

## BGP — General

Path Vector Routing Protocol  
 Exterior Gateway Protocol (EGP)  
 RFC 4271 — BGP version 4  
 TCP Port 179 (reliable transport)  
 Admin Distance: eBGP = 20 | iBGP = 200  
 Routes between Autonomous Systems (AS)  
 Supports IPv4, IPv6, MPLS/VPN (MP-BGP)  
 Policy-based routing via path attributes  
 No automatic neighbour discovery  
 Incremental / triggered updates only  
 Supports route aggregation & summarisation

## eBGP vs iBGP

<b>eBGP</b>	Between different Autonomous Systems
<b>iBGP</b>	Within the same Autonomous System
<b>eBGP AD</b>	20
<b>iBGP AD</b>	200
<b>eBGP TTL</b>	1 (directly connected, default)
<b>iBGP TTL</b>	255
<b>eBGP multihop</b>	ebgp-multihop <ttl>
<b>iBGP full mesh</b>	Required OR use Route Reflector
<b>Next-hop-self</b>	Required when iBGP peers can't reach NH

## BGP Message Types

<b>OPEN</b>	Establish session & exchange capabilities
<b>KEEPALIVE</b>	Maintain session (default every 60 s)
<b>UPDATE</b>	Advertise / withdraw prefixes
<b>NOTIFICATION</b>	Report error — tear down session
<b>ROUTE-REFRESH</b>	Request prefix re-advertisement

## BGP Best Path Selection (evaluated in order)

- 1 Highest WEIGHT (Cisco proprietary — local only)
- 2 Highest LOCAL\_PREF (higher = preferred)
- 3 Locally originated (network / aggregate / redistrib)
- 4 Shortest AS\_PATH length
- 5 Lowest ORIGIN code (i < e < ?)
- 6 Lowest MED (Multi-Exit Discriminator)
- 7 eBGP preferred over iBGP path
- 8 Lowest IGP metric to NEXT\_HOP
- 9 Oldest eBGP path (most stable)
- 10 Lowest BGP Router-ID of advertising peer
- 11 Lowest neighbour IP address (final tiebreaker)

## BGP Neighbour FSM States (in order)

<b>Idle</b>	Not started / TCP connection blocked
<b>Connect</b>	TCP SYN sent — connection in progress
<b>Active</b>	TCP failed — actively retrying
<b>OpenSent</b>	TCP up — OPEN message sent
<b>OpenConfirm</b>	OPEN rcvd — waiting for KEEPALIVE
<b>Established</b>	Session up — exchanging routes

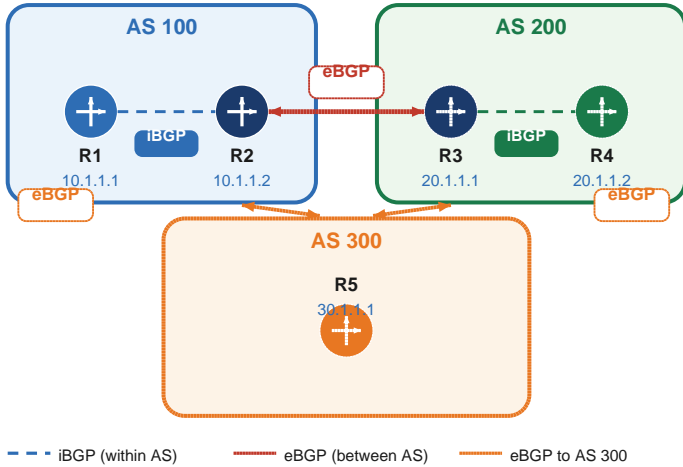
## BGP Path Attributes

<b>WEIGHT</b>	Cisco only — higher=better — not advertised
<b>LOCAL_PREF</b>	Well-known Discret. — higher=better
<b>AS_PATH</b>	Well-known Mandatory — shorter=better
<b>NEXT_HOP</b>	Well-known Mandatory — must be reachable
<b>ORIGIN</b>	Well-known Mandatory — i / e / ?
<b>MED</b>	Optional Non-transitive — lower=better
<b>COMMUNITY</b>	Optional Transitive — tag & filter routes
<b>AGGREGATOR</b>	Optional Transitive — identifies aggregator
<b>ATOMIC_AGGR</b>	Well-known Discret. — summary flag
<b>ORIGINATOR_ID</b>	Optional Non-trans — Route Reflector use
<b>CLUSTER_LIST</b>	Optional Non-trans — RR loop prevention

