

NETWORKERS 2003

THE POWER TO TRANSFORM BUSINESS. **now.**



Troubleshooting MPLS VPN Networks

Session RST-3061

RST-3061
8186_05_2003_c1


© 2003, Cisco Systems, Inc. All rights reserved.

2

Recommended Reading

Cisco.com

- **MPLS and VPN Architectures “Vol.2”**
By Jim Guichard, Jeff Apcar et al



RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

3

Agenda

Cisco.com

- **Prerequisites**
- **MPLS VPN Troubleshooting**
 - Control Plane
 - Forwarding Plane
- **Conclusion**

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

4

Prerequisites

Cisco.com

Targeted Audience

- Anybody who has either **deployed or is deploying MPLS VPNs**
- Anybody who **understands MPLS VPN** and played with it

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

5

Prerequisites

Cisco.com

- Routing protocols especially **BGPv4**
- **MPLS VPN** in general
- LDP
- **Other Sessions**

Intro to MPLS

RST-1061

Deploying MPLS VPN

RST-2061

Deployment of BGP

RST-2003

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

6

Before We Begin...

Cisco.com

- **This presentation will not cover**
Cell-mode MPLS, TE, AToM etc.
- **What can you expect to get from this presentation?**
Overview of MP-BGP/VPN
Learn how to use show commands and
debugs to troubleshoot MPLS/VPN problems
Troubleshooting tips; Real world examples

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

7

Agenda

Cisco.com

- **Prerequisites**
- **MPLS VPN (L3 VPN) Troubleshooting**
Control Plane
Forwarding Plane
- **Conclusion**

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

8

MPLS VPN Troubleshooting

Cisco.com

- Oh...the VPN traffic is not getting through
- Ooops...I don't see VPN routes in the table
- Routes are there, but labels aren't
- Labels are there in BGP, but not in LFIB
-
- What do we do now—Call somebody?

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

9

MPLS VPN Control Plane—MPBGP

Cisco.com

- MPLS VPN is based on **RFC2547**
- The whole MPLS VPN concept revolves around MP-BGP
- **MP-BGP stands for Multi Protocol BGP**
- Multi-protocol refers to the ability of BGP to exchange information about multiple protocols such as IPv4, **VPNv4**, IPv6, multicast etc

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

10

MPLS VPN Control Plane—MPBGP

Cisco.com

- Why/What/How MP-BGP is different from the typical BGP?
 - It is not really different ☺
- It is just other capability(s) that peers have to negotiate during BGP session setup
- In the context of VPN, MP-BGP refers to the BGP session in VPNv4 address-family

```
router bgp 1
  bgp router-id 10.13.1.61
  neighbor 10.13.1.21 remote-as 1
  neighbor 10.13.1.21 update-source Lo0
  !
  address-family ipv4
    neighbor 10.13.1.21 activate
  !
```

```
router bgp 1
  bgp router-id 10.13.1.61
  neighbor 10.13.1.21 remote-as 1
  neighbor 10.13.1.21 update-source Lo0
  !
  address-family vpnv4
    neighbor 10.13.1.21 activate
  !
```

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

11

MPLS VPN Control Plane—MPBGP

Cisco.com

- MP-BGP session facilitates the advertisement of VPNv4* prefixes + Labels between MP-BGP peers
- **On the advertising PE**, BGP allocates labels for VPN prefixes and installs them in the LFIB
- **On the receiving PE**, (if) BGP accepts VPN prefixes with labels, (then) installs them in the VRF FIB

* VPNv4 = RD:IPv4

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

12

MPLS VPN Control Plane—

Cisco.com

• IPv4 BGP session between PE-CE

```
PE1#
*May 9 02:46:00.139: BGP: 200.1.61.6 sending OPEN, version 4, my as: 1
*May 9 02:46:00.139: BGP: 200.1.61.6 rcv OPEN w/ OPTION parameter len: 24
*May 9 02:46:00.139: BGP: 200.1.61.6 rcvd OPEN w/ optional parameter type 2 (Capability) len 6
*May 9 02:46:00.139: BGP: 200.1.61.6 OPEN has CAPABILITY code: 1, length 4
*May 9 02:46:00.139: BGP: 200.1.61.6 OPEN has MP_EXT CAP for afi/safi: 1/1
```

IPv4 capability

```
.....
Apr 30 01:25:31.416 EDT: %BGP-5-ADJCHANGE: neighbor 200.1.61.6 vpn vrf v1 Up
PE1#
```

```
CE1#
*May 9 02:45:59.557: BGP: 200.1.61.5 sending OPEN, version 4, my as: 65000
*May 9 02:45:59.557: BGP: 200.1.61.5 rcv OPEN w/ OPTION parameter len: 16
*May 9 02:45:59.557: BGP: 200.1.61.5 rcvd OPEN w/ optional parameter type 2 (Capability) len 6
*May 9 02:45:59.557: BGP: 200.1.61.5 OPEN has CAPABILITY code: 1, length 4
*May 9 02:45:59.557: BGP: 200.1.61.5 OPEN has MP_EXT CAP for afi/safi: 1/1
*May 9 02:45:59.557: BGP: 200.1.61.5 rcvd OPEN w/ optional parameter type 2 (Capability) len 2
```


IPv4 capability

```
.....
*May 9 02:45:59.649: %BGP-5-ADJCHANGE: neighbor 200.1.61.5 Up
CE1#
```

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

13



MPLS VPN Control Plane—

IPv4 BGP Session between PE-CE

Cisco.com

It is a Normal IPv4
BGP Session with CE

```
PE1#sh ip bgp vpnv4 vrf v1 neighbors
BGP neighbor is 200.1.61.6, vrf v1, remote AS 65000, external link
BGP version 4, remote router ID 5.5.5.5
BGP state = Established, up for 01:01:32
Last read 00:00:32, hold time is 180, keepalive interval is 60 seconds
neighbor capabilities:
  Route refresh: advertised and received(old & new)
  Address family IPv4 Unicast: advertised and received
  vpnv4 MPLS label capability: received
Message statistics:
  InQ depth is 0
  OutQ depth is 0

      Sent          Rcvd
Opens:             2           2
Notifications:     0           0
Updates:           9           6
Keepalives:        76          76
Route Refresh:     0           2
Total:             87          85
Default minimum time between advertisement runs is 30 seconds
```

But PE Stores the
CE-Sent Routes in
the VPNv4 Table

```
For address family: VPNv4 Unicast
Translates address family IPv4 Unicast for VRF v1
BGP table version 23, neighbor version 23
Index 3, Offset 0, Mask 0x8

      Sent          Rcvd
Prefix activity:
  Prefixes Current:  2           4 (Consumes 256 bytes)
  Prefixes Total:    3           4
Implicit Withdraw:   0           0
Explicit Withdraw:   1           0
Used as bestpath:    n/a         3
Used as multipath:    n/a         0
//////////deleted//////////
```




RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

14

MPLS VPN Control Plane—

• MP-iBGP session between RE1-RR (both IPv4 and VPNv4)



PE1#

```
*May 9 03:34:27.399: BGP: 10.13.1.21 rcv OPEN, version 4
*May 9 03:34:27.399: BGP: 10.13.1.21 sending OPEN, version 4, my as:1
*May 9 03:34:27.399: BGP: 10.13.1.21 rcvd OPEN w/ optional parameter type 2 (Capability) len 6
```

```
*May 9 03:34:27.399: BGP: 10.13.1.21 OPEN has CAPABILITY code: 1, length 4
*May 9 03:34:27.399: BGP: 10.13.1.21 OPEN has MP_EXT CAP for afi/safi: 1/1
```

IPv4 capability

```
*May 9 03:34:27.399: BGP: 10.13.1.21 rcvd OPEN w/ optional parameter type 2 (Capability) len 6
```

```
*May 9 03:34:27.399: BGP: 10.13.1.21 OPEN has CAPABILITY code: 1, length 4
*May 9 03:34:27.399: BGP: 10.13.1.21 OPEN has MP_EXT CAP for afi/safi: 1/128
```

VPNv4 capability

```
*May 9 03:34:27.647: %BGP-5-ADJCHANGE: neighbor 10.13.1.21 Up
```

RR1#

```
*May 9 03:34:26.808: BGP: 10.13.1.61 rcv OPEN, version 4
*May 9 03:34:26.808: BGP: 10.13.1.61 sending OPEN, version 4, my as: 1
*May 9 03:34:26.808: BGP: 10.13.1.61 rcvd OPEN w/ optional parameter type 2 (Capability) len 6
```

```
*May 9 03:34:26.808: BGP: 10.13.1.61 OPEN has CAPABILITY code: 1, length 4
*May 9 03:34:26.808: BGP: 10.13.1.61 OPEN has MP_EXT CAP for afi/safi: 1/1
```

IPv4 capability

```
*May 9 03:34:26.808: BGP: 10.13.1.61 rcvd OPEN w/ optional parameter type 2 (Capability) len 6
```

```
*May 9 03:34:26.808: BGP: 10.13.1.61 OPEN has CAPABILITY code: 1, length 4
*May 9 03:34:26.808: BGP: 10.13.1.61 OPEN has MP_EXT CAP for afi/safi: 1/128
```

VPNv4 capability

```
*May 9 03:34:26.808: BGP: 10.13.1.61 rcvd OPEN w/ optional parameter type 2 (Capability) len 2
*May 9 03:34:27.148: %BGP-5-ADJCHANGE: neighbor 10.13.1.61 Up
```

