使用相同的路由區分器和「cef encap-sharing disabled」排除內聯BGP VPNv4 RR故障

目錄

簡介

本檔案介紹用作邊界閘道通訊協定(BGP)VPNv4內嵌路由反射器(RR)和提供商邊緣(PE)的思科系統網路聚合系統(NCS)540的行為。

背景資訊

本文檔重點介紹NCS 540的行為，如使用Cisco IOS® XR軟體版本7.3.1在實驗室環境中驗證的那樣。本文檔中描述的行為適用於所有基於NCS5500或NCS500系列DNX的平台和軟體版本。

問題

請考慮以下場景：NCS540配置了虛擬路由和轉發(VRF)例項，並使用與遠端PE RR客戶端節點所使用的路由區分器(RD)值相同的路由區分器(RD)值。當作為內聯RR和PE角色使用與遠端PE RR客戶端節點相同的RD值時，最上面的標籤在轉發之前不會彈出，從而導致在出口PE處丟包。

方案1.具有不同RD值的內聯RR和PE

該圖顯示了配置為BGP VPNv4內聯RR和PE的受測試裝置(DUT)的方案，其中VRF vrf1的RD值與其他PE節點使用的不同。

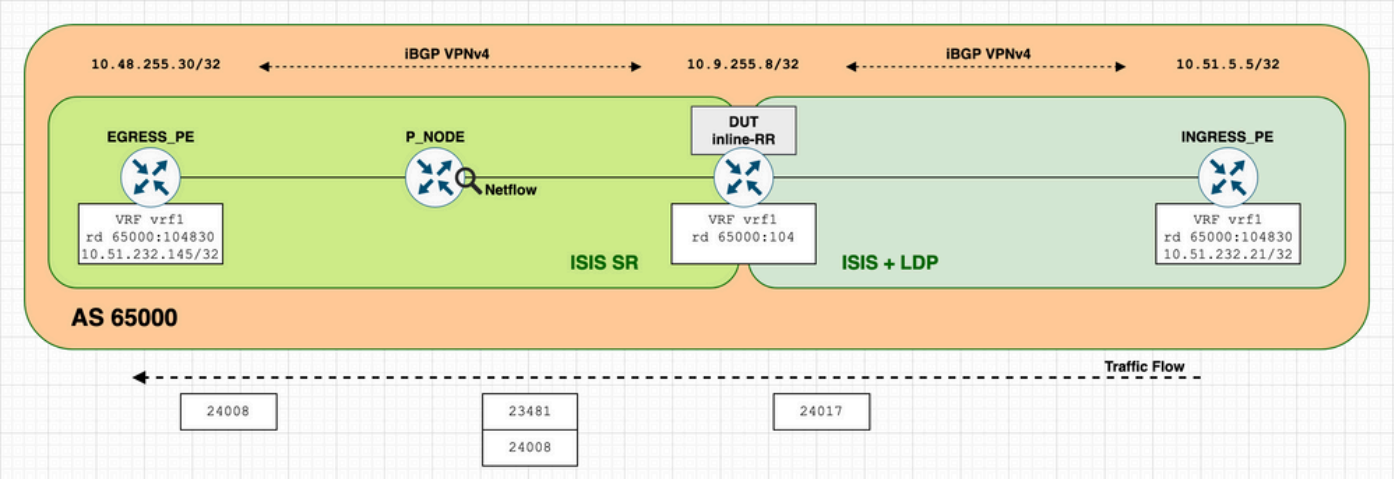

影象1 - DUT — 具有不同RD值的內聯RR和PE。

在VRF vrf1內的輸入PE(IP 10.51.232.21)和輸出PE(10.51.232.145)之間建立IP連線，並在PE節點之間成功轉發資料包，如Ping和Traceroute命令輸出所示：

```
INGRESS_PE#ping vrf vrf1 10.51.232.145 source 10.51.232.21
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.51.232.145, timeout is 2 seconds:
Packet sent with a source address of 10.51.232.21
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

INGRESS_PE#traceroute vrf vrf1 10.51.232.145 source 10.51.232.21
Type escape sequence to abort.
Tracing the route to 10.51.232.145
VRF info: (vrf in name/id, vrf out name/id)
 1 10.51.5.161 [MPLS: Label 24017 Exp 0] 4 msec 8 msec 0 msec
 2 10.9.255.207 [MPLS: Labels 23481/24008 Exp 0] 4 msec 0 msec 0 msec
 3 10.9.255.217 0 msec *  4 msec
```

