含MPLS設計的Cisco IOS-XR BGP

目錄

簡介

本檔案介紹幾個案例，其中針對Cisco IOS®-XR中的多重通訊協定標籤交換(MPLS)和邊界閘道通訊協定(BGP)的組合，具有特殊的行為和組態。

1. Inter-AS MPLS VPN選項B & C需要靜態路由

此圖顯示Inter-AS選項B設定。

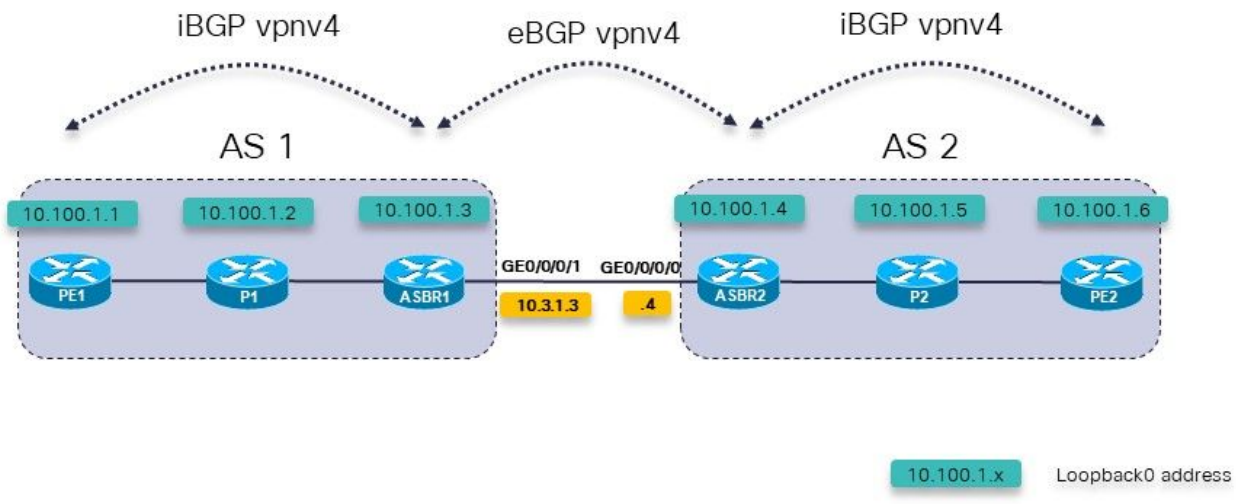

圖1.

提供商邊緣(PE)路由器PE1具有用於VRF字首10.200.1.2/32的路由，但它未解析。

```
RP/0/0/CPU0:PE1#show cef vrf one 10.200.1.2
10.200.1.2/32, version 3, internal 0x5000001 0x0 (ptr 0xa140be74) [1], 0x0 (0x0), 0x208
(0xa14a7118)
 Updated Apr  7 14:36:45.628
 Prefix Len 32, traffic index 0, precedence n/a, priority 3
   via 10.3.1.4/32, 0 dependencies, recursive [flags 0x6000]
    path-idx 0 NHID 0x0 [0xa0d87468 0x0]
    recursion-via-/32
    next hop VRF - 'default', table - 0xe0000000
    unresolved
    labels imposed {24004}
```

PE1沒有10.3.1.4/32的路由，但有10.3.1.0/24的路由。

```
RP/0/0/CPU0:PE1#show route 10.3.1.4

Routing entry for 10.3.1.0/24
  Known via "ospf 1", distance 110, metric 3, type intra area
  Installed Apr  7 14:07:01.140 for 00:32:48
  Routing Descriptor Blocks
    10.1.1.2, from 10.100.1.3, via GigabitEthernet0/0/0/0
      Route metric is 3
  No advertising protos.
```

自治系統邊界路由(ASBR)上必須有下一跳的靜態路由。您必須在每個ASBR上配置此靜態路由，並將其重新分配到內部網關協定(IGP)中。

```
router static
 address-family ipv4 unicast
  10.3.1.4/32 GigabitEthernet0/0/0/1
 !
!

router ospf 1
 redistribute static
```

路由現在已解析。

```
RP/0/0/CPU0:PE1#show cef vrf one 10.200.1.2
10.200.1.2/32, version 3, internal 0x5000001 0x0 (ptr 0xa140be74) [1], 0x0 (0x0), 0x208
(0xa14a7118)
 Updated Apr  7 14:36:45.628
 Prefix Len 32, traffic index 0, precedence n/a, priority 3
   via 10.3.1.4/32, 3 dependencies, recursive [flags 0x6000]
    path-idx 0 NHID 0x0 [0xa150f9f4 0x0]
    recursion-via-/32
    next hop VRF - 'default', table - 0xe0000000
    next hop 10.3.1.4/32 via 24005/0/21
     next hop 10.1.1.2/32 Gi0/0/0/0    labels imposed {24003 24004}
```

ASBR1為VPNv4/6字首安裝指向ASBR2的POP傳出標籤：

```
RP/0/0/CPU0:ASBR1#show mpls forwarding prefix 10.3.1.4/32
Local  Outgoing    Prefix            Outgoing     Next Hop        Bytes
```

```
Label  Label       or ID              Interface
Switched
------ ----------- ------------------ ------------ -------------- -----
-------
24005  Pop         10.3.1.4/32        Gi0/0/0/1    10.3.1.4       2506
```

即使ASBR上使用指向iBGP鄰居的下一跳自身，如果ASBR上未配置靜態路由，ASBR之間的標籤轉發也將中斷。

在ASBR1上使用指向PE1的下一跳self，並且沒有靜態路由：

```
RP/0/0/CPU0:ASBR1#show mpls forwarding labels 24006 detail
Local  Outgoing    Prefix           Outgoing     Next Hop        Bytes
Label  Label       or ID            Interface
Switched
------ ----------- ---------------- ------------ --------------- -----
-------
24006  24004       2:2:10.200.1.2/32               10.3.1.4
0
     Updated: Apr  7 14:49:58.190
     Path Flags: 0x6000 [  ]
     Label Stack (Top -> Bottom): { }
     MAC/Encaps: 0/0, MTU: 0
     Packets Switched: 0
```

請注意，「傳出介面」列中缺少傳出介面。Inter-AS選項B和C的ASBR上需要靜態路由。

2.在ASBR上保留Inter-AS選項B的路由目標

需要命令來確保ASBR儲存/保留vpnv4/6路由，然後通告它們。如果沒有此命令，如果ASBR上沒有配置匯入路由的任何路由目標的本地VRF，或者它不是地址系列vpnv4/6的路由反射器(RR)，則ASBR不會儲存路由。

```
router bgp 1
 address-family ipv4 unicast
 !
 address-family vpnv4 unicast
  retain route-target all
 !
```

3. ASBR不通告IPv4標籤的單點傳播地址

Inter-AS選項C或無縫MPLS（統一MPLS）網路中需要IPv4 labeled-unicast。這是因為預設情況下會標籤vpnv4/6字首，但IPv4(IPv6)單播的情況並非如此。如果情況並非如此，則標籤交換路徑(LSP)端對端會中斷，流量會端對端失敗。

檢視圖2，其中顯示Inter-AS Option C設定。

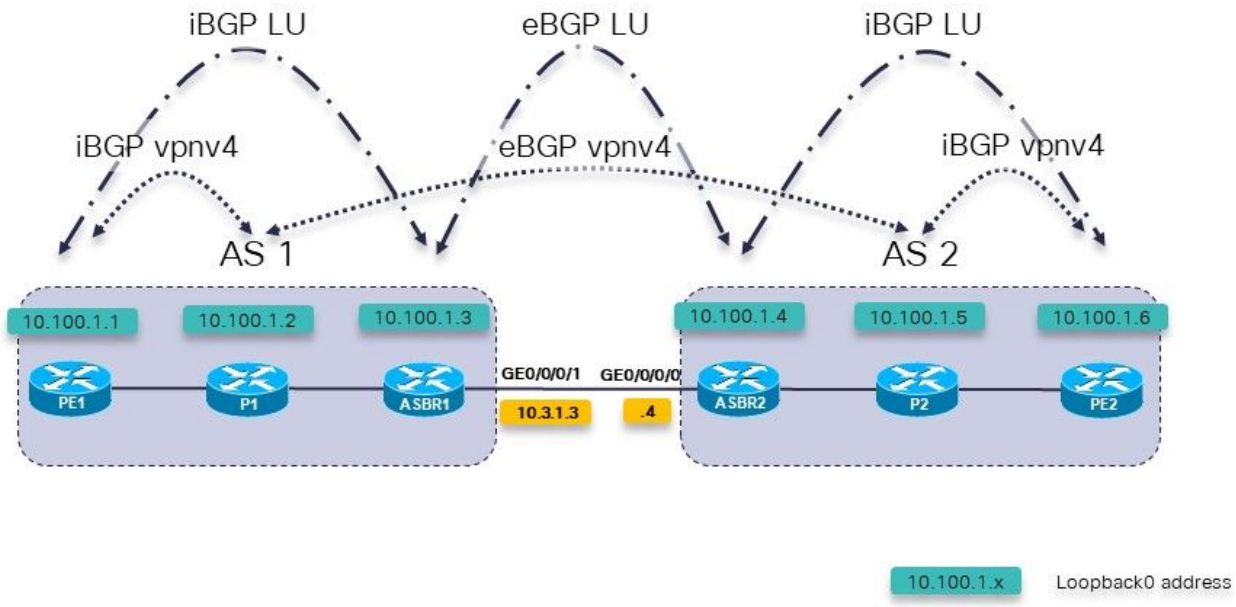

圖2.

P1和P2路由器也是vpnv4的自治系統(AS)中的路由反射器。

標籤單播(LU)用於將環回字首從一個AS傳輸到另一個AS。

ASBR1已配置此地址系列，但其中沒有路由：

```
RP/0/0/CPU0:ASBR1#show bgp ipv4 labeled-unicast
RP/0/0/CPU0:ASBR1#

RP/0/0/CPU0:ASBR1#show bgp ipv4 labeled-unicast summary
BGP router identifier 10.100.1.3, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000000   RD version: 41
BGP main routing table version 41
BGP NSR Initial initsync version 2 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
BGP scan interval 60 secs

BGP is operating in STANDALONE mode.


Process       RcvTblVer   bRIB/RIB   LabelVer  ImportVer  SendTblVer
StandbyVer
Speaker              41         41         41         41
41           0

Neighbor        Spk    AS MsgRcvd MsgSent   TblVer  InQ OutQ  Up/Down
St/PfxRcd
10.3.1.4          0     2     150     151       41    0    0
00:06:29          0
10.100.1.2        0     1      52      52       41    0    0
```

```
00:06:42           0
```

原因是ASBR必須使用以下命令，以便可以為每個路由分配多協定標籤交換(MPLS)標籤，然後通告路由。

```
RP/0/0/CPU0:ASBR1#show run router bgp
router bgp 1
 address-family ipv4 unicast
  redistribute ospf 1
  allocate-label all
 !
```

