

In This Chapter

This section describes advanced BGP anycast configurations.

Topics in this section include:

- [Applicability on page 938](#)
- [Summary on page 939](#)
- [Overview on page 940](#)
- [Configuration on page 945](#)
- [Conclusion on page 979](#)

Applicability

The configuration in this chapter is based on SR OS 9.0R1. The feature is supported on 7750 SR-c4/c12, SR-7 and SR-12 in chassis mode **d**, and 7450 ESS-7 and ESS-12 in mixed-mode.

Summary

Release 8.0.R4 and higher provide a resiliency mechanism to protect the data flow of Layer 2/ Layer 3 traffic in the event of a complete nodal failure of the PE or ABR routers. The intent is to allow traffic to continue to flow during the convergence process and as a result reduce the amount of traffic lost during this process. This is achieved by allowing two designated routers to back each other up such that both routers store the forwarding information learned from the alternate router in a secondary forwarding table, also referred to as context-specific label space. In the event that the primary router fails, the secondary router will continue to forward traffic to the ultimate destination using the forwarding information stored in the secondary forwarding table (context-specific label space). This will continue until the rest of the network can converge to use an alternate route through the network.

BGP anycast can be used in two different scenarios.

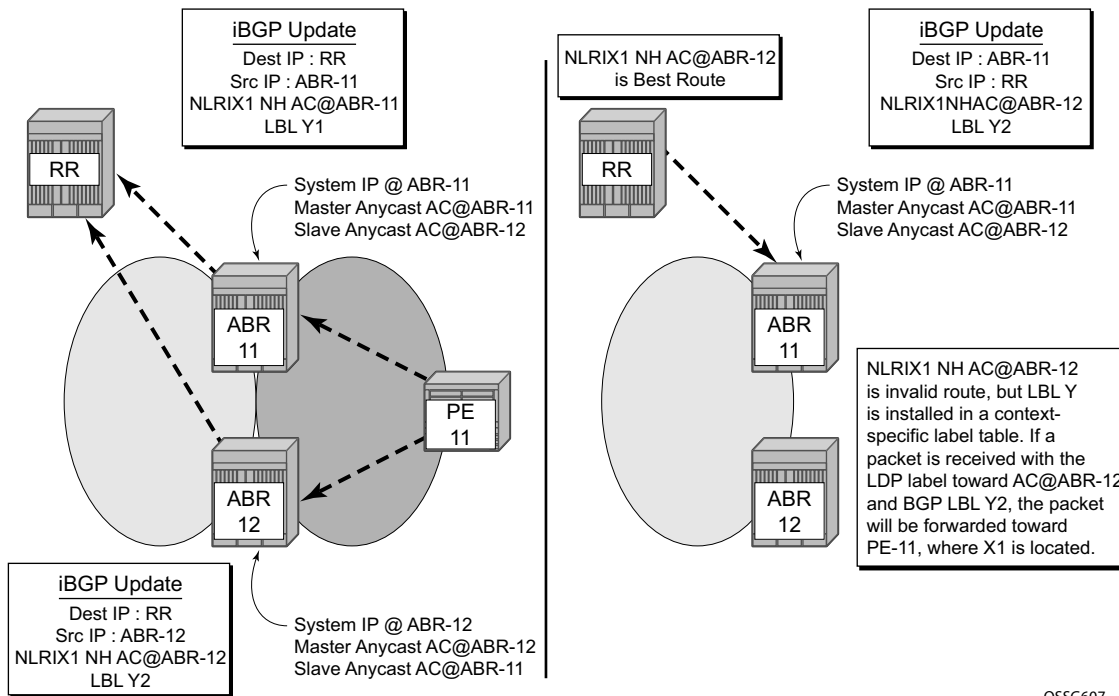
1. BGP to BGP label swapping in GRT — Where the LDP and BGP labels of an incoming IP packet will be swapped, and the packet being forwarded towards a destination in the related region/area.
2. BGP to VRF — Where an incoming packet, encapsulated with the redundant BGP service label will be forwarded towards the related VRF where an (egress) IP lookup will be performed so that the packet can be forwarded towards the correct CPE.

Overview

GRT

Assume that an access node (or similar) is connected to PE-11, represented by prefix X1/32, illustrated in Figure 130).

1. PE-11 will advertise a BGP labelled route towards its area border routers (ABRs) with its own system-IP as next hop (NH) (regular behavior).
2. Both redundant ABR (11 and 12), acting as route reflectors (RRs) for their own region, will advertise the route X1 towards the RR of the core, with their master anycast (AC) address as an NH.
3. The RR (of the core) will select the best route out of both, in this case originated from ABR-12 and reflects the route to all of the ABRs in the core.
4. The redundant ABR-11 will receive this BGP route and marks it as invalid since the NH will be a local interface, being the slave anycast interface.
5. ABR-11 will, however, install the BGP label in the context specific label space, so that incoming packets with the BGP label assigned by ABR-12 (step 2) and kept by ABR-11 (following upon the outer LDP label) can be forwarded towards destination X1.



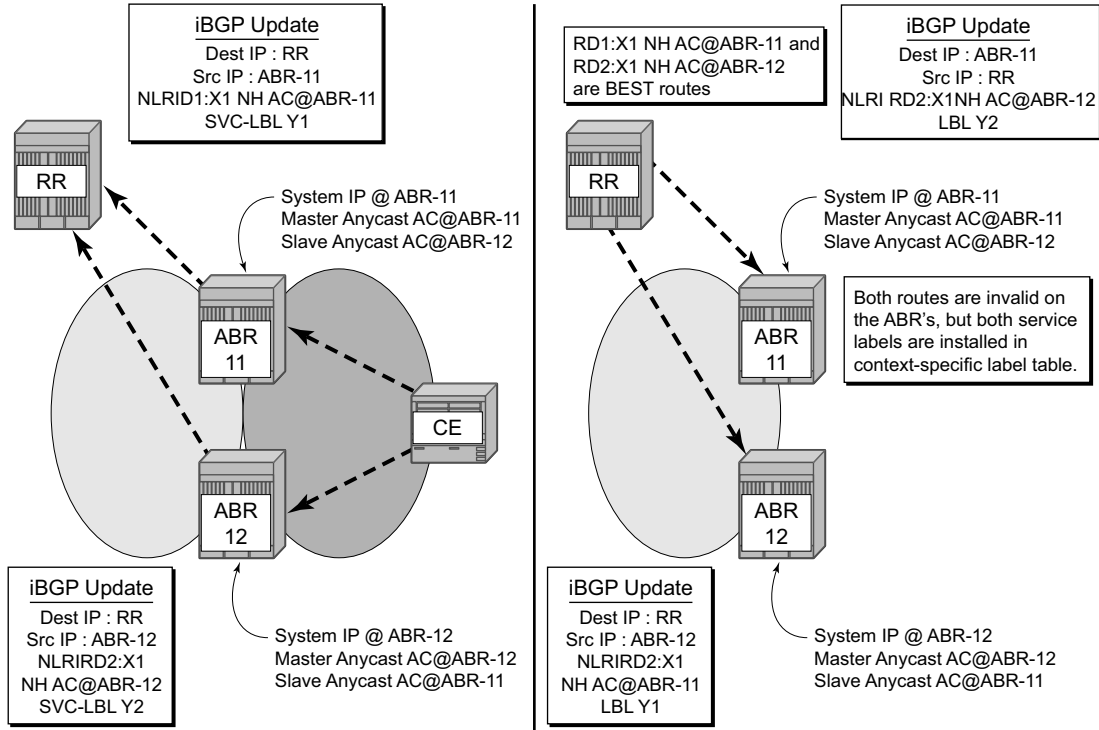
OSSG607

Figure 130: BGP Anycast Operation in GRT

Note that [Figure 130](#) shows that X1 will be advertised by both (redundant) ABRs, each assigning their master anycast address as the NH. RR will select one of them as best and will reflect the route towards both ABRs. Only one ABR (ABR-11 in this case) will install the BGP anycast route towards the destination X1. This is fine since ABR-12 will be selected by the other ABR (connected to remote regions) in the core as NH for the route towards X1/32, hence ABR-12 will receive the traffic from other regions towards X1, and if there is a failure on ABR-12 (link/nodal) traffic will be diverted to ABR-11. ABR-11 can interpret the LDP and BGP labels, which are advertised by ILDP and the BGP label (in the context-specific label space), which results in forwarding traffic to X1.

IP-VPN

In case of IP-VPN routes, the scenario is different since unique route distinguisher (RD) will result in both IP-VPN routes being active at the RR, as illustrated [Figure 131](#).



