



BGP Diverse Path Using a Diverse-Path Route Reflector

The BGP Diverse Path Using a Diverse-Path Route Reflector feature allows Border Gateway Protocol (BGP) to distribute an alternative path other than the best path between BGP speakers when route reflectors are deployed. This feature is meant to provide path diversity within an autonomous system (AS), within a single cluster only. That is, a route reflector is allowed to advertise the diverse path to its client peers only.

- [Prerequisites for BGP Diverse Path Using a Diverse-Path Route Reflector, on page 1](#)
- [Restrictions for BGP Diverse Path Using a Diverse-Path Route Reflector, on page 1](#)
- [Information About BGP Diverse Path Using a Diverse-Path Reflector, on page 2](#)
- [How to Configure a BGP Diverse-Path Route Reflector, on page 5](#)
- [Configuration Examples for BGP Diverse Path Using a Diverse-Path Route Reflector, on page 8](#)
- [Additional References, on page 10](#)
- [Feature Information for BGP Diverse Path Using a Diverse-Path Route Reflector, on page 11](#)

Prerequisites for BGP Diverse Path Using a Diverse-Path Route Reflector

You should understand the BGP Best External feature.

Restrictions for BGP Diverse Path Using a Diverse-Path Route Reflector


- A diverse path can be configured on a route reflector only.
- Only one shadow route reflector is allowed per existing route reflector, which will calculate one additional best path (the second best path). That is, only one additional plane (topology) is configured.
- Path diversity is configured within an AS, within a single route reflector cluster. That is, the route reflector will advertise the diverse path to its route reflector client peers only.
- Diverse path functionality is not supported on a route server.

Information About BGP Diverse Path Using a Diverse-Path Reflector

Limitation that a BGP Diverse Path Overcomes

As a path vector routing protocol, BGP-4 requires a router to advertise to its neighbors only the best path for a destination. However, multiple paths to the same destination would allow mechanisms that can improve resilience, quickly recover from failures, and load balance, for example.

The use of route reflectors is one of the main reasons for poor path diversity within an autonomous system (AS). In a network with route reflectors, even if a prefix is learned from multiple egress routers, the route reflector reflects only the best path to its clients. The figure below shows how deploying route reflectors might reduce path diversity in an AS, even when the BGP Best External feature is deployed.



In the figure above, P1 and P2 are diverse paths for prefix p. Assume Path 2 (P2) has a lower MED and higher local preference than P1. The BGP Best External feature on PE1 will make sure that P1 is propagated to the route reflectors, regardless of P2 having a lower MED and higher local preference. The route reflectors will have path diversity; they will learn both P1 and P2 with different exit points PE1 and PE2 (assuming that PE1 and PE2 have the **set ip next-hop self** command configured). However, both route reflectors select the best path as P2 due to its lower MED/higher local preference and advertise it to PE3. PE3 will not learn P1 (that is, PE3 will not learn about existing path diversity).

The BGP Diverse Path Using a Diverse-Path Route Reflector feature is a way to resolve that limitation and achieve path diversity.

BGP Diverse Path Using a Diverse-Path Route Reflector


The BGP Diverse Path Using a Diverse-Path Route Reflector feature overcomes the lack of path diversity in an AS containing route reflectors. This feature is meant to provide path diversity within an AS, within a single cluster only. That is, a route reflector is allowed to advertise the diverse path to its client peers only.

For each route reflector in the AS, a *shadow route reflector* is added to distribute the *second best path*, also known as the *diverse path*. The figure below shows the shadow route reflector for RR2. The shadow route reflector improves path diversity because PE3 can now learn both P1 (from RR1/RR2) and learn P2 from the shadow route reflector.




Note The primary route reflector and shadow route reflector must have the exact same connections (physical/control plane) to the rest of the routers in the network.

Shadow route reflectors can be both control plane route reflectors and data plane route reflectors.



The figure below shows a diverse path in greater detail, indicating the next hops:

- BR2 announces to RR1 and shadow RR2 that R2 (BR2) is the Next Hop for those who want to reach Prefix Z. Likewise, BR3 announces to RR1 and shadow RR2 that R3 (BR3) is the Next Hop for those who want to reach Prefix Z
- RR1 sends a packet to BR1 announcing that the Next Hop is R2 if BR1 wants to reach Prefix Z. The second best path (or diverse path) comes from shadow RR2, which sends a packet to BR1 announcing that the Next Hop is R3 if BR1 want to reach Prefix Z.
- At BR1 (far right), we see there are two (diverse) paths to Prefix Z.