RR1#


RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

15

MPLS VPN Control Plane—

MP-iBGP Session between PE-RR



Information about IPv4 Routes First

```
PE1#sh ip bgp vpnv4 all neighbors 10.13.1.21
BGP neighbor is 10.13.1.21, remote AS 1, internal link
BGP version 4, remote router ID 10.13.1.21
BGP state = Established, up for 00:17:35
Last read 00:00:35, hold time is 180, keepalive interval is 60 seconds
```

Neighbor capabilities:

```
Route refresh: advertised and received(old & new)
Address family IPv4 Unicast: advertised and received
Address family VPNv4 Unicast: advertised and received
```

Message statistics:

```
InQ depth is 0
OutQ depth is 0

Opens:          Sent      Rcvd
Notifications:  0         0
Updates:        9         10
Keepalives:     80        80
Route Refresh:  0         0
Total:          91        92
Default minimum time between advertisement runs is 5 seconds
```

For address family: IPv4 Unicast

```
BGP table version 4, neighbor version 4
Index 1, Offset 0, Mask 0x2
NEXT_HOP is always this router
```

```
Prefix activity:          Sent      Rcvd
-----
Prefixes Current:         0         1 (Consumes 48 bytes)
Prefixes Total:           0         1
Implicit Withdraw:        0         0
Explicit Withdraw:        0         0
Used as bestpath:         n/a        1
Used as multipath:         n/a        0
```

```
Local Policy Denied Prefixes:  Outbound  Inbound
-----
Bestpath from this peer:       1         n/a
Total:                         1         0
Number of NLRI in the update sent: max 0, min 0
```

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

16

MPLS VPN Control Plane— PE-RR MP-iBGP Session (Cont.)

Cisco.com

Information about VPNv4 Routes

```

...confd
For address family: VPNv4 Unicast
BGP table version 23, neighbor version 23
Index 1, Offset 0, Mask 0x2

Prefix activity:
Prefixes Current:      4      Sent      Rcvd
Prefixes Total:        4      1
Implicit Withdraw:      0      0
Explicit Withdraw:      0      0
Used as bestpath:      n/a     1
Used as multipath:      n/a     0

Local Policy Denied Prefixes:
ORIGINATOR loop:      n/a     4
Bestpath from this peer: 2      n/a
Total:                  2      4
Number of NLRI in the update sent: max 3, min 0

Connections established 2; dropped 1
Last reset 00:19:50, due to User reset
..//////////<deleted>//////////
    
```

- When/Who allocates the label for a VPN prefix? **Let's go through the routing flow—**

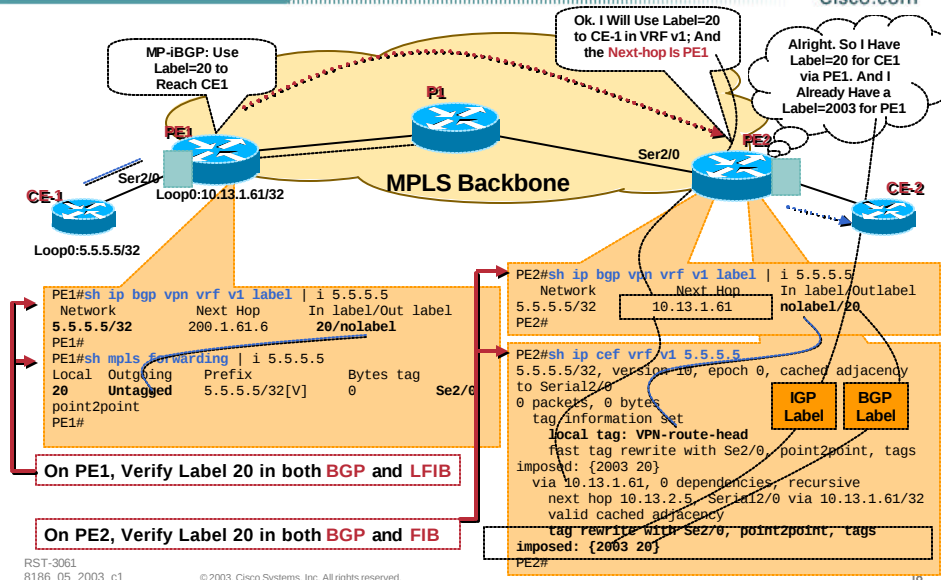
RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

17

MPLS VPN Control Plane—MPBGP Flow

Cisco.com



RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

18

MPLS VPN Control Plane

Cisco.com

- MPLS-VPN requires the IP network to be MPLS enabled first
- **One of the applications that enable MPLS is LDP;** Others are RSVP, BGPipv4+label etc
- LDP is used to exchange the label for the PEs i.e. next-hop of VPN prefixes
(BGP is used to exchange the label for the VPNv4 prefixes)

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

19

MPLS VPN Ctrl Plane— How/What to Configure

Cisco.com

The common Q regarding MPLS-VPN are:

- What is needed to configure MPLS VPN
- How do I configure?
- What show commands to look at?
- Do I need to be a rocket scientist? 😊

Next Two Slides Summarize the Configuration Steps (Additional Slides at the End of the Preso Has the Detailed Steps

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

20

MPLS VPN Ctrl Plane— How/What to Configure

Cisco.com

On PEs:

1. Configure a VRF i.e. **vrf**, RD, RT.
2. Attach a PE-CE interface to the **vrf**
3. Configure the PE-CE routing protocol in the “address-family ipv4 vrf <**vrf**>”
4. If (3) is not eBGP, then **also** redistribute the respective IGP in “address-family ipv4 vrf <**vrf**>” under BGP and vice-versa
5. Configure the MP-iBGP neighbor i.e. RR and activate it in the “address-family vpnv4” within BGP

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

21

MPLS VPN Ctrl Plane— How/What to Configure

Cisco.com

On each routers i.e. PE, P and RR:

- Usual IGP+LDP configurations

On RRs:

1. Configure the MP-iBGP neighbors i.e. PEs and activate them in the “address-family vpnv4” within BGP
2. RRs **shouldn't** be kept in the forwarding path

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

22

MPLS VPN Control Plane— Show Commands on PE

Cisco.com

1. **“sh ip bgp vpn all summary”**
Analogous to ‘sh ip bgp summary’; Lists all the MPBGP and CE peers
2. **“sh ip bgp vpn all”**
Lists **all** the VPN prefixes advertised/rcvd by the router
3. **“sh ip bgp vpn vrf <vrf> summary”**
Similar to the first one, but for a specific VRF
4. **“sh ip bgp vpn vrf <vrf>”**
Lists all the VPN prefixes received in a specific VRF
5. **“sh ip bgp vpn vrf <vrf> labels”**
List labels for the VPN prefixes in a VRF

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

23

MPLS VPN Control Plane— Show Commands on PE

Cisco.com

If OSPF on PE-CE -

- **“sh ip ospf neighbors”**
Lists both VPN(s) and non-VPN(s) OSPF neighbors
- **“sh ip ospf <process-id> ”**
Select the VRF associated process-id to see relevant OSPF info (a lot of info)
- **“sh ip ospf <process-id> database”**
Select the VRF associated process-id to see the OSPF database for that VRF
- **“clear ip ospf <process-id>”**
Clear OSPF neighbors in the VRF if VRF associated process-id is chosen

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

24

MPLS VPN Control Plane— Show Commands on PE

Cisco.com

If EIGRP on PE-CE –

- “sh ip eigrp vrf <vrf> topology”
Lists VRF specific EIGRP topology
- “sh ip eigrp vrf <vrf> neighbor|interface”
Lists EIGRP neighbors or interfaces in the VRF
- “sh ip eigrp vrf <vrf> events”
Shows VRF specific EIGRP events
- “clear ip eigrp vrf <vrf> neighbors”
Clears VRF specific EIGRP neighbors

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

25

MPLS VPN Control Plane— Clear Commands on PE

Cisco.com

Relevant towards RR (or remote PE) peers:

- “clear ip bgp * vpnv4 unicast in”
Route-refresh request is sent to all the MP-BGP peers
- “clear ip bgp <MP-BGP peer> vpnv4 unicast in”
Route-refresh request is sent to a specific MP-BGP peer

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

26

MPLS VPN Control Plane— Clear Commands on **PE**

Cisco.com

Relevant towards CEs:

- “clear ip bgp * vrf <vrf>”
Clear all PE-CE eBGP sessions in that **vrf**
- “clear ip bgp * vrf <vrf> in”
Route-refresh message is sent to all the CEs in that **vrf**
- “clear ip bgp * vrf <vrf> out”
Send respective VPN routes to all the CEs in that **vrf**
- “clear ip bgp <CE> vrf <vrf> soft in|out”

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

27

MPLS VPN Control Plane— Show Commands on **RR**

Cisco.com

- **Route-reflector know nothing about VRF**
Following commands come quite handy
(especially on RR)
 1. “sh ip bgp vpn all”
 2. “sh ip bgp vpn rd <RD>”
Lists all VPNv4 prefixes that have **RD** in them
 3. “sh ip bgp vpn rd <RD> label”
Lists labels for VPNv4 prefixes that have **RD**

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

28

MPLS VPN Control Plane— Debugs on PE

Cisco.com

Be Careful on the Production Routers

1. “debug ip bgp vpnv4”

Useful while troubleshooting label related problems in BGP (could spit a lot of output)

2. “debug mpls lfib cef [acl]”

Useful troubleshooting label mismatch in FIB/LFIB

3. “debug ip bgp vpnv4 import”

Useful when VPN prefixes don’t get imported in the VRF table (could spit a lot of output)

4. “debug ip routing vrf <vrf> [acl]”

