



BGP Path Selection and Route Preference

BGP path selection is critical for efficient and policy-driven network routing. This chapter explains how BGP determines the best paths across complex networks and describes the configuration options available to influence route preference.

- [BGP best path algorithm, on page 1](#)
- [How BGP path selection works, on page 1](#)
- [Order of comparisons and nontransitivity, on page 3](#)
- [Suppressing best path changes, on page 3](#)
- [Tuning BGP path selection, on page 4](#)
- [Advanced path selection features, on page 15](#)

BGP best path algorithm

A BGP best path algorithm is a routing decision process that

- systematically evaluates multiple available BGP routes to the same destination based on a sequence of well-defined path attributes,
- selects the optimal route for installation in the routing table and advertisement to peers, and
- allows for administrative policy and configuration to influence path selection at various steps.

How BGP path selection works

BGP path selection occurs whenever multiple valid routes exist to the same destination. The router applies a series of steps, taking into account both protocol defaults and administrator-configured attributes to determine which path is installed in the routing table.

Summary

The key components involved in the process are:

- **BGP router:** A router that evaluates and compares multiple available routes using protocol-defined criteria.
- **Configured attributes:** Attributes that include weight, local preference, AS path, origin, MED, and community values, which influence path selection at each stage.

- BGP neighbors: External (eBGP) or internal (iBGP/ confederation), whose routes are candidates for selection.

When a router receives multiple BGP routes for the same destination, it selects the best path by first discarding invalid paths and preferring those with lower pre-bestpath cost communities if configured. It then compares attributes like weight, local preference, and whether the path is locally originated. Next, it evaluates AS path length, origin type, and MED among paths from the same AS. Advanced criteria favor eBGP-learned paths, lower IGP metrics to the next hop, and lower IP cost communities. If still tied, it uses tie-breakers such as lowest router ID, shortest cluster list, and lowest neighbor IP address. If all attributes tie, the current best path is retained until a change triggers re-selection.

Workflow

These stages describe the BGP path selection process.

1. The router receives multiple BGP route advertisements for the same destination from different peers.
2. The router performs preliminary checks:
 - removes invalid paths, such as those with unreachable next hops or exceeded MED values.
 - prefers the path with the lowest pre-bestpath cost community value.
3. The router compares attribute-based preferences and selects as follows:
 - Weight: chooses the path with the highest weight.
 - Local preference: prefers the path with the highest local preference value.
 - Locally originated path or aggregate: prefers any paths that are locally originated or created with the network or aggregate-address commands.
4. The router evaluates path characteristics:
 - AS path length: selects the path with the shortest AS path (unless the relevant configuration is set to ignore AS path).
 - Origin type: chooses the path with the lowest origin type (IGP over EGP over Incomplete).
 - Multi-Exit Discriminator (MED): selects the one with the lowest MED among paths from the same neighboring AS.
5. The router applies advanced criteria:
 - eBGP over iBGP: prefers a path learned via eBGP over those learned via iBGP or confederation peers.
 - IGP metric to next hop: chooses the path with the lowest IGP metric to the BGP next hop.
 - IP cost community: if present, selects the path with the lowest IP cost community value.
6. The router uses final tie-breakers when paths are still equal:
 - Router ID: selects the path received from the peer with the lowest router ID (using originator ID if available).
 - Cluster list length: prefers the path with the shortest cluster list.

- Neighbor IP address: As a last resort, chooses the path from the peer with the lowest neighbor IP address (locally generated paths are considered to have IP address 0).
7. If all these attributes result in a tie, the router maintains the current best path until a configuration change forces re-selection.

Example of path comparison order

Suppose a router receives three routes (A, B, C) to the same prefix. The router compares routes in sequence, stopping at the first attribute where one route is superior. The first attribute to break the tie determines the best path.

Order of comparisons and nontransitivity

The order in which BGP compares route attributes impacts routing decisions, especially because some attributes like MED are nontransitive and require special handling. These key points summarize this behavior:

- **MED nontransitivity:**

- BGP compares the MED attribute only among routes received from the same neighboring autonomous system (AS). If route A is better than B, and B is better than C, route A is not necessarily better than C unless all three routes are from the same AS.

