



# BGP Notes

# Overview

## Definition

Border Gateway Protocol version 4 is the routing protocol currently used on the Internet. It is an Exterior Gateway Protocol (EGP) unlike RIP, OSPF, IS-IS and EIGRP that are Interior Gateway Protocol (IGP).

BGP exchanges routing information by forming a unicast neighbor relationship with other devices running BGP. This connection is formed using the TCP protocol on port 179.

BGP can establish internal or external peering. BGP peers established in the same autonomous system are called iBGP peers, BGP peers established between different autonomous systems are called eBGP peers.

BGP neighbors in the same autonomous system SHOULD be fully meshed.

## Rules of prefixes advertising

A prefix is advertised to a neighbor based on the type of neighbor the prefix was learned from.

- Route learned from an eBGP peer, the route will be propagated to all peers.
- Route learned from an iBGP peer, the route will be propagated to eBGP peers only. This is why the BGP enabled devices within a AS have to be fully meshed or to use route-reflectors or confederations.
- Route originated locally, the route will be propagated to all peers.
- A route is advertised only if it has been selected as the best for a specific destination.

## Rules of synchronization (enabled by default on Cisco platforms)

It only applies to iBGP learned prefixes.

- Before BGP advertises a prefix to an external peer (eBGP) it must check that the prefix is available through IGP and the next hop router exists.
- Next hop addresses that are only reachable via a default route are not valid.
- Next hop addresses that are only reachable via another BGP route are also not valid.

## The three ways to add prefixes into BGP

- network statement, 200 entries maximum for Cisco devices before 12.0 and ? for Juniper
- aggregate statement
- redistribute an IGP into BGP

## Rules of path selection process

- Ignore the path that are not synchronized, if the synchronization is enabled (currently Cisco default). See above rule of synchronization.
- Ignore the path where the NEXT\_HOP is inaccessible.
- Ignore the path from an eBGP neighbor if the local AS appears in the AS path.
- If the BGP best path enforce-first-as is setup and the UPDATE does not contain the AS of the neighbor as the first AS number in the AS\_SEQUENCE, send a NOTIFICATION and close the peer session.
- Ignore path marked as received-only. You can view these entries with the “show ip bgp <ip address>” command.
- Ignore path with the next-hop metric marked as inaccessible.

## Rules of best path selection process with Cisco IOS

- Prefer the path with the largest **weight**. This is a Cisco specific parameter. The information is local to the router.
- Prefer the path with the largest LOCAL\_PREF.
- Prefer the path that was locally originated, i.e. through **network** or **aggregate** BGP command or redistribution from an IGP.
- Prefer the path locally sourced, network, redistributed over locally generated aggregate.
- Prefer the path with the shortest AS path.
- Prefer the path with the lowest origin path, internal (IGP) 0, external (EGP) 1, incomplete 2.
- Prefer the path with the lowest Multi-Exist Discriminator, MED.
- Prefer the path learned via eBGP over iBGP.
- Prefer the path with the lowest IGP metric to the next hop.
- If the **maximum-paths** N is configured and there are multiple path from the same neighbor, add up to the N path to the routing table. This is for eBGP multi path load-sharing.
- Prefer the oldest path
- Prefer the path originated from the lowest BGP router ID, if the router ID is the same prefer the path with the minimum cluster ID. This is for route-reflector environment only
- Prefer the path originated from the lowest neighbor address. This is the address the remote peer use to connect to the local device.

**Gotcha** : When the iBGP speaker are NOT FULLY meshed and there is no route reflector or confederation, any prefix learned from an iBGP neighbor will only be advertised to eBGP neighbors.

## Configuration best practice

Always set the BGP router ID.

Gotcha, OSPF to BGP redistribution will only work if the OSPF router ID and BGP router ID are the same.

# ANNEX

## Default administrative Distance Values

	Values	Manufacturers
Connected interface	0	Cisco, Juniper
Static route	1	Cisco, Juniper
EIGRP summary route	5	Cisco
External BGP	20	Cisco, Juniper
Internal EIGRP	90	Cisco
IGRP	100	Cisco
OSPF	110	Cisco, Juniper
ISIS	115	Cisco, Juniper
RIP	120	Cisco, Juniper
EGP	140	Cisco
ODR	160	Cisco
External EIGRP	170	Cisco
Internal BGP	200	Cisco, Juniper
Unknown	255	Cisco, Juniper

As per as the document available on [www.cisco.com](http://www.cisco.com) and [www.juniper.net](http://www.juniper.net).