



BGP Best External

The BGP Best External feature provides the network with a backup external route to avoid loss of connectivity of the primary external route. The BGP Best External feature advertises the most preferred route among those received from external neighbors as a backup route. This feature is beneficial in active-backup topologies, where service providers use routing policies that cause a border router to choose a path received over an Interior Border Gateway Protocol (iBGP) session (of another border router) as the best path for a prefix even if it has an Exterior Border Gateway Protocol (eBGP) learned path. This active-backup topology defines one exit or egress point for the prefix in the autonomous system and uses the other points as backups if the primary link or eBGP peering is unavailable. The policy causes the border router to hide the paths learned over its eBGP sessions from the autonomous system because it does not advertise any path for such prefixes. To cope with this situation, some devices advertise one externally learned path called the best external path.

- [Prerequisites for BGP Best External, on page 1](#)
- [Restrictions for BGP Best External, on page 1](#)
- [Information About BGP Best External, on page 2](#)
- [How to Configure BGP Best External, on page 7](#)
- [Configuration Examples for BGP Best External, on page 22](#)
- [Additional References, on page 27](#)
- [Feature Information for BGP Best External, on page 28](#)

Prerequisites for BGP Best External

- The Bidirectional Forwarding Detection (BFD) protocol must be enabled to quickly detect link failures.
- Ensure that the BGP and the Multiprotocol Label Switching (MPLS) network is up and running with the customer site connected to the provider site by more than one path (multihomed).
- The backup path must have a unique next hop that is not the same as the next hop of the best path.
- BGP must support lossless switchover between operational paths.

Restrictions for BGP Best External

- The BGP Best External feature will not install a backup path if BGP Multipath is installed and a multipath exists in the BGP table. One of the multipaths automatically acts as a backup for the other paths.
- The BGP Best External feature is not supported with the following features:

- MPLS VPN Carrier Supporting Carrier
- MPLS VPN Per Virtual Routing and Forwarding (VRF) Label
- The BGP Best External feature cannot be configured with Multicast or L2VPN VRF address families.
- The BGP Best External feature cannot be configured on a route reflector, unless it is running Cisco IOS XE Release 3.4S or later.
- The BGP Best External feature does not support NSF/SSO. However, ISSU is supported if both Route Processors have the BGP Best External feature configured.
- The BGP Best External feature can only be configured on VPNv4, VPNv6, IPv4 VRF, and IPv6 VRF address families.
- When you configure the BGP Best External feature using the **bgp advertise-best-external** command, you need not enable the BGP PIC feature with the **bgp additional-paths install** command. The BGP PIC feature is automatically enabled by the BGP Best External feature.
- When you configure the BGP Best External feature, it will override the functionality of the "MPLS VPN--BGP Local Convergence" feature. However, you do not have to remove the **protection local-prefixes** command from the configuration.
- BGP Best External Path with MPLS VPN Inter-AS Option C is supported with deployments that use IPv4 addresses and labels. The configuration is not supported with IPv6 addresses and labels.

Information About BGP Best External

BGP Best External Overview

Service providers use routing policies that cause a border router to choose a path received over an iBGP session (of another border router) as the best path for a prefix even if it has an eBGP learned path. This practice is popularly known as active-backup topology and is done to define one exit or egress point for the prefix in the autonomous system and to use the other points as backups if the primary link or eBGP peering is unavailable.

The policy, though beneficial, causes the border router to hide the paths learned over its eBGP sessions from the autonomous system because the border router does not advertise any path for such prefixes. To cope with this situation, some routers advertise one externally learned path called the best external path. The best external behavior causes the BGP selection process to select two paths to every destination:

- The best path is selected from the complete set of routes known to that destination.
- The best external path is selected from the set of routes received from its external peers.

BGP advertises the best path to external peers. Instead of withdrawing the best path from its internal peers when it selects an iBGP path as the best path, BGP advertises the best external path to the internal peers.

The BGP Best External feature is an essential component of the Prefix-Independent Convergence (PIC) edge for both Internet access and MPLS VPN scenarios and makes alternate paths available in the network in the active-backup topology.

What the Best External Route Means

The BGP Best External feature uses a “best external route” as a backup path, which, according to draft-marques-idr-best-external, is the most preferred route among those received from external neighbors. The most preferred route from external neighbors can be the following:


- Two routers in different clusters that have an Interior Border Gateway Protocol (iBGP) session between them.
- Two routers in different autonomous systems of a confederation that have an External Border Gateway Protocol (eBGP) session between them.

The best external route might be different from the best route installed in the Routing Information Base (RIB). The best route could be an internal route. By allowing the best external route to be advertised and stored, in addition to the best route, networks gain faster restoration of connectivity by providing additional paths that may be used if the primary path fails.

How the BGP Best External Feature Works

The BGP Best External feature is based on Internet Engineering Task Force (IETF) draft-marques-idr-best-external.txt. The BGP Best External feature advertises a best external route to its internal peers as a backup route. The backup route is stored in the RIB and Cisco Express Forwarding. If the primary path fails, the BGP PIC functionality enables the best external path to take over, enabling faster restoration of connectivity.

Figure 1: MPLS VPN: Best External at the Edge of MPLS VPN



The figure above shows an MPLS VPN using the BGP Best External feature. The network includes the following components:

- eBGP sessions exist between the provider edge (PE) and customer edge (CE) routers.
- PE1 is the primary router and has a higher local preference setting.
- Traffic from CE2 uses PE1 to reach router CE1.
- PE1 has two paths to reach CE1.
- CE1 is dual-homed with PE1 and PE2.

- PE1 is the primary path and PE2 is the backup path.

In the figure above, traffic in the MPLS cloud flows through PE1 to reach CE1. Therefore, PE2 uses PE1 as the best path and PE2 as the backup path.

PE1 and PE2 are configured with the BGP Best External feature. BGP computes both the best path (the PE1-CE1 link) and a backup path (PE2) and installs both paths into the RIB and Cisco Express Forwarding. The best external path (PE2) is advertised to the peer routers, in addition to the best path.

When Cisco Express Forwarding detects a link failure on the PE1-CE1 link, Cisco Express Forwarding immediately switches to the backup path PE2. Traffic is quickly rerouted due to local Fast Convergence in Cisco Express Forwarding using the backup path. Thus, traffic loss is minimized and fast convergence is achieved.

Configuration Modes for Enabling BGP Best External

You can enable the BGP Best External feature in different modes, each of which protects Virtual Routing and Forwarding (VRF) in its own way:

- If you issue the **bgp advertise-best-external** command in VPNv4 address family configuration mode, it applies to all IPv4 VRFs. If you issue the command in this mode, you need not issue it for specific VRFs.
- If you issue the **bgp advertise-best-external** command in IPv4 address family configuration mode, it applies only to that VRF.

