48 activate
neighbor 10.13.1.48 send-community extended
neighbor 10.13.1.61 activate
neighbor 10.13.1.61 send-community extended
no auto-summary
exit-address-family
!
ip classless
no ip http server
ntp clock-period 17179973
end

```

### 2621-CE-4

```

2611-CE-4#sh run
hostname 2611-CE-4
!
ip vrf vrflitel
rd 81:81
route-target export 81:81
route-target import 81:81
!
ip vrf vrflite2
rd 82:82
route-target export 82:82
route-target import 82:82
!
ip vrf vrflite3
rd 83:83
route-target export 83:83
route-target import 83:83
!
ip vrf vrflite4
rd 84:84
route-target export 84:84
route-target import 84:84
!
ip vrf vrflite5
rd 85:85
route-target export 85:85
route-target import 85:85
ip cef
frame-relay switching
cns event-service server
!
interface Loopback0
description Router ID
ip address 10.13.1.74 255.255.255.255
!
interface Serial0/0
description T1 connection to PE - VRF_Lite
no ip address
encapsulation frame-relay
no fair-queue

```



```
service-module t1 clock source internal
service-module t1 timeslots 1-24 speed 56
frame-relay intf-type dce
!
interface Serial0/0.1 point-to-point
  description VRF_Lite CE to PE connection 1
  ip vrf forwarding vrflite1
  ip address 220.1.65.6 255.255.255.252
  frame-relay interface-dlci 21
!
interface Serial0/0.2 point-to-point
  description VRF_Lite CE to PE connection 2
  ip vrf forwarding vrflite2
  ip address 220.1.65.10 255.255.255.252
  frame-relay interface-dlci 22
!
interface Serial0/0.3 point-to-point
  description VRF_Lite CE to PE connection 3
  ip vrf forwarding vrflite3
  ip address 220.1.65.14 255.255.255.252
  frame-relay interface-dlci 23
!
interface Serial0/0.4 point-to-point
  description VRF_Lite CE to PE connection 4
  ip vrf forwarding vrflite4
  ip address 220.1.65.18 255.255.255.252
  frame-relay interface-dlci 24
!
interface Serial0/0.5 point-to-point
  description VRF_Lite CE to PE connection 5
  ip vrf forwarding vrflite5
  ip address 220.1.65.22 255.255.255.252
  frame-relay interface-dlci 25
!
interface Ethernet0/1
  description Subinterfaces to Host CE
  no ip address
  half-duplex
!
interface Ethernet0/1.11
  description VRF_Lite CE to host 1 (dup addr)
  encapsulation dot1Q 11
  ip vrf forwarding vrflite1
  ip address 192.1.1.1 255.255.255.0
!
interface Ethernet0/1.12
  description VRF_Lite CE to host 2
  encapsulation dot1Q 12
  ip vrf forwarding vrflite2
  ip address 192.1.2.1 255.255.255.0
!
interface Ethernet0/1.13
  description VRF_Lite CE to host 3
  encapsulation dot1Q 13
  ip vrf forwarding vrflite3
  ip address 192.1.3.1 255.255.255.0
!
interface Ethernet0/1.14
  description VRF_Lite CE to host 4
  encapsulation dot1Q 14
```

```

ip vrf forwarding vrflite4
ip address 192.1.4.1 255.255.255.0
!
interface Ethernet1/0
  description VRF_Lite CE to host 5 (dup addr)
  ip vrf forwarding vrflite5
  ip address 192.1.1.1 255.255.255.0
  half-duplex
!
router ospf 11 vrf vrflitel
  log-adjacency-changes
  area 11 virtual-link 220.1.65.5
  network 192.1.1.0 0.0.0.255 area 0
  network 220.1.65.4 0.0.0.3 area 11
!
router ospf 12 vrf vrflite2
  log-adjacency-changes
  redistribute rip subnets
  network 220.1.65.8 0.0.0.3 area 12
!
router ospf 13 vrf vrflite3
  log-adjacency-changes
  redistribute bgp 13 subnets
  network 220.1.65.12 0.0.0.3 area 13
!
router ospf 14 vrf vrflite4
  log-adjacency-changes
  redistribute connected
  redistribute static subnets
  network 220.1.65.16 0.0.0.3 area 14
!
router ospf 15 vrf vrflite5
  log-adjacency-changes
  area 15 virtual-link 220.1.65.21
  network 192.1.1.0 0.0.0.255 area 0
  network 220.1.65.20 0.0.0.3 area 15
!
router rip
  version 2
  !
  address-family ipv4 vrf vrflite2
  version 2
  redistribute ospf 12
  network 192.1.2.0
  default-metric 1
  no auto-summary
  exit-address-family
!
router bgp 13
  bgp log-neighbor-changes
  !
  address-family ipv4 vrf vrflite5
  no auto-summary
  no synchronization
  exit-address-family
  !
  address-family ipv4 vrf vrflite4
  no auto-summary
  no synchronization
  exit-address-family

```