DUT是本地配置了VRF的BGP VPNv4內聯RR和PE，但它使用的RD值(65000:104)與PE RR客戶端節點使用的RD值(65000:104830)不同。 如輸出所示，DUT匯入所有路由並交換VPNv4標籤：

```
RP/0/RP0/CPU0:DUT-N540#show bgp vpnv4 unicast
BGP router identifier 10.9.255.8, local AS number 65000
<snip>
  Network            Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 65000:104 (default for vrf vrf1)
*>i0.0.0.0/0          10.48.255.30             10    100      0 ?
*>i10.51.232.20/30    10.51.5.5                 0    100      0 ?
*>i10.51.232.145/32   10.48.255.30              0    100      0 ?
Route Distinguisher: 65000:104830
*>i0.0.0.0/0          10.48.255.30             10    100      0 ?
*>i10.51.232.20/30    10.51.5.5                 0    100      0 ?
*>i10.51.232.145/32   10.48.255.30              0    100      0 ?
<snip>

RP/0/RP0/CPU0:DUT-N540#show bgp vpnv4 unicast rd 65000:104 labels
BGP router identifier 10.9.255.8, local AS number 65000
<snip>
  Network            Next Hop        Rcvd Label      Local Label
Route Distinguisher: 21497:104 (default for vrf vrf1)
*>i0.0.0.0/0          10.48.255.30    24008           nolabel
*>i10.51.232.20/30    10.51.5.5       17              nolabel
*>i10.51.232.145/32   10.48.255.30    24008           nolabel
<snip>
```

```
RP/0/RP0/CPU0:DUT-N540#show bgp vpnv4 unicast rd 65000:104830 labels
BGP router identifier 10.9.255.8, local AS number 65000
<snip>
  Network            Next Hop        Rcvd Label      Local Label
Route Distinguisher: 21497:104830
*>i0.0.0.0/0          10.48.255.30    24008           24018
*>i10.51.232.20/30    10.51.5.5       17              24019
*>i10.51.232.145/32   10.48.255.30    24008           24017
<snip>
```

作為本場景中使用不同RD值的內RR和PE的附加參考，從EGRESS_PE節點接收的字首10.51.232.145/32的完整輸出如下所示：

```
RP/0/RP0/CPU0:DUT-N540#show bgp vpnv4 unicast rd 65000:104 10.51.232.145
BGP routing table entry for 10.51.232.145/32, Route Distinguisher: 65000:104
Versions:
 Process           bRIB/RIB  SendTblVer
 Speaker                115         115
Last Modified: Feb  8 11:00:27.032 for 2w6d
Paths: (1 available, best #1)
 Not advertised to any peer
 Path #1: Received by speaker 0
 Not advertised to any peer
 Local, (received & used)
   10.48.255.30 (metric 20) from 10.48.255.30 (10.48.255.30)
     Received Label 24008
     Origin incomplete, metric 0, localpref 100, valid, internal, best, group-best, import-
candidate, imported
     Received Path ID 1, Local Path ID 1, version 115
     Extended community: RT:65000:104830 RT:65000:105130
     Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 65000:104830

RP/0/RP0/CPU0:DUT-N540#show bgp vpnv4 unicast rd 65000:104830 10.51.232.145
BGP routing table entry for 10.51.232.145/32, Route Distinguisher: 65000:104830
Versions:
 Process           bRIB/RIB  SendTblVer
 Speaker                113         113
   Local Label: 24017
Last Modified: Feb  8 11:00:22.032 for 2w6d
Paths: (1 available, best #1)
 Advertised to peers (in unique update groups):
   10.51.5.5
 Path #1: Received by speaker 0
 Advertised to peers (in unique update groups):
   10.51.5.5
 Local, (received & used)
   10.48.255.30 (metric 20) from 10.48.255.30 (10.48.255.30)
     Received Label 24008
     Origin incomplete, metric 0, localpref 100, valid, internal, best, group-best, import-
candidate, not-in-vrf
     Received Path ID 1, Local Path ID 1, version 113
     Extended community: RT:65000:104830 RT:65000:10513
```