BGP Best External Path on RR for Intercluster

Beginning with Cisco IOS XE Release 3.4S, BGP Best External is extended to BGP Best External for Intercluster RRs. This feature provides path diversity between RR clusters, providing best external functionality toward non-client iBGP peers. The feature is also known as the “intercluster best external path.”

Best external path at an RR means the best path within the RR’s cluster. This path might also be referred to as the best internal path.

When an RR (RR1) chooses a non-client iBGP path (that is, a path learned from another RR, let’s say RR2) as its overall best, with the BGP Best External for Intercluster RRs feature, RR1 will be able to advertise its best internal path to the non-client iBGP peers. This will help RR2 to learn an additional path, providing a diverse path.

Best external functionality at RRs is only for non-client iBGP peers. An RR cannot advertise best external paths to its clients because it has to advertise its overall bestpath (which can be either a client path or non-client or eBGP path).

The best external path calculated by the RR is the best internal path for the cluster. It will be advertised to the non-client iBGP peers only when the overall best path at this RR is a non-client iBGP path.

When there are multiple RRs, each in its own cluster, each RR must have the **neighbor advertise best-external** command configured for each of its neighbor RRs.

If the RR is in the forwarding plane, the **bgp additional paths install** command is necessary.

CLI Differences for Best External Path on an RR for Intercluster

Prior to Cisco IOS XE Release 3.4S, the BGP Best External feature was allowed on a PE only, and it was configured by the **bgp advertise-best-external** command. The calculation of the backup path, installation, and advertisement were tied together in one command.

Beginning with Cisco IOS XE Release 3.4S, the BGP Best External feature is allowed on PEs and RRs. The functionality of the **bgp advertise-best-external** command is divided among the following three commands that calculate, install, and advertise the best external path:

- **bgp additional-path select best-external**
- **bgp additional-path install**
- **neighbor advertise diverse-path best-external**

If the **bgp additional-path select best-external** command is not configured, the system will calculate and install the best external path, but not advertise it.

The **neighbor advertise diverse-path best-external** command enables the advertisement of the best external path to the specified neighbor.

Rules Used to Calculate the BGP Best External Path for Intercluster RRs

The best internal path implementation on an RR toward non-clients (different cluster RRs) is calculated based on the following rules:

1. Calculate the overall primary bestpath on the RR per the normal bestpath selection rules.
2. If a backup path configuration is enabled, calculate the second bestpath (which is a different path from the primary bestpath selected in Rule 1 and has a different nexthop from this bestpath), which is marked as the backup path. Backup path selection is enabled using the **bgp additional-paths install** or **bgp additional-paths select [best-external] [backup]** command.
3. If the overall best path on the RR is a non-client iBGP path and not an eBGP path, calculate the best external/internal path from the remaining paths after excluding results from Rule 1 and Rule 2 and by ignoring all the other paths from the other clusters and run normal bestpath rules by including all the remaining eBGP and iBGP paths. Select the newly obtained bestpath and mark it as the best internal path.
4. Advertise this best internal path, which is either eBGP (received from CE peers for RR/ASBR) or iBGP (received from RR clients) toward the non-client RRs when **neighbor advertise best-external** is configured towards the non-client RRs.
5. If the overall bestpath is a path received from either an RR client or eBGP peer (in case of RR/ASBR) either an iBGP or an eBGP path will be chosen as bestpath per the normal bestpath algorithm. Because the overall bestpath is an internal client path, the normal advertisement rules will automatically advertise this path to non-client iBGP peers/RRs. This behavior is the same as the existing behavior (when best external is not enabled on RRs) when an RR client's path is chosen as the overall bestpath.
6. We do not allow a best external path to be configured on an RR towards RR-clients. The **neighbor advertise best-external** command can be configured on RR/ASBR only for non-clients or peering with RRs in the other clusters.
7. When multipath is enabled on the RR and only when the overall bestpath is from a non-client and if some of the intracluster client paths are also marked as multipaths, when best external is enabled on the RR

(**neighbor advertise best-external** towards the RR non-client), the algorithm selects the older multipath among the intra-cluster client multipaths (paths obtained from RR clients and eBGP peers within the cluster) and marks it as best internal path and announces it to the non-clients as best external path, so that the non-clients get path diversity from this cluster. If there are no intra-cluster multipaths found, we choose the best external path per Rules 3 through 5.

BGP Best External Path with MPLS Inter-AS Options B and C

BGP Best External Path with MPLS VPN Inter-AS Option B

With this feature, you can configure the border routers in an MPLS VPN Inter-AS Option B deployment to compute and advertise a best external path based on routes received from EBGP peers. The border routers advertise the best external path so computed to Internal BGP (IBGP) peers. The best external path is advertised in addition to the best path and serves as a back-up path. If a link in the best path fails, traffic flows along the best external path.

Consider the MPLS VPN Inter-AS Option B deployment shown in the following diagram:

Figure 2: MPLS VPN Inter-AS Option B Deployment

ASBR1 and ASBR2 exchange VPNv4/VPNv6 addresses, PE node loopback addresses, and labels using External BGP (EBGP). Similarly, ASBR3 and ASBR4 exchange VPNv4/VPNv6 addresses, PE node loopback addresses, and labels using EBGP.

Suppose that ASBR3 is configured so that the routes received from ASBR1 have a higher Local Preference, and that ASBR1 is the primary router and ASBR3 is the back-up router. With this configuration, traffic for AS2 exits AS1 through ASBR1 along the best path through ASBR2. This path through ASBR1 and ASBR2 is the best path that ASBR1 advertises to IBGP peers after assigning a label to the path. The VPNv4/VPNv6 labels and PE loopback addresses from AS2 are also advertised along with the label for the best path. ASBR3 installs the best path as the route for any traffic that it must send to AS2. In doing so, ASBR3 ignores the path through ASBR4 that was learned using EBGP.

You can configure the BGP Best External Path on ASBR1 and ASBR3 to provide a back-up path for the traffic between AS1 and AS2 if the link between ASBR1 and ASBR2 fails.

When you configure BGP Best External Path on ASBR1 and ASBR3, ASBR3 calculates a best external path to AS2, through ASBR4, assigns a label to this path, and installs this path into the Label Forwarding Information Base (LFIB) along with the best path. ASBR3 advertises both the best path and the best external path to ASBR1.

ASBR1 uses the best external path received from ASBR3 to calculate a back-up path and assigns a label to it. ASBR1 installs this best external path in the LFIB along with the best path through ASBR2. ASBR1 advertises the best path and best external path to IBGP peers.


