



CCIE Enterprise Infrastructure v1.0 Bootcamp

BGP



BGP

- + Core BGP Goals
 - + Establish Inter-AS Connectivity
 - + Establish Intra-AS Connectivity
 - + Originate NLRI to establish IPv4/IPv6 Connectivity

Inter-AS Connectivity

- + What is my ASN?
 - + Global AS
 - + Private AS
 - + Confederation AS
 - + Local AS
- + How are the peers connected?
 - + Directly connected EBGP peers
 - + Multi-hop EBGP peers
 - + Directly connected loopback EBGP peers

Inter-AS Connectivity (cont.)

- + What are the peering requirements?
 - + TTL Security
 - + Authentication
 - + Timers
 - + BFD
 - + Max Prefixes
 - + Dampening
 - + Communities

BGP Transport Considerations

- + BGP uses TCP 179 unicast
- + Implies IP reachability between peers must already exist
 - + Typically directly connected routes for EBGp
 - + Typically IGP connectivity for iBGP
- + What else can affect BGP transport?
 - + Path MTU Discovery
 - + Active vs. passive peers
 - + Data plane (ACL) filters
 - + Stateful firewall filters
 - + Control plane (CoPP) filters

BGP Update Source Considerations

- + BGP neighbor statement binds a socket
 - + SRC IP
 - + DST IP
 - + TCP Port
- + Packets not matching socket are discarded
 - + TCP RST sent by default
- + Implies update source must match from client to server
 - + Not necessarily server to client

EBGP Loop Prevention

- + AS-Path used to prevent looping updates
- + Can be modified with...
 - + allow-as-in
 - + enforce-first-as
 - + as-override on MPLS PE to CE

BGP Next-Hop Processing

- + Outbound EBGP updates are set to next-hop-self by default
- + Inbound EBGP updates sent to iBGP neighbors are next-hop-unchanged
 - + Implies PE-CE link is advertised into internal IGP/BGP
 - + Allows for BGP Next-Hop Tracking (NHT)
- + Inbound EBGP update sent to iBGP neighbors can be set to next-hop-self
 - + Implies PE-CE link does not need to be in internal IGP/BGP
 - + Breaks BGP Next-Hop Tracking
 - + i.e. convergence implications

BGP Next-Hop Processing (cont.)

- + What else can affect BGP next-hop processing?
 - + Manual route-map set ip next-hop
 - + BGP Selective Next-Hop Tracking
 - + Route reflectors?
 - + Confederation?

iBGP Peering Considerations

- + Supported peering designs
 - + iBGP full mesh
 - + Route Reflectors
 - + Confederation
 - + Mix of three is supported
 - + E.g. full mesh within a cluster, RR between clusters

iBGP Full Mesh

- + Negative is scaling limitations
 - + $n * ((n-1)/2)$ peering relationships
- + Positive is full topology info
 - + All peers learn all possible exits
 - + i.e hot potato vs. cold potato

iBGP Route Reflectors

- + RRs change route advertisement rules
 - + EBGP learned routes...
 - + Pass to EBGP peers, Clients, & Non-Clients
 - + Client learned routes...
 - + Pass to EBGP peers, Clients, & Non-Clients
 - + Non-Client learned routes...
 - + Pass to EBGP peers & Clients
 - + RR placement based upon these rules

iBGP Route Reflectors (cont.)

- + iBGP Route Reflection Loop Prevention
 - + Loop prevention through Cluster-ID
 - + RR discards routes received with its own Cluster-ID
 - + Sets Originator-ID attribute to the router-id of RR client on routes received from the client
 - + Client uses Originator-ID for loop prevention
 - + Implies that duplicate cluster-ids/router-ids could be a problem

iBGP Confederation

- + Confeds change advertisement rules through Sub-ASes
 - + Normal iBGP rules within Sub-AS
 - + EBGP rules between Sub-ASes
 - + Next-hops not modified between Sub-ASes
- + Confed requires changing the global BGP ASN
 - + i.e. deleting the process and re-creating it

Other iBGP Considerations

- + Are all devices in the transit path running iBGP
 - + Synchronization problem
 - + Tunnel solutions?
- + Does iBGP and IGP agree on exit points?
 - + Mismatch can cause traffic black holes
 - + Route deflection problem

Originating NLRI Considerations

- + Peering is separate from NLRI advertisements
 - + In IGP they are the same process
 - + i.e. enable OSPF and it automatically advertises the routes on your links
- + NLRI can be originated a number of ways and in a number of places
 - + Network statement
 - + Redistribution
 - + Aggregation
 - + Conditional advertisement
 - + Conditional route injection
- + What are the implications of different methods and where they are implemented?

BGP Path Selection Considerations

- + Understanding path selection order is a must
 - + You can't manipulate traffic flow if you don't know why it's going the way it is currently
- + General path selection order is...
 - + Weight
 - + Local preference
 - + AS-Path
 - + Origin
 - + MED
 - + EBGp over iBGP
 - + Lowest IGP metric to next-hop

Misc BGP Considerations

- + Communities
 - + Standard, extended, transitive, non-transitive
- + Filtering
 - + What methods does BGP support?
- + Convergence
 - + What affects BGP convergence?

