

# Troubleshoot In-line BGP VPNv4 RR with Same Route-Distinguisher and 'cef encaps-sharing disabled'

## Contents

[Introduction](#)

[Background Information](#)

[Problem](#)

[Scenario 1. In-line RR and PE with Distinct RD Value](#)

[Scenario 2. In-line RR and PE with Same RD Value](#)

[Scenario 3. In-line RR and PE with Same RD Value and 'cef encaps-sharing disable' Configured Solution](#)

## Introduction

This document describes the behavior of Cisco Systems Network Convergence System (NCS) 540 while used as Border Gateway Protocol (BGP) VPNv4 in-line Route Reflector (RR) and Provider Edge (PE).

## Background Information

This document focus on NCS 540 behavior as verified in a lab environment with Cisco IOS® XR software release 7.3.1. behavior described in this document applies to all NCS5500 or NCS500 series DNX-based platforms and software releases.

## Problem

Consider the scenario where NCS540 is configured with a Virtual Routing and Forwarding (VRF) instance and uses the same Route-Distinguisher (RD) value as the one used by remote PE RR client nodes. When used as an in-line RR and PE role and configured with the same RD value as remote PE RR client nodes the topmost label is not popped before forwarding which leads to packet drop at egress PE.

### Scenario 1. In-line RR and PE with Distinct RD Value

The image shows the scenario of Device Under Test (DUT) configured as BGP VPNv4 in-line RR and PE with a distinct RD value for VRF vrf1 than the one used by the other PE nodes.

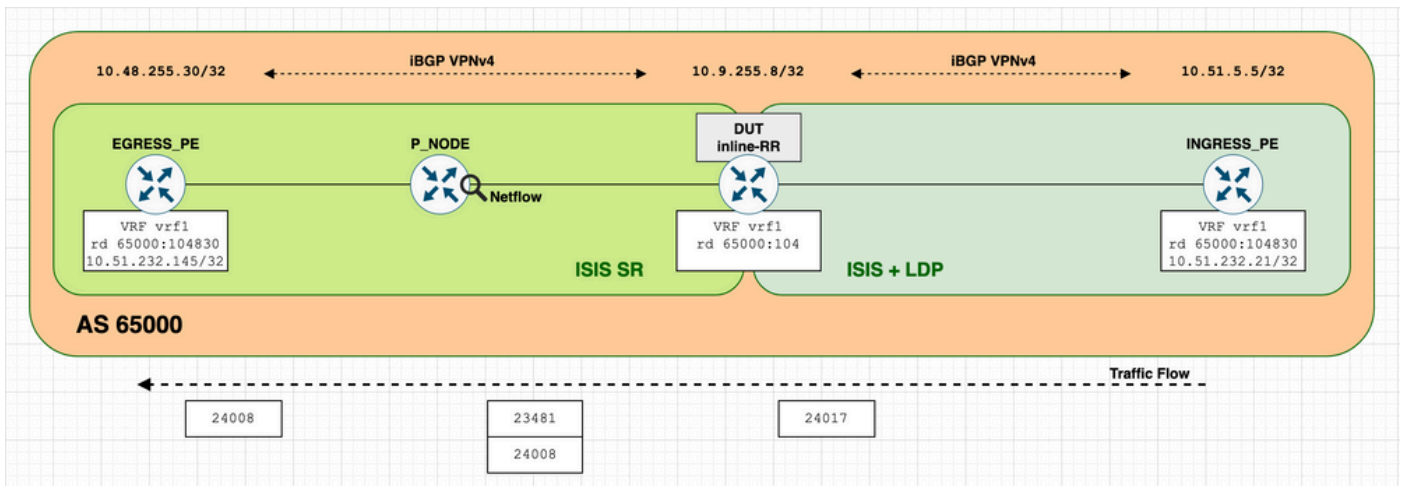


Image 1 - DUT - in-line RR and PE with distinct RD value.

IP Connectivity between Ingress PE (IP 10.51.232.21) and Egress PE (10.51.232.145) within VRF vrf1 is established with packets successfully forwarded between PE nodes as shown by Ping and Traceroute commands output:

```
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
```

The DUT is a BGP VPNv4 in-line RR and PE with VRF locally configured, but that uses a distinct RD value (65000:104) than the one used by the PE RR client nodes (65000:104830). As exhibited by the outputs, the DUT imports all routes and exchanges VPNv4 labels accordingly:

```
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: <b>21497:104830</b>			
*>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>

As an additional reference in the present scenario of **in-line RR and PE that uses a distinct RD value**, the complete output for prefix 10.51.232.145/32 as received from EGRESS\_PE node is shown:

