



BGP Best External

Last Updated: April 13, 2012

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 internal BGP (iBGP) session (of another border router) as the best path for a prefix even if it has an external BGP (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 routers advertise one externally learned path called the best external path.

- [Finding Feature Information, page 1](#)
- [Contents, page 2](#)
- [Prerequisites for BGP Best External, page 2](#)
- [Restrictions for BGP Best External, page 2](#)
- [Information About BGP Best External, page 2](#)
- [How to Configure BGP Best External, page 5](#)
- [Configuration Examples for BGP Best External, page 9](#)
- [Additional References, page 10](#)
- [Feature Information for BGP Best External, page 12](#)

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest feature information and caveats, see the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the [Feature Information for BGP Best External, page 12](#).

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



Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Contents

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 Inter-Autonomous Systems, option B
 - 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 route reflectors.
- 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.

Information About BGP Best External

- [BGP Best External Overview, page 3](#)
- [What the Best External Route Means, page 3](#)
- [How the BGP Best External Feature Works, page 3](#)
- [Configuration Modes for Enabling BGP Best External, page 4](#)

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 iBGP session between them.
- Two routers in different autonomous systems of a confederation that have an 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*



[How the BGP Best External Feature Works, page 3](#) 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 [How the BGP Best External Feature Works, page 3](#), 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 VRFs 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 that VRF.

How to Configure BGP Best External

- [Enabling the BGP Best External Feature, page 5](#)
- [Verifying the BGP Best External Feature, page 7](#)

Enabling the BGP Best External Feature

Perform the following task to enable the BGP Best External feature. This task shows how to configure the BGP Best External feature in either IPv4 or VPNv4 address family. In VPNv4 address family configuration mode, the BGP Best External feature applies to all IPv4 VRFs; you do not have to 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 that VRF.

- Configure the MPLS VPN and verify that it is working properly before configuring the BGP Best External feature. See "Configuring MPLS Layer 3 VPNs" for more information.
- Configure multiprotocol VRFs, which 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 "MPLS VPN--VRF CLI for IPv4 and IPv6 VPNs".
- Ensure that the 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 vpnv4** [**unicast**]
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** [**b fd**|**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.

Command or Action	Purpose
Step 2 configure terminal Example: <pre>Router# configure terminal</pre>	Enters global configuration mode.
Step 3 router bgp <i>autonomous-system-number</i> Example: <pre>Router(config)# router bgp 40000</pre>	Enters router configuration mode for the specified routing process.
Step 4 Do one of the following: <ul style="list-style-type: none"> address-family ipv4 [unicast vrf <i>vrf-name</i>] or address-family vpnv4 [unicast] Example: <pre>Router(config-router)# address-family ipv4 unicast</pre> Example: <pre>Router(config-router)# address-family vpnv4</pre>	Specifies the IPv4 or VPNv4 address family and enters address family configuration mode. <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 bgp advertise-best-external Example: <pre>Router(config-router-af)# bgp advertise-best-external</pre>	Calculates and uses an external backup path and installs it into the RIB and Cisco Express Forwarding.
Step 6 neighbor <i>ip-address</i> remote-as <i>autonomous-system-number</i> Example: <pre>Router(config-router-af)# neighbor 192.168.1.1 remote-as 45000</pre>	Adds the IP address of the neighbor in the specified autonomous system to the IPv4 multiprotocol BGP neighbor table of the local router. <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.

Command or Action	Purpose
Step 7 neighbor <i>ip-address</i> activate Example: <pre>Router(config-router-af)# neighbor 192.168.1.1 activate</pre>	Enables the neighbor to exchange prefixes for the IPv4 unicast address family with the local router.
Step 8 neighbor <i>ip-address</i> fall-over [b fd route-map <i>map-name</i>] Example: <pre>Router(config-router-af)# neighbor 192.168.1.1 fall-over bfd</pre>	Configures the BGP peering to use fast session deactivation and enables BFD protocol support for failover. <ul style="list-style-type: none"> BGP will remove all routes learned through this peer if the session is deactivated.
Step 9 end Example: <pre>Router(config-router-af)# end</pre>	(Optional) Exits address family configuration mode and returns to privileged EXEC mode.

Verifying the BGP Best External Feature

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

SUMMARY STEPS

- enable
- show vrf detail
- show ip bgp ipv4 {mdt {all | rd | vrf} | multicast | tunnel| unicast} or show ip bgp vpnv4 {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]
- show bgp vpnv4 unicast vrf vrf-name ip-address
- show ip route vrf vrf-name repair-paths ip-address
- 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 vpnv4 {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          0      0 1 ?
*bx              10.0.0.1           0          0      0 1 ?
*                 10.0.0.1           0          0      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

```

Configuration Examples for BGP Best External

- [Example Configuring the BGP Best External Feature, page 10](#)

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
```

Additional References

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Commands List, All Releases
BGP commands: complete command syntax, command mode, command history, defaults, usage guidelines, and examples	<i>Cisco IOS IP Routing: BGP Command Reference</i>

Related Topic	Document Title
Basic MPLS VPNs	"Configuring MPLS Layer 3 VPNs"
Multiprotocol VRFs	"MPLS VPN--VRF CLI for IPv4 and IPv6 VPNs"
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	BGP Best External, Advertisement of the best external route to iBGP

MIBs

MIB	MIBs Link
None	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 2547	<i>BGP/MPLS VPNs</i>
RFC 1771	A Border Gateway Protocol 4 (BGP-4)

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	12.2(33)SRE 15.2(3)T	<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 Release 12.2(33)SRE, 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>

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

© 2012 Cisco Systems, Inc. All rights reserved.