OSSG608

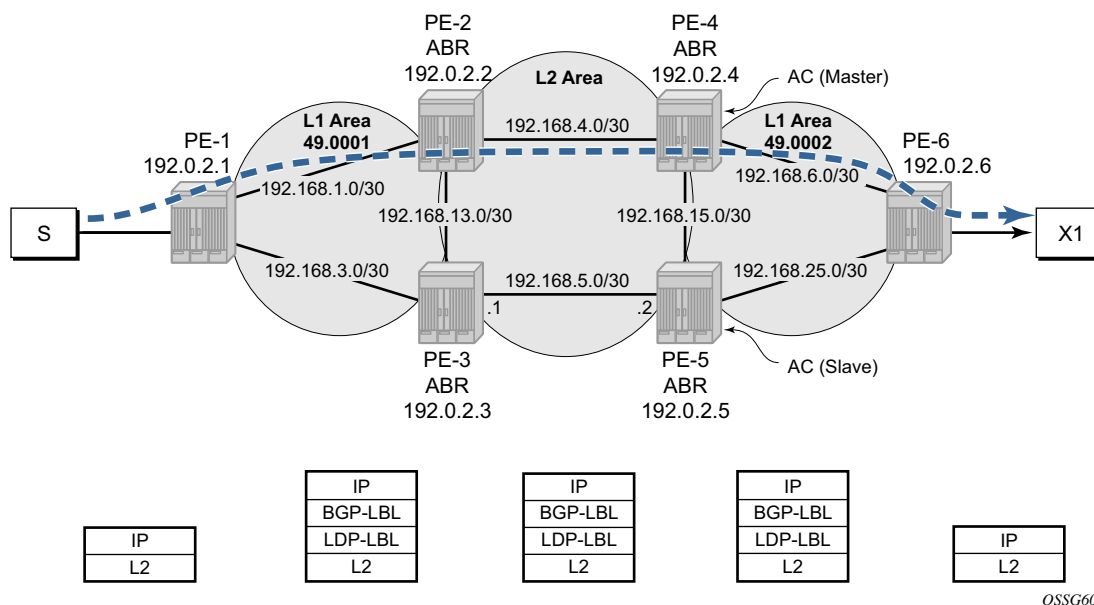
Figure 131: BGP Anycast Operation with IP-VPN

Context-Specific Label Space

A key-point of the BGP anycast feature is the presence of a context-specific label space on the Label Switch Router (LSR)/ABR. Where in normal MPLS scenarios, an LSR selects a free label from its own label-range, installs this in the Label Forwarding Information Base (LFIB) and signals this by a control protocol (BGP/LDP/RSVP), the LSR/ABR will now learn a label from a redundant LSR/ABR and installs this as an ingress label into a dedicated label space. This mechanism is different from the MPLS mechanisms used to date in the SR OS, which are based on downstream label allocation.

Data Path in GRT

Figure 132 illustrates the data path from source S to destination X1.



OSSG609

Figure 132: BGP Anycast Data Path (BGP to BGP Swapping)

An incoming IP packet on PE-1 will be encapsulated in a dual MPLS label packet, where;

1. BGP label will assure the forwarding towards the remote PE, being PE-6 in this case, PE-6 is part of the tunnel-table on PE-1, with BGP as MPLS control protocol.
2. LDP label will forward the packet to the anycast address of area 49.0002, where PE-4 is the master (in normal operations).

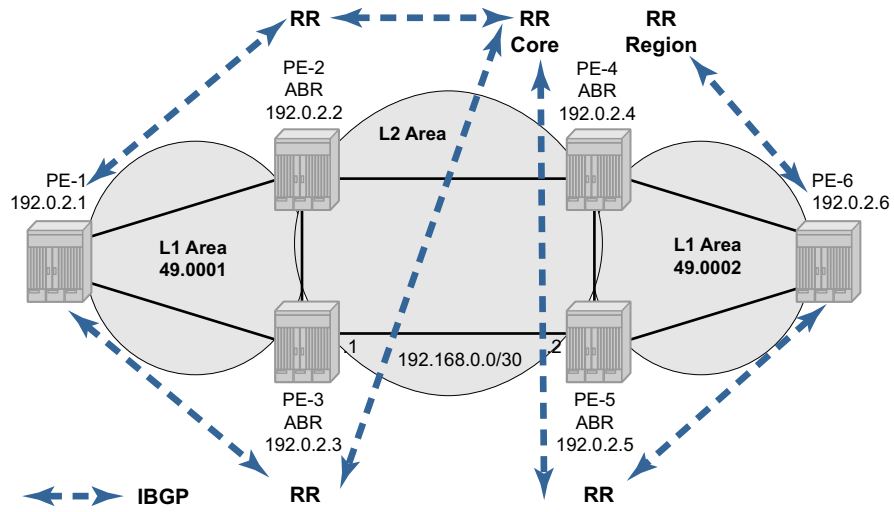
- From PE-4 onwards, both the BGP and LDP labels are swapped, and forwarded out on the interface towards PE-6.

The data path for IP-VPN routes is described in the IP-VPN routes section.

BGP Control Plane

Additional attention is required for the BGP control plane, where the ABR needs to act as RR for its related region. This is mandatory since the ABR needs to perform a NHS (next hop self) action to the BGP routes advertising the system addresses of the PE (or AN) in the region.

Note that different RR hierarchies can be created based upon different address families.



OSSG610

Figure 133: BGP Anycast BGP Topology

In this particular case, there is a cluster in the core area with PE-4 acting as RR. PE-4 is RR for 2 different clusters, 1.1.1.1 towards the core and 2.2.2.2 towards the regional PE-6. All other PEs (PE-2-3-5) in the core act as an RR for their related region.

Configuration

The following steps must be taken to configure the BGP anycast scenario.

1. Create the anycast interfaces.
2. IGP configuration in the MPLS domain.
3. Enable LDP to assure reach ability of the ABR.
4. Configure BGP to advertise the system addresses of PE-1 and PE-6 with a BGP label.
5. Configure a service to set-up end-to-end connectivity.

Anycast Interface

Each ABR will be configured with two anycast addresses, a master and slave address. The redundant ABR within the same area will be configured with the reverse, as seen in the following configuration;

On PE-2:

```
A:PE-2>config>router# info
-----
#-----
echo "IP Configuration"
#-----

      mh-primary-interface "masterAnycast"
        address 10.20.0.1/32
      exit
      mh-secondary-interface "slaveAnycast"
        address 10.30.0.1/32
```

On PE-3:

```
*A:PE-3>config>router# info
-----
#-----
echo "IP Configuration"
#-----

      mh-primary-interface "masterAnycast"
        address 10.30.0.1/32
      exit
      mh-secondary-interface "slaveAnycast"
        address 10.20.0.1/32
      exit
```

Both PE-2 and PE-3 will have 10.20.0.1 as one of their anycast addresses, but only PE-2 has it configured as the master anycast address. Conversely both PE-2 and PE-3 have 10.30.0.1/32 as

Anycast Interface

one of their anycast addresses, but only PE3 has it configured as the master anycast address. Both interfaces will be considered as a kind of virtual interface, with a virtual port-id;

```
*A:PE-3>config>router# show router interface "masterAnycast" detail | match Port
Port Id          : vport-7
*A:PE-3>config>router#
```

Both anycast addresses, master and slave will be visible in the GRT (global routing table) of the PE (ABR);

```
*A:PE-2>config>router# show router route-table
=====
Route Table (Router: Base)
=====
Dest Prefix          Type      Proto    Age      Pref
  Next Hop[Interface Name]          Metric
-----
10.20.0.1/32        Local    Local    01d14h16m  0
    masterAnycast                0
10.30.0.1/32        Local    Local    01d14h16m  0
    slaveAnycast                  0
```

As similar configuration is performed on PE-4 and PE-5 with the anycast addresses being (Figure 134):

```
PE-4: masterAnycast = 10.40.0.1/32
      slaveAnycast  = 10.50.0.1/32
PE-5: masterAnycast = 10.50.0.1/32
      slaveAnycast  = 10.40.0.1/32
```

Therefore, if PE-4 fails, for example, PE-5 will take over the responsibility for PE-4's master anycast address (10.40.0.1); this will cause PE-2 and PE-3 to continue to send traffic which is directed towards the BGP routes with a next hop of 10.40.0.1 (PE-4's masterAnycast address) thinking its next hop is PE-4 when in reality it is PE-5. The failover initially only involves an IGP (ISIS) re-convergence, and hence an LDP re-convergence, instead of a BGP re-convergence which would take a longer time and therefore cause a larger outage (the IGP re-convergence may be reduced if ECMP or LDP FRR is used and an alternate label is already installed in the LFIB of PE-2/PE-3). BGP will subsequently re-converge such that the BGP next hop on PE-2 and PE-3 becomes 10.50.0.1 when they receive the withdraw of PE-4's 10.40.0.1/32 from the RR.

