

OSPF - Things to Remember

1. The IP header of an OSPF packet specifies protocol number 89. 2. To establish OSPF neighbor adjacency, hello/dead timers, MTU (otherwise have to use "*ip ospf mtu-ignore*") must match. Unique router-id is also required. 3. Routers in stub area can only be adjacent with the routers in stubs or totally stubby area. Routers in NSSA can only be adjacent with the routers in NSSA or totally NSSA. 4. OSPF sees secondary networks as stub networks and cannot make adjacencies over secondary addresses. OSPF will advertise a secondary network or subnet only if it is also running on the primary network or subnet and OSPF routes of secondary addresses must be in same area as the primary address to be advertised. To learn routes from a neighbor connected to the secondary network, another routing protocol such as RIP should be running and redistributed into OSPF. Another solution to this kind of problem is to create dot1q subinterfaces. 5. The only time that OSPF will form adjacencies between neighbors that are not on the same subnet is when the neighbors are connected through point-to-point links using "ip unnumbered". 6. The primary interface and IP unnumbered interface will have OSPF enabled if a network statement matches the IP address of the primary interface. 7. An OSPF external route cannot use another OSPF external route as its next hop. 8. Inside an area, OSPF uses Link State logic, but between areas OSPF acts much like a Distance Vector (DV) protocol in some regard. For example, the advertisement of a Type 3 LSA from one area to another hides the topology in the original area from the second area, just listing a destination subnet, metric (cost), and the ABR through which the subnet can be reached?all DV concepts. 9. Only broadcast and non-broadcast network elect DR/BDR based on priority or router-id (in case of a tie in the priority). 10. In non-broadcast network, DR/BDR must have layer 2 connectivity to all other routers in the same area. 11. With OSPF network types broadcast and non-broadcast, next hop values are not modified when updates are transmitted across an NBMA media. Both point-to-multipoint and point-to-multipoint non-broadcast network type update the next-hop value of routes learned on partially meshed networks to the directly connected neighbor, and advertise the network as a set of endpoints instead of a transit network. 12. OSPF network point-to-point is the default option for point-to-point interfaces such as HDLC, PPP, or point-to-point NBMA subinterfaces. 13. As only broadcast and non-broadcast network type elects DR/BDR, they are compatible with each other, but they are not compatible with any other network types. 14. OSPF cost can be modified using (i) interface "*bandwidth ...*" command, (ii) interface "*ip ospf cost ...*" command, (iii) process "*auto-cost reference-bandwidth ...*" command, or (iv) "*neighbor ... cost ...*" command on point-to-multipoint non-broadcast network. 15. Only OSPF point-to-multipoint and point-to-multipoint non-broadcast network types support OSPF cost value on a per neighbor basis. On point-to-multipoint broadcast networks, if the "*neighbor...*" command is used, a cost to that neighbor must be specified. But on point-to-multipoint non-broadcast networks, the "*neighbor ...*" command must be used to identify neighbors, assigning a cost to a neighbor is optional. 16. The internal OSPF routes can only be summarized on ABRs whereas the external (redistributed) routes can only be summarized on ASBRs. 17. "*area default-cost ...*" command is used to specify a cost for the default summary route (default cost 1) that is sent into a stub area or NSSA. 18. In NSSA, ABR with the highest router-id does the LSA 7 to 5 conversion. 19. In NSSA, ?*default-information originate*? command cannot be used, since it generates Type-5 LSA, which is prohibited in NSSA area. 20. NSSA ASBR can generate a default only when it has a default route in its routing table whereas NSSA ABR can generate a default route with or without a default route in its own routing table. 21. Virtual links are not allowed in the stubby area or NSSA. In this case OSPF can be tunneled over a stub area using GRE tunnel (tunnel must be connected to area 0). 22. If the authentication is wrong on the virtual-link, the virtual-link interface will not go down immediately. As the virtual-link does not support periodic hellos, ?*clear ip ospf process*? command should be issued if the authentication is enabled on the virtual link. 23. The virtual link will not come up if the only interface to reach the other end of the virtual link has a cost that is maximized (65535). 24. For BGP to redistribute routes into OSPF, the router-id must be identical, in OSPF and in BGP. 25. OSPF filtering using "*distribute-list ...*", "*route-map ...*" (match route-type, match ip route-source, match ip next-hop), and "*distance ...*" commands can only block route from entering into local RIB, but cannot stop LSAs propagation into the OSPF database. 26. OSPF filtering using "*area ... filter-list prefix ...*", "*area ... range ... not-adv*", ?*summary-address ? not-adv?*, ?*ip ospf database-filter all out?*, or ?*neighbor ? database-filter all out?* commands can filter LSAs from OSPF database. 27. If the ?*area ? range ??* and "*area ... filter-list prefix ... out*" both commands are configured for an area, then type 3 LSAs that correspond to the area range are sent to all other areas, only if at least one prefix in the area range matches an entry in the prefix list. 28. OSPF defaults to cost 20 when redistributing from an IGP, and 1 when redistributing from BGP. 29. ?*neighbor ? database-filter all out?* only works on point-to-multipoint network types. 30. If ?*distribute-list out?* command is configured on an ASBR, then the ASBR generates Type 5 external LSAs only for those networks that are explicitly permitted in the distribute list. 31. OSPF demand circuit sets ?do not age? flag on all LSAs learned and will only send updates when there is a change in the OSPF topology. The command must be configured in a point-to-point link and is needed only on one side. If the router is part of a point-to-multipoint topology,

only the multipoint end must be configured with this command. 32. The main difference between flooding reduction ("*ip ospf flood-reduction*") and demand circuits ("*ip ospf demand-circuit*") is that former suppresses only periodic LSA refreshes; it does not suppress periodic hello packets. Thus, the flooding reduction feature does not impair the detection of a neighbor router going down. 33. OSPF stub router (*?max-metric router-lsa?*) advertises all non self-originated routes/LSAs with maximum metric. 34. When "*redistribute maximum-prefix ...*" command is configured, the redistribution limit does not apply to default routes or prefixes that are generated as a result of Type-7 to Type-5 translation. *OSPF - Things to remember . Pdf | Download By Zakir A. Khan