

BGP - Things to Remember

1. BGP uses TCP port 179 for transport. Router with the higher BGP router-id initiates BGP session from a random port. 2. The interface from which the BGP router ID is taken does not have to be running BGP. Any valid IP address can be used as BGP router-id, even an address that is not locally configured on the router. 3. The BGP router-id must be the same as the OSPF router-id for redistributing the routes from OSPF to BGP or vice versa. 4. If the *'network ?'* command is configured with the 'mask' option under the BGP process, then an exact match (network/mask) must exist in the IP routing table in order to advertise this route into BGP regardless of '*'auto-summary'* / '*'no auto-summary'*' command. But the *'network ?'* command configured without the 'mask' assumes the default classful mask and if '*'auto-summary'*' is configured then BGP will advertise a classful network only if any subnets of the classful network exist in the IP routing table. Again if the *'network ?'* command is configured without the 'mask' option and if '*'no auto-summary'*' is configured, then that router must have the exact classful network in the IP routing table in order to advertise it in BGP. 5. To accept and attempt BGP connections to the external peers residing on networks that are not directly connected, we need to use either '*neighbor ebgp-multihop ?'* or '*neighbor ttl-security ?'* command. These two commands are mutually exclusive. We can use another command '*neighbor disable-connected-check*' to accomplish the same task if the BGP neighbor is one-hop away. 6. The synchronization rule states that an iBGP learned prefix cannot be considered best unless there is a matching IGP route for that BGP prefix. BGP only advertises what it considers the best path. This issue can be resolved (1) by redistributing BGP routes into the IGP, (2) by creating a full-mesh of IBGP routers and disabling the synchronization, or (3) by creating a GRE tunnel. When BGP is synchronizing with OSPF, the router ID must match in both protocols in order to make it work. 7. When a prefix is received from an eBGP neighbor, it is advertised to both eBGP & iBGP neighbors. When a prefix is received from an iBGP neighbor, it is advertised ONLY to eBGP neighbors and not to any iBGP neighbors. To advertise iBGP learned routes to other iBGP peers requires the use of route-reflectors or confederations or a full-mesh of iBGP peers. 8. While sending BGP updates, EBGP peers modify the next-hop value to its own IP address. But iBGP peers do not modify it. 9. The *'default-information originate'* command, however, requires explicit redistribution of the route 0.0.0.0. Default routes can be injected into BGP in one of three ways: (1) using the '*network ?'* command (default route must exist in the local routing table), (2) using the '*default-information originate*' command (a redistribution statement must also be configured to redistribute the default route from the local routing table to the BGP table), and (3) using the '*neighbor ? default-originat [route-map route-map-name]*' command (this method does not even check for the existence of a default route in the IP routing table). The '*default-information originate*' command should not be configured with the '*neighbor ? default-originat*' command on the same router. 10. '*weight*' and '*local-preference*' are set inbound and they affect outbound traffic. But '*as-path*' and '*med*' are set outbound and they affect inbound traffic. 11. The weights assigned with the '*set weight ?*' route-map command overrides the weights assigned using the '*neighbor ? weight ?*' command. 12. Origin code 'i' is default on the BGP routes advertised by '*network ...', 'aggregate-address ...'* (if all subnet has 'i'), and '*neighbor ? default-originat*' commands. And origin code '?' is default on the BGP routes advertised by '*redistribute ...', 'aggregate-address ...'* (if any single subnet has '?', but can be changed using '*attribute-map*' option), '*default-information originate*', and '*bgp inject-map ...*' commands. 13. When BGP originates a route with the *'network ?'* command, MED is copied from the metric of the original route. 14. BGP MED values are not passed beyond the receiving (neighbor) AS. 15. Enabling the *'bgp deterministic-med'* command ensures the comparison of the MED variable when choosing routes advertised by different peers in the same autonomous system. Enabling the *'bgp always-compare-med'* command ensures the comparison of the MED for paths from neighbors in different autonomous systems. 16. The default behavior of BGP routers that run Cisco IOS software is to treat routes without the MED attribute as having a MED of 0, making the route that lacks the MED variable the most preferred. The '*bgp bestpath med missing-as-worst*' command can be configured to treat the route that missing MED as the least preferred one. 17. *'bgp bestpath as-path ignore'* is a hidden command in Cisco IOS which allows BGP to not consider the AS path during best path route selection. 18. There are two ways to create an aggregate address under BGP. The first is to create a static route to null interface in the routing table for the aggregate address and then advertise it with the *'network ?'* command. The second way is to use the *'aggregate-address ?'* command. 19. By default when aggregation is configured in BGP, the '*atomic-aggregate*' attribute is attached to the aggregate address if the '*as-set*' argument is not used in the *'aggregate-address ?'* command. The '*as-set*' argument reveals the AS numbers which can prevent a routing loop, and once '*as-set*' is configured along with the *'aggregate-address ?'* command, the '*atomic-aggregate*' attribute is automatically removed. 20. A router reflector and its clients are known collectively as a cluster. If the cluster contains a single route reflector, the cluster ID is the router ID of the route reflector. If the cluster contains multiple route reflectors, each RR must be manually configured with a cluster ID. 21. A client router in a route reflection cluster can peer with external neighbors, but the only internal neighbor it can peer with is a route reflector in its cluster or other clients in the cluster. Clients cannot peer with routers outside of their own cluster.

