驗證MPLS第3層VPN轉發

目錄

簡介

本文檔介紹在MPLS第3層VPN核心網路中驗證端到端連線的過程。

必要條件

需求

思科建議您瞭解以下主題：

- 基本IP路由知識
- Cisco IOS® XE和Cisco IOS® XR命令列知識

採用元件

本文中的資訊係根據以下軟體和硬體版本：

- 使用Cisco IOS XR軟體的路由器
- 使用Cisco IOS XE軟體的路由器

本文中的資訊是根據特定實驗室環境內的裝置所建立。文中使用到的所有裝置皆從已清除（預設）的組態來啟動。如果您的網路運作中，請確保您瞭解任何指令可能造成的影響。

背景資訊

本文檔的目的是演示基本的驗證和故障排除步驟，這些步驟用於檢查透過混合使用Cisco IOS XE和

Cisco IOS XR路由器作為PE（提供商邊緣）和P（提供商）路由器的MPLS第3層VPN核心網路與BGP（邊界網關協定）互連的兩台CE（客戶邊緣）路由器之間的連線和轉發。

慣例

如需文件慣例的詳細資訊，請參閱思科技術提示慣例。

拓撲

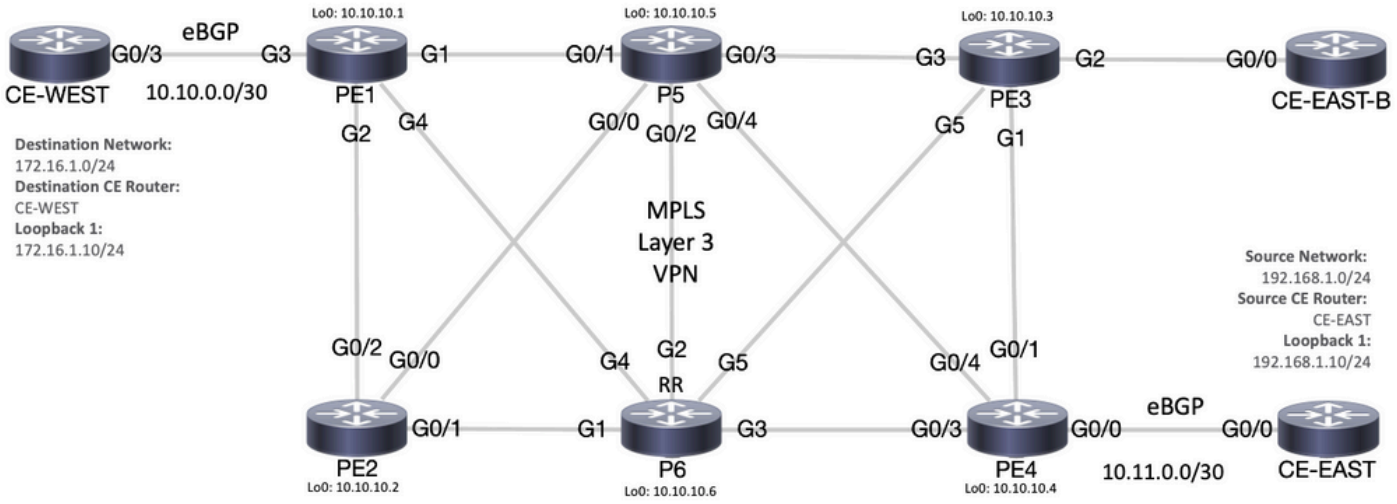

MPLS拓撲圖

疑難排解

初始資訊

源網路：192.168.1.0/24

源CE路由器：CE-EAST

目的網路：172.16.1.0/24

目的CE路由器：CE-WEST

根據初始資訊和拓撲，在路由器CE-EAST上Loopback1表示的源地址192.168.1.10和路由器CE-WEST上Loopback1表示的目的地地址172.16.1.10之間必須能夠成功連通：

<#root>

CE-EAST#

show run interface loopback1

Building configuration...

Current configuration : 66 bytes

```
!
interface Loopback1
 ip address 192.168.1.10 255.255.255.0
end

CE-WEST#
```

show run interface loopback 1

```
Building configuration...

Current configuration : 65 bytes
!
interface Loopback1
 ip address 172.16.1.10 255.255.255.0
end
```

ICMP可達性和traceroute用於開始檢查這些源地址和目標地址之間的連線，但從接下來的輸出中可以看到此操作不成功：

<#root>

CE-EAST#

ping 172.16.1.10 source loopback1

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.1.10, timeout is 2 seconds:
Packet sent with a source address of 192.168.1.10
.....
Success rate is 0 percent (0/5)

CE-EAST#
```

traceroute 172.16.1.10 source loop1 probe 1 numeric

```
Type escape sequence to abort.
Tracing the route to 172.16.1.10
VRF info: (vrf in name/id, vrf out name/id)
  1 10.11.0.2 2 msec
  2  *
  3 10.10.0.2 [MPLS: Label 16 Exp 0] 9 msec
  4  *
  5  *
  6  *
  7  *
  8  *
  9  *
 10  *
 11  *
 12  *
 13  *
 14  *
 15  *
 16  *
 17  *
 18  *
 19  *
 20  *
 21  *
```

