

## Chapter 3

### Determining IP Routes

Routers perform two main functions: switching and routing. The switching function is the process of moving packets from an inbound interface to an outbound interface. The switching function is also responsible for stripping the data link layer information from a packet it receives and encapsulating it with the data link layer information (MAC address) of the router's exiting interface. Routing is a relay function in which packets are forwarded from one location to another. The routing mechanism is responsible for learning and maintaining awareness of the network topology. A router functions as a hop-by-hop paradigm and performs best-effort packet delivery. The packet is delivered to the downstream router that the sending router feels is closest to the final destination. For a router to be an effective relay device, it must perform both routing and switching functions.

There are two ways a router learns about routes a packet must take: statically and dynamically. Static routes are entered manually by a network administrator. Dynamic routes are learned by a routing protocol. Routing protocols can be classified into two categories: interior and exterior. Interior routing protocols learn about routes and route packets within an autonomous system (AS). Exterior routing protocols learn about routes and route packets between autonomous systems. Routing protocols use metrics to determine what path is best for a packet to travel. A metric is a standard measurement, such as distance, used by routing protocols to determine the optimum path to a destination. Each routing protocol uses a different algorithm to determine its metric. This chapter focuses on the interior routing protocols covered on the ICND and CCNA exams.

## Distance Vector Routing Protocols

RIP and IGRP are distance vector routing protocols. The name is derived from the fact that routes are advertised as vectors of distance and direction. Routers running distance vector protocols learn routes from neighboring routers' perspectives and then advertise the routes from their own perspective. Because each router depends on its neighbors for information, distance vector routing is sometimes called "routing by rumor." Routers running distance vector protocols store the route that has the best administrative distance and metric in its routing table. RIP and IGRP are also classful routing protocols. Classful routing protocols exchange routes to subnetworks in the same major network (Class A, B, or C). In other words, all networks in the internetwork must have the same subnet mask. This is because RIP and IGRP do not send the subnet mask in their routing updates, thus assuming that all networks use the same subnet mask assigned to the exiting interface.

## Link-State Routing Protocols

Also known as shortest path first (SPF) protocols, link-state routing protocols learn about every router in an AS and flood information about the state of their directly connected links to each router in the AS. A link is an interface on the router. The link's state is a description of that interface. Link-state protocols use this information to build a link-state database that gives the router a picture of the network. Link-state protocols have fast convergence because of the link-state database. On the other hand, they require more CPU power and memory. OSPF is a very scalable link-state protocol. It is the interior routing protocol of choice in large internetworks. OSPF is scalable because larger networks can be divided into areas. Generally, each router in an area is concerned with the link state of the other routers in its area; it does not store the link

states of routers in different areas in its link-state database, thus conserving resources. OSPF is considered a classless routing protocol; it sends information about the interface's subnet mask in its updates. The ICND and CCNA exams cover OSPF in a single area. OSPF in multiple areas is covered on the CCNP exams.

## Hybrid Routing Protocols

Enhanced IGRP (EIGRP) is a hybrid routing protocol developed by Cisco. EIGRP makes decisions based on metric information it receives from neighbors, and it employs some of the mechanisms of link-state protocols. It establishes adjacencies with neighboring routers, sends only changes in its updates, and keeps a topology database of the network. EIGRP uses Diffusing Update Algorithm (DUAL) to achieve rapid convergence, conserve bandwidth, and support multiple network protocols.

### Question 1

**What must a router have to correctly route a packet?**

### Question 2

**What are the three types of routes you can use in a Cisco router?**

### Question 1 Answer

A router needs the network address of the destination network and a routing table with valid information about the destination to correctly route a packet.

### Question 2 Answer

The three types of routes are static, dynamic, and default.

### Question 3

**What is the difference between static and dynamic routes?**

### Question 4

**How do you configure a static route on a Cisco router?**

### Question 3 Answer

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.

### Question 4 Answer

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.