However, the RR itself can peer with both internal and external neighbors outside of the cluster and can reflect their routes to its clients. 22. In case of route reflection, (1) routes from EBGP are advertised to EBGP, client, non-client (2) routes from client are advertised to EBGP, client, non-client (3) routes from non-client are advertised to EBGP, client. 23. When the '*no bgp client-to-client reflection*' command is configured the RR does not reflect routes from one client to another. It does, however, continue to reflect routes from clients to peers outside of the cluster, and from peers outside of the cluster to clients. 24. Standard and extended BGP communities are removed from the reflected routes unless the '*neighbor ... send-community [both]*' is configured on the route reflector. The link bandwidth community is removed from reflected route if the route-reflector performs IBGP multipath load-sharing for that route. 25. The '*neighbor ? nexthop-self*' on router reflectors only affects the next hop of eBGP learned routes because the next hop of reflected routes should not be changed. To avoid a common configuration error for reflected routes, the '*set ip next-hop*' command should not be used in a route map to BGP route reflector clients. 26. Unlike route reflector environments in which only the route reflector itself has to support route reflection, all routers within a confederation must support the confederation functionality. 27. EBGP routes external to the confederation are preferred over EBGP routes to member autonomous systems, which are preferred over iBGP routes. 28. AS_PATH types are AS_SEQUENCE, AS_CONFED_SEQUENCE, AS_SET, and AS_CONFED_SET. AS_SEQUENCE is an ordered set of AS numbers, and AS_SET is an unordered set of AS numbers. AS_CONFED_SEQUENCE and AS_CONFED_SET are the same as AS_SEQUENCE and AS_SET but are used only within BGP confederations. 29. When '*bgp bestpath med confed*' command is configured, the router picks the confederation-internal path with the lowest MED and ignores the path with the external AS number. 30. BGP private autonomous system numbers are from 64,512 to 65,535 31. BGP prefixes can be filtered using (1) 'distribute-list', (2) 'prefix-list', (3) 'filter-list', (4) 'policy-list', (5) community/extended community lists, (6) 'route-map'. 32. For BGP, the '*distance ?*' command sets the administrative distance of the External BGP (eBGP) route. This command only affects the routing table and not the BGP table. 33. The '*network ? backdoor*' command has the same effect as the '*network ?*' command. The EBGP route is treated as a local BGP route, and the administrative distance is changed to 200. The difference is that the address specified by the network backdoor command is not advertised to EBGP peers. 34. iBGP routes are not redistributed into an IGP unless you use "*bgp redistribute-internal*" command under BGP routing process. 35. '*bgp inject-map ... exist-map ...*' command injects prefixes in the local BGP RIB when a valid parent route exists. Only prefixes that are equal to or more specific than the aggregate route (existing prefix) can be injected. exist-map (route-map) must contain a '*match ip address prefix-list ...*' command statement to specify the aggregate prefix and a '*match ip route-source prefix-list ...*' command statement to specify the route source. If the parent route is a default route, we can inject any route out of it. 36. A BGP neighbor cannot be configured to work with both peer groups and peer templates. BGP peer templates and BGP peer groups are mutually exclusive. 37. Peer session template can inherit only one session template directly, but peer policy template can inherit multiple policy templates. 38. When the maximum number (as set by the '*neighbor ? maximum-prefix ...*' command) of prefixes are reached, the string "PfxRcd" appears in the entry, the neighbor goes to shutdown state, and the connection becomes idle. 39. No penalty is applied to a BGP peer reset when route dampening is enabled. Although the reset withdraws the route, no penalty is applied in this instance. 39. In case of iBGP multipath load sharing, when multiple iBGP paths installed in a routing table, a route reflector will advertise only one of the paths (one next hop). 40. For multiple paths to the same destination to be considered as multipaths, all attributes including weight, local preference, autonomous system path (entire attribute and not just length), origin code, MED, and IGP distance must be same. But if '*bgp bestpath as-path multipath-relax*' command is configured, the AS paths still have to be the same length, but don't have to be identical. 41. Though BGP Multipath allows the installation of multiple BGP paths (for load sharing purpose) into the IP routing table for the same prefix, it does not affect the bestpath selection. A router still designates one of the paths as the best path and advertises this best path to its neighbors. 42. '*neighbor ? dmzlink-bw*' command can be used with eBGP and iBGP multipath features to enable unequal cost load balancing over multiple links. BGP can originate the link bandwidth community only for directly connected links to eBGP neighbors. 43. The '*bgp update-delay ...*' command is used to tune the maximum time the software will wait after the first neighbor is established until it starts calculating best paths and sending out advertisements. 44. The '*neighbor ? local-as ??*' command is valid only if the peer is a true eBGP peer. It does not work for two peers in different sub-ASs in a confederation. 45. In a route-map, a continue clause can be executed, without a successful match, if a route map entry does not contain a match clause. But if a match clause exists, the continue clause is executed only if a match occurs. If no successful matches occur, the continue clause is ignored. The continue statement proceeds to the specified route map entry only after configured set actions (if any) are performed. 46. When multiple values are configured in the same community list statement, a logical AND condition is created. All community values must match to satisfy an AND condition. When multiple values are configured in separate community list statements, a logical OR condition is created. The first list that matches a condition is processed. 47. While redistributing OSPF into BGP, by default only

OSPF intra-area and inter-area routes are redistributed into BGP. 48. When a BGP router with synchronization enabled has also a OSPF route (redistributed from BGP) for a iBGP-learned route, then the OSPF ASBR router-id must match the originating BGP router-id in order to synchronize BGP route with OSPF route. 49. An ?update group? is a group of peers with a common outbound policy which will be converged as if they are in a peer-group. By Zakir A. Khan * BGP - Things to Remember . Pdf | Download