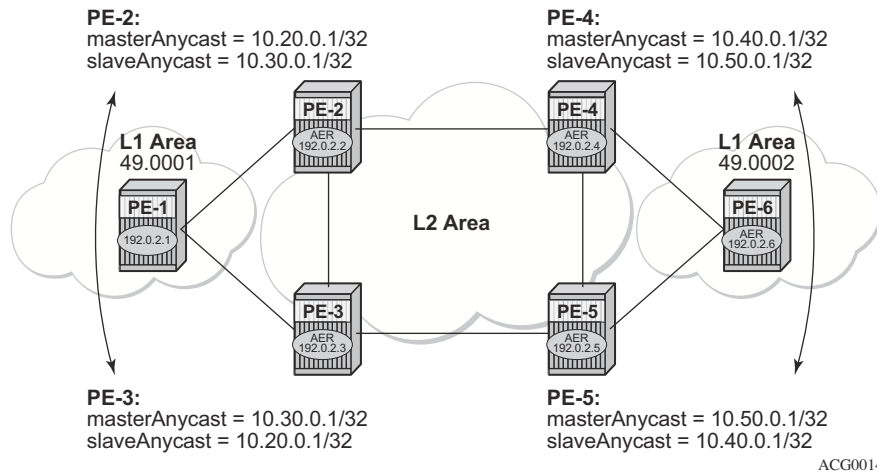


Figure 134: Anycast Address Configuration

IGP

Since both anycast addresses (master and slave) need to be reachable from the remote ABR and all remote PEs, they have to be advertised into the IGP (ISIS in this case). This will result in following ISIS configuration;

```
*A:PE-2>config>router>isis# info
-----
    area-id 49.0001
    export "remoteABR"
    traffic-engineering
    interface "system"
        passive
    exit
    interface "masterAnycast"
        level-capability level-2
    exit
    interface "slaveAnycast"
        level-capability level-2
    exit
    interface "int-PE-2-PE-3"
        level-capability level-2
        interface-type point-to-point
    exit
    interface "int-PE-2-PE-4"
        level-capability level-2
        interface-type point-to-point
    exit
    interface "int-PE-2-PE-1"
        level-capability level-1
        interface-type point-to-point
    exit
-----
*A:PE-2>config>router>isis#
```

Here, notice that the master and slave interface will be part of the Layer 2 area. Both interfaces will automatically be assigned with the correct ISIS metric, being 10 for the master and 30 for the slave (default values).

```
*A:PE-2>config>router>isis# show router isis interface
=====
ISIS Interfaces
=====
```

Interface	Level	CircID	Oper State	L1/L2 Metric
system	L1L2	1	Up	0/0
masterAnycast	L2	2	Up	-/10
slaveAnycast	L2	3	Up	-/30

The IGP (ISIS) metric can be configured to a different value if required.

The ISIS export policy is used to control redistribution of prefixes and essentially serves two purposes.

1. Ensure that anycast addresses and system IP addresses of other ABRs learned in Layer 2 LSPs through the core (from the perspective of PE-2, this would include PE-4 and PE-5) are re-advertised to PEs in the same area as Layer 1 LSPs.
2. Ensure that system IP addresses of PEs in the same region as the ABR learned through L1 LSPs are not advertised into the core as Layer 2 LSPs. The reason for this is that it is mandatory that a remote ABR has a best route to a remote PE (in a different area) through BGP and not by an IGP route in the FIB, otherwise it would not advertise the labeled BGP route toward PEs in its own region.

```
*A:PE-2>config>router>policy-options# info
```

```
-----
prefix-list "PE"
  prefix 192.0.2.1/32 exact
exit
prefix-list "remoteABR"
  prefix 10.40.0.1/32 exact
  prefix 10.50.0.1/32 exact
  prefix 192.0.2.4/32 exact
  prefix 192.0.2.5/32 exact
exit
prefix-list "masterAnycast"
  prefix 10.20.0.1/32 exact
exit

policy-statement "remoteABR"
  entry 10
    from
      prefix-list "remoteABR"
    exit
    action accept
  exit
  entry 20
    description "reject system-ip of own region"
    from
      prefix-list "PE"
    exit
    action reject
  exit
exit
```

LDP

All physical interfaces need to be part of ILDP so that the BGP label can be encapsulated in an MPLS packet with an LDP label. The configuration on PE-2 looks like the following.

```
*A:PE-2>config>router>ldp# info
-----
      export "exp-anycast"
      interface-parameters
        interface "int-PE-2-PE-3"
        exit
        interface "int-PE-2-PE-4"
        exit
        interface "int-PE-2-PE-1"
        exit
      exit
      targeted-session
      exit
-----
* A:PE-2>config>router>ldp# show router policy "exp-anycast"
  entry 10
    description "advertise master anycast address"
    from
      prefix-list "masterAnycast"
    exit
    to
      protocol ldp
    exit
    action accept
    exit
  exit
  entry 20
    description "advertise slave anycast address as well"
    from
      prefix-list "slaveAnycast"
    exit
    to
      protocol ldp
    exit
    action accept
    exit
  exit
A:PE-2>config>router>ldp#
```

By default, only the system address will be advertised by ILDP; hence the need for the export policy. This policy will advertise both the master and slave anycast address. The slave address advertisement is needed in both ABRs so that PE-2 can take over the role from the master (PE-3) in case of link or node failures (and vice versa).

To verify the LDP bindings, check if PE-2 can reach the remote anycast addresses through LDP, but not its own master and slave anycast address.

```

A:PE-2>config>router>ldp# show router ldp bindings
=====
LDP LSR ID: 192.0.2.2
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        S - Status Signaled Up, D - Status Signaled Down
        E - Epipe Service, V - VPLS Service, M - Mirror Service
        A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
        P - Ipipe Service, WP - Label Withdraw Pending, C - Cpipe Service
        TLV - (Type, Length: Value)
=====
LDP Prefix Bindings
=====
Prefix          Peer           IngLbl      EgrLbl  EgrIntf/  EgrNextHop
                Peer           Ingress     Egress   LspId
                -----
10.20.0.1/32    192.0.2.1      262140U     --      --        --
10.20.0.1/32    192.0.2.4      262140U     --      --        --
10.30.0.1/32    192.0.2.1      262139U     --      --        --
10.30.0.1/32    192.0.2.4      262139U     --      --        --
10.40.0.1/32    192.0.2.1      262143U     262143  --        --
10.40.0.1/32    192.0.2.3      262143U     262135  --        --
10.40.0.1/32    192.0.2.4      --          262139  1/1/3:0  192.168.4.2
10.50.0.1/32    192.0.2.1      262130U     262138  --        --
10.50.0.1/32    192.0.2.3      262130N     262142  1/1/2:0  192.168.13.2
10.50.0.1/32    192.0.2.4      --          262143  --        --
192.0.2.1/32    192.0.2.1      --          262142  1/1/1:0  192.168.2.1
192.0.2.1/32    192.0.2.3      262134U     262132  --        --

```

PE-2 can reach the remote anycast addresses 10.40.0.1 and 10.50.0.1. In this case both addresses are reachable by a different interface. Note that it can be via the same interface as well, due to the ECMP nature in the (Layer 2 area) square topology, where 10.50.0.1 can be reached via PE-4 as well as PE-2. There is no outgoing label for 10.20.0.1 and 10.30.0.1 which is to be expected since they are local addresses on PE-2.

To verify the actual datapath, following OAM command can be used.

```

*A:PE-2# oam lsp-trace prefix 10.40.0.1/32
lsp-trace to 10.40.0.1/32: 0 hops min, 0 hops max, 104 byte packets
1 192.0.2.4 rtt=2.02ms rc=3(EgressRtr)
*A:PE-2# oam lsp-trace prefix 10.50.0.1/32
lsp-trace to 10.50.0.1/32: 0 hops min, 0 hops max, 104 byte packets
1 192.0.2.3 rtt=2.76ms rc=8(DSRtrMatchLabel)
2 192.0.2.5 rtt=3.57ms rc=3(EgressRtr)
*A:PE-2#

```

BGP Configuration

The configuration of BGP will be split up in following parts;

1. BGP configuration on the ABR
2. BGP configuration on the RR
3. BGP configuration on the PE

BGP on the ABR

The following BGP configuration is present on PE-2;

```
*A:PE-2# configure router bgp
*A:PE-2>config>router>bgp# info
-----
      min-route-advertisement 10
      group "region"
        description "all PE of the region will peer with this ABR"
        family ipv4 vpn-ipv4
        type internal
        cluster 2.2.2.2
        neighbor 192.0.2.1
          advertise-label ipv4
        exit
      exit
      group "fullmesh"
        description "PE-4 is RR"
        family ipv4 vpn-ipv4
        type internal
        export "exp-anycast-nhs"
        neighbor 192.0.2.4
          advertise-label ipv4
        exit
      exit
-----
*A:PE-2>config>router>bgp#
```

The two groups identify the PE in the local region (part of L1 area) and the RR in the core (part of Layer 2 area).