### Question 5

**What is a default route?**

Determining IP  
Routes

### Question 6

**How do you configure the default route on a Cisco router?**

Determining IP  
Routes



## Question 5 Answer

Also known as the gateway of last resort, a default route is a special type of static route with an all-0s network and network mask. The default route is used to direct any packets for which a next hop is not specifically listed in the routing table. 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.

## Question 6 Answer

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
```

### NOTE

All Cisco routers before Cisco IOS software Release 12.0 are classful. They expect a default subnet mask on each interface of the router. If a router receives a packet for a destination subnet not in the routing table, the router forwards the packet to the best available major Class A, B, or C network. If the router does not have specific information about the major network, it drops the packet. When you are using an IOS version before Release 12.0 and are using a default route, you must use the **ip classless** global command so that the packets go to the default route.

### Question 7

**What is administrative distance?**

Determining IP  
Routes

### Question 8

**What are the three classes of routing protocols?**

Determining IP  
Routes

## Question 7 Answer

Administrative distance (AD) is an integer from 0 to 255 that rates the trustworthiness of the source of the IP routing information. It is important only when a router learns about a destination route from more than one source. The path with the lower AD is the one given priority.

### NOTE

If a routing protocol has multiple paths within the same routing protocol to the same destination, the metric is used as the tiebreaker. The route with the lowest metric is the path taken.

## Question 8 Answer

The three classes of routing protocols are

- Distance vector
- Link-state
- Balanced hybrid

Distance vector protocols use a vector of distance and direction to find the best path, where distance is defined in terms of a metric and direction is defined as the next-hop router. Distance vector protocols broadcast the entire routing table at periodic intervals and are slow to converge because of hold-down timers. Examples include RIP and IGRP.

Link-state protocols use a topological database that is created on each router. This database keeps track of directly attached neighbors, the entire network, and the routing table. Link-state updates are typically multicast to all neighbors. (OSPF is a link-state protocol.)

Balanced hybrid protocols combine aspects of distance vector and link-state protocols. An example is EIGRP.

**Question 9**

**What is the AD for each of the following?**

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

**Question 10**

**Name two classful routing protocols.**

### Question 9 Answer

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

### Question 10 Answer

Two classful routing protocols are RIP and IGRP. Classful routing protocols require all interfaces to have the same subnet mask. Classful routing protocols do not include the subnet mask in their routing updates. Thus, they assume that all remote networks have the subnet mask of the exiting interface.

**Question 11**

**For different VLANs to communicate with each other, they need to be routed. To perform inter-VLAN routing, what two things must occur?**

**Question 12**

**How do you route between VLANs with ISL trunks?**

## Question 11 Answer

To perform inter-VLAN routing, the following must occur:

- The router must know how to reach all VLANs being interconnected.
- The router must have a separate physical connection on the router for each VLAN, or trunking must be enabled on a single physical connection.

### NOTE

Remember that each VLAN is a subnet and that a router is needed to route from one subnet to another.

## Question 12 Answer

To route between VLANs using ISL on a Cisco router, you must do the following:

- Step 1** Create a subinterface for each VLAN to be routed.
- Step 2** Enable ISL encapsulation on each subinterface being configured. (ISL trunking works only on Fast Ethernet and Gigabit Ethernet interfaces.)
- Step 3** Assign an IP address to the subinterface.
- Step 4** Enable ISL encapsulation for the particular VLAN with the **encapsulate isl *vlan#*** subinterface command.

Here's an example:

```
RouterB(config)#int f0/0.10
RouterB(config-subif)#encapsulate isl 10
RouterB(config-subif)#ip address 172.16.0.1 255.255.0.0
```

**Question 13**

**What commands enable routing between VLANs with 802.1Q trunks?**

**Question 14**

**How do distance vector routing protocols function?**



## Question 13 Answer

The **encapsulation dot1q *vlan-id*** command enables 802.1Q on a Cisco router.

The native VLAN in 802.1Q does not carry a tag; therefore, a trunk's major interface has an IP address. Any other configuration information for the native VLAN subinterfaces is configured with the **dot1q** encapsulation and IP address. Remember that the major interface of a router using ISL cannot have an IP address. The following example enables 802.1Q encapsulation on a router:

```
RouterB(config)#int f0/0
RouterB(config-if)#ip address 192.168.10.1 255.255.255.0
RouterB(config-if)#encapsulation dot1q 10
RouterB(config-if)#int f0/0.20
RouterB(config-if)#ip address 192.168.20.1 255.255.255.0
RouterB(config-if)#encapsulation dot1q 20
```

### NOTE

To enable 802.1Q on a Catalyst 1900 interface, use the **switchport trunk encapsulation dot1q** command. Catalyst 2950 switches support only 802.1Q encapsulation and are enabled with the **switchport mode trunk interface** command.

## Question 14 Answer

Also known as Bellman-Ford 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.

**Question 15**

**How do distance vector routing protocols keep track of any changes to the internetwork?**

**Question 16**

**Slow convergence of distance vector routing protocols can cause inconsistent routing tables and routing loops. What are some mechanisms that distance vector protocols implement to prevent these problems?**

### Question 15 Answer

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.”

### Question 16 Answer

Here are some of the ways distance vector routing protocols prevent routing loops and inconsistent routing tables:

- Maximum hop count (count to infinity)
- Split horizon
- Route poisoning
- Holddowns
- TTL

### Question 17

**What is maximum hop count?**

Determining IP  
Routes

### Question 18

**What is split horizon?**

Determining IP  
Routes

### Question 17 Answer

If a loop is in an internetwork, a packet loops around the internetwork until the TTL in the IP packet reaches 0 and is removed. Maximum hop counts prevent routing loops by defining the maximum number of times a packet loops 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.

### Question 18 Answer

The split-horizon rule prohibits a router from advertising a route through an interface that the router itself is using to reach the destination.

### Question 19

**What is convergence?**

Determining IP  
Routes

### Question 20

**What is route poisoning?**

Determining IP  
Routes

### Question 19 Answer

Convergence is when all routers have consistent knowledge and correct routing tables.

### Question 20 Answer

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.

**NOTE**

A poison reverse is an update that a router sends to the router it received the route poison from, specifying that all routers on the segment have received the poisoned route information.

### Question 21

**What are hold-down timers?**

### Question 22

**What are triggered updates?**



### Question 21 Answer

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.

### Question 22 Answer

Also known as flash updates, triggered updates are routing updates sent immediately out a router's interface when it notices that a directly connected subnet has changed state.

**Question 23**

Which of the following are used for loop avoidance?

- A. Link-state advertisements
- B. Poison reverse
- C. Route discovery
- D. Split horizon

**Question 24**

How is routing information maintained with link-state routing protocols?

### Question 23 Answer

B and D. Poison reverse and split horizon are used for loop avoidance. OSPF uses link-state advertisements to advertise its links. Route discovery is the process of discovering all available routes.

---

### Question 24 Answer

Link-state routing protocols use link-state advertisements (LSAs), a database describing the entire area, and the shortest path first (SPF) algorithm to maintain routing information within an internetwork.

**Question 25**

**What are the functions of areas and autonomous systems in link-state protocols?**

**Question 26**

**What are four advantages that link-state protocols have over distance vector protocols?**

## Question 25 Answer

Areas are a grouping of contiguous networks. They are logical subdivisions of an autonomous system.

### NOTE

An autonomous system is a collection of networks under a common administration that share a common routing strategy. An autonomous system is sometimes called a domain. It can be logically subdivided into multiple areas to reduce routing updates.

## Question 26 Answer

- Link-state protocols send routing updates only when they detect a topology change.
- Fast convergence.
- Support for classless addressing.
- Networks can be segmented into area hierarchies, limiting where routing updates are flooded to.

### Question 27

List three disadvantages of link-state protocols as compared to distance vector protocols.

### Question 28

What is balanced hybrid routing? Give an example of a balanced hybrid routing protocol.

### Question 27 Answer

- Significant demands on resources. Because link-state protocols require a topology database of the internetwork, they require a significant amount of memory and CPU cycles to run the SPF algorithm.
- Link-state protocol networks are more complex, making them more difficult to troubleshoot than distance vector protocols.
- All areas must connect to a backbone area. Therefore, implementing a link-state network requires much planning.

### Question 28 Answer

Balanced hybrid routing protocols combine aspects of distance vector and link-state protocols. Balanced hybrid routing protocols use distance vectors, which are more accurate, to determine the best path to a destination network, and they use topology changes to trigger routing updates.

EIGRP is a balanced hybrid protocol that is Cisco-proprietary.

### Question 29

**What is IP RIP?**

Determining IP  
Routes

### Question 30

**What four timers does IP RIP use to regulate its performance?**

Determining IP  
Routes



## Question 29 Answer

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. RIP has two versions. Version 1 is classful, and version 2 is classless. IP RIP can load-balance over as many as six equal-cost paths.

## Question 30 Answer

The four timers that IP RIP uses to regulate its performance are

- **Route update timer**—The time between router updates. The default is 30 seconds.
- **Route invalid timer**—The 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**—The time from when a route becomes invalid to when it is removed from the routing table. The default is 240 seconds.

**Question 31**

**What are the commands to enable RIP on a Cisco router?**

**Question 32**

**How do you stop RIP updates from propagating out an interface on a router?**

## Question 31 Answer

The commands to enable RIP on a Cisco router are

- **router rip**
- **network** *connected-network-address*

For example, the following commands enable RIP and advertise networks 192.168.1.0 and 192.168.2.0:

```
RouterB(config)#router rip  
RouterB(config-router)#network 192.168.1.0  
RouterB(config-router)#network 192.168.2.0
```

## Question 32 Answer

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** router configuration command.

**Question 33**

**Which Cisco IOS command displays values associated with routing timers, the administrative distance, and network information associated with the entire router?**

**Question 34**

**What commands enable RIP on a router and advertise network 172.16.0.0?**

## Question 33 Answer

The IOS command **show ip protocols** displays values associated with routing timers, the administrative distance, and network information associated with the entire router:

```
RouterB#show ip protocols
Routing Protocol is "rip"
  Sending updates every 30 seconds, next due in 2 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Outgoing update filter list for all interfaces is
  Incoming update filter list for all interfaces is
  Redistributing: rip
  Default version control: send version 1, receive any version
    Interface          Send Recv  Key-chain
    Serial0             1     1 2
    Serial1             1     1 2
  Routing for Networks:
    192.168.1.0
    192.168.2.0
  Routing Information Sources:
    Gateway           Distance      Last Update
  Distance: (default is 120)
```

## Question 34 Answer

```
router rip
network 172.16.0.0
```

**Question 35**

**How do you display the contents of a Cisco IP routing table?**

**Question 36**

**What IOS command displays RIP routing updates as they are sent?**

## Question 35 Answer

`show ip route` displays the Cisco routing table's contents:

```
RouterA#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile,
       B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter
       area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external
       type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2,
       E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * -
       candidate default
       U - per-user static route, o - ODR

