

# Experiment 3a - Comparing Received Prefixes

DE-CIX Academy

Version 2.0 - Students Copy - November 15, 2019

## 1 Introduction

In this experiment you will receive some prefix announcements. Figure out why the one selected as “best” prefix is considered best by your router.

## 2 Network Setup




Figure 1: Network Setup

Figure 1 shows the network topology for this experiment. It is the same as for experiment 2. All participants are using AS64501 upwards and are connected to AS65550, AS64496 (and if required to AS286 for peering).

## 3 Experiments

You are sent various prefixes via your configured peering connections.

### 3.1 Tasks:

- Gather information about what you receive

- Explain why one of the announcement is selected as “best” by your router.

### 3.1.1 Commands you need:

**show bgp ipv4** lists all prefixes in the BGP table.

**show bgp ipv4 <prefix>** lists prefix <prefix> including all details the router uses to make its best path selection.

**show bgp ipv4 summary** lists all neighbors and the number of prefixes they announce

**show bgp ipv4 neighbor <address> routes** shows the accepted prefixes from neighbor <address>

**show bgp ipv4 neighbor <address> received-routes** shows all prefixes from neighbor <address>, even the ones filtered out

Remark: The commands are the same for IPv6 - just use *ipv6* instead of *ipv4*.

## 3.2 Configuration commands to influence best path selection

These commands have to be entered in configuration mode ("conf term") and in context of the BGP router ("router bgp 645xx").

### 3.2.1 AS path length

**bgp bestpath as-path ignore** disables the use of the AS path length as criterium for best path selection. Do not use.

### 3.2.2 MED

**bgp always-compre-med** this enables the comparison of MEDs if the first AS in the AS path is *not* the same. Use with care.

**bgp bestpath missing-as-worst** changes the default behavior that a missing MED attribute is treated as *best* - now a missing MED is treated as *worst*.

## 4 Slides

## BGP Best Path Selection: Local Preference

1. **Local Preference: Higher wins**
2. AS-Path length: Shorter wins
3. Origin Type: *IGP* over *EGP* over *Incomplete*
4. MED: Lower wins
5. eBGP, iBGP: eBGP wins
6. Age of route: Older wins
7. Router ID: Lowest wins
8. Neighbor IP address: Lowest wins

80.81.192.1	Peer
10.230.xx.1	Upstream
10.200.xx.1	Upstream

Show commands:

- show bgp ipv4 summary
- show bgp ipv4
- show bgp ipv4 129.13.0.0

## BGP Best Path Selection: AS path length

1. Local Preference: Higher wins
2. **AS-Path length: Shorter wins**
3. Origin Type: *IGP* over *EGP* over *Incomplete*
4. MED: Lower wins
5. eBGP, iBGP: eBGP wins
6. Age of route: Older wins
7. Router ID: Lowest wins
8. Neighbor IP address: Lowest wins

### Show commands:

- show bgp ipv4 summary
- show bgp ipv4
- show bgp ipv4 129.13.0.0

## BGP Best Path Selection: Origin Type

1. Local Preference: Higher wins
2. AS-Path length: Shorter wins
3. **Origin Type: *IGP* over *EGP* over *Incomplete***
4. MED: Lower wins
5. eBGP, iBGP: eBGP wins
6. Age of route: Older wins
7. Router ID: Lowest wins
8. Neighbor IP address: Lowest wins

i	IGP	Injected using <i>network</i> statement
e	EGP	Injected from EGP (obsolete)
?	Incomplete	Injected using <i>redistribute</i> from other routing protocol

Show commands:

- show bgp ipv4 summary
- show bgp ipv4
- show bgp ipv4 129.13.0.0

## BGP Best Path Selection: Multi Exit Discriminator (MED)

1. Local Preference: Higher wins
2. AS-Path length: Shorter wins
3. Origin Type: *IGP* over *EGP* over *Incomplete*
4. **MED: Lower wins**
5. eBGP, iBGP: eBGP wins
6. Age of route: Older wins
7. Router ID: Lowest wins
8. Neighbor IP address: Lowest wins

**MED is only compared if the first AS in the path is the same!**

Show commands:

- show bgp ipv4 summary
- show bgp ipv4
- show bgp ipv4 193.141.4.0

## BGP Best Path Selection: Older (more stable) prefix wins

1. Local Preference: Higher wins
2. AS-Path length: Shorter wins
3. Origin Type: *IGP* over *EGP* over *Incomplete*
4. MED: Lower wins
5. eBGP, iBGP: eBGP wins
6. **Age of route: Older wins**
7. Router ID: Lowest wins
8. Neighbor IP address: Lowest wins

### Show commands:

- show bgp ipv4 summary
- show bgp ipv4
- show bgp ipv4 193.141.4.0

## **BGP Best Path Selection: If own AS is in path, discard**

1. Local Preference: Higher wins
2. AS-Path length: Shorter wins
3. Origin Type: *IGP* over *EGP* over *Incomplete*
4. MED: Lower wins
5. eBGP, iBGP: eBGP wins
6. Age of route: Older wins
7. Router ID: Lowest wins
8. Neighbor IP address: Lowest wins

### Show commands:

- show bgp ipv4 summary
- show bgp ipv4
- show bgp ipv4 neighbor 10.200.xx.1 received-routes
- show bgp ipv4 neighbor 10.200.xx.1 routes



## BGP Best Path Selection: General routing rules still take priority

1. Local Preference: Higher wins
2. AS-Path length: Shorter wins
3. Origin Type: *IGP* over *EGP* over *Incomplete*
4. MED: Lower wins
5. eBGP, iBGP: eBGP wins
6. Age of route: Older wins
7. Router ID: Lowest wins
8. Neighbor IP address: Lowest wins

### **o. More specific wins**

Show commands:

- show bgp ipv4 summary
- show bgp ipv4