```
 22   *
 23   *
 24   *
 25   *
 26   *
 27   *
 28   *
 29   *
 30   *
CE-EAST#
```

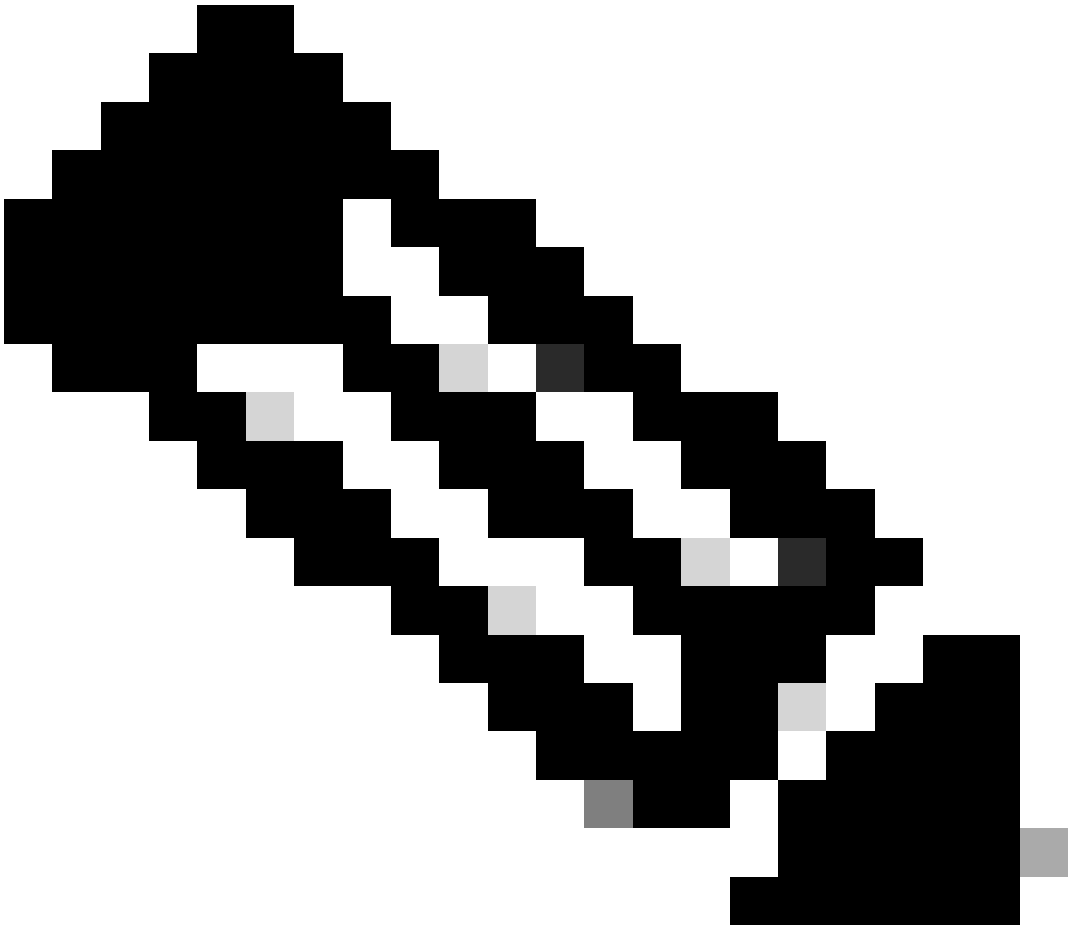

> 注意：在排除故障時，在連線到MPLS網路時使用跟蹤路由可能不太有效，因為某些服務提供商傾向於在Cisco IOS XE中配置no mpls ip propagate-ttl forward命令或在Cisco IOS XR中配置mpls ip-ttl-propagate disable forwarded命令，以隱藏核心中的所有LSR（標籤交換機路由器）（但入口和出口PE路由器除外）。

在檢視源CE路由器的狀態時，由於此路由器沒有任何VRF（虛擬路由轉發）並且不支援MPLS感知，您需要驗證RIB（路由資訊庫）、CEF（思科快速轉發）和BGP。在接下來的輸出中，可以觀察

到有一個路由條目，該條目透過BGP獲知到目標子網172.16.1.0/24，可透過介面GigabitEthernet0/0到達：

<#root>

CE-EAST#

show ip route 172.16.1.10

Routing entry for 172.16.1.0/24

Known via "bgp 65001", distance 20, metric 0

```
          <<<<<
  Tag 65500, type external
  Last update from 10.11.0.2 3d01h ago
  Routing Descriptor Blocks:
  * 10.11.0.2, from 10.11.0.2, 3d01h ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
      Route tag 65500
      MPLS label: none
```

CE-EAST#

show ip cef 172.16.1.10

172.16.1.0/24

nexthop 10.11.0.2 GigabitEthernet0/0

```
          <<<<<
CE-EAST#
```

由於源CE-EAST路由器在RIB中安裝了通往目標的路由，因此檢視提供商邊緣路由器PE4（入口PE）所花的時間如拓撲所示。此時，配置了VRF和路由區分符以及路由目標導入和導出，如下面的輸出所示：

<#root>

RP/0/0/CPU0:PE4#

show run vrf EAST

```
Mon Sep 11 20:01:54.454 UTC
vrf EAST
 address-family ipv4 unicast
```

import route-target 65000:1 65001:1 65001:2 ! export route-target 65001:1

```
  !
 !
!
```

```
RP/0/0/CPU0:PE4#
```

show run router bgp

```
Mon Sep 11 20:06:48.164 UTC
router bgp 65500
 address-family ipv4 unicast
 !
 address-family vpnv4 unicast
 !
 neighbor 10.10.10.6
  remote-as 65500
  update-source Loopback0
  address-family vpnv4 unicast
  !
 !
 vrf EAST
```

rd 65001:1

```
  address-family ipv4 unicast
  !
  neighbor 10.11.0.1
   remote-as 65001
   address-family ipv4 unicast
    route-policy PASS in
    route-policy PASS out
   !
  !
 !
!

RP/0/0/CPU0:PE4#
```

從先前的輸出中可以看到，VRF名稱「EAST」是使用65000：1的路由目標導入定義的，現在可以檢查VRF路由表，這有助於確定PE4是否具有通向目標IP地址172.16.1.10的路由：

<#root>

```
RP/0/0/CPU0:PE4#
```

show route vrf EAST 172.16.1.10

```
Mon Sep 11 19:58:28.128 UTC

Routing entry for 172.16.1.0/24
  Known via "bgp 65500", distance 200, metric 0
  Tag 65000, type internal
  Installed Sep  8 18:28:46.303 for 3d01h
  Routing Descriptor Blocks
    10.10.10.1, from 10.10.10.6
      Nexthop in Vrf: "default", Table: "default", IPv4 Unicast, Table Id: 0xe0000000
      Route metric is 0
  No advertising protos.
RP/0/0/CPU0:PE4#
```

由於此PE是思科IOS XR裝置，因此可在show route vrf <name>命令末尾使用「detail」關鍵字檢視某些其他資訊，例如由MP-BGP（多協定BGP）和源RD（路由區分器）強加的VPNv4標籤以及字首：

<#root>

RP/0/0/CPU0:PE4#

show route vrf EAST 172.16.1.10 detail

```
Mon Sep 11 20:21:48.492 UTC

Routing entry for 172.16.1.0/24
  Known via "bgp 65500", distance 200, metric 0
  Tag 65000, type internal
  Installed Sep  8 18:28:46.303 for 3d01h
  Routing Descriptor Blocks
    10.10.10.1, from 10.10.10.6
      Nexthop in Vrf: "default", Table: "default", IPv4 Unicast, Table Id: 0xe0000000
      Route metric is 0

Label: 0x10 (16)

          <<<<<
      Tunnel ID: None
      Binding Label: None
      Extended communities count: 0

Source RD attributes: 0x0000:65000:1

         <<<<<
      NHID:0x0(Ref:0)
  Route version is 0x5 (5)
  No local label
  IP Precedence: Not Set
  QoS Group ID: Not Set
  Flow-tag: Not Set
  Fwd-class: Not Set
  Route Priority: RIB_PRIORITY_RECURSIVE (12) SVD Type RIB_SVD_TYPE_REMOTE
  Download Priority 3, Download Version 36
  No advertising protos.
RP/0/0/CPU0:PE4#
```

現在，我們來看一下導入到VRF中的BGP VPNv4字首，發現這與之前輸出中的標籤16相同，並且它還具有擴展社群65000：1。此外，必須注意的是，10.10.10.1是PE4為其執行路由遞迴所需的下一跳，下一個地址「來自10.10.10.6」是PE4用於學習此字首的BGP對等體（在此場景中是路由反射器P6）：

<#root>

RP/0/0/CPU0:PE4#

show bgp vpnv4 unicast vrf EAST 172.16.1.10