Useful when VPN prefixes don’t get installed in the VRF routing table

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

29

MPLS VPN Ctrl Plane—Outgoing Labels

Cisco.com

RSP-PE	SOUTH-5#	sh	mpls forwarding 10.13.1.11			
Local	Outgoing		Prefix	Bytes tag	Outgoing	Next Hop
tag	tag or VC		or Tunnel Id	switched	interface	
59	46		10.13.1.11/32	0	Se10/0/0	point2point
RSP-PE	SOUTH-5#					

- **Outgoing** label also conveys what treatment the packet is going to get; it could also be:

Untagged —Untag the incoming MPLS packet

Aggregate —Untag and then do a FIB lookup

Pop —Pops the topmost label

0 —Nullify the top label (first 20bits)

- Label values 0–15 are reserved

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

30

MPLS VPN—Outgoing Labels

Cisco.com

```
PE1#sh mpls forwarding-table
```

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes tag switched	Outgoing interface	Next Hop
16	2002	10.13.1.22/32	0	Et0/0	10.13.1.5
	2002	10.13.1.22/32	0	Et1/0	10.13.1.9
17	2001	10.13.1.62/32	0	Et0/0	10.13.1.5
	2001	10.13.1.62/32	0	Et1/0	10.13.1.9
18	Pop tag	10.13.1.101/32	0	Et1/0	10.13.1.9
	Pop tag	10.13.1.101/32	0	Et0/0	10.13.1.5
19	Pop tag	10.13.2.4/30	0	Et1/0	10.13.1.9
	Pop tag	10.13.2.4/30	0	Et0/0	10.13.1.5
→ 20	Untagged	5.5.5.5/32[V]	0	Se2/0	point2point
21	Pop tag	10.13.21.4/30	0	Et1/0	10.13.1.9
	Pop tag	10.13.21.4/30	0	Et0/0	10.13.1.5
22	Pop tag	10.13.22.4/30	0	Et1/0	10.13.1.9
	Pop tag	10.13.22.4/30	0	Et0/0	10.13.1.5
23	Aggregate	0.0.0.0/0[V]	0		
→ 24	Aggregate	200.1.61.4/30[V]	0		
26	Untagged	30.30.30.1/32[V]	0	Se2/0	point2point

PE1#

Connected
VPN Prefix

V Means It is a VPN Prefix

No Outgoing Interface for the
Aggregate Entries; an
Additional FIB Lookup Is Done

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

31

MPLS VPN—OSPF/EIGRP on PE-CE

Cisco.com

- OSPF and EIGRP on PE-CE brings few new stuff
- Let's go over OSPF first...and then EIGRP

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

32

MPLS VPN—OSPF on PE-CE

Cisco.com

- By default, OSPF->BGP redistributes “internal” i.e intra and inter-area routes only, **no external**
- Configure “match internal external” within BGP VRF to redistribute OSPF **externals as well**
- All intra-area routes (type1 and type2) are advertised as inter-area (type3) by PE to CE
- OSPF information i.e. area#, LSA type, router-id, domain-id etc are carried across MPLS/VPN backbone by BGP in **new extended communities**

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

33

MPLS VPN—OSPF on PE-CE

Cisco.com

```
PE1#sh ip bgp vpnv4 vrf v1 30.1.61.4/30
BGP routing table entry for 1:1:30.1.61.0/30, version 8
Paths: (1 available, best #1, table v1)
  Not advertised to any peer
    Local
      200.1.1.1 (metric 435200) from 200.1.1.1 (200.1.1.1)
        Origin incomplete, metric 0, localpref 100, valid, internal, best
        Extended Community: RT:1:1 OSPF DOMAIN ID:0.0.0.1 OSPF RT:1:2:0
        OSPF ROUTER ID:10.13.100.1
```

router-id within the OSPF VRF instance

domain-id = OSPF process-id (default) = 1

OSPF Route Type : area 1, type 2, no options

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

34

MPLS VPN—OSPF on PE-CE

Cisco.com

- Only on **type3 LSA**, PE sets “down bit” to avoid the looping while advertising to CE
- Only on **type5 LSA**, PE sets “tag” to avoid the looping while advertising to CE (can be tuned via “domain-tag”)
- By default, BGP MED is set to the value of the OSPF metric and vice versa (can be tuned)

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

35

MPLS VPN—EIGRP on PE-CE

Cisco.com

- PE->PE: EIGRP routes are advertised into MP-BGP preserving EIGRP info i.e. route-type, bandwidth, delay, reliability, MTU etc in the new **extended-community** attributes.
- PE->CE: BGP redistributes routes into EIGRP using route-type and metric information extracted from BGP extended-communities.
 - If EIGRP sites are in the same AS#, then PE will keep the route-type and metric of the route unchanged. Otherwise, route-type will be external and metric will be set to default.
- PE->CE: EIGRP adds PE-CE link cost to compute new VecMetric to avoid routing loop (for dual-homed sites)

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

36

MPLS VPN—EIGRP on PE-CE

```
7206-PE-SOUTH-1#sh ip eigrp vrf v19 topology 61.19.6.12 255.255.255.252
IP-EIGRP topology entry for 61.19.6.12/30
State is Passive, Query origin flag is 1, 1 Successor(s), FD is 40537600
Routing Descriptor Blocks:
  200.19.61.6 (Serial1/3:19), from 200.19.61.6, Send flag is 0x0
    Composite metric is (40537600/51200), Route is External
    Vector metric:
      Minimum bandwidth is 64 Kbit
      Total delay is 21000 microseconds
      Reliability is 255/255
      Load is 1/255
      Minimum MTU is 1500
      Hop count is 1
    External data:
      Originating router is 200.19.61.6
      AS number of route is 130
      External protocol is BGP, external metric is 0
      Administrator tag is 1300 (0x00000514)
7206-PE-SOUTH-1#
```

Cisco.com

```
7206-PE-SOUTH-1#sh ip bgp vpnv4 vrf v19
BGP routing table entry for 19:1:61.19.6.12/30, version 15528
Paths: (1 available, best #1, table v19)
  Advertised to update-groups:
    9
    Local
    200.19.61.6 (via v19) from 0.0.0.0 (10.13.1.61)
      Origin incomplete, metric 1, localpref 100, weight 32768, valid, sourced, best
      Extended Community: RT:19:1 0x8800:0:1300 0x8801:6119:537600
      0x8802:65281:40537600 0x8803:65281:1500 0x8804:130:3356704006 0x8805:9:0
7206-PE-SOUTH-1#
```

RST-3061
8186_05_2003_c1

Please check the notes for the explanation.

External route info

37

MPLS VPN—Troubleshooting Tips

Cisco.com

1. Make sure that “export RT <X>” on the advertising router matches with “import RT <X>” on the received router
“sh ip vrf detail <vrf> | inc Export|import|RT”
2. If export or import-map are configured in the VRF, then validate the RT in the “set clause”
“sh ip vrf de <vrf> | inc route-map”;
“sh route-map <map>”
3. If BGP is not used as the PE-CE protocol, then make sure the redistribution between BGP’s VRF instance and respective IGP’s VRF instance”

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

38

MPLS VPN—Troubleshooting Tips

Cisco.com

4. On RRs, PEs must be configured as the **rr-client** in the “**address-family vpnv4**” under BGP
5. MP-BGP neighbors i.e. PE and RR must be configured to send extended-community

“sh run | inc send-community”

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

39

MPLS VPN—Troubleshooting Tips

Cisco.com

6. Make sure that the label in BGP VPN table matches with label in FIB table for a **received** VPN prefix
7. Make sure that the label in BGP VPN table matches with label in LFIB table for an **advertised** VPN prefix

“sh ip bgp vpn vrf <vrf> label | inc <prefix>”

“sh ip cef vrf <vrf> <prefix>”

“sh ip bgp vpn vrf <vrf> label | inc <prefix>”

“sh mpls forwarding vrf <vrf> | inc <prefix>”

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

40

MPLS VPN—Troubleshooting Tips

Cisco.com

8. **Outer (or IGP) label** in the label stack provides a LSP from ingress PE to egress PE via MPLS cloud
9. **Inner (or BGP) label** refers to the VPNv4 prefix at the egress PE

tag rewrite with Se2/0, point2point, tags imposed: {2003 20}

10. If the router doesn't change the next-hop attribute of VPNv4 prefix, then no LFIB entry is created; so don't panic

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

41

MPLS VPN Control Plane—Troubles

Cisco.com

- Let's do some MPLS VPN troubles(hooting) 😊

RST-3061
8186_05_2003_c1


© 2003, Cisco Systems, Inc. All rights reserved.

42

MPLS VPN Ctrl Plane—Trouble #1

Cisco.com

#1: VPN prefix doesn't have any label in the LFIB on the local PE



TIP: Label allocation is done by BGP. So make sure the prefix is in the BGP VRF table. Hint—**“redistribute connected”**

RST-3061
8186_05_2003_c1

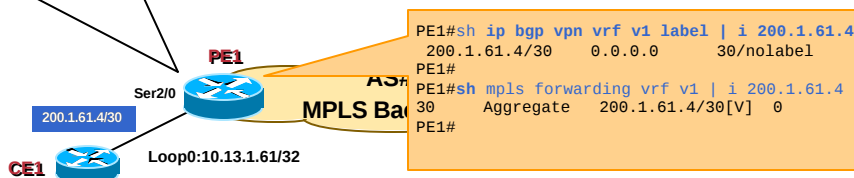
© 2003, Cisco Systems, Inc. All rights reserved.

43

MPLS VPN Ctrl Plane—Trouble #1 (Cont.)