```
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
```

The packets sourced from INGRESS\_PE (10.51.232.21) and destined to EGRESS PE (10.51.232.145) are label switched, and at the DUT top label {24017} swapped by {23481 24008} as per programmed forwarding details:

```

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

```

At P-NODE, with Netflow configured to match ingress traffic sourced from INGRESS\_PE node, the expected label stack {23481 24008} is observed as shown in the flow monitor output:

```

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>

```

## Scenario 2. In-line RR and PE with Same RD Value

The image shows the problematic scenario with DUT configured as BGP VPNv4 in-line RR and PE but now configured with the same RD value for VRF vrf1 - 65000:10430 - as the other PE nodes.

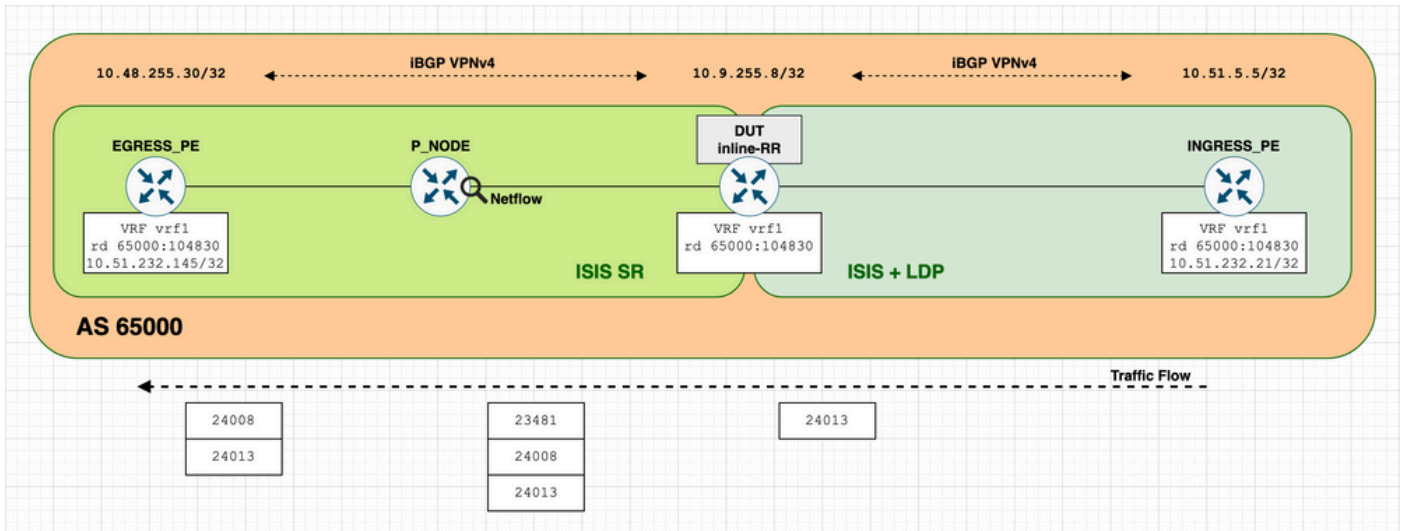


Image 2 - DUT - in-line RR and PE with same RD value.

In this scenario, the IP Connectivity between Ingress PE (IP 10.51.232.21) and Egress PE (10.51.232.145) within VRF vrf1 fails as shown with Ping and Traceroute commands output:

```
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>
```

At DUT, problem origin is not clearly perceived from either BGP or programmed forwarding outputs with all seen as expected:

```
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>
```

Similar to the previous section and as an additional reference in the present scenario of in-line RR and PE using the same RD value, the complete output for prefix 10.51.232.145/32 as received

from EGRESS\_PE node is shown:

```
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
```

The packets sourced from INGRESS\_PE (10.51.232.21) and destined to EGRESS PE (10.51.232.145), are label switched and the expectation is that at the DUT these packets have their top label {24013} swapped by {23481 24008} as per programmed forwarding details:

```
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
```

Further, triage at P\_NODE with Netflow configured to match ingress traffic sourced from INGRESS\_PE reveals the problem origin to be with how DUT forwards the packet. As highlighted in the output shown, the local-label 24013 is not popped by the DUT before forwarding the packet. Hence, P\_NODE receives a three-label stacked MPLS frame {23481 24008 24013} instead of the expected two-label stacked {23481 24008}. At P\_NODE top label 23481 is stripped and the packet forwarded towards EGRESS\_PE carries as label stack {24008 24013} which leads to packet drop at EGRESS\_PE.