```
Mon Sep 11 22:42:28.114 UTC
BGP routing table entry for 172.16.1.0/24, Route Distinguisher: 65001:1
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                48          48
Last Modified: Sep  8 18:28:46.314 for 3d04h
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  65000
```

10.10.10.1 (metric 20) from 10.10.10.6 (10.10.10.1)

<<<<<

Received Label 16

```
      Origin IGP, metric 0, localpref 100, valid, internal, best, group-best, import-candidate, importe
      Received Path ID 0, Local Path ID 0, version 48
```

Extended community: RT:65000:1

<<<<<

```
      Originator: 10.10.10.1, Cluster list: 10.10.10.6
```

Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 65000:1

<<<<<

透過在VRF級別使用exact-route關鍵字檢視CEF，您可以瞭解資料包的送出介面。此命令還可以提供一些重要的詳細資訊，因為它顯示強加到字首24001和16的兩個標籤，原因是標籤16來自BGP VPNv4，標籤24001來自LDP（標籤分發協定）：

<#root>

RP/0/0/CPU0:PE4#

show cef vrf EAST exact-route 192.168.1.10 172.16.1.10

```
Mon Sep 11 22:48:15.241 UTC
172.16.1.0/24, version 36, internal 0x5000001 0x0 (ptr 0xa12dc74c) [1], 0x0 (0x0), 0x208 (0xa155b1b8)
 Updated Sep  8 18:28:46.323
 local adjacency 10.0.0.16
 Prefix Len 24, traffic index 0, precedence n/a, priority 3
   via GigabitEthernet0/0/0/4
   via 10.10.10.1/32, 3 dependencies, recursive [flags 0x6000]
    path-idx 0 NHID 0x0 [0xa15c3f54 0x0]
    recursion-via-/32
    next hop VRF - 'default', table - 0xe0000000
    next hop 10.10.10.1/32 via 24010/0/21
```

next hop 10.0.0.16/32 Gi0/0/0/4 labels imposed {24001 16}

<<<<<

下一步，show bgp vpnv4 unicast命令用於檢查此PE正在獲取的VPNv4路由。此輸出顯示VPNv4字首導入VRF之前的資訊，請記住，配置的RT（路由目標）(例如，導入的RT是65000：1、65001：1和65001：2)指示要導入的VRF的路由和目標：

<#root>

RP/0/0/CPU0:PE4#

show bgp vpnv4 unicast

```
Fri Sep 15 02:15:15.463 UTC
BGP router identifier 10.10.10.4, local AS number 65500
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0   RD version: 0
BGP main routing table version 85
BGP NSR Initial initsync version 1 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            Next Hop            Metric LocPrf Weight Path
```

Route Distinguisher: 65000:1

***>i172.16.1.0/24 10.10.10.1 0 100 0 65000 i**

```
         <<<<<
*>i172.16.2.0/24      10.10.10.1               0    100      0 65000 i
Route Distinguisher: 65001:1 (default for vrf EAST)
* i0.0.0.0/0          10.10.10.3               0    100      0 65001 i
*>                    10.11.0.1                0             0 65001 i
*>i172.16.1.0/24      10.10.10.1               0    100      0 65000 i
*>i172.16.2.0/24      10.10.10.1               0    100      0 65000 i
*> 192.168.1.0/24     10.11.0.1                0             0 65001 i
*>i192.168.2.0/24     10.10.10.3               0    100      0 65001 i
*> 192.168.3.0/24     10.11.0.1                0             0 65001 i
Route Distinguisher: 65001:2
*>i0.0.0.0/0          10.10.10.3               0    100      0 65001 i
*>i192.168.2.0/24     10.10.10.3               0    100      0 65001 i

Processed 10 prefixes, 11 paths
```

在本示例中，VPNv4表可以很小，但在生產環境中，您可以使用下一個命令將驗證範圍縮小到特定RD和字首，而不是檢視所有VPNv4字首：

<#root>

RP/0/0/CPU0:PE4#

show bgp vpnv4 unicast rd 65000:1 172.16.1.10

```
Mon Sep 11 22:54:04.967 UTC
BGP routing table entry for 172.16.1.0/24, Route Distinguisher: 65000:1
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                46          46
Last Modified: Sep  8 18:28:46.314 for 3d04h
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  65000
    10.10.10.1 (metric 20) from 10.10.10.6 (10.10.10.1)
      Received Label 16
      Origin IGP, metric 0, localpref 100, valid, internal, best, group-best, import-candidate, not-in-v
      Received Path ID 0, Local Path ID 0, version 46
      Extended community: RT:65000:1
      Originator: 10.10.10.1, Cluster list: 10.10.10.6
```

此時，MP-BGP控制平面分別具有目標字首和LDP和VPNv4標籤{24001 16}，此流量的送出介面似乎是Gi0/0/0/4，需要轉發流量的下一跳是10.10.10.1。但是，是否有其他選項檢驗首選送出介面？現在是時候瞭解MPLS轉發表或LFIB（標籤轉發資訊庫）了。使用命令show mpls forwarding可顯示指向10.10.10.1目標（Pe1的Loopback0）的兩個條目，一條路徑具有傳出介面Gi0/0/0/4和下一跳10.0.0.16（路由器P5），其中強加的傳出標籤為24001，另一條路徑透過Gi0/0/0/3，其中下一跳為10.0.0.13（路由器P6），傳出標籤為23：

<#root>

RP/0/0/CPU0:PE4#

show mpls forwarding

```
Mon Sep 11 23:28:33.425 UTC
Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface                    Switched
------ ----------- ------------------ ------------ --------------- ------------
24000  Unlabelled  192.168.1.0/24[V]  Gi0/0/0/0    10.11.0.1       1096
24001  Unlabelled  192.168.3.0/24[V]  Gi0/0/0/0    10.11.0.1       56056
24002  Unlabelled  0.0.0.0/0[V]       Gi0/0/0/0    10.11.0.1       0
24003  Pop         10.10.10.6/32      Gi0/0/0/3    10.0.0.13       7778512
24004  Pop         10.0.0.4/31        Gi0/0/0/3    10.0.0.13       0
24005  Pop         10.0.0.8/31        Gi0/0/0/3    10.0.0.13       0
24006  Pop         10.10.10.5/32      Gi0/0/0/4    10.0.0.16       3542574
24007  Pop         10.0.0.10/31       Gi0/0/0/3    10.0.0.13       0
       Pop         10.0.0.10/31       Gi0/0/0/4    10.0.0.16       0
24008  Pop         10.0.0.6/31        Gi0/0/0/4    10.0.0.16       0
24009  Pop         10.0.0.0/31        Gi0/0/0/4    10.0.0.16       0
```

24010 23 10.10.10.1/32 Gi0/0/0/3 10.0.0.13 22316

<<<<<

24001 10.10.10.1/32 Gi0/0/0/4 10.0.0.16 42308

<<<<<