Cisco.com

```
PE1(conf)#router bgp 1
PE1(conf-router)#address-family ipv4 vrf v1
PE1(conf-router-af)#redistribute connected
PE1(conf-router-af)#end
```



- As soon as BGP gets the VPN prefix, it allocates the local label, and installs the prefix+label in both BGP and LFIB

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

44

MPLS VPN Ctrl Plane—Trouble #2

Cisco.com

- **Prob#2:** VPN prefix doesn't have any label in the LFIB on the local PE, though BGP now does
- **TIP:** “clear ip route vrf <vrf> <prefix>”
If the above doesn't fix, then (soft) reset the BGP session

RST-3061
8186_05_2003_c1


© 2003, Cisco Systems, Inc. All rights reserved.

45

MPLS VPN Ctrl Plane—Trouble #3

Cisco.com

- **#3:** Remote PE (PE2) doesn't get the VPNv4 prefix from PE1



- **TIP:** Validate “route-target import” config on PE2. If not present, then configure it; Check for import-map as well

RST-3061
8186_05_2003_c1


© 2003, Cisco Systems, Inc. All rights reserved.

46

MPLS VPN Ctrl Plane—Trouble #4

Cisco.com

#4: Remote PE (PE2) **still** doesn't get the VPNv4 prefix from PE1



- We already fixed PE2; so let's go to PE1
- Validate "Route-target export" in the VRF on the PE1


RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

47

MPLS VPN Ctrl Plane—Trouble #4 (Cont.)

Cisco.com



- **TIP:** Configure Route-target export in the VRF on the local PE i.e. PE1
- Let's make sure that RT is getting tagged to the VPNv4 prefix


RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

48

MPLS VPN Ctrl Plane—Trouble #4 (Cont.)

Cisco.com



```
PE1#sh ip bgp vpnv4 vpn vrf v1 200.1.61.4
BGP routing table entry for 1:1:200.1.61.4/30, version 10
Paths: (2 available, best #2, table v1)
  Advertised to non peer-group peers:
    10.13.1.21 200.1.61.6
  Local
    0.0.0.0 from 0.0.0.0 (10.13.1.61)
    Origin incomplete, metric 0, localpref 100, weight
    32768, valid, sourced, best
    Extended Community: RT:1:1
PE1#
```

Extra-TIP 😊

- If export or import map are also configured, then check the RT in “set clause”, along with the “match clause”

RST-3061
8186_05_2003_c1


© 2003, Cisco Systems, Inc. All rights reserved.

49

MPLS VPN Ctrl Plane—Trouble#5

Cisco.com

#5: Remote PE (PE2) **STILL** doesn't get the VPNv4 prefix from PE1



```
RR1#sh ip bgp vpnv4 rd 1:1 200.1.61.4
BGP routing table entry for 1:1:200.1.61.4/30, version 14
Paths: (1 available, best #1, no table)
  Advertised to non peer-group peers:
    10.13.1.62
  Local, (Received from a RR-client)
    10.13.1.61 (metric 75) from 10.13.1.61 (10.13.1.61)
    Origin incomplete, metric 0, localpref 100, valid, internal, best
    Extended Community: RT:1:1
RR1#
```

- We have confirmed that PE1 is advertising the prefix; let's check the RR
- Let's make sure that RR is configured with “neighbor <PE2> send-community extended” under vpnv4 af


RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

50

MPLS VPN Ctrl Plane—Trouble #5 (Cont.)

Cisco.com



TIP:

- All the MP-BGP peers **must** be configured with “send-community extended|both”
- Make sure that PE1 and PE2 are configured as rr-client under “vpnv4 af” on the RR1

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.


51

MPLS VPN Control Plane—Trouble #6

Cisco.com

#6: Remote PE (PE2) **STILL** doesn't get the VPNv4 prefix from PE1

```
PE2#sh ip vrf detail v1 | i import
Import route-map: raj-import
PE2#
PE2#sh route-map raj-import
route-map raj-import, permit, sequence 10
Match clauses:
  extcommunity (extcommunity-list filter):1
Set clauses:
  Policy routing matches: 0 packets, 0 bytes
PE2#
PE2#sh ip extcommunity-list 1
Extended community standard list 1
  deny RT:1:1
  deny RT:2:2
PE2#
```



Oh no....who did that
&^%#@#%#@^%

That's ok. Let's
Remove RT 1:1
from the Filter.

- Hmm... we have already verified PE1 and RR1; something must be missing on PE2
- Let's check for the import-map on PE2 again

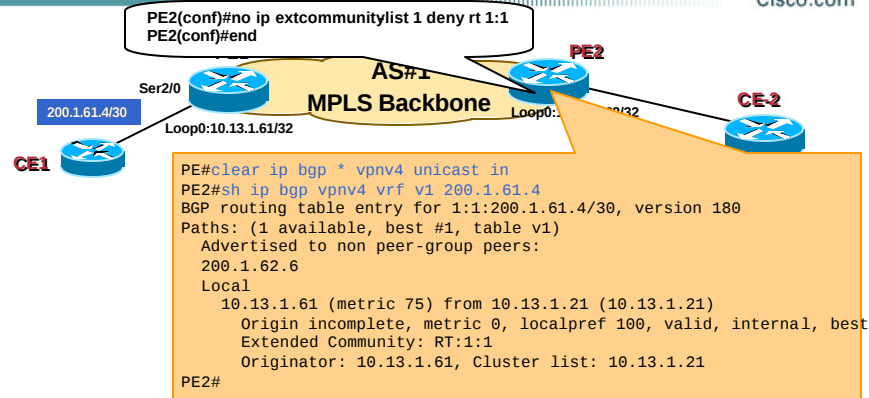
RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

52

MPLS VPN Control Plane—Trouble #6 (Cont.)

Cisco.com



TIP: If import-map is configured within the VRF, then “import route-target <rt>” **must** be configured

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.


53

MPLS VPN Control Plane—Trouble #7

Cisco.com

#7: Label mismatch between BGP and FIB

```
PE2#sh ip bgp vpnv4 vrf v1 labels | i 200.1.61.4
200.1.61.4/30 10.13.1.61 nlabel/25
PE2#
PE2#sh ip cef vrf v1 200.1.61.4
200.1.61.4/30, version 64, epoch 0, cached adjacency to Serial2/0
0 packets, 0 bytes
tag information set
  local tag: VPN-route-head
  fast tag rewrite with Se2/0, point2point, tags imposed: {2003 20}
  via 10.13.1.61, 0 dependencies, recursive
  next hop 10.13.2.5, Serial2/0 via 10.13.1.61/32
  valid cached adjacency
  tag rewrite with Se2/0, point2point, tags imposed: {2003 20}
PE2#
```



- **Fix:** “clear ip route vrf <vrf> <prefix>”. If the mismatch doesn’t go away, then “debug ip bgp vpn” and “debug mpls lfib cef” to dig in. 😊

RST-3061
8186_05_2003_c1


© 2003, Cisco Systems, Inc. All rights reserved.

54

MPLS VPN Control Plane—Trouble #8

Cisco.com

#8: Remote PE receives the route, but remote CE doesn't



- **TIP:** If eBGP on PE-CE and VPN sites use the same ASN, then configure “as-override” on the BGP VRF af on both PEs
- If IGP on PE-CE, then validate BGP->IGP redistribution (within IGP VRF) on the PE

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

55

Agenda

Cisco.com

- Prerequisites
- **MPLS VPN (L3 VPN) Troubleshooting**
 - Control Plane
 - Forwarding Plane
- Conclusion

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

56

MPLS VPN Fwd Plane—Show Command

Cisco.com

```
RSP-PE-WEST-4#sh mpls forward 10.13.1.11 detail
Local Outgoing Prefix Bytes tag Outgoing Next Hop
tag tag or VC or Tunnel Id switched interface
45 51 10.13.1.11/32 0 Fa1/1/1 10.13.7.33
MAC/Encaps=14/18, MRU=1500, Tag Stack{51}
0003FD1C828100044E7547298847 00033000
No output feature configured
Per-packet load-sharing
RSP-PE-WEST-4#
```

Detail is optional

Only one outgoing label in the label stack

MRU – Max Receivable Unit. The **received packet** will be transmitted unfragmented on Fa1/1/1, if its size is not more than 1500B.

MAC header = 0003FD1C828100044E754829
MPLS Ethertype= 0x8847
Label = 0x00033000 = 51

Although MAC header is of 14 bytes, actual encapsulation i.e MAC+MPLS header is of 18 bytes (one label is 4 bytes)

```
PE1#sh mpls for vrf v1 30.30.30.1 detail
Local Outgoing Prefix Bytes tag Outgoing Next Hop
tag tag or VC or Tunnel Id switched interface
27 Untagged 30.30.30.1/32[V] 0 Se2/0 point2point
MAC/Encaps=0/0, MRU=1504, Tag Stack{}
VPN route: v1
No output feature configured
Per-packet load-sharing
PE1#
```

Se2/0 is a PE-CE interface which is under VRF v1

MPLS VPN Fwd Plane—Loadsharing

Cisco.com

- Loadsharing in MPLS VPN network is same as that of the IP network
- Hence, FIB per-destination loadsharing is the default
- IP src and dest inside the MPLS packet are hashed

Let's Go through PE-P and P-P Loadsharing

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.