來源為INGRESS_PE(10.51.232.21)且目的地為EGRESS PE(10.51.232.145)的封包會進行標籤交換，且在DUT頂部的標籤{24017}上，根據程式設計的轉送詳細資訊，交換為{23481 24008}:

```
RP/0/RP0/CPU0:DUT-N540#show cef vrf vrf1 10.51.232.145/32 detail
10.51.232.145/32, version 96, internal 0x5000001 0x0 (ptr 0x8ce0d034) [1], 0x0 (0x8b941ee0),
0xa08 (0x8cacb5f8)
 Updated Feb 25 12:18:36.885
 Prefix Len 32, traffic index 0, precedence n/a, priority 3
```

```
 gateway array (0x8b7b6fd0) reference count 2, flags 0x38, source rib (7), 0 backups
               [3 type 1 flags 0x8441 (0x8cb11e28) ext 0x0 (0x0)]
 LW-LDI[type=1, refc=1, ptr=0x8b941ee0, sh-ldi=0x8cb11e28]
 gateway array update type-time 1 Feb 25 12:18:36.885
 LDI Update time Feb 25 12:18:36.885
 LW-LDI-TS Feb 25 12:18:36.885
  via 10.48.255.30/32, 7 dependencies, recursive [flags 0x6000]
   path-idx 0 NHID 0x0 [0x8d37e3b8 0x0]
   recursion-via-/32
   next hop VRF - 'default', table - 0xe0000000
   next hop 10.48.255.30/32 via 23481/0/21
    next hop 10.9.255.207/32 BE100        labels imposed {23481 24008}

   Load distribution: 0 (refcount 3)

   Hash  OK  Interface                 Address
   0     Y   recursive                 23481/0
```

在P-NODE上，如果將Netflow配置為匹配源自INGRESS_PE節點的輸入流量，則會觀察到預期的標籤堆疊{23481 24008}，如流監控器輸出所示：

```
RP/0/RP0/CPU0:P_NODE#show flow monitor MONITOR_MAP_MPLS cache location 0/RP0/CPU0
<snip>
LabelType Prefix/Length      Label1-EXP-S     Label2-EXP-S Label3-EXP-S    InputInterface
OutputInterface ForwardStatus        FirstSwitched    LastSwitched     ByteCount     PacketCount
Dir SamplerID  IPV4SrcAddr      IPV4DstAddr      IPV4TOS  IPV4Prot L4SrcPort  L4DestPort
L4TCPFlags    InputVRFID                        OutputVRFID
BGPNextHopV4
Unknown 10.48.255.30/32       23481-0-0        24008-0-1           -
BE100           BE1             Fwd                  33 17:49:08:468 33 17:49:11:765
108000       1000         Ing 1          10.51.232.21     10.51.232.145    0         icmp
0          0          0            default
default                           0.0.0.0
<snip>
```