- **Grouping for MED comparison:**

- Before evaluating the MED attribute, BGP groups routes by their neighboring AS. Within each group, BGP selects the route with the lowest MED value as the group's best route.

- **Iterative best path selection:**

- Once a best-path is chosen from each group, BGP resumes the standard best-path selection process among these group winners to determine the overall best route.

Suppressing best path changes

BGP includes suppression logic to reduce unnecessary routing changes (churn) and enhance network stability.

Suppression logic:

Suppression logic operates under these conditions:

- If the new and current best paths are identical up to the decision point where tie-breakers (such as router ID) apply, BGP may retain the current best path to prevent route instability.
- Suppression generally applies only when both routes are learned from external peers and their routing attributes are otherwise equivalent.

Configuration note:

- Suppression can be disabled with the **bgp bestpath compare-routerid** command. When this command is used, any change in route attributes triggers a new best path selection.

Tuning BGP path selection

While BGP's default algorithm provides a robust method for selecting the best path, you must influence this decision to meet business or technical objectives. This section covers the major attributes and configuration tools available for tuning path selection, including the cost community, local preference, weight, MED, and the AIGP attribute.

BGP cost communities

A BGP cost community is a nontransitive extended community attribute that

- passes to internal BGP (iBGP) and confederation peers but not to external BGP (eBGP) peers,
- enables customization of local route preference by assigning cost values to specific routes, and
- influences the BGP best-path selection process at defined points in the algorithm.

The cost community feature allows network operators to prefer certain paths over others by assigning lower cost values. By default, the cost community affects the best-path algorithm after the IGP metric comparison. If multiple paths have equal preference, the cost community serves as a tie-breaker.

Points of insertion

The extended community format introduces generic points of insertion (POI) within the best-path selection process, allowing granular control over route preference. This mechanism gives administrators the flexibility to adjust the decision logic at key stages for more sophisticated routing policies.

How BGP cost communities affect best path selection

Cost communities can be configured to act as tiebreakers when several paths have equal preference by traditional BGP metrics. You can compare these attributes using a deterministic and hierarchical process.

Summary

The key components involved in the process are:

- **BGP router:** A router that analyzes multiple received paths and applies the cost community attribute during best-path computation.
- **Cost community attribute:** A path attribute that is used for fine-tuning the selection among equal-cost paths.
- **Point of Insertion (POI):** The decision stage where cost community is evaluated before other standard metrics.

After IGP metric evaluation, the BGP router uses cost communities to compare paths. It sorts and compares attributes by point of insertion, community ID, and cost, preferring lower values. The router selects the path that ranks highest as the best path in the routing table.

Workflow

These stages describe how BGP cost communities affect best path selection:

1. Path reception: The BGP router receives multiple routes to the same destination.
2. POI evaluation: The router applies the point of insertion (POI), where cost community attributes take effect, after IGP metric comparisons.
3. Cost community comparison:

When multiple paths have the cost community attribute, the BGP router compares them as follows:

- Sort and prepare:
 - For each path, the router sorts cost community attributes by Point of Insertion (POI) and then by Community ID.
 - If a path does not include a cost community for the given POI and community ID, the router assigns it the default cost value (2147483647).
- Pairwise comparison:
 - The router compares two paths at a time, examining the cost community attributes for the POI and community ID in order.
 - At each step:
 - If neither path has a relevant cost community, the router declares a tie for this comparison.
 - If only one path has a cost community, the router selects that path as the best path.
 - If both paths have cost communities with different community IDs, the path with the lower Community ID is chosen.
 - If both have the same Community ID but different cost values, the path with the lower cost value is selected.
 - If both have the same Community ID and cost value, the router continues to the next lowest Community ID and repeats the process.
- The comparison ends when a best path is determined or all comparisons result in a tie.

Result

The BGP router installs the path with the most favorable cost community, so administrators can control routing decisions when standard BGP metrics do not produce a unique result.

Aggregate routes and multipaths

Aggregate routes and multipaths are extensions of the BGP cost community feature that

- allow the cost community attribute to be applied to aggregate and multipath routes
- ensure that unique cost community IDs from component routes are passed to the aggregate or multipath route, and
- apply only the highest cost value from all component routes for each unique ID to the aggregate or multipath route.