The export policy will set the next-hop of the IPv4 and IPv4-VPN routes equal to the master anycast address;

```
A:PE-2>config>router>bgp# show router policy "exp-anycast-nhs"
  entry 10
    description "set NH of PE to master anycast"
    from
      prefix-list "PE"
      family ipv4
    exit
```



```
to
  protocol bgp
exit
action accept
  local-preference 150 (*)
  next-hop 10.20.0.1
exit
exit
entry 20
  description "set NH of IP-VPN routes to master anycast"
  from
    family vpn-ipv4
  exit
  to
    protocol bgp-vpn
  exit
  action accept
    local-preference 150 (*)
    next-hop 10.20.0.1
  exit
exit
A:PE-2>config>router>bgp#
```

(*) optional, not mandatory for functional testing.

For testing purpose, this policy is sufficient. In reality more restrictions will be needed to avoid obsolete advertisements of routes by the ABR, but this is out-of-scope for this document.

BGP on the RR

BGP on RR (PE-4) will look like the following:

```
A:PE-4>config>router>bgp# info
-----
vpn-apply-import
vpn-apply-export
min-route-advertisement 1
group "region"
  description "all PE of the region will peer with this ABR"
  family ipv4 vpn-ipv4
  type internal
  cluster 4.4.4.4
  neighbor 192.0.2.6
    advertise-label ipv4
  exit
exit
group "fullmesh"
  description "RR in cluster 1.1.1.1"
  family ipv4 vpn-ipv4
  type internal
  cluster 1.1.1.1
  export "exp-anycast-nhs"
  neighbor 192.0.2.2
    advertise-label ipv4
  exit
  neighbor 192.0.2.3
    advertise-label ipv4
  exit
  neighbor 192.0.2.5
    advertise-label ipv4
  exit
exit
-----

A:PE-4>config>router>bgp# show router policy "exp-anycast-nhs"
entry 10
  description "set NH of PE to master anycast"
  from
    prefix-list "PE"
    family ipv4
  exit
  to
    protocol bgp
  exit
  action accept
    local-preference 150
    next-hop 10.40.0.1
  exit
exit

A:PE-4# show router policy prefix-list "PE"
prefix 192.0.2.6/32 exact
prefix 192.168.0.6/32 exact
A:PE-4#
```

Where the anycast address is set as the BGP NH in all IPv4 prefixes for all PEs in the attached area/region. The BGP NH is not set for IP-VPN routes, as this would override the NH in advertised IP-VPN routes of PE-2 and PE-3 for which PE-4 acts as RRs and should therefore not modify the NH of the IP-VPN routes.

BGP on the PE

Each PE will have a peering towards each ABR of its own region. In this case PE-1 will peer with PE-2 and PE-3, which in turn act as RRs so there are no inter-PE peering sessions required within the region.

```
*A:PE-1# configure router bgp
*A:PE-1>config>router>bgp# info
-----
      min-route-advertisement 10
      group "abr"
        description "peering to ABR, which act as RR"
        family ipv4 vpn-ipv4
        type internal
        export "expbgp"
        neighbor 192.0.2.2
          advertise-label ipv4
        exit
        neighbor 192.0.2.3
          advertise-label ipv4
        exit
      exit
-----
A:PE-1>config>router>bgp#
```

The referenced export policy triggers the advertisement of the PE-1 system IP address with a BGP-label.

```
A:PE-1# show router policy "expbgp"
  entry 10
    from
      prefix-list "PE"
    exit
  to
    protocol bgp
  exit
  action accept
  exit
exit

A:PE-1# show router policy prefix-list "PE"
prefix 192.0.2.1/32 exact
A:PE-1#
```

Data Path Verification

The best way to verify the data path is between both (remote) PEs.

Looking at PE-1, verify the following:

1. Was a BGP route received toward PE-6 with an NH equal to an anycast address?

```
A:PE-1# show router bgp routes 192.0.2.6/32
=====
BGP Router ID:192.0.2.1          AS:65536          Local AS:65536
=====
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  192.0.2.6/32              150        10
      10.40.0.1                -
      No As-Path
*i    192.0.2.6/32              150        10
      10.50.0.1                -
      No As-Path
-----
Routes : 2
=====
A:PE-1#
```

Note that both routes are in fact equal from a BGP point of view. If ECMP and (BGP) multipath would have been set to 2, both routes would be active in the RTM.

2. Can the anycast addresses be reached through LDP?

```
*A:PE-1# show router ldp bindings prefix 10.40.0.1/32
=====
LDP LSR ID: 192.0.2.1
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
      WP - Label Withdraw Pending
=====
LDP Prefix Bindings
=====
Prefix                Peer                IngLbl  EgrLbl  EgrIntf  EgrNextHop
-----
10.40.0.1/32          192.0.2.2           262143N 262143 1/1/1:0  192.168.2.2
10.40.0.1/32          192.0.2.3           262143U 262135  --      --
-----
No. of Prefix Bindings: 2
=====
*A:PE-1# oam lsp-trace prefix 10.40.0.1/32
lsp-trace to 10.40.0.1/32: 0 hops min, 0 hops max, 104 byte packets
```

```

1 192.0.2.2 rtt=2.18ms rc=8(DSRtrMatchLabel)
2 192.0.2.4 rtt=2.88ms rc=3(EgressRtr)
*A:PE-1#

```

There are two labels learned and one is active since the shortest path to PE-4 is via PE-2, as seen in the OAM command.

3. Verify that the ABR has learned the anycast advertisements of the neighboring master. This is needed to evaluate the incoming labelled packets correctly after failures in the Layer 2 area (link or node).

On PE-2, observe the presence of the anycast label in the context-specific label space, with a link to VPRN-id 1. This is described in [IP-VPN Routes on page 964](#).

```

A:PE-2# show router bgp anycast-label
=====
BGP Anycast-MH labels
=====
Secondary-MH-Addr          ABR-Lbl   Cfg-Time  VPRN-ID
PE-Addr                    PE-Lbl    Rem-Time  Ref-Count
-----
10.30.0.1                  262130    30        1
-                          -         -         3
=====
A:PE-2#

```

On the redundant ABR PE-3, an additional label-mapping is present because the route 192.168.0.1/32 (part of GRT) is received on PE-3 with the anycast NH of PE-2. PE-2 does not receive this route from PE-3 because the RR (PE-4) has selected the route 192.168.0.1/32 from PE-2 as the best one, which means that PE-2 will not receive the route to 192.168.0.1/32 with PE-3 anycast address as NH.

```

A:PE-3# show router bgp anycast-label
=====
BGP Anycast-MH labels
=====
Secondary-MH-Addr          ABR-Lbl   Cfg-Time  VPRN-ID
PE-Addr                    PE-Lbl    Rem-Time  Ref-Count
-----
10.20.0.1                  262131    30        -
192.0.2.1                  262134    -         1

10.20.0.1                  262133    30        1
-                          -         -         3
=====
A:PE-3#

```

This command can be used to check the mapping of BGP labels in the context-specific label space.

More details are provided [Deployment Options on page 959](#).

Data Path Verification

4. Verify data path towards PE-6, where some basic CLI commands can be used;

```
*A:PE-1# show router route-table 192.0.2.6/32
=====
Route Table (Router: Base)
=====
Dest Prefix                               Type   Proto   Age           Pref
  Next Hop[Interface Name]                Metric
-----
192.0.2.6/32                               Remote BGP     00h30m50s    170
  10.40.0.1 (tunneled)                      0
-----
No. of Routes: 1
=====
*A:PE-1#
*A:PE-1# show router fib 1 192.0.2.6/32
=====
FIB Display
=====
Prefix                                     Protocol
  NextHop
-----
192.0.2.6/32                               BGP
  10.40.0.1 (Transport:LDP)
-----
Total Entries : 1
=====
...
*A:PE-1# show router tunnel-table
=====
Tunnel Table (Router: Base)
=====
Destination      Owner Encap TunnelId  Pref   Nexthop      Metric
-----
192.0.2.6/32     bgp   MPLS   -         10     10.40.0.1    1000
-----
*A:PE-1# traceroute 192.0.2.6 no-dns
traceroute to 192.0.2.6, 30 hops max, 40 byte packets
 1 192.0.2.6    28.9 ms  4.53 ms  5.00 ms
*A:PE-1#
```

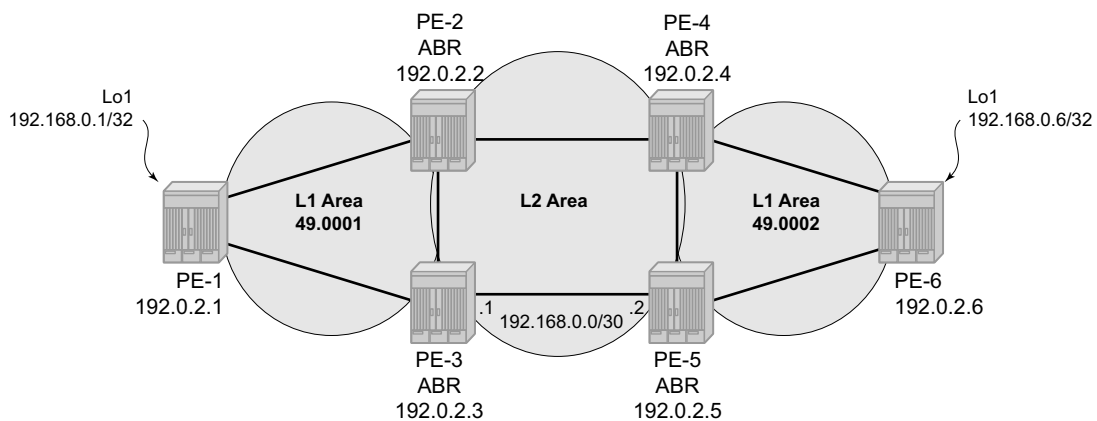
Deployment Options

There are two main deployment options with BGP anycast.