附註：如果指定了route-policy，該命令可以將標籤分配給特定字首。

此命令的結果如下：

```
RP/0/0/CPU0:ASBR1#show bgp ipv4 labeled-unicast
BGP router identifier 10.100.1.3, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000000   RD version: 52
BGP main routing table version 52
BGP NSR Initial initsync version 2 (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
*> 10.1.1.0/24        10.1.2.2                 2         32768 ?
*> 10.1.2.0/24        0.0.0.0                  0         32768 ?
*> 10.2.1.0/24        10.3.1.4                 0             0 2 ?
*> 10.2.2.0/24        10.3.1.4                 2             0 2 ?
*> 10.3.1.0/24        0.0.0.0                  0         32768 ?
*                     10.3.1.4                 0             0 2 ?
*> 10.100.1.1/32      10.1.2.2                 3         32768 ?
*> 10.100.1.2/32      10.1.2.2                 2         32768 ?
*> 10.100.1.3/32      0.0.0.0                  0         32768 ?
*> 10.100.1.4/32      10.3.1.4                 0             0 2 ?
*> 10.100.1.5/32      10.3.1.4                 2             0 2 ?
*> 10.100.1.6/32      10.3.1.4                 3             0 2 ?

Processed 11 prefixes, 12 paths

RP/0/0/CPU0:ASBR1#show bgp ipv4 labeled-unicast 10.100.1.6/32
BGP routing table entry for 10.100.1.6/32
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                48          48
    Local Label: 24008
Last Modified: Apr  7 16:20:04.509 for 00:00:49
Paths: (1 available, best #1)
  Advertised to peers (in unique update groups):
    10.100.1.2
  Path #1: Received by speaker 0
  Advertised to peers (in unique update groups):
    10.100.1.2
  2
```

```
    10.3.1.4 from 10.3.1.4 (10.100.1.4)
      Received Label 24002
     Origin incomplete, metric 3, localpref 100, valid, external, best, group-best
      Received Path ID 0, Local Path ID 1, version 48
      Origin-AS validity: not-found
```

所以，簡而言之：

- 必須配置標籤分配（在所有ASBR和PE路由器上）
- 必須為BGP鄰居配置地址系列ipv4 labeled-unicast

4. eBGP-eBGP（背對背的eBGP）Vpnv4需要虛擬iBGP鄰居

看圖3。

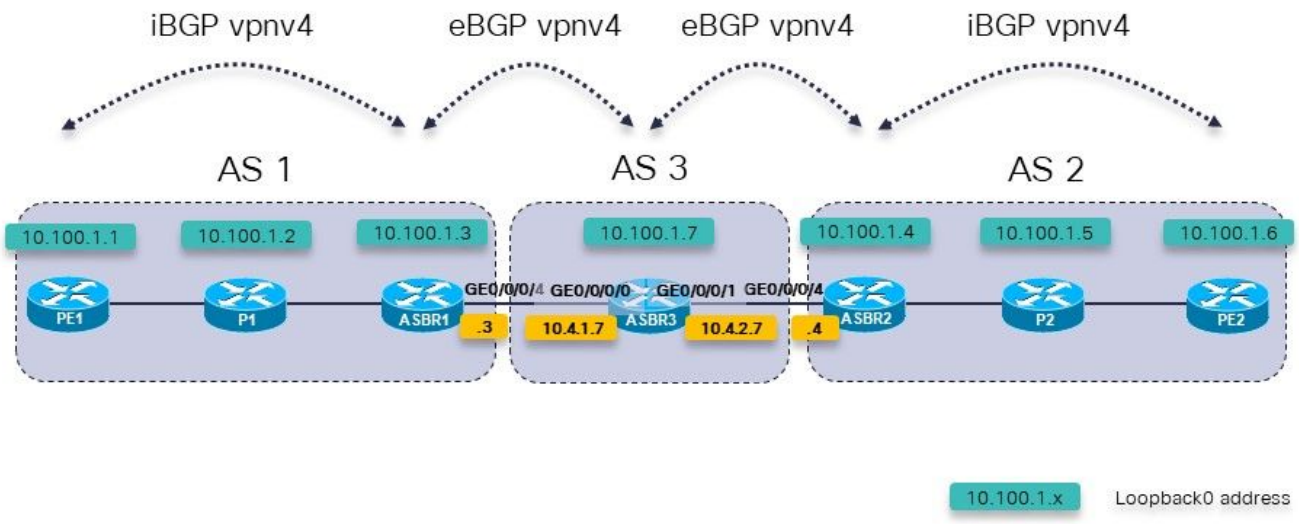

圖3.

連續有三個ASBR。ASBR3運行eBGP vpnv4單播到ASBR1和ASBR2。

附註：您還必須在ASBR3上配置靜態路由。

```
RP/0/0/CPU0:ASBR3#show bgp vpnv4 unicast
BGP router identifier 10.100.1.7, local AS number 3
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 3
BGP NSR Initial initsync version 2 (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: 1:1
*> 10.200.1.1/32      10.4.1.3                               0 1 ?
Route Distinguisher: 2:2
*> 10.200.1.2/32      10.4.2.4                               0 2 ?

Processed 2 prefixes, 2 paths

RP/0/0/CPU0:ASBR3#show bgp vpnv4 unicast rd 1:1 10.200.1.1/32
BGP routing table entry for 10.200.1.1/32, Route Distinguisher: 1:1
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 2           2
Last Modified: Apr  7 18:45:21.510 for 00:03:30
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  1
    10.4.1.3 from 10.4.1.3 (10.100.1.3)
      Received Label 24009
      Origin incomplete, localpref 100, valid, external, best, group-
best, import-candidate, not-in-vrf
      Received Path ID 0, Local Path ID 1, version 2
      Extended community: RT:1:1
```

從ASBR3通告vpnv4路由時出現問題：ASBR3不會向外通告外部vpnv4路由。

解決方案是在ASBR3上配置虛擬iBGP鄰居並啟用next-hop-self:虛構iBGP鄰居無需啟動。

```
router bgp 3
 address-family vpnv4 unicast
  retain route-target all
 !
 neighbor 10.4.1.3
  remote-as 1  address-family vpnv4 unicast
   route-policy PASS in
   route-policy PASS out
  !
 !
 neighbor 10.4.2.4
  remote-as 2
  address-family vpnv4 unicast
   route-policy PASS in
   route-policy PASS out
  !
 !
neighbor 10.99.99.99
  remote-as 3
  description dummy-iBGP neighbor for back-to-back eBGP vpnv4
  update-source Loopback0
  address-family vpnv4 unicast
   next-hop-self
  !
 !
!
```

結果是vpnv4路由現在被通告：

```
RP/0/0/CPU0:ASBR3#show bgp vpnv4 unicast rd 1:1 10.200.1.1/32
BGP routing table entry for 10.200.1.1/32, Route Distinguisher: 1:1
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 12          12
    Local Label: 24002
Last Modified: Apr  7 18:58:04.510 for 00:01:46
Paths: (1 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.2
  Path #1: Received by speaker 0
  Advertised to update-groups (with more than one peer):
   0.2
  1
    10.4.1.3 from 10.4.1.3 (10.100.1.3)
      Received Label 24009
      Origin incomplete, localpref 100, valid, external, best, group-
best, import-candidate, not-in-vrf
      Received Path ID 0, Local Path ID 1, version 12
      Extended community: RT:1:1
```

5. Inter-AS選項C — 首選BGP標籤而非LDP標籤

請參閱此映像，檢視通過多條鏈路連線的兩個ASBR的設定。為了讓此起作用，ASBR之間的eBGP ipv4 LU會話必須是多跳，因為它們之間具有並行鏈路。

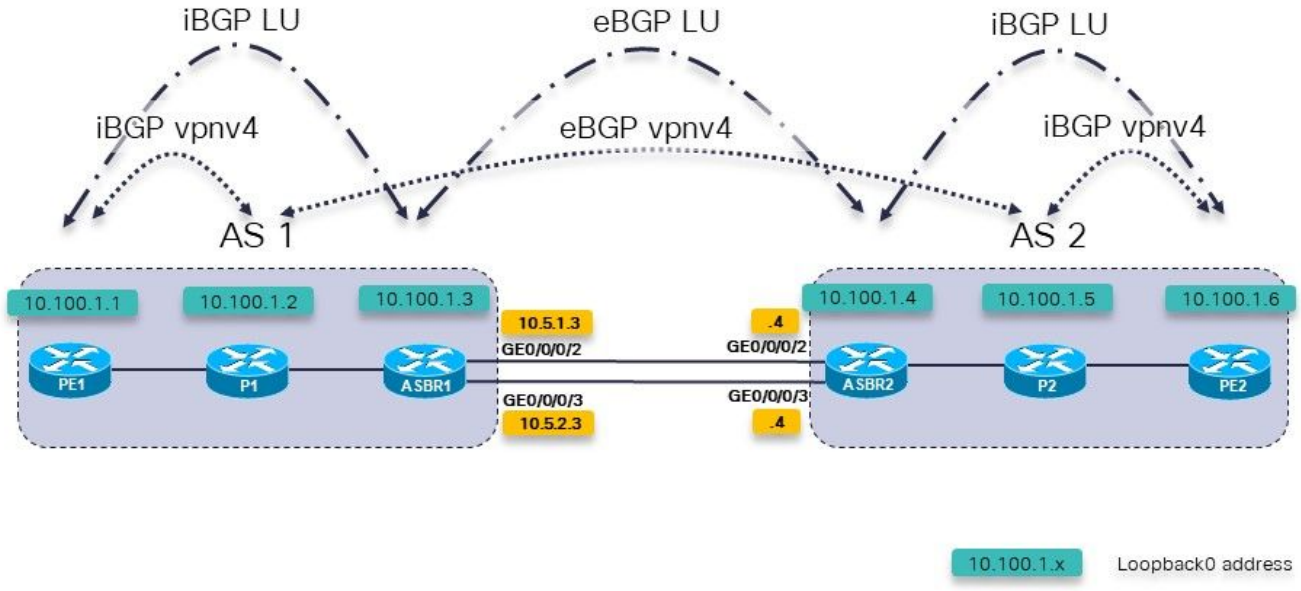

圖4.

這是Inter-AS選項C。路由器P1和P2也是vpnv4的路由反射器。

PE路由器和ASBR之間有IPv4標籤的單播。ASBR通過多條鏈路直接連線。

在ASBR上，您會看到：

```
router bgp 1
…
neighbor 10.100.1.4
  remote-as 2
  ebgp-multihop 2
  update-source Loopback0
  address-family ipv4 labeled-unicast
   route-policy PASS in
   route-policy PASS out
```

ASBR之間不需要標籤分發協定(LDP)。BGP將處理ASBR之間鏈路上的MPLS轉發。

```
RP/0/0/CPU0:ASBR1#show mpls interfaces

Interface                  LDP       Tunnel   Static   Enabled
-------------------------- --------  -------- -------- --------
GigabitEthernet0/0/0/0     Yes       No       No       Yes
GigabitEthernet0/0/0/1     No        No       No       Yes
GigabitEthernet0/0/0/2     No        No       No       Yes
GigabitEthernet0/0/0/3     No        No       No       Yes
GigabitEthernet0/0/0/4     No        No       No       Yes
```

