排除同一路由区分器和“cef encap-sharing disabled”的串行BGP VPNv4 RR故障

目录

简介

本文档介绍用作边界网关协议(BGP)VPNv4串接路由反射器(RR)和提供商边缘(PE)时思科系统网络融合系统(NCS)540的行为。

背景信息

本文档重点介绍在Cisco IOS® XR软件版本7.3.1的实验室环境中验证的NCS 540行为。本文档中描述的行为适用于所有基于NCS5500或NCS500系列DNX的平台和软件版本。

问题

考虑NCS540配置了虚拟路由和转发(VRF)实例并使用与远程PE RR客户端节点使用的相同路由区分器(RD)值的场景。当用作串行RR和PE角色并配置了与远程PE RR客户端节点相同的RD值时，转发前不会弹出最顶部的标签，从而导致出口PE丢包。

场景1.具有不同RD值的串联RR和PE

该图显示了配置为BGP VPNv4串行RR和PE的“测试设备(DUT)”的场景，其RD值与其他PE节点使用的不同。

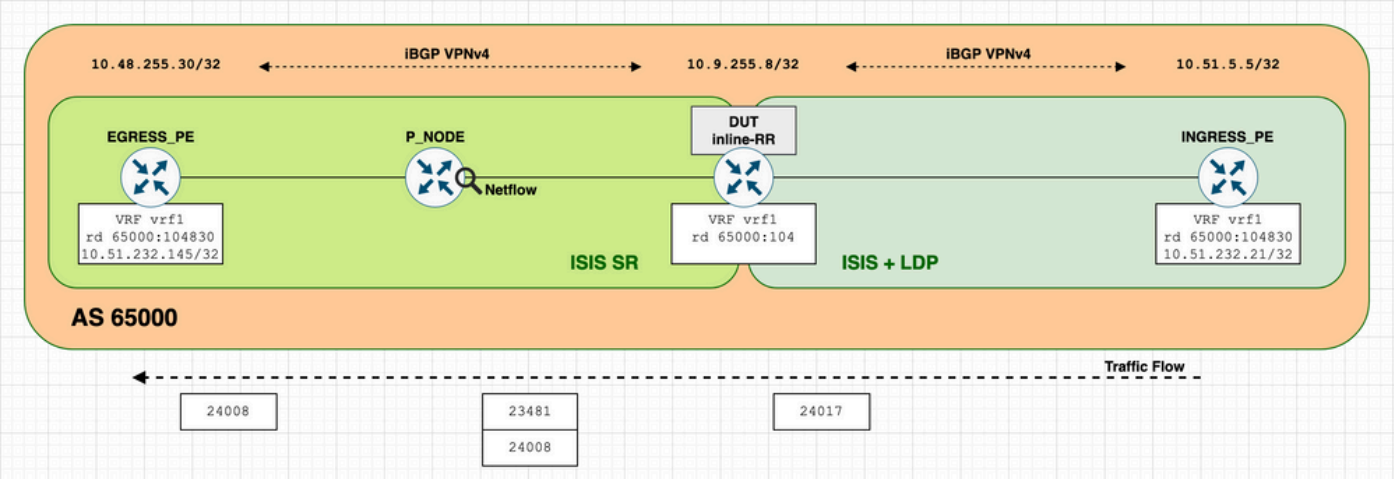

图1 - DUT — 行内RR和PE，RD值不同。

在VRF vrf1内，入口PE(IP 10.51.232.21)和出口PE(10.51.232.145)之间建立IP连接，数据包在PE节点之间成功转发，如Ping和Traceroute命令输出所示：

```
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串联RR和PE，但使用的RD值(65000:104)与PE RR客户端节点(65000:104830)不同。 如输出所示，DUT导入所有路由并相应地交换VPNv4标签：

```
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>
```

作为使用不同**RD值的串行RR和PE的当前场景的附加参考**，从EGRESS_PE节点收到的前缀10.51.232.145/32的完整输出如下所示：

```
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
```

源自INGRESS_PE(10.51.232.21)且发往EGRESS PE(10.51.232.145)的数据包是标签交换的，在DUT顶部标签{24017}交换的数据包是{23481 24008}，根据编程的转发详细信息：

```
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，Netflow配置为匹配来自INGRESS_PE节点的入口流量，将观察预期的标签堆栈{23481 24008}，如流监控器输出所示：

```
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>
```