1. Providing an end-to-end (E2E) MPLS tunnel, between access nodes (ANs) connected to the PE. In this case, VLLs can be supported between AN over a BGP signalled SDP. This solves scaling issues and also provides resiliency.
2. Use the anycast address as the NH for IP-VPN routes, which improves NH tracking in large MPLS domains, where IGP convergence can not contribute to the NH convergence.

E2E MPLS Between Access Nodes

The AN will be simulated by loopback interfaces on PE-1 and PE-6. From a functional point of view this is equivalent to a connected AN that is not part of the IGP domain.



OSSG615

Figure 135: E2E MPLS Between Access Nodes

The loopback interface Lo1 is not part of the IGP (ISIS) but will be advertised by BGP towards the ABR with a BGP label. The configuration looks like the following.

```
*A:PE-1>config>router# info
-----
#-----
echo "IP Configuration"
#-----

interface "lo1"
  address 192.168.0.1/32
  description "simulate an AN (access node) "
```

Deployment Options

```
        loopback
    exit
    interface "system"
        address 192.0.2.1/32
    exit
    autonomous-system 64496
```

BGP Configuration

The BGP configuration on the regional PE will look like the following:

```
A:PE-1# configure router bgp
A:PE-1>config>router>bgp# info
-----
    vpn-apply-import
    vpn-apply-export
    min-route-advertisement 1
    group "abr"
        description "peering to ABR, which act as RR"
        family ipv4 vpn-ipv4
        type internal
        export "expbgp"
        neighbor 192.0.2.2
            advertise-label ipv4
        exit
        neighbor 192.0.2.3
            advertise-label ipv4
        exit
    exit
-----

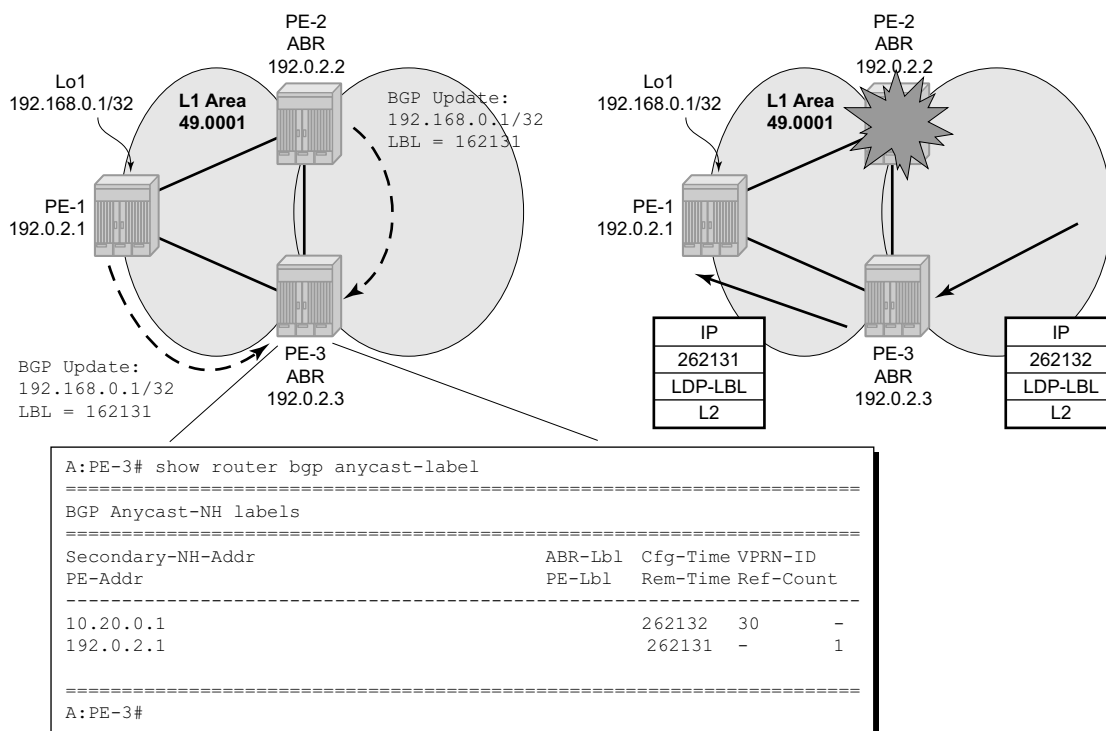
A:PE-1>config>router>bgp# show router policy "expbgp"
    entry 10
        description "advertise the PE system-ip through BGP"
        from
            prefix-list "PE"
        exit
        to
            protocol bgp
        exit
        action accept
            local-preference 200
        exit
    exit
    entry 20
        description "advertise connect AN through BGP"
        from
            prefix-list "AN"
        exit
        to
            protocol bgp
        exit
        action accept
        exit
    exit
```



```
A:PE-1>config>router>bgp# show router policy prefix-list "AN"
prefix 192.168.0.1/32 exact
A:PE-1>config>router>bgp# show router policy prefix-list "PE"
prefix 192.0.2.1/32 exact
A:PE-1>config>router>bgp#
```

The prefix list should include all the addresses of connected ANs that need to be reachable through labelled BGP.

Note that the ABR will also learn the BGP route advertisement of its neighboring ABR and installs the BGP label accordingly, as shown in [Figure 136](#).



OSSG611

Figure 136: Data Path with Failing ABR

In case PE-2 fails or becomes unreachable, PE-3 can still recognize the BGP label of the packet and perform the correct BGP (and LDP) label swap so that the packet will not be dropped and can continue its way towards the destination.

Note that for every AN a new entry in the BGP anycast-label table will be added. Refer to the following output.

Deployment Options

```

BGP anycast for A:PE-3# show router bgp anycast-label
=====
BGP Anycast-MH labels
=====
Secondary-MH-Addr          ABR-Lbl    Cfg-Time  VPRN-ID
PE-Addr                   PE-Lbl     Rem-Time  Ref-Count
-----
10.20.0.1                 262131     30        -
192.0.2.1                 262134     -         1

10.20.0.1                 262132     30        -
192.0.2.1                 262131     -         1

```

Important note: If an E2E MPLS transport tunnel is needed between PE-1 and PE-6, additional loopback interfaces need to be created on both PE since the system address cannot be used with the BGP anycast feature. The reason for this is that on the ABR, the best route towards the PE system-ip is an IGP route, not a (labeled) BGP route. BGP anycast can only map a BGP label against another BGP label, not towards an IGP (hence LDP) route¹. This is important to keep in mind, if BGP anycast is deployed in networks where BGP anycast is required for E2E MPLS tunnels between AN and remote PEs.

Figure 137 illustrates the SDP between PE-1 and PE-6.




Figure 137: End-to-End Transport Tunnel Using Additional Loopback Interfaces

1. Although PE-1 will advertise the system IP as a BGP-labeled route, it cannot be used for the BGP anycast feature since the route to system IP of PE-1 in the GRT of PE-2 and PE-3 will always be an IGP route, which makes it impossible to use the BGP anycast feature since this requires an E2E BGP tunnel.

