Internet Access from an MPLS VPN Using a Global Routing Table

Document ID: 24508

Contents

Introduction
Prerequisites
Requirements
Components Used
Background Theory
Conventions
Configure
Network Diagram
Configurations
Verify
VPN Connectivity Between CE 1 and CE 2
Connectivity to the Internet from CE 1
Troubleshoot
Related Information

Introduction

The purpose of this document is to demonstrate the sample configuration used to access the Internet from a Multiprotocol Label Switching (MPLS)−based VPN using a global routing table.

In certain network scenarios, it is required to access the Internet from an MPLS−based VPN in addition to continuing to maintain the VPN connectivity among corporate sites. This sample configuration focuses on providing Internet access from the VPN routing and forwarding (VRF) that contains the default route to the Internet gateway router (IGW).

Prerequisites

Requirements

A basic understanding of MPLS forwarding and MPLS VPN is required to fully understand the contents of this document.

Components Used

The information in this document is based on the software and hardware versions below.

- Cisco IOS® Software Release 12.1(3)T. Release 12.0(5)T includes the MPLS VPN feature
- Any Cisco router from the 3600 series or later, such as the Cisco 3660 or 7206

The information presented in this document was created from devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If you are working in a live network, ensure that you understand the potential impact of any command before using it.

Background Theory

In this example configuration, these policies were in place:

- A router with connectivity to the Internet is attached to the MPLS network. It may or may not inject Border Gateway Protocol (BGP) routes into the global routing table.

  Note: PE routers understand BGP. Routers such as the Gigabit Switch Router (GSR) (which performs as a Provider Core router) do not run BGP at all.
- There is no requirement for a VRF to have a full routing table from the Internet (global BGP table), so a static default route is put in a VRF pointing to the global next hop address of the IGW.
- A VPN customer uses a registered unique address range that is routable in the global Internet routing table. The method of access discussed in this document is not recommended where customers have only private addresses in their network.

Conventions

These acronyms are used in this document:

- CE − Customer Edge router
- PE − Provider Edge router
- P − Provider core router

For more information on document conventions, refer to Cisco Technical Tips Conventions.

Configure

- You can refer to the Network Diagram for an illustration of this configuration. In this example, CE 1 and CE 2 are in the same VPN. They are configured under the customer1 VRF, since there is no requirement for a VRF to have a full routing table from the Internet (as per the policies in the Background Theory section of this document).
- A static default route is configured in the customer1 VRF on CE 1 pointing to the IGW. By placing a static default route within the customer1 VRF, packets that do not match any of the routes contained within customer1 VRF will be sent to the IGW.

Note: Since the Internet gateway next hop 192.168.67.1 is not a part of the customer1 VRF, a default route is configured under the customer1 VRF pointing to the Internet gateway interface s8/0 IP 192.168.67.1. The route to 192.168.67.1 does not lie within customer1 VRF, so you need to have a global keyword within the static default route configured under customer1 VRF. The global keyword specifies that the next hop address of the static route is resolved within the global routing table, not within the the customer1 VRF.

The following is an example of the static route.

```
ip route vrf customer1 0.0.0.0 0.0.0.0 192.168.67.1 global
```

Having a static route with a global keyword in the customer1 VRF ensures that all packets destined to the Internet are routed to the Internet gateway and subsequently to the Internet.

Note: The default route in PE 1 is configured to point to the serial interface IP address of the Internet gateway (192.168.67.1) and not to the loopback address (10.1.1.6). This avoids blackholing the routes in the event of connectivity failure between the Internet gateway and the Internet (R7). If the default route is pointed to the loopback address of the Internet gateway and the connectivity between the Internet gateway−R7 breaks, all the packets would continue to route to the Internet gateway. This happens because the loopback address

remains up (unlike 192.168.67.1 which is withdrawn from the global routing table when interface s8/0 goes down) and the default route continues to exist in the routing table.

The next step is to ensure that packets coming back from the Internet to destination CE 1 network 11.11.11.0/24, are routed from the Internet gateway to PE 1 and to CE 1 through the MPLS core. This is achieved by configuring a static route for the CE 1 network pointing to the Serial 8/0 interface in the global routing table on PE 1. Redistribute it into the Open Shortest Path First (OSPF) so that the Internet gateway has that route in its global routing table. This allows the Internet gateway to route all packets coming from the Internet to PE 1, and to the final destination beyond CE 1.