Gateway of last resort is 192.168.1.1 to network 0.0.0.0

    1.0.0.0/32 is subnetted, 1 subnets
C       1.1.1.1 is directly connected, Loopback0
R       192.168.0.0/24 [120/1] via 192.168.1.1, 00:00:21, Serial0
C       192.168.1.0/24 is directly connected, Serial0
C       192.168.2.0/24 is directly connected, Ethernet0
R*      0.0.0.0/0 [120/1] via 192.168.1.1, 00:00:21, Serial0
```

[120/1] indicates that 120 is the AD and 1 is the number of hops to the remote network.

## Question 36 Answer

`debug ip rip` displays routing updates as they are sent and received.

### NOTE

Use the `no debug all` IOS command to turn off all debugging commands.

### Question 37

After enabling IP RIP debugging on your router, you see the following message: “RIP:broadcasting general request on Ethernet1”. What causes this message?

### Question 38

What is Interior Gateway Routing Protocol (IGRP)?



### Question 37 Answer

“RIP:broadcasting general request on Ethernet1” is usually displayed upon startup or if a user manually clears the routing table.

### Question 38 Answer

IGRP is a Cisco-proprietary distance vector routing protocol. It 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.

### Question 39

**What four timers does IGRP use to regulate its performance?**

### Question 40

**How many paths can IGRP support?**

### Question 39 Answer

The four timers that IGRP uses to regulate its performance are as follows:

- **Route update timer**—The time between router updates. The default is 90 seconds.
- **Route invalid timer**—The 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**—The time from when a route becomes invalid to when it is removed from the routing table. The default is 630 seconds.

### Question 40 Answer

IGRP can support up to six multiple equal or unequal paths. (Four is the default.)

#### NOTE

For IGRP to support multiple paths, the following rules must apply:

- The next-hop router in any of the paths must be closer to the destination than the local router is. This ensures that no routing loops occur.
- The alternative path metric must be within the specified variance of the best local metric.

### Question 41

**How do you enable IGRP on a Cisco router?**

### Question 42

**Router A has two unequal paths to Router B. Interface S0 is connected to Router B's S0 interface at 1544 kbps, and interface S1 is connected to Router B's S1 interface at 256 kbps. By default, IGRP uses interface S0 to send traffic to Router B because it has the better metric, leaving S1 unused. How do you enable Router A to use interfaces S0 and S1 to route traffic to Router B?**

## Question 41 Answer

The way you enable IGRP on a Cisco router is similar to the way you enable RIP, except that you specify IGRP as the protocol and add an autonomous system number:

```
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
```

## Question 42 Answer

To enable unequal-cost routing in Router A, you use the **variance** command. It defines a multiplier by which a metric may vary from the metric of the lowest-cost route. Any route whose metric exceeds the metric of the lowest-cost route multiplied by the variance is not used. The variance must be specified in whole numbers.

For example, suppose that Router A's metric through S0 is 8576, and its metric through S1 is 41162.  $41162 / 8576 = 4.8$ . Thus, the S1 metric is 4.8 times larger than S0, so to enable unequal-cost routing on Router A, the variance is 5. The configuration is as follows:

```
RouterA(config)#router igrp 10
RouterA(config-router)#network 192.168.0.0
RouterA(config-router)#network 192.168.1.0
RouterA(config-router)#variance 5
```

**Question 43**

**What are some of EIGRP's improvements over IGRP?**

Determining IP  
Routes

**Question 44**

**By default, what does EIGRP use to calculate routes?**

Determining IP  
Routes

### Question 43 Answer

Here are some of EIGRP's improvements over IGRP:

- DUAL
- Incremental updates
- Loop-free networks
- Reduced bandwidth usage
- Support for multiple network layer protocols (IP, IPX, AppleTalk)
- Support for variable-length subnet masks (VLSMs), discontinuous networks, and classless routing
- Advanced distance vector capabilities
- Automatic route summarization on major network boundaries

### Question 44 Answer

EIGRP uses bandwidth and delay by default to calculate its metric. It can also be configured to use reliability, load, and MTU. EIGRP's metric is the same as IGRP's metric, except that it is multiplied by 256 for improved granularity.

### Question 45

**In EIGRP, what is a successor?**

Determining IP  
Routes

### Question 46

**In EIGRP, what is the feasible successor?**

Determining IP  
Routes



### Question 45 Answer

A successor is a route selected as the primary route used to reach a destination. It is the route kept in the routing table.

### Question 46 Answer

The feasible successor is the backup route. These routes are selected at the same time the successors are identified, but they are kept only in the topology table, not in the routing table. They are used for fast convergence. If the successor fails, the router can immediately route through the feasible successor. A destination can have multiple feasible successors.

**Question 47**

**What IOS commands enable EIGRP on a Cisco router and advertise 192.168.3.0 and 192.168.4.0 as its directly connected networks?**

**Question 48**

**What command would you use to see EIGRP adjacencies?**

## Question 47 Answer