Triggers to Compute a BGP Diverse Path

Computation of a diverse path per address family is triggered by any of the following commands:

- **bgp additional-paths install**
- **bgp additional-paths select**
- **maximum-paths ebgp**
- **maximum-paths ibgp**

The **bgp additional-paths install** command will install the type of path that is specified in the **bgp additional-paths select** command. If the **bgp additional-paths select** command specifies both keyword options (**best-external** and **backup**), the system will install a backup path.

The **maximum-paths ebgp** and **maximum-paths ibgp** commands trigger a multipath computation, and multipaths are automatically installed as primary paths.

On the other hand, the **bgp additional-paths install** command triggers computation of a backup path or best-external path.

If the **bgp additional-paths select** command is not configured, the **bgp additional-paths install** command will trigger both computation and installation of a backup path (as is done with the BGP PIC feature).

IGP Metric Check

Disabling the Interior Gateway Protocol (IGP) metric check and configuring the BGP Diverse Path feature are independent of each other. One does not imply the other. That is, configuring **bgp bestpath igp-metric ignore** does not imply that the BGP Diverse Path feature is enabled. Conversely, enabling the BGP Diverse Path feature might not require that **bgp bestpath igp-metric ignore** be configured (because, for example, the route reflector and shadow route reflector are co-located).

The **bgp bestpath igp-metric ignore** command can be configured at route reflectors and provider edges (PEs).



Note Per-VRF functionality for the **bgp bestpath igp-metric ignore** command is not supported. If you use it anyway, it is at your own risk.

Route Reflector Determination

If a router's configuration includes either one of the following commands, the router is a route reflector:

- **bgp cluster-id**
- **neighbor route-reflector-client**

How to Configure a BGP Diverse-Path Route Reflector

Determining Whether You Need to Disable the IGP Metric Check

Before you configure a shadow route reflector in order to get a BGP diverse path, determine whether you need to disable the IGP metric check. The IGP metric is a configurable value indicating physical distance, and is used by an Interior Gateway Protocol, such as Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), or Routing Information Protocol (RIP). A smaller IGP metric is preferred over a larger IGP metric.

The locations of the route reflector and shadow route reflector determine whether or not you need to disable the IGP metric check, as follows:

- When the route reflector and shadow route reflector are colocated—They have the same IP subnetwork address and are connected to the Ethernet switch with different links. Failure of such a link is equivalent to the route reflector going down. When RRs are colocated, their IGP metrics cannot be different from each other; and therefore there is no need to disable the IGP metric check during the best path calculation at any route reflector. Because there is no need to disable the IGP metric check, the first plane route reflectors do not need to be upgraded to Cisco IOS XE Release 3.4S.
- When the shadow route reflector is in a different IGP place from the route reflector (it is not colocated with its best path route reflector)--In this case, the IGP metric check is ignored on both the best path route reflector and shadow route reflector when the best path and second best path are being calculated. The IGP metric check must be disabled on the primary route reflector by configuring the **bgp bestpath igp-metric ignore** command. This command is available beginning with Cisco IOS XE Release 3.4S, which means you need to upgrade to that release.

Configuring the Route Reflector for BGP Diverse Path

Perform this task to configure a route reflector for the BGP Diverse Path feature. This task specifies the IPv4 address family, but other address families are also supported.