SDP Configuration

The SDP configuration consists of two parts, the service component and LDP component. In the service context, the SDP will be created with BGP as the tunneling protocol.

```
*A:PE-1# configure service sdp 100
*A:PE-1>config>service>sdp# info
```

```
-----
      far-end 192.168.0.6
      bgp-tunnel
      keep-alive
      shutdown
      exit
      no shutdown
-----
```

```
*A:PE-1>config>service>sdp#
```

Since the T-LDP session is set up between two loopback interfaces instead of the system IP of PE-1 and PE-6, a dedicated configuration is required at the LDP level.

```
*A:PE-1# show router interface "lo1"
```

```
-----
Interface Table (Router: Base)
-----
```

Interface-Name IP-Address	Adm	Opr (v4/v6)	Mode	Port/SapId PfxState
lo1 192.168.0.1/32	Up	Up/Down	Network	loopback n/a

```
-----
Interfaces : 1
-----
```

```
*A:PE-1#
*A:PE-1# configure router ldp
*A:PE-1>config>router>ldp# info
```

```
-----
      interface-parameters
      interface "int-PE-1-PE-2"
      exit
      interface "int-PE-1-PE-3"
      exit
      exit
      targeted-session
      peer 192.168.0.6
      local-lsr-id "lo1"
      exit
      exit
-----
```

```
*A:PE-1>config>router>ldp#
```

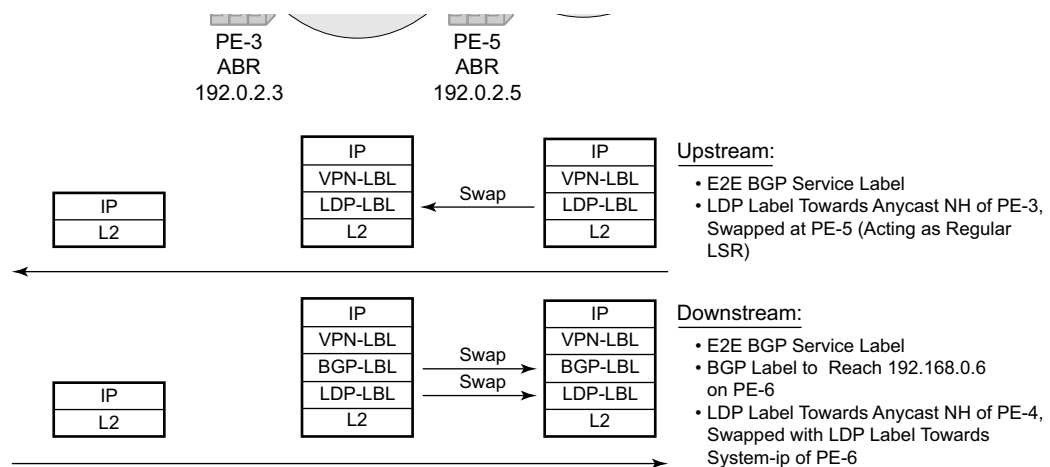
Note that where the destination address of the T-LDP session is set to the remote loopback address and the local address is set to the local loopback interface.

IP-VPN Routes

Important note: In the SR-OS 9.0R1 software release, BGP anycast for IP-VPN routes is only supported with a RR cluster in the core. A full mesh of iBGP peers or confederations are not supported in SROS 9.0R1.

The BGP anycast mechanism can also be used to advertise IP-VPN routes. The advertising PE sets the NH equal to the anycast master address instead of the regular system IP. The configuration is based upon PE-1 which will now act as a customer edge (CE) device. This can be achieved by creating a local IP-VPN on PE-1, without any MPLS connectivity, due to the use of a unique RT (route-target) and hybrid ports towards the ABR². As such, both PE-2 and PE-3 can advertise the same route with the AC address as a NH.

Figure 138 illustrates the logical IP-VPN topology.



OSSG614

Figure 138: IP-VPN with Anycast NH

- The unique RT and hybrid ports are required for this specific set-up where PE-1 is combined as a regular PE and as CE. The CE functionality is achieved by creating a local VPRN, and the unique RT will assure that the route is not imported at PE-2 and PE-3 by MP-BGP. The hybrid port is needed to combine a PE-CE interface and network interface (MPLS capable) on the same physical port.

Configure IP-VPN at the PE (PE-2 and PE-3)

The following output displays the VPRN configuration. The change of the next-hop to the BGP anycast address is done at the BGP group level in this case.

```
A:PE-2>config>service>vprn# info
-----
description "IP-VPN PE, with anycast NH set"
vrf-import "vrfImpPolBgpVpnRts"
vrf-export "adv_vpn"
router-id 192.0.2.2
autonomous-system 64498
route-distinguisher 1:2
auto-bind ldp
vrf-target target:1:1
interface "to_CE" create
  address 172.16.2.2/30
  sap 1/1/1:1 create
  exit
exit
bgp
  group "CE"
    type external
    export "exp_bgp_vpn_rts_to_ce"
    peer-as 64499
    neighbor 172.16.2.1
    exit
  exit
exit
no shutdown
-----

A:PE-2>config>service>vprn#

A:PE-2>config>service>vprn# show router policy "vrfImpPolBgpVpnRts"
description "Policy From bgpVpn To none"
entry 10
  description "Entry 10 - From Prot. bgpVpn To none"
  from
    protocol bgp-vpn
    community "vprn1"
  exit
  to
  exit
  action accept
  exit
exit

A:PE-2>config>service>vprn# show router policy "adv_vpn"
entry 10
  description "local routes"
  from
    protocol direct
  exit
  to
    protocol bgp-vpn
  exit
  action accept
```

Deployment Options

```
        community add "vprn1"
    exit
exit
entry 20
    description "PE-CE BGP routes"
    from
        protocol bgp
    exit
    to
        protocol bgp-vpn
    exit
    action accept
        community add "vprn1"
    exit
exit

A:PE-2# show router policy community "vprn1"
community "vprn1" members "target:1:1"
A:PE-2#
A:PE-2# show router policy "exp_bgp_vpn_rts_to_ce"
    description "Policy From bgpVpn To none"
    entry 10
        description "Entry 10 - From Prot. bgpVpn To none"
        from
            protocol bgp-vpn
        exit
        to
        exit
        action accept
        exit
    exit
A:PE-2#
```

Note that the route distinguisher of VPRN 1 at PE-3 is set to 1:3 instead of 1:2 as on PE-2.

As stated above, the NH will be set to the anycast address in the BGP group at the global level.

```
*A:PE-2# configure router bgp
*A:PE-2>config>router>bgp# info
-----
vpn-apply-export
min-route-advertisement 1
group "region"
    description "all PE of the region will peer with this ABR"
    family ipv4 vpn-ipv4
    type internal
    cluster 1.1.1.1
    neighbor 192.0.2.1
        advertise-label ipv4
    exit
exit
group "fullmesh"
    description "PE-4 is RR"
    family ipv4 vpn-ipv4
    type internal
    export "exp-anycast-nhs"
    neighbor 192.0.2.4
        advertise-label ipv4
```

```
        exit
    exit
-----
*A:PE-2>config>router>bgp# show router policy "exp-anycast-nhs"
  entry 10
    description "set NH of PE to master anycast"
    from
      prefix-list "PE"
      family ipv4
    exit
    to
      protocol bgp
    exit
    action accept
      local-preference 150 (*)
      next-hop 10.20.0.1
    exit
  exit
  entry 20
    description "set NH of IP-VPN routes to master anycast"
    from
      family vpn-ipv4
    exit
    to
      protocol bgp-vpn
    exit
    action accept
      local-preference 150 (*)
      next-hop 10.20.0.1
    exit
  exit
*A:PE-2>config>router>bgp#
(*) optional
```

Configure IP-VPN at the Remote PE (PE-6)

On PE-6, IP-VPN routes need to be advertised with the address of lo1 as NH, instead of the regular system-ip. This is required to use a BGP tunnel (and on top of that BGP anycast) to reach PE-6 from PE-2 and PE-3.

```
A:PE-6>config>service>vprn# info
-----
vrf-import "vrfImpPolBgpVpnRts"
router-id 172.16.6.6
route-distinguisher 1:6
auto-bind mpls
vrf-target target:1:1
interface "lo1" create
    address 172.16.6.6/32
    loopback
exit
no shutdown
-----
A:PE-6>config>service>vprn#
```

Note that the NH of the IP-VPN routes will be set to 192.168.0.6 instead of the system-ip;

```
A:PE-6# configure router bgp
A:PE-6>config>router>bgp# info
-----
vpn-apply-import
vpn-apply-export
min-route-advertisement 1
outbound-route-filtering
exit
group "abr"
    description "peering to ABR, which act as RR"
    family ipv4 vpn-ipv4
    type internal
    export "expbgp"
    neighbor 192.0.2.4
        advertise-label ipv4
    exit
    neighbor 192.0.2.5
        advertise-label ipv4
    exit
exit
-----
A:PE-6>config>router>bgp# show router policy
policy          policy-edits
A:PE-6>config>router>bgp# show router policy "expbgp"
entry 10
description "advertise the PE system-ip through BGP"
from
    prefix-list "PE"
exit
to
    protocol bgp
exit
action accept
```



```

        local-preference 200
    exit
exit
entry 20
    description "advertise connect AN through BGP"
    from
        prefix-list "AN"
    exit
    to
        protocol bgp
    exit
    action accept
    exit
exit
entry 30
    description "set NH for vpn-routes to 192.168.0.6"
    from
        community "vprn1"
    exit
    action accept
        local-preference 333
        next-hop 192.168.0.6
    exit
exit
A:PE-6>config>router>bgp# exit all
A:PE-6# show router policy prefix-list "AN"
prefix 192.168.0.6/32 exact
A:PE-6# show router policy prefix-list "PE"
prefix 192.0.2.6/32 exact
A:PE-6#

```

The local preference is set to 333 for troubleshooting purposes, primarily to find the routes again more easily.

