

排除MPLS VPN中的LSP故障

目錄

[簡介](#)

[必要條件](#)

[需求](#)

[採用元件](#)

[慣例](#)

[網路圖表](#)

[路由器配置](#)

[問題](#)

[LSP故障的原因](#)

[解決方案](#)

[相關資訊](#)

簡介

本檔案假設您已預先瞭解基本的多重協定標籤交換(MPLS)概念。基於標籤轉發資訊庫(LFIB)中包含的資訊轉發MPLS交換資料包。通過標籤交換介面離開路由器的資料包將接收LFIB指定的值的標籤。根據轉發等價類(FEC)，標籤與LFIB中的目標關聯。FEC是一組IP資料包，它們通過同一路徑傳輸並接收相同的轉發處理。FEC的最簡單示例是到達特定子網的所有資料包。另一個範例可能是所有具有給定IP優先順序的封包都前往與邊界閘道通訊協定(BGP)路由群組相關聯的內部閘道通訊協定(IGP)下一躍點。

標籤資訊庫(LIB)是一種用於儲存從標籤分發協定(LDP)或標籤分發協定(TDP)鄰居處接收的標籤的結構。對於思科實施，會將指定路由器的路由表中的所有路由 (BGP路由除外) 的標籤傳送到所有LDP或TDP鄰居。無論是否使用鄰居標籤，從鄰居接收的所有標籤都會保留在LIB中。如果標籤從下游鄰居接收到其FEC，則儲存在LIB中的標籤將用於LFIB轉發資料包。根據路由器的Cisco Express Forwarding(CEF)和路由表，表示用於轉發的標籤是從路由器到目的地的下一跳收到的標籤。

如果從下游鄰居收到標籤繫結，查詢未在路由器的路由表和CEF表中顯示的字首 (包括子網掩碼)，則不會使用這些繫結。同樣地，如果路由器通告的子網/子網掩碼對標籤不對應於此路由器也針對同一子網/子網掩碼對通告的路由更新，則上游鄰居不會使用這些標籤，這些裝置之間的標籤交換路徑(LSP)將失敗。

本文檔提供了此類LSP故障的示例和幾個可能的解決方案。本文檔介紹一種情況，其中路由器接收的標籤繫結不用於轉發MPLS交換資料包。但是，用於診斷和糾正此問題的步驟適用於涉及為MPLS配置的路由器上的標籤繫結和LFIB的任何問題。