```
24011  18          10.10.10.2/32      Gi0/0/0/3    10.0.0.13       0
       24003       10.10.10.2/32      Gi0/0/0/4    10.0.0.16       0
24012  17          10.0.0.2/31        Gi0/0/0/3    10.0.0.13       0
```

```
       24005       10.0.0.2/31        Gi0/0/0/4    10.0.0.16       0
24013  Pop         10.10.10.3/32      Gi0/0/0/1    10.0.0.20       3553900
24014  Pop         10.0.0.14/31       Gi0/0/0/1    10.0.0.20       0
       Pop         10.0.0.14/31       Gi0/0/0/4    10.0.0.16       0
24015  Pop         10.0.0.18/31       Gi0/0/0/1    10.0.0.20       0
       Pop         10.0.0.18/31       Gi0/0/0/3    10.0.0.13       0

RP/0/0/CPU0:PE4#
```

show mpls forwarding prefix 10.10.10.1/32

```
Mon Sep 11 23:30:54.685 UTC
Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface                    Switched
------ ----------- ------------------ ------------ --------------- ------------
24010  23          10.10.10.1/32      Gi0/0/0/3    10.0.0.13       3188
       24001       10.10.10.1/32      Gi0/0/0/4    10.0.0.16       6044


RP/0/0/CPU0:PE4#
```

show mpls forwarding prefix 10.10.10.1/32 detail hardware egress

```
Mon Sep 11 23:36:06.504 UTC
Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface                    Switched
------ ----------- ------------------ ------------ --------------- ------------
24010  23          10.10.10.1/32      Gi0/0/0/3    10.0.0.13       N/A
     Updated: Sep  8 20:27:26.596
     Version: 39, Priority: 3
     Label Stack (Top -> Bottom): { 23 }
     NHID: 0x0, Encap-ID: N/A, Path idx: 0, Backup path idx: 0, Weight: 0
     MAC/Encaps: 14/18, MTU: 1500
     Outgoing Interface: GigabitEthernet0/0/0/3 (ifhandle 0x000000a0)
     Packets Switched: 0

       24001       10.10.10.1/32      Gi0/0/0/4    10.0.0.16       N/A
     Updated: Sep  8 20:27:26.596
     Version: 39, Priority: 3
     Label Stack (Top -> Bottom): { 24001 }
     NHID: 0x0, Encap-ID: N/A, Path idx: 1, Backup path idx: 0, Weight: 0
     MAC/Encaps: 14/18, MTU: 1500
     Outgoing Interface: GigabitEthernet0/0/0/4 (ifhandle 0x000000c0)
     Packets Switched: 0
```

從先前的輸出可以清楚地看到，流量可以在兩個路徑選項中進行負載均衡，但是有幾種方法可以幫助確定哪一條是首選路徑。一種方式是使用show cef exact-route <source IP> <destination IP>命令，增加來自源PE的Loopback0和來自目標PE的Loopback0。如下一個輸出所示，首選路徑是透過Gi0/0/0/4：

<#root>

```
RP/0/0/CPU0:PE4#
```

show cef exact-route 10.10.10.4 10.10.10.1

```
Mon Sep 11 23:49:44.558 UTC
```

```
10.10.10.1/32, version 39, internal 0x1000001 0x0 (ptr 0xa12dbdbc) [1], 0x0 (0xa12c18c0), 0xa28 (0xa185
 Updated Sep  8 20:27:26.596
 local adjacency 10.0.0.16
 Prefix Len 32, traffic index 0, precedence n/a, priority 3
   via GigabitEthernet0/0/0/4


via 10.0.0.16/32, GigabitEthernet0/0/0/4, 9 dependencies, weight 0, class 0 [flags 0x0]

         <<<<<
    path-idx 1 NHID 0x0 [0xa16765bc 0x0]
    next hop 10.0.0.16/32
    local adjacency
     local label 24010      labels imposed {24001}
```

另一個選項是首先驗證LIB（標籤資訊庫）並透過使用show mpls ldp bindings <prefix/mask>命令獲取目標Loopback0（屬於出口PE的10.10.10.1）的LDP繫結，從該輸出中找到本地繫結標籤後，使用命令show mpls forwarding exact-route label <label> ipv4 <source IP> <destination IP> detail 中的標籤值查詢首選路徑：

<#root>

RP/0/0/CPU0:PE4#

show mpls ldp bindings 10.10.10.1/32

```
Wed Sep 13 17:18:43.007 UTC
10.10.10.1/32, rev 29


Local binding: label: 24010

         <<<<<
        Remote bindings: (3 peers)
            Peer                Label
            -----------------   ---------
            10.10.10.3:0        24
            10.10.10.5:0        24001
            10.10.10.6:0        23
```

RP/0/0/CPU0:PE4#

show mpls forwarding exact-route label 24010 ipv4 10.10.10.4 10.10.10.1 detail

```
Wed Sep 13 17:20:06.342 UTC
Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface                    Switched
------ ----------- ------------------ ------------ --------------- ------------

24010 24001 10.10.10.1/32 Gi0/0/0/4 10.0.0.16 N/A

                  <<<<<
     Updated: Sep 12 14:15:37.009
     Version: 198, Priority: 3
     Label Stack (Top -> Bottom): { 24001 }
     NHID: 0x0, Encap-ID: N/A, Path idx: 1, Backup path idx: 0, Weight: 0
     Hash idx: 1
     MAC/Encaps: 14/18, MTU: 1500
     Outgoing Interface: GigabitEthernet0/0/0/4 (ifhandle 0x000000c0)
```

```
     Packets Switched: 0

     Via: Gi0/0/0/4, Next Hop: 10.0.0.16
     Label Stack (Top -> Bottom): { 24001 }
     NHID: 0x0, Encap-ID: N/A, Path idx: 1, Backup path idx: 0, Weight: 0
     Hash idx: 1
     MAC/Encaps: 14/18, MTU: 1500
     Outgoing Interface: GigabitEthernet0/0/0/4 (ifhandle 0x000000c0)
```

接下來，務必檢查資料平面中的下一跳路由器，在本例中，要檢驗的路由器是P5（介面為10.0.0.16）。首先要瞭解的是MPLS轉發表，其中字首10.10.10.1的本地標籤必須24001：

<#root>

RP/0/0/CPU0:P5#

show mpls forwarding