Inheritance of cost community attributes in aggregate and multipath routes

Aggregate or multipath routes inherit the BGP cost community attribute from the individual component routes. This means that any cost community attribute present in a component route is transferred to the aggregate or multipath route as part of the aggregation process.

Selection of cost values for each community ID

For each unique cost community ID found among the component routes, the aggregate or multipath route uses only the highest cost value for that ID. This ensures that, when multiple component routes share the same ID but have different cost values, the most significant (highest) value is selected and advertised.

Handling missing or mismatched cost community attributes

If a component route does not contain the cost community attribute, or if the component routes possess different cost community IDs, the default value of 2147483647 is advertised for those IDs.

Aggregate route with the same cost community ID

- Route 10.0.0.1 has POI=IGP, cost community ID=1, and cost number=100.
- Route 192.168.0.1 has POI=IGP, cost community ID=1, and cost number=200.

After aggregation, the aggregate route advertises cost community ID=1 with a cost of 200, which is the highest value among the component routes.

Aggregate route with different cost community IDs

- Route 10.0.0.1 has POI=IGP, cost community ID=1, and cost number=100
- Route 172.16.0.1 has POI=IGP, cost community ID=2, and cost number=100
- Route 192.168.0.1 has POI=IGP, cost community ID=1, and cost number=200

In this scenario, the aggregate route advertises the default cost value 2147483647 for both ID=1 and ID=2, since multiple IDs are present and cannot be consolidated under a single value.

If any component route does not include the cost community attribute, the aggregate or multipath route advertises the default cost value (2147483647) for that community ID.

Example: Adding routes to the Routing Information Base

The Routing Information Base (RIB) uses specific criteria for adding routes and selecting the best path. These criteria ensure the most optimal route is chosen based on BGP's cost community and best-path algorithms.

The RIB uses these criteria when adding routes and selecting the best path:

- If a nonsourced path becomes the best path after a best-path calculation, BGP adds the route to the RIB and passes cost communities along with other IGP extended communities.
- When a protocol adds a route with multiple paths to the RIB, the RIB checks both current best paths and incoming paths for the presence of cost extended communities.
- If cost extended communities are present:

- The RIB compares the sets of cost communities between paths.
- If there is a tie in cost community comparison, the RIB continues with the remaining steps of the best-path algorithm.
- If a cost community is missing from either current best paths or newly added paths, the RIB also continues with the remaining best-path evaluation steps.

Table 1: Summary:

Condition	RIB Action
Nonsourced path is best after calculation	Adds route to RIB with cost or IGP communities
Cost community found in compared paths	Compares cost communities, selects best
Tie in cost community comparison	Runs next steps of best-path algorithm
Missing cost community in any compared path	Runs next steps of best-path algorithm

Example: Influencing route preference in a multiexit IGP network

In large enterprise and service provider networks, it is common to have multiple exit points (border routers) connecting an internal network to one or more external networks or service providers. These are known as multiexit points. If interior gateway protocols (IGPs) such as OSPF or IS-IS provide equal-cost paths to these exit points, BGP treats all exit points equally.

Default behavior:

- If **multipath load sharing** is configured in BGP, both equal-cost paths are installed in the routing table and used for balancing traffic.
- If **multipath load balancing** is not configured, BGP selects the path that was learned first and installs only this path in the routing table. You might see suboptimal routing if, for example, the first-learned path uses a lower-speed link.


Using cost community to influence preference:

The BGP cost community attribute allows you to customize route preference within the AS, even when IGP metrics are equal. By assigning a lower cost community value to a preferred route, you can ensure that BGP selects that route over others.

Example:

This figure shows an IGP network with two autonomous system boundary routers (ASBRs) on the edge. Each ASBR has an equal cost path to network 10.8/16

Figure 1: Multiexit point IGP network



Suppose you want BGP to prefer the path learned from ASBR2 over ASBR1. You can apply a lower cost community value to the route received via ASBR2. For example:

```
Router(config)# route-policy ISP2_PE1
Router(config-rpl)# set extcommunity cost (1:1)
```

This configuration assigns a cost community number of 1 to the 10.8.0.0 route learned from ASBR2. By default, BGP sets the cost community value of the path learned from ASBR1 to 2147483647. Because 1 is lower than 2147483647, the path via ASBR2 is selected as the best path by BGP.