必要條件

需求

本文件沒有特定需求。

採用元件

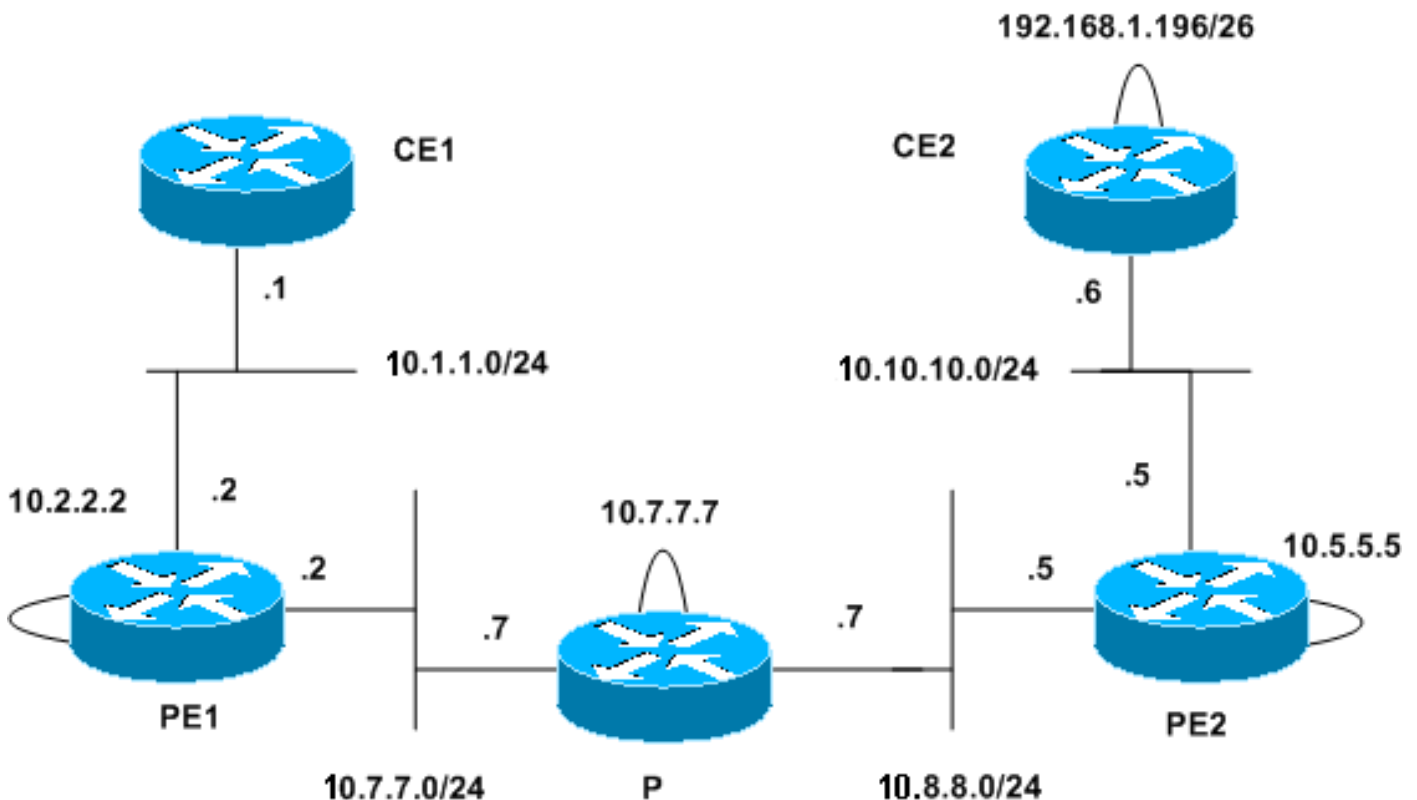
本檔案中的資訊是根據以下軟體版本：

- Cisco IOS®軟體版本12.0(21)ST2

慣例

如需文件慣例的詳細資訊，請參閱[思科技術提示慣例](#)。

網路圖表



路由器配置

PE1路由器配置

```
ip vrf aqua
 rd 100:1
  route-target export 1:1
  route-target import 1:1
!
interface Loopback0
 ip address 10.2.2.2 255.255.255.255
 no ip directed-broadcast
!
interface Ethernet2/0/1
 ip vrf forwarding aqua
 ip address 10.1.1.2 255.255.255.0
```

```

no ip directed-broadcast
ip route-cache distributed
!--- The VPN Routing and Forwarding (VRF) interface !---
toward the customer edge (CE) router. interface
Ethernet2/0/2 ip address 10.7.7.2 255.255.255.0 no ip
directed-broadcast ip route-cache distributed tag-
switching ip ! router ospf 1 log-adjacency-changes
network 0.0.0.0 255.255.255.255 area 0 ! router bgp 1
bgp log-neighbor-changes neighbor 10.5.5.5 remote-as 1
neighbor 10.5.5.5 update-source Loopback0 no auto-
summary ! address-family vpnv4 neighbor 10.5.5.5
activate neighbor 10.5.5.5 send-community extended exit-
address-family ! address-family ipv4 neighbor 10.5.5.5
activate no auto-summary no synchronization exit-
address-family ! address-family ipv4 vrf aqua
redistribute connected no auto-summary no
synchronization exit-address-family

```

P路由器配置

```

interface Loopback0
ip address 10.7.7.7 255.255.255.255
no ip directed-broadcast
!
interface Ethernet2/0
ip address 10.8.8.7 255.255.255.0
no ip directed-broadcast
tag-switching ip
!
interface Ethernet2/1
ip address 10.7.7.7 255.255.255.0
no ip directed-broadcast
tag-switching ip
!
router ospf 1
log-adjacency-changes
network 0.0.0.0 255.255.255.255 area 0

!--- BGP is not run on this router.