58

MPLS VPN Fwd Plane—Loadsharing (I)

Cisco.com

• PE-P Loadsharing (Cont.)

```
PE1#sh ip cef vrf v1 200.1.62.4
200.1.62.4/30, version 13, epoch 0, per-destination sharing
0 packets, 0 bytes
tag information set
  local tag: VPN-route-head
  fast tag rewrite with
    Recursive rewrite via 10.13.1.62/32, tags imposed {25}
  via 10.13.1.62, 0 dependencies, recursive
  next hop 10.13.1.9, Ethernet1/0 via 10.13.1.62/32
  valid adjacency
  tag rewrite with
    Recursive rewrite via 10.13.1.62/32, tags imposed {25}
  Recursive load sharing using 10.13.1.62/32.
PE1#
```



Because There Are Loadshared Paths
to the Egress PE i.e. 10.13.1.62/32

Only VPN Label Is Shown

- **Don't panic...**IGP label is chosen during the forwarding (depending on the hash-bucket)

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.


59

MPLS VPN Fwd Plane—Loadsharing (I)

Cisco.com

• PE-P Loadsharing (Cont.)

```
PE1#sh ip cef 10.13.1.62
10.13.1.62/32, version 30, epoch 0, per-destination sharing
0 packets, 0 bytes
tag information set, shared
  local tag: 18
  via 10.13.1.5, Ethernet0/0, 1 dependency
  traffic share 1
  next hop 10.13.1.5, Ethernet0/0
  valid adjacency
  tag rewrite with Et0/0 10.13.1.5, tags imposed: {2001}
  via 10.13.1.9, Ethernet1/0, 1 dependency
  traffic share 1
  next hop 10.13.1.9, Ethernet1/0
  valid adjacency
  tag rewrite with Et1/0 10.13.1.9, tags imposed: {2001}
  0 packets, 0 bytes switched through the prefix
  tmstats: external 0 packets, 0 bytes
  internal 0 packets, 0 bytes
PE1#
```



IGP Label Is Right Here

- IGP Label and the outgoing interface are derived after the hash-bucket is decided ☺

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

60

MPLS VPN Fwd Plane—Loadsharing (I)


Cisco.com

• PE-P Loadsharing (cont.)

```
PE1#sh ip cef vrf v1 exact-route 30.1.61.4 200.1.62.4 internal
30.1.61.4      -> 200.1.62.4      [ Ethernet1/0 (next hop 10.13.1.9)
                                   Bucket 7 from 16, total 2 paths
```

```
PE1#
PE1#sh ip cef vrf v1 exact-route 200.1.61.4 30.1.62.4
200.1.61.4      -> 30.1.62.4      : Null0 (attached)
PE1#
```

Because the Destination 30.1.62.4
Is Not in the VRF FIB Table



- In summary, the show-output in load-sharing case gets bit tricky; but the fundamental is the same

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.


61

MPLS Fwd Plane—Loadsharing (II)

Cisco.com

• P-P Loadsharing

```
P1#sh mpls for 10.13.1.62
Local  Outgoing  Prefix  Bytes tag  Outgoing  Next Hop
tag    tag or VC   or Tunnel Id  switched interface
52      21          10.13.1.62/32  0          Eth0/0    point2point
27      27          10.13.1.62/32  0          Eth1/0    point2point
P1#
```



- For VPN traffic, P router **hashes** the IP src+dest to apply the packet to the correct hash bucket
- “sh ip cef exact-route” command can’t be used on the P router since it doesn’t know the VPN addresses ☹
- Hence, rely on “(LFIB) counters” to make sure the traffic is getting loadshared

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

62

MPLS VPN Fwd Plane—MPLS TTL

Cisco.com

- Just like TTL in the IP header, MPLS header also has a 1-byte TTL
- When an IP packet is first labelled, the (IP TTL -1) is copied to the MPLS TTL
- When the label is removed, the MPLS TTL value of removed label is copied to the either MPLS TTL of inner label or IP TTL field (if no inner label), provided

$\text{MPLS TTL} < \text{IP TTL}$

And then, IP TTL is decremented

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

63

MPLS VPN Fwd Plane—Show Commands

Cisco.com

- “sh mpls forwarding”
Shows all LFIB entries (vpn, non-vpn, TE etc.)
- “sh mpls forwarding | inc <prefix>”
Whether the prefix is present in the LFIB or not
- “sh mpls forwarding vrf <vrf> <prefix>”
LFIB lookup based on a VPN prefix
- “sh mpls forwarding label <label>”
LFIB lookup based on an incoming label

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

64

MPLS VPN Fwd Plane—Sh Commands

Cisco.com

- “sh ip arp vrf <vrf>”
Lists ARP entries relevant to the <vrf> only
- “sh ip cef vrf <vrf> <prefix>”
Displays the label stack, outgoing interface etc
- “sh mpls forwarding vrf <vrf> ”
Lists labels for the VPN prefixes learned from the CE(s)

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

65

MPLS VPN Fwd Plane—Debugs

Cisco.com

Be Careful on the Production Routers

- “debug arp”
Useful for VPN prefixes as well
- “debug mpls lfib cef [acl]”
Useful when VPN prefixes have label mismatch among BGP, FIB and LFIB.

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

66

MPLS VPN Fwd—Troubleshooting Tips

Cisco.com



1. On PE, verify the PE-to-PE Labeled Switched Path (LSP) via
“ping <PE’s IP>”; “traceroute <PE’s IP>”
2. On PE, we could also verify the LSP via
“traceroute vrf <vrf> <prefix>”
Be careful with the traceroute output,
absence of label could mean either “pop”
or “untagged”

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

67

MPLS VPN Fwd—Troubleshooting Tips

Cisco.com



3. **VRF ping** PE-to-PE to verify that MPLS backbone is working fine
“Ping vrf <vrf> <prefix>”
Have “deb ip icmp” enabled on both local and remote PEs while VRF pinging PE-to-PE;
4. Ping CE-to-CE to verify that PEs are correctly switching the traffic
Have “deb ip icmp” enabled on both CEs while pinging; helps to find the broken LSP

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

68

MPLS VPN Fwd—Troubleshooting Tips

Cisco.com

5. If LFIB doesn't have VPN entries, then check that the FIB doesn't have "punt" adj for those prefixes. Unless adj is resolved, LFIB won't have such entries.

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

69

MPLS VPN Fwd—Troubleshooting Steps

Cisco.com

So you have been reported about the VPN traffic outage:

1. First, verify "VRF ping" from PE1 to PE2
2. **If passed**, then either CE->PE or PE->CE may be the problem=>not a MPLS core prob; STOP and Check whether the packets are getting dropped by ingress LC on PE
3. **If failed**, then MPLS core may be the problem; PROCEED
4. "Ping" ingress PE to egress PE to verify the IP reachability
5. **If failed**, then STOP and verify egress PE's route hop-by-hop
6. **If passed**, then traceroute PE1->PE2 and PE2->PE1 to ensure the PE-to-PE LSP setup
7. Also check for the labels in the each line of the traceroute output (watch out for the PHP)
8. If traceroute fails for some reason, then STOP and verify the label on every hop
9. **If good**, then the problem may be very specific to the HW on either PE or P routers; Find out that HW is dropping the packets

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

70

MPLS VPN Fwd Plane—Troubles

Cisco.com

- Let's do some **more** troubles(hooting) 😊

RST-3061
8186_05_2003_c1


© 2003, Cisco Systems, Inc. All rights reserved.

71

MPLS VPN Fwd Plane—Trouble #1

Cisco.com

#1: VPN connectivity is broken bet CEs



- Check the control plane information first
`PE1#sh ip cef vrf v1 6.6.6.6; PE1#sh mpls for vrf v1 | inc 5.5.5.5`
`PE2#sh ip cef vrf v1 5.5.5.5; PE2#sh mpls for vrf v1 | inc 6.6.6.6`
Make sure that the label information is **correct**
- Turn on “deb ip icmp” on both PEs
- Issue “ping vrf v1 <remote_PE-CE_address>” on both PEs
- If they pass, then we have verified that the problem is not in the MPLS core.


RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

72

MPLS VPN Fwd Plane—Trouble #1 (Cont.)

Cisco.com



```
PE1#sh ip cef vrf v1 200.1.62.4
200.1.62.4/30, version 10, epoch 0, per-destination sharing
0 packets, 0 bytes
tag information set
  local tag: VPN-route-head
  fast tag rewrite with
    Recursive rewrite via 10.13.1.62/32, tags imposed:
    {25}
    via 10.13.1.62, 0 dependencies, recursive
  next hop 10.13.1.9, Ethernet1/0 via 10.13.1.62/32
  valid adjacency
  tag rewrite with
    Recursive rewrite via 10.13.1.62/32, tags imposed:
    {25}
  Recursive load sharing using 10.13.1.62/32.
PE1#
```

```
PE2#sh mpls for vrf v1 | inc 200.1.62.4
25 Aggregate 200.1.62.4/30[V] 0
PE2#

PE1#sh ip cef 10.13.1.62
10.13.1.62/32, version 56, epoch 0, per-destination sharing
0 packets, 0 bytes
tag information set
  local tag: 18
  via 10.13.1.5, Ethernet0/0, 1 dependency
  traffic share 1
  next hop 10.13.1.5, Ethernet0/0
  valid adjacency
  tag rewrite with Et0/0, 10.13.1.5, tags imposed: {2001}
  via 10.13.1.9, Ethernet1/0, 2 dependencies
  traffic share 1
  next hop 10.13.1.9, Ethernet1/0
  valid adjacency
  tag rewrite with Et1/0, 10.13.1.9, tags imposed: {2001}
0 packets, 0 bytes switched through the prefix
PE1#
```

PE1->PE2 Validated for the Labels


RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

73

MPLS VPN Fwd Plane—Trouble #1 (Cont.)