目前為止一切順利。問題出在此圖所示的情境。

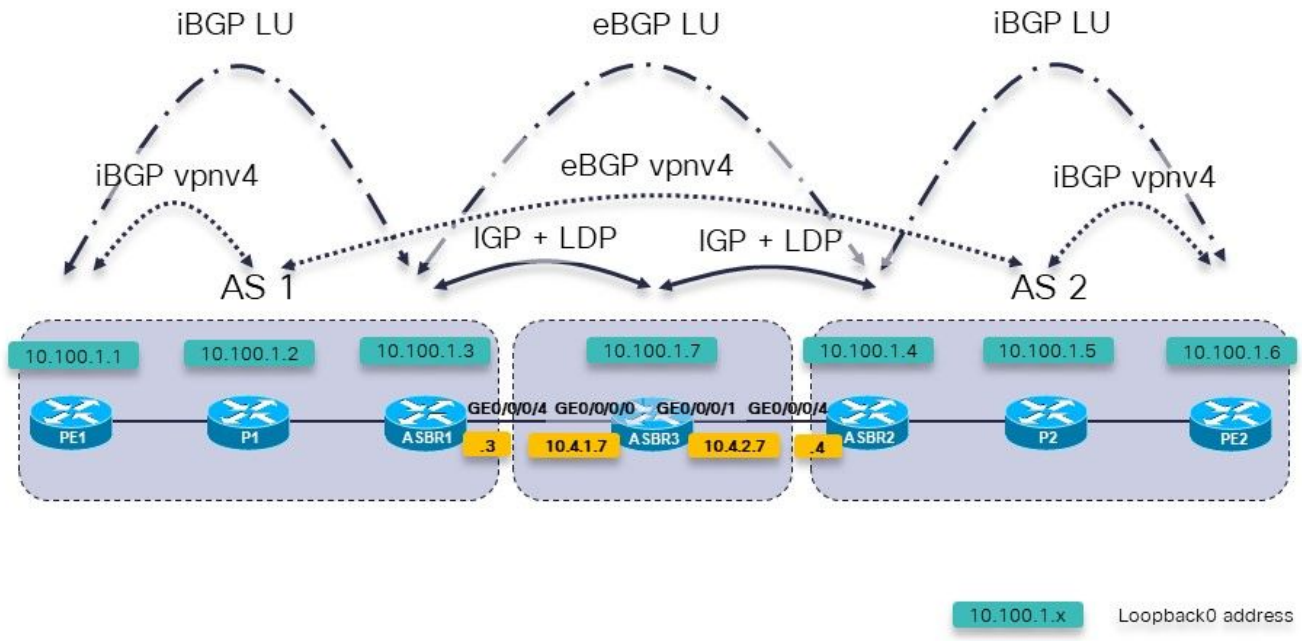

圖5.

這是Inter-AS選項C。路由器P1和P2也是vpnv4的路由反射器。

PE路由器和ASBR之間有IPv4標籤的單播。ASBR1和ASBR2未直接連線。它們通過運行IGP和LDP的網路進行多跳連線。在映像5中，此中間網路由路由器ASBR3表示，該路由器使用ASBR1和ASBR2運行IGP和LDP。

在ASBR上使用eBGP多躍點時，出現問題。每個AS中的RR之間的BGP會話甚至不啟動。

```
RP/0/0/CPU0:P1#show cef 10.100.1.5
10.100.1.5/32, version 263, internal 0x1000001 0x0 (ptr 0xa13bde74) [1], 0x0 (0xa1389560), 0xa28
(0xa14a72a8)
 Updated Apr  8 09:38:02.551
 local adjacency 10.1.2.3
 Prefix Len 32, traffic index 0, precedence n/a, priority 3
   via 10.1.2.3/32, GigabitEthernet0/0/0/1, 5 dependencies, weight 0,
class 0 [flags 0x0]
    path-idx 0 NHID 0x0 [0xa0e8b2a4 0x0]
    next hop 10.1.2.3/32
    local adjacency
     local label 24004      labels imposed {24007}
```

為了從P1、AS 1中的RR到P2、AS 2中的RR，傳出標籤為24007。在ASBR1上，此標籤交換為標籤24000。

```
RP/0/0/CPU0:ASBR1#show mpls forwarding labels 24007
Local  Outgoing    Prefix           Outgoing     Next Hop        Bytes
Label  Label       or ID            Interface                    Switched
------ ----------- ---------------- ------------ --------------- -----------
24007  24000       10.100.1.5/32                 10.100.1.4      1404

RP/0/0/CPU0:ASBR1#show cef 10.100.1.5
10.100.1.5/32, version 155, internal 0x5000001 0x0 (ptr 0xa13be174) [1],
0x0 (0xa138965c), 0xa08 (0xa14a72d0)
 Updated Apr  8 10:02:38.101
 Prefix Len 32, traffic index 0, precedence n/a, priority 4
   via 10.100.1.4/32, 5 dependencies, recursive, bgp-ext [flags 0x6020]
    path-idx 0 NHID 0x0 [0xa150f874 0x0]
    recursion-via-/32
    next hop 10.100.1.4/32 via 24004/0/21
     local label 24007
     next hop 10.4.1.7/32 Gi0/0/0/4     labels imposed {ImplNull 24000}
```

標籤24000是BGP LU從ASBR2在ASBR1上接收的標籤。

```
RP/0/0/CPU0:ASBR1#show bgp ipv4 labeled-unicast 10.100.1.5
BGP routing table entry for 10.100.1.5/32
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 76          76
    Local Label: 24007
Last Modified: Apr  8 09:37:57.509 for 00:04:05
Paths: (1 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.3
  Advertised to peers (in unique update groups):
    10.100.1.1      10.100.1.2
  Path #1: Received by speaker 0
  Advertised to update-groups (with more than one peer):
    0.3
  Advertised to peers (in unique update groups):
    10.100.1.1      10.100.1.2
  2
    10.100.1.4 from 10.100.1.4 (10.100.1.4)
      Received Label 24000
     Origin incomplete, metric 2, localpref 100, valid, external, best, group-best
      Received Path ID 0, Local Path ID 1, version 76
      Origin-AS validity: not-found
```

但是，介於兩者之間的ASBR路由器不運行BGP，因此無法轉發它收到的帶有此標籤的資料包，因為它沒有分配標籤24000。應該用於使資料包到達10.100.1.5的標籤是來自LDP的標籤：

```
RP/0/0/CPU0:ASBR1#show route 10.100.1.5/32

Routing entry for 10.100.1.5/32
  Known via "bgp 1", distance 20, metric 2, [ei]-bgp, labeled unicast
(3107)
  Tag 2, type external
  Installed Apr  8 10:02:38.082 for 01:24:37
  Routing Descriptor Blocks
    10.100.1.4, from 10.100.1.4, BGP external
      Route metric is 2
  No advertising protos.
```

這會遞迴到下一跳10.100.1.4，即ASBR2的環回。

應使用LDP從ASBR3接收的標籤，但不應使用。

新增的標籤堆疊為{ImplNull 24000}，而不是{24002 24000}。

```
RP/0/0/CPU0:ASBR1#show mpls ldp bindings 10.100.1.4/32
10.100.1.4/32, rev 146
        Local binding: label: 24004
        Remote bindings: (2 peers)
            Peer                 Label
            -----------------    ---------
            10.100.1.2:0         24003
            10.100.1.7:0         24002
```

ASBR1應該從ASBR3路由器24002收到的LDP標籤強制。若要停用BGP MPLS轉送，請將**mpls**關鍵字新增到eBGP multi-hop命令中。

ASBR1:

```
router bgp 1
…
neighbor 10.100.1.4
 remote-as 2
 ebgp-multihop 2 mpls
  update-source Loopback0
 address-family ipv4 labeled-unicast
  route-policy PASS in
  route-policy PASS out
 !
```

ASBR1現在具有正確的標籤重寫：

```
RP/0/0/CPU0:ASBR1#show cef 10.100.1.5
10.100.1.5/32, version 155, internal 0x5000001 0x0 (ptr 0xa13be174) [1], 0x0 (0xa138965c), 0xa08
(0xa14a72d0)
 Updated Apr  8 10:02:38.102
 Prefix Len 32, traffic index 0, precedence n/a, priority 4
  via 10.100.1.4/32, 5 dependencies, recursive, bgp-ext [flags 0x6020]
   path-idx 0 NHID 0x0 [0xa150f874 0x0]
   recursion-via-/32
   next hop 10.100.1.4/32 via 24004/0/21
    local label 24007
```

```
    next hop 10.4.1.7/32 Gi0/0/0/4    labels imposed {24002 24000}
```

在命令引用中：

在**ebgp-multihop**命令中使用**mpls**選項會阻止BGP在對等介面上啟用MPLS，還會阻止為從對等裝置獲取的下一跳地址分配隱式 — NULL重寫標籤。這在已通過BGP標籤 — 單播或LDP學習到MPLS轉發標籤到下一跳的某些場景中很有用。

換句話說，在IOS-XR中，當BGP提供向LFIB分配標籤時，它將優先於LDP。在ASBR路由器之間有多個躍點的Inter-AS選項C就是這樣一種場景。

6. Inter-AS選項B — 首選BGP標籤而非LDP標籤

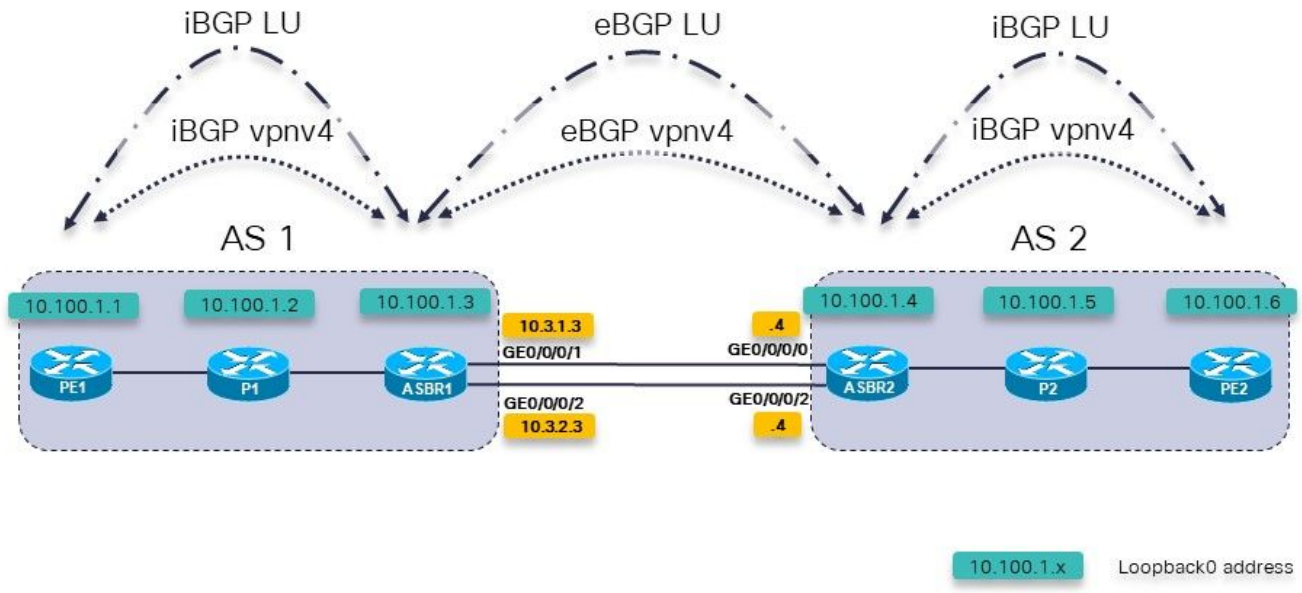

圖6.