```

PE2路由器配置

```

ip vrf aqua
rd 100:1
route-target export 1:1
route-target import 1:1
!
interface Loopback0
ip address 10.5.5.5 255.255.255.0
no ip directed-broadcast
!
interface Ethernet0/0
ip vrf forwarding aqua
ip address 10.10.10.5 255.255.255.0
no ip directed-broadcast
!--- The VRF interface toward the CE router. ! interface
Ethernet0/3 ip address 10.8.8.5 255.255.255.0 no ip
directed-broadcast tag-switching ip ! router ospf 1 log-
adjacency-changes network 0.0.0.0 255.255.255.255 area 0
! router rip version 2 ! address-family ipv4 vrf aqua
version 2 network 10.0.0.0 no auto-summary exit-address-
family ! router bgp 1 bgp log-neighbor-changes neighbor
10.2.2.2 remote-as 1 neighbor 10.2.2.2 update-source

```

```
Loopback0 no auto-summary ! address-family vpnv4
neighbor 10.2.2.2 activate neighbor 10.2.2.2 send-
community extended exit-address-family ! address-family
ipv4 neighbor 10.2.2.2 activate no auto-summary no
synchronization exit-address-family ! address-family
ipv4 vrf aqua redistribute connected redistribute rip no
auto-summary no synchronization exit-address-family
```

CE2路由器配置

```
interface Loopback0
 ip address 192.168.1.196 255.255.255.192
 no ip directed-broadcast
!
interface Ethernet1
 ip address 10.10.10.6 255.255.255.0
 no ip directed-broadcast
!
router rip
 version 2
 network 10.0.0.0
 network 192.168.1.0
 no auto-summary
!--- Routing Information Protocol (RIP) is used for the
advertisement !--- of routes between the CE and the
provider edge (PE) router. ! ip route 0.0.0.0 0.0.0.0
10.10.10.5
```

註：CE1配置已被省略。該配置僅包含乙太網介面上的IP編址和通往10.2.2.2的靜態預設路由。

問題

CE1與CE2的環回介面之間的連線已丟失，如下例所示。

```
CE1#ping 192.168.1.196
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 192.168.1.196, timeout is 2 seconds:
```

```
.....
```

```
Success rate is 0 percent (0/5)
```

但是，CE1具有此目的地的有效路由條目，如以下示例所示。

```
CE1#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
```

```
Redistributing via ospf 100
```

```
Routing Descriptor Blocks:
```

```
* 10.1.1.2
```

```
Route metric is 0, traffic share count is 1
```

在PE1 (連線到CE1的PE路由器)，您可以檢查MPLS VPN特定資訊。以下示例說明此VPN的VRF表中存在到達目的地的有效路由。

```
PE1#show ip route vrf aqua 192.168.1.196
```

```
Routing entry for 192.168.1.192/26
```

```
Known via "bgp 1", distance 200, metric 1, type internal
```

```
Last update from 10.5.5.5 00:09:52 ago
```

Routing Descriptor Blocks:

```
* 10.5.5.5 (Default-IP-Routing-Table), from 10.5.5.5, 00:09:52 ago
  Route metric is 1, traffic share count is 1
  AS Hops 0, BGP network version 0
```

PE1#show tag-switching forwarding-table vrf aqua 192.168.1.196 detail

```
Local  Outgoing  Prefix          Bytes tag  Outgoing  Next Hop
tag    tag or VC  or Tunnel Id    switched   interface
None   16         192.168.1.192/26  0          Et2/0/2   10.7.7.7
      MAC/Encaps=14/22, MTU=1496, Tag Stack{16 32}
      00603E2B02410060835887428847 0001000000020000
      No output feature configured
```

