

CCNA Quick Notes - Routing

1.How do OSPF-enabled routers build adjacencies and exchange their routing tables? OSPF-enabled routers build adjacencies by sending Hello packets out through all OSPF-enabled interfaces. If these routers share a common link and agree on parameters set within their Hello packets then they become neighbors. If these parameters differ then the routers do not become neighbors and communication stops. OSPF routers form adjacencies with certain routers. These routers are determined by the layer 2 (data link) media type and as soon as the adjacencies are formed each router sends LSAs (Link State Advertisements) to all adjacent routers. The LSAs describe the state of each router's links. There are multiple LSA types and a route that receives an LSA from a neighbor records the LSA in a link-state database and floods a copy of the LSA to all its other neighbors. When all databases are complete - then each router uses the SPF (Shortest-Path First) algorithm to calculate a loop-free topology and builds its routing table based on this topology. It is important to note that the Hello protocol is bidirectional and is the means by which neighbors are discovered and acts as keepalives between neighboring routers. It also establishes and maintains neighbor relationships and elects the DR (Designated Router) and BDR (Backup Designated Router) to represent the segment on Broadcast and NBMA (nonbroadcast multiaccess) networks. Note: Hello protocols are sent periodically out each OSPF-enabled interface using IP multicast address 224.0.0.5. The default interval on NBMA (nonbroadcast multiaccess) networks is 30 seconds. The default interval on Broadcast, Point-to-point, and point-to-multipoint networks is 10 seconds.

2.What are LSAs (link-state advertisements)? LSAs are sent out all OSPF-enabled router interfaces describing the state of the router's links. They are also packets that OSPF uses to advertise changes in the condition of a link or other OSPF routers.

3.Name two LSA (link-state advertisement) types? Type 1 LSAs are router LSAs and are generated by each router for the area to which the router belongs. These LSAs describe the states of the router's links to the area (area 0 for example) and are flooded within a single area (area 0 for example). Type 2 LSAs are network LSAs and are generated by the DR (Designated Router) and the BDR (Backup Designated Router). They describe the routers attached to a particular network and are flooded within a single area (area 0 for example).

4.What is the routing metric OSPF is based on? Bandwidth. Formula: Cost = $100,000,000 / \text{bandwidth in bits per seconds}$ The cost of a 100 MBbps connection would be: $1000,000,000 / 100,000,000 = 1$

5.Based on the schema above -- if adjacencies are established with only the DR (Designated Router) and BDR (Backup Designated Router)- what is the circuit count? Formula: $2(n - 1)$ where n is the number of routers in the network. $2(5 - 1) = 8$ circuits. A circuit can also be thought of as an adjacency or connection. Count four going into the DR and 4 going into the BDR for a total of 8.

Note: OSPF avoids synchronizing between every pair of routers in the network by using a DR and BDR. This way adjacencies are formed only to the DR and BDR, and the number of LSAs sent over the network is reduced. Now only the DR and BDR have four adjacencies, and all the other routers have two.

6.On an OSPF-enabled router -- what is the router ID and where does an OSPF-enabled router receive its router ID? To initialize - OSPF must be able to define a router ID. The most common and stable source for a router ID is the IP address set on the logical loopback interface that is always available. If no logical interface is defined -- then the router receives its ID from the highest IP address on the physical interfaces. Note: If two loopback addresses are defined -- it will use the highest loopback address. Think highest logical or highest physical interface address.

Name five OSPF network types: Broadcast networks: Ethernet/TOKEN Ring. OSPF-enabled routers on broadcast networks elect a DR (Designated Router) and BDR (Backup Designated Router). All the routers in the network form adjacencies with the DR and BDR. Note: OSPF packets are multicast to the DR and BDR.

NBMA (nonbroadcast multiaccess) networks: Frame Relay/X.25/ATM. NBMA networks can connect more than two routers but have no broadcast functionality. These networks elect an DR and BDR. Note: OSPF packets are unicast.

Point-to-point networks: A physical DS1 (T1) for example. Point-to-point networks connect a pair of routers and always becomes adjacent.