SUMMARY STEPS

1. **enable**

2. **configure terminal**
3. **router bgp** *autonomous-system-number*
4. **neighbor** *ip-address* **remote-as** *autonomous-system-number*
5. **address-family ipv4 unicast**
6. **neighbor** *ip-address* **activate**
7. **maximum-paths ibgp** *number-of-paths*
8. **bgp bestpath igp-metric ignore**
9. **bgp additional-paths select** [**backup**]
10. **bgp additional-paths install**
11. **neighbor** *ip-address* **route-reflector-client**
12. **neighbor** *ip-address* **advertise diverse-path** [**backup**] [**mpath**]
13. **end**
14. **show ip bgp neighbor** *ip-address* **advertised-routes**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	router bgp <i>autonomous-system-number</i> Example: Device(config)# router bgp 1	Enters router configuration mode for the BGP routing process.
Step 4	neighbor <i>ip-address</i> remote-as <i>autonomous-system-number</i> Example: Device(config-router)# neighbor 10.1.1.1 remote-as 1	Adds an entry to the BGP or multiprotocol BGP neighbor table.
Step 5	address-family ipv4 unicast Example: Device(config-router)# address-family ipv4 unicast	Specifies the address family and enters address family configuration mode. <ul style="list-style-type: none"> • Supported address families are IPv4 unicast, VPNv4 unicast, IPv6 unicast, VPNv6 unicast, IPv4+label, and IPv6+label.
Step 6	neighbor <i>ip-address</i> activate Example:	Enables the exchange of information with a BGP neighbor.

	Command or Action	Purpose
	Device(config-router-af)# neighbor 10.1.1.1 activate	
Step 7	maximum-paths ibgp <i>number-of-paths</i> Example: Device(config-router-af)# maximum-paths ibgp 4	Controls the maximum number of parallel Internal BGP (IBGP) routes that can be installed in a routing table.
Step 8	bgp bestpath igp-metric ignore Example: Device(config-router-af)# bgp bestpath igp-metric ignore	Configures the system to ignore the Interior Gateway Protocol (IGP) metric during BGP best path selection.
Step 9	bgp additional-paths select [backup] Example: Device(config-router-af)# bgp additional-paths select backup	Configures the system to calculate a second BGP best path.
Step 10	bgp additional-paths install Example: Device(config-router-af)# bgp additional-paths install	Enables BGP to calculate a backup path for a given address family and to install it into the routing information base (RIB) and Cisco Express Forwarding (CEF).
Step 11	neighbor ip-address route-reflector-client Example: Device(config-router-af)# neighbor 10.1.1.1 route-reflector-client	Configures the router as a BGP route reflector and configures the specified neighbor as its client.
Step 12	neighbor ip-address advertise diverse-path [backup] [mpath] Example: Device(config-router-af)# neighbor 10.1.1.1 advertise diverse-path backup	(Optional) Configures a neighbor to receive the diverse path in an advertisement.
Step 13	end Example: Device(config-router-af)# end	(Optional) Exits address family configuration mode and returns to privileged EXEC mode.
Step 14	show ip bgp neighbor ip-address advertised-routes Example: Device# show ip bgp neighbor 10.1.1.1 advertised-routes	(Optional) Displays the routes advertised to the specified neighbor.


Configuration Examples for BGP Diverse Path Using a Diverse-Path Route Reflector

Example: Configuring BGP Diverse Path Where Additional Path Is the Backup Path

Diverse path functionality is contained within a single cluster; that is, only the clients of a route reflector can be configured to advertise the diverse path. A diverse path is advertised to the clients of a route reflector only if the client is configured to get the additional path.

A shadow route reflector can be added to calculate and advertise the additional path, or an existing route reflector can be configured to calculate and advertise the additional path. In the figure below, instead of adding a shadow route reflector, RR2 (the existing backup RR) is configured to calculate the additional path and advertise it to a particular neighbor.

In the figure below, assume that from the route reflectors, the path to CE1 via PE1 is preferred over the path via PE2. Without the diverse path feature, both route reflectors will advertise to PE3 that the path to CE1 is via PE1. If the connection between RR1 and PE1 fails (or the path between PE1 and CE1 fails), there is no other path.



In the following configuration example based on the figure above, RR2 is configured with an additional path, which is a backup path.

If RR1 and RR2 are not colocated, you must configure the **bgp bestpath igp-metric ignore** command before the additional path is calculated. (If RR1 and RR2 are colocated, do not configure that command.)

The **bgp additional-paths select backup** command triggers calculation of the backup path at RR2, which is the path via PE2.

The **bgp additional-paths install** command installs the backup path if RR2 is in the forwarding plane. (Do not configure this command if RR2 is in the control plane.)

The address of PE3 is 10.1.1.1, and that address is used in the **neighbor advertise diverse-path backup** command on RR2. This command triggers advertisement of the backup path to PE3. PE3 will learn the best path, (which is the path via PE1) from RR1, and it will learn the backup path from RR2.

