



Configuring BGP Additional Paths

This chapter describes how to configure the BGP Additional Paths allowing the advertisement of multiple paths through the same peering session for the same prefix without the new paths implicitly replacing any previous paths. This behavior promotes path diversity and reduces multi-exit discriminator (MED) oscillations.

This chapter includes the following sections:

- [Information About BGP Additional Paths, on page 1](#)
- [How to Configure BGP Additional Paths, on page 4](#)
- [Verifying the BGP Additional Paths Configuration, on page 9](#)
- [Feature History for BGP Additional Paths, on page 9](#)

Information About BGP Additional Paths

This section includes the following topics:

Problem That Additional Paths Can Solve

BGP routers and route reflectors (RRs) propagate only their best path over their sessions. The advertisement of a prefix replaces the previous announcement of that prefix (this behavior is known as an implicit withdraw). The implicit withdraw can achieve better scaling, but at the cost of path diversity.


Path hiding can prevent efficient use of BGP multipath, prevent hitless planned maintenance, and can lead to MED oscillations and suboptimal hot-potato routing. Upon nexthop failures, path hiding also inhibits fast and local recovery because the network has to wait for BGP control plane convergence to restore traffic. The BGP Additional Paths feature provides a generic way of offering path diversity; the Best External or Best Internal features offer path diversity only in limited scenarios.

The BGP Additional Paths feature provides a way for multiple paths for the same prefix to be advertised without the new paths implicitly replacing the previous paths. Thus, path diversity is achieved instead of path hiding.

Path-Hiding Scenario

This section describes in more detail how path hiding can occur. In the following figure, we have prefix p with paths p1 and p2 advertised from BR1 and BR4 to RR1. RR1 selects the bestpath of the two and then advertises to PE only p1.

Figure 1: RR Hiding an Additional Path




In the figure **RR Hiding an Additional Path**, we also see prefix x with path x1 being advertised from BR2 to BR3 (which has path x2) with local preference 100. BR3 also has path x2, but due to routing policy, BR3 will advertise to the RRs x1 (not shown) instead of x2, and x2 will be suppressed. A user could turn on the advertisement of best external on BR3 and thereby advertise x2 to the RRs, but, again, the RRs advertise only the best path.

Suboptimal Hot-Potato Routing Scenario

In order to minimize internal transport costs, transit ISPs try to forward packets to the closest exit point (according to IGP cost). This behavior is known as hot-potato routing. In the distributed RR cluster model of the following figure, assume traffic coming from LA must go to Mexico. All links have the same IGP cost. If there are two exit points toward Mexico—one toward Austin and one toward Atlanta—the border router will try to send traffic to Austin based on the lower IGP cost from LA toward Austin than toward Atlanta. In a centralized RR model where the central RR resides where RR3 is (and RR1, RR2, RR4, and RR5 do not exist), the closest exit point toward Mexico, as seen from RR3, might be Atlanta. Sending the traffic from LA toward the Atlanta BR and thus resulting in suboptimal hot-potato routing is not desirable.

Figure 2: Distributed RR Cluster



344259

Benefits of Additional BGP Paths

BGP routers and route reflectors (RR) propagate only their best path over their sessions. The advertisement of a prefix replaces the previous announcement of that prefix (this is known as an implicit withdraw).

While this behavior may achieve better scaling, it can prevent path diversity, which tends to be poor or completely lost. The behavior in turn prevents efficient use of BGP multipath, prevents hitless planned maintenance, and can lead to multi-exit discriminator (MED) oscillations and suboptimal hot-potato routing. It also inhibits fast and local recovery upon nexthop failures, because the network has to wait for BGP control plane convergence to restore traffic.

The BGP Additional Paths feature is a BGP extension that allows the advertisement of multiple paths for the same prefix without the new paths implicitly replacing any previous paths. This behavior promotes path diversity and reduces MED oscillations.

BGP Additional Paths Functionality

The BGP Additional Paths feature is implemented by adding a path identifier to each path in the NLRI. The path identifier (ID) can be considered as something similar to a route distinguisher (RD) in VPNs, except that a path ID can apply to any address family. Path IDs are unique to a peering session and are generated for each network. The path identifier is used to prevent a route announcement from implicitly withdrawing the previous one. The Additional Paths feature allows the advertisement of more paths, in addition to the bestpath. The Additional Paths feature allows the advertisement of multiple paths for the same prefix, without the new paths implicitly replacing any previous paths.