Point-to-multipoint networks: Point-to-multipoint networks are a special configuration of NBMA networks in which networks are treated as a collection of point-to-point links. Point-to-multipoint networks do not elect a DR or BDR. Note: OSPF packets are multicast.

Virtual links: Virtual links are a special configuration that the router interprets as unnumbered point-to-point networks. The network administrator creates/defines virtual links.

7.What is routing? Routing is the process in which items are forwarded from one location to another. Routing is a hop-by-hop paradigm. A Cisco router performs routing and switching functions. Describe what each function does. Routing is a way to learn and maintain awareness of the network topology. Each router maintains a routing table in which it looks up the destination Layer 3 address to get the packet one step closer to its destination. The switching function is the actual movement of temporary traffic through the router, from an inbound interface to an outbound interface.

8.What are the three types of routes you can use in a Cisco router? The three types of routes are static routes, dynamic routes, and default routes.

9.What is the difference between static and dynamic routes? Static routes are routes that an administrator manually enters into a router. Dynamic routes are routes that a router learns automatically through a routing protocol.

10.How do you configure a static route on

a Cisco router? To configure a static route on a Cisco router, enter the ip route destination-network [mask] {next-hop-address outbound-interface} [distance] [permanent] global command. Here's an example: RouterB(config)#ip route 172.17.0.0 255.255.0.0 172.16.0.1 This example instructs the router to route to 172.16.0.1 any packets that have a destination of 172.17.0.0 to 172.17.255.255 11.What is a default route? Also known as the gateway of last resort, a default route is a special type of static route with an all-zeros network and network mask. The default route is used to route any packets to a network that a router does not directly know about to a next-hop router. By default, if a router receives a packet to a destination network that is not in its routing table, it drops the packet. When a default route is specified, the router does not drop the packet. Instead, it forwards the packet to the IP address specified in the default route. How do you configure the default route on a Cisco router? To configure a default route on a Cisco router, enter the following global configuration command: ip route 0.0.0.0 0.0.0.0 [ip-address of the next-hop router outbound-interface] For example: RouterB(config)#ip route 0.0.0.0 0.0.0.0 172.16.0.2 12.What is a routing protocol? A routing protocol defines the set of rules used by a router when it communicates with neighboring routers. Routing protocols listen for packets from other participants in order to learn and maintain a routing table. What are the two major types of routing protocols? The two major types of routing protocols are

- Interior Gateway Protocol (IGP)
- Exterior Gateway Protocol (EGP)

IGP is used to exchange routing information among routers in the same autonomous system (AS). EGP is used to communicate between ASs. Note: For more information about autonomous systems, see Interconnecting Cisco Network Device (Cisco Press).

13.What is administrative distance? Administrative distance (AD) is an integer from 0 to 255 that rates the trustworthiness of routing information received on a router from a neighboring router. The AD is used as the tiebreaker when a router has multiple paths from different routing protocols to the same destination. The path with the lower AD is the one given priority. 14.What are the three classes of routing protocols? The three classes of routing protocols are

- Distance vector
- Link-state
- Balanced hybrid

15.What is the AD for each of the following?

- Directly connected interface
- Static route
- EIGRP
- IGRP
- OSPF
- RIP
- External EIGRP
- Unknown

The ADs are as follows:

- Directly connected interface 0
- Static route 1
- EIGRP 90
- IGRP 100
- OSPF 110
- RIP 120
- External EIGRP 170
- Unknown 255

16.How do distance vector routing protocols function? Also known as Bellman-Ford-Fulkerson algorithms, distance vector routing protocols pass complete routing tables to neighboring routers. Neighboring routers then combine the received routing table with their own routing table. Each router receives a routing table from its directly connected neighbor. Distance vector routing tables include information about the total cost and the logical address of the first router on the path to each network they know about. 17.How do distance vector routing protocols keep track of any changes to the internetwork? Distance vector routing protocols keep track of an internetwork by periodically broadcasting updates out all active interfaces. This broadcast contains the entire routing table. This method is often called routing by rumor. Slow convergence of distance vector routing protocols can cause inconsistent routing tables and routing loops. 18.What are some mechanisms that distance vector protocols implement to prevent routing loops and inconsistent routing tables? Here are some of the ways distance vector routing protocols prevent routing loops and inconsistent routing tables:

- Maximum hop count
- Split horizon
- Route poisoning
- Holddowns

19.What is maximum hop count? If a loop is in an internetwork, a packet loops around the internetwork forever. Maximum hop counts prevent routing loops by defining the maximum number of times a packet will loop around the internetwork. RIP uses a hop count of up to 15, so anything that requires 16 hops is unreachable. Anytime a packet passes through a router, it is considered one hop. 20.What is split horizon? The rule of split horizon is that it is never useful to send information about a route back in the direction from which the original update came. 21.What is convergence? Convergence is when all routers have consistent knowledge and correct routing tables. 22.What is route poisoning? With route poisoning, when a

distance vector routing protocol notices that a route is no longer valid, the route is advertised with an infinite metric, signifying that the route is bad. In RIP, a metric of 16 is used to signify infinity. Route poisoning is used with holddowns. 23.What are hold-down timers? Hold-down timers prevent regular update messages from reinstating a route that might have gone bad. Hold-down timers also tell routers to hold for a period of time any changes that might affect routes. 24.What are triggered updates? When a router notices that a directly connected subnet has changed state, it immediately sends another routing update out its other interfaces rather than waiting for the routing update timer to expire. Triggered updates are also known as Flash updates. 25.What is IP RIP? IP RIP is a true distance vector routing protocol that sends its complete routing table out all active interfaces every 30 seconds. IP RIP uses a hop count as its metric to determine the best path to a remote network. The maximum allowable hop count is 15, meaning that 16 is unreachable. There are two versions of RIP. Version 1 is classful, and version 2 is classless. IP RIP can load-balance over as many as six equal-cost paths. 26.What four timers does IP RIP use to regulate its performance? Here are the four timers that IP RIP uses to regulate its performance:

- Route update timer Time between router updates. The default is 30 seconds.
- Route invalid timer Time that must expire before a route becomes invalid. The default is 180 seconds.
- Route hold-down timer If IP RIP receives an update with a hop count higher than the metric recorded in the routing table, the router goes into holddown for 180 seconds.

Route flush timer Time from when a route becomes invalid to when it is removed from the routing table. The default is 240 seconds.

27.How do you enable RIP on a Cisco router? To enable RIP on a Cisco router, start by using the router global configuration command, followed by the rip protocol. This selects RIP as the routing protocol. Then you assign the network command, followed by the directly connected network number(s) you want to activate RIP on. Here's an example: RouterB(config)#router rip RouterB(config-router)#network 192.168.1.0 RouterB(config-router)#network 192.168.2.0 28.How do you stop RIP updates from propagating out an interface on a router? Sometimes you do not want RIP updates to propagate across the WAN, wasting valuable bandwidth or giving out valuable information about your internetwork. The easiest way to stop RIP updates from propagating out an interface is to use the passive-interface global configuration command. How do you display the contents of a Cisco IP routing table? The show ip route command displays the Cisco routing table's contents. 29.What is Interior Gateway Routing Protocol (IGRP)? IGRP is a Cisco proprietary distance vector routing protocol. IGRP has a default hop count of 100 hops, with a maximum hop count of 255. IGRP uses bandwidth and line delay as its default metric, but it can also use reliability, load, and MTU. 30.How do you enable IGRP on a Cisco router? The way you enable IGRP on a Cisco router is similar to the way you enable RIP, except you specify IGRP as the protocol and add an autonomous system number. For example:

RouterA(config)#router igrp 10 (10 is the AS number) RouterA(config-router)#network 192.168.0.0

RouterA(config-router)#network 192.168.1.0 RouterA(config-router)#network 192.168.2.0 31.What four timers does IGRP use to regulate its performance? The four timers IGRP uses to regulate its performance are as follows: · Route update timer Time between router updates The default is 90 seconds. · Route invalid timer Time that must expire before a route becomes invalid . The default is 270 seconds. · Route hold-down timer If a destination becomes unreachable, or if the next-hop router increases the metric recording in the routing table, the router goes into holddown for 280 seconds. · Route flush timer[md]Time from when a route becomes invalid to when it is removed from the routing table. The default is 630 seconds. [PDF Version](#) | [Download](#)