場景2.內聯RR和PE具有相同的RD值

該圖顯示了問題場景：DUT配置為BGP VPNv4內聯RR和PE，但現在VRF vrf1的RD值配置與其他PE節點相同 — 65000:10430。

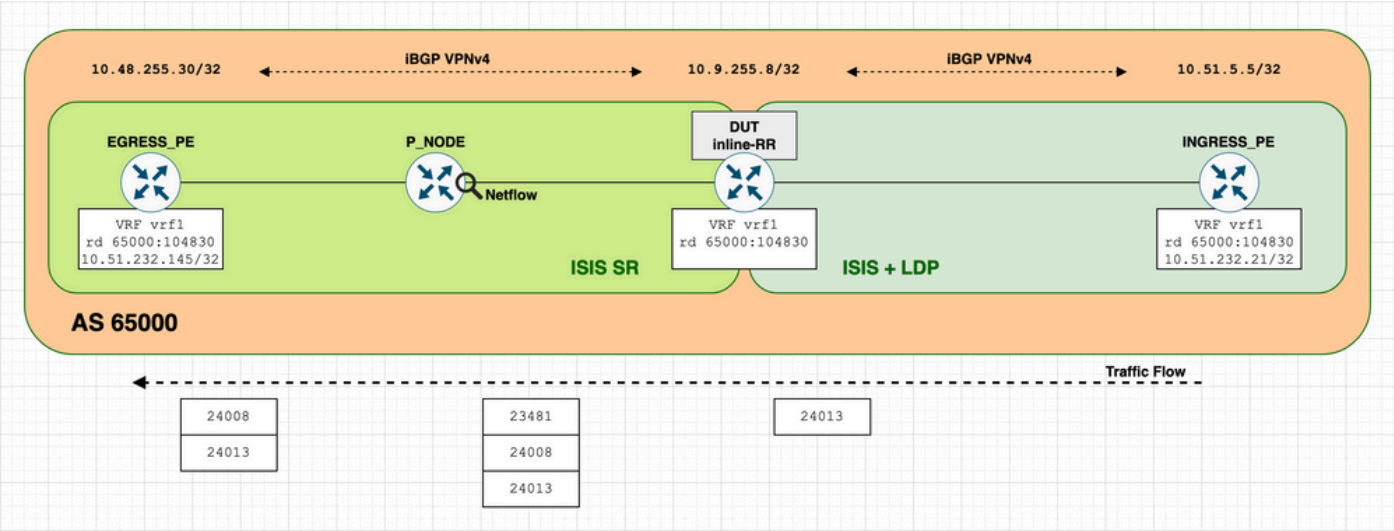

影象2 — 冗餘 — 內聯RR和PE具有相同的RD值。

在此案例中，VRF vrf1中的輸入PE(IP 10.51.232.21)和輸出PE(10.51.232.145)之間的IP連線失敗，如使用Ping和Traceroute指令輸出所示：

```
INGRESS_PE#ping vrf vrf1 10.51.232.145 source 10.51.232.21
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.51.232.145, timeout is 2 seconds:
Packet sent with a source address of 10.51.232.21
.....
Success rate is 0 percent (0/5)

INGRESS_PE#traceroute vrf vrf1 10.51.232.145 source 10.51.232.21
Type escape sequence to abort.
Tracing the route to 10.51.232.145
VRF info: (vrf in name/id, vrf out name/id)
 1 10.51.5.161 [MPLS: Label 24013 Exp 0] 4 msec 4 msec 0 msec
 2  *  *  *
<snip>
```

在DUT時，從BGP或程式設計的轉發輸出中均未明確發現問題根源，且所有輸出均被視為預期輸出：

```
RP/0/RP0/CPU0:DUT-N540#show bgp vpnv4 unicast
BGP router identifier 10.9.255.8, local AS number 65000
<snip>
  Network            Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 65000:104830 (default for vrf vrf1)
*>i0.0.0.0/0          10.48.255.30            10    100      0 ?
*>i10.51.232.20/30    10.51.5.5                0    100      0 ?
*>i10.51.232.145/32   10.48.255.30             0    100      0 ?
<snip>

RP/0/RP0/CPU0:DUT-N540#show bgp vpnv4 unicast rd 65000:104830 labels
BGP router identifier 10.9.255.8, local AS number 65000
<snip>
  Network            Next Hop        Rcvd Label      Local Label
Route Distinguisher: 21497:104830 (default for vrf vrf1)
*>i0.0.0.0/0          10.48.255.30    24008           24020
*>i10.51.232.20/30    10.51.5.5       17              24016
*>i10.51.232.145/32   10.48.255.30    24008           24013
<snip>
```

與上一節類似，並且作為使用相同RD值的內RR和PE的當前方案中的附加參考，顯示從EGRESS_PE節點接收的字首10.51.232.145/32的完整輸出：

```
RP/0/RP0/CPU0:DUT-N540#show bgp vpnv4 unicast rd 65000:104830 10.51.232.145
BGP routing table entry for 10.51.232.145/32, Route Distinguisher: 65000:104830
Versions:
 Process           bRIB/RIB  SendTblVer
 Speaker               134         134
   Local Label: 24013
Last Modified: Feb 28 18:03:20.032 for 00:04:50
Paths: (1 available, best #1)
 Advertised to peers (in unique update groups):
   10.51.5.5
 Path #1: Received by speaker 0
 Advertised to peers (in unique update groups):
   10.51.5.5
 Local, (received & used)
   10.48.255.30 (metric 20) from 10.48.255.30 (10.48.255.30)
     Received Label 24008
     Origin incomplete, metric 0, localpref 100, valid, internal, best, group-best, import-
candidate, imported
     Received Path ID 1, Local Path ID 1, version 134
     Extended community: RT:65000:104830 RT:65000:105130
```

```
    Source AFI: VPNv4 Unicast, Source VRF: vrf1, Source Route Distinguisher: 65000:10483
```

來源為INGRESS_PE(10.51.232.21)且目的地為EGRESS PE(10.51.232.145)的封包會進行標籤交換，且預期在DUT時，這些封包會根據程式化的轉送詳細資訊，將其頂部標籤{24013}交換為{23481 24008}:

```
RP/0/RP0/CPU0:DUT-N540#show cef vrf vrf1 10.51.232.145/32 detail
10.51.232.145/32, version 107, internal 0x1000001 0x0 (ptr 0x8ce0d13c) [1], 0x0 (0x8b946be8),
0xa08 (0x8cacb7d8)
 Updated Feb 28 18:03:19.778
 Prefix Len 32, traffic index 0, precedence n/a, priority 3
 gateway array (0x8b7b71a0) reference count 6, flags 0x78, source rib (7), 0 backups
               [3 type 5 flags 0x8441 (0x8cb125d8) ext 0x0 (0x0)]
 LW-LDI[type=5, refc=3, ptr=0x8b946be8, sh-ldi=0x8cb125d8]
 gateway array update type-time 1 Feb 28 18:03:19.778
 LDI Update time Feb 28 18:03:19.778
 LW-LDI-TS Feb 28 18:03:19.778
  via 10.48.255.30/32, 7 dependencies, recursive [flags 0x6000]
   path-idx 0 NHID 0x0 [0x8d37e3b8 0x0]
   recursion-via-/32
   next hop VRF - 'default', table - 0xe0000000
   next hop 10.48.255.30/32 via 23481/0/21
    local label 24013
    next hop 10.9.255.207/32 BE100        labels imposed {23481 24008}

   Load distribution: 0 (refcount 3)

   Hash  OK  Interface                 Address
   0     Y   recursive                 23481/0
```