```

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>

```

### Scenario 3. In-line RR and PE with Same RD Value and 'cef encap-sharing disable' Configured

On top of the previous scenario, **cef encap-sharing disable** is configured and separate hardware resources are allocated for every prefix.

```

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

```

The image exhibits the scenario of in-line RR and PE configured with the same RD and configured with **cef encap-sharing disabled**.

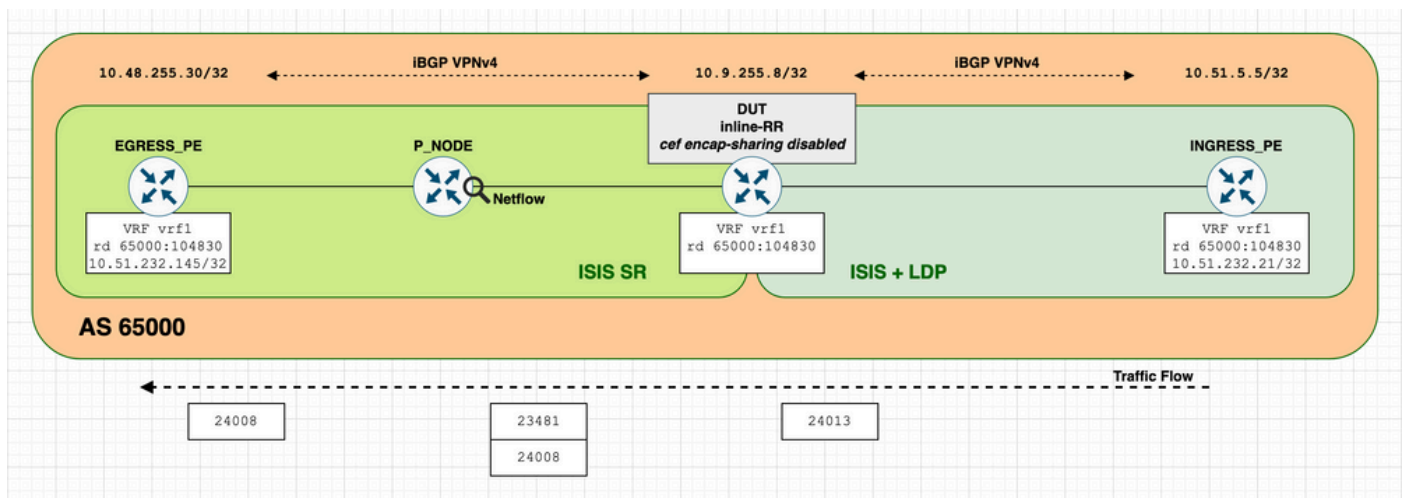


Image 3 - DUT - in-line RR and PE with same RD value and 'cef encap-sharing disable' configured.

IP Connectivity between Ingress PE (IP 10.51.232.21) and Egress PE (10.51.232.145) within VRF vrf1 is restored with packets once more successfully forwarded between PE nodes as shown by Ping and Traceroute command outputs:

```

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
```

At P\_NODE, the Netflow output reveals that DUT now forwards the packet with the expected label stack {23481 24008}. P\_NODE strips the top label 23481 and forwards the packet towards EGRESS\_PE node with label stack {24008} which EGRESS\_PE is able to match and successfully forward encapsulated packet towards its final destination within VRF1.

```
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>
```

With **cef encap-sharing disable** configuration, additional hardware resources are allocated for every prefix, and required forwarding information is in place at DUT to correctly forward packets in this particular scenario of in-line RR with the same RD value. With the purpose to highlight the distinction with regards to programmed forwarding details, refer to the two outputs shown of **show cef vrf vrf1 10.51.232.145/32 hardware egress location 0/RP0/CPU0** and note the additional information included when **cef encap-sharing disable** is configured.

As seen in the scenario of **in-line RR and PE with same RD value but without 'cef encap-sharing disable'** configured (default):

```
!
! --- 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

```

As seen in the scenario of **in-line RR and PE with same RD value and cef encap-sharing disable** configured:

```

!
! --- 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

## Solution

As described throughout this document, the solution for improper label action of the topmost label not popped by in-line RR with the same RD value is to configure **cef encap-sharing disable**. This configuration is mandatory in this particular scenario and enforces that separate hardware resources are allocated for every prefix which guarantees that proper label action and forwarding happen at the in-line RR node.

Before configuration commit, available resources usage must be assessed to anticipate resources condition once the command is committed. To verify and confirm actual resources consumption these commands can be used:

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

**Note:** With Cisco bug ID [CSCvw20873](#) - L3VPN LSP path (label swap) optimization - introduces an optimization on resources consumed and allocated with **cef encap-sharing disable** configured.