Configure a BGP cost community

Configure BGP with the cost community feature to enable tiebreaking among equal paths during the best-path selection process.

Before you begin

- Ensure you are in privileged EXEC mode on the router.
- Gather the desired route policy name and cost community information.

Follow these steps to configure a BGP cost community:

Procedure

Step 1 Enter route policy configuration mode, and set the BGP extended community attribute for cost in the route-policy.

Example:

```
Router# configure
Router(config)# route-policy costA
Router(config)# set extcommunity cost cost_A
```

Step 2 Configure BGP and apply the route policy to a BGP neighbor.

- Use the **default-information originate** command to originate default-information.
- Use the **aggregate-address [address | mask-length] route-policy route-policy-name** to configure an aggregate-address.

Example:


```
Router(config)# router bgp 120
Router(config-bgp)# default-information originate
Router(config-bgp)# neighbor 172.168.40.24
Router(config-bgp-nbr)# address-family ipv4 unicast
Router(config-bgp-nbr-af)# route-policy costA in
```

Step 3 Use the **showbgp ip-address** command to verify the configuration.

Ensure that the cost community is displayed in this format:

Example:

```
Cost: POI : cost-community-ID : cost-number
```

Local preference and weight

Local preference and weight are key BGP attributes used to influence route selection within an AS, allowing for straightforward prioritization of certain paths.

Change the BGP default local preference value

Adjust the default local preference value for BGP paths to influence route selection.

Use this task when you need to make certain BGP paths more or less preferred in your network routing policy by changing the default preference value.

Before you begin

- Ensure you have administrative access to the router.

Follow this step to change the BGP default local preference value:

Procedure

Enter BGP router configuration mode for your autonomous system, and configure the default local preference value.

Example:

```
Router# configure
Router(config)# router bgp 120
Router(config-bgp)# bgp default local-preference 200
```

This value overrides the default (100) for all BGP paths learned unless set otherwise.

Configure BGP weights

Assign a weight to routes received from a BGP neighbor to influence the best-path selection process.

BGP weights allow you to prefer specific neighbors for routing most of your traffic. By assigning a higher weight to routes learned from a particular neighbor, you control which path is chosen as the best for outgoing traffic.

Before you begin

- Determine the desired autonomous system (AS) number for your BGP router.
- Identify the IP address and remote AS number of the BGP neighbor whose routes you want to modify.

Follow this step to configure BGP weights:

Procedure

Configure the neighbor and assign a weight to all routes learned from the neighbor.

Example:

```
router# configure
Router(config)# router bgp 120
Router(config-bgp)# neighbor 172.168.40.24
Router(config-bgp-nbr)# remote-as 2002
Router(config-bgp-nbr)# address-family ipv4 unicast
Router(config-bgp-nbr-af)# weight 41150
```

Multi-Exit Discriminator

The Multi Exit Discriminator (MED) is an attribute used in BGP to influence the path selection process when multiple exit points exist from one autonomous system (AS) to another. Understanding its comparison logic and configuration is important for controlling external routing behavior.

Conditions for BGP MED comparison

BGP compares the Multi-Exit Discriminator (MED) attribute under specific conditions controlled by default behavior and configuration options.

Table 2: MED comparison summary

Condition	MED compared?
Paths from same neighbor AS	Yes
bgp bestpath med always configured	Yes, for all paths
Paths from different neighbor AS (by default)	No
Paths with only confederation segments (default)	No
bgp bestpath med confed configured	Yes, among confederation peers

Default comparison rules

- BGP compares MED values only between paths received from the same neighboring AS by default.
- If two paths are from different neighboring ASes, MED is not considered unless specific configuration options are set.

