

# BGP Attributes and Path Selection

BGP (Border Gateway Protocol) routers usually receive multiple paths to the same destination. Like how our IGPs (RIP, EIGRP, OSPF) work, we need to select the best path to each destination.

IGPs select the path with the lowest metric. For example:

- RIP selects the path with the lowest hop count.
- OSPF selects the path with the lowest cost.
- EIGRP selects the path with the highest bandwidth and lowest delay (unless you change the K values).

BGP however, selects the best path based on a **list of attributes**. On the Internet, it's more important that you have granular control over how you forward your traffic and to which autonomous systems instead of just going for the shortest path based on a metric.

Let's look at a quick example. Below I have the output of the BGP table of a [looking glass server](#):

```
oute-views.optus.net.au>show ip bgp
BGP table version is 781755060, local router ID is 203.202.125.6
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
                r RIB-failure, S Stale, m multipath, b backup-path, x best-
external
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*  1.0.0.0/24       203.202.143.34           0  7474 4826 13335
i
*                  192.65.89.161           1           0  7474 4826 13335
i
*                  202.139.124.130          1           0  7474 4826 13335
i
*                  203.13.132.7             10          0  7474 4826 13335
i
*>                 203.202.143.33           0  7474 4826 13335
i
```

This BGP router has 5 paths for network 1.0.0.0/24. Look at the **> symbol** at the bottom left. The **>** symbol means that BGP has selected this path as the **best path**. This path will be installed in the routing table.

Out of all those 5 paths, why did BGP select this path as the best path?

# Attributes

This path was selected based on the following attributes:

## Priority Attribute

- 1 Weight
- 2 Local Preference
- 3 Originate
- 4 AS path length
- 5 Origin code
- 6 MED
- 7 eBGP path over iBGP path
- 8 Shortest IGP path to BGP next hop
- 9 Oldest path
- 10 Router ID
- 11 Neighbor IP address

Let me give you a quick overview of each attribute. We will cover these in other lessons in detail.

## Weight

Prefer the path with the **highest weight**. This is a value that is **local** to the router and it's Cisco proprietary. The default value is 0 for all routes that are not originated by the local router. You can learn how it works in the [BGP weight attribute lesson](#).

## Local Preference

The local preference is used within an autonomous system and exchanged between iBGP routers. We prefer the path with the **highest local preference**. The default value is 100. To learn more, take a look at the [BGP local preference attribute lesson](#).

## Originate

Prefer the path that the local router originated. In the BGP table, you will see **next hop 0.0.0.0**. You can get a path in the BGP table through the BGP network command, aggregation, or redistribution. A BGP router will prefer routes that it installed into BGP itself over a route that another router installed in BGP.

## AS path length

Prefer the path with the **shortest AS path length**. For example, AS path 1 2 3 is preferred over AS path 1 2 3 4 5. You can learn more about [AS path length here](#).

## Origin code

Prefer the **lowest origin code**. There are three origin codes:

- IGP
- EGP
- INCOMPLETE

IGP is lower than EGP and EGP is lower than INCOMPLETE. You can learn how it works in the [origin code lesson](#).

## MED

Prefer the path with the **lowest MED**. The MED is exchanged between autonomous systems. For a detailed explanation, take a look at the [MED lesson](#).

## eBGP path over iBGP path

Prefer eBGP (external BGP) over iBGP (internal BGP) paths.

## Shortest IGP path to BGP next hop

Prefer the path within the autonomous system with the **lowest IGP metric** to the BGP next hop.

## Oldest Path

Prefer the path that **we received first**, in other words, the oldest path.

## Router ID

Prefer the path with the **lowest BGP neighbor router ID**. The router ID is based on the highest IP address. If you have a loopback interface, then the IP address on the loopback will be used. The router ID can also be manually configured.

## Neighbor IP address

Prefer the path with the **lowest neighbor IP address**. If you have two eBGP routers and two links in between then the router ID will be the same. In this case, the neighbor IP address is the tiebreaker.

## Path Selection

When BGP has multiple paths to a destination they are stored in the BGP table. All paths are in the BGP table but **only one** gets installed in the routing table.

Which path do we select? We start at the **top of the list** with BGP attributes and **work our way to the bottom**:

1. We start with weight because it's at the top of the BGP attributes list. We now have two options:
  1. If one path has a better weight then we select this path as the best path.
  2. If the weight is equal, we move down to the next attribute.
2. The next attribute is local preference. Once again, we have two options:
  1. If one path has a better local preference then we select this path as the best path.
  2. If the local preference is equal, we move down to the next attribute.
3. We work our way down this attribute list until we have a tiebreaker to select the best path. If all paths have the same BGP attributes, then we end up with the neighbor IP address.

## BGP Best Path Selection