Cisco.com



```
PE1#sh mpls for vrf v1 | i 200.1.61.4
28 Aggregate 200.1.61.4/30[V] 0
PE1#
```

```
PE2#sh ip cef vrf v1 200.1.61.4
200.1.61.4/30, version 73, epoch 0, cached adjacency to
Serial2/0
0 packets, 0 bytes
tag information set
  local tag: VPN-route-head
  fast tag rewrite with Se2/0, point2point, tags imposed:
  {2003 28}
  via 10.13.1.61, 0 dependencies, recursive
  next hop 10.13.2.5, Serial2/0 via 10.13.1.61/32
  valid cached adjacency
  tag rewrite with Se2/0, point2point, tags imposed:
  {2003 28}
PE2#
```

PE1->PE2 Validated for the Labels


RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

74

MPLS VPN Fwd Plane—Trouble #1 (Cont.)

Cisco.com



- Ok... Although the pings failed on PE1, ICMP debugs on PE2 confirms that **PE1->PE2 LSP is error free**
- Let's ping in the other direction to find out the opposite path

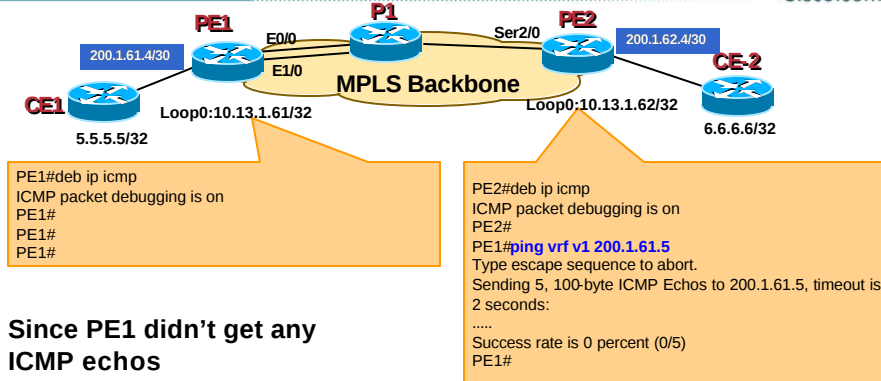
RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

75

MPLS VPN Fwd Plane—Trouble #1 (Cont.)

Cisco.com



Since PE1 didn't get any ICMP echos

- either PE2->PE1 LSP is broken
- ~~or PE1 doesn't have the LFIB entry for 200.1.61.5~~
- or PE1 is dropping the received MPLS packets for some reason

We Already
Verified this Earlier

Ok... so let's troubleshoot for (a) first.


RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

76

MPLS VPN Fwd Plane—Trouble #1 (Cont.)

Cisco.com



- IP reachability is confirmed between PE1 and PE2. GOOD. But that doesn't validate the LSP.

We will have to check P1's LFIB to confirm whether it has correct label for PE1.

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

77

MPLS VPN Fwd Plane—Trouble #1 (Cont.)

Cisco.com

- Remember—"untagged" outgoing label means that get rid of the label stack; Hence, the VPN label is lost at P1
- "untagged" label for /32 routes inside the MPLS core is almost always bad
- To fix this "untagged" problem,

Validate LIB bindings on P1

If fine, then "clear ip route 10.13.1.61" on P1

If the above doesn't fix, then "deb mpls lfib cef" to dig further

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

78

MPLS VPN Fwd Plane—Trouble #1 (Cont.)

Cisco.com

- **Although less reliable**, traceroute can also be used to identify “untagged” problem, as shown:

```
RSP-PE-SOUTH-3#traceroute 10.13.1.23
Type escape sequence to abort.
Tracing the route to 10.13.1.23
 1 10.13.6.25 [MPLS: Label 66 Exp 0] 0 msec 4 msec 0 msec
 2 10.13.2.65 [MPLS: Label 75 Exp 0] 4 msec 0 msec 0 msec
 3 10.13.2.77 [MPLS: Label 70 Exp 0] 48 msec 80 msec 268 msec
 4 10.13.2.38 0 msec 4 msec 0 msec
 5 10.13.3.94 0 msec * 0 msec
RSP-PE-SOUTH-3#
```

No Label

After the problem got fixed, the traceroute output correctly showed the label at the step 4 -

```
RSP-PE-SOUTH-3#traceroute 10.13.1.23
Type escape sequence to abort.
Tracing the route to 10.13.1.23
 1 10.13.6.25 [MPLS: Label 66 Exp 0] 0 msec 4 msec 0 msec
 2 10.13.2.65 [MPLS: Label 75 Exp 0] 4 msec 0 msec 0 msec
 3 10.13.2.77 [MPLS: Label 70 Exp 0] 48 msec 80 msec 268 msec
 4 10.13.2.38 [MPLS: Label 41 EXP 0] 0 msec 4 msec 0 msec
 5 10.13.3.94 0 msec * 0 msec
RSP-PE-SOUTH-3#
```

RST-3061
8186_05_2003_c1


© 2003, Cisco Systems, Inc. All rights reserved.

79

MPLS VPN Fwd Plane—Trouble #2

Cisco.com

#2: VPN connectivity is broken bet CEs



- Similar to #1, traffic could be dropped due to incorrect label(s)
- **IGP Label Mismatch on P/PE** (can be detected by PE-to-PE traceroute)
- Let's troubleshoot using the steps outlined in the slide#71

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

80

MPLS VPN Fwd Plane—Trouble #2 (Cont.)

Cisco.com

(step 1) Try VRF pinging PE-to-PE -

```
RSP-PE-WEST-4#ping vrf v39 30.39.130.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 30.39.130.4, timeout is 2
seconds:
.....
Success rate is 0 percent (0/5)
RSP-PE-WEST-4#
```

Ooops... ping failed. [Jump to Step 3](#). Let's verify the label information in the FIB for the VPN prefix

```
RSP-PE-WEST-4#sh ip cef vrf v39 30.39.130.4
30.39.130.4/30, version 16, epoch 0, cached adjacency to POS4/0/0
0 packets, 0 bytes
Flow: AS 0, mask 30
tag information set
local tag: VPN-route-head
fast tag rewrite with P04/0/0, point2point, tags imposed: {154 19}
via 217.60.217.3, 0 dependencies, recursive
next hop 10.13.7.37, POS4/0/0 via 217.60.217.3/32
valid cached adjacency
tag rewrite with P04/0/0, point2point, tags imposed: {154 19}
RSP-PE-WEST-4#
```

Looks good. Let's check the IP connectivity to the BGP next-hop 217.60.217.3

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

81

MPLS VPN Fwd Plane—Trouble #2 (Cont.)

Cisco.com

(Step 4) Try pinging the remote PE

```
RSP-PE-WEST-4#ping 217.60.217.3 source 10.13.1.74
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 217.60.217.3, timeout is 2
seconds:
Packet sent with a source address of 10.13.1.74
.....
Success rate is 0 percent (0/5)
RSP-PE-WEST-4#
```

Ooops,..we can't even ping. ☹️ [So, we get to Step \(5\)](#).

(Step 6) Let's check for the route and the label to 217.60.217.3 hop-by-hop. Pick up the physical next-hop and telnet to it.

```
RSP-PE-WEST-4#sh ip cef 217.60.217.3
217.60.217.3/32, version 115235, epoch 0, cached adjacency to POS4/0/0
0 packets, 0 bytes
Flow: AS 0, mask 32
tag information set
local tag: 224
fast tag rewrite with P04/0/0, point2point, tags imposed: {154}
via 10.13.7.37, POS4/0/0, 1 dependency
next hop 10.13.7.37, POS4/0/0
valid cached adjacency
tag rewrite with P04/0/0, point2point, tags imposed: {154}
RSP-PE-WEST-4#
```

154 is what this router forwards the packet with

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

82

MPLS VPN Fwd Plane—Trouble #2 (Cont.)

Cisco.com

(Step 6) So let's telnet to the first next-hop 10.13.7.37, and check the LFIB entry for 217.60.217.3 on it.

```
RSP-PE-WEST-4#telnet 10.13.7.37
Open...
```

```
GSR-P-WEST-B#sh mpls for 217.60.217.3
```

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes switched	tag interface	Outgoing interface	Next Hop
213	160	217.60.217.3/32	0		AT7/1.248	point2point

```
GSR-P-WEST-B#
```

That's it. You see the problem !!!! ☹

The previous router is ending packet destined to 217.60.217.3 with label=154, but this router expects label 213 for it.

So what happens to the MPLS packet with label=154 ? See below -

```
GSR-P-WEST-B#sh mpls for label 154
```

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes switched	tag interface	Outgoing interface	Next Hop
154	186	10.13.1.45/32	2683949106	G11/0		10.13.2.81

```
GSR-P-WEST-B#sh deb
```

MPLS packets received with label=154 are forwarded on Gig1/0 with label=186. Probably, the next-hop router is either dropping or sending the packets somewhere else.

In the absence of any entry for label=154, packets will be dropped right on this router.

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

83

MPLS VPN Fwd Plane—Trouble #2 (Cont.)

Cisco.com

- **To fix the label mismatch problem:**
 - Validate LIB bindings for the prefix
 - If LIB has correct binding, then “clear ip route <prefix>” should fix
 - If not, then LDP neighbors are out-of-sync, flap the LDP neighbor

RST-3061
8186_05_2003_c1


© 2003, Cisco Systems, Inc. All rights reserved.