- When the AS path starts with an AS_SEQUENCE, the neighbor AS is the first AS number in that sequence. MED is compared only with other paths from the same neighbor AS.
- If a path lacks an AS path or its AS path starts with an AS_SET, the path is considered internal, and MED is compared only with other internal paths.
- When a path contains only confederation segments or starts with confederation segments followed by an AS_SET, MED is not compared by default unless the **bgp bestpath med confed** command is configured.
- If the AS path starts with confederation segments followed by an AS_SEQUENCE, MED is compared with other paths whose neighbor AS matches the first AS in the AS_SEQUENCE.

Configuration overrides

- If the **bgp bestpath med always** command is set, MED is compared among all paths, regardless of the neighbor AS.
- If the **bgp bestpath med confed** command is set, MED is also compared for paths from confederation peers.

Handling missing MED values

- If a route does not include a MED attribute, BGP assigns it a default value of 0.
- When the **bgp bestpath med missing-as-worst** command is configured, a missing MED is treated as the highest possible value, making the path least preferred.

Key points

- Lower MED values are preferred. BGP chooses paths with the lowest MED when other path attributes are equal.
- MED comparison is only relevant if higher-priority path attributes (weight, local preference, AS path length, origin) are the same.

Configure MED for BGP routes

Set the MED value to advertise to BGP peers for routes that do not already have a MED attribute.

Before you begin

Ensure you are in global configuration mode on the router.

Procedure

Set the default metric (MED) for routes without an existing MED attribute:

Example:

```
Router# configure
Routing(config)# router bgp 120
Routing(config-bgp)# default metric 10
```

Accumulated IGP attributes for BGP

Accumulated IGP attributes for BGP are optional BGP path attributes that

- record and propagate the cumulative IGP metric along a BGP route
- enable BGP to make routing decisions based on IGP metrics, and
- support multi-AS environments by simulating IGP metric behavior across BGP domains.

Table 3: Feature History Table

Feature Name	Release Information	Feature Description
Accumulated IGP Attribute for BGP	Release 7.3.2	This feature enables you to implement multiple contiguous BGP Autonomous Systems under a single administration. You can allow BGP to make its routing decisions based on the IGP metric just as an IGP would do.

The Accumulated IGP (AIGP) attribute is assigned an IANA type code and is included in BGP updates as a set of Type/Length/Value elements (TLVs) that specifically carry the accumulated IGP metric.

Purpose and benefits of AIGP

The AIGP attribute allows BGP to carry IGP cost equivalents between provider edge devices. It enables optimal path selection and OSPF-like behavior for redistributed routes within BGP.

Use case for AIGP in network design

AIGP is especially useful in networks where OSPF or LDP only carry metric information locally, but BGP redistributes prefixes or labels at area boundaries to remote areas. This design lets BGP make path decisions that include information about IGP costs from various network segments.

- When an OSPF or static route is redistributed into BGP, the AIGP attribute carries the total IGP cost from the originating router.
- Using a route-policy, the AIGP metric can be set during route redistribution to inform BGP of the correct path cost.

If the AIGP attribute is not set, BGP bases its routing decisions only on attributes like AS_PATH or LOCAL_PREF, without considering the underlying IGP costs. This can lead to suboptimal path selection in complex or multi-AS environments.

Configure accumulated IGP attributes for BGP

Enable BGP to use accumulated IGP (AIGP) metrics for route selection, simulating IGP metric-based decision-making and supporting seamless redistribution of IGP routes into BGP.

Use this task when you want BGP to compute routes based on IGP-derived metrics, especially in networks where prefixes are redistributed from IGP (such as OSPF) into BGP, and you need to preserve and propagate cost information throughout the BGP domain.

Before you begin

- Ensure BGP and the relevant IGP (such as OSPF) are configured and operational.
- Confirm that route-policies and address families relevant to redistribution are defined.

Follow these steps to configure accumulated IGP attributes for BGP:

Procedure

Step 1 Define a route-policy to set the AIGP metric using IGP cost.

Example:

```
Router(config)# route-policy aip_policy
Router(config-rpl)# set aigp-metric igp-cost
Router(config-rpl)# exit
```

Step 2 Configure redistribution of IGP routes (such as OSPF) into BGP using the defined route-policy.