```
!  
address-family ipv4 vrf vrflite3  
redistribute ospf 13 match internal  
neighbor 192.1.3.2 remote-as 3  
neighbor 192.1.3.2 activate  
no auto-summary  
no synchronization  
network 192.1.3.0  
exit-address-family  
!  
address-family ipv4 vrf vrflite2  
no auto-summary  
no synchronization  
exit-address-family  
!  
address-family ipv4 vrf vrflite1  
no auto-summary  
no synchronization  
exit-address-family  
!  
ip classless  
ip route vrf vrflite4 4.4.4.0 255.255.255.0 192.1.4.2
```

#### 2611-CE-5

```
2611-CE-5#sh run  
hostname 2611-CE-5  
!  
ip cef  
!  
interface Loopback0  
  description Router ID  
  ip address 10.13.1.75 255.255.255.255  
!  
interface Ethernet1/0  
  description Host to VRF_Lite CE 5 (dup addr)  
  ip address 192.1.1.2 255.255.255.0  
  half-duplex  
!  
router ospf 5  
  log-adjacency-changes  
  network 192.1.1.0 0.0.0.255 area 0  
!  
ip classless
```

#### 2621-CE-6

```
hostname 2621-CE-6  
!  
memory-size iomem 30  
ip subnet-zero  
ip cef  
!  
interface Loopback0  
  description Router ID  
  ip address 10.13.1.76 255.255.255.255  
!  
interface Loopback41  
  description Host 4 loopback 1  
  ip address 4.4.4.1 255.255.255.252  
!
```

```

interface Loopback42
  description Host 4 loopback 2
  ip address 4.4.4.5 255.255.255.252
!
interface Loopback43
  description Host 4 loopback 3
  ip address 4.4.4.9 255.255.255.252
!
interface Ethernet0/1
  description Subinterfaces to Multi-VRF CE CE
  no ip address
  half-duplex
!
interface Ethernet0/1.11
  description Host to VRF_Lite CE 1 (dup addr)
  encapsulation dot1Q 11
  ip address 192.1.1.2 255.255.255.0
!
interface Ethernet0/1.12
  description Host to VRF_Lite CE 2
  encapsulation dot1Q 12
  ip address 192.1.2.2 255.255.255.0
!
interface Ethernet0/1.13
  description Host to VRF_Lite CE 3
  encapsulation dot1Q 13
  ip address 192.1.3.2 255.255.255.0
!
interface Ethernet0/1.14
  description Host to VRF_Lite CE 4
  encapsulation dot1Q 14
  ip address 192.1.4.2 255.255.255.0
!
router ospf 1
  log-adjacency-changes
  network 192.1.1.0 0.0.0.255 area 0
!
router rip
  version 2
  network 192.1.2.0
!
router bgp 3
  bgp log-neighbor-changes
  neighbor 192.1.3.1 remote-as 13
  neighbor 192.1.3.1 update-source Ethernet0/1.13
!
ip classless
ip route 220.1.65.16 255.255.255.252 192.1.4.1

```

## Network Connectivity

To verify that there is connectivity between the VXR-CE-EAST and 3640-PE-WEST-1 and 2611-CE-4, several show commands are necessary.

For more detail explanations on show commands for a MPLS network, see <http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120t/120t5/vpn.htm>



#### Show IP route vrf v15 | include 200.15.44.4

```
3640-PE-WEST-1#sh ip route vrf v15 | include 200.15.44.4
B      200.15.44.4 [200/0] via 10.13.1.44, 00:16:59
```

#### Show ip route vrf v15 200.15.44.4

```
3640-PE-WEST-1#sh ip route vrf v15 200.15.44.4
Routing entry for 200.15.44.4/30
  Known via "bgp 1", distance 200, metric 0, type internal
  Redistributing via ospf 15
  Advertised by ospf 15 subnets
  Last update from 10.13.1.44 00:17:10 ago
  Routing Descriptor Blocks:
  * 10.13.1.44 (Default-IP-Routing-Table), from 10.13.1.48, 00:17:10 ago
    Route metric is 0, traffic share count is 1
    AS Hops 0
```

#### Show ip cef vrf v15 200.15.44.4

