

Understanding BGP MED Attribute

Document ID: 112965

Contents

Introduction

Prerequisites

- Requirements

- Components Used

- Conventions

Case Study

- Scenario 1

- Scenario 2

Related Information

Introduction

The purpose of this document is to provide a better understanding of the Border Gateway Protocol (BGP) Multi Exit Discriminator (MED) Attribute when crossing over an autonomous system (AS) boundary by implementing it in different scenarios.

The MED provides a dynamic way to influence another AS in the way to reach a certain route when there are multiple entry points for that AS. BGP follows a systematic procedure for choosing the best path. There are other important attributes such as weight, local preference, originate route, and AS path that are taken in to account before considering the MED attribute. So, if any of these criteria matches, the MED attribute will not be considered.

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

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

Conventions


Refer to Cisco Technical Tips Conventions for more information on document conventions.

Case Study

Scenario 1

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

Consider this network setup:



Here, 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:

Router 1

```
(Config)#interface Loopback10
(Config-if)#ip address 10.10.10.10 255.255.255.255
(Config-if)#interface FastEthernet0/0
(Config-if)#ip address 192.1.12.1 255.255.255.0
(Config)#router bgp 100
(Config-router)#no synchronization
(Config-router)#bgp router-id 10.10.10.10
(Config-router)#bgp log-neighbor-changes
(Config-router)#network 10.10.10.10 mask 255.255.255.255 route-map ATTACH_MED
(Config-router)#neighbor 192.1.12.2 remote-as 100
(Config-router)#no auto-summary
(Config)#access-list 10 permit 10.10.10.10
(Config)#route-map ATTACH_MED permit 10
(Config)#match ip address 10
(Config)#set metric 100
```

Router 2

```
(Config)#interface FastEthernet0/0
(Config-if)#ip address 192.1.12.2 255.255.255.0
(Config-if)#interface Serial1/0
(Config-if)#ip address 192.1.23.2 255.255.255.0
(Config-if)#encapsulation frame-relay IETF
(Config-if)#no fair-queue
(Config-if)#frame-relay map ip 192.1.23.3 203 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 22.22.22.22
(Config-router)#bgp log-neighbor-changes
(Config-router)#neighbor 192.1.12.1 remote-as 100
(Config-router)#neighbor 192.1.23.3 remote-as 101
```

```
(Config-router)#neighbor 192.1.23.3 ebgp-multihop 3
(Config-router)#no auto-summary
```

Router R3 configuration is given here:


```
Router 3
(Config)#interface Serial1/0
(Config-if)#ip address 192.1.23.3 255.255.255.0
(Config-if)#encapsulation frame-relay IETF
(Config-if)#no fair-queue
(Config-if)#frame-relay map ip 192.1.23.2 302 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 192.1.23.2 remote-as 100
(Config-router)#neighbor 192.1.23.2 ebgp-multihop 3
(Config-router)#no auto-summary
```

In this setup, R1 and R2 have iBGP running. 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 show you the metric value for 10.10.10.10, which is learnt via iBGP neighbor 192.1.12.1 and will have a MED value of 100.

The output from R2 is shown here:

```
R2#show ip bgp
BGP table version is 3, local router ID is 22.22.22.22
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop        Metric LocPrf Weight Path
  *>i10.10.10.10/32 192.1.12.1          100   100    0  i
```



On R2, the prefix 10.10.10.10 learned via iBGP neighbor 192.1.12.1 carries 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 will have its metric removed, because 10.10.10.10 crosses the AS boundary(101).

The output from R3 is shown here:

```

R3#show ip bgp neighbors 192.1.23.2 routes
BGP table version is 2, local router ID is 192.1.23.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 10.10.10.10/32   192.1.23.2

```

↑


When the prefix (10.10.10.10) crosses the AS boundary (101), the MED attribute is removed. This is when the prefix was learned via an IBGP neighbor.