Example:

```
Router(config)# router bgp 100
Router(config-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# redistribute ospf route-policy aip_policy
```

Step 3 (Optional) Repeat for static routes or other sources as needed, applying a route-policy that sets or modifies the AIGP metric as appropriate.

Step 4 Use the **show bgp prefix** command to verify the configuration.

Example:

```
Router# show bgp 10.0.0.1
Thu Sep 30 21:21:15.279 EDT
BGP routing table entry for 10.0.0.1/32
Versions:
Process bRIB/RIB SendTblVer
Speaker 4694 4694
Last Modified: Sep 30 21:20:09.000 for 00:01:06
Paths: (2 available, best #1)
Not advertised to any peer
Path #1: Received by speaker 0
Not advertised to any peer
Local
192.168.0.1 (metric 2) from 192.168.0.1 (192.168.0.6)
Received Label 24000
Origin IGP, localpref 80, aigp metric 900, valid, internal, best, group-best, labeled-unicast
Received Path ID 1, Local Path ID 1, version 4694
Originator: 192.168.0.6, Cluster list: 192.168.0.1
Total AIGP metric 902 <-- AIGP attribute received.
```

Administrative distances

A administrative distance is a routing protocol metric that:

- quantifies the trustworthiness of a routing information source,
- determines which route is preferred when multiple protocols learn the same route, and
- assigns lower values to more trusted routes.

Role of administrative distance in route selection

When a router learns about a network route from more than one routing protocol (such as eBGP, OSPF, or EIGRP), it uses administrative distance to decide which route to install in the IP routing table. The protocol with the lowest administrative distance is favored.

Administrative distance and BGP path selection

Administrative distance does not influence the BGP path selection algorithm, but it does determine whether BGP-learned routes are installed in the IP routing table.


Default administrative distances for common protocols

In most cases, routes learned through eBGP (distance 20) are preferred over routes learned from IGP (such as OSPF, distance 110; EIGRP, distance 90; IGRP, distance 100; and RIP, distance 120).

BGP default administrative distances

Distance	Default Value	Function
External	20	Applied to routes learned from eBGP.
Internal	200	Applied to routes learned from iBGP.
Local	200	Applied to routes originated by the router.

Figure 2: Back door example



Suppose Router A learns about network 160.10.0.0 from both eBGP (distance 20) and EIGRP (distance 90). The eBGP route is preferred because it has a lower administrative distance. If you want Router A to instead use the EIGRP path, you can configure a BGP back-door route:

```
Router(config)# router bgp 100
Router(config-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# network 160.10.0.0/16 backdoor
```

With this configuration, Router A treats the eBGP-learned route as a local route (distance 200), making the EIGRP-learned route (distance 90) the preferred path. If the EIGRP route fails, the eBGP route becomes active.

Configure BGP back-door routes

Configure a BGP back-door route so that the router prefers an IGP-learned path over an eBGP-learned path for specific network prefixes.

When a prefix is reachable via both IGP (internal) and eBGP (external), routers usually prefer the eBGP path due to its lower administrative distance. Use the BGP back-door feature to override this preference, so the router uses the IGP route and uses the eBGP route as a backup if the IGP path fails.

Before you begin

- Make sure BGP is enabled and running.
- Identify the network prefix you want to configure as a back-door route.

Follow this step to configure a BGP back-door route:

Procedure

Enter BGP configuration mode, and configure the network prefix with the backdoor option.

Example:

```
Router# configure
Router(config)# router bgp 120
Router(config-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# network 172.20.0.0/16 backdoor
```

Advanced path selection features

Modern BGP implementations provide advanced features for path diversity and resiliency. The BGP additional paths feature allows a BGP speaker to advertise more than just the best path, supporting use cases such as fast failover, route oscillation mitigation, and improved bandwidth utilization.

BGP best paths and multipaths

A best path and multipath prefix is a BGP routing outcome that

- identifies the optimal route (best path) for each destination prefix based on established selection rules
- enables the installation of multiple qualified paths (multipaths) in the forwarding table to enhance load balancing and redundancy, and
- determines which routes are advertised to BGP peers and how traffic is forwarded throughout the network.