```
3640-PE-WEST-1#sh ip cef vrf v15 200.15.44.4
200.15.44.4/30, version 283, cached adjacency 10.13.4.17
0 packets, 0 bytes
  tag information set
    local tag: VPN-route-head
    fast tag rewrite with Fa1/0, 10.13.4.17, tags imposed: {56 5899}
  via 10.13.1.44, 0 dependencies, recursive
    next hop 10.13.4.17, FastEthernet1/0 via 10.13.1.44/32
    valid cached adjacency
    tag rewrite with Fa1/0, 10.13.4.17, tags imposed: {56 5899}
```

#### Show ip route | include 200.15.44.4

```
2611-CE-5#sh ip route | include 200.15.44.4
O E2   200.15.44.4 [110/1] via 192.1.1.1, 00:16:16, Ethernet1/0
2611-CE-5#sh ip route 200.15.44.4
Routing entry for 200.15.44.4/30
  Known via "ospf 5", distance 110, metric 1
  Tag Complete, Path Length == 1, AS 1, , type extern 2, forward metric 84
  Last update from 192.1.1.1 on Ethernet1/0, 00:02:12 ago
  Routing Descriptor Blocks:
  * 192.1.1.1, from 220.1.65.21, 00:02:12 ago, via Ethernet1/0
    Route metric is 1, traffic share count is 1
```

#### Show ip cef 200.15.44.4

```
2611-CE-5#sh ip cef 200.15.44.4
200.15.44.4/30, version 406, cached adjacency 192.1.1.1
0 packets, 0 bytes
  via 192.1.1.1, Ethernet1/0, 0 dependencies
    next hop 192.1.1.1, Ethernet1/0
    valid cached adjacency
```

## Verifying Network Connectivity

To verify that a host 2611-CE-5 is connected to VXR-CE-EAST as simple ping and trace route are necessary.

A ping command to the 200.15.44.4 /30 network from our host CE (2611-CE-5):

```
2611-CE-5#ping 200.15.44.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.15.44.4, timeout is 2 seconds:
!!!!!
```

Success rate is 100 percent (5/5), round-trip min/avg/max = 4/6/8 ms

And a trace route from the destination network to the host CE shows the path taken and the tags used along the way:


```
RSP-PE-EAST-4#traceroute vrf v15 192.1.1.2
Type escape sequence to abort.
Tracing the route to 192.1.1.2

 0 10.13.6.161 [MPLS: Labels 90/117 Exp 0] 4 msec 8 msec 4 msec
 1 10.13.3.161 0 msec 4 msec 4 msec
 2 10.13.3.137 [MPLS: Labels 31/117 Exp 0] 4 msec 8 msec 4 msec
 3 10.13.2.25 [MPLS: Labels 62/117 Exp 0] 4 msec 4 msec 4 msec
 4 * * *
 5 10.13.2.6 [MPLS: Labels 19/117 Exp 0] 4 msec 4 msec 4 msec
 6 220.1.65.21 4 msec 4 msec 0 msec
 7 220.1.65.22 4 msec 4 msec 4 msec
 8 192.1.1.2 4 msec * 4 msec
```

## Application Study: Multi-Access Application

In this case study, a multi-dwelling environment discussed where one physical infrastructure can server several type customers. This application can be used to segment Internet and Intranet traffic in a either a local branch office or a multi-dwelling unit where the CE router is the entry and exit point to a Service Provider's network.

Figure 9 Multi-access Application Model





In Figure 9, only one Fast Ethernet is configured on the Host CE router, which is segmented into two sub-interfaces. ATM or Frame Relay can be used to connect the Host CE router to the PE router. The PE router then connects to the MPLS core as stated in the MPLS-VPN overview. RIPv2, OSPF and static routing can be used as a routing protocol between the Host CE router and the PE router.

Segmentation occurs in two places in this design:

1. Separate Sub-interfaces either via Frame Relay or ATM.

The traffic from the Host CE router to the PE router is separated via separate sub-interfaces so that data can maintain security. Each sub-interface is associated with its own VRF, which is forwarded to the PE router

2. Fast Ethernet Sub-interfaces

The traffic from the local LAN is separated into separate sub-interfaces. There is no need for a switch as each sub-interface is designated its own VRF.

In this application, one Fast Ethernet sub-interface is segmented for DMZ traffic where users are allowed an Internet connection. NAT is used to keep the local LAN's ip address space private while connecting to the Internet. One Fast Ethernet sub-interface is segmented for customer Intranet traffic. This sub-interface is to access only a company's Intranet; thereby, controlling who accesses the privately maintained network. The local LAN ip address space is also kept private without the use of NAT.

This type of solution is a combined Internet and Intranet service offering that is a more elegant and easier to configure and maintain solution than policy routing or the use of a switch.

#### Sample Configuration for Multi-access application for VRF