With the best external path installed, suppose PE1 has traffic for PE2 and the ASBR1-ASBR2 link fails, the traffic flows through the best external path along PE1-ASBR1-ASBR3-ASBR4-PE2.

BGP Best External Path with MPLS VPN Inter-AS Option C

With this feature, you can configure the border routers in an MPLS VPN Inter-AS Option C deployment to compute and advertise a best external path based on routes received from EBGP peers. The border routers advertise the best external path so computed to Internal BGP (IBGP) peers through a route reflector. The best external path is advertised in addition to the best path and serves as a back-up path. If a link in the best path fails, traffic flows along the best external path.

Consider the MPLS VPN Inter-AS Option C deployment shown in the following diagram:

Figure 3: MPLS VPN Inter-AS Option C Deployment



The route reflectors, RR1 and RR2, exchange VPNv4 routes and labels using multiprotocol EBGP. The routes preserve VPN labels and next-hop information between ASs. After RR1 reflects the exchanged routes, the next-hop for VPN traffic received by PE1 is either PE2 or PE3.

ASBR1 and ASBR2 exchange IPv4 loopback addresses and labels for PE2 and PE3 using EBGP. Similarly, ASBR3 and ASBR4 IPv4 addresses and labels using EBGP.

Suppose that ASBR3 is configured so that the routes received from ASBR1 have a higher Local Preference, and that ASBR1 is the primary router and ASBR3 is the back-up router. With this configuration, traffic for AS2 exits AS1 through ASBR1 along the best path through ASBR2. This path through ASBR1 and ASBR2 is the best path that ASBR1 advertises to the route reflector RR1 after assigning a label. The PE loopback addresses from AS2 are also advertised to RR1 along with the label for the best path. ASBR3 installs the best path as the route for any traffic that it must send to AS2. In doing so, ASBR3 ignores the path through ASBR4 that was learned using EBGP.

When you configure BGP Best External Path on ASBR1 and ASBR3, ASBR3 calculates a best external path to AS2, through ASBR4, assigns a label to this path, and installs this path into the Label Forwarding Information Base (LFIB) along with the best path. ASBR3 advertises both the best path and the best external path to ASBR1.

ASBR1 uses the best external path received from ASBR3 to calculate a back-up path and assigns a label to it. ASBR1 installs this best external path in the LFIB along with the best path through ASBR2. ASBR1 advertises the best path and the best external path to RR1 for reflection to clients in AS1.

With the best external path installed, suppose PE1 has traffic for PE2 and the ASBR1-ASBR2 link fails, the traffic flows through the best external path along PE1-ASBR1-ASBR3-ASBR4-PE2.

How to Configure BGP Best External

Configuring the BGP Best External Feature

Perform the following task to configure the BGP Best External feature. This task shows how to configure the BGP Best External feature in either an IPv4 or VPNv4 address family. In VPNv4 address family configuration mode, the BGP Best External feature applies to all IPv4 Virtual Routing Forwarding (VRF); you need not

configure it for specific VRFs. If you issue the **bgp advertise-best-external** command in IPv4 VRF address family configuration mode, the BGP Best External feature applies only to that VRF.

Before you begin

- Configure the MPLS VPN and verify that it is working properly before configuring the BGP Best External feature. See the "Configuring MPLS Layer 3 VPNs" section for more information.
- Configure multiprotocol VRFs to allow you to share route-target policies (import and export) between IPv4 and IPv6 or configure separate route-target policies for IPv4 and IPv6 VPNs. For information about configuring multiprotocol VRFs, see the "MPLS VPN--VRF CLI for IPv4 and IPv6 VPNs section".
- Ensure that the customer edge (CE) router is connected to the network by at least two paths.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router bgp** *autonomous-system-number*
4. Do one of the following:
 - **address-family ipv4** [**unicast** | **vrf** *vrf-name*]
 - or
 - **address-family vpv4** [**unicast**]
 - or
5. **bgp advertise-best-external**
6. **neighbor** *ip-address* **remote-as** *autonomous-system-number*
7. **neighbor** *ip-address* **activate**
8. **neighbor** *ip-address* **fall-over** [**bfd** | **route-map** *map-name*]
9. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Router# configure terminal	Enters global configuration mode.
Step 3	router bgp <i>autonomous-system-number</i> Example: Router(config)# router bgp 40000	Enters router configuration mode for the specified routing process.

	Command or Action	Purpose
Step 4	<p>Do one of the following:</p> <ul style="list-style-type: none"> • address-family ipv4 [unicast vrf vrf-name] • or • address-family vpnv4 [unicast] • or <p>Example:</p> <pre>Router(config-router)# address-family ipv4 unicast</pre> <p>Example:</p> <pre>Router(config-router)# address-family vpnv4</pre>	<p>Specifies the IPv4 or VPNv4 address family and enters address family configuration mode.</p> <ul style="list-style-type: none"> • The unicast keyword specifies the IPv4 or VPNv4 unicast address family. • The vrf keyword and <i>vrf-name</i> argument specify the name of the VRF instance to associate with subsequent IPv4 address family configuration mode commands.
Step 5	<p>bgp advertise-best-external</p> <p>Example:</p> <pre>Router(config-router-af)# bgp advertise-best-external</pre>	<p>Calculates and uses an external backup path and installs it into the RIB and Cisco Express Forwarding.</p>
Step 6	<p>neighbor ip-address remote-as autonomous-system-number</p> <p>Example:</p> <pre>Router(config-router-af)# neighbor 192.168.1.1 remote-as 45000</pre>	<p>Adds the IP address of the neighbor in the specified autonomous system to the IPv4 multiprotocol BGP neighbor table of the local router.</p> <ul style="list-style-type: none"> • By default, neighbors that are defined using the neighbor remote-as command in router configuration mode exchange only IPv4 unicast address prefixes. To exchange other address prefix types, neighbors must also be activated using the neighbor activate command in address family configuration mode for the other prefix types.
Step 7	<p>neighbor ip-address activate</p> <p>Example:</p> <pre>Router(config-router-af)# neighbor 192.168.1.1 activate</pre>	<p>Enables the neighbor to exchange prefixes for the IPv4 unicast address family with the local router.</p>
Step 8	<p>neighbor ip-address fall-over [bfd route-map map-name]</p> <p>Example:</p> <pre>Router(config-router-af)# neighbor 192.168.1.1 fall-over bfd</pre>	<p>Configures the BGP peering to use fast session deactivation and enables BFD protocol support for failover.</p> <ul style="list-style-type: none"> • BGP will remove all routes learned through this peer if the session is deactivated.
Step 9	<p>end</p> <p>Example:</p> <pre>Router(config-router-af)# end</pre>	<p>(Optional) Exits address family configuration mode and returns to privileged EXEC mode.</p>

Verifying the BGP Best External Feature

Perform the following task to verify that the BGP Best External feature is configured correctly.

