

# Understand BGP MED Attribute

## Contents

---

### [Introduction](#)

### [Prerequisites](#)

#### [Requirements](#)

#### [Components Used](#)

#### [Conventions](#)

### [Background Information](#)

### [Case Study](#)

#### [Scenario 1](#)

#### [Scenario 2](#)

### [Related Information](#)

---

## Introduction

This document describes the Border Gateway Protocol (BGP) MED Attribute when it crosses over an AS boundary by implementation in different scenarios.

## Prerequisites

### Requirements

Cisco recommends that you have basic knowledge of BGP.

### Components Used

This document is not restricted to specific software and hardware versions. The scenarios discussed in this document use these hardware and software versions:

- Scenario 1: Cisco 2600 routers on Cisco IOS® Software Release 12.4 or later
- Scenario 2: Cisco 2600 routers on Cisco IOS® Software Release 12.4 or later

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

### Conventions


Refer to [Cisco Technical Tips Conventions](#) for more information on document conventions.

## Background Information

The Multi Exit Discriminator (MED) provides a dynamic way to influence another autonomous system (AS) in the way to reach a certain route when there are multiple entry points for that AS. BGP uses a systematic

procedure for the best path choice. There are other important attributes such as weight, local preference, originate route, and AS path that are taken in to account before you consider the MED attribute. So, if any of these criteria match, the MED attribute is not considered.

---



**Note:** When all other factors are equal, the exit point with the lowest MED is preferred.

---

## Case Study

### Scenario 1

When a BGP speaker learns a route from a peer, the route MED is passed to other interior BGP (iBGP) peers, but not to exterior BGP (eBGP) peers.

Router R1 and Router R2 are considered in the same AS, for example AS#100, and Router R3 belongs to AS#101. For easy convention, IP addresses in the /24 block are used.

Routers R1 and R2 are configured as follows:

<b>Router 1</b>
-----------------

```
(Config)#interface Loopback10
(Config-if)#ip address xx.xx.xx.xx xxx.xxx.xxx.xxx
(Config-if)#interface FastEthernet0/0
(Config-if)#ip address xx.xx.xx.xx xxx.xxx.xxx.xxx
(Config)#router bgp 100
(Config-router)#no synchronization
(Config-router)#bgp router-id xx.xx.xx.xx
(Config-router)#bgp log-neighbor-changes
(Config-router)#network xx.xx.xx.xx mask xxx.xxx.xxx.xxx route-map ATTACH_MED
(Config-router)#neighbor xxx.x.xx.x remote-as xxx
(Config-router)#no auto-summary
(Config)#access-list 10 permit xx.xx.xx.xx
(Config)#route-map ATTACH_MED permit xx
(Config)#match ip address xx
(Config)#set metric xxx
```

### Router 2

```
(Config)#interface FastEthernet0/0
(Config-if)#ip address xxx.x.xx.x xxx.xxx.xxx.x
(Config-if)#interface Serial1/0
(Config-if)#ip address xxx.x.xx.x xxx.xxx.xxx.x
(Config-if)#encapsulation frame-relay IETF
(Config-if)#no fair-queue
(Config-if)#frame-relay map ip xxx.x.xx.x 203 broadcast
(Config-if)#no frame-relay inverse-arp
(Config-if)#frame-relay lmi-type ansi
(Config)#router bgp 100
(Config-router)#no synchronization
(Config-router)#bgp router-id xx.xx.xx.xx
(Config-router)#bgp log-neighbor-changes
(Config-router)#neighbor xxx.x.xx.x remote-as 100
(Config-router)#neighbor xxx.x.xx.x remote-as 101
(Config-router)#neighbor xxx.x.xx.x ebgp-multihop 3
(Config-router)#no auto-summary
```

Router R3 configuration is shown here:

### Router 3

```
(Config)#interface Serial1/0
(Config-if)#ip address xxx.x.xx.x xxx.xxx.xxx.x
(Config-if)#encapsulation frame-relay IETF
(Config-if)#no fair-queue
(Config-if)#frame-relay map ip xxx.x.xx.x 302 broadcast
(Config-if)#no frame-relay inverse-arp
(Config-if)#frame-relay lmi-type ansi
```

```
(Config)#router bgp 101
(Config-router)#no synchronization
(Config-router)#bgp log-neighbor-changes
(Config-router)#neighbor xxx.x.xx.x remote-as 100
(Config-router)#neighbor xxx.x.xx.x ebgp-multihop 3
(Config-router)#no auto-summary
```

In this setup, R1 and R2 run iBGP. Therefore, when an update enters the AS with a certain metric, that metric is used to make decisions inside the AS.