這是Inter-AS選項B。但是，兩個ASBR之間存在多個並行鏈路。ASBR之間有RFC3107（交換IPv4路由和MPLS標籤），而不是使用IGP和LDP。

為了在ASBR1和ASBR2的環回介面之間啟動eBGP多跳會話，兩個ASBR之間需要eBGP LU。ASBR之間有兩條鏈路，因此需要兩個eBGP LU會話。地址系列IPv4需要命令allocate-label。

```
router bgp 65001
 address-family ipv4 unicast
  network 10.100.1.3/32
  allocate-label all
!
 neighbor 10.3.1.4
  remote-as 65002
  address-family ipv4 labeled-unicast
   route-policy pass in
   route-policy pass out
  !
 !
```

```
 neighbor 10.3.2.4
  remote-as 65002
  address-family ipv4 labeled-unicast
   route-policy pass in
   route-policy pass out
  !
 !
```

仍然需要來自第1部分的靜態路由：

```
router static
 address-family ipv4 unicast
  10.3.1.4/32 GigabitEthernet0/0/0/1
  10.3.2.4/32 GigabitEthernet0/0/0/2
 !
!
```

ASBR之間的eBGP vpnv4會話：

```
router bgp 65001
 address-family ipv4 unicast
  network 10.100.1.3/32
  allocate-label all
 !
 address-family vpnv4 unicast
  retain route-target all
 !
 neighbor 10.100.1.4
  remote-as 65002
 ebgp-multihop 255
  update-source Loopback0
  address-family vpnv4 unicast
   route-policy pass in
   route-policy pass out
  !
 !
```

請注意，此處不需要mpls關鍵字，如第5節所述。此外，如果為iBGP vpnv4會話配置了**next-hop-self**，則不需要在PE和ASBR之間執行iBGP LU會話。ASBR2為10.100.1.4/32通告的標籤為標籤3:

```
RP/0/0/CPU0:ASBR1#show bgp ipv4 labeled-unicast 10.100.1.4/32
Fri Jun  2 11:50:16.178 UTC
BGP routing table entry for 10.100.1.4/32
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 8           8
    Local Label: 24005
Last Modified: Jun  2 11:48:39.920 for 00:01:36
Paths: (4 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.3
  Advertised to peers (in unique update groups):
    10.100.1.7
  Path #1: Received by speaker 0
  Advertised to update-groups (with more than one peer):
    0.3
  Advertised to peers (in unique update groups):
    10.100.1.7
  65002
    10.3.1.4 from 10.3.1.4 (10.100.1.4)
      Received Label 3
```

```
      Origin IGP, metric 0, localpref 100, valid, external, best, group-best
      Received Path ID 0, Local Path ID 1, version 8
      Origin-AS validity: not-found
  Path #2: Received by speaker 0
  Not advertised to any peer
  65002
    10.3.2.4 from 10.3.2.4 (10.100.1.4)
      Received Label 3
      Origin IGP, metric 0, localpref 100, valid, external
      Received Path ID 0, Local Path ID 0, version 0
      Origin-AS validity: not-found
  Path #3: Received by speaker 0
  Not advertised to any peer
  65003 65002
    10.3.3.9 from 10.3.3.9 (10.100.1.9)
      Received Label 24001
      Origin IGP, localpref 100, valid, external, group-best
      Received Path ID 0, Local Path ID 0, version 0
      Origin-AS validity: not-found
  Path #4: Received by speaker 0
  Not advertised to any peer
  65003 65002
    10.3.4.9 from 10.3.4.9 (10.100.1.9)
      Received Label 24001
      Origin IGP, localpref 100, valid, external
      Received Path ID 0, Local Path ID 0, version 0
      Origin-AS validity: not-found

RP/0/0/CPU0:ASBR1#show cef 10.100.1.4
Fri Jun  2 11:51:06.994 UTC
10.100.1.4/32, version 254, internal 0x1000001 0x0 (ptr 0xa13be474) [1],
0x0 (0xa13896ec), 0xa20 (0xa14a70f0)
 Updated Jun  2 11:48:39.634
 local adjacency 10.3.1.4
 Prefix Len 32, traffic index 0, precedence n/a, priority 4
   via 10.3.1.4/32, GigabitEthernet0/0/0/1, 5 dependencies, weight 0,
class 0 [flags 0x0]
    path-idx 0 NHID 0x0 [0xa0e8b1fc 0xa0e8b34c]
    next hop 10.3.1.4/32
    local adjacency
     local label 24005      labels imposed {ImplNull}

RP/0/0/CPU0:ASBR1#show mpls forwarding labels 24005
Fri Jun  2 11:51:20.204 UTC
Local  Outgoing    Prefix            Outgoing     Next Hop        Bytes
Label  Label       or ID             Interface                    
Switched
------ ----------- ----------------- ------------ --------------- -----
-------
24005  Pop         10.100.1.4/32     Gi0/0/0/1    10.3.1.4        610
```

當ASBR之間存在另一個路徑，並且該路徑使用IGP + LDP或MPLS TE時，eBGP多跳命令需要mpls關鍵字。

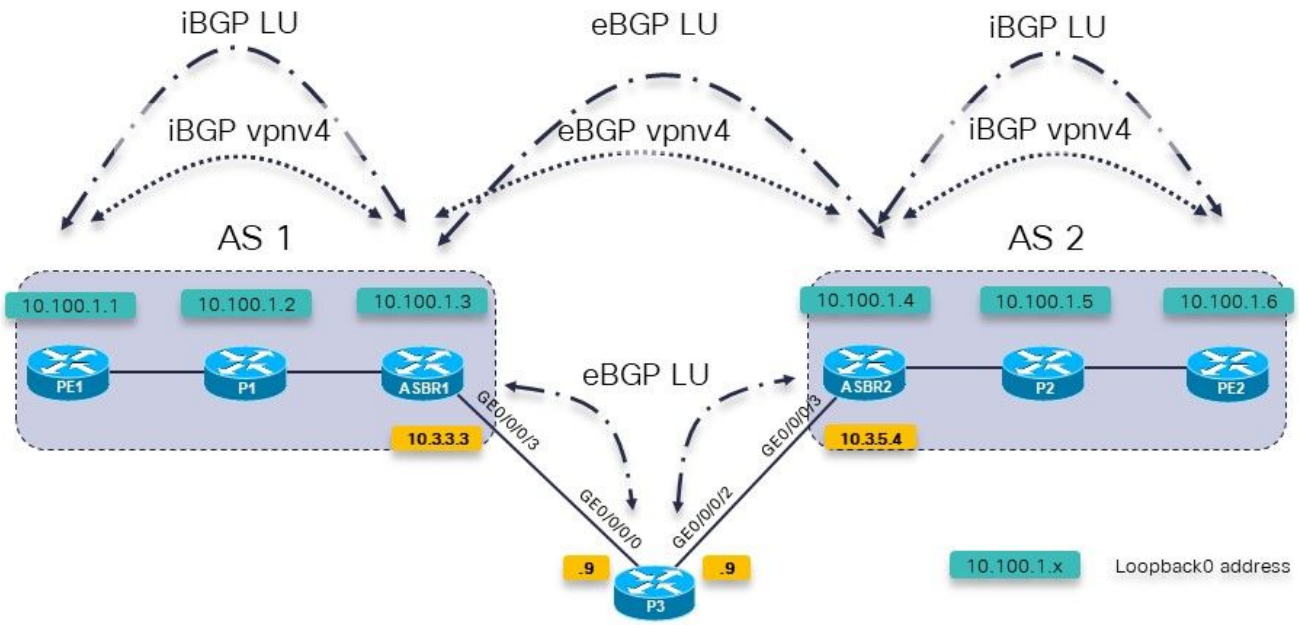

圖7.

ASBR1上指向P3的BGP路由策略用於設定非常高的權重，因此來自P3的字首優先於來自ASBR2的字首。

```
RP/0/0/CPU0:ASBR1#show bgp ipv4 labeled-unicast 10.100.1.4/32
Fri Jun  2 11:57:23.789 UTC
BGP routing table entry for 10.100.1.4/32
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 9           9
    Local Label: 24005
Last Modified: Jun  2 11:51:58.920 for 00:05:24
Paths: (4 available, best #3)
  Advertised to update-groups (with more than one peer):
    0.3
  Advertised to peers (in unique update groups):
    10.100.1.7
  Path #1: Received by speaker 0
  Not advertised to any peer
  65002
    10.3.1.4 from 10.3.1.4 (10.100.1.4)
      Received Label 3
      Origin IGP, metric 0, localpref 100, valid, external, group-best
      Received Path ID 0, Local Path ID 0, version 0
      Origin-AS validity: not-found
  Path #2: Received by speaker 0
  Not advertised to any peer
  65002
    10.3.2.4 from 10.3.2.4 (10.100.1.4)
      Received Label 3
      Origin IGP, metric 0, localpref 100, valid, external
      Received Path ID 0, Local Path ID 0, version 0
      Origin-AS validity: not-found
  Path #3: Received by speaker 0
  Advertised to update-groups (with more than one peer):
```

```
    0.3
  Advertised to peers (in unique update groups):
    10.100.1.7
  65003 65002
    10.3.3.9 from 10.3.3.9 (10.100.1.9)
     Received Label 24001
      Origin IGP, localpref 100, weight 65535, valid, external, best, group-best
      Received Path ID 0, Local Path ID 1, version 9
      Origin-AS validity: not-found
  Path #4: Received by speaker 0
  Not advertised to any peer
  65003 65002
    10.3.4.9 from 10.3.4.9 (10.100.1.9)
      Received Label 24001
      Origin IGP, localpref 100, valid, external
      Received Path ID 0, Local Path ID 0, version 0
      Origin-AS validity: not-found
```

ASBR1現在應使用label 24001作為10.100.1.4/32的傳出標籤。它不：

```
RP/0/0/CPU0:ASBR1#show cef 10.100.1.4
Fri Jun  2 11:59:46.519 UTC
10.100.1.4/32, version 255, internal 0x1000001 0x0 (ptr 0xa13be474) [1],
0x0 (0xa13896ec), 0xa20 (0xa14a7140)
 Updated Jun  2 11:51:58.741
 local adjacency 10.3.3.9
 Prefix Len 32, traffic index 0, precedence n/a, priority 4
   via 10.3.3.9/32, GigabitEthernet0/0/0/3, 7 dependencies, weight 0,
class 0 [flags 0x0]
    path-idx 0 NHID 0x0 [0xa0e8b544 0xa0e8b5ec]
    next hop 10.3.3.9/32
    local adjacency
     local label 24005      labels imposed {ImplNull}
```

解決方案與第5部分相同：將mpls關鍵字用於eBGP多跳命令。

```
RP/0/0/CPU0:ASBR1# conf t
Fri Jun  2 13:56:45.618 UTC
RP/0/0/CPU0:ASBR1(config)#router bgp 65001
RP/0/0/CPU0:ASBR1(config-bgp)# neighbor 10.100.1.4
RP/0/0/CPU0:ASBR1(config-bgp-nbr)#ebgp-multihop 255 mpls
RP/0/0/CPU0:ASBR1(config-bgp-nbr)#commit
```

ASBR1現在將標籤24001用作10.100.1.4/32的傳出標籤。

```
RP/0/0/CPU0:ASBR1#show cef 10.100.1.4
Fri Jun  2 13:58:13.402 UTC
10.100.1.4/32, version 200, internal 0x5000001 0x0 (ptr 0xa13be474) [1],
0x0 (0xa13895cc), 0xa08 (0xa14a71b8)
 Updated Jun  2 13:56:59.378
 Prefix Len 32, traffic index 0, precedence n/a, priority 15
   via 10.3.3.9/32, 3 dependencies, recursive, bgp-ext [flags 0x6020]
    path-idx 0 NHID 0x0 [0xa15102f4 0x0]
    recursion-via-/32
    next hop 10.3.3.9/32 via 24014/0/21
     local label 24005
     next hop 10.3.3.9/32 Gi0/0/0/3    labels imposed {ImplNull 24001}
```