SUMMARY STEPS

1. **enable**
2. **show vrf detail**
3. **show ip bgp ipv4 mdt all | rd vrf} | multicast | tunnel unicast** or **show ip bgp vpn4 all rd route-distinguisher | vrf vrf-name rib-failure ip-prefix/length longer-prefixes]] network-address mask longer-prefixes]] cidr-only community community-list dampened-paths filter-list [flap-statistics inconsistent-as neighbors paths line]] peer-group quote-regexp regexp [summary labels**
4. **show bgp vpn4 unicast vrf vrf-name ip-address**
5. **show ip route vrf vrf-name repair-paths ip-address**
6. **show ip cef vrf vrf-name ip-address detail**

DETAILED STEPS

Step 1 enable

Use this command to enable privileged EXEC mode. Enter your password, if prompted. For example:

Example:

```
Router> enable
Router#
```

Step 2 show vrf detail

Use this command to verify that the BGP Best External feature is enabled. The following **show vrf detail** command output shows that the BGP Best External feature is enabled.

Example:

```
Router# show vrf detail
VRF test1 (VRF Id = 1); default RD 400:1; default VPNID <not set>
  Interfaces:
    Se4/0
Address family ipv4 (Table ID = 1 (0x1)):
  Export VPN route-target communities
    RT:100:1                RT:200:1                RT:300:1
    RT:400:1
  Import VPN route-target communities
    RT:100:1                RT:200:1                RT:300:1
    RT:400:1
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix

  Prefix protection with additional path enabled
  Address family ipv6 not active.
```

Step 3 show ip bgp ipv4 mdt all | rd vrf} | multicast | tunnel unicast or show ip bgp vpn4 all rd route-distinguisher | vrf vrf-name rib-failure ip-prefix/length longer-prefixes]] network-address mask longer-prefixes]] cidr-only

**community community-list dampened-paths filter-list] [flap-statistics inconsistent-as neighbors paths line]]
peer-group quote-regexp regexp] [summary labels**

Use this command to verify that the best external route is advertised. In the command output, the code b indicates a backup path and the code x designates the best external path.

Example:

```
Router# show ip bgp vpnv4 all
BGP table version is 1104964, local router ID is 10.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, multipath,
b backup-path, x best-external
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 11:12 (default for vrf blue)
*>i1.0.0.1/32      10.10.3.3          0      200    0 1 ?
* i                10.10.3.3          0      200    0 1 ?
*                  10.0.0.1           0      200    0 1 ?
*bx               10.0.0.1           0      200    0 1 ?
*                  10.0.0.1           0      200    0 1 ?
```

Step 4 **show bgp vpnv4 unicast vrf vrf-name ip-address**

Use this command to verify that the best external route is advertised.

Example:

```
Router# show bgp vpnv4 unicast vrf vpn1 10.10.10.10
BGP routing table entry for 10:10:10.10.10/32, version 10
Paths: (2 available, best #1, table vpn1)
  Advertise-best-external
    Advertised to update-groups:
      1          2
    200
      10.6.6.6 (metric 21) from 10.6.6.6 (10.6.6.6)
      Origin incomplete, metric 0, localpref 200, valid, internal, best
      Extended Community: RT:1:1
      mpls labels in/out 23/23
    200
      10.1.2.1 from 10.1.2.1 (10.1.1.1)
      Origin incomplete, metric 0, localpref 100, valid,
external, backup/repair, advertise-best-external
      Extended Community: RT:1:1 , recursive-via-connected
      mpls labels in/out 23/nolabel
```

Step 5 **show ip route vrf vrf-name repair-paths ip-address**

Use this command to display the repair route.

Example:

```
Router# show ip route vrf vpn1 repair-paths

Routing Table: vpn1
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP
```

```

+ - replicated route, % - next hop override
Gateway of last resort is not set
10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
B   10.1.1.0/24 [200/0] via 10.6.6.6, 00:38:33
    [RPR][200/0] via 10.1.2.1, 00:38:33
B   10.1.1.1/32 [200/0] via 10.6.6.6, 00:38:33
    [RPR][200/0] via 10.1.2.1, 00:38:33
10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
C   10.1.2.0/24 is directly connected, Ethernet0/0
L   10.1.2.2/32 is directly connected, Ethernet0/0
B   10.1.6.0/24 [200/0] via 10.6.6.6, 00:38:33
    [RPR][200/0] via 10.1.2.1, 00:38:33

```

Step 6 `show ip cef vrf vrf-name ip-address detail`

Use this command to display the best external route.

Example:

```

Router# show ip cef vrf test 10.71.8.164 detail


10.71.8.164/30, epoch 0, flags rib defined all labels
 recursive via 10.249.0.102 label 35
  nexthop 10.249.246.101 Ethernet0/0 label 25
 recursive via 10.249.0.104 label 28,
repair
  nexthop 10.249.246.101 Ethernet0/0 label 24

```

Configuring Best External Path on an RR for an Intercluster

Perform the following task to configure a best external path on an RR for an intercluster. The steps in this particular task configure RR1 in the figure below, in the IPv4 address family. The step that configures address family lists the other address families supported.

Figure 4: Scenario for Configuring a BGP Best External Path on a RR for an Intercluster



SUMMARY STEPS

1. enable
2. configure terminal
3. router bgp *autonomous-system-number*
4. neighbor *ip-address* remote-as *autonomous-system-number*
5. neighbor *ip-address* remote-as *autonomous-system-number*
6. address-family ipv4 unicast
7. neighbor *ip-address* activate
8. neighbor *ip-address* activate
9. bgp additional-paths select best-external
10. bgp additional-paths install
11. neighbor *ip-address* advertise best-external
12. neighbor *ip-address* advertise best-external
13. end

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	configure terminal Example: Router# configure terminal	Enters global configuration mode.
Step 3	router bgp <i>autonomous-system-number</i> Example: Router(config)# router bgp 1	Enters router configuration mode for the specified routing process.
Step 4	neighbor <i>ip-address</i> remote-as <i>autonomous-system-number</i> Example: Router(config-router)# neighbor 10.5.1.1 remote-as 1	Adds an entry to the BGP or multiprotocol BGP neighbor table. • This step is for RR3.
Step 5	neighbor <i>ip-address</i> remote-as <i>autonomous-system-number</i> Example: Router(config-router)# neighbor 10.5.1.2 remote-as 1	Adds an entry to the BGP or multiprotocol BGP neighbor table. • This step is for RR6.
Step 6	address-family ipv4 unicast Example: Router(config-router)# address-family ipv4 unicast	Specifies the address family and enters address family configuration mode. • Supported address families are ipv4 unicast, vpv4 unicast, ipv6 unicast, vpv6 unicast, ipv4+label, and ipv6+label.
Step 7	neighbor <i>ip-address</i> activate Example: Router(config-router-af)# neighbor 10.5.1.1 activate	Enables the exchange of information with a BGP neighbor. • This step is for RR3.
Step 8	neighbor <i>ip-address</i> activate Example: Router(config-router-af)# neighbor 10.5.1.2 activate	Enables the exchange of information with a BGP neighbor. • This step is for RR6.
Step 9	bgp additional-paths select best-external Example:	Configures the system to calculate a best external path (external to RR cluster).