BGP best path selection and advertisement

BGP routers receive multiple paths to the same destination prefix. By default, the BGP best path algorithm compares all valid paths and selects one as the best path to install in the local IP routing table and use for traffic forwarding. Only this best path is typically advertised to other BGP peers, which ensures consistent routing decisions across the network.

Multipath selection for load balancing and redundancy

However, to improve bandwidth utilization and provide redundancy, BGP implementations can select additional paths-called multipaths-if certain conditions are met. These multiple eligible paths are installed in the forwarding table, and incoming packets are load-balanced across the best path and the multipaths.

Community strings and the BGP add-path feature

Route reflectors in BGP can use community strings to signal local path decisions, differentiating between best paths and multipaths. The introduction of the BGP add-path feature allows a router to advertise more than the best path, signaling all equivalent or backup paths as per network policy and multipath rules.

Key considerations and limitations

- By default, BGP only advertises the best path to its peers. This behavior may not show you the entire routing state available on a BGP speaker.
- If Add-Path is not configured, installing multipath does not change what is advertised to peers.

A BGP router receives three paths to destination 10.1.1.0/24 from different peers. It selects one best path for routing and advertises only that path to other peers. If multipath is enabled and two paths qualify, both are used for forwarding traffic, providing load balancing and redundancy, though only the best path is advertised unless BGP add-path feature is configured.

Modify BGP best-path selection criteria

Change the default selection behavior of the BGP best-path algorithm to suit specific routing requirements.

BGP routers may learn multiple paths to a destination. By default, they use a specific algorithm to choose the best path. You can tune this behavior to influence routing and traffic flow between autonomous systems (ASes).

Before you begin

- Ensure you have access to global configuration mode on your router.
- Confirm that BGP is already enabled on the router.

Follow these steps to modify BGP best-path selection criteria:

Procedure

Step 1 Enter BGP configuration mode, and adjust MED comparison behavior.

- Use the **bgp bestpath med missing-as-worst** command to treat missing MED as least desirable.
- Use the **bgp bestpath med always** command to compare MED across all paths, regardless of AS.
- Use the **bgp bestpath med confed** command to compare MED values for confederation peers.

Example:

```
Router(config)# router bgp 126
Router(config-bgp)# bgp bestpath med missing-as-worst
```

Step 2 Customize path attribute evaluation.

- Use the **bgp bestpath as-path ignore** command to ignore AS path length for best-path selection.
- Use the **bgp bestpath compare-routerid** command to compare router IDs among otherwise equal paths.

Example:

```
Router(config)# router bgp 126
Router(config-bgp)# bgp bestpath as-path ignore
```

The BGP best-path algorithm selects routes based on your modified criteria.

BGP additional paths

A BGP additional path is a BGP feature that

- enables a BGP speaker to advertise multiple paths for the same network prefix
- provides path diversity to improve convergence and operational flexibility, and
- supports granular control of path selection and advertisement types per neighbor based on outbound policy configuration.

Table 4: Feature history table

Feature Name	Release Information	Feature Description
Additonal path control per neighbor	Release 7.3.15	This features allows flexibility and granular control of the advertisement of additional paths based on the neighbor outbound policy configuration. This is done by allowing configuration of combinations different path selection procedures unlike singular path selection, and extending neighbor outpound policy to have finer control of the path types to be advertised. This feature enables operational efficiency to manage additional paths and reduce scale of the paths in a typical clustered network architecture. Without this feature, the path scale limitation of the memory is impacted, and control plane convergence issues develop because of the excessive number of paths.

The BGP additional paths feature modifies BGP protocol behavior so a BGP speaker can send more than one path for a given prefix.

Benefits for path diversity and convergence

This modification provides network path diversity and enables edge routers to achieve prefix-independent convergence (PIC), especially in iBGP networks.

Enhanced policy control and scalability

By supporting multiple path types and expanded policy controls per neighbor, the feature helps prevent issues like path scale limitations and enables better control plane convergence.

BGP additional paths can advertise the following types of paths for a single prefix:

- Backup paths - enable fast convergence and rapid restoration of connectivity.
- Group-best paths - help resolve route oscillation scenarios.
- All paths - emulate an iBGP full mesh by advertising every possible path.