RR2

```
router bgp 1
 neighbor 10.1.1.1 remote-as 1
 address-family ipv4 unicast
 neighbor 10.1.1.1 activate
 maximum-paths ibgp 4
 bgp bestpath igp-metric ignore
 bgp additional-paths select backup
 bgp additional-paths install
 neighbor 10.1.1.1 route-reflector-client
 neighbor 10.1.1.1 advertise diverse-path backup
```

Example: Configuring BGP Diverse Path Where Additional Path Is the Multipath

In the following example based on the figure above, assume that paths toward CE1 via PE1 and PE2 are multipaths. The **maximum-paths ibgp** command will trigger calculation of multipaths.

The address of PE3 is 10.1.1.1, and that address is used in the **neighbor advertise diverse-path mpath** command on RR2. This command will trigger advertisement of the multipath, that is, the second best path, to PE3. PE3 will learn the best path, path via PE1 from RR1, and will learn second best path from RR2.

RR2

```
router bgp 1
 neighbor 10.1.1.1 remote-as 1
 address-family ipv4 unicast
 neighbor 10.1.1.1 activate
 maximum-paths ibgp 4
 neighbor 10.1.1.1 remote-as 1
 neighbor 10.1.1.1 route-reflector-client
 neighbor 10.1.1.1 advertise diverse-path mpath
```

Example: Configuring BGP Diverse Path Where Both Multipath and Backup Path Calculations Are Triggered

The following example is based on the figure above. The **maximum-paths ibgp** command will trigger calculation of multipaths. When both multipath and backup path calculations are triggered, the backup path and the second multipath (which is the second best path) are the same paths and it will be installed as the active path, regardless of whether the route reflector is in the control plane or forwarding plane.

The address of PE3 is 10.1.1.1, and that address is used in the **neighbor advertise diverse-path backup mpath** command on RR2. This command causes RR2 to advertise the second best path, which is the second multipath, to PE3.

RR2

```
router bgp 1
 neighbor 10.1.1.1 remote-as 1
 address-family ipv4 unicast
 neighbor 10.1.1.1 activate
 maximum-paths ibgp 4
 bgp additional-paths select backup
 neighbor 10.1.1.1 remote-as 1
```

Example: Configuring Triggering Computation and Installation of a Backup Path

```
neighbor 10.1.1.1 route-reflector-client
neighbor 10.1.1.1 advertise diverse-path backup mpath
```

Example: Configuring Triggering Computation and Installation of a Backup Path

When the **bgp additional-paths install** command is configured without configuring **bgp additional-paths select backup**, the former command will trigger both computation and installation of the backup path (as it is with the existing BGP PIC feature).

The address of PE3 is 10.1.1.1, and that address is used in the **neighbor advertise diverse-path backup** command on RR2. This command will trigger advertisement of a backup path to PE3. PE3 will learn the best path, a path via PE1 from RR1, and it will learn a backup path from RR2.

RR2

```
router bgp 1
neighbor 10.1.1.1 remote-as 1
address-family ipv4 unicast
neighbor 10.1.1.1 activate
maximum-paths ibgp 4
bgp additional-paths install
neighbor 10.1.1.1 remote-as 1
neighbor 10.1.1.1 route-reflector-client
neighbor 10.1.1.1 advertise diverse-path backup
```

Additional References

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Commands List, All Releases
BGP commands	Cisco IOS IP Routing: BGP Command Reference
Configuring BGP Best External Path on a Route Reflector for Intercluster	BGP Best External module
BGP configuration tasks	Cisco IOS XE IP Routing: BGP Configuration Guide

Standards

Standard	Title
draft-ietf-grow-diverse-bgp-path-dist-02.txt	<i>Distribution of Diverse BGP Paths</i>

MIBs

MIB	MIBs Link
No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified by this feature.	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 4271	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 Diverse Path Using a Diverse-Path Route Reflector

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 Diverse Path Using a Diverse-Path Route Reflector

Feature Name	Releases	Feature Information
BGP Diverse Path Using a Diverse-Path Route Reflector		<p>This feature allows BGP to distribute an alternative path other than the best path between BGP speakers when route reflectors are deployed.</p> <p>The following commands were introduced:</p> <ul style="list-style-type: none">• bgp additional-paths select• bgp bestpath igp-metric ignore• debug ip bgp igp-metric ignore• neighbor advertise best-external• neighbor advertise diverse-path