ASBR1推送此額外標籤。從PE1到PE2的虛擬路由和轉發(VRF)中的traceroute顯示推送的額外標籤。

```
RP/0/0/CPU0:PE1#trace vrf one 10.99.1.2
Fri Jun  2 13:49:38.959 UTC

Type escape sequence to abort.
Tracing the route to 10.99.1.2

 1  10.1.1.5 [MPLS: Labels 24002/24012 Exp 0] 29 msec  39 msec  39 msec
 2  10.1.2.3 [MPLS: Label 24012 Exp 0] 29 msec  29 msec  39 msec
 3  10.3.1.4 [MPLS: Label 24007 Exp 0] 39 msec  39 msec  39 msec
 4  10.2.1.6 [MPLS: Labels 24001/24005 Exp 0] 39 msec  39 msec  29 msec
 5  10.2.2.2 39 msec  *  239 msec
```

ASBR1和P3之間使用IGP和LDP，ASBR2和P3之間使用IGP和LDP。在這些路由器之間使用MPLS流量工程(TE)時，存在相同的問題和解決方案。

從ASBR1到P3沒有LDP，但有MPLS TE。

在eBGP multihop命令上沒有mpls關鍵字，相同的問題會再次出現：

轉發到10.100.1.4的資料包不會獲得BGP LU標24000被推。

```
RP/0/0/CPU0:ASBR1#show cef 10.100.1.4
Tue Jun  6 10:36:56.528 UTC
10.100.1.4/32, version 50, internal 0x1000001 0x0 (ptr 0xa12cc1fc) [1],
0x0 (0xa12b18c0), 0xa20 (0xa14a7258)
 Updated Jun  6 10:36:32.930
 Prefix Len 32, traffic index 0, precedence n/a, priority 4
   via 10.3.3.9/32, tunnel-te1, 7 dependencies, weight 0, class 0 [flags
0x0]
    path-idx 0 NHID 0x0 [0xa15d58f8 0xa15d5840]
    next hop 10.3.3.9/32
    local adjacency
     local label 24012      labels imposed {ImplNull}
```

而使用mpls關鍵字時，標籤頁24000會出現：

```
RP/0/0/CPU0:ASBR1#show cef 10.100.1.4
Tue Jun  6 10:36:03.241 UTC
10.100.1.4/32, version 34, internal 0x5000001 0x0 (ptr 0xa12cc1fc) [1],
0x0 (0xa12b15a8), 0xa08 (0xa14a70f0)
 Updated Jun  6 09:39:24.56
 Prefix Len 32, traffic index 0, precedence n/a, priority 15
 Extensions: context-label:24012
   via 10.3.3.9/32, 3 dependencies, recursive, bgp-ext [flags 0x6020]
    path-idx 0 NHID 0x0 [0xa150fecc 0x0]
    recursion-via-/32
    next hop 10.3.3.9/32 via 24011/0/21
     local label 24012
     next hop 10.3.3.9/32 tt1          labels imposed {ImplNull 24000}
```

使用mpls關鍵字時，重寫如下所示：

```
RP/0/0/CPU0:ASBR1#show mpls forwarding labels 24012
Tue Jun  6 10:43:50.559 UTC
Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface                    Switched
------ ----------- ------------------ ------------ --------------- -----
-------
```

```
24012  24000       10.100.1.4/32      tt1          10.3.3.9        0
```
如果不使用mpls關鍵字，則重寫如下所示：

```
RP/0/0/CPU0:ASBR1#show mpls forwarding labels 24012
Tue Jun  6 10:45:08.734 UTC
Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface                    
Switched
------ ----------- ------------------ ------------ --------------- -----
-------
24012  Pop         10.100.1.4/32      tt1          10.3.3.9        0
```
此標籤14012不用於從VRF到VRF或PE到PE的流量，但是如果遇到此標籤，則可能表示標籤轉發例項庫(LFIB)條目錯誤或錯誤。

```
RP/0/0/CPU0:PE1# trace vrf one 10.99.1.2

Type escape sequence to abort.
Tracing the route to 10.99.1.2

 1  10.1.1.5 [MPLS: Labels 24001/24015 Exp 0] 129 msec  229 msec  129
msec
 2  10.1.2.3 [MPLS: Label 24015 Exp 0] 219 msec  439 msec  349 msec
 3  10.3.3.9 [MPLS: Labels 24000/24011 Exp 0] 169 msec  249 msec  139
msec
 4  10.3.5.4 [MPLS: Label 24011 Exp 0] 89 msec  129 msec  109 msec
 5  10.2.1.6 [MPLS: Labels 24004/24008 Exp 0] 139 msec  99 msec  139
msec
 6  10.2.2.2 129 msec  *  219 msec
```
在eBGP multihop命令上切換關鍵字mpls可能會導致BGP標籤衝突的系統日誌消息：

```
bgp[1051]: %ROUTING-BGP-4-LABEL_COLLISION : Label 24012 collision: prev: [T: 3 RD:0:0:0
PFX/NHID:10.100.1.4/32] curr: [T: 13 RD:0:0:0 PFX/NHID:10.100.1.4/32]
```
此消息用於本地標籤24012。

檢查是為了確保BGP擁有的活動標籤不會被BGP重新分配到任何其他內容。此檢查僅適用於每個字首標籤。

此郵件是症狀而非導致本文出現任何問題的原因。

7.使用VPN（或6PE或EVPN）底層的多跳BGP會話

如果存在eBGP多重躍點作業階段，則無法透過vpnv4/6或6PE（使用MPLS的IPv6）或乙太網路虛擬私人網路(EVPN)路由得知下一個躍點位址的路由，除非路由器具有Cisco IOS®-XR 6.3.2或更新版本。請參閱此映像。

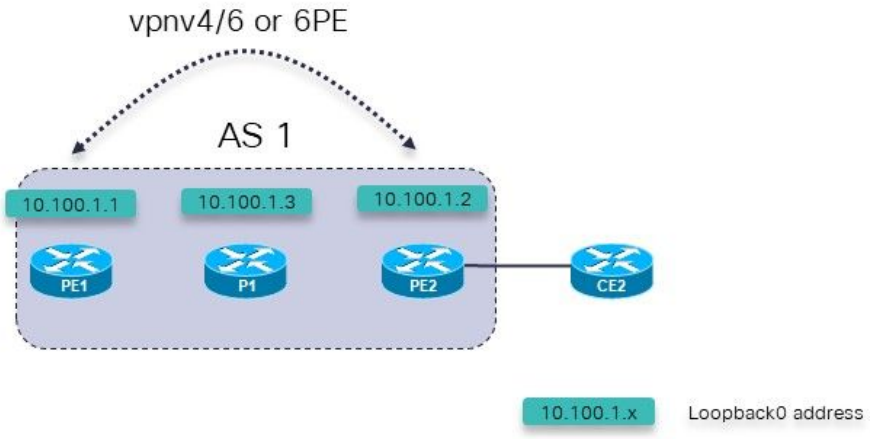

圖8.

可能的故障情形：

1. 從PE1（VRF內）到PE2（VRF內）的eBGP多跳會話
2. 從PE1（VRF內部）到CE2的eBGP多跳會話

這適用：

eBGP多跳會話配置在PE路由器上路由器BGP下的VRF部分下。

從PE1（在VRF內）到PE2（在VRF內）的eBGP多跳會話，或從PE1（在VRF內）到CE2的eBGP多跳會話，只支援從Cisco IOS®-XR 6.3.2開始。

可通過包含vnpv4的底層到達eBGP對等體地址。vpnv6、6PE或EVPN。

在6.3.2之前的Cisco IOS®版本中，eBGP會話將處於空閒狀態。

例如，在VRF one中配置了eBGP多跳會話PE1到PE2。

從PE1到PE2的eBGP多跳會話的相關配置：

```
interface Loopback100
 vrf one
 ipv4 address 10.2.100.1 255.255.255.255

router bgp 1
 address-family vpnv4 unicast
 !
 neighbor 10.100.1.2
 remote-as 1
 update-source Loopback0
 address-family vpnv4 unicast
```

```
 !
 !
 vrf one
 rd 1:1
 address-family ipv4 unicast
 redistribute connected
 !
 neighbor 10.2.100.2
remote-as 65002
 ebgp-multihop 255
local-as 65001
 update-source Loopback100
 address-family ipv4 unicast
 route-policy pass in
 route-policy pass out
 !
 !
 !
!
```

eBGP會話保持空閒：

```
RP/0/0/CPU0:PE1#show bgp vrf one neighbors

BGP neighbor is 10.2.100.2, vrf one
 Remote AS 65002, local AS 65001, external link
 Remote router ID 0.0.0.0
 BGP state = Idle (No route to multi-hop neighbor)
```

eBGP對等體地址的路由存在於VRF one路由表中：

```
RP/0/0/CPU0:PE1# show route vrf one

Codes: C - connected, S - static, R - RIP, B - BGP, (>) - Diversion path
 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 - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
 ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
 U - per-user static route, o - ODR, L - local, G - DAGR
 A - access/subscriber, a - Application route, (!) - FRR Backup path

Gateway of last resort is not set

L 10.2.100.1/32 is directly connected, 00:23:25, Loopback100
B 10.2.100.2/32 [200/0] via 10.100.1.2 (nexthop in vrf default), 00:19:28
```

```
RP/0/0/CPU0:PE1# show route vrf one 10.2.100.2/32

Routing entry for 10.2.100.2/32
 Known via "bgp 1", distance 200, metric 0, type internal
 Installed May 29 09:07:53.368 for 00:19:36
 Routing Descriptor Blocks
 10.100.1.2, from 10.100.1.2
 Nexthop in Vrf: "default", Table: "default", IPv4 Unicast, Table Id: 0xe0000000
Route metric is 0
 No advertising protos.
```

問題的根本原因是對等地址的路由是匯入的路由：

```
RP/0/0/CPU0:PE1# show bgp vpnv4 unicast vrf one 10.2.100.2/32

BGP routing table entry for 10.2.100.2/32, Route Distinguisher: 1:1
Versions:
 Process bRIB/RIB SendTblVer
 Speaker 7 7
Last Modified: May 29 09:07:53.524 for 00:21:20
Paths: (1 available, best #1)
 Not advertised to any peer
 Path #1: Received by speaker 0
 Not advertised to any peer
 Local
 10.100.1.2 (metric 2) from 10.100.1.2 (10.100.1.2)
 Received Label 16001
 Origin incomplete, metric 0, localpref 100, valid, internal, best,
group-best, import-candidate, imported
Received Path ID 0, Local Path ID 1, version 7
 Extended community: RT:1:1
Source VRF: one, Source Route Distinguisher: 1:1
```

Cisco IOS®-XR 6.3.2後支援此操作。

8.將BGP重新分配到LDP

這就是統一或無縫MPLS的本質及其使用IOS-XR配置的方式：使用IOS-XR的統一MPLS

對於常規的統一MPLS，所有PE路由器和ABR路由器之間都存在BGP LU，如下圖所示。

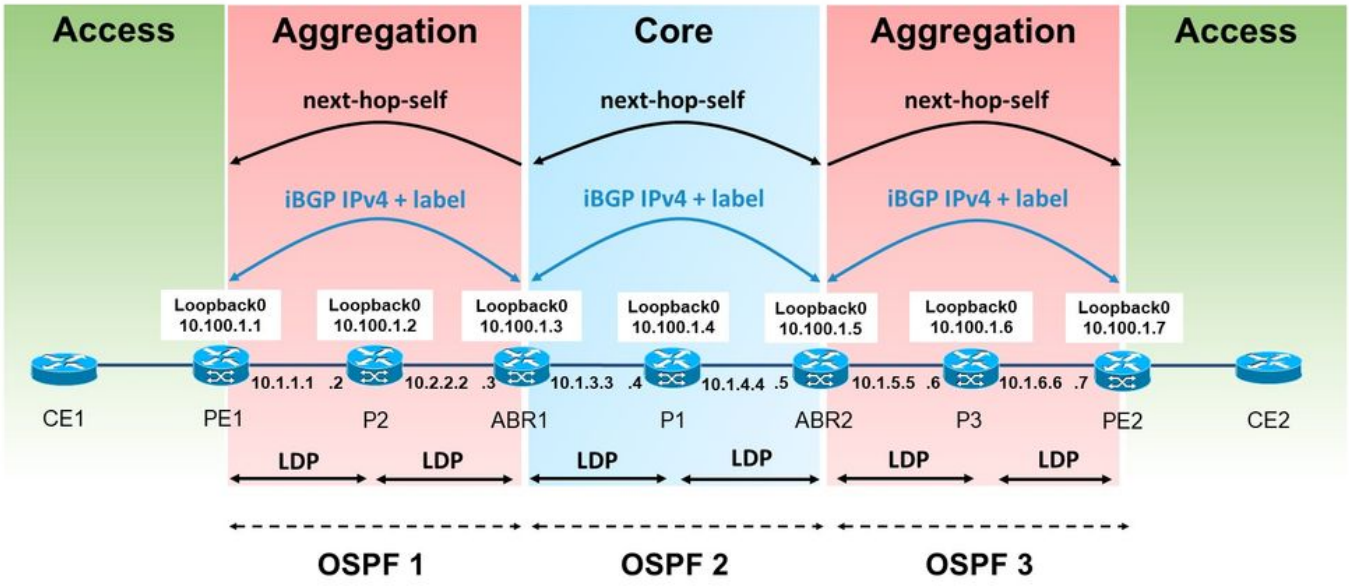

圖9.