The following example is the **ip route** command used in configuration on PE 1.

```
ip route 11.11.11.0 255.255.255.0 Serial8/0 192.168.10.1
```

Note: The above static route configured in the global routing table is in addition to the static route configured within the customer1 VRF, which is used for VPN Network Layer Reachability Information (NLRI). On PE 1, it is configured as shown as below.

```
ip route vrf customer1 11.11.11.0 255.255.255.0 192.168.10.1
```

Note: To find additional information on the commands used in this document, use the Command Lookup Tool (registered customers only) .

Network Diagram

This document uses the network setup shown in the diagram below.

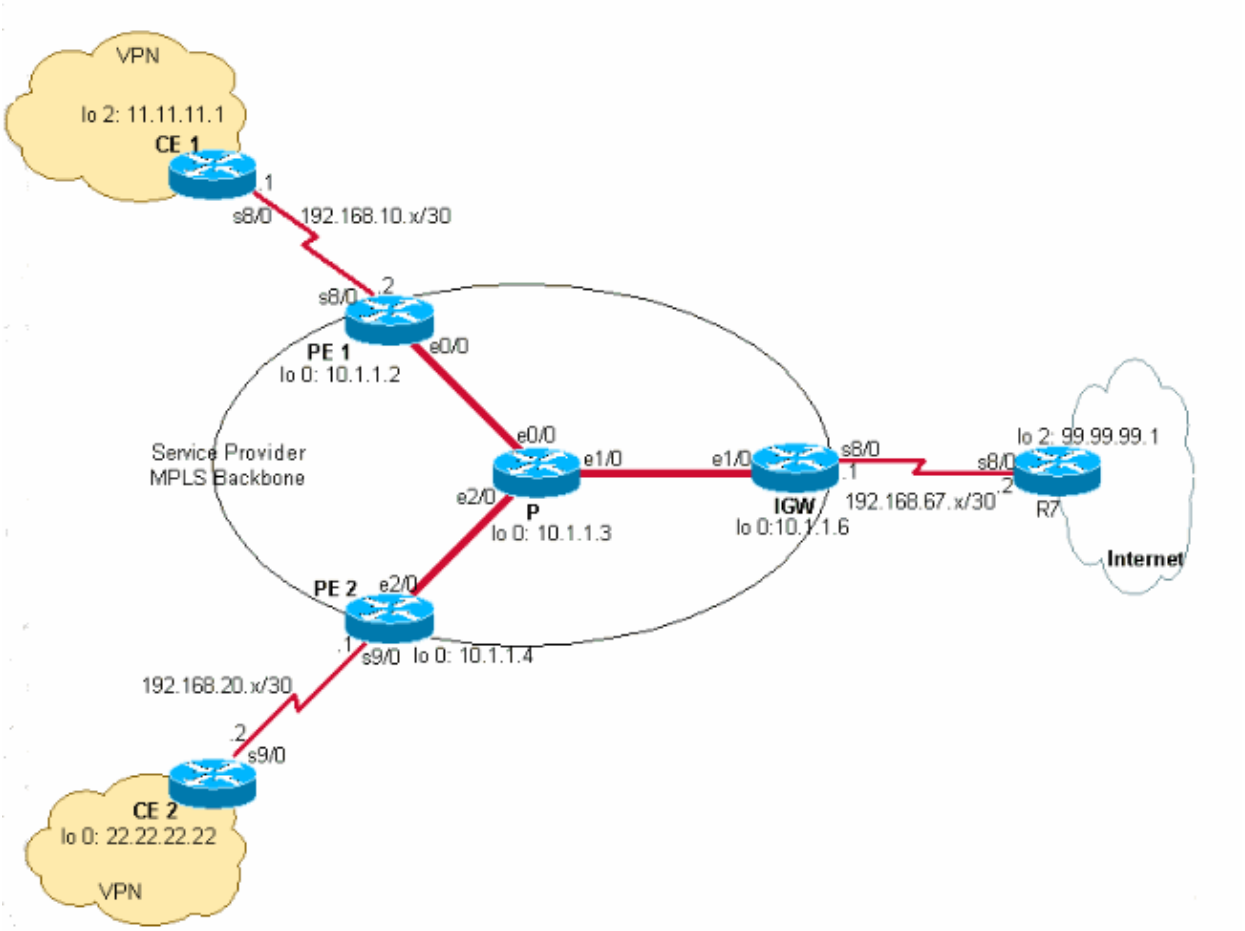

Configurations

This document uses the configurations shown below.