	Command or Action	Purpose
	<pre>Router(config-router-af)# bgp additional-paths select best-external</pre>	
Step 10	bgp additional-paths install Example: <pre>Router(config-router-af)# bgp additional-paths install</pre>	Enables BGP to calculate a backup path for a given address family and to install it into the RIB and CEF. <ul style="list-style-type: none"> • This step is necessary if the RR is enabled for forwarding (the RR is in the forwarding plane). Otherwise, this step is unnecessary.
Step 11	neighbor ip-address advertise best-external Example: <pre>Router(config-router-af)# neighbor 10.5.1.1 advertise best-external</pre>	(Optional) Configures a neighbor to receive the best external path in an advertisement. <ul style="list-style-type: none"> • This step is for RR3.
Step 12	neighbor ip-address advertise best-external Example: <pre>Router(config-router-af)# neighbor 10.5.1.2 advertise best-external</pre>	(Optional) Configures a neighbor to receive the best external path in an advertisement. <ul style="list-style-type: none"> • This step is for RR6.
Step 13	end Example: <pre>Router(config-router-af)# end</pre>	(Optional) Exits address family configuration mode and returns to privileged EXEC mode.

In the scenario shown above, the following paths are selected as best path, backup bath, and best internal path on the three RRs located in the three different clusters:

On RR1:

On RR3:

On RR6:

To Reach Prefix 10/8	Next Hop:
	PE5 (best path, local preference = 200)
	PE3 (backup path, local preference = 150)
	PE3 (best internal path, local preference = 150)
To Reach Prefix 10/8	Next Hop:
	PE5 (best path, local preference = 200)
	PE6 (backup path, local preference = 50)
	PE3 (received as best external path from RR1, local preference = 150)

To Reach Prefix 10/8	Next Hop:
	PE5 (best path, local preference = 200)
	PE6 (backup path, local preference = 50)
	PE3 (received as best external path from RR1, local preference = 150)

Configure BGP Best External Path with MPLS VPN Inter-AS Option B

The configuration instructions in this section are limited to the additional configuration required to enable the computation and advertisement of BGP Best External Path in an MPLS VPN Inter-AS Option B deployment.

To configure BGP Best External Path,

- *Configure the Primary ASBR to Compute and Install a Back-up Path*
- *Configure the Secondary ASBR to Compute, Install, and Advertise Best External Path*



Note An ASBR becomes the Primary ASBR when you configure the routes the ASBR propagates to have a higher local preference on other ASBRs.

Configure the Primary ASBR to Compute and Install a Back-up Path

Step 1 enable

Example:

```
Device> enable
```

Enables privileged EXEC mode. Enter your password if prompted.

Step 2 configure terminal

Example:

```
Device# configure terminal
```

Enters global configuration mode.

Step 3 router bgp *autonomous-system-number*

Example:

```
Device(config)# router bgp 1
```

Configures a BGP routing process, and enters router configuration mode for the specified routing process.

Step 4 address-family vpnv4

Example:

```
Device(config-router)# address-family vpnv4
```

Specifies the VPNv4 address family and enters address family configuration mode.

Step 5 **bgp additional-paths install****Example:**

```
Device(config-router-af)# bgp additional-paths install
```

Computes and installs the best external path in the LFIB, and advertises the path to IBGP peers.

Step 6 **exit-address-family****Example:**

```
Device(config-router-af)# exit-address-family
```

Exits address family configuration mode.

Step 7 **end****Example:**

```
Device(config-router)# end
```

(Optional) Exits to privileged EXEC mode.

Configure the Secondary ASBR to Compute, Install, and Advertise Best External Path

Step 1 **enable****Example:**

```
Device> enable
```

Enables privileged EXEC mode. Enter your password if prompted.

Step 2 **configure terminal****Example:**

```
Device# configure terminal
```

Enters global configuration mode.

Step 3 **router bgp *autonomous-system-number*****Example:**

```
Device(config)# router bgp 1
```

Configures a BGP routing process, and enters router configuration mode for the specified routing process.

Step 4 **address-family vpnv4****Example:**

```
Device(config-router)# address-family vpnv4
```

Specifies the VPNv4 address family and enters address family configuration mode.

Step 5 **bgp advertise-best-external****Example:**

```
Device(config-router-af)# bgp advertise-best-external
```

Computes and installs the best external path in the LFIB, and advertises the path to the primary ASBR.

Step 6 **exit-address-family**

Example:

```
Device(config-router-af)# exit-address-family
```

Exits address family configuration mode.

Step 7 **end**

Example:

```
Device(config-router)# end
```

(Optional) Exits to privileged EXEC mode.

Configure BGP Best External Path with MPLS VPN Inter-AS Option C

The configuration instructions in this section are limited to the additional configuration required to enable the computation and advertisement of BGP Best External Path in an MPLS VPN Inter-AS Option C deployment.

To configure BGP Best External Path,

- *Configure the Primary ASBR to Compute and Install a Back-up Path*
- *Configure the Secondary ASBR to Compute, Install, and Advertise Best External Path*



Note An ASBR becomes the Primary ASBR when you configure the routes the ASBR propagates to have a higher local preference on other ASBRs.

Configure the Primary ASBR to Compute and Install a Back-up Path

Step 1 **enable**

Example:

```
Device> enable
```

Enables privileged EXEC mode. Enter your password if prompted.

Step 2 **configure terminal**

Example:

```
Device# configure terminal
```

Enters global configuration mode.

Step 3 **router bgp** *autonomous-system-number*

Example:

```
Device(config)# router bgp 1
```

Configures a BGP routing process, and enters router configuration mode for the specified routing process.

Step 4 **address-family ipv4 unicast**

Example:

```
Device(config-router)# address-family ipv4 unicast
```

Specifies the IPv4 unicast address family and enters address family configuration mode.