```
Thu Sep 14 20:07:16.455 UTC
Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface                    Switched
------ ----------- ------------------ ------------ --------------- ------------
24000  Pop         10.10.10.6/32      Gi0/0/0/2    10.0.0.11       361906

24001 Pop 10.10.10.1/32 Gi0/0/0/1 10.0.0.0 361002

               <<<<<
24002  Pop         10.0.0.4/31        Gi0/0/0/1    10.0.0.0        0
       Pop         10.0.0.4/31        Gi0/0/0/2    10.0.0.11       0
24003  Pop         10.10.10.2/32      Gi0/0/0/0    10.0.0.6        360940
24004  Pop         10.0.0.8/31        Gi0/0/0/0    10.0.0.6        0
       Pop         10.0.0.8/31        Gi0/0/0/2    10.0.0.11       0
24005  Pop         10.0.0.2/31        Gi0/0/0/0    10.0.0.6        0
       Pop         10.0.0.2/31        Gi0/0/0/1    10.0.0.0        0
24006  Pop         10.10.10.4/32      Gi0/0/0/4    10.0.0.17       361230
24007  Pop         10.0.0.12/31       Gi0/0/0/2    10.0.0.11       0
       Pop         10.0.0.12/31       Gi0/0/0/4    10.0.0.17       0
24008  Pop         10.10.10.3/32      Gi0/0/0/3    10.0.0.15       361346
24009  Pop         10.0.0.20/31       Gi0/0/0/3    10.0.0.15       0
       Pop         10.0.0.20/31       Gi0/0/0/4    10.0.0.17       0
24010  Pop         10.0.0.18/31       Gi0/0/0/2    10.0.0.11       0
       Pop         10.0.0.18/31       Gi0/0/0/3    10.0.0.15       0
```

RP/0/0/CPU0:P5#

show mpls forwarding labels 24001

```
Thu Sep 14 20:07:42.584 UTC
Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface                    Switched
------ ----------- ------------------ ------------ --------------- ------------

24001 Pop 10.10.10.1/32 Gi0/0/0/1 10.0.0.0 361060
```

RP/0/0/CPU0:P5#

從先前的輸出中可以看到，字首10.10.10.1/32的LFIB條目顯示「Pop」作為傳出標籤，這意味著此路由器是倒數第二跳躍點(PHP)。它還顯示流量必須根據LFIB資訊透過Gi0/0/0/1傳送，這也可以透過檢視CEF進行驗證。下一個CEF exact-route輸出顯示隱式Null標籤作為強加的標籤，這同樣是由於在Gi0/0/0/1連線的下一跳是標籤交換機路徑中的最後一個路由器，也是面向目標站點(CE-WEST)的PE。這也是為什麼路由器P5會刪除資料包而不對資料包施加其他標籤的原因，由於此過程，出口路由器PE1將接收不帶LDP標籤的資料包：

<#root>

RP/0/0/CPU0:P5#

show cef exact-route 10.10.10.4 10.10.10.1

```
Thu Sep 14 20:25:57.269 UTC
10.10.10.1/32, version 192, internal 0x1000001 0x0 (ptr 0xa1246394) [1], 0x0 (0xa122b638), 0xa20 (0xa15
 Updated Sep 12 14:15:38.009
 local adjacency 10.0.0.0
 Prefix Len 32, traffic index 0, precedence n/a, priority 3
   via GigabitEthernet0/0/0/1
   via 10.0.0.0/32, GigabitEthernet0/0/0/1, 9 dependencies, weight 0, class 0 [flags 0x0]
    path-idx 0 NHID 0x0 [0xa166e280 0xa166e674]
    next hop 10.0.0.0/32
    local adjacency
```

local label 24001 labels imposed {ImplNull}

<<<<<

檢驗標籤交換機路徑的最後一點是PE1。在檢視MPLS轉發表時，可以注意到LFIB中沒有字首10.10.10.1/32的條目：

<#root>

PE1#

show mpls forwarding-table

Local Label	Outgoing Label	Prefix or Tunnel Id	Bytes Label Switched	Outgoing interface	Next Hop
16	No Label	172.16.1.0/24[V]	12938	Gi3	10.10.0.1
17	No Label	172.16.2.0/24[V]	0	Gi3	10.10.0.1
18	Pop Label	10.0.0.6/31	0	Gi1	10.0.0.1
	Pop Label	10.0.0.6/31	0	Gi2	10.0.0.3
19	Pop Label	10.0.0.8/31	0	Gi2	10.0.0.3
	Pop Label	10.0.0.8/31	0	Gi4	10.0.0.5
20	Pop Label	10.0.0.10/31	0	Gi1	10.0.0.1
	Pop Label	10.0.0.10/31	0	Gi4	10.0.0.5
21	Pop Label	10.0.0.12/31	0	Gi4	10.0.0.5
22	Pop Label	10.0.0.14/31	0	Gi1	10.0.0.1
23	Pop Label	10.0.0.16/31	0	Gi1	10.0.0.1
24	Pop Label	10.0.0.18/31	0	Gi4	10.0.0.5
25	24009	10.0.0.20/31	0	Gi1	10.0.0.1
	22	10.0.0.20/31	0	Gi4	10.0.0.5
26	Pop Label	10.10.10.2/32	0	Gi2	10.0.0.3

```
27          24008      10.10.10.3/32    0             Gi1        10.0.0.1
            24         10.10.10.3/32    0             Gi4        10.0.0.5
28          24006      10.10.10.4/32    0             Gi1        10.0.0.1
            25         10.10.10.4/32    0             Gi4        10.0.0.5
29          Pop Label  10.10.10.5/32    0             Gi1        10.0.0.1
Local       Outgoing   Prefix           Bytes Label   Outgoing   Next Hop
Label       Label      or Tunnel Id     Switched      interface
30          Pop Label  10.10.10.6/32    0             Gi4        10.0.0.5
31     [T]  Pop Label  1/1[TE-Bind]     0             drop

[T]     Forwarding through a LSP tunnel.
        View additional labelling info with the 'detail' option
```

正如您已經瞭解的那樣，此行為的原因是字首(10.10.10.1/32)屬於PE1，並且路由器也為此連線的字首分配了一個隱式null標籤。這可以透過show mpls ldp bindings命令進行驗證：

<#root>

PE1#

show run interface loopback 0

```
Building configuration...