The BGP Additional Paths feature requires you to take three general steps:

1. Specify whether the device can send, receive, or send and receive additional paths. This is done at the address family level or the neighbor level. During session establishment, two BGP neighbors negotiate the additional path capabilities (whether they can send and/or receive) between them.
2. Select a set or sets of candidate paths for advertisement by specifying selection criteria.
3. Advertise for a neighbor a set or sets of additional paths from the candidate paths marked.

To send or receive additional paths, the additional path capability must be negotiated. If it is not negotiated, even if the selection criteria are such that more than the bestpath is marked and the neighbor is configured to advertise the marked paths, the selections would be useless because without the capability negotiated, only the bestpath can be sent.

Configuring BGP to send or receive additional paths triggers negotiation of additional path capability with the device's peers. Neighbors that have negotiated the capability will be grouped together in an update group (if other update group policies allow), and in a separate update group from those peers that have not negotiated the capability. Therefore, additional path capability causes the neighbor's update group membership to be recalculated.

Additional Path Selection

Only the best-path is advertised to peers unless you configure the **set path-selection all advertise** command which advertises all BGP paths as additional paths to peers if the receive capability is enabled.

Advertise a Subset of the Paths Selected

Take care when you select a set of paths but want to advertise a different set of paths. If the set of paths you want to advertise is not a subset of the selected paths, then you will not advertise the paths you want advertised.

Guidelines and Limitations

Configuring BGP Additional Paths has the following guidelines and limitations:

- BGP add-path is not supported as dynamic capability. It is included in OPEN but not in CAPABILITY message. The configuration will take effect upon the next session establishment and does not cause established sessions to get torn down.

How to Configure BGP Additional Paths

This section includes the following topics:

Configuring BGP Additional Paths per Address Family

Perform the following steps to specify whether the device can send and receive additional paths to and from all neighbors within an address family.

Before you begin

Ensure that you have enabled the BGP feature.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters configuration mode.
Step 2	router bgp <i>as-number</i> Example: switch(config)# router bgp 65000 switch(config-router)#	Enables BGP and assigns the autonomous system number to the local BGP speaker.
Step 3	address-family ipv4 unicast Example: switch(config-router)# address family ipv4 unicast	Enters address family configuration mode.
Step 4	(Optional) additional-paths receive Example: switch(config-router-af)# additional-paths receive	Enables BGP additional paths for a prefix to be received from a capable peer. Note This capability applies to all neighbors under the specified address family unless the capability is explicitly disabled with the neighbor additional-paths receive disable command, which overrides the configuration for the address family.
Step 5	(Optional) additional-paths send Example: switch(config-router-af)# additional-paths send	Enables BGP additional paths for a prefix to be sent to a capable peer. Note This capability applies to all neighbors under the specified address family unless the capability is explicitly disabled with the neighbor additional-paths send disable command, which overrides the configuration for the address family.
Step 6	(Optional) additional-paths selection route-map Example: switch(config-router-stmp)# exit switch(config-router)#	Configures additional paths selection capability for a prefix.
Step 7	(Optional) end Example:	Exits to privileged EXEC mode.

	Command or Action	Purpose
	<code>switch(config-router-af)# end</code>	

Configuring BGP Additional Paths per Neighbor

Perform the following steps to configure whether a particular neighbor can send or receive additional paths.

Before you begin

Ensure that you have enabled the BGP feature (see the [Enabling the BGP Feature](#) section).

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <code>switch# configure terminal</code> <code>switch(config)#</code>	Enters configuration mode.
Step 2	router bgp <i>as-number</i> Example: <code>switch(config)# router bgp 65000</code> <code>switch(config-router)#</code>	Enables BGP and assigns the autonomous system number to the local BGP speaker.
Step 3	neighbor { <i>ipv4-address</i> <i>ipv4-prefix/length</i> } [remote-as { <i>as-num</i> } [. <i>as-num</i>]]	Configures a BGP neighbor (router, VRF) and enters the neighbor configuration mode.
Step 4	address-family ipv4 unicast Example: <code>switch(config-router)# address family ipv4 unicast</code>	Enters address family configuration mode.
Step 5	(Optional) capability additional-paths receive [disable] Example: <code>switch(config-router-af)# capability additional-paths receive</code>	Configures the receive additional paths capability for the specified neighbor. Note This command overrides any send or receive capability that is configured at the address-family level.
Step 6	(Optional) capability additional-paths send [disable] Example: <code>switch(config-router-af)# capability additional-paths send</code>	Configures the send additional paths capability for the specified neighbor. Note This command overrides any send or receive capability that is configured at the address-family level.