Step 5 **bgp additional-paths install**

Example:

```
Device(config-router-af)# bgp additional-paths install
```

Computes and installs the best external path in the LFIB, and advertises the path to IBGP peers.

Step 6 **exit-address-family**

Example:

```
Device(config-router-af)# exit-address-family
```

Exits address family configuration mode.

Step 7 **end**

Example:

```
Device(config-router)# end
```

(Optional) Exits to privileged EXEC mode.

Configure the Secondary ASBR to Compute, Install, and Advertise Best External Path

Perform this procedure to configure BGP Best External Path on the Secondary ASBR in an MPLS VPN Inter-AS Option C deployment.

Step 1 **enable**

Example:

```
Device> enable
```

Enables privileged EXEC mode. Enter your password if prompted.

Step 2 **configure terminal**

Example:

```
Device# configure terminal
```

Enters global configuration mode.

Step 3 **router bgp *autonomous-system-number***

Example:

```
Device(config)# router bgp 1
```

Configures a BGP routing process, and enters router configuration mode for the specified routing process.

Step 4 **address-family ipv4 unicast**

Example:

```
Device(config-router)# address-family ipv4 unicast
```

Specifies the IPv4 unicast address family and enters address family configuration mode.

Step 5 **bgp advertise-best-external****Example:**

```
Device(config-router-af)# bgp advertise-best-external
```

Computes and installs the best external path in the LFIB, and advertises the path to the primary ASBR.

Step 6 **exit-address-family****Example:**

```
Device(config-router-af)# exit-address-family
```

Exits address family configuration mode.

Step 7 **end****Example:**

```
Device(config-router)# end
```

(Optional) Exits to privileged EXEC mode.

Verify BGP Best External Path with MPLS VPN Inter-AS Option B or MPLS VPN Inter-AS Option C

Verify VPNv4 routes

To check for BGP external paths among the installed VPNv4 routes, issue the command **show bgp vpnv4 unicast all**.

In the command output, the status code **x best-external** indicates that a path is a best external path, as shown in the following example.

```
Device#show bgp vpnv4 unicast all
BGP table version is 12211, local router ID is 44.44.44.44
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
               t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

      Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 200:200
*>i 19.1.1.1/32        33.33.33.33          0      600          0 200 2001 i
*b x                21.1.1.2              0      600          0 200 2001 i
```

```
Device#show bgp vpnv4 unicast all 19.1.1.1/32
BGP routing table entry for 200:200:19.1.1.1/32, version 7682
Paths: (2 available, best #1, no table)
Advertise-best-external
Advertised to update-groups:
```

```

      1          2
Refresh Epoch 1
200 2001
  33.33.33.33 (via default) from 33.33.33.33 (33.33.33.33)
    Origin IGP, metric 0, localpref 600, valid, internal, best
    Extended Community: RT:100:100 , recursive-via-host
    mpls labels in/out 2036/2036
    rx pathid: 0, tx pathid: 0x0
    Updated on Jun 22 2020 20:54:43 PST
Refresh Epoch 1
200 2001
  21.1.1.2 (via default) from 21.1.1.2 (22.22.22.22)
    Origin IGP, localpref 100, valid, external, backup/repair, advertise-best-external
    Extended Community: RT:100:100 , recursive-via-connected
    mpls labels in/out 2036/2035
    rx pathid: 0, tx pathid: 0
    Updated on Jun 22 2020 20:54:43 PST
Device#

```

Verify IPv4 LU routes

To check for BGP external paths among the installed IPv4 Labelled Unicast (LU) routes, issue the command **show bgp ipv4 unicast**.

In the command output, the status code **x best-external** indicates that a path is a best external path, as shown in the following example.

```

Device#show bgp ipv4 unicast
BGP table version is 5007, local router ID is 44.44.44.44
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

      Network          Next Hop          Metric LocPrf Weight Path
*bi  11.11.11.11/32    33.33.33.33          11     600      0 i
*>   23.1.1.1          23.1.1.1            11           32768 i
*>i  66.66.66.66/32    33.33.33.33          11     600      0 200 i
*b x 21.1.1.2          21.1.1.2            11           0 200 i

```

```

Device#show bgp ipv4 unicast 66.66.66.66/32
BGP routing table entry for 66.66.66.66/32, version 7
Paths: (2 available, best #1, table default)
Advertise-best-external
  Advertised to update-groups:
    1          2          3
Refresh Epoch 1
200
  33.33.33.33 from 33.33.33.33 (33.33.33.33)
    Origin IGP, metric 11, localpref 600, valid, internal, best
    mpls labels in/out 19/19
    rx pathid: 0, tx pathid: 0x0
    Updated on Jun 22 2020 20:38:02 PST
Refresh Epoch 1
200
  21.1.1.2 from 21.1.1.2 (22.22.22.22)
    Origin IGP, metric 11, localpref 100, valid, external, backup/repair,
advertise-best-external , recursive-via-connected
    mpls labels in/out 19/18
    rx pathid: 0, tx pathid: 0

```

```
Updated on Jun 22 2020 20:37:58 PST
Device#
```

Verify BGP Best External Path Advertisement

To confirm that an ASBR is advertising a best external path for a VPNv4 unicast route, issue the command **show bgp vpnv4 unicast all neighbor *neighbour-ip-address* advertise-route**.

```
Device#show bgp vpnv4 unicast all neighbors 77.77.77.77 advertised-routes
BGP table version is 16792, local router ID is 44.44.44.44
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
               t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found
```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 200:200					
*b x 8.88.88.88/32	21.1.1.2		0	200	2001 i
*b x 88.88.88.88/32	21.1.1.2		0	200	2001 i

```
Total number of prefixes 2
Device#
```

To confirm that an ASBR is advertising a best external path for an IPv4 LU route, issue the command **show bgp ipv4 unicast all neighbor *neighbour-ip-address* advertise-route**.