PE1#show ip bgp vpnv4 vrf aqua 192.168.1.192

```
BGP routing table entry for 100:1:192.168.1.192/26, version 43
Paths: (1 available, best #1, table aqua)
  Not advertised to any peer
  Local
    10.5.5.5 (metric 21) from 10.5.5.5 (10.5.5.5)
      Origin incomplete, metric 1, localpref 100, valid, internal, best
      Extended Community: RT:1:1
```

PE1#show tag-switching forwarding-table 10.5.5.5 detail

```
Local  Outgoing  Prefix          Bytes tag  Outgoing  Next Hop
tag    tag or VC  or Tunnel Id    switched   interface
18     16         10.5.5.5/32     0          Et2/0/2   10.7.7.7
      MAC/Encaps=14/18, MTU=1500, Tag Stack{16}
      00603E2B02410060835887428847 00010000
      No output feature configured
      Per-packet load-sharing, slots: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
```

如本例所示，PE1沒有具有正確掩碼的BGP下一跳路由。

PE1#

PE1#show ip route 10.5.5.5 255.255.255.0

% Subnet not in table

PE1#show ip route 10.5.5.5 255.255.255.255

Routing entry for 10.5.5.5/32

Known via "ospf 1", distance 110, metric 21, type intra area

Last update from 10.7.7.7 on Ethernet2/0/2, 00:38:55 ago

Routing Descriptor Blocks:

* 10.7.7.7, from 10.5.5.5, 00:38:55 ago, via Ethernet2/0/2

Route metric is 21, traffic share count is 1

從P路由器接收PE1用於到達此BGP下一跳的IGP路由資訊。如以下示例所示，此路由器的PE2環回掩碼也不正確，而且該字首的路由沒有正確的掩碼。

P#show ip route 10.5.5.5

Routing entry for 10.5.5.5/32

Known via "ospf 1", distance 110, metric 11, type intra area

Last update from 10.8.8.5 on Ethernet2/0, 00:47:48 ago

Routing Descriptor Blocks:

* 10.8.8.5, from 10.5.5.5, 00:47:48 ago, via Ethernet2/0

Route metric is 11, traffic share count is 1

P#show ip route 10.5.5.5 255.255.255.0

% Subnet not in table

LSP故障的原因

P路由器上的LFIB和標籤繫結顯示了此路由器和PE2之間的LSP故障的原因。10.5.5.5沒有傳出標籤。當資料包離開PE1時，它攜帶兩個標籤，即P路由器(16)生成的BGP下一跳標籤和PE2(32)生成的

VPN標籤。因為P路由器上的此條目顯示該目的地的無標籤標籤交換資料包，所以它將在沒有任何標籤的情況下傳送。因為VPN標籤32丟失，所以PE2永遠不會收到它，並且PE2將沒有正確的資訊將資料包轉發到正確的VPN目標。

```
P#show tag-switching forwarding-table 10.5.5.5 detail
```

```
Local   Outgoing   Prefix           Bytes tag   Outgoing   Next Hop
tag     tag or VC  or Tunnel Id     switched    interface
16      Untagged   10.5.5.5/32      5339        Et2/0      10.8.8.5
        MAC/Encaps=0/0, MTU=1504, Tag Stack{}
        No output feature configured
        Per-packet load-sharing, slots: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
```

如下例所示，P路由器的標籤繫結表顯示PE2(tsr:10.8.8.5:0)僅通告使用/24掩碼對10.5.5.5的繫結。P路由器和PE1(tsr:10.2.2.2:0)，而不是PE2。由於PE2通告的繫結與其也通告的路由不匹配，因此P路由器的LFIB中沒有將資料包轉發到此目的地的標籤。

```
P#show tag-switching tdp bindings detail
```

```
tib entry: 10.5.5.0/24, rev 67(no route)
  remote binding: tsr: 10.8.8.5:0, tag: imp-null
tib entry: 10.5.5.5/32, rev 62
  local binding: tag: 16
  Advertised to:
    10.2.2.2:0           10.8.8.5:0
  remote binding: tsr: 10.2.2.2:0, tag: 18
```