	Command or Action	Purpose
Step 7	(Optional) end Example: switch(config-router-af)# end	Exits to privileged EXEC mode.

Configuring Additional Paths Using a Peer Policy Template

In this configuration task, the capability to send and receive additional paths and the selection criteria are configured for the address family, and then the template is configured.

Before you begin

Ensure that you have enabled the BGP feature (see the [Enabling the BGP Feature](#) section).

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters configuration mode.
Step 2	router bgp <i>as-number</i> Example: switch(config)# router bgp 65000 switch(config-router)#	Enables BGP and assigns the autonomous system number to the local BGP speaker.
Step 3	template peer-policy <i>template-name</i> Example: switch(config-router)# template peer-policy rr-client-ptl #	Enters policy-template configuration mode and creates a peer policy template.
Step 4	(Optional) capability additional-paths receive [disable] Example: switch(config-router-af)# capability additional-paths receive	Configures the receive additional paths capability for the specified neighbor. Note This command overrides any send or receive capability that is configured at the address-family level.
Step 5	(Optional) capability additional-paths send [disable] Example: switch(config-router-af)# capability additional-paths send	Configures the send additional paths capability for the specified neighbor. Note This command overrides any send or receive capability that is configured at the address-family level.

	Command or Action	Purpose
Step 6	exit Example: switch(config-router-ptmp)# exit	Exits policy-template configuration mode and returns to router configuration mode.
Step 7	neighbor { <i>ipv4-address</i> <i>ipv4-prefix/length</i> } [remote-as { <i>as-num</i> } [<i>as-num</i>]]	Configures a BGP neighbor (router, VRF) and enters the neighbor configuration mode.
Step 8	(Optional) address-family ipv4 unicast Example: switch(config-router)# address family ipv4 unicast	Enters address family configuration mode.
Step 9	inherit peer-policy <i>template-name</i> <i>sequence-number</i> Example: switch(config-router-neighbor-af)# inherit peer-policy rr-client-ptl 10	Sends a peer policy template to a neighbor so that the neighbor can inherit the configuration.
Step 10	(Optional) end Example: switch(config-router-af)# end	Exits to privileged EXEC mode.

Filtering and Setting Actions for Additional Paths

You can optionally use a route map to filter the paths to be advertised by matching on the prefix of additional paths that are candidates to be advertised. (These prefixes are configured with the **additional-paths selection** command.)

You can also optionally set one or more actions to take for those paths that pass through the route map. This task happens to use the **set metric** command. Other **set** commands are available that are not shown in this task.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters configuration mode.
Step 2	route-map <i>map-name</i> [deny permit] [<i>sequence-number</i>] Example: switch(config)# route-map add_path4 permit 10	Defines a route map and the conditions for redistributing routes from one routing protocol into another.

	Command or Action	Purpose
Step 3	set path-selection all advertise Example: <pre>switch(config-route-map) # set path-selection all advertise</pre>	Advertises all BGP paths as additional paths to peers if the receive capability is enabled.
Step 4	set metric <i>metric-value</i> Example: <pre>switch(config-route-map) # set metric 500</pre>	Sets the metric of the additional paths that pass the match criteria. <ul style="list-style-type: none"> Note that other set commands can be used to take action on the paths that pass the route map.
Step 5	(Optional) end Example: <pre>switch(config-router-af) # end</pre>	Exits to privileged EXEC mode.

Verifying the BGP Additional Paths Configuration

To display information about the BGP additional paths configuration, use the following commands:

Command	Purpose
show ip bgp [<i>ip-address</i>]	Displays entries in the BGP table.
show ip bgp neighbors [<i>ip-address</i> [advertise-routes]]	Displays the configured neighbors and the other information specific to individual neighbor.

Feature History for BGP Additional Paths

Table below lists the release history for this feature.

Table 1: Feature History for BGP

Feature Name	Releases	Feature Information
BGP Additional Paths	6.0(2)U1(1)	<p>The BGP Additional Paths feature allows the advertisement of multiple paths for the same prefix without the new paths implicitly replacing any previous paths.</p> <p>The following commands were introduced:</p> <ul style="list-style-type: none"> • additional-paths receive • additional-paths selection • additional-paths send • capability additional-paths receive • capability additional-paths send • set path-selection all advertise <p>The following commands were modified:</p> <ul style="list-style-type: none"> • show ip bgp • show ip bgp neighbors