此外，在P_NODE進行分類時，如果將Netflow配置為匹配源自INGRESS_PE的輸入流量，則會發現DUT如何轉發資料包的問題來源。如圖所示輸出中突出顯示，在轉發24013包之前，DUT不會彈出本地標簽欄位。因此，P_NODE接收三標籤堆疊MPLS幀{23481 24008 24013}，而不是預期的雙標籤堆疊{23481 24008}。在P_NODE，剝離頂部標籤23481，並將轉發到EGRESS_PE的資料包作為標籤堆疊{24008 24013}，這導致在EGRESS_PE處丟棄資料包。

```
RP/0/RP0/CPU0:P_NODE#show flow monitor MONITOR_MAP_MPLS cache location 0/RP0/CPU0
<snip>
LabelType Prefix/Length      Label1-EXP-S     Label2-EXP-S     Label3-EXP-S     InputInterface
OutputInterface ForwardStatus        FirstSwitched   LastSwitched    ByteCount    PacketCount
Dir SamplerID  IPV4SrcAddr      IPV4DstAddr      IPV4TOS  IPV4Prot L4SrcPort  L4DestPort
L4TCPFlags   InputVRFID                         OutputVRFID
BGPNextHopV4
Unknown 10.48.255.30/32       23481-0-0        24008-0-0        24013-0-1       BE100
BE1             Fwd                  33 17:51:40:181 33 17:51:41:521 112000       1000
Ing 1          10.51.232.21     10.51.232.145    0        icmp     0          0
0            default                            default
0.0.0.0
<snip>
```

案例3.已設定具有相同RD值和「cef encap-sharing disable」的串聯RR和PE

在前一個場景之上，配置cef encap-sharing disable，並為每個字首分配單獨的硬體資源。

```
RP/0/RP0/CPU0:DUT-N540(config)#cef encap-sharing disable
Warning: The command will clear the forwarding table.Traffic loss is expected during rebuilding.
RP/0/RP0/CPU0:DUT-N540(config)#commit
```