- CE 1
- PE 1
- P
- IGW
- PE 2
- CE 2

CE 1

```
version 12.2
!
hostname CE-1
!
ip subnet-zero
!
interface Loopback0
 ip address 10.1.1.1 255.255.255.255
!
interface Loopback2
 ip address 11.11.11.1 255.255.255.0
!
interface Serial8/0
 ip address 192.168.10.1 255.255.255.252

 !--- The interface is connected to PE 1.

 !
ip classless
ip route 0.0.0.0 0.0.0.0 192.168.10.2

!--- This is the default route to route all packets to PE 1.

!
```

PE 1

```
version 12.2
!
hostname PE-1
!
ip subnet-zero
!
ip vrf customer1

!--- This configured VRF customer1.

 rd 100:1

!--- This configured the route distiguisher for VRF.

 route-target export 1:1
 route-target import 1:1

!--- This configured the export and import policies into VRF.

!
ip cef

!--- This enabled Cisco Express Forwarding (CEF) switching.
```

```
!
interface Loopback0
 ip address 10.1.1.2 255.255.255.255
!
interface Ethernet0/0

!--- It is connected to P router.

 ip address 10.10.23.2 255.255.255.0
 tag-switching ip

!--- MPLS switching is enabled.

!
interface Serial8/0
! Connected to CE-1
 ip vrf forwarding customer1

!--- Route forwarding based on customer1 VRF is enabled.

 ip address 192.168.10.2 255.255.255.252
!
router ospf 1
 log-adjacency-changes
 redistribute static subnets
 network 0.0.0.0 255.255.255.255 area 0
!
router bgp 100
 no synchronization
 bgp log-neighbor-changes
 neighbor 10.1.1.4 remote-as 100

!--- Neighbor relationship with PE 2 is established.

 neighbor 10.1.1.4 update-source Loopback0
 neighbor 10.1.1.4 next-hop-self
 no auto-summary
!
 address-family ipv4 vrf customer1

!--- The address-family configuration mode specifies IPv4 unicast
!---address prefixes for customer1 VRF.

 no auto-summary
 no synchronization
 network 11.11.11.0 mask 255.255.255.0

!--- CE 1 network 11.11.11.0/24 to PE 2 is announced.

network 192.168.10.0 mask 255.255.255.252
 exit-address-family
!
address-family vpnv4

!--- This is the address-family VPNV4 configuration mode for
!--- configuring BGP sessions.

 neighbor 10.1.1.4 activate
 neighbor 10.1.1.4 send-community extended
 no auto-summary
 exit-address-family
!
ip classless
ip route 11.11.11.0 255.255.255.0 Serial8/0 192.168.10.1
```

```
!--- The static route in the global routing table is pointing to
!--- the interface connected to CE 1.

ip route vrf customer1 0.0.0.0 0.0.0.0 192.168.67.1 global

!--- The static default route under customer1 VRF, routing packets
!--- outside of VPN to the Internet gateway.

! routes
ip route vrf customer1 11.11.11.0 255.255.255.0 192.168.10.1

!--- The static route for network 11.11.11.0/24 (CE-1 Network) under
!---customer1 VRF ensures the reachability of CE 1 network from the
!--- other VPN sites.
```