PE2通告的路由更新和標籤繫結之間出現差異的原因可以在該路由器的路由表和標籤繫結表中看到。直連環回顯示正確的/24掩碼，路由器用它來生成標籤繫結。由於此網路使用開放最短路徑優先(OSPF)，因此路由器使用/32遮罩通告此介面，如下例所示。

```
PE2#show ip route 10.5.5.5
```

```
Routing entry for 10.5.5.0/24
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Routing Descriptor Blocks:
  * directly connected, via Loopback0
    Route metric is 0, traffic share count is 1
```

```
PE2#show tag-switching tdp bindings detail
```

```
tib entry: 10.5.5.0/24, rev 142
  local binding: tag: imp-null
  Advertised to:
    10.7.7.7:0
tib entry: 10.5.5.5/32, rev 148
  remote binding: tsr: 10.7.7.7:0, tag: 16
```

```
PE2#show ip ospf interface loopback 0
```

```
Loopback0 is up, line protocol is up
  Internet Address 10.5.5.5/24, Area 0
  Process ID 1, Router ID 10.5.5.5, Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host
```

```
!--- OSPF advertises all interfaces of Network Type LOOPBACK as host !--- routes (/32).
```

解決方案

由於P路由器和PE1之間的LSP故障是由通告環回的路由與PE1生成的標籤繫結不匹配引起的，因此最簡單的解決方案是更改環回的掩碼，使其與OSPF通告的所有LOOPBACK型別網路的掩碼一致。

解決方案1:PE2上的子網掩碼更改

```
PE2#configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
PE2(config)#int lo 0
PE2(config-if)#ip add 10.5.5.5 255.255.255.255
PE2(config-if)#end
PE2#
```

PE1上的資訊與LSP發生故障的情況相同，如下例所示。

```
PE1#show tag-switching forwarding-table vrf aqua 192.168.1.196 detail
```

```
Local   Outgoing   Prefix                               Bytes tag  Outgoing   Next Hop
tag     tag or VC  or Tunnel Id                         switched  interface
None    16         192.168.1.192/26                     0         Et2/0/2    10.7.7.7
MAC/Encaps=14/22, MTU=1496, Tag      Stack{16 32}
00603E2B02410060835887428847 0001000000020000
No output feature configured
```

```
PE1#show tag-switching forwarding-table 10.5.5.5 detail
```

```
Local   Outgoing   Prefix                               Bytes tag  Outgoing   Next Hop
tag     tag or VC  or Tunnel Id                         switched  interface
18      16         10.5.5.5/32                          0         Et2/0/2    10.7.7.7
MAC/Encaps=14/18, MTU=1500, Tag      Stack{16}
00603E2B02410060835887428847 00010000
No output feature configured
```

```
Per-packet load-sharing, slots: 0 1 2 3 4 5 6 7 8 9 10      11 12 13 14 15
```

P路由器顯示，導致LSP故障的情況不再出現。傳出標籤現在是彈出標籤。這表示當封包經過路由器時，BGP下一躍點的頂部標籤會彈出，但封包仍具有第二個VPN標籤（封包不再以未標籤的方式傳送出去）。

標籤繫結表顯示PE2通告了一個標籤(imp-null)(tsr:10.8.8.5:0)。

```
P#show tag-switching forwarding-table 10.5.5.5 detail
```

```
Local   Outgoing   Prefix                               Bytes tag  Outgoing   Next Hop
tag     tag or VC  or Tunnel Id                         switched  interface
16      Pop tag    10.5.5.5/32                          3493     Et2/0      10.8.8.5
MAC/Encaps=14/14, MTU=1504, Tag      Stack{}
006009E08B0300603E2B02408847
No output feature configured
```

```
Per-packet load-sharing, slots: 0 1 2 3 4 5 6 7 8 9 10 11    12 13 14 15
```

```
P#show tag-switching tdp bindings detail
```

```
tib entry: 10.5.5.5/32, rev 71
  local binding: tag: 16
    Advertised to:
      10.2.2.2:0                10.8.8.5:0
  remote binding: tsr: 10.2.2.2:0, tag: 18
  remote binding: tsr: 10.8.8.5:0, tag: imp-null
```

