

BGP Path Selection Decision Process

BGP bases its decision process on the **attribute** values.

When faced with **multiple** routes to the **same** destination, BGP chooses the **best** route for routing traffic toward the destination. The following process summarizes how BGP chooses the best route.

1. If the next hop is inaccessible, do not consider it.

This decision is why it is important to have an IGP route to the next hop.

2. If the path is internal, synchronization is enabled, and the route is not in the IGP, do not consider the route.

3. Prefer the path with the **largest weight** (weight is a Cisco proprietary parameter).

4. If the routes have the same weight, prefer the route with the **largest local preference**.

5. If the routes have the same local preference, prefer the route that was **originated** by the **local router**.

For example, a route might be originated by the local router using the **network bgp** router configuration command, or through redistribution from an IGP.

6. If the local preference is the same, or if no route was originated by the local router, prefer the route with the **shortest autonomous system path** (i.e., AS path).

7. If the autonomous system path length is the same, prefer the route with the **lowest origin code** (IGP < EGP < INCOMPLETE).

8. If the origin codes are the same, prefer the route with the **lowest MED** metric attribute.

This comparison is only made if the neighboring autonomous system is the same for all routes considered, unless the **bgp always-compare-med** router configuration command is enabled.

Note The most recent Internet Engineering Task Force (IETF) decision regarding BGP MED assigns a value of infinity to the missing MED, making the route lacking the MED variable the least preferred. The default behavior of BGP routers running Cisco IOS software is to treat routes without the MED attribute as having a MED of 0, making the route lacking the MED variable the most preferred. To configure the router to conform to the IETF standard, use the **bgp bestpath med missing-as-worst** router configuration command.

9. Prefer the external BGP (eBGP) path over the iBGP path.

All confederation paths are considered internal paths.

10. Prefer the route that can be reached through the closest IGP neighbor (the lowest IGP metric).

The router will prefer the shortest internal path within the autonomous system to reach the destination (the shortest path to the BGP next hop).

11. If the following conditions are all true, insert the route for this path into the IP routing table:

- Both the best route and this route are external.
 - Both the best route and this route are from the same neighboring autonomous system.
 - The **maximum-paths** router configuration command is enabled.
-

Note eBGP load sharing can occur at this point, which means that multiple paths can be installed in the forwarding table.

12. If multipath is not enabled, prefer the route with the lowest IP address value for the BGP router ID.

The router ID is usually the highest IP address on the router or the loopback (virtual) address, but might be implementation-specific.

BGP Multipath Support

When a BGP speaker learns two identical eBGP paths for a prefix from a neighboring autonomous system, it will choose the path with the lowest route ID as the best path. This best path is installed in the IP routing table. If BGP multipath support is enabled and the eBGP paths are learned from the same neighboring autonomous system, instead of one best path being picked, multiple paths are installed in the IP routing table.

During packet switching, depending on the switching mode, either per-packet or per-destination load balancing is performed among the multiple paths. A maximum of six paths is supported. The **maximum-paths** router configuration command controls the number of paths allowed. By default, BGP will install only one path to the IP routing table.

Reference

https://www.cisco.com/c/en/us/td/docs/ios/12_2/ip/configuration/guide/fipr_c/1cfbgp.html