VPRN Inter-AS VPN Model C

In This Chapter

This section provides following information.

Topics in this section include:

Applicability

This example is applicable to all of the 7750 and 7710 SR series and was tested on release 7.0R5. There are no pre-requisites for this configuration. This is supported on 7450 ESS-7 or ESS-12 in mixed-mode since 8.0R1. The 7750 SR-c4 is supported from 8.0R4 and higher.

Overview

Introduction

Section 10 of RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*, describes three potential methods for service providers to interconnect their IP-VPN (Internet Protocol — Virtual Private Network) backbones in order to provide an end-to-end MPLS-VPN where one or more sites of the VPN are connected to different service provider autonomous systems. The purpose of this section is to describe the configuration and troubleshooting for inter-AS VPN model C.

In this architecture, VPN prefixes are neither held, nor re-advertised by the Autonomous System Border Router — Provider Edge (ASBR-PE) routers. The ASBR-PE does however maintain labeled IPv4 /32 routes to other PE routers within its own AS. It then redistributes these /32 IPv4 prefixes in external Border Gateway Protocol (eBGP) to the ASBR-PE in other service providers ASs. Using this methodology, it is possible for PE routers in different ASs to establish multi-hop Multi Protocol — external Border Gateway Protocol (MP-eBGP) sessions to each other in order to exchange customer VPN prefixes over those connections.

To be more specific, the /32 IPv4 routes for the PE routers in the other service providers AS will need to be redistributed into the interior Gateway Protocol (IGP) in the local AS together with an assigned label. As most service providers do not like redistribution of loop-back addresses from another service provider into the local IGP, a potential solution can be found by imposing a three-level label stack on the ingress PE. The bottom-level label would be assigned by the egress PE (advertised in multi-hop MP-eBGP without next-hop override) and is commonly referred to as the VPN-label. The middle label would be assigned by the local ASBR-PE and would correspond to the /32 route of the egress PE (in a different AS) using BGP-LBL (RFC 3107, *Carrying Label Information in BGP-4*). The top level label would then be the label assigned by the local ASBR-PE(s) /32 loop-back address, which would be assigned by the IGP next-hop of the ingress PE. This label is referred to as the LDP-LBL. Figure 125 reflects this mechanism. The VPN-LBL is assigned by PE-5, the BGP-LBL is assigned by PE-4 and the LDP-LBL is also assigned by PE-4.

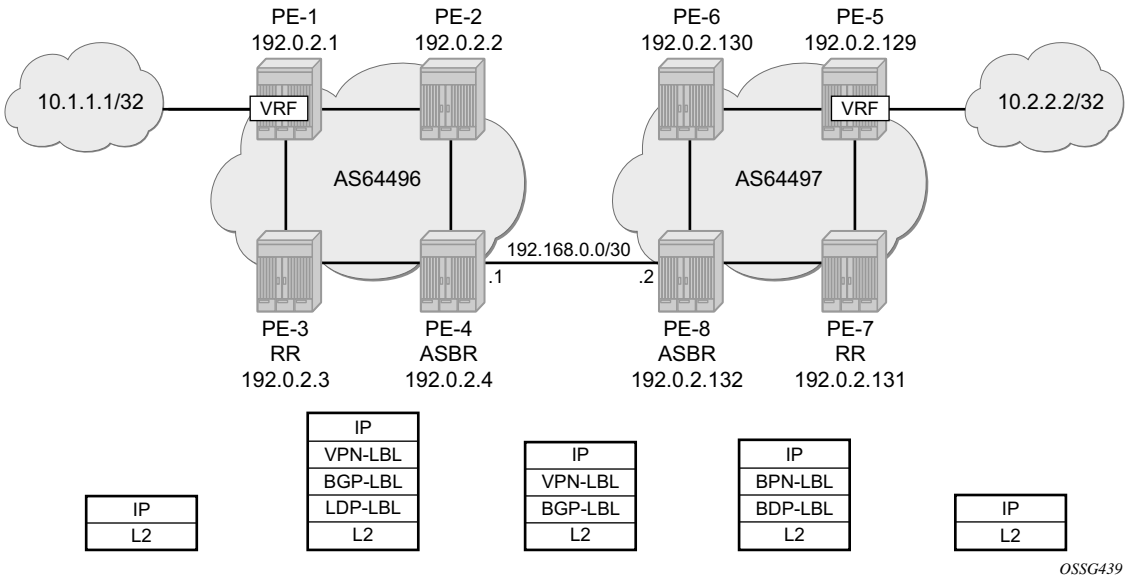

Figure 125: Inter-AS VPN Model C

The VPN connectivity is established using Labeled VPN route exchange using MP-eBGP without next-hop override. The PE connectivity will be established as described below.