The **show ip bgp** command, when checked from R2, shows the metric value for xx.xx.xx.xx, which comes from iBGP neighbor xxx.x.xx.x and has a MED value of 100.

eBGP runs between R2 and R3 because these are in a different AS. When the same update passes to a third AS, for example AS#101, that metric returns to 0.

The **show ip bgp** command, when checked from R3, has its metric removed, because xx.xx.xx.xx crosses the AS boundary(101).

From this scenario, it is evident that the MED attribute can influence the inbound traffic from neighbor autonomous systems.

The MED attribute cannot influence the route decisions of more remote autonomous systems. When a BGP speaker learns a route from a peer, it can pass the MED of the route to any iBGP peers, but not to eBGP peers.

As a result, the MED has relevance only between neighbor autonomous systems.

## Scenario 2

If the route injected into the BGP (either through the **network or redistribute** command) comes from an IGP (RIP or EIGRP or OSPF), the MED is derived from the IGP metric and the route is advertised to an eBGP neighbor with this MED.

In this network, R1 is configured to run in a RIP network. Routers R2 and R3 run BGP, where R2 is configured with AS 100 while R3 is with AS 101.

Router R1 is configured here:

Router R1
<pre>(Config)#interface Loopback10 (Config-if)#ip address xx.xx.xx.xx xxx.xxx.xxx.xxx (Config-if)#interface FastEthernet0/0 (Config-if)#ip address xxx.x.xx.x xxx.xxx.xxx.x (Config)#router rip (Config-router)#network xx.x.x.x (Config-router)#network xxx.x.xx.x (Config-router)#no auto-summary</pre>

Routers R2 and R3 are configured for BGP, where redistribution is done in R2 in order to inject the RIP networks to a BGP.

## Router R2

```
(Config)#interface FastEthernet0/0
(Config-if)#ip address xxx.x.xx.x xxx.xxx.xxx.x
(Config-if)#interface Serial1/0
(Config-if)#ip address xxx.x.xx.x xxx.xxx.xxx.x
(Config-if)#encapsulation frame-relay IETF
(Config-if)#no fair-queue
(Config-if)#frame-relay map ip xxx.x.xx.x 203 broadcast
(Config-if)#no frame-relay inverse-arp
(Config-if)#frame-relay lmi-type ansi
(Config)#router rip
(Config-router)# network xxx.x.xx.x
(Config-router)#no auto-summary
(Config-router)#router bgp 100
(Config-router)#no synchronization
(Config-router)#bgp router-id xx.xx.xx.xx
(Config-router)#bgp log-neighbor-changes
(Config-router)#neighbor xxx.x.xx.x remote-as 101
(Config-router)#neighbor xxx.x.xx.x ebgp-multihop 3
(Config-router)#redistribute rip metric 1
(Config-router)#no auto-summary
```

## Router R3

```
(Config)#interface Serial1/0
(Config-if)#ip address xxx.x.xx.x xxx.xxx.xxx.x
(Config-if)#encapsulation frame-relay IETF
(Config-if)#no fair-queue
(Config-if)#frame-relay map ip xxx.x.xx.x 302 broadcast
(Config-if)#no frame-relay inverse-arp
(Config-if)#frame-relay lmi-type ansi
(Config)#router bgp 101
(Config-router)# no synchronization
(Config-router)#bgp router-id xx.xx.xx.xx
(Config-router)#bgp log-neighbor-changes
(Config-router)#neighbor xxx.x.xx.x remote-as 100
(Config-router)#neighbor xxx.x.xx.x ebgp-multihop 3
(Config-router)#no auto-summary
```

Both RIP and BGP run on R2. If you check with the `show ip bgp` command, you can see that the prefix `xx.x.x.x` network is shown with a metric of 1, which is derived from the RIP.

However, in R3 which runs on eBGP, the network is advertised by consideration of the MED value derived from the IGP. In this case it is RIP. The prefix `10.0.0.0` is advertised with the IGP MED value, which is metric 1 of RIP. This can be seen in this output.

From this scenario, the behavior of the MED, in the case that networks are injected to the BGP router via the `network` or `redistribute` command, is seen where the actual MED value is replaced with that of the IGP metric.

Given that this attribute is a hint to external neighbors about the path preference into an AS, as stated earlier, it is not always considered if there are other more important attributes to determine the best route.

In order to have the same effect with a more deterministic attribute, use the `set as-path prepend` command under

the route map.

If you prepend the AS path for certain routes, it continues to be seen by other AS. For more information on the usage of AS-path prepend, refer to [Use of set-aspah prepend Command](#).

## Related Information

- [BGP: Frequently Asked Questions](#)
- [BGP Case Studies](#)
- [BGP Support Page](#)
- [BGP Multi-Homing: Design and Troubleshooting - Video from Live Webcast](#)
- [Technical Support & Documentation - Cisco Systems](#)