Verify the Presence of a Best Internal Path and a Best External Path

To confirm that both a best internal path and a best external path are available for a node, issue the command **show ip cef *dest-ip-address* internal**.

```
Device#show ip cef 66.66.66.66 internal
66.66.66.66/32, epoch 0, flags [rlbls], RIB[B], refcnt 6, per-destination sharing
sources: RIB, RR
feature space:
  IPRM: 0x00018000
  LFD: 66.66.66.66/32 0 local labels
      contains path extension list
      dflt label switch chain 0x7FBDE68C24D0
      loadinfo 7FBDE68C24D0, per-session, 2 choices, flags 0493, 5 locks
subblocks:
  1 RR source [non-eos indirection, heavily shared]
    non-eos chain loadinfo 7FBDE1787F60, per-session, flags 0591, 3 locks
ifnums:
  Ethernet0/1(3): 11.1.1.3
  Ethernet1/3(9): 23.1.1.4
```

Configuration Examples for BGP Best External

Example: Configuring the BGP Best External Feature

The following example shows how to configure the BGP Best External feature in VPNv4 mode:

```
vrf definition test1
rd 400:1
```

```

route-target export 100:1
route-target export 200:1
route-target export 300:1
route-target export 400:1
route-target import 100:1
route-target import 200:1
route-target import 300:1
route-target import 400:1
address-family ipv4
exit-address-family
exit
!
interface Ethernet1/0
vrf forwarding test1
ip address 10.0.0.1 255.0.0.0
exit
!
router bgp 64500
no synchronization
bgp log-neighbor-changes
neighbor 10.5.5.5 remote-as 64500
neighbor 10.5.5.5 update-source Loopback0
neighbor 10.6.6.6 remote-as 64500
neighbor 10.6.6.6 update-source Loopback0
no auto-summary
!
address-family vpnv4

bgp advertise-best-external
neighbor 10.5.5.5 activate
neighbor 10.5.5.5 send-community extended
neighbor 10.6.6.6 activate
neighbor 10.6.6.6 send-community extended
exit-address-family
!
address-family ipv4 vrf test1
no synchronization
bgp recursion host
neighbor 192.168.13.2 remote-as 64511
neighbor 192.168.13.2 fall-over bfd
neighbor 192.168.13.2 activate
neighbor 192.168.13.2 as-override
exit-address-family

```

Example: Configuring a Best External Path on an RR for an Intercluster

The following example configures RR1 in the figure shown in the “Configuring a Best External Path on an RR for an Intercluster” section. RR1 is configured to calculate, install, and advertise the best external path to its intercluster RR neighbors.

RR1

```

router bgp 1
neighbor 10.5.1.1 remote-as 1
neighbor 10.5.1.2 remote-as 1
address-family ipv4 unicast
neighbor 10.5.1.1 activate
neighbor 10.5.1.2 activate
bgp additional-paths select best-external
bgp additional-paths install
neighbor 10.5.1.1 advertise best-external

```

```
neighbor 10.5.1.2 advertise best-external
end
```

Example: Configuring BGP Best External Path with MPLS VPN Inter-AS Option B



In this sample topology, EBGP VPNv4 peering exists between ASBR3 and ASBR5, and similarly between ASBR4 and ASBR2.

ASBR3 is configured as the primary ASBR for ASN 100. Thus, traffic flows from CE9 to CE8 through the path CE9 -> PE1 -> RR7 -> ASBR3 -> ASBR5 -> PERR6 -> CE8. In other words, the primary path is through RR7, ASBR3, and ASBR5.

The secondary ASBR, ASBR4, receives the VPNv4 route 88.88.88.88 (CE8) from ASBR3 (best internal path) and ASBR2 (best external path).

With BGP External Path configured, ASBR4 advertises the best external path to both ASBR3 and RR7. ASBR3 advertises the best external path as a back-up path along with the best path through ASBR3 and ASBR5. If the link between ASBR3 and ASBR5 fails, traffic from CE9 to CE8 flows through ASBR4 -> ASBR2 -> PERR6.

ASBR3 Configuration:

```
router bgp 100
  bgp log-neighbor-changes
  no bgp default route-target filter
  neighbor 12.1.1.5 remote-as 200
  neighbor 12.1.1.5 fall-over bfd
  neighbor 44.44.44.44 remote-as 100
  neighbor 44.44.44.44 update-source Loopback0
  neighbor 44.44.44.44 fall-over
  neighbor 77.77.77.77 remote-as 100
  neighbor 77.77.77.77 update-source Loopback0
  neighbor 77.77.77.77 fall-over
  !
  address-family vpnv4
    bgp additional-paths install
    neighbor 12.1.1.5 activate
    neighbor 12.1.1.5 send-community both
    neighbor 12.1.1.5 route-map rt-rewrite-map in
    neighbor 33.33.33.33 activate
    neighbor 33.33.33.33 send-community both
    neighbor 33.33.33.33 next-hop-self
    neighbor 33.33.33.33 route-map LOCPREF-600 in
    neighbor 77.77.77.77 activate
    neighbor 77.77.77.77 send-community both
    neighbor 77.77.77.77 next-hop-self
  exit-address-family
```


ASBR4 Configuration:

```

router bgp 100
  bgp log-neighbor-changes
  no bgp default route-target filter
  neighbor 21.1.1.2 remote-as 200
  neighbor 21.1.1.2 fall-over bfd
  neighbor 33.33.33.33 remote-as 100
  neighbor 33.33.33.33 update-source Loopback0
  neighbor 33.33.33.33 fall-over
  neighbor 77.77.77.77 remote-as 100
  neighbor 77.77.77.77 update-source Loopback0
  neighbor 77.77.77.77 fall-over
  !
  address-family vpnv4
    bgp advertise-best-external
    neighbor 21.1.1.2 activate
    neighbor 21.1.1.2 send-community both
    neighbor 21.1.1.2 route-map rt-rewrite-map in
    neighbor 33.33.33.33 activate
    neighbor 33.33.33.33 send-community both
    neighbor 33.33.33.33 next-hop-self
    neighbor 33.33.33.33 route-map LOCPREF-600 in
    neighbor 77.77.77.77 activate
    neighbor 77.77.77.77 send-community both
    neighbor 77.77.77.77 next-hop-self
  exit-address-family

```

Verification of Best External Path:

ASBR4 advertises the best external path to RR7. This advertisement can be verified using the **show bgp vpnv4 unicast all neighbor neighbour-ip-address advertise-route** command.

```

ASBR4#show bgp vpnv4 unicast all neighbors 77.77.77.77 advertised-routes
BGP table version is 16792, local router ID is 44.44.44.44
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
               t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

      Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 200:200
  *b x 8.88.88.88/32    21.1.1.2          0 200 2001 i
  *b x 88.88.88.88/32  21.1.1.2          0 200 2001 i

Total number of prefixes 2
ASBR4#

```

Example: Configuring BGP Best External Path with MPLS VPN Inter-AS Option C

In this sample topology, EBGP IPv4 Labelled Unicast (LU) peering is configured between ASBR3 and ASBR5, and similarly, between ASBR4 and ASBR2. EBGP VPNv4 multihop peering is configured between PERR1 and PERR6.