## Key Cisco IOS Commands

```
router bgp <ASN>
neighbor <ip> remote-as <ASN>
neighbor <ip> update-source loopback0
neighbor <ip> ebgp-multihop <ttl>
neighbor <ip> next-hop-self
neighbor <ip> route-map <name> in | out
neighbor <ip> prefix-list <name> in | out
neighbor <ip> weight <value>
neighbor <ip> send-community
neighbor <ip> shutdown
network <prefix> mask <mask>
aggregate-address <prefix> <mask> summary-only
maximum-paths <n> / maximum-paths ibgp <n>
bgp bestpath as-path multipath-relax
clear ip bgp * soft
show ip bgp summary
show ip bgp <prefix>
show ip bgp neighbors <ip> advertised-routes
show ip bgp neighbors <ip> received-routes
```

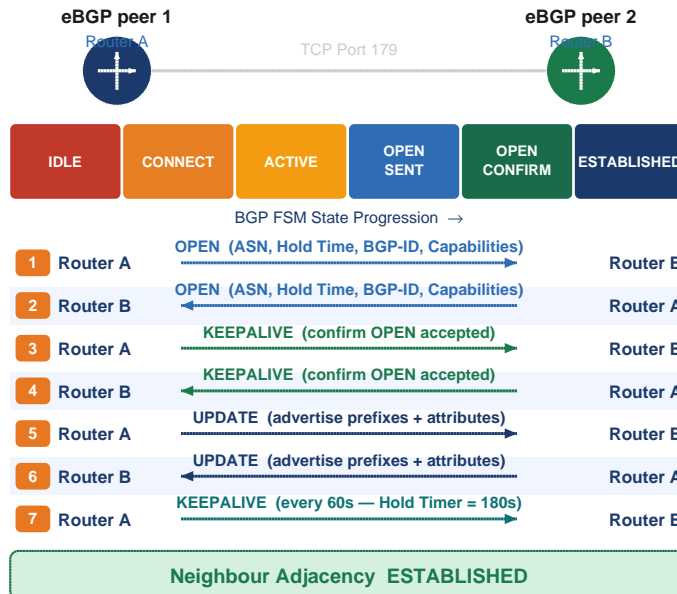
BGP Network Topology (eBGP & iBGP)



Route Reflector — avoids iBGP full mesh

RR reflects iBGP routes to all RR-clients.  
 Clients only peer with the RR — not each other.  
 cluster-id attribute prevents routing loops.  
 Confederation is an alternative to full mesh.

BGP Session Establishment & Message Exchange



BGP Session Requirements

ASN must match (e.g. router bgp 100)  
 K-values do not apply — TCP-based protocol  
 Routers do NOT need to be directly connected  
 Loopback source recommended for iBGP stability  
 Authentication password must match if configured  
 Hold time negotiated — lower value is used

NOTIFICATION Message — Sent on errors

Sent when an error is detected:  
 Bad AS number / unsupported capability  
 Hold timer expired / FSM error  
 Session is torn down immediately after.

BGP Packet & UPDATE Message Structure

TCP / IP Wrapper (BGP runs over TCP port 179)

Eth Hdr 14B	IP Hdr 20B	TCP Hdr 20B	BGP Marker 16B	BGP Length 2B	BGP Type 1B	BGP Data Variable
----------------	---------------	----------------	-------------------	------------------	----------------	----------------------

BGP Common Header (19 bytes minimum)

Marker 16 bytes — all 1s (synchronisation / auth)	Length 2 bytes — total message length byte — msg type (1-)	Type
--	---	------

BGP OPEN Message (Type = 1)

Version 1B	My AS 2B	Hold Time 2B	BGP Identifier 4B	Opt Param Len 1B	Optional Parameters Variable
---------------	-------------	-----------------	----------------------	---------------------	---------------------------------

BGP UPDATE Message (Type = 2) — carries prefixes & attributes

Withdrawn Routes 2B	Withdrawn Routes Variable	Path Attributes 2B	Path Attributes Variable	NLRIPrefixes Variable
------------------------	------------------------------	-----------------------	-----------------------------	--------------------------

NLRI Encoding (prefix list inside UPDATE message)

Length 1B — prefix bits	Prefix Variable — network address	Length 1B — next prefix	Prefix Variable
----------------------------	--------------------------------------	----------------------------	--------------------

Path Attribute Flags Byte (1st byte of every attribute)

O Optional	T Transitive	P Partial	E Ext Len	0 Unused	0 Unused	0 Unused	0 Unused
---------------	-----------------	--------------	--------------	-------------	-------------	-------------	-------------

BGP Message Type Codes

1	OPEN	Initiate BGP session
2	UPDATE	Advertise / withdraw routes
3	NOTIFICATION	Error — session teardown
4	KEEPALIVE	Session keepalive heartbeat
5	ROUTE-REFRESH	Request re-advertisement