84

MPLS VPN Fwd Plane—Trouble #3

Cisco.com

#3: VPN traffic of 1492B (or more) fails



- Default MTU size of Ethernet is 1500 Bytes
- Presence of 2 labels (8 Bytes) reduces the IP data from 1500B to 1492B in the packet
- Hence, PE will drop any (CE-sent) packet that has DF bit set and exceeds 1492B

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

85

MPLS VPN Fwd Plane—Trouble #3

Cisco.com

- This could be confirmed by doing pings with sweeping sizes and DF bit=1

```
GSR-PE-NTHWEST-4#ping vrf v29
Protocol [ip]:
Target IP address: 200.29.75.1
Repeat count [5]: 1
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: yes
Source address or interface:
Type of service [0]:
Set DF bit in IP header? [no]: yes
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]: yes
Sweep min size [36]: 1400
Sweep max size [18024]: 1500
Sweep interval [1]:
Type escape sequence to abort.
Sending 101, [1400..1500]-byte ICMP Echos to 200.29.75.1, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 92 percent (93/101), round-trip min/avg/max = 1/22/200 ms
GSR-PE-NTHWEST-4#
```

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

86

MPLS VPN Fwd Plane—Trouble #3

Cisco.com

To fix this problem:

- Increase the MPLS MTU size to 1508 (for 2 labels) on **all the MPLS enabled interfaces**

`"Router(config-if)#mpls mtu 1508"`

- Also make sure that the trunks on the L2 switches (if present) are enabled with jumbo frame

`"Switch(config)#set port jumbo <mod/port> enable"`

`"Switch#show port jumbo"`

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

87

Agenda

Cisco.com

- Prerequisites
- MPLS VPN Troubleshooting
- **Conclusion**

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

88

Conclusion

Cisco.com

- MPLS seems cryptic, but it is not 😊
- Whether to look at FIB or LFIB?
- Whether it is a BGP or MPLS problem?
- Whether the problem is within the core or outside the core?
- Ongoing IETF work to ease operators
LSP ping, MPLS MIBs etc.


RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

89

Q & A

Cisco.com



- Feel free to download the powerpoint preso here
<ftp://ftpeng.cisco.com/rajiva/Networkers>

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.


90

Recommended Reading

Cisco.com

**MPLS and VPN
Architectures, CCIP Edition**
ISBN: 1587050811

**MPLS and VPN
Architectures, Vol II**
ISBN: 1587051125



Available on-site at the Cisco Company Store

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

91

Visit the World of Solutions

Cisco.com

- **Learn more about products and services surrounding the technologies covered in this session in the World of Solutions.**
- **The World of Solutions is open:**
 - Tuesday: 11:00am ? 2:00pm
5:00pm ? 8:00pm**
 - Wednesday: 11:00am ? 2:00pm
5:00pm ? 7:00pm**

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

92



Cisco.com

Troubleshooting MPLS VPN Networks

Session RST-3061

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

93

NETWORKERS 2003

THE POWER TO TRANSFORM BUSINESS. **now.**

Please Complete Your Evaluation Form

Session RST-3061



RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

94

CISCO SYSTEMS



RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

95

Additional Slides

Cisco.com

- **Additional slides**

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

96

MPLS VPN Ctrl Plane— How/What to Configure

Cisco.com

- Let's go through the configuration steps and understand the **relevant show commands** at each step
- eBGP is the chosen PE-CE protocol

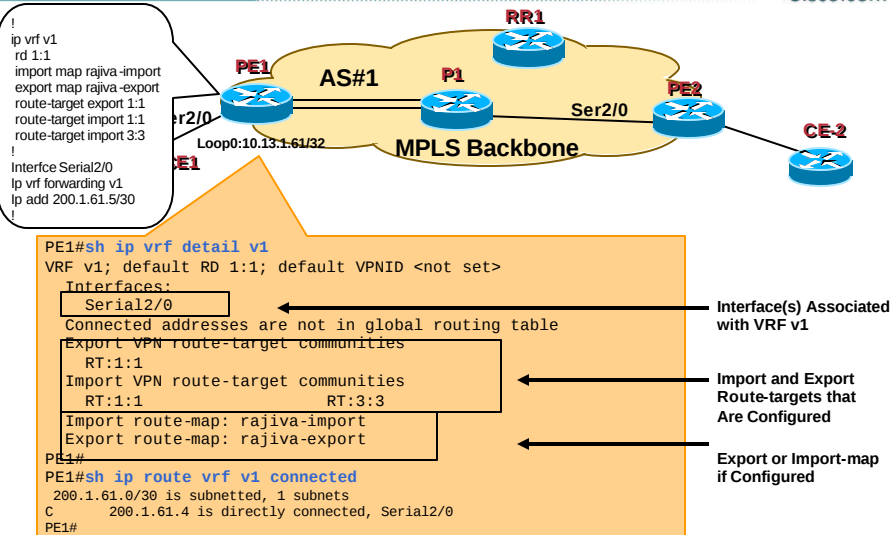
RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

97

MPLS VPN Control Plane—PE-CE int

Cisco.com



RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

98

MPLS VPN Control Plane— PE-CE Protocol

Cisco.com

router bgp 1

```
!
address-family ipv4 vrf v1
 redistribute connected
 neighbor 200.1.61.6 remote-as 65000
 neighbor 200.1.61.6 activate
 neighbor 200.1.61.6 as-override
 no auto-summary
 exit-address-family
!
```

PE1#sh ip bgp vpnv4 vrf v1 summary

```
BGP router identifier 10.13.1.61, local AS number 1
BGP table version is 2818, main routing table version 2818
3 network entries using 363 bytes of memory
3 path entries using 192 bytes of memory
8 BGP path attribute entries using 480 bytes of memory
1 BGP extended community entries using 24 bytes of memory
.....rest is deleted.....
BGP activity 19/12 prefixes, 1402/1394 paths, scan interval 15 secs
```

CE1

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
200.1.61.6	4	65000	5544	5540	2818	0	0	00:04:39	4

PE1#

- All the eBGP neighbors i.e. CE in VRF v1

RST-3061

8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

99

MPLS VPN Control Plane— PE-CE Protocol

Cisco.com

router bgp 1

```
!
address-family ipv4 vrf v1
 redistribute connected
 neighbor 200.1.61.6 remote-as 65000
 neighbor 200.1.61.6 activate
 neighbor 200.1.61.6 as-override
 no auto-summary
 exit-address-family
!
```

PE1#sh ip bgp vpnv4 vrf v1 neighbors 200.1.61.6 routes

```
BGP table version is 2835, local router ID is 10.13.1.61
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 1:1 (default for vrf v1)					
*> 5.5.5.5/32	200.1.61.6	0		0	65000 ?
*> 30.1.61.4/30	200.1.61.6	0		0	65000 ?
*> 30.30.30.1/32	200.1.61.6	0		0	65000 ?
* 200.1.61.4/30	200.1.61.6	0		0	65000 ?

Total number of prefixes 4

PE1#

- BGP routes received from the CE in VRF v1

RST-3061


8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

100

MPLS VPN Control Plane— PE-CE Protocol

```
router bgp 1
!
address-family ipv4 vrf v1
 redistribute connected
 neighbor 200.1.61.6 remote-as 65000
 neighbor 200.1.61.6 activate
 neighbor 200.1.61.6 as-override
 no auto-summary
 exit-address-family
!
```



```
PE1#sh ip bgp vpn vrf v1
BGP table version is 26, local router ID is 10.13.1.61
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 1:1 (default for vrf v1)					
*> 5.5.5.5/32	200.1.61.6	0	0	65000	?
*> 30.1.61.4/30	200.1.61.6	0	0	65000	?
*> 30.30.30.1/32	200.1.61.6	0	0	65000	?
* 200.1.61.4/30	200.1.61.6	0	0	65000	?
*>	0.0.0.0	0	32768	?	

- BGP routes in VRF v1 (from CE and PEs)


RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

101

MPLS VPN Control Plane— PE-CE Protocol

```
router bgp 1
!
address-family ipv4 vrf v1
 redistribute connected
 neighbor 200.1.61.6 remote-as 65000
 neighbor 200.1.61.6 activate
 neighbor 200.1.61.6 as-override
 no auto-summary
 exit-address-family
!
```



```
PE1#sh ip bgp vpnv4 vrf v1 200.1.61.4
BGP routing table entry for 1:1:200.1.61.4/30, version 24
Paths: (2 available, best #2, table v1)
Advertised to non peer-group peers:
```

CE1	200.1.61.6	65000
	200.1.61.6 from 200.1.61.6 (20.20.20.1)	
	Origin incomplete, metric 0, localpref 100, valid, external	
	Extended Community: RT:1:1	
	Local	
	0.0.0.0 from 0.0.0.0 (10.13.1.61)	
	Origin incomplete, metric 0, localpref 100, weight 32768, valid, sourced, best	
	Extended Community: RT:1:1	
Export RT		

```
PE1#
```

- Routes in BGP table of VRF v1 (from CEs+PEs)

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

102

MPLS VPN Control Plane— PE-CE Protocol

Cisco.com

```
router bgp 1
!
address-family ipv4 vrf v1
 redistribute connected
 neighbor 200.1.61.6 remote-as 65000
 neighbor 200.1.61.6 activate
 neighbor 200.1.61.6 as-override
 no auto-summary
 exit-address-family
```

```
PE1#sh ip bgp vpnv4 vrf v1 labels
```

Network	Next Hop	In label/Out label
Route Distinguisher: 1:1 (v1)		
0.0.0.0	0.0.0.0	26/aggregate(v1)
5.5.5.5/32	200.1.61.6	27/nolabel
30.1.61.4/30	200.1.61.6	28/nolabel
30.30.30.1/32	200.1.61.6	29/nolabel
200.1.61.4/30	200.1.61.6	30/nolabel
	0.0.0.0	30/aggregate(v1)

```
PE1#
```