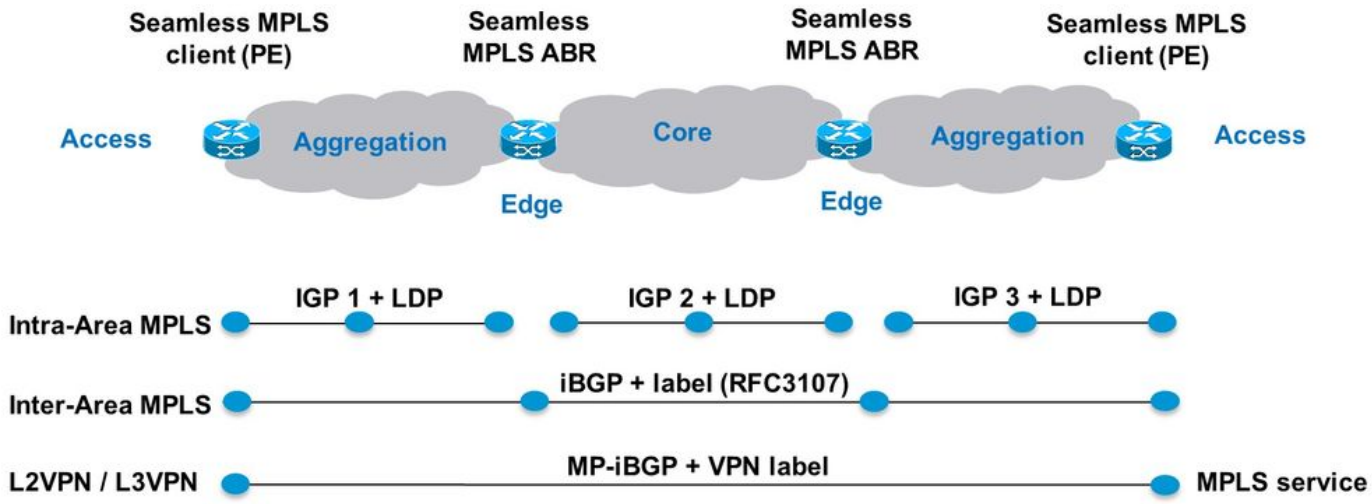

圖10.

在本例中，有一個IGP區域/級別沒有BGP LU。在左側，聚合區域實際上是開放最短路徑優先(OSPF)進程1，該進程在核心沒有使用OSPF進程2進行重分發。在使用OSPF 1的網路部分，PE路由器與區域邊界路由器(ABR)路由器之間沒有BGP LU。

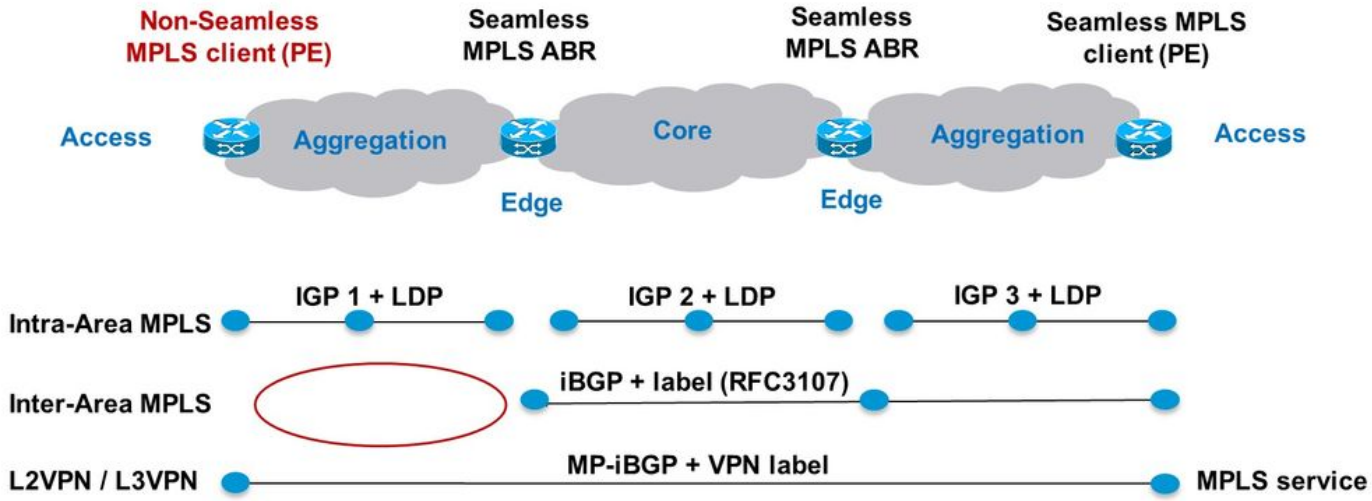

圖11.

BGP LU字首將重分佈到ABR1上的IGP OSPF 1中，如下圖所示。

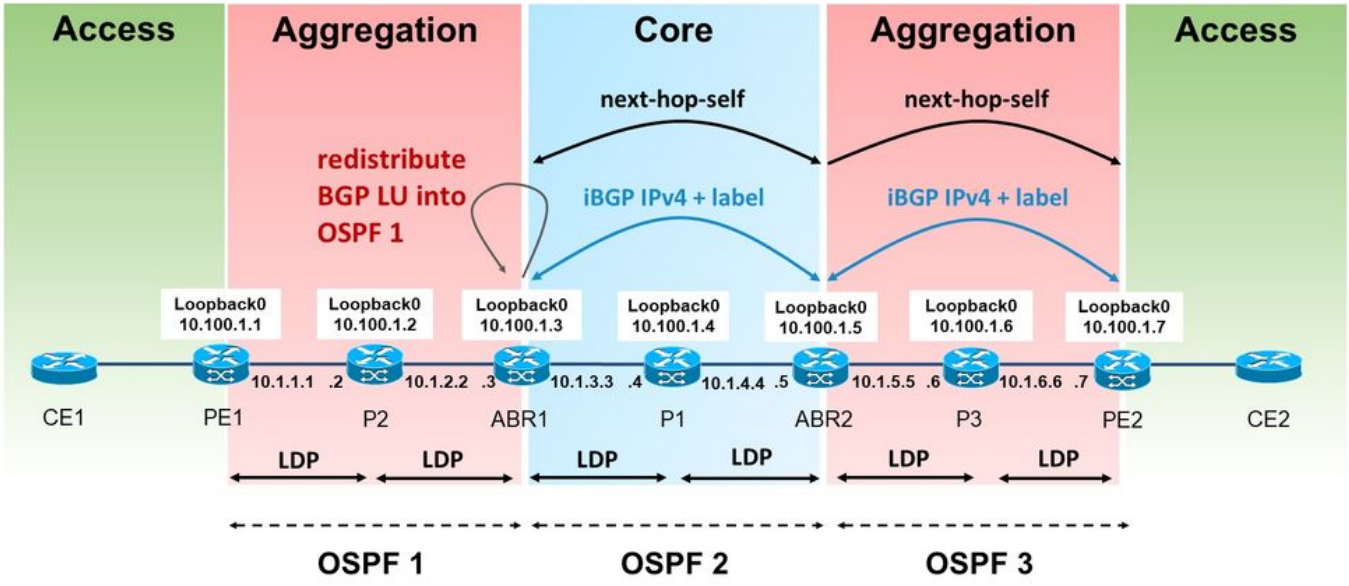

圖12.

您需要BGP為接收的iBGP LU字首分配標籤。但是，LDP不會在標籤繫結中自動通告此標籤的重新分發字首。預設情況下，IOS(-XE)會執行此操作。

請注意，ABR正在將內部BGP路由重分佈到左側區域的IGP中。這表示路由器bgp下需要命令**bgp redistribute-internal**。

```
router bgp 1
bgp redistribute-internal

router ospf 1
 router-id 10.100.1.3
redistribute bgp 1 metric 10 route-policy select-to-allocate
 area 0
  interface Loopback0
  !
  interface GigabitEthernet0/0/0/0
   network point-to-point
  !
 !
!
route-policy select-to-allocate
  if destination in (10.100.1.7/32) then
    pass
  else
    drop
  endif
end-policy
```

啟用本地標籤分配時，ABR會將本地標籤分配給接收到的iBGP LU路由。

```
router bgp 1
 bgp redistribute-internal
 ibgp policy out enforce-modifications
```

```
 address-family ipv4 unicast
  redistribute ospf 1 metric 10 route-policy ospf-1-loopbacks-PE
  allocate-label route-policy select-to-allocate
```

route-policy select-to-allocate可用於指定將哪些接收的BGP LU字首分配給本地標籤。

```
route-policy select-to-allocate
  if destination in (10.100.1.7/32) then
    pass
  else
    drop
  endif
end-policy
!
```

ABR1上使用本地標籤顯示PE2的環回字首，但LDP看不到此本地標籤：

```
RP/0/0/CPU0:ABR1#show bgp ipv4 labeled-unicast 10.100.1.7/32

BGP routing table entry for 10.100.1.7/32
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 6           6
    Local Label: 24006
Last Modified: Sep  5 06:55:47.368 for 06:40:23
Paths: (1 available, best #1)
  Advertised IPv4 Labeled-unicast paths to update-groups (with more than
one peer):
    0.2
  Path #1: Received by speaker 0
  Advertised IPv4 Labeled-unicast paths to update-groups (with more than
one peer):
    0.2
  Local, (Received from a RR-client)
    10.100.1.5 (metric 20) from 10.100.1.5 (10.100.1.7)
      Received Label 24003
      Origin IGP, metric 0, localpref 100, valid, internal, best, group-
best, labeled-unicast
      Received Path ID 0, Local Path ID 1, version 6
      Originator: 10.100.1.7, Cluster list: 10.100.1.5

RP/0/0/CPU0:ABR1#show mpls ldp bindings 10.100.1.7/32

10.100.1.7/32, rev 0 (no route)
        No local binding
        Remote bindings: (1 peers)
            Peer                Label
            -----------------   ---------
            10.100.1.2:0        18
```