场景2.具有相同RD值的串联RR和PE

该图显示了DUT配置为BGP VPNv4串联RR和PE，但现在为VRF vrf1 - 65000:10430配置了与其他PE节点相同的RD值的问题场景。

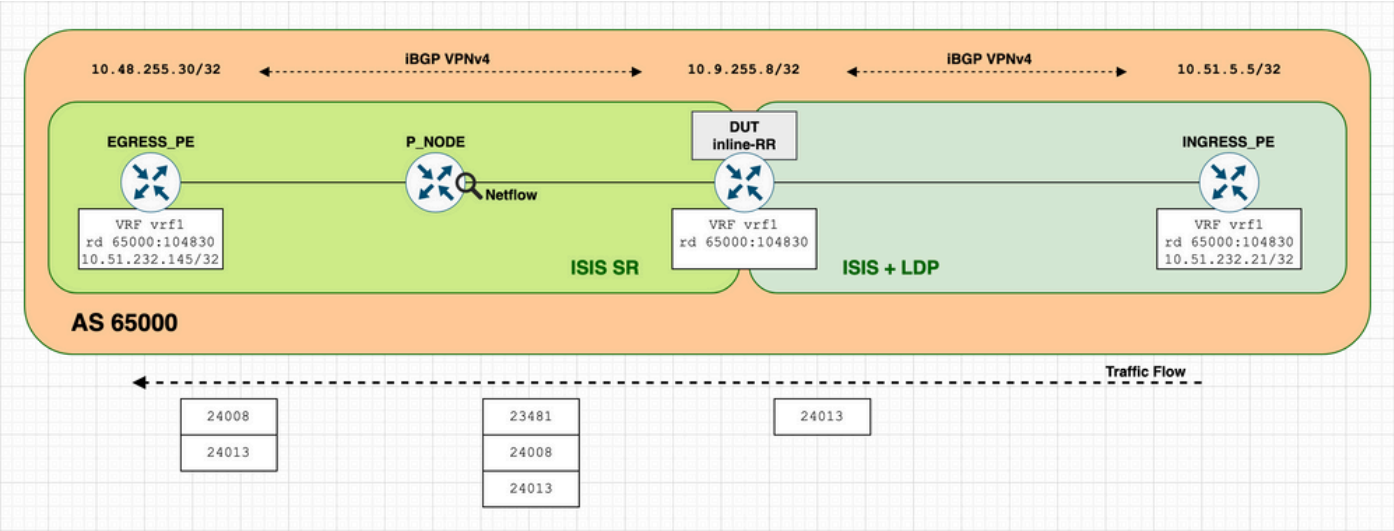

图2 - DUT — 内联RR和PE，具有相同RD值。

在此场景中，VRF vrf1内入口PE(IP 10.51.232.21)和出口PE(10.51.232.145)之间的IP连接失败，如Ping和Traceroute命令输出所示：

```
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中，从BGP或编程转发输出中无法清楚地了解问题根源，所有输出都被视为预期：

```
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>
```

与上一节类似，作为使用相同RD值的**串行RR和PE的当前场景的附加参考**，显示从EGRESS_PE节点收到的前缀10.51.232.145/32的完整输出：

```
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
```

源自INGRESS_PE(10.51.232.21)并发往EGRESS PE(10.51.232.145)的数据包被标签交换，预期在DUT，这些数据包的顶标签{24013}会根据编程转发详细信息由{23481}交换：

```
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，P_NODE上的Netflow分类配置为与来自INGRESS_PE的入口流量匹配，这揭示了DUT如何转发数据包的问题根源。如图所示输出中突出显示的，在转发数据包之前，本地标签24013不会被DUT弹出。因此，P_NODE接收一个三标签堆叠MPLS帧{23481 24008 24013}，而不是预期的双标签堆叠{23481 24008}。在P_NODE上标签23481被剥离，转发到EGRESS_PE的数据包作为标签栈{24008 24013}携带，从而导致在EGRESS_PE上丢弃数据包。

```
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>
```

场景3.配置了相同RD值和“cef encap-sharing disable”的串联RR和PE

在上一个方案之上，配置了cef encap-sharing disable，并为每个前缀分配单独的硬件资源。

```
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
```