P

```
version 12.2
!
hostname P
!
ip subnet-zero
!
ip cef

!--- CEF switching is enabled.

!
interface Loopback0
 ip address 10.1.1.3 255.255.255.255
!
interface Ethernet0/0

!--- This is connected to PE 1.

 ip address 10.10.23.3 255.255.255.0
 tag-switching ip

!--- MPLS switching is enabled.

!
interface Ethernet1/0

!--- This is connected to PE 2.

 ip address 10.10.34.3 255.255.255.0
 tag-switching ip
!
interface Ethernet2/0

!--- This is connected to the Internet gateway.

 ip address 10.10.36.3 255.255.255.0
 tag-switching ip
!
router ospf 1
 log-adjacency-changes
network 0.0.0.0 255.255.255.255 area 0
```

IGW

```
version 12.2
!
hostname IGW
```

```
!
ip subnet-zero
!
ip cef

!--- This enabled CEF switching.

!
interface Loopback0
 ip address 10.1.1.6 255.255.255.255
!
interface Ethernet2/0

!--- This is connected to P router.

 ip address 10.10.36.6 255.255.255.0
tag-switching ip
!
interface Serial8/0

!--- This is connected to Internet R7.

 ip address 192.168.67.1 255.255.255.252
!
router ospf 1
 log-adjacency-changes
 network 0.0.0.0 255.255.255.255 area 0
 !
 router bgp 100
 no synchronization
 bgp log-neighbor-changes
 network 11.11.11.0 mask 255.255.255.0
 network 22.22.22.0 mask 255.255.255.0
 neighbor 192.168.67.2 remote-as 200
 no auto-summary
```

PE 2

```
version 12.2
!
hostname PE-2
!
ip subnet-zero
!
ip vrf customer1

!--- Customer1 VRF is configured.

 rd 100:1

!--- Route Distinguisher for VRF is configured.

 route-target export 1:1
 route-target import 1:1

!--- This configured the import and export policies for customer1
!--- VRF.

!
ip cef

!--- This enabled CEF switching.

!
interface Loopback0
```

```
 ip address 10.1.1.4 255.255.255.255
interface Ethernet1/0

!--- Connected to P router.

 ip address 10.10.34.4 255.255.255.0
 tag-switching ip

!--- MPLS switching is enabled.

!
interface Serial9/0

!--- Connected to CE 2 router.

ip vrf forwarding customer1

!--- This enables VRF forwarding on the interface.

 ip address 192.168.20.1 255.255.255.252
!
router ospf 1
 log-adjacency-changes
 redistribute static subnets
 network 0.0.0.0 255.255.255.255 area 0
!
router bgp 100
 no synchronization
 bgp log-neighbor-changes
 neighbor 10.1.1.2 remote-as 100
 neighbor 10.1.1.2 update-source Loopback0
 neighbor 10.1.1.2 next-hop-self
no auto-summary
!
 address-family ipv4 vrf customer1

!--- This is the address-family IPv4 configuration of customer1 VRF.

 no auto-summary
 no synchronization
 network 22.22.22.0 mask 255.255.255.0

!--- This announces the CE 2 network to PE 1.

 exit-address-family
!
 address-family vpnv4

!--- This is the address-family VPNV4 configuration for BGP Sessions
!--- with PE 1.

 neighbor 10.1.1.2 activate
 neighbor 10.1.1.2 send-community extended
 no auto-summary
 exit-address-family
!
ip classless
ip route 22.22.22.0 255.255.255.0 Serial9/0 192.168.20.2

!--- This is the static route for network 22.22.22.0/24 in the global
!--- routing table pointing to the interface connected to CE 2.

ip route vrf customer1 0.0.0.0 0.0.0.0 192.168.67.1 global

!--- This is the static default route for customer VRF
!--- for destinations outside the VPN.
```

```
ip route vrf customer1 22.22.22.0 255.255.255.0 192.168.20.2

!--- This is the static route within customer1 VRF for CE 2
!--- network for VPN connectivity.
```

CE 2

```
version 12.2
!
hostname CE-2
!
ip subnet-zero
!
interface Loopback0
 ip address 22.22.22.22 255.255.255.0
!
interface Serial9/0

!--- This is connected to PE 2.

 ip address 192.168.20.2 255.255.255.252
!
ip classless
ip route 0.0.0.0 0.0.0.0 192.168.20.1

!--- This is the default route pointing to PE 2.
```

Verify

This section provides information you can use to confirm your configuration is working properly.