EBGP PE /32 loopback leaking routing exchange using eBGP LBL (RFC 3107) at the ASBR-PE. The /32 PE routes learned from the other AS through the ASBR-PE are further distributed into the local AS using iBGP and optionally Route Reflectors (RRs). This model uses a three label stack and is referred to as Model C. Resilience for ASBR-PE failures is dependent on BGP.

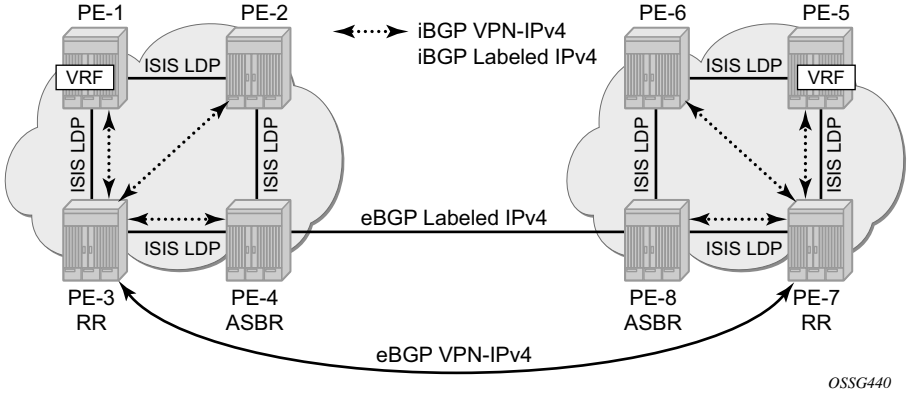

Figure 126: Protocol Overview

Figure 126 gives an overview of all protocols used when implementing Inter-AS Model C. Inside each AS there is an ISIS adjacency and a link LDP session between each pair of adjacent nodes. As an alternative, OSPF can be used as IGP. Also there is an iBGP session between each PE and the RR. The address family is both VPN-IPv4 for the exchange of customer VPN prefixes and Labeled IPv4 for the exchange of labeled IPv4 prefixes. Note that as an alternative, a full mesh of iBGP sessions can be used in each AS.

Between the ASBRs there is an eBGP sessions for the exchange of labeled IPv4 prefixes. The ASBRs will override the next-hop for those prefixes. Between the RRs in the different ASs there is an eBGP session for the exchange of VPN customer prefixes. The RRs will not override the next-hop for those prefixes.

The big advantage of this model is that no VPN routes need to be held on the ASBR-PEs and as such it scales the best among all the three Inter-AS IP-VPN models. However, leaking /32 PE addresses between service providers creates some security concerns. As such we see Model C typically deployed within a service provider network.

The network topology is displayed in Figure 125. The setup consists of two times four (2 x 4) 7750/7710 nodes located in different autonomous systems. There is an AS interconnection from ASBR PE-4 to ASBR PE-8. PE-3 and PE-7 will act as RRs for their AS. It is assumed that an IP-VPN is already configured in each AS. Following configuration tasks should be done first:

- ISIS or OSPF on all interfaces within each of the ASs.
- LDP on all interfaces within each of the ASs.
- MP-iBGP sessions between the PE routers and the RRs in each of the ASs.
- IP-VPN on PE-1 and on PE-5 with identical route targets.
- A loopback interface in the VRF on PE-1 and PE-5.

Configuration

The first step is to configure a MP-eBGP session between the ASBRs in both ASs. This session will be used to redistribute labelled IPv4 routes for the /32 system IP addresses between the AS□. These MP-BGP extensions are described in RFC 3107.