該影象顯示了使用同一RD配置並禁用cef encap-sharing的內聯RR和PE的場景。

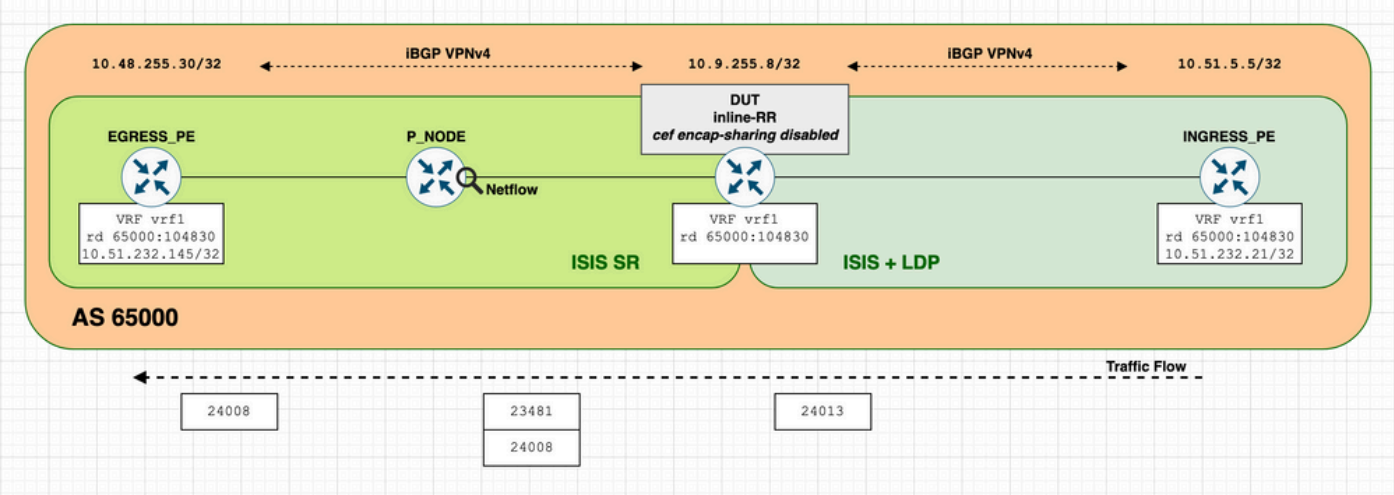

映像3 - DUT — 配置了同一RD值和「cef encap-sharing disable」的內聯RR和PE。

在VRF vrf1內的輸入PE(IP 10.51.232.21)和輸出PE(10.51.232.145)之間的IP連線會復原，而封包會再次在PE節點之間成功轉送，如Ping和Traceroute指令輸出所示：

```
INGRESS_PE#ping vrf vrf1 10.51.232.145 source 10.51.232.21 repeat 1000
Type escape sequence to abort.
Sending 1000, 100-byte ICMP Echos to 10.51.232.145, timeout is 2 seconds:
Packet sent with a source address of 10.51.232.21
<snip>
Success rate is 100 percent (1000/1000), round-trip min/avg/max = 1/3/12 ms

INGRESS_PE#traceroute vrf vrf1 10.51.232.145 source 10.51.232.21
Type escape sequence to abort.
Tracing the route to 10.51.232.145
VRF info: (vrf in name/id, vrf out name/id)
 1 10.51.5.161 [MPLS: Label 24013 Exp 0] 0 msec 0 msec 0 msec
 2 10.9.255.207 [MPLS: Labels 23481/24008 Exp 0] 4 msec 4 msec 0 msec
 3 10.9.255.217 0 msec *  0 msec
```

在P_NODE，Netflow輸出顯示，DUT現在轉發標籤堆疊預期為{23481 24008}的資料包。P_NODE去除頂部標籤23481，並將資料包轉發到標籤堆疊為{24008}的EGRESS_PE節點，EGRESS_PE可以匹配該節點，並將封裝的資料包成功轉發到VRF1中的最終目標。

```
RP/0/RP0/CPU0:P_NODE#show flow monitor MONITOR_MAP_MPLS cache location 0/RP0/CPU0
Cache summary for Flow Monitor MONITOR_MAP_MPLS:
<snip>
LabelType Prefix/Length      Label1-EXP-S     Label2-EXP-S     Label3-EXP-S     InputInterface
OutputInterface ForwardStatus        FirstSwitched   LastSwitched    ByteCount    PacketCount
Dir SamplerID  IPV4SrcAddr      IPV4DstAddr      IPV4TOS  IPV4Prot L4SrcPort  L4DestPort
L4TCPFlags   InputVRFID                       OutputVRFID
BGPNextHopV4
Unknown 10.48.255.30/32       23481-0-0        24008-0-1           -           BE100
BE1             Fwd                  33 18:03:14:211 33 18:03:17:505 108000       1000
Ing 1          10.51.232.21     10.51.232.145    0        icmp     0          0
0            default                       default
0.0.0.0
<snip>
```

使用cef encap-sharing disable配置時，會為每個字首分配額外的硬體資源，並且在DUT處有所需的轉發資訊，以便在此具有相同RD值的內聯RR的特定場景中正確轉發資料包。為了突出有關程式設計轉發詳細資訊的區別，請參閱show cef vrf vrf1 10.51.232.145/32 hardware egress location 0/RP0/CPU0所示的兩個輸出，並注意在配置cef encap-sharing disable時包含的其他資訊。