VPN Connectivity Between CE 1 and CE 2

To verify the VPN connectivity between CE 1 and CE 2, CE 1 should be able to reach CE 2's network 22.22.22.0/24 and the other way around. To check this, verify the route to network 22.22.22.0/24 in the customer1 VRF at PE 1.

Certain **show** commands are supported by the Output Interpreter Tool (registered customers only) , which allows you to view an analysis of **show** command output.

1. The **show ip route vrf customer1** command confirms the route to network 22.22.22.0/24 learned from 10.1.1.4 (PE 2's loopback address) shown highlighted in the output below.

```
PE-1# show ip route vrf customer1
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is 192.168.67.1 to network 0.0.0.0

192.168.10.0/30 is subnetted, 1 subnets
C       192.168.10.0 is directly connected, Serial8/0
       22.0.0.0/24 is subnetted, 1 subnets
```

```
B       22.22.22.0 [200/0] via 10.1.1.4, 01:00:50
       11.0.0.0/24 is subnetted, 1 subnets
S      11.11.11.0 [1/0] via 192.168.10.1
S*   0.0.0.0/0 [1/0] via 192.168.67.1
```

2. Similarily, at PE 2, the route to network 11.11.11.0/24 in the customer1 VRF is shown in the example below.

```
PE−2# show ip route vrf customer1
Codes: C − connected, S − static, I − IGRP, R − RIP, M − mobile, B − BGP
       D − EIGRP, EX − EIGRP external, O − OSPF, IA − OSPF inter area
       N1 − OSPF NSSA external type 1, N2 − OSPF NSSA external type 2
       E1 − OSPF external type 1, E2 − OSPF external type 2, E − EGP
       i − IS−IS, L1 − IS−IS level−1, L2 − IS−IS level−2, ia − IS−IS inter area
       * − candidate default, U − per−user static route, o − ODR
       P − periodic downloaded static route

Gateway of last resort is 192.168.67.1 to network 0.0.0.0

192.168.10.0/30 is subnetted, 1 subnets
B       192.168.10.0 [200/0] via 10.1.1.2, 01:00:09
     22.0.0.0/24 is subnetted, 1 subnets
S       22.22.22.0 [1/0] via 192.168.20.2
     192.168.20.0/30 is subnetted, 1 subnets
C       192.168.20.0 is directly connected, Serial9/0
     11.0.0.0/24 is subnetted, 1 subnets
B       11.11.11.0 [200/0] via 10.1.1.2, 01:00:09
S*   0.0.0.0/0 [1/0] via 192.168.67.1
```

3. Now check the connectivity between CE 1 and CE 2 by pinging a host 22.22.22.22 on CE 2 using the source IP address of 11.11.11.1 from CE 1.

```
CE−1# ping
Protocol [ip]:
Target IP address: 22.22.22.22
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 11.11.11.1
Type of service [0]:
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 5, 100−byte ICMP Echos to 22.22.22.22, timeout is 2 seconds:
!!!!!

Success rate is 100 percent (5/5), round−trip min/avg/max = 20/20/20 ms
```

Connectivity to the Internet from CE 1

Follow the steps below to verify connectivity to the Internet from CE1.

1. All packets destined to the Internet or VPN from CE 1 will route using a default route configured in CE 1 pointing to PE 1, as shown below.

```
CE−1# show ip route 0.0.0.0
Routing entry for 0.0.0.0/0, supernet
  Known via "static", distance 1, metric 0, candidate default path
  Routing Descriptor Blocks:
  * 192.168.10.2
Route metric is 0, traffic share count is 1
```