The configuration for ASBR PE-4 is displayed below. The **advertise-label ipv4** command is required to enable the advertising of labelled IPv4 routes. Note that this command is also required on the RR neighbor in order to propagate the labelled IPv4 routes towards the other PEs in the AS. The address family for labelled IPv4 routes is IPv4 so this family must be enabled for the peering with the RR.

```
configure router bgp
            group "rr"
                family ipv4 vpn-ipv4
                neighbor 192.0.2.3
                    advertise-label ipv4
                exit
            exit
            group "remote-as"
                family ipv4
                type external
                peer-as 64497
                neighbor 192.168.0.2
                    advertise-label ipv4
                exit
            exit
exit all
```

Note that address **family vpn-ipv4** is also required to advertise IPv4 customer routes within the AS. On the RR, the **advertise-label ipv4** command must be specified for each PE neighbor. Also note that address family IPv4 must be enabled. The configuration for RR PE-3 is displayed below.

```
configure router bgp
  group "rr-clients"
                family ipv4 vpn-ipv4
                neighbor 192.0.2.1
                    advertise-label ipv4
                exit
                neighbor 192.0.2.2
                    advertise-label ipv4
                exit
                neighbor 192.0.2.4
                    advertise-label ipv4
                exit
            exit
exit all
```

On the remaining PE nodes in AS 64496, the **advertise-label ipv4** command must be specified on the RR neighbor. Also the IPv4 family must be enabled.

```
configure router bgp
     group rr
         family ipv4 vpn-ipv4
         neighbor 192.0.2.3
             advertise-label ipv4
         exit
      exit
exit all
```

The configuration for the nodes in AS64497 is very similar. The IP addresses can be derived from Figure 125.

On ASBR PE-4, verify that the BGP session with ASBR PE-8 is up:

```
A:PE-4# show router bgp neighbor 192.168.0.2
===============================================================================
BGP Neighbor
===============================================================================
-------------------------------------------------------------------------------
Peer  : 192.168.0.2
Group : remote-as
-------------------------------------------------------------------------------
Peer AS               : 64497            Peer Port            : 179
Peer Address          : 192.168.0.2
Local AS              : 64496            Local Port           : 51262
Local Address         : 192.168.0.1
Peer Type             : External
State                 : Established      Last State           : Active
Last Event            : recvKeepAlive
Last Error            : Cease
Local Family          : IPv4
Remote Family         : IPv4
Hold Time             : 90               Keep Alive           : 30
Active Hold Time      : 90               Active Keep Alive    : 30
Cluster Id            : None
Preference            : 170              Num of Flaps         : 4
Recd. Paths           : 0
IPv4 Recd. Prefixes   : 0                IPv4 Active Prefixes : 0
IPv4 Suppressed Pfxs  : 0                VPN-IPv4 Suppr. Pfxs : 0
VPN-IPv4 Recd. Pfxs   : 0                VPN-IPv4 Active Pfxs : 0
Mc IPv4 Recd. Pfxs.   : 0                Mc IPv4 Active Pfxs. : 0
Mc IPv4 Suppr. Pfxs   : 0                IPv6 Suppressed Pfxs : 0
IPv6 Recd. Prefixes   : 0                IPv6 Active Prefixes : 0
VPN-IPv6 Recd. Pfxs   : 0                VPN-IPv6 Active Pfxs : 0
VPN-IPv6 Suppr. Pfxs  : 0                L2-VPN Suppr. Pfxs   : 0
L2-VPN Recd. Pfxs     : 0                L2-VPN Active Pfxs   : 0
MVPN-IPv4 Suppr. Pfxs: 0                 MVPN-IPv4 Recd. Pfxs : 0
MVPN-IPv4 Active Pfxs: 0
Input Queue           : 0                Output Queue         : 0
i/p Messages          : 37               o/p Messages         : 39
i/p Octets            : 891              o/p Octets           : 891
i/p Updates           : 4                o/p Updates          : 4
TTL Security          : Disabled         Min TTL Value        : n/a
```

```
Graceful Restart     : Disabled         Stale Routes Time    : n/a
Advertise Inactive   : Disabled         Peer Tracking        : Disabled
Advertise Label      : ipv4
Auth key chain       : n/a
Bfd Enabled          : Disabled
Local Capability     : RtRefresh MPBGP 4byte ASN
Remote Capability    : RtRefresh MPBGP 4byte ASN
Import Policy        : None Specified / Inherited
Export Policy        : None Specified / Inherited

-------------------------------------------------------------------------------
Neighbors : 1
===============================================================================
A:PE-4#
```