Current configuration : 66 bytes
!
interface Loopback0
 ip address 10.10.10.1 255.255.255.255
end
```

PE1#

show mpls ldp bindings 10.10.10.1 32

```
  lib entry: 10.10.10.1/32, rev 24
```

local binding: label: imp-null

```
        remote binding: lsr: 10.10.10.6:0, label: 23
        remote binding: lsr: 10.10.10.5:0, label: 24001
        remote binding: lsr: 10.10.10.2:0, label: 24000
```

由於PE1是Cisco IOS XE路由器，使用命令show bgp vpnv4 unicast all或show bgp vpnv4 unicast rd <value> <destination IP>有助於辨識和確認透過MP-BGP正確瞭解目標字首172.16.1.0/24。這些命令的輸出顯示了導出後的字首：

<#root>

PE1#

show bgp vpnv4 unicast all

```
BGP table version is 61, local router ID is 10.10.10.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
```

```
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 65000:1 (default for vrf WEST)
 *>i  0.0.0.0          10.10.10.3               0    100      0 65001 i
 *bi                   10.10.10.4               0    100      0 65001 i

 *> 172.16.1.0/24 10.10.0.1 0 0 65000 i

        <<<<<
 *>   172.16.2.0/24    10.10.0.1                0             0 65000 i
 *>i  192.168.1.0      10.10.10.4               0    100      0 65001 i
 *>i  192.168.2.0      10.10.10.3               0    100      0 65001 i
 *>i  192.168.3.0      10.10.10.4               0    100      0 65001 i
Route Distinguisher: 65001:1
 *>i  0.0.0.0          10.10.10.4               0    100      0 65001 i
 *>i  192.168.1.0      10.10.10.4               0    100      0 65001 i
 *>i  192.168.3.0      10.10.10.4               0    100      0 65001 i
Route Distinguisher: 65001:2
     Network          Next Hop            Metric LocPrf Weight Path
 *>i  0.0.0.0          10.10.10.3               0    100      0 65001 i
 *>i  192.168.2.0      10.10.10.3               0    100      0 65001 i

PE1#
```

```
show bgp vpnv4 unicast rd 65000:1 172.16.1.10

BGP routing table entry for 65000:1:172.16.1.0/24, version 2
Paths: (1 available, best #1, table WEST)
  Additional-path-install
  Advertised to update-groups:
     6
  Refresh Epoch 2
  65000


10.10.0.1 (via vrf WEST) from 10.10.0.1 (172.16.2.10)

         <<<<<
      Origin IGP, metric 0, localpref 100, valid, external, best


Extended Community: RT:65000:1 , recursive-via-connected

         <<<<<
      mpls labels in/out 16/nolabel
      rx pathid: 0, tx pathid: 0x0
      Updated on Sep 15 2023 18:27:23 UTC
```

類似地，檢視VRF中的BGP VPNv4字首，即CE-WEST接收的字首，使用命令show bgp vpnv4 unicast vrf <name> <prefix>，輸出顯示一直傳輸到入口PE4的MP-BGP標籤16以及配置的RT導出65000：1：

<#root>

PE1#

show bgp vpnv4 unicast vrf WEST 172.16.1.10

```
BGP routing table entry for 65000:1:172.16.1.0/24, version 2
Paths: (1 available, best #1, table WEST)
  Additional-path-install
  Advertised to update-groups:
     6
  Refresh Epoch 2
  65000
    10.10.0.1 (via vrf WEST) from 10.10.0.1 (172.16.2.10)
      Origin IGP, metric 0, localpref 100, valid, external, best
```

Extended Community: RT:65000:1 , recursive-via-connected

```
         <<<<<
```

mpls labels in/out 16/nolabel

```
        <<<<<
      rx pathid: 0, tx pathid: 0x0
      Updated on Sep 15 2023 18:27:23 UTC

PE1#
```

show run vrf WEST

```
Building configuration...

Current configuration : 478 bytes
vrf definition WEST
 rd 65000:1
```

route-target export 65000:1

```
         <<<<<
 route-target import 65000:1
 route-target import 65001:1
 route-target import 65001:2
 !
 address-family ipv4
 exit-address-family
!
!
interface GigabitEthernet3
 vrf forwarding WEST
 ip address 10.10.0.2 255.255.255.252
 negotiation auto
 no mop enabled
 no mop sysid
!
router bgp 65500
 !
 address-family ipv4 vrf WEST
  neighbor 10.10.0.1 remote-as 65000
  neighbor 10.10.0.1 activate
 exit-address-family
!
end
```

在此PE上要檢查的最後一個資訊是到目標IP的VRF級別的RIB和CEF條目，與在PE4上看到的條目相比，RIB上沒有字首172.16.1.0/24的標籤，原因是這是來自CE的路由，透過eBGP獲知並在將此字首導出到VPNv4之前插入VRF路由表中。這可以透過使用命令show ip route vrf <name> <prefix>和show ip cef vrf <name> <prefix>進行驗證，如下所示：

<#root>

PE1#

show ip route vrf WEST 172.16.1.10

```
Routing Table: WEST
Routing entry for 172.16.1.0/24
  Known via "bgp 65500", distance 20, metric 0
  Tag 65000, type external
  Last update from 10.10.0.1 1w0d ago
  Routing Descriptor Blocks:
  * 10.10.0.1, from 10.10.0.1, 1w0d ago, recursive-via-conn
      opaque_ptr 0x7F8B4E3E1D50
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 65000
      MPLS label: none
```

PE1#

show ip cef vrf WEST 172.16.1.10

```
172.16.1.0/24
  nexthop 10.10.0.1 GigabitEthernet3
```

此時，已確認流量CE (CE-EAST)的源已正確獲知目標字首172.16.1.0/24，該字首已透過MP-BGP正確傳播，並且已跨標籤交換機路徑獲知來自PE和Ps環回的標籤。但是，源/目標之間的可達性仍然不成功，並且還有最後一台路由器需要檢驗CE-WEST。要在此路由器中檢查的第一件事是路由表，請記住，源IP字首192.168.1.0/24應出現在其中：

<#root>

CE-WEST#

show ip route 192.168.1.10

% Network not in table

CE-WEST#

「Network not in table」顯然是問題，BGP表也可以驗證，但在查詢字首後它也不存在：

<#root>

CE-WEST#

show ip bgp

```
BGP table version is 41, local router ID is 172.16.2.10
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>   172.16.1.0/24    0.0.0.0                  0         32768 i
 *>   172.16.2.0/24    0.0.0.0                  0         32768 i
CE-WEST#
```

下一步，您可以驗證此提供商邊緣路由器(PE1)是否正在向eBGP鄰居CE-WEST通告字首，這可以透過使用命令show bgp vpnv4 unicast vrf <name> neighbors <neighbor IP> advertised-routes完成，如下所示：

<#root>

PE1#

show bgp vpnv4 unicast vrf WEST neighbors 10.10.0.1 advertised-routes

```
BGP table version is 61, local router ID is 10.10.10.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 65000:1 (default for vrf WEST)
 *>i  0.0.0.0          10.10.10.3               0    100      0 65001 i

*>i 192.168.1.0 10.10.10.4 0 100 0 65001 i

         <<<<<
 *>i  192.168.2.0      10.10.10.3               0    100      0 65001 i
 *>i  192.168.3.0      10.10.10.4               0    100      0 65001 i