```
!
ip vrf NMS
rd 100100
route-target both 100100
!
interface FastEthernet 0/0.1
encapsulation isl 1
ip address 172.19.0.1 255.255.192.0
ip nat inside
!
interface Fast Ethernet 0/0.2
encapsulation isl 2
ip vrf forwarding NMS
ip address 172.32.1.1 255.255.0.0
!
interface serial 0/0
encapsulation frame-relay
!
interface serial 0/0.1 point-to-point
ip address 172.198.254.2 255.255.255.252
ip nat inside
frame-relay interface-dlci 100
!
interface serial 0/0.2 point-to-point
ip vrf forwarding NMS
ip address 10.1.100.2 255.255.255.0
frame-relay interface-dlci 101
```

This configuration shows only one VRF but more can be easily created for further segmentation. All design restrictions apply.

#### Under Investigation

While the Multi-VRF CE feature creates privacy and confidentiality by allowing the CE router to segment data from hosts on the traditional LAN network into separate VRFs, this data is not encrypted. IPSec over each VRF is a growing concern for many customers to add security to their already private network; however, IPSec with the Multi-VRF CE feature is not supported currently.

This functionality is currently being investigated.

## Conclusions

In order to ensure that their data is kept private while traveling across a Service Provider's network, customers are presented many VPN options to suit their needs. This paper has focused on one particular type of VPNs: MPLS-VPNs. A general description was outlined for MPLS-VPNs in order to discuss the new feature in Cisco IOS release 12.2: Multi-VRF CE. Multi-VRF CE extends limited PE functionality to CE devices by allowing the traditional LAN network behind a CE router to be segmented into separate VRFs. With this feature, the CE router is now able to segment their LAN traffic into a maximum of 5 separate VRFs. Also, the Cisco 2600 series now has the ability to provide limited PE functionality which overall extends the MPLS-VPN network to the branch office.

## Additional Resources

MPLS Virtual Private Networks

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120t/120t5/vpn.htm>

Cisco IOS MPLS

<http://www.cisco.com/go/mpls>



Corporate Headquarters  
Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
www.cisco.com  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 526-4100

European Headquarters  
Cisco Systems Europe  
11, Rue Camille Desmoulins  
92782 Issy-les-Moulineaux  
Cedex 9  
France  
www-europe.cisco.com  
Tel: 33 1 58 04 60 00  
Fax: 33 1 58 04 61 00

Americas Headquarters  
Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
www.cisco.com  
Tel: 408 526-7660  
Fax: 408 527-0883

Asia Pacific Headquarters  
Cisco Systems, Inc.  
Capital Tower  
168 Robinson Road  
#22-01 to #29-01  
Singapore 068912  
www.cisco.com  
Tel: 65 317 7777  
Fax: 65 317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the

**Cisco Web site at [www.cisco.com/go/offices](http://www.cisco.com/go/offices)**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia  
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia  
Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru  
Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa  
Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

All contents are Copyright © 1992–2001 Cisco Systems, Inc. All rights reserved. AccessPath, AtmDirector, Browse with Me, CCIP, CCSI, CD-PAC, *CiscoLink*, the Cisco *Powered* Network logo, Cisco Systems Networking Academy, the Cisco Systems Networking Academy logo, Cisco Unity, Fast Step, Follow Me Browsing, FormShare, FrameShare, IGX, Internet Quotient, IP/VC, iQ Breakthrough, iQ Expertise, iQ FastTrack, the iQ logo, iQ Net Readiness Scorecard, MGX, the Networkers logo, ScriptBuilder, ScriptShare, SMARTnet, TransPath, Voice LAN, Wavelength Router, and WebViewer are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and Discover All That's Possible are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, FastHub, FastSwitch, GigaStack, IOS, IP/TV, LightStream, MICA, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, SlideCast, StrataView Plus, Stratm, SwitchProbe, TeleRouter, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0110R)

Printed in the USA