Note that both ASBRs have MPBGP capabilities. At this time, no prefixes have been received from the remote ASBR. To enable the advertising of labelled IPv4 routes for the system loopback interfaces, an export policy must be created and applied to the BGP session on both ASBRs. The policy configuration is displayed below for ASBR PE-4. Note that the configuration for ASBR PE-8 is very similar, the IP addresses can be derived from Figure 125.

```
configure router policy-options
     prefix-list "pe_sys"
         prefix 192.0.2.128/25 longer
     exit
     policy-statement "pe-sys-to-bgp"
         entry 10
             from
                 prefix-list "pe-sys"
             exit
             to
                 protocol bgp
             exit
             action accept
             exit
         exit
     exit
exit all
configure router bgp
     group remote-as
         neighbor 192.168.0.2
             export "pe-sys-to-bgp"
         exit
     exit
exit all
```

After creating and applying the export policies on both ASBRs, labelled IPv4 routes will be advertised towards the remote AS for system IP addresses of the PE nodes in the local AS.

On ASBR PE-4, verify if labelled IPv4 routes have been received from ASBR PE-8:

```
A:PE-4# show router bgp routes
===============================================================================
 BGP Router ID:192.0.2.4        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
BGP IPv4 Routes
===============================================================================
Flag  Network                                           LocalPref   MED
      Nexthop                                                       VPNLabel
      As-Path
-------------------------------------------------------------------------------
u*>i  192.0.2.129/32                                    None        20
      192.168.0.2                                                   -
      64497
u*>i  192.0.2.130/32                                    None        10
      192.168.0.2                                                   -
      64497
u*>i  192.0.2.131/32                                    None        10
      192.168.0.2                                                   -
      64497
u*>?  192.0.2.132/32                                    None        None
      192.168.0.2                                                   -
      64497
-------------------------------------------------------------------------------
Routes : 4
===============================================================================
A:PE-4#
```

As can be seen from the output above, 4 labelled IPv4 routes have been received. One route for every system IP address in the remote AS with a label attached.

The actual labels can be seen with following command:

```
A:PE-4# show router bgp routes 192.0.2.129/32 hunt
===============================================================================
 BGP Router ID:192.0.2.4        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
BGP IPv4 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : 192.0.2.129/32
Nexthop        : 192.168.0.2
From           : 192.168.0.2
Res. Nexthop   : 192.168.0.2
Local Pref.    : None                  Interface Name : int-PE-4-PE-8
Aggregator AS  : None                  Aggregator     : None
```

```
Atomic Aggr.   : Not Atomic             MED            : 20
Community      : No Community Members
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.132
IPv4 Label     : 131065
Flags          : Used  Valid  Best  IGP
AS-Path        : 64497
-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Network        : 192.0.2.129/32
Nexthop        : 192.0.2.4
To             : 192.0.2.3
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : 20
Community      : No Community Members
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.3
IPv4 Label     : 131062
Origin         : IGP
AS-Path        : 64497
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
A:PE-4#
```

Note that in the RIB In entries, the received label from PE-8 can be seen (131065). In the RIB Out entries, the locally assigned label for this prefix can be seen (131062). The label mapping can also be seen with following command:

```
A:PE-4# show router bgp inter-as-label
===============================================================================
BGP Inter-AS labels
===============================================================================
NextHop                       Received       Advertised     Label
                              Label          Label          Origin
-------------------------------------------------------------------------------
192.0.2.1                     0              131065         Internal
192.168.0.2                   131064         131061         External
192.168.0.2                   131065         131062         External
192.168.0.2                   131066         131060         External
192.168.0.2                   131067         131063         External
192.0.2.2                     0              131064         Internal
192.0.2.3                     0              131066         Internal
192.0.2.4                     0              131067         Edge
===============================================================================
A:PE-4#
```