Total number of prefixes 4
```

根據上一步驟，可以確認PE1路由器已正確將字首通告給CE-WEST，因此現在應該檢視CE端的BGP鄰居：

<#root>

CE-WEST#

show ip bgp neighbors

```
BGP neighbor is 10.10.0.2,  remote AS 65500, external link
  BGP version 4, remote router ID 10.10.10.1
  BGP state = Established, up for 1w4d
  Last read 00:00:40, last write 00:00:43, hold time is 180, keepalive interval is 60 seconds
  Neighbor sessions:
    1 active, is not multisession capable (disabled)
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Four-octets ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
    Enhanced Refresh Capability: advertised and received
    Multisession Capability:
    Stateful switchover support enabled: NO for session 1
  Message statistics:
    InQ depth is 0
    OutQ depth is 0

                         Sent       Rcvd
    Opens:                  1          1
    Notifications:          0          0
    Updates:                3         17
    Keepalives:         19021      18997
    Route Refresh:          2          0
    Total:              19029      19019
  Do log neighbor state changes (via global configuration)
  Default minimum time between advertisement runs is 30 seconds

 For address family: IPv4 Unicast
  Session: 10.10.0.2
  BGP table version 41, neighbor version 41/0
  Output queue size : 0
  Index 3, Advertise bit 0
  3 update-group member
  Inbound path policy configured
```

Route map for incoming advertisements is FILTER

```
        <<<<<
  Slow-peer detection is disabled
  Slow-peer split-update-group dynamic is disabled
                                Sent       Rcvd
  Prefix activity:              ----       ----
    Prefixes Current:              2          0
    Prefixes Total:                4         23
    Implicit Withdraw:             2         13
    Explicit Withdraw:             0         10
    Used as bestpath:            n/a          0
    Used as multipath:           n/a          0
    Used as secondary:           n/a          0

                                  Outbound    Inbound
  Local Policy Denied Prefixes:   --------    -------
    route-map:                           0          4
    Bestpath from this peer:            18        n/a
    Total:                              18          4
  Number of NLRIs in the update sent: max 2, min 0
  Last detected as dynamic slow peer: never
  Dynamic slow peer recovered: never
  Refresh Epoch: 3
```

```
  Last Sent Refresh Start-of-rib: 4d23h
  Last Sent Refresh End-of-rib: 4d23h
  Refresh-Out took 0 seconds
  Last Received Refresh Start-of-rib: 4d23h
  Last Received Refresh End-of-rib: 4d23h
  Refresh-In took 0 seconds
                                     Sent       Rcvd
        Refresh activity:            ----       ----
          Refresh Start-of-RIB        1          2
          Refresh End-of-RIB          1          2

  Address tracking is enabled, the RIB does have a route to 10.10.0.2
  Route to peer address reachability Up: 1; Down: 0
    Last notification 1w5d
  Connections established 3; dropped 2
  Last reset 1w4d, due to Peer closed the session of session 1
  External BGP neighbor configured for connected checks (single-hop no-disable-connected-check)
  Interface associated: GigabitEthernet0/3 (peering address in same link)
  Transport(tcp) path-mtu-discovery is enabled
  Graceful-Restart is disabled
  SSO is disabled
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Mininum incoming TTL 0, Outgoing TTL 1
Local host: 10.10.0.1, Local port: 179
Foreign host: 10.10.0.2, Foreign port: 39410
Connection tableid (VRF): 0
Maximum output segment queue size: 50

Enqueued packets for retransmit: 0, input: 0  mis-ordered: 0 (0 bytes)

Event Timers (current time is 0x4D15FD56):
Timer          Starts    Wakeups            Next
Retrans         19027          1             0x0
TimeWait            0          0             0x0
AckHold         19012      18693             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0
Linger              0          0             0x0
ProcessQ            0          0             0x0

iss: 1676751051  snduna: 1677112739  sndnxt: 1677112739
irs: 2109012892  rcvnxt: 2109374776

sndwnd:  16061  scale:      0  maxrcvwnd:  16384
rcvwnd:  15890  scale:      0  delrcvwnd:    494

SRTT: 1000 ms, RTTO: 1003 ms, RTV: 3 ms, KRTT: 0 ms
minRTT: 0 ms, maxRTT: 1000 ms, ACK hold: 200 ms
uptime: 1036662542 ms, Sent idletime: 40725 ms, Receive idletime: 40925 ms
Status Flags: passive open, gen tcbs
Option Flags: nagle, path mtu capable
IP Precedence value : 6

Datagrams (max data segment is 1460 bytes):
Rcvd: 37957 (out of order: 0), with data: 19014, total data bytes: 361883
Sent: 37971 (retransmit: 1, fastretransmit: 0, partialack: 0, Second Congestion: 0), with data: 19027,

 Packets received in fast path: 0, fast processed: 0, slow path: 0
 fast lock acquisition failures: 0, slow path: 0
```

```
TCP Semaphore      0x0F3194AC  FREE
```

先前的輸出顯示有應用於名為「FILTER」的傳入通告的路由對映，在檢視路由對映配置後，它顯示了一個指向192.168.0.0/16的permit語句的字首清單的匹配子句，但是這是不正確的，因為prefix-list僅允許該特定字首，而不是可以包含在此範圍中的所有字首：

```
<#root>

CE-WEST#

show route-map FILTER

route-map FILTER, permit, sequence 10
  Match clauses:

ip address prefix-lists: FILTER


  Set clauses:
  Policy routing matches: 0 packets, 0 bytes

CE-WEST#

show ip prefix-list FILTER

ip prefix-list FILTER: 1 entries

seq 5 permit 192.168.0.0/16

          <<<<<

CE-WEST#

show run | i ip prefix-list

ip prefix-list FILTER seq 5 permit 192.168.0.0/16
```

對字首清單配置稍作更改後，到達192.168.1.10的路由現在已安裝在RIB中：

```
<#root>

CE-WEST#

show run | i ip prefix-list

ip prefix-list FILTER seq 5 permit 192.168.0.0/16 le 32

          <<<<<

CE-WEST#

show ip bgp

BGP table version is 44, local router ID is 172.16.2.10
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
```

```
              x best-external, a additional-path, c RIB-compressed,
              t secondary path,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>   172.16.1.0/24    0.0.0.0                  0         32768 i
 *>   172.16.2.0/24    0.0.0.0                  0         32768 i


*> 192.168.1.0 10.10.0.2 0 65500 65001 i

         <<<<<
 *>   192.168.2.0      10.10.0.2                              0 65500 65001 i
 *>   192.168.3.0      10.10.0.2                              0 65500 65001 i