ASBR3 is configured as the primary ASBR for ASN 100. Thus, traffic flows from CE7 to CE8 through the path CE7 -> PERR1 -> ASBR3 -> ASBR5 -> PERR6 -> CE8. In other words, the primary path is through RR7, ASBR3, and ASBR5.

The secondary ASBR, ASBR4, receives the BGP route 66.66.66.66 (PERR6) from ASBR3 (best internal path) and ASBR2 (best external path).

With BGP External Path configured, ASBR4 advertises the best external path to both ASBR3 and PERR1. ASBR3 advertises the best external path as a back-up path along with the best path through ASBR3 and ASBR5. If the link between ASBR3 and ASBR5 fails, traffic from CE7 to CE8 flows through ASBR4 -> ASBR2 -> PERR6.

ASBR3 Configuration:

```
router bgp 100
  bgp log-neighbor-changes
  neighbor 11.11.11.11 remote-as 100
  neighbor 11.11.11.11 update-source Loopback0
  neighbor 12.1.1.5 remote-as 200
  neighbor 12.1.1.5 fall-over bfd
  neighbor 44.44.44.44 remote-as 100
  neighbor 44.44.44.44 update-source Loopback0
  neighbor 44.44.44.44 fall-over
  !
  address-family ipv4
    bgp redistribute-internal
    bgp additional-paths install
    network 11.11.11.11 mask 255.255.255.255
    neighbor 11.11.11.11 activate
    neighbor 11.11.11.11 next-hop-self all
    neighbor 11.11.11.11 route-map SET-MPLS-LABEL out
    neighbor 11.11.11.11 send-label
    neighbor 12.1.1.5 activate
    neighbor 12.1.1.5 route-map SET-MPLS-LABEL out
    neighbor 12.1.1.5 send-label
    neighbor 44.44.44.44 activate
    neighbor 44.44.44.44 route-reflector-client
    neighbor 44.44.44.44 next-hop-self
    neighbor 44.44.44.44 route-map SET-MPLS-LABEL out
    neighbor 44.44.44.44 send-label
  exit-address-family
```

ASBR4 Configuration:

```
router bgp 100
  bgp log-neighbor-changes
  neighbor 11.11.11.11 remote-as 100
  neighbor 11.11.11.11 update-source Loopback0
  neighbor 11.11.11.11 fall-over
  neighbor 21.1.1.2 remote-as 200
  neighbor 21.1.1.2 fall-over bfd
  neighbor 33.33.33.33 remote-as 100
  neighbor 33.33.33.33 update-source Loopback0
  neighbor 33.33.33.33 fall-over
  !
  address-family ipv4
    bgp advertise-best-external
    network 11.11.11.11 mask 255.255.255.255
    neighbor 11.11.11.11 activate
    neighbor 11.11.11.11 next-hop-self all
    neighbor 11.11.11.11 route-map SET-MPLS-LABEL out
    neighbor 11.11.11.11 send-label
    neighbor 21.1.1.2 activate
    neighbor 21.1.1.2 send-label
    neighbor 33.33.33.33 activate
    neighbor 33.33.33.33 next-hop-self
    neighbor 33.33.33.33 route-map LOCPREF-600 in
    neighbor 33.33.33.33 route-map SET-MPLS-LABEL out
```

```
neighbor 33.33.33.33 send-label
exit-address-family
```

Verification of Best External Path:

On PERR1, we can use the command **show ip cef <> internal** to verify that there are two paths to 66.66.66.66 (PERR6).

```
PERR1#show ip cef 66.66.66.66 internal
66.66.66.66/32, epoch 0, flags [rlbls], RIB[B], refcnt 6, per-destination sharing
sources: RIB, RR
feature space:
  IPRM: 0x00018000
  LFD: 66.66.66.66/32 0 local labels
    contains path extension list
    dflt label switch chain 0x7FBDE68C24D0
    loadinfo 7FBDE68C24D0, per-session, 2 choices, flags 0493, 5 locks
subblocks:
  1 RR source [non-eos indirection, heavily shared]
    non-eos chain loadinfo 7FBDE1787F60, per-session, flags 0591, 3 locks
ifnums:
  Ethernet0/1(3): 11.1.1.3
  Ethernet1/3(9): 23.1.1.4
```

Additional References

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Command List, All Releases
BGP commands	Cisco IOS IP Routing: BGP Command Reference
Basic MPLS VPNs	“Configuring MPLS Layer 3 VPNs” module in the <i>MPLS: Layer 3 VPNs Configuration Guide</i>
Multiprotocol VRFs	“MPLS VPN VRF CLI for IPv4 and IPv6 VPNs” module in the <i>MPLS: Layer 3 VPNs Configuration Guide</i>
A failover feature that creates a new path after a link or node failure	MPLS VPN--BGP Local Convergence

Standards

Standard	Title
draft-marques-idr-best-external	<i>BGP Best External, Advertisement of the best external route to iBGP</i>

MIBs

MIB	MIBs Link
—	To locate and download MIBs for selected platforms, Cisco software releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFC	Title
RFC 1771	<i>A Border Gateway Protocol 4 (BGP-4)</i>
RFC 2547	<i>BGP/MPLS VPNs</i>

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for BGP Best External

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1: Feature Information for BGP Best External

Feature Name	Releases	Feature Information
BGP Best External	Cisco IOS XE Release 3.2S	<p>The BGP Best External feature provides the network with a backup external route to avoid loss of connectivity of the primary external route. This feature advertises the most preferred route among those received from external neighbors as a backup route.</p> <p>In Cisco IOS XE Release 3.2S, this feature was introduced.</p> <p>The following commands were introduced or modified: bgp advertise-best-external, bgp recursion host, show ip bgp, show ip bgp vpnv4, show ip cef, show ip cef vrf, show ip route, show ip route vrf</p>
BGP Best External Path on an RR for Intercluster	Cisco IOS XE Release 3.4S	<p>The BGP Best External Path on RR for Intercluster feature provides path diversity between RR clusters. The feature provides best external functionality toward non-client iBGP peers, and is also known as "intercluster best external path."</p> <p>The following commands were introduced: bgp additional-pathsselect, neighbor advertise best-external.</p>
BGP Best External Path with MPLS Inter-AS Options B and C	Cisco IOS XE Amsterdam 17.3.1	<p>In MPLS VPN Inter-AS Option B and MPLS VPN Inter-AS Option C deployments, you can configure border routers to compute and advertise best external paths based on routes received from EBGP peers. The border routers advertise best external paths to Internal BGP (iBGP) peers as back-up paths. If a link in the best path fails, traffic flows along the best external path.</p> <p>This feature is introduced on Cisco ASR 1000 Series Aggregation Services Routers and the Cisco CSR 1000v Cloud Services Router.</p>