Verify that the routes have been installed in the routing table:

```
A:PE-4# show router route-table
===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix                                     Type    Proto     Age         Pref
       Next Hop[Interface Name]                                       Metric
-------------------------------------------------------------------------------
192.0.2.1/32                                    Remote  ISIS      02h24m15s   18
       192.168.3.1                                                    20
192.0.2.2/32                                    Remote  ISIS      02h24m15s   18
       192.168.3.1                                                    10
192.0.2.3/32                                    Remote  ISIS      02h27m29s   18
       192.168.4.1                                                    10
192.0.2.4/32                                    Local   Local     02h27m35s   0
       system                                                         0
192.0.2.129/32                                  Remote  BGP       00h03m54s   170
       192.168.0.2                                                    0
192.0.2.130/32                                  Remote  BGP       00h03m54s   170
       192.168.0.2                                                    0
192.0.2.131/32                                  Remote  BGP       00h03m54s   170
       192.168.0.2                                                    0
192.0.2.132/32                                  Remote  BGP       00h03m54s   170
       192.168.0.2
...
===============================================================================
A:PE-4#
```

Verify that the BGP routes are further advertised towards all the PEs in the AS (through the RR) and are installed in the routing table on all PEs by using the above command on the other PEs.

At this point, all PEs in one AS have the /32 system IPs of the remote PEs in their routing table. All PEs in one AS have also received labels for all /32 system IPs of the remote PEs. Now a MP-eBGP session can be created between the RRs in the different ASs to exchange VPN-IPv4 routes.

The configuration for RR PE-3 is displayed below. The configuration for RR PE-7 is very similar. The IP addresses can be derived from Figure 126.

```
configure router bgp
     group "remote-as-rr"
         family vpn-ipv4
         multihop 10
         peer-as 64497
         neighbor 192.0.2.131
         exit
     exit
exit all
```

On the RRs, verify that the MP-eBGP session is up:

```
A:PE-3# show router bgp neighbor 192.0.2.131
===============================================================================
BGP Neighbor
===============================================================================
-------------------------------------------------------------------------------
Peer  : 192.0.2.131
Group : remote-as-rr
-------------------------------------------------------------------------------
Peer AS               : 64497            Peer Port            : 179
Peer Address          : 192.0.2.131
Local AS              : 64496            Local Port           : 49714
Local Address         : 192.0.2.3
Peer Type             : External
State                 : Established      Last State           : Active
Last Event            : recvKeepAlive
Last Error            : Unrecognized Error
Local Family          : VPN-IPv4
Remote Family         : VPN-IPv4
Hold Time             : 90               Keep Alive           : 30
Active Hold Time      : 90               Active Keep Alive    : 30
Cluster Id            : None
Preference            : 170              Num of Flaps         : 0
Recd. Paths           : 1
IPv4 Recd. Prefixes   : 0                IPv4 Active Prefixes : 0
IPv4 Suppressed Pfxs  : 0                VPN-IPv4 Suppr. Pfxs : 0
VPN-IPv4 Recd. Pfxs   : 1                VPN-IPv4 Active Pfxs : 0
Mc IPv4 Recd. Pfxs.   : 0                Mc IPv4 Active Pfxs. : 0
Mc IPv4 Suppr. Pfxs   : 0                IPv6 Suppressed Pfxs : 0
IPv6 Recd. Prefixes   : 0                IPv6 Active Prefixes : 0
VPN-IPv6 Recd. Pfxs   : 0                VPN-IPv6 Active Pfxs : 0
VPN-IPv6 Suppr. Pfxs  : 0                L2-VPN Suppr. Pfxs   : 0
L2-VPN Recd. Pfxs     : 0                L2-VPN Active Pfxs   : 0
MVPN-IPv4 Suppr. Pfxs: 0                 MVPN-IPv4 Recd. Pfxs : 0
MVPN-IPv4 Active Pfxs: 0
Input Queue           : 0                Output Queue         : 0
i/p Messages          : 14               o/p Messages         : 14
i/p Octets            : 370              o/p Octets           : 370
i/p Updates           : 1                o/p Updates          : 1
TTL Security          : Disabled         Min TTL Value        : n/a
Graceful Restart      : Disabled         Stale Routes Time    : n/a
Advertise Inactive    : Disabled         Peer Tracking        : Disabled
Advertise Label       : None
Auth key chain        : n/a
Bfd Enabled           : Disabled
Local Capability      : RtRefresh MPBGP ORFSendExComm ORFRecvExComm 4byte ASN
Remote Capability     : RtRefresh MPBGP ORFSendExComm ORFRecvExComm 4byte ASN
Import Policy         : None Specified / Inherited
Export Policy         : None Specified / Inherited
-------------------------------------------------------------------------------
Neighbors : 1
===============================================================================
A:PE-3#
```

The BGP session is established. Note that 1 VPN-IPv4 prefix has been received for the remote AS.

