

Attributes:	
Well-known Mandatory (Optional bit 0)	AS-Path Next Hop Origin Type (Junos only "I")
Well-known Discretionary (Optional bit 0)	Localpref (def 100, only in AS) Atomic Aggregate (marks that ASs were dropped in aggreg.) If using AS_SET, NO Atom_aggr
Optional Transitive - Tbit 1 (marked as partial if unSUP.)	Aggregator Community Ext-Community
Option Non-Transitive Tbit0 (deleted if unsupported)	Cluster List - RR Originator ID - RR MED (def. compare same AS, 0)
AFI: IPv4,IPv6,MPLS_VPN	MP_REACH_NLRI
SAFI: unicast,mcast	MP_UNREACH_NLRI

Best Path Algorithm - WLLA OMNI
0. Ignore ASPathLoop, !NextHop, not synchronized
1. Weight Cisco
2. Local Pref.
3. Locally Orig. or redistrib. (netw.+redist > agrgr) Cisco
4. AS-PATH (AS-SET = 1, AS_CONFED_SEQUENCE and AS_CONFED_SET ignored)
5. Origin type (lowest): IGP < EGP < Incomplete
6. MED (if same first AS / any preceding AS_CONFED_SEQUENCE is ignored)
7. Neighbor: eBGP > iBGP (confed treat as iBGP)
8. IGP metric to NH (lowest)
9. Multipath check (if Yes, skip RID,PeerIR)
10. Oldest path for external paths for stability Cisco
11. Lowest RID (originator ID if RR) JunOS
12. Minimum cluster list (in RR environment)
13. Lowest peer address (if here then multipath)

Private ASs: 64512 - 65535 (1024)
BGP Link BW feature: advert. link bandwidth in iBGP for multipath over links with unequal BW
NoExport (only to neigh AS) / **NoAdvertise** (only to neigh)
NoExportSubConf (only in Sub_AS)

Graceful Restart: (new capability negot + AFI/SAFI). After restart, estab. peer & set "restart-bit" ro resume (FW-bit "I can still fwd"). Routes marked stale (used). Routes adv by peer. **End-of-RIB** used to mark end of Updates. (empty **Update** packet). Restart def. **120sec**.

BGP Header:
 - **Marker** 16B - all 1s/auth
 - **Length** 2B
 - **Type** 1B:
 1 - OPEN
 2 - Update
 3 - Notification
 4 - Keepalive
 5 - RRefresh

BGP Open Message:
 - **Version**
 - **AS**
 - **HoldTime** (lowest, 90s/3 = KA)
 - **RID**
 - **Optional Capabill.** & params. (RRefresh, Auth, NLRI, .. TLVs)

BGP Update Message:
 - Unfeasible routes len
 - Withdrawn routes
 - Total Path Attrs Len
 - Path Attrs.
 - NLRI
 Attrs:
 - **Optional** Bit
 - **Trazitive** Bit
 - **Partial** Bit

BGP States:
 1. **Idle** (after Start, SYN)
 2. **Connect** (waitTCP, bad->Act)
 3. **Active** (Send SYN->Conn)
 4. **OpenSent** (TCP OK, -> open)
 5. **OpenConfirm** (->KeepA.)
 6. **Established** (when keepA. received in OpenConfirm)
 - dual connections => the one established by high RID **OK**
 - "passive" can be used
 - generally problem means flap between **Connect** and **Active**.

BGP Keepalive: 19 bytes header and no data (replaceable by Update)
BGP Notification: Error Code + Subcode + Data (6 - Cease, 4 Hold expired)
BGP Route Refresh: AFI + reserved + SAFI - request readvertisement of Adj-RIB-Out

RRs
 - Only advertise Best Routes
 - do not change Attrs
 - **cluster list + originator ID**
 - **origID:** 1st router that sent route to RR
 - CL > RR > CL + NONCL All except NONCL to NONCL
 - NONCL > RR > CL
RRs operation
 - on reflect, if no Originator ID, add RID of originator router in local AS
 - add clusID to clusList or create clusList
Loop detection:
 If Orig recv. = RID, ignore update.
 If ClusterID in Cluster-list recv, ignore

MED
Non-Deterministic (Cisco def., not rec)
 - meds chosen in order of arrival
Deterministic-MED
 - group by AS then choose
Always compare MEDs
 - compare MEDs from any source

Confederations:
 - adds AS_CONFED_SEQ, CBGP
 - CBGP does not modify Attrs.
IPv6 requires multiprotocol extensions capability in the OPEN + MP_REACH_NLRI attr.
Dampening
 1000p for Adv/Wdraw
 500p for Attrib Change

 3000 - suppress
 750 - reuse
 15 min - decay
ORF - Outbound Route Filter