图像显示了配置了相同RD并配置了禁用cef封装共享的串行RR和PE的场景。

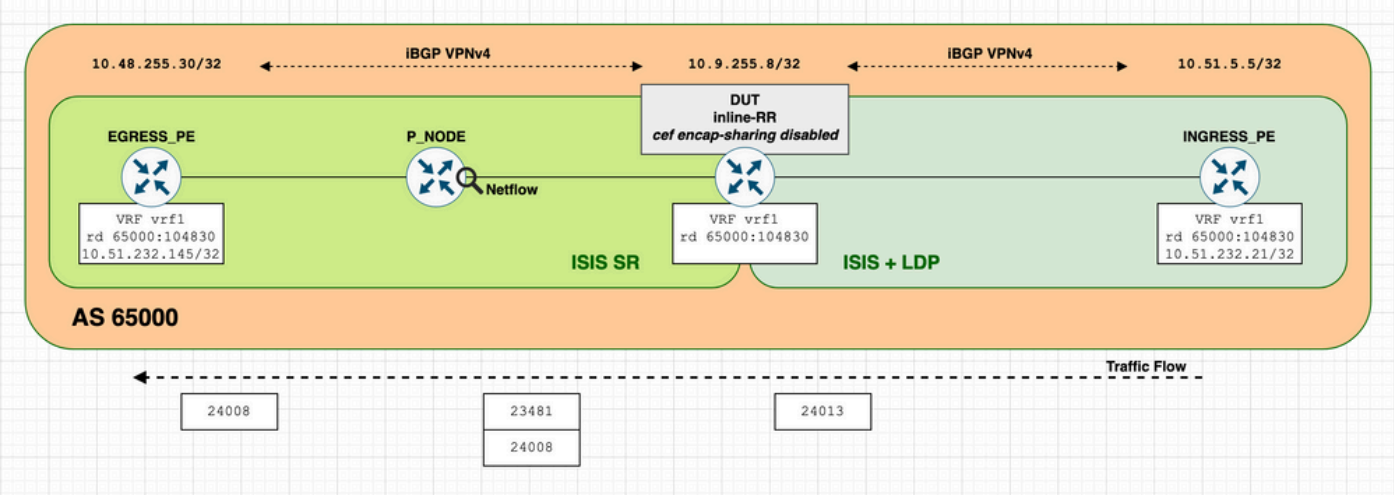

映像3 - DUT — 配置了相同RD值和"cef encap-sharing disable"的串行RR和PE。

VRF vrf1内的入口PE(IP 10.51.232.21)和出口PE(10.51.232.145)之间的IP连接会通过数据包恢复，一旦数据包在PE节点之间成功转发，如Ping和Traceroute命令输出所示：

```
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上，Netflow输出显示，DUT现在转发具有预期标签堆栈{23481 24008}的数据包。P_NODE删除顶部标签23481，并将数据包转发到标签栈{24008}的EGRESS_PE节点，EGRESS_PE能够匹配并成功将封装的数据包转发到VRF1内的最终目标。

```
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**配置时，会为每个前缀分配额外的硬件资源，并在DUT处设置所需的转发信息，以便在具有相同RD值的串行RR的这一特定场景中正确转发数据包。为了强调编程转发详细信息的区别，请参阅**show cef vrf vrf1 10.51.232.145/32硬件出口位置0/RP0/CPU0**的两个输出，并注意配置**cef encap-sharing disable时包含的其他信息。**

如在RD值相**同但未配置"cef encap-sharing disable"的串行RR和PE的场**景中**所示**（默认）：

```
!
! --- 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
```

如配置了相同RD值和CEF封装共享禁用的串行RR和PE的场景所示：

```
!
! --- 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
```

解决方案

如本文档中所述，不由具有相同RD值的行内RR弹出的最顶层标签的不正确标签操作的解决方案是配置**cef encap-sharing disable**。此配置在此特定方案中是强制性的，并强制为每个前缀分配单独的硬件资源，以确保在行内RR节点执行正确的标签操作和转发。

在提交配置之前，必须评估可用资源使用情况，以在提交命令后预测资源状况。要验证并确认实际资源消耗，可使用以下命令：

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

注意：使用Cisco Bug ID CSCvw20873 - L3VPN LSP路径（标签交换）优化 — 对已配置cef encap-sharing disable的消耗和分配的资源**引入优化**。