這意味著從PE1到PE2的LSP中斷：

```
RP/0/0/CPU0:PE1#traceroute 10.100.1.7 source 10.100.1.1

Type escape sequence to abort.
Tracing the route to 10.100.1.7

 1  10.1.1.2 [MPLS: Label 18 Exp 0] 9 msec  0 msec  0 msec
 2  10.1.2.3 0 msec  0 msec  0 msec <<< no MPLS labels
 3  10.1.3.4 [MPLS: Labels 16/24003 Exp 0] 29 msec  19 msec  29 msec
 4  10.1.4.5 [MPLS: Label 24003 Exp 0] 9 msec  9 msec  9 msec
```

```
 5    *   *   *
 6   10.1.6.7 9 msec   *   19 msec
```

LSP在P2被中斷，因為它沒有從ABR1通過LDP獲取遠端標籤。ABR1沒有在LDP中為字首10.100.1.7/32本地分配**的標籤**。

在BGP路由被重分發到IGP的路由器上，ABR上需要配置將BGP重分發到LDP。

ABR1不會向P2路由器通告字首10.100.1.7/32的LDP標籤繫結。

為了使ABR1通告重新分發的iBGP字首的LDP標籤繫結，ABR1必須具有以下配置（必須配置AS編號）。

```
mpls ldp
 mldp
  address-family ipv4
  !
 !
 router-id 10.100.1.3
 address-family ipv4
  redistribute
   bgp
    as 1
    !
   !
 !
```

可以讓LDP過濾通告。例如，您可以按如下方式配置過濾器：

```
mpls ldp
 mldp
  address-family ipv4
  !
 !
 router-id 10.100.1.3
 address-family ipv4
  redistribute
   bgp
    as 1
    advertise-to 1
   !

ipv4 access-list 1
 10 permit ipv4 host 10.100.1.2 any
```

您可以在訪問清單中指定LDP router-ID。

在本示例中，ABR只向LDP鄰居P2（而不是向P1）通告重分佈的iBGP路由的LDP繫結，因為10.100.1.2是P2的LDP路由器ID。

從PE1到PE2的LSP現在不受干擾：

```
RP/0/0/CPU0:PE1#traceroute 10.100.1.7 source 10.100.1.1

Type escape sequence to abort.
Tracing the route to 10.100.1.7

 1  10.1.1.2 [MPLS: Label 20 Exp 0] 39 msec  49 msec  29 msec
 2  10.1.2.3 [MPLS: Label 24006 Exp 0] 29 msec  49 msec  39 msec
```

```
 3  10.1.3.4 [MPLS: Labels 16/24003 Exp 0] 29 msec  19 msec  29 msec
 4  10.1.4.5 [MPLS: Label 24003 Exp 0] 29 msec  19 msec  29 msec
 5   *  *  *
 6  10.1.6.7 19 msec  *  19 msec
```

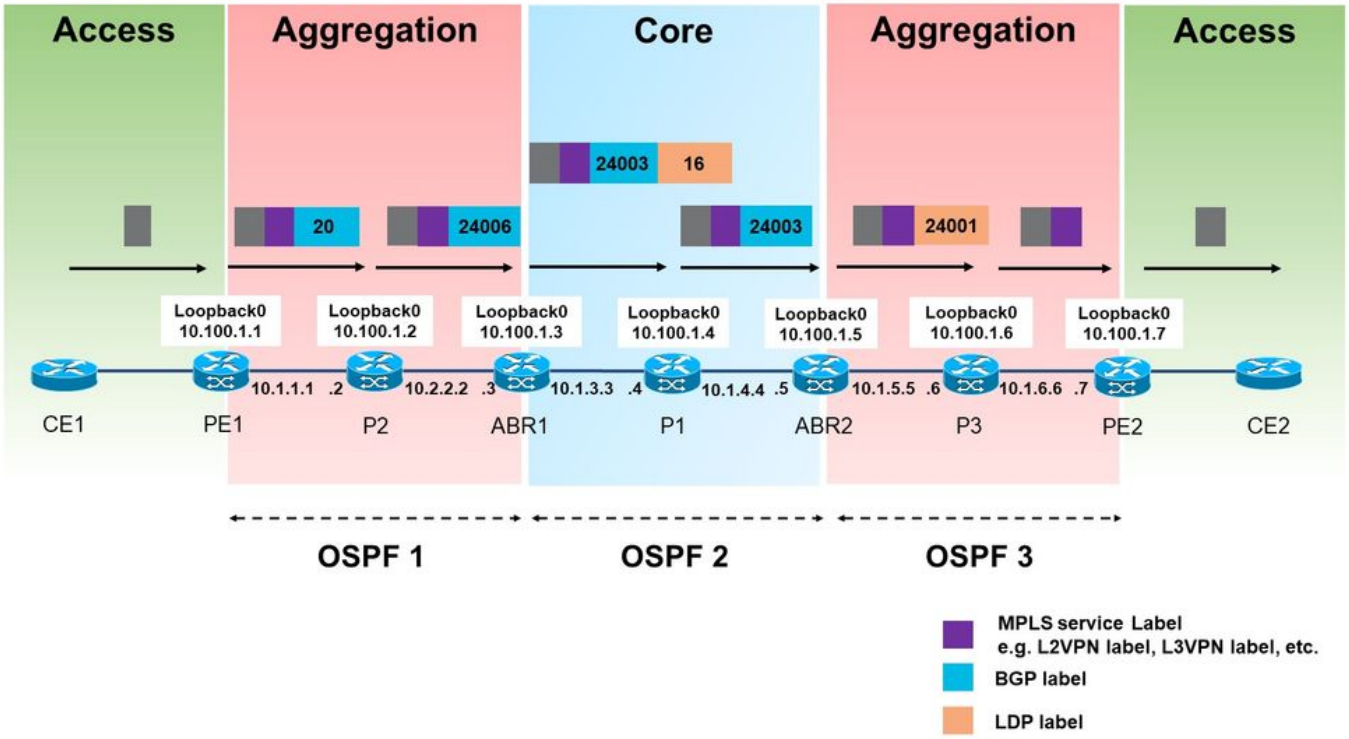

圖13.

現在，由LDP通告到左側聚合區域的BGP分配標籤(24006)用於從PE1到PE2的流量。

請注意，左側聚合區域中僅使用一個MPLS標籤。如果這是常規統一MPLS，則使用兩個標籤。

此時，您無法過濾哪些重分發的LU iBGP路由進入LDP、接收本地標籤以及哪些沒有。啟用iBGP LU路由到LDP的再分配後，它們都會獲得本地標籤。

PE2還在BGP LU中通告字首10.100.1.99/32。ABR1不會將此字首重新分發到OSPF 1。但是，一旦開啟將iBGP LU路由重新分發到LDP，字首10.100.1.99/32也會獲得本地標籤。

```
RP/0/0/CPU0:ABR1#show mpls ldp bindings 10.100.1.99/32

10.100.1.99/32, rev 24
        Local binding: label: 24007
        No remote bindings
```

9. MPLS啟用介面命令

示例1. IGP，但沒有LDP

如果有負責內部路由的IGP，但沒有用於通告標籤繫結的LDP，則需要使用**mpls activate**命令。如果每個躍點**執行BGP**，則BGP LU可用於通告字首和標籤。當它在鏈路上為**iBGP**時，需要使用**mpls activate**命令在路由器BGP下啟用該鏈路。請參閱此映像。

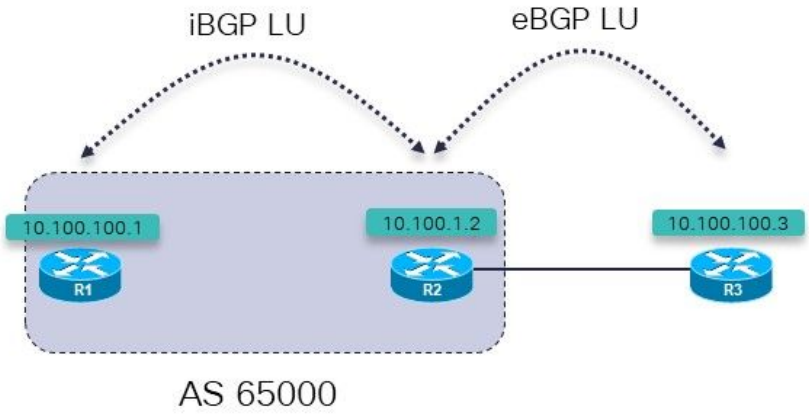

圖14.

R1和R2運行它們之間的IGP和iBGP LU。R1和R2直接連線。R2具有到R3的eBGP LU會話。

R3通過eBGP LU會話將字首10.100.100.3/2通告給R2。R2通過iBGP LU會話向R1通告此字首。

目標是有一個從R1到R3的不間斷LSP。它在嗎？

```
RP/0/0/CPU0:R1#trace 10.100.100.3 so 10.100.100.1

Type escape sequence to abort.
Tracing the route to 10.100.100.3

1 100.1.1 !N *  !N
```

第一跳處沒有此字首的標籤。

```
RP/0/0/CPU0:R1#traceroute mpls ipv4 10.100.100.3/32 ttl 5

Tracing MPLS Label Switched Path to 10.100.100.3/32, timeout is 2
seconds

Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
  'L' - labeled output interface, 'B' - unlabeled output interface,
  'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
  'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
  'P' - no rx intf label prot, 'p' - premature termination of LSP,
  'R' - transit router, 'I' - unknown upstream index,
  'X' - unknown return code, 'x' - return code 0

Type escape sequence to abort.

  0 0.0.0.0 MRU 0 [No Label]
Q 1 *
```

所以，沒有標籤。這並不令人意外，因為到R2的介面上未啟用MPLS:

```
RP/0/0/CPU0:R1#show mpls interfaces
RP/0/0/CPU0:R1#
```

但是，R3通告的LU字首存在於R1上：

```
RP/0/0/CPU0:R1#show bgp ipv4 labeled-unicast 10.100.100.3/32
BGP routing table entry for 10.100.100.3/32
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 7           7
    Local Label: 24001
Last Modified: Sep 13 14:27:17.510 for 00:11:39
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  65001
    10.100.1.2 (metric 2) from 10.100.1.2 (10.100.1.2)
      Received Label 24002
      Origin IGP, metric 0, localpref 100, valid, internal, best, group-
best, labeled-unicast
      Received Path ID 0, Local Path ID 1, version 7
```

您可以在R1上為通向R2的介面配置**mpls active**命令：

```
router bgp 65000
mpls activate
  interface GigabitEthernet0/0/0/0
 !
 address-family ipv4 unicast
  network 10.100.100.1/32
  allocate-label all
 !
 neighbor 10.100.1.2
  remote-as 65000
  update-source Loopback0
  address-family ipv4 labeled-unicast
  !
 !
!
```