From this scenario it is evident that the MED attribute can influence the incoming traffic from neighboring autonomous systems. The MED attribute cannot influence the routing decisions of more-remote autonomous systems. When a BGP speaker learns a route from a peer, it can pass the route's MED to any iBGP peers, but not to eBGP peers. As a result, the MED has relevance only between neighboring autonomous systems.

Scenario 2

If the route injected into the BGP (either using 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.

This scenario uses this network setup:



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 as follows:

Router R1
<pre> (Config)#interface Loopback10 (Config-if)#ip address 10.10.10.10 255.255.255.255 (Config-if)#interface FastEthernet0/0 (Config-if)#ip address 192.1.12.1 255.255.255.0 (Config)#router rip (Config-router)#network 10.0.0.0 (Config-router)#network 192.1.12.0 (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 192.1.12.2 255.255.255.0
(Config-if)#interface Serial1/0
(Config-if)#ip address 192.1.23.2 255.255.255.0
(Config-if)#encapsulation frame-relay IETF
(Config-if)#no fair-queue
(Config-if)#frame-relay map ip 192.1.23.3 203 broadcast
(Config-if)#no frame-relay inverse-arp
(Config-if)#frame-relay lmi-type ansi
(Config)#router rip
(Config-router)# network 192.1.12.0
(Config-router)#no auto-summary
(Config-router)#router bgp 100
(Config-router)#no synchronization
(Config-router)#bgp router-id 22.22.22.22
(Config-router)#bgp log-neighbor-changes
(Config-router)#neighbor 192.1.23.3 remote-as 101
(Config-router)#neighbor 192.1.23.3 ebgp-multihop 3
(Config-router)#redistribute rip metric 1
(Config-router)#no auto-summary
```

Router R3

```
(Config)#interface Serial1/0
(Config-if)#ip address 192.1.23.3 255.255.255.0
(Config-if)#encapsulation frame-relay IETF
(Config-if)#no fair-queue
(Config-if)#frame-relay map ip 192.1.23.2 302 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 33.33.33.33
(Config-router)#bgp log-neighbor-changes
(Config-router)#neighbor 192.1.23.2 remote-as 100
(Config-router)#neighbor 192.1.23.2 ebgp-multihop 3
(Config-router)#no auto-summary
```

Both RIP and BGP run on R2. If you check using the **show ip bgp** command, you can see that the prefix 10.0.0.0 network is shown with a metric of 1, which is derived from the RIP.

The output from R2 is shown here:

```
R2#show ip route rip
R   10.0.0.0/8 [120/1] via 192.1.12.1, 00:00:20, FastEthernet0/0
R2#show ip bgp
BGP table version is 6, local router ID is 22.22.22.22
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop        Metric LocPrf Weight Path
*> 10.0.0.0       192.1.12.1      1         32768 ?
*> 192.1.12.0    0.0.0.0         0         32768 ?
```

On R2, the prefix 10.0.0.0 has a metric of 1

However, in R3 which runs on eBGP, the network is advertised by considering 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 RIP's metric 1.

This can be seen in this output:

```
R3#sh ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       I - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODP, P - periodic downloaded static route


Gateway of last resort is not set

B  192.1.12.0/24 [30/0] via 192.1.23.2, 00:10:24
C  192.1.23.0/24 is directly connected, Serial1/0
B  10.0.0.0/8 [20/1] via 192.1.23.2, 00:10:24
```

The prefix 10.0.0.0 is advertised to an eBGP (R3) neighbor with the original IGP MED value (1).

From this scenario the behavior of the MED, in the case of networks being injected to the BGP router via the **network** or **redistribute** command, is clearly seen where the actual MED value is being replaced with that of the IGP metric. Now, 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 will continue to be seen by other AS. For more information on the usage of As-path prepend, refer to Use of Set-aspath 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](#)

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2013 – 2014 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Nov 07, 2014

Document ID: 112965