The NH 192.168.0.6 is a local loopback, and reachable from PE-2 and PE-3 by a BGP tunnel with the anycast address of PE-4 as NH.

On PE-6:

```

A:PE-6# show router interface
=====
Interface Table (Router: Base)
=====
Interface-Name      Adm      Opr (v4/v6)  Mode      Port/SapId
IP-Address          PfxState
-----
int-PE-6-PE-4      Up       Up/Down     Network  1/1/1:0
    192.168.6.2/30
int-PE-6-PE-5      Up       Up/Down     Network  1/1/3:0
    192.168.25.2/30
lol                 Up       Up/Down     Network  loopback
    192.168.0.6/32
system              Up       Up/Down     Network  system
    192.0.2.6/32
-----
Interfaces : 4
=====

```

Deployment Options

A:PE-6#

On PE-2 (and PE-3):

A:PE-2# show router tunnel-table

=====
Tunnel Table (Router: Base)
=====

Destination	Owner	Encap	TunnelId	Pref	Nexthop	Metric
10.40.0.1/32	ldp	MPLS	-	9	192.168.4.2	20
10.50.0.1/32	ldp	MPLS	-	9	192.168.13.2	30
192.0.2.1/32	ldp	MPLS	-	9	192.168.2.1	10
192.0.2.3/32	ldp	MPLS	-	9	192.168.2.1	20
192.0.2.4/32	ldp	MPLS	-	9	192.168.4.2	10
192.0.2.5/32	ldp	MPLS	-	9	192.168.13.2	20
192.0.2.6/32	bgp	MPLS	-	10	10.40.0.1	1000
192.168.0.6/32	bgp	MPLS	-	10	10.40.0.1	1000

=====
A:PE-2#

On PE-4 (and PE-5):

The BGP anycast labels are active.

A:PE-4# show router bgp anycast-label

=====
BGP Anycast-MH labels
=====

Secondary-MH-Addr	ABR-Lbl	Cfg-Time	VPRN-ID
PE-Addr	PE-Lbl	Rem-Time	Ref-Count
10.50.0.1	262132	30	-
192.0.2.6	262132	-	1

=====
A:PE-4#

Configure IP-VPN at the CE

On PE-1, VPRN 1 is a local IP-VPN acting as CE. The IP-VPN is dual homed towards PE-2 and PE-3, including EBGP sessions that will advertise the directly connected links toward the PE.

```
A:PE-1>config>service>vprn# info
-----
description "local VPN, simulating CE"
router-id 172.168.1.1
autonomous-system 64499
route-distinguisher 1:1
interface "to_ABR_PE3" create
  address 172.16.3.1/30
  sap 1/1/3:1 create
  exit
exit
interface "to_ABR_PE2" create
  address 172.16.2.1/30
  sap 1/1/1:1 create
  exit
exit
interface "lo1" create
  address 172.168.1.1/32
  loopback
exit
bgp
  min-route-advertisement 1
  export "adv_direct"
  group "ABR"
    type external
    peer-as 64498
    neighbor 172.16.2.2
    exit
    neighbor 172.16.3.2
    exit
  exit
exit
no shutdown
-----
A:PE-1>config>service>vprn#

A:PE-1>config>service>vprn# show router policy "adv_direct"
  entry 10
    description "advertise local interfaces to PE"
    from
      protocol direct
    exit
    to
      protocol bgp
    exit
    action accept
    exit
  exit
A:PE-1>config>service>vprn#
```

Deployment Options

To verify the BGP peering sessions;

```
A:PE-1# show router 1 bgp summary
```

```
=====
BGP Router ID:192.0.2.1          AS:64499          Local AS:64499
=====
BGP Admin State      : Up          BGP Oper State      : Up
Total Peer Groups    : 1           Total Peers          : 2
Total BGP Paths       : 3           Total Path Memory    : 384
Total IPv4 Remote Rts : 4           Total IPv4 Rem. Active Rts : 0
Total IPv6 Remote Rts : 0           Total IPv6 Rem. Active Rts : 0
Total Supressed Rts  : 0           Total Hist. Rts      : 0
Total Decay Rts      : 0
=====
```

```
BGP Summary
=====
```

```
Neighbor
```

```
AS PktRcvd InQ Up/Down State|Rcv/Act/Sent (Addr Family)
PktSent OutQ
-----
172.16.2.2
        64498 28695 0 05d23h53m 2/0/3 (IPv4)
        28671 0
172.16.3.2
        64498 28618 0 06d00h03m 2/0/3 (IPv4)
        28611 0
-----
```

```
A:PE-1#
```

The received routes are not inserted in the FIB since they are already known as local routes.

Verify Anycast Labels

On PE-2, first check the BGP service label that has been advertised by PE-3 as follows.

```
A:PE-2# show router bgp routes vpn-ipv4 1:3:172.168.1.1/32
=====
BGP Router ID:192.0.2.2      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
=====
Network       : 172.168.1.1/32
Nexthop       : 10.30.0.1
Route Dist.   : 1:3                VPN Label     : 262130
From          : 192.0.2.4
Res. Nexthop  : n/a
Local Pref.   : 150                Interface Name : slaveAnycast
Aggregator AS : None                Aggregator    : None
Atomic Aggr.  : Not Atomic         MED           : None
Community     : target:1:1
Cluster       : 1.1.1.1
Originator Id : 192.0.2.3            Peer Router Id : 192.0.2.4
Flags         : Invalid Incomplete (*)
AS-Path       : 64499
VPRN Imported : None
-----
Routes : 1
```

(*) This means that the NH is resolved to its own local interface. Nevertheless, the service label will be inserted in the context-specific label space so that incoming packets with this service label can be linked to VPRN 1.

Notice the loopback route of the CE with label 262130, which is the BGP label that represents VPRN 1 at PE-3.

```
A:PE-2# show router bgp anycast-label
=====
BGP Anycast-MH labels
=====
Secondary-MH-Addr      ABR-Lbl   Cfg-Time VPRN-ID
PE-Addr                PE-Lbl    Rem-Time Ref-Count
-----
10.30.0.1              262130    30       1
-                      -         -        3
=====
A:PE-2#
```

PE-3 will forward the traffic received with BGP service label 262130 towards VPRN 1 where an egress IP lookup will be performed. In the VRF (FIB of VPRN-id 1) at PE-2, a route toward 172.168.1.1/32 with VPRN 1 at PE-1 as an NH is found.

Deployment Options

```
A:PE-2# show router 1 route-table
=====
Route Table (Service: 1)
=====
Dest Prefix                               Type   Proto   Age           Pref
  Next Hop[Interface Name]                Metric
-----
172.16.2.0/30                             Local  Local   06d00h35m    0
      to_CE                                0
172.16.3.0/30                             Remote BGP     05h01m41s   170
      172.16.2.1                            0
172.168.1.1/32                           Remote BGP     06d00h35m   170
      172.16.2.1                            0
-----
No. of Routes: 3
=====
A:PE-2#
```

Verify Data Path Between VPRN on PE-6 and the CE

At PE-6, a ping and traceroute are performed towards the loopback interface of the CE (VPRN-id 1 at PE-1) with the following results.

```
A:PE-6# ping router 1 172.168.1.1 source 172.16.6.6
PING 172.168.1.1 56 data bytes
64 bytes from 172.168.1.1: icmp_seq=1 ttl=63 time=3.79ms.
^C
ping aborted by user

---- 172.168.1.1 PING Statistics ----
1 packet transmitted, 1 packet received, 0.00% packet loss
round-trip min = 3.79ms, avg = 3.79ms, max = 3.79ms, stddev = 0.000ms

A:PE-6# traceroute no-dns router 1 172.168.1.1 source 172.16.6.6
traceroute to 172.168.1.1 from 172.16.6.6, 30 hops max, 40 byte packets
 1  0.0.0.0  * * *
 2  172.168.1.1    4.62 ms  8.05 ms  4.28 ms
A:PE-6#
```