如在內聯RR和PE的場景中所見，它們具有相同的RD值，但未配置「cef encap-sharing disable」

（預設）：

```
!
! --- without 'cef encap-sharing disable' (default)
! --- note highlighted (bold) sections
!
RP/0/RP0/CPU0:DUT-N540#show cef vrf vrf1 10.51.232.145/32 hardware egress location 0/RP0/CPU0
10.51.232.145/32, version 107, internal 0x1000001 0x0 (ptr 0x8ce0d13c) [1], 0x0 (0x8b946be8),
0xa08 (0x8cacb7d8)
 Updated Feb 28 18:03:19.778
 Prefix Len 32, traffic index 0, precedence n/a, priority 3
 gateway array (0x8b7b71a0) reference count 6, flags 0x78, source rib (7), 0 backups
               [3 type 5 flags 0x8441 (0x8cb125d8) ext 0x0 (0x0)]
 LW-LDI[type=5, refc=3, ptr=0x8b946be8, sh-ldi=0x8cb125d8]
 gateway array update type-time 1 Feb 28 18:03:19.778
 LDI Update time Feb 28 18:03:19.778
 LW-LDI-TS Feb 28 18:03:19.778
  via 10.48.255.30/32, 7 dependencies, recursive [flags 0x6000]
   path-idx 0 NHID 0x0 [0x8d37e3b8 0x0]
   recursion-via-/32
   next hop VRF - 'default', table - 0xe0000000
   next hop 10.48.255.30/32 via 23481/0/21
    local label 24013
    next hop 10.9.255.207/32 BE100        labels imposed {23481 24008}

 Show-data Print at RPLC

 LEAF - HAL pd context :
 sub-type : IPV4, ecd_marked:0, has_collapsed_ldi:0
 collapse_bwalk_required:0, ecdv2_marked:0,
HW Walk:
LEAF:
   PI:0x8ce0d13c PD:0x8ce0d1dc rev:892768 type: IPV4 (0)
   LEAF location: LEM
   FEC key: 0x57f40001104

   LWLDI:
        PI:0x8b946be8 PD:0x8b946c28 rev:892767 p-rev:892766  ldi type:IMP_EOS0_EOS1
        FEC key: 0x57f40001104 fec index: 0x0(0) num paths:1, bkup paths: 0

 REC-SHLDI HAL PD context :
ecd_marked:0, collapse_bwalk_required:0, load_shared_lb:0

   RSHLDI:
        PI:0x8cb125d8 PD:0x8cb126a8 rev:892766 dpa-rev:41494702 flag:0x1
        FEC key: 0x57f40001104 fec index: 0x2000ffcc(65484) num paths: 1
        p-rev:854950
        Path:0 fec index: 0x2000ffcc(65484) DSP fec index: 0x2000ffca(65482)
               MPLS EEI push label: 24008

 LEAF - HAL pd context :
 sub-type : MPLS, ecd_marked:0, has_collapsed_ldi:0
 collapse_bwalk_required:0, ecdv2_marked:0,
HW Walk:
LEAF:
   PI:0x8d37e3b8 PD:0x8d37e458 rev:854953 type: MPLS (2)
   LEAF location: LEM
   FEC key: 0

   LWLDI:
       PI:0x8b945288 PD:0x8b9452c8 rev:854950 p-rev:854949  ldi type:IMP_EOS0_EOS1
       FEC key: 0x51140001104 fec index: 0x0(0) num paths:1, bkup paths: 0
```

```
       IMP LDI:
       IMP pattern:3
       PI:0x8b945288 PD:0x8b9452c8 rev:854950 p-rev:854949
       FEC key: 0x51240001104 fec index: 0x2000ffca(65482) num paths:1
       Path:0  fec index: 0x2000ffca(65482) DSP:0xc000001
               MPLS encap key: 0xf1b0000040014822 MPLS encap id: 0x40014822 Remote: 0

       SHLDI:
           PI:0x8cb10718 PD:0x8cb107e8 rev:854949 dpa-rev:39755988 flag:0x0
           FEC key: 0x51140001104 fec index: 0x2000ffcb(65483) num paths: 1 bkup paths: 0
           p-rev:72522
           Path:0 fec index: 0x2000ffcb(65483) DSP:0xc000001 Dest fec index: 0x0(0)

           TX-NHINFO:
               PI: 0x8d11fad0 PD: 0x8d11fb50 rev:72522 dpa-rev:3303803 Encap hdl: 0x8cd16098
               Encap id: 0x40010003 Remote: 0 L3 int: 1579 flags: 0x407
               npu_mask: 0x1 DMAC: 5c:5a:c7:ff:78:84

   Load distribution: 0 (refcount 3)

   Hash  OK  Interface                 Address
   0     Y   recursive                 23481/0
```