2. Packets coming into PE 1 interface s8/0 get routed using the customer1 VRF routing table. PE 1 has a default route in the customer1 VRF pointing to the IGW IP address 192.168.67.1, as shown below in the output for the show ip route vrf customer1 on PE 1.

```
PE−1# show ip route vrf customer1
Codes: C − connected, S − static, I − IGRP, R − RIP, M − mobile, B − BGP
       D − EIGRP, EX − EIGRP external, O − OSPF, IA − OSPF inter area
       N1 − OSPF NSSA external type 1, N2 − OSPF NSSA external type 2
       E1 − OSPF external type 1, E2 − OSPF external type 2, E − EGP
       i − IS−IS, L1 − IS−IS level−1, L2 − IS−IS level−2, ia − IS−IS inter area
       * − candidate default, U − per−user static route, o − ODR
       P − periodic downloaded static route

Gateway of last resort is 192.168.67.1 to network 0.0.0.0

     192.168.10.0/30 is subnetted, 1 subnets
C       192.168.10.0 is directly connected, Serial8/0
     22.0.0.0/24 is subnetted, 1 subnets
B       22.22.22.0 [200/0] via 10.1.1.4, 01:21:11
     11.0.0.0/24 is subnetted, 1 subnets
S       11.11.11.0 [1/0] via 192.168.10.1
S*   0.0.0.0/0 [1/0] via 192.168.67.1
```

3. Because the default route on PE 1 is configured with a global keyword, it looks for next hop 192.168.67.1 in its global routing table and routes to the IGW, as shown below.

```
PE−1# show ip route 192.168.67.1
Routing entry for 192.168.67.0/30
  Known via "ospf 1", distance 110, metric 84, type intra area
  Last update from 10.10.23.3 on Ethernet0/0, 00:21:54 ago
  Routing Descriptor Blocks:
  * 10.10.23.3, from 10.1.1.6, 00:21:54 ago, via Ethernet0/0
 Route metric is 84, traffic share count is 1
```

4. The packets reaching IGW get routed over to the Internet based on the BGP routes it learned from R7. In this case, you can look at the BGP route learned from R7 to demonstrate the connectivity to the Internet. Shown below is the BGP route (network 99.99.99.0/24) learned from R7 in the IGW routing table.

```
IGW# show ip route 99.99.99.0
Routing entry for 99.99.99.0/24
  Known via "bgp 100", distance 20, metric 0
  Tag 200, type external
  Last update from 192.168.67.2 01:37:25 ago
  Routing Descriptor Blocks:
  * 192.168.67.2, from 192.168.67.2, 01:37:25 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
```

   The packets that originated from CE−1 get routed to the Internet.

5. For packets coming back from the Internet destined to CE 1 network 11.11.11.0/24, IGW should have a route pointing to PE 1 in its global routing table. A static route in PE 1's global routing table pointing to s8/0 interface on PE 1 connecting to CE 1 and redistributed it into OSPF is configured. This ensures that the IGW has a route in its global routing table pointing to PE 1. The static route on PE 1 and the OSPF learned route on IGW is shown below.

```
IGW# show ip route 11.11.11.0
Routing entry for 11.11.11.0/24
  Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 20
  Last update from 10.10.36.3 on Ethernet2/0, 00:34:34 ago
  Routing Descriptor Blocks:
  * 10.10.36.3, from 10.1.1.2, 00:34:34 ago, via Ethernet2/0
      Route metric is 20, traffic share count is 1
```

```
PE−1# show ip route 11.11.11.0
Routing entry for 11.11.11.0/24
  Known via "static", distance 1, metric 0
  Redistributing via ospf 1
  Advertised by ospf 1 subnets
  Routing Descriptor Blocks:
  * 192.168.10.1, via Serial8/0
      Route metric is 0, traffic share count is 1
```

6. Now check the connectivity to the Internet from CE 1 by pinging the R7 IP address 99.99.99.1 with the CE 1 source address of 11.11.11.1.

```
CE−1# ping
Protocol [ip]:
Target IP address: 99.99.99.1
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 11.11.11.1
Type of service [0]:
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 5, 100−byte ICMP Echos to 99.99.99.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round−trip min/avg/max = 20/24/32 ms
CE−1#
```

Troubleshoot

There is currently no specific troubleshooting information available for this configuration.

Related Information

- **Configuring a Basic MPLS VPN**
- **Configuring Basic MPLS Using OSPF**
- **How to Troubleshoot the MPLS VPN**
- **MPLS Troubleshooting**
- **MPLS FAQ For Beginners**
- **MPLS (Multiprotocol Label Switching) Support Page**
- **MPLS for VPNs (Multiprotocol Label Switching for VPNs) Support Page**
- **Technical Support − Cisco Systems**

Updated: Aug 10, 2005 Document ID: 24508