Now the VPRNs on PE-1 in AS64496 and PE-5 in AS64497 are interconnected. Packets originating in AS 64496 with a destination in AS 64497 will have 3 labels in AS 64496. Originate a VPRN ping on PE-1 towards the VPRN loopback IP address on PE-5:

```
A:PE-1# ping router 1 10.2.2.2
PING 10.2.2.2 56 data bytes
64 bytes from 10.2.2.2: icmp_seq=1 ttl=64 time=7.50ms.
64 bytes from 10.2.2.2: icmp_seq=2 ttl=64 time=3.77ms.
64 bytes from 10.2.2.2: icmp_seq=3 ttl=64 time=3.80ms.
64 bytes from 10.2.2.2: icmp_seq=4 ttl=64 time=3.77ms.
64 bytes from 10.2.2.2: icmp_seq=5 ttl=64 time=3.78ms.

---- 10.2.2.2 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 3.77ms, avg = 4.52ms, max = 7.50ms, stddev = 1.49ms
```

The top label is the LDP label to reach the exit point of the AS (PE-4). This label can be seen with following command on PE-1:

```
A:PE-1# show router ldp bindings prefix 192.0.2.4/32 active
===============================================================================
Legend:  (S) - Static
===============================================================================
LDP Prefix Bindings (Active)
===============================================================================
Prefix                Op   IngLbl    EgrLbl    EgrIntf/LspId  EgrNextHop
-------------------------------------------------------------------------------
192.0.2.4/32          Push   --      131069    1/1/2          192.168.1.2
192.0.2.4/32          Swap 131068    131069    1/1/2          192.168.1.2
-------------------------------------------------------------------------------
No. of Prefix Bindings: 2
===============================================================================
A:PE-1#
```

The middle label is the label assigned by MP-BGP on the local ASBR-PE to reach the remote PE in the remote AS. This label can be seen with following command on PE-1:

```
A:PE-1# show router bgp routes 192.0.2.129/32 hunt
===============================================================================
 BGP Router ID:192.0.2.1        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
BGP IPv4 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : 192.0.2.129/32
Nexthop        : 192.0.2.4
From           : 192.0.2.3
Res. Nexthop   : 192.168.1.2
Local Pref.    : 100                   Interface Name : int-PE-1-PE-2
```

```
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : 20
Community      : No Community Members
Cluster        : 1.1.1.1
Originator Id  : 192.0.2.4              Peer Router Id : 192.0.2.3
IPv4 Label     : 131062
Flags          : Used  Valid  Best  IGP
AS-Path        : 64497
-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
A:PE-1#
```

The bottom label is the VPN label assigned by the remote PE in the remote AS for the destination network. This label can be seen with following command on PE-1:

```
A:PE-1# show router bgp routes vpn-ipv4 10.2.2.2/32 hunt
===============================================================================
 BGP Router ID:192.0.2.1        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
-------------------------------------------------------------------------------
RIB In Entries
-------------------------------------------------------------------------------
Network        : 10.2.2.2/32
Nexthop        : 192.0.2.129
Route Dist.    : 64497:1               VPN Label      : 131070
From           : 192.0.2.3
Res. Nexthop   : n/a
Local Pref.    : 100                   Interface Name : NotAvailable
Aggregator AS  : None                  Aggregator     : None
Atomic Aggr.   : Not Atomic            MED            : None
Community      : target:64496:1
Cluster        : No Cluster Members
Originator Id  : None                  Peer Router Id : 192.0.2.3
Flags          : Used  Valid  Best  IGP
AS-Path        : 64497
VPRN Imported  :  1
-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
A:PE-1#
```

Conclusion

Inter-AS option C allows the delivery of Layer 3 VPN services to customers who have sites connected multiple ASs. This example shows the configuration of inter-AS option C (specific to this feature) together with the associated show output which can be used verify and troubleshoot it.