解決方案2:OSPF網路型別更改

第二種解決方案是更改環回介面的OSPF網路型別。當PE2的環回介面的OSPF網路型別更改為點對點時，環回字首不再使用/32掩碼自動通告。這意味著PE2在其路由表中引用直連子網時（包含/24子網掩碼），生成的標籤繫結將匹配從PE2接收的P路由器上的OSPF路由（包含用於此字首的/24子網掩碼）。

ip ospf network point-to-point命令可用於更改PE2環回介面上的網路型別，如下例所示。

```
PE2#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
PE2(config)#interface loopback 0
PE2(config-if)#ip ospf network point-to-point
PE2(config-if)#
```

如下圖所示，PE1上的標籤轉發表包含BGP下一跳的條目，該條目與PE2上環回介面的實際掩碼一致。路由表顯示了與此轉發表相關聯的OSPF路由也正確。

```
PE1#show tag-switching forwarding-table 10.5.5.5 detail
Local   Outgoing   Prefix          Bytes tag   Outgoing   Next Hop
tag     tag or VC  or Tunnel Id   switched   interface
22      17         10.5.5.0/24    0          Et2/0/2    10.7.7.7
        MAC/Encaps=14/18, MTU=1500, Tag Stack{17}
        00603E2B02410060835887428847 00011000
        No output feature configured
        Per-packet load-sharing, slots: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
```

```
PE1#show ip route 10.5.5.5
Routing entry for 10.5.5.0/24
  Known via "ospf 1", distance 110, metric 21, type intra area
  Last update from 10.7.7.7 on Ethernet2/0/2, 00:36:53 ago
  Routing Descriptor Blocks:
  * 10.7.7.7, from 10.5.5.5, 00:36:53 ago, via Ethernet2/0/2
    Route metric is 21, traffic share count is 1
```

在以下示例中，P路由器的標籤轉發條目將傳出標籤顯示為pop標籤，如解決方案1所示，如以下示例所示。同樣，當資料包經過此路由器時，BGP下一跳的頂部標籤會彈出，但第二個VPN標籤會保留，而LSP不會失敗。還顯示正確子網掩碼的繫結。

```
P#show tag-switching forwarding-table 10.5.5.5 detail
Local   Outgoing   Prefix          Bytes tag   Outgoing   Next Hop
tag     tag or VC  or Tunnel Id   switched   interface
17      Pop tag    10.5.5.0/24    4261       Et2/0      10.8.8.5
        MAC/Encaps=14/14, MTU=1504, Tag Stack{}
        006009E08B0300603E2B02408847
        No output feature configured
        Per-packet load-sharing, slots: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
```

```
P#show tag-switching tdp bindings detail

tib entry: 10.5.5.0/24, rev 68
  local binding: tag: 17
  Advertised to:
    10.2.2.2:0          10.8.8.5:0
  remote binding: tsr: 10.8.8.5:0, tag: imp-null
  remote binding: tsr: 10.2.2.2:0, tag: 22
```

如下圖所示，此命令的輸出確認網路型別已更改為點對點。從CE1到CE2的環回介面存在完全連線。

```
PE2#show ip ospf interface loopback 0
Loopback0 is up, line protocol is up
  Internet Address 10.5.5.5/24, Area 0
  Process ID 1, Router ID 10.5.5.5, Network Type POINT_TO_POINT, Cost: 1
  Transmit Delay is 1 sec, State POINT_TO_POINT,
```



```
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Index 3/3, flood queue length 0
Next 0x0(0)/0x0(0)
Last flood scan length is 0, maximum is 0
Last flood scan time is 0 msec, maximum is 0 msec
Neighbor Count is 0, Adjacent neighbor count is 0
Suppress hello for 0 neighbor(s)
```

CE1#**ping 192.168.1.196**

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 192.168.1.196, timeout is 2 seconds:

!!!!

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

CE1.

[相關資訊](#)

- [IP路由技術支援](#)
- [如何排除MPLS VPN故障](#)
- [MPLS技術支援](#)
- [技術支援與文件 - Cisco Systems](#)