現在，在傳出介面上啟用了MPLS。

```
RP/0/0/CPU0:R1#show mpls interfaces
Interface                  LDP      Tunnel   Static   Enabled
-------------------------- -------- -------- -------- --------
GigabitEthernet0/0/0/0     No       No       No       Yes
```

traceroute現在顯示LSP未中斷。

```
RP/0/0/CPU0:R1#trace 10.100.100.3 so 10.100.100.1

Type escape sequence to abort.
Tracing the route to 10.100.100.3

 1  10.1.2.2 [MPLS: Label 24002 Exp 0] 39 msec  9 msec  9 msec
 2 10.2.3.3 19 msec *  9 msec
```

```
RP/0/0/CPU0:R1#traceroute mpls ipv4 10.100.100.3/32 ttl 5 source 10.100.100.1

Tracing MPLS Label Switched Path to 10.100.100.3/32, timeout is 2
seconds

Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
  'L' - labeled output interface, 'B' - unlabeled output interface,
  'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
  'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx labl,
  'P' - no rx intf label prot, 'p' - premature termination of LSP,
  'R' - transit router, 'I' - unknown upstream index,
  'X' - unknown return code, 'x' - return code 0

Type escape sequence to abort.

  0 10.1.2.1 MRU 1500 [Labels: implicit-null/24002 Exp: 0/0]
L 1 10.1.2.2 MRU 1500 [Labels: implicit-null/implicit-null Exp: 0/0] 0
ms
! 2 10.2.3.3 10 ms
```

範例2.聯盟

此範例說明使用BGP LU(RFC 3107)而不使用LDP時，eBGP（AS間）聯盟連結上需要**mpls activate**命令。

此映像中的網路是包含子65000自治系統65501、65502、65503和介面的聯65504。

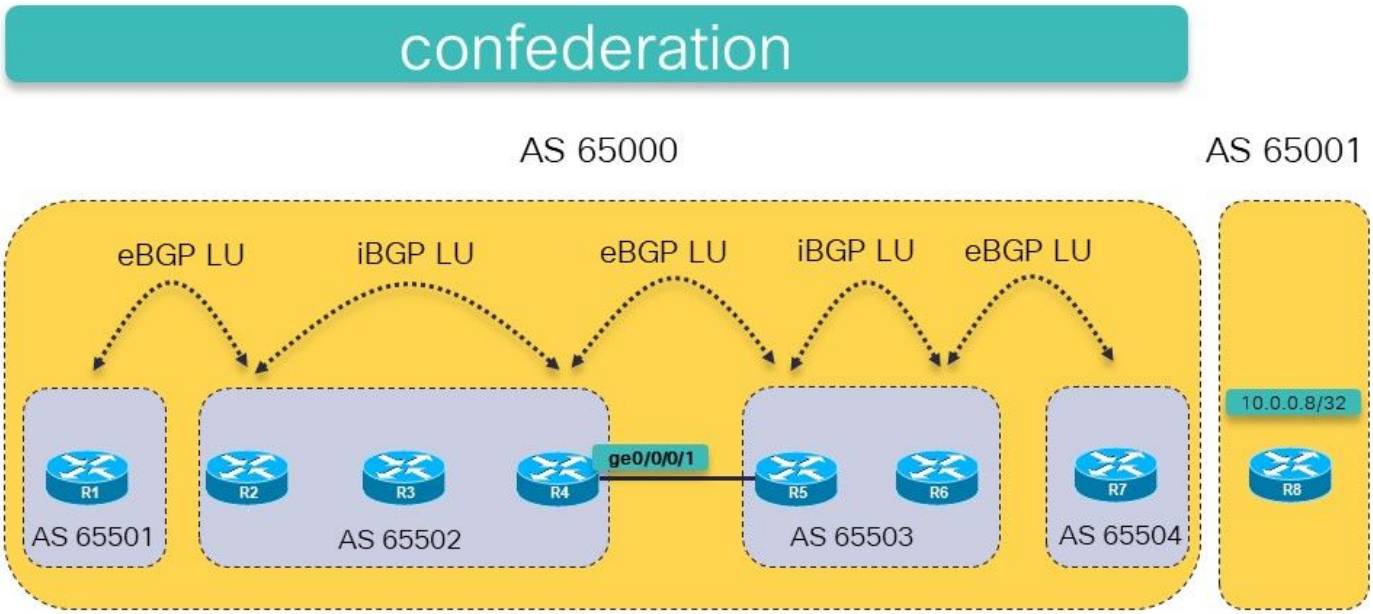

圖15.

其思想是通過在兩個自治系統中使用BGP LU，將一個MPLS LSP從R1到R8(10.0.0.8/32由R8在BGP LU中通告)。

R7和R8之間有常規的eBGP LU，R2和R4之間以及R5和R6之間有iBGP，R1和R2、R4和R5之間以及R6和R7之間有eBGP。每個eBGP會話上都有下一跳自身。

需要到達eBGP對等體的下一躍點的靜態路由（通常用於AS BGP間會話），因為聯盟內的子自治系統之間存在eBGP。

這是否足以實現R1和R8之間的連線？這意味著目標是有一個從R1到R8的不間斷LSP。

看看這個。

```
RP/0/0/CPU0:R1#traceroute 10.0.0.8

Type escape sequence to abort.
Tracing the route to 10.0.0.8

 1   *  *  *
 2   *  *  *
 3   *  *  *
 4   *  *  *
 5   *  *  *
```

traceroute不會返回任何躍點/標籤，且如果命令上未提供TTL限制，則會繼續。路由器可能應答traceroute，但資料包可能不會返回R1。執行mpls traceroute是一個更安全的選擇。

附註：只有在路徑沿途的每台路由器上啟用MPLS OAM時，MPLS路徑才起作用。

```
RP/0/0/CPU0:R1#trace mpls ipv4 10.0.0.8/32

Tracing MPLS Label Switched Path to 10.0.0.8/32, timeout is 2 seconds

Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
  'L' - labeled output interface, 'B' - unlabeled output interface,
  'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
  'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
  'P' - no rx intf label prot, 'p' - premature termination of LSP,
  'R' - transit router, 'I' - unknown upstream index,
  'X' - unknown return code, 'x' - return code 0

Type escape sequence to abort.

  0 10.1.2.1 MRU 1500 [Labels: implicit-null/24015 Exp: 0/0]
L 1 10.1.2.2 MRU 1500 [Labels: 24003/24014 Exp: 0/0] 10 ms
L 2 10.2.3.3 MRU 1500 [Labels: implicit-null/24014 Exp: 0/0] 10 ms
N 3 10.3.4.4 MRU 0 [No Label] 10 ms
```

您發現問題出在R4上。LFIB中缺少傳出介面：

```
RP/0/0/CPU0:R4#show mpls forwarding prefix 10.0.0.8/32
Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface
Switched
------ ----------- ------------------ ------------ --------------- -----
-------
24014  24014       10.0.0.8/32                     10.4.5.5        5140
```

未解析CEF中的條目：

```
RP/0/0/CPU0:R4#show cef 10.0.0.8/32
10.0.0.8/32, version 109, drop adjacency, internal 0x5000001 0x0 (ptr
0xa14160e4) [1], 0x0 (0xa13f83c8), 0xa08 (0xa16cd370)
```

```
 Updated Sep 13 12:43:30.252
 Prefix Len 32, traffic index 0, precedence n/a, priority 4
   via 10.4.5.5/32, 0 dependencies, recursive [flags 0x6000]
    path-idx 0 NHID 0x0 [0xa0f182d8 0x0]
    recursion-via-/32
    unresolved
     local label 24014
     labels imposed {24014}
```

GE0/0/0/1介面上未啟用MPLS:

```
RP/0/0/CPU0:R4#show mpls interfaces
Interface                  LDP      Tunnel   Static   Enabled
-------------------------- -------- -------- -------- --------
GigabitEthernet0/0/0/0     Yes      No       No       Yes
```

此問題可通過在R4和R5之間的鏈路上啟用BGP的MPLS命令得到解決。R4和R5在此鏈路上具有eBGP聯盟會話。65000實際上，這是聯盟域內的iBGP作業階段。因此，需要使用啟用MPLS的命令來確保R4上的首碼解析到下一個躍點R5。在其他常規網路中，會有LDP處理此作業，但在這裡，R4和R5之間沒有LDP，因為它是在聯盟域內的eBGP作業階段。

為R4上的介面ge 0/0/0/1新增**mpls activate**命令：

```
router bgp 65502
 bgp confederation peers
  65501
  65503
  65504
 !
 bgp confederation identifier 65000
 mpls activate
  interface GigabitEthernet0/0/0/1
 !
…

RP/0/0/CPU0:R4#show mpls interfaces
Interface                  LDP      Tunnel   Static   Enabled
-------------------------- -------- -------- -------- --------
GigabitEthernet0/0/0/0     Yes      No       No       Yes
GigabitEthernet0/0/0/1     No       No       No       Yes
```

現在traceroute顯示從R1到R8的不連續LSP。

```
RP/0/0/CPU0:R1#trace mpls ipv4 10.0.0.8/32

Tracing MPLS Label Switched Path to 10.0.0.8/32, timeout is 2 seconds

Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
  'L' - labeled output interface, 'B' - unlabeled output interface,
  'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
  'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label
  'P' - no rx intf label prot, 'p' - premature termination of LSP,
  'R' - transit router, 'I' - unknown upstream index,
  'X' - unknown return code, 'x' - return code 0

Type escape sequence to abort.

  0 10.1.2.1 MRU 1500 [Labels: implicit-null/24015 Exp: 0/0]
L 1 10.1.2.2 MRU 1500 [Labels: 24003/24014 Exp: 0/0] 10 ms
L 2 10.2.3.3 MRU 1500 [Labels: implicit-null/24014 Exp: 0/0] 10 ms
```

```
L 3 10.3.4.4 MRU 1500 [Labels: implicit-null/24014 Exp: 0/0] 10 ms
L 4 10.4.5.5 MRU 1500 [Labels: implicit-null/24014 Exp: 0/0] 20 ms
L 5 10.5.6.6 MRU 1500 [Labels: implicit-null/24014 Exp: 0/0] 30 ms
L 6 10.6.7.7 MRU 1500 [Labels: implicit-null/implicit-null Exp: 0/0] 30
ms
! 7 10.7.8.8 30 ms

RP/0/0/CPU0:R1#traceroute 10.0.0.8

Type escape sequence to abort.

Tracing the route to 10.0.0.8

 1  10.1.2.2 [MPLS: Label 24015 Exp 0] 69 msec  29 msec  29 msec
 2  10.2.3.3 [MPLS: Labels 24003/24014 Exp 0] 49 msec  29 msec  29 msec
 3  10.3.4.4 [MPLS: Label 24014 Exp 0] 19 msec  19 msec  19 msec
 4  10.4.5.5 [MPLS: Label 24014 Exp 0] 49 msec  19 msec  29 msec
 5  10.5.6.6 [MPLS: Label 24014 Exp 0] 19 msec  19 msec  29 msec
 6  10.6.7.7 [MPLS: Label 24014 Exp 0] 29 msec  29 msec  29 msec
 7  10.7.8.8 29 msec  *  29 msec
```

LFIB中現在有一個用於此條目的傳出介面：

```
RP/0/0/CPU0:R4#show mpls forwarding prefix 10.0.0.8/32

Local  Outgoing    Prefix            Outgoing     Next Hop        Bytes
Label  Label       or ID             Interface
Switched
------ ----------- ----------------- ------------ --------------- -----
-------
24014  24014       10.0.0.8/32       Gi0/0/0/1    10.4.5.5        2890
```

R4上存在用於字首的傳出標籤，並且CEF顯示解析的字首：

```
RP/0/0/CPU0:R4#show cef 10.0.0.8/32
Updated Sep 13 12:43:30.252
 Prefix Len 32, traffic index 0, precedence n/a, priority 4
   via 10.4.5.5/32, 3 dependencies, recursive [flags 0x6000]
    path-idx 0 NHID 0x0 [0xa17420e4 0x0]
   recursion-via-/32
   next hop 10.4.5.5/32 via 24016/0/21
     local label 24014
     next hop 10.4.5.5/32 Gi0/0/0/1    labels imposed {ImplNull 24014}
```