如內聯RR和**PE配置了相同的RD值和cef封鎖共用禁用情形所**示：

```
!
! --- with 'cef encap-sharing disable'
! --- note highlighted (bold) sections for the extra and additional forwarding information
included
!
RP/0/RP0/CPU0:DUT-N540#show cef vrf vrf1 10.51.232.145/32 hardware egress location 0/RP0/CPU0
10.51.232.145/32, version 127, internal 0x1000001 0x0 (ptr 0x8ce0ffa4) [1], 0x0 (0x8b948630),
0xa08 (0x8cacb5f8)
 Updated Feb 28 18:26:25.775
 Prefix Len 32, traffic index 0, precedence n/a, priority 3
 gateway array (0x8b7b5f80) reference count 3, flags 0x78, source rib (7), 0 backups
               [2 type 5 flags 0x8441 (0x8cb14c48) ext 0x0 (0x0)]
 LW-LDI[type=5, refc=3, ptr=0x8b948630, sh-ldi=0x8cb14c48]
 gateway array update type-time 1 Feb 28 18:26:25.775
 LDI Update time Feb 28 18:26:25.775
 LW-LDI-TS Feb 28 18:26:25.779
  via 10.48.255.30/32, 11 dependencies, recursive [flags 0x6000]
   path-idx 0 NHID 0x0 [0x8d37e3b8 0x0]
   recursion-via-/32
   next hop VRF - 'default', table - 0xe0000000
   next hop 10.48.255.30/32 via 23481/0/21
    local label 24013
    next hop 10.9.255.207/32 BE100        labels imposed {23481 24008}

 Show-data Print at RPLC

 LEAF - HAL pd context :
 sub-type : IPV4, ecd_marked:0, has_collapsed_ldi:0
 collapse_bwalk_required:0, ecdv2_marked:0,
HW Walk:
LEAF:
   PI:0x8ce0ffa4 PD:0x8ce10044 rev:893768 type: IPV4 (0)
   LEAF location: LEM
   FEC key: 0x5ae40001104

   LWLDI:

        LSP pattern:3
```

```
        PI:0x8b948630 PD:0x8b948670 rev:893767 p-rev:893766  ldi type:IMP_EOS0_EOS1
        FEC key: 0x5af40001104 fec index: 0x2000ffbf(65471) num paths:1, bkup paths: 0
        Path:0  fec index: 0x2000ffbf(65471) DSP fec index:0x20000001(1)
                MPLS encap key: 0xf1b000004001482f MPLS encap id: 0x4001482f Remote: 0
        IMP LDI:
        IMP pattern:3
        PI:0x8b948630 PD:0x8b948670 rev:893767 p-rev:893766
        FEC key: 0x5ae40001104 fec index: 0x2000ffc0(65472) num paths:1
        Path:0  fec index: 0x2000ffc0(65472) DSP fec index: 0x20000001(1)
                MPLS encap key: 0xf1b000004001482e MPLS encap id: 0x4001482e Remote: 0

 REC-SHLDI HAL PD context :
ecd_marked:0, collapse_bwalk_required:0, load_shared_lb:0

   RSHLDI:
        PI:0x8cb14c48 PD:0x8cb14d18 rev:893766 dpa-rev:41503635 flag:0x1
        FEC key: 0x5ac40001104 fec index: 0x2000ffe0(65504) num paths: 1
        p-rev:893704
        Indirection ECMP FEC key: 0x5ad20001104 fec index: 0x20000001(1)
        Path:0 fec index: 0x2000ffe0(65504) DSP fec index: 0x2000ffca(65482)

 LEAF - HAL pd context :
 sub-type : MPLS, ecd_marked:0, has_collapsed_ldi:0
 collapse_bwalk_required:0, ecdv2_marked:0,
HW Walk:
LEAF:
   PI:0x8d37e3b8 PD:0x8d37e458 rev:893707 type: MPLS (2)
   LEAF location: LEM
   FEC key: 0

   LWLDI:
       PI:0x8b9451a0 PD:0x8b9451e0 rev:893704 p-rev:893703  ldi type:IMP_EOS0_EOS1
       FEC key: 0x59f40001104 fec index: 0x0(0) num paths:1, bkup paths: 0
       IMP LDI:
       IMP pattern:3
       PI:0x8b9451a0 PD:0x8b9451e0 rev:893704 p-rev:893703
       FEC key: 0x5a040001104 fec index: 0x2000ffca(65482) num paths:1
       Path:0  fec index: 0x2000ffca(65482) DSP:0xc000001
               MPLS encap key: 0xf1b0000040014822 MPLS encap id: 0x40014822 Remote: 0

       SHLDI:
           PI:0x8cb112a0 PD:0x8cb11370 rev:893703 dpa-rev:41503599 flag:0x0
           FEC key: 0x59f40001104 fec index: 0x2000ffcb(65483) num paths: 1 bkup paths: 0
           p-rev:72522
           Path:0 fec index: 0x2000ffcb(65483) DSP:0xc000001 Dest fec index: 0x0(0)

           TX-NHINFO:
               PI: 0x8d11fad0 PD: 0x8d11fb50 rev:72522 dpa-rev:3303803 Encap hdl: 0x8cd16098
               Encap id: 0x40010003 Remote: 0 L3 int: 1579 flags: 0x407
               npu_mask: 0x1 DMAC: 5c:5a:c7:ff:78:84

   Load distribution: 0 (refcount 2)

   Hash  OK  Interface                Address
   0     Y   recursive                23481/0
```

解決方案

如本文檔中所述，對於具有相同路由回報值的內聯RR未彈出的最頂端標籤的不當標籤操作，解決方案是配置**cef encap-sharing disable**。此配置在此特定方案中是強制性的，強制將單獨的硬體資源分配給每個字首，以確保在內聯RR節點上發生正確的標籤操作和轉發。

提交配置之前，必須評估可用資源使用情況，以便在提交命令後預測資源狀況。要驗證和確認實際資源消耗，可以使用以下命令：

```
show controllers npu resources all location all
show controllers fia diagshell 0 "diag alloc all" location all
```

附註：使用思科錯誤ID CSCvw20873 - L3VPN LSP路徑（標籤交換）最佳化 — 引入對已使用和已分配資源的最佳化，並配置了**cef encap-sharing** disable。