CE-WEST#

show ip route 192.168.1.10

Routing entry for 192.168.1.0/24

         <<<<<
  Known via "bgp 65000", distance 20, metric 0
  Tag 65500, type external
  Last update from 10.10.0.2 00:00:37 ago
  Routing Descriptor Blocks:
  * 10.10.0.2, from 10.10.0.2, 00:00:37 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
      Route tag 65500
      MPLS label: none
```

驗證

現在，源和目標之間的可接通性是成功的，並且可以確認traceroute是否透過與MPLS網路相同的標籤交換路徑：

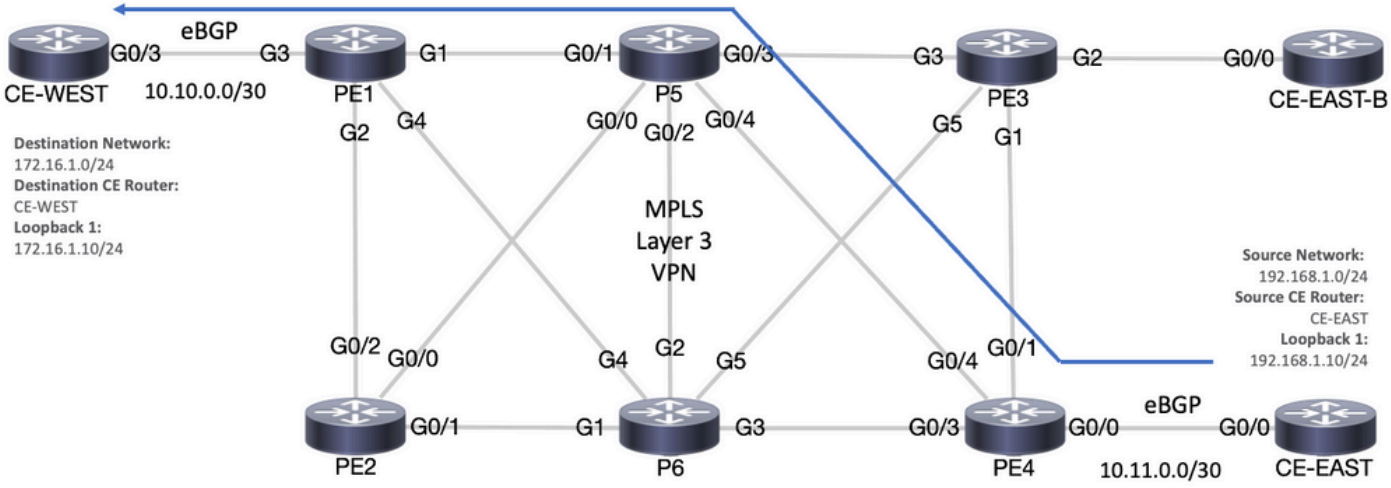

轉發路徑

CE-EAST#

ping 172.16.1.10 source loopback 1

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.1.10, timeout is 2 seconds:
Packet sent with a source address of 192.168.1.10
!!!!!
```

Success rate is 100 percent (5/5), round-trip min/avg/max = 7/7/9 ms

<<<<<

CE-EAST#

traceroute 172.16.1.10 source loop1 probe 1 numeric

```
Type escape sequence to abort.
Tracing the route to 172.16.1.10
VRF info: (vrf in name/id, vrf out name/id)
  1 10.11.0.2 2 msec
  2 10.0.0.16 [MPLS: Labels 24001/16 Exp 0] 9 msec
  3 10.10.0.2 [MPLS: Label 16 Exp 0] 8 msec
  4 10.10.0.1 9 msec
```

RP/0/0/CPU0:P5#

show ipv4 interface brief

```
Wed Sep 20 18:23:47.158 UTC

Interface                      IP-Address      Status          Protocol Vrf-Name
Loopback0                      10.10.10.5      Up              Up       default
MgmtEth0/0/CPU0/0              unassigned      Shutdown        Down     default
GigabitEthernet0/0/0/0         10.0.0.7        Up              Up       default
```

GigabitEthernet0/0/0/1 10.0.0.1 Up Up default

<<<<<

```
GigabitEthernet0/0/0/2         10.0.0.10       Up              Up       default
GigabitEthernet0/0/0/3         10.0.0.14       Up              Up       default
```

GigabitEthernet0/0/0/4 10.0.0.16 Up Up default

<<<<<

RP/0/0/CPU0:P5#

Cisco IOS XE驗證命令

<#root>

MPLS/LDP

```
show mpls interfaces
show mpls forwarding-table
show mpls ldp bindings [destination prefix]
```

```
show mpls ldp neighbor [neighbor address]
clear mpls ldp neighbor [neighbor address|*]
```

RIB and CEF

```
show ip vrf [detail]
show run vrf
show ip route [destination prefix]
show ip route vrf <name> [destination prefix]
show ip cef vrf <name> [destination prefix]
show ip cef exact-route <source IP> <destination IP>
show ip cef vrf <name> exact-route <source IP> <destination IP>
```

BGP/VPNv4

```
show ip bgp [neighbors] <neighbor address>
show bgp vpnv4 unicast all [summary|destination prefix]
show bgp vpnv4 unicast all neighbor <neighbor address> advertised-routes
show bgp vpnv4 unicast vrf <name> neighbors <neighbor IP> advertised-routes
show bgp vpnv4 unicast vrf <name> <prefix>
show bgp vpnv4 unicast rd <value> <destination IP>
```

Cisco IOS XR驗證命令

<#root>

MPLS/LDP

```
show mpls interfaces
show mpls forwarding
show mpls ldp bindings [destination prefix/mask]
show mpls ldp neighbor [neighbor address]
show mpls forwarding prefix [destination prefix/mask]
show mpls forwarding prefix [destination prefix/mask] detail hardware egress
clear mpls ldp neighbor [neighbor address]
```

RIB and CEF

```
show vrf [name|all]
show run vrf [name]
show route [destination prefix]
show route vrf <name> [destination prefix]
show cef vrf <name> [destination prefix]
show cef exact-route <source IP> <destination IP>
show cef vrf <name> exact-route <source IP> <destination IP>
```

BGP/VPNv4

```
show bgp vpnv4 unicast [summary|destination prefix/mask]
show bgp vpnv4 unicast neighbors <neighbor address> advertised-routes
show bgp vpnv4 unicast vrf <name> [prefix]
show bgp vrf <name> neighbors <neighbor IP> advertised-routes
show bgp vpnv4 unicast rd [value|all] [destination IP]
```

相關資訊

- MPLS基本MPLS配置指南
- 設定基本 MPLS VPN 網路
- 如何排除MPLS VPN故障
- 驗證網段路由SP的端到端連線

關於此翻譯

思科已使用電腦和人工技術翻譯本文件，讓全世界的使用者能夠以自己的語言理解支援內容。請注意，即使是最佳機器翻譯，也不如專業譯者翻譯的內容準確。Cisco Systems, Inc. 對這些翻譯的準確度概不負責，並建議一律查看原始英文文件（提供連結）。