- Routes and labels in BGP table of VRF v1

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

103

MPLS VPN Control Plane— PE-CE Protocol

Cisco.com

```
router bgp 1
!
address-family ipv4 vrf v1
 redistribute connected
 neighbor 200.1.61.6 remote-as 65000
 neighbor 200.1.61.6 activate
 neighbor 200.1.61.6 as-override
 no auto-summary
 exit-address-family
```

```
PE1#sh mpls forwarding vrf v1
```

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes switched	Outgoing interface	Next Hop
27	Untagged	5.5.5.5/32[V]	0	Se2/0	point2point
28	Untagged	30.1.61.4/30[V]	0	Se2/0	point2point
29	Untagged	30.30.30.1/32[V]	0	Se2/0	point2point
30	Aggregate	200.1.61.4/30[V]	0		

```
PE1#
```

- Routes learned from CEs go into the LFIB

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

104

MPLS VPN Control Plane—PE-RP

```
router bgp 1
!
bgp router-id 10.13.1.61
neighbor 10.13.1.21 remote-as 1
neighbor 10.13.1.21 update-source Lo0
!
address-family vpnv4
neighbor 10.13.1.21 activate
neighbor 10.13.1.21 send-comm both
!
```

```
router bgp 1
!
bgp router-id 10.13.1.21
neighbor 10.13.1.61 remote-as 1
neighbor 10.13.1.61 update-source Lo0
!
address-family vpnv4
neighbor 10.13.1.61 activate
neighbor 10.13.1.61 send-comm both
!
```

```
PE1#sh ip bgp vpnv4 all summary
BGP router identifier 10.13.1.61, local AS number 1
BGP table version is 26, main routing table version 26
5 network entries using 605 bytes of memory
6 path entries using 384 bytes of memory
.....deleted.....
BGP using 1361 total bytes of memory
BGP activity 26/20 prefixes, 1428/1421 paths, scan interval 15 secs
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.13.1.21	4	1	6240	7612	26	0	0	00:28:00	1
200.1.61.6	4	65000	5594	5596	26	0	0	00:31:22	4

PE1#

- Lists PE's not only MP-BGP peers, but also CE peers

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

105

MPLS VPN Control Plane—RR

```
router bgp 1
!
bgp router-id 10.13.1.21
neighbor 10.13.1.62 remote-as 1
neighbor 10.13.1.62 update-source Lo0
!
address-family vpnv4
neighbor 10.13.1.62 activate
neighbor 10.13.1.62 send-comm both
!
```

```
router bgp 1
!
bgp router-id 10.13.1.62
neighbor 10.13.1.21 remote-as 1
neighbor 10.13.1.21 update-source Lo0
!
address-family vpnv4
neighbor 10.13.1.21 activate
neighbor 10.13.1.21 send-comm both
!
```

```
RR1#sh ip bgp vpnv4 all labels
Network      Next Hop
Route Distinguisher: 1:1
5.5.5.5/32    10.13.1.61    nolabel/27
30.1.61.4/30  10.13.1.61    nolabel/28
30.30.30.1/32 10.13.1.61    nolabel/29
200.1.61.4/30 10.13.1.61    nolabel/30
200.1.62.4/30 10.13.1.62    nolabel/25
```

```
RR1#
RR1#sh mpls forwarding
Local Outgoing Prefix Bytes tag Outgoing Next Hop
tag tag or VC or Tunnel Id switched interface
RR1#
```

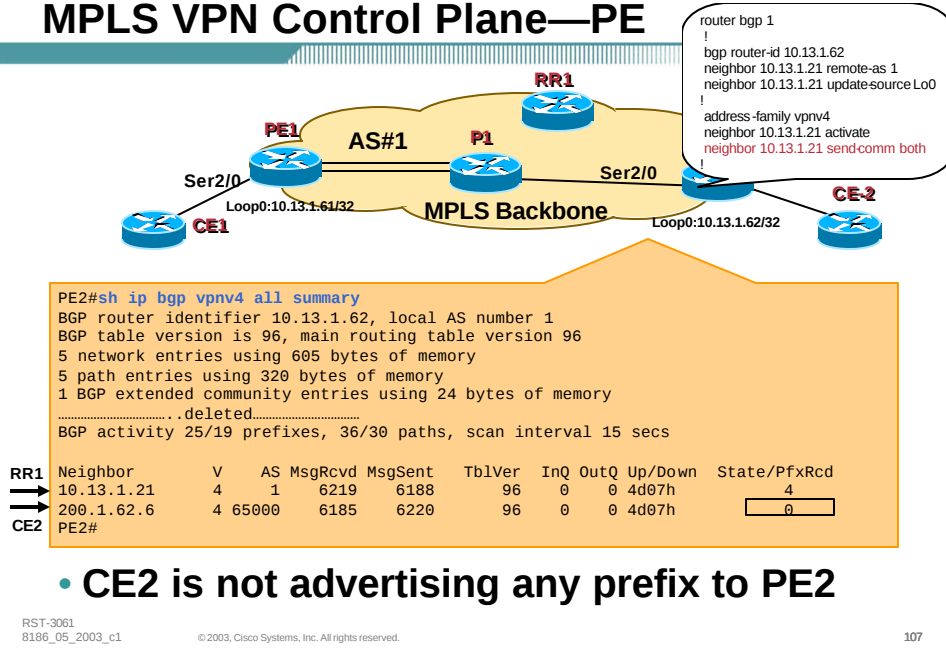
- RR **shouldn't** allocate any local label; hence, LFIB **shouldn't** have any VPN prefix

RST-3061
8186_05_2003_c1


© 2003, Cisco Systems, Inc. All rights reserved.

106


MPLS VPN Control Plane—PE



MPLS VPN Control Plane—PE



MPLS VPN Control Plane—PE




RST-3061

8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

109

MPLS VPN Control Plane—PE



- 200.1.61.4 is accepted since its RT=1:1 matches with “import RT” of VRF v1 on PE2


RST-3061

8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

110

MPLS VPN Control Plane—PE

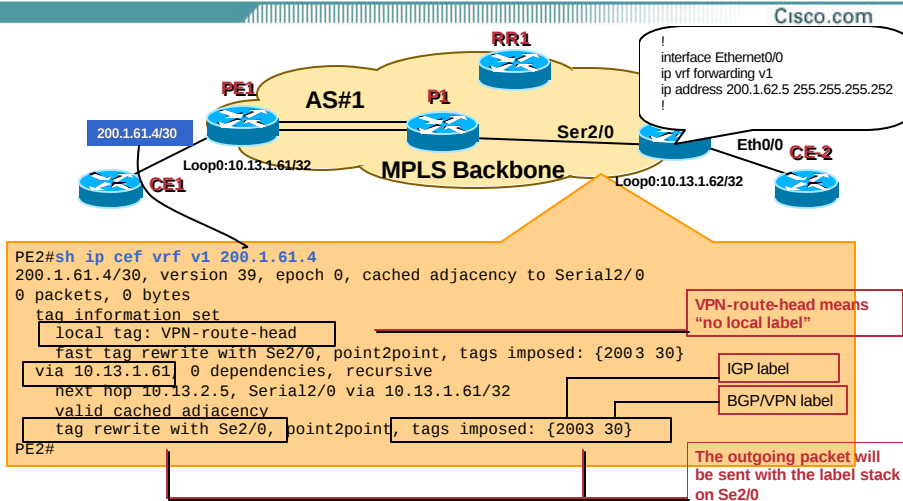


RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

111

MPLS VPN Control Plane—PE




- Traffic received on Eth0/0 will be an IP traffic, hence PE2 will do a CEF lookup in the VRF v1

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

112

MPLS VPN Control Plane—PE



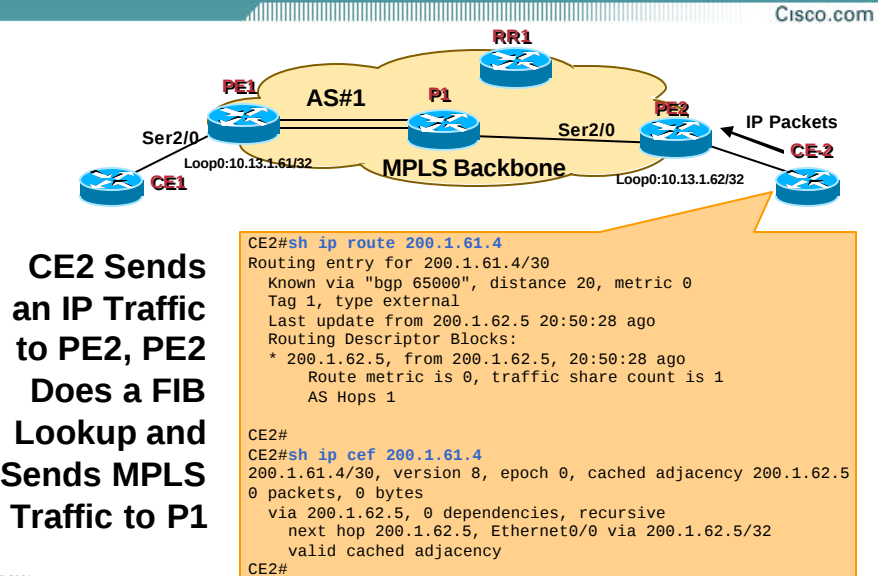
- CE learned VPN routes must be in the LFIB
- PE1's advertised VPN routes shouldn't be in the PE2's LFIB; **no need**

RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

113

MPLS VPN Control Plane—PE



RST-3061
8186_05_2003_c1

© 2003, Cisco Systems, Inc. All rights reserved.

114