Upstream

From PE-6, the NH of the IP-VPN route will be the anycast address (in this case 10.30.0.1) which will be reachable through an LDP tunnel.

```
A:PE-6# show router 1 route-table
=====
Route Table (Service: 1)
=====
Dest Prefix                               Type   Proto   Age           Pref
  Next Hop[Interface Name]                Metric
-----
172.16.2.0/30                             Remote BGP VPN 16h21m20s    170
  10.20.0.1 (tunneled)                    0
172.16.3.0/30                             Remote BGP VPN 16h21m20s    170
  10.30.0.1 (tunneled)                    0
172.16.6.6/32                             Local  Local   10d21h52m    0
  lo1                                       0
172.168.1.1/32                            Remote BGP VPN 16h21m20s    170
  10.30.0.1 (tunneled)                    0
-----
No. of Routes: 4
=====
A:PE-6# show router 1 fib 1
=====
FIB Display
=====
Prefix                                     Protocol
  NextHop
-----
172.16.2.0/30                             BGP_VPN
  10.20.0.1 (VPRN Label:262133 Transport:LDP)
172.16.3.0/30                             BGP_VPN
  10.30.0.1 (VPRN Label:262130 Transport:LDP)
172.16.6.6/32                             LOCAL
```

Deployment Options

```
172.16.6.6 (lo1)
172.168.1.1/32
10.30.0.1 (VPRN Label:262130 Transport:LDP)
-----
Total Entries : 4
=====
A:PE-6#

A:PE-6# show router ldp bindings prefix 10.30.0.1/32
=====
LDP LSR ID: 192.0.2.6
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
      WP - Label Withdraw Pending
=====
LDP Prefix Bindings
=====
Prefix          Peer          IngLbl    EgrLbl  EgrIntf    EgrNextHop
-----
10.30.0.1/32    192.0.2.4      262135U   262139   --          --
10.30.0.1/32    192.0.2.5      262135N   262134   1/1/3:0     192.168.25.1
-----
No. of Prefix Bindings: 2
=====
A:PE-6#
```

Downstream

First, verify which PE is selected by the CE to reach the 172.16.6.6/32 (in VPN 1 at PE-6);

```
A:PE-1# show router 1 route-table
=====
Route Table (Service: 1)
=====
Dest Prefix          Type   Proto   Age          Pref
  Next Hop[Interface Name]          Metric
-----
172.16.2.0/30        Local  Local   10d16h54m   0
   to_ABR_PE2                0
172.16.3.0/30        Local  Local   10d16h53m   0
   to_ABR_PE3                0
172.16.6.6/32        Remote BGP     00h54m07s  170
   172.16.2.2                0
172.168.1.1/32       Local  Local   10d16h52m   0
   lo1                        0
-----
No. of Routes: 4
=====
A:PE-1#
```

In this case, PE-2 is selected as the NH.

At PE-2, the path towards PE-6 is based upon a BGP tunnel with the BGP anycast feature.


```

A:PE-2# show router 1 route-table
=====
Route Table (Service: 1)
=====
Dest Prefix                               Type   Proto   Age           Pref
  Next Hop[Interface Name]                Metric
-----
172.16.2.0/30                             Local  Local   06d16h45m    0
    to_CE                                  0
172.16.3.0/30                             Remote BGP     21h12m08s    170
    172.16.2.1                             0
172.16.6.6/32                             Remote BGP VPN 00h08m18s    170
    192.168.0.6 (tunneled)                 0
172.168.1.1/32                            Remote BGP     06d16h45m    170
    172.16.2.1                             0
-----
No. of Routes: 4
=====
A:PE-2#

A:PE-2# show router 1 fib 1
=====
FIB Display
=====
Prefix                                     Protocol
  NextHop
-----
172.16.2.0/30                             LOCAL
    172.16.2.0 (to_CE)
172.16.3.0/30                             BGP
    172.16.2.1 Indirect (to_CE)
172.16.6.6/32                             BGP_VPN
    192.168.0.6 (VPRN Label:262134 Transport:BGP)
172.168.1.1/32                            BGP
    172.16.2.1 Indirect (to_CE)
-----
Total Entries : 4
=====
A:PE-2#

A:PE-2# show router tunnel-table
=====
Tunnel Table (Router: Base)
=====
Destination      Owner Encap TunnelId Pref   Nexthop      Metric
-----
10.40.0.1/32     ldp   MPLS   -      9      192.168.4.2  20
10.50.0.1/32     ldp   MPLS   -      9      192.168.13.2 30
192.0.2.1/32     ldp   MPLS   -      9      192.168.2.1  10
192.0.2.3/32     ldp   MPLS   -      9      192.168.2.1  20
192.0.2.4/32     ldp   MPLS   -      9      192.168.4.2  10
192.0.2.5/32     ldp   MPLS   -      9      192.168.13.2 20
192.0.2.6/32     bgp   MPLS   -     10     10.40.0.1    1000
192.168.0.6/32   bgp   MPLS   -     10     10.40.0.1    1000
=====
A:PE-2#

```

The end-to-end label stack for the ping will look like the following (in this particular case);

Deployment Options

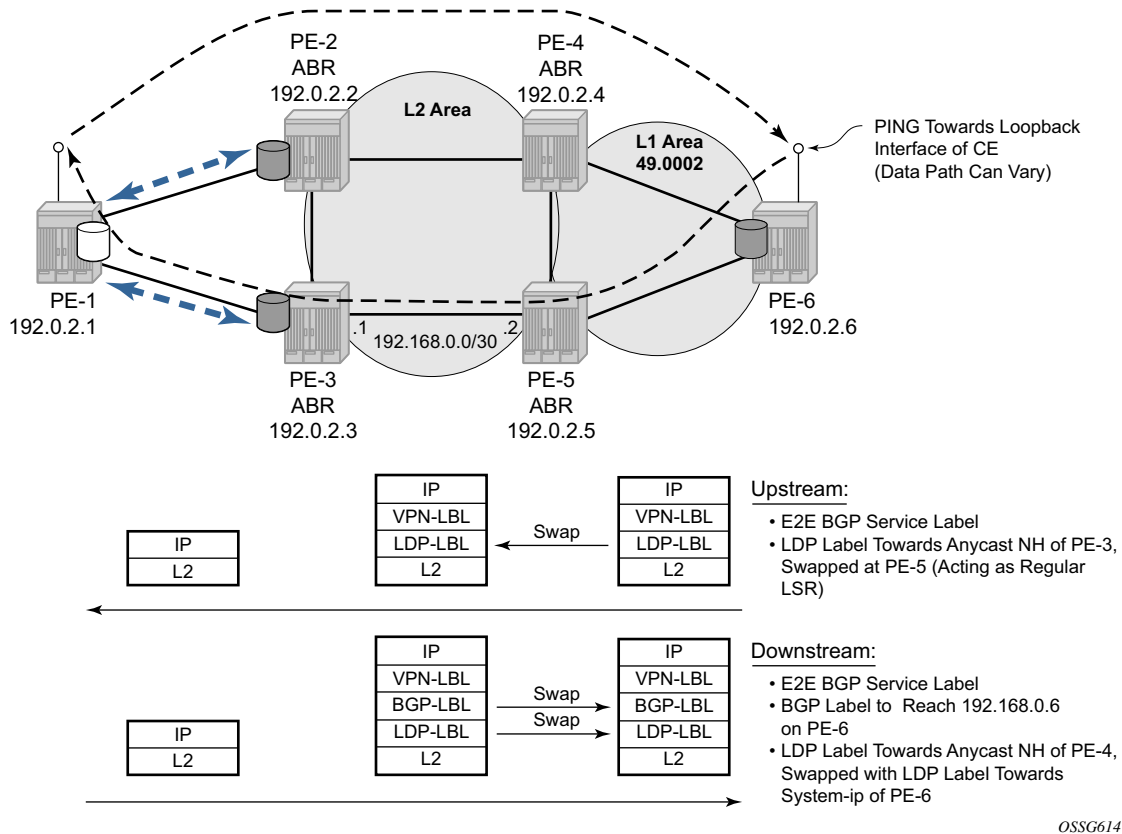


Figure 139: IP-VPN with Anycast NH, Data Path

Notice that the data path is not symmetric. By modifying BGP attributes, the path can be influenced so that the upstream and downstream directions follow the same path, but this is out-of-scope for this document.

Conclusion

BGP anycast provides an end-to-end MPLS reach ability in large MPLS domains, where a single IGP area cannot be deployed or even when a single IGP is not sufficient.

The feature provides a transport layer which can be used for any kind of service on the 7750 SR. On top of that, BGP anycast can also be used to advertise redundant IP-VPN routes into large MPLS domains. In this case, BGP anycast is not offering a redundant transport layer, but redundancy at the service layer.

The (context-specific label-switching) mechanism where a neighboring anycast ABR will install advertised labels from its anycast neighbor in the LFIB contributes to a fast convergence after nodal or link failures.

Conclusion