

# BGPv4 - Border Gateway Protocol

Advanced Distance-Vector protocol

Exterior Routing Protocol that connects Autonomous Systems (AS) via advertised policies - (Network Paths)

Uses **TCP port 179** - Triggerred **UPDATE** and **KEEPALIVE** messages

**BGPv4** - (RFCs: 1771-1774, 1893, 1965-1966 ...)

In BGP terminology, a route is made up of **Network Layer Reachability Information (NLRI) and path attributes**. In BGP-4 The NLRI consists of an IP address and a prefix length which is the number of bits that make up the range of addresses. BGP-4 does not respect the traditional class distinctions. This means BGP-4 can aggregate many networks into a single advertisement and indicates this with the two **Aggregate Attributes**

## Common BGP Header

at least 19 and no greater than 4096



Value that the receiver of the message can predict.

**Stub AS network** (One ISP in and out) - **Single**  
**Transit AS** (More than one AS to enter network)  
**non-Transit AS** (Two ISP servicing network)

**Ingress filtering** - select routes to other BGP neighbors or peers

## PATH ATTRIBUTES

(which act as metrics)


Value / Description

- |                         |                           |
|-------------------------|---------------------------|
| 1 ORIGIN                | 10 CLUSTER_LIST           |
| 2 AS_PATH               | 11 DPA                    |
| 3 NEXT_HOP              | 12 ADVERTISER             |
| 4 MULTI_EXIT_DISC (MED) | 13 RCID_PATH / CLUSTER_ID |
| 5 LOCAL_PREF            | 14 MP_REACH_NLRI          |
| 6 ATOMIC_AGGREGATE      | 15 MP_UNREACH_NLRI        |
| 7 AGGREGATOR            | 16 EXTENDED COMMUNITIES   |
| 8 COMMUNITY             | 17                        |
| 9 ORIGINATOR_ID         | 255 Reserved              |
| 10 CLUSTER_LIST         |                           |

- 1 - OPEN (first msg. sent)
- 2 - UPDATE
- 3 - NOTIFICATION
- 4 - KEEPALIVE (19 byte header sent every one-third of hold time)
- 5 - ROUTER-REFRESH (AFI - Address Family Identifier) RFC2981

## OPEN Message (1)

(First message sent after a TCP session has been established)  
Used to identify the AS that the router is advertising



## UPDATE Message

(Actual topology)

**MED** - informs other external AS routers as to which route to use in order to receive traffic.

BGP's **NEXT\_HOP** is always the IP address of the first router in the next autonomous system, even though this may actually be several hops away. **NEXT\_HOPs** are only changed across EBGP sessions, but left intact across IBGP sessions.

## CISCO BGP Configuration

```
router bgp <1-65,535> ASN number
bgp dampening
neighbor 1.10.21.1 remote-as 12015
(You MUST explicitly specify the IP addresses of the remote
or neighbor BGP router - eBGP uses different AS)
network 172.16.31.0 mask 255.255.255.0
(network-number represents the network that is advertised)
```

```
sh ip bgp
sh ip bgp neighbors
sh ip bgp routes
sh ip bgp paths
sh ip bgp summary
```

**eBGP** - External BGP  
**iBGP** - Internal BGP

Full-Mesh

Autonomous

Prefix

(0-16 = 00000000-11111111) = 16-bit number