```
RouterA(config)#router eigrp 100 (100 is the AS number)
RouterA(config-router)#network 192.168.3.0
RouterA(config-router)#network 192.168.4.0
```

## Question 48 Answer

The **show ip eigrp neighbors** command displays EIGRP adjacencies and directly connected neighbors.

**Question 49**

**What IOS command would you use to view the EIGRP neighbor states?**

**Question 50**

**What routing metric is OSPF based on?**

## Question 49 Answer

You use the **debug eigrp neighbors** command to check the EIGRP neighbor states. This command displays the contents of the hello packet used in EIGRP as well as the neighbors discovered by EIGRP.

## Question 50 Answer

Bandwidth. OSPF's metric is a cost value based on bandwidth or the speed of its connection. The default formula used to calculate OSPF cost is

$$\text{Cost} = 100,000,000 / \text{bandwidth in bps}$$

For example, OSPF assigns the cost of 10 to a 10 MB Ethernet line ( $100,000,000 / 10,000,000 = 10$ ).

### NOTE

This is a tricky question in some ways. For example, with this question, you might have to choose between cost and bandwidth. Bandwidth is the more correct answer, because OSPF metric uses a cost based on bandwidth.

### Question 51

**How do OSPF-speaking routers build adjacencies and exchange routing tables?**

### Question 52

**What does the Hello protocol do in an OSPF network?**

### Question 51 Answer

OSPF-speaking routers build adjacencies by sending Hello packets out all OSPF-enabled interfaces. If the routers share a common data link and agree on certain parameters set in their Hello packets, they become neighbors. If these parameters are different, the routers do not become neighbors, and communication stops. OSPF routers form adjacencies with certain routers. These routers are determined by the data link media type. As soon as adjacencies are formed, each router sends LSAs to all adjacent routers. These LSAs describe the state of each of the router's links. Because of the varying types of link-state information, OSPF defines multiple LSA types. Finally, a router 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, each router uses the SPF algorithm to calculate a loop-free topology and builds its routing table based on this topology.

### Question 52 Answer

In OSPF, the Hello protocol ensures that communication between OSPF-speaking routers is bidirectional. It is the means by which neighbors are discovered and acts as keepalives between neighbors. It also establishes and maintains neighbor relationships and elects the designated router (DR) and backup designated router (BDR) to represent the segment on broadcast and nonbroadcast multiaccess (NBMA) networks.

**Question 53**

**What are the five network types that OSPF defines?**

Determining IP  
Routes

**Question 54**

**What is the IP multicast address of Hello protocols?**

Determining IP  
Routes



## Question 53 Answer

The five network types that OSPF defines are

- Broadcast networks
- NBMA networks
- Point-to-point networks
- Point-to-multipoint networks
- Virtual links

Examples of broadcast networks are Ethernet and Token Ring. OSPF routers on broadcast networks elect a DR and BDR. All the routers form adjacencies with the DR and BDR. On broadcast networks, all OSPF packets are multicast to the DR and BDR.

NBMA networks are Frame Relay, X.25, and ATM. They can connect more than two routers but have no broadcast capability. NBMA networks elect a DR and BDR, and all OSPF packets are unicast.

Point-to-point networks, such as a T1, connect a single pair of routers and always become adjacent.

Point-to-multipoint networks are a special configuration of NBMA networks in which networks are treated as a collection of point-to-point links. Routers on these networks do not elect a DR or BDR, and because all links are seen as point-to-point, all OSPF packets are multicast.

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

## Question 54 Answer

Hello protocols are periodically sent out each interface using IP multicast address 224.0.0.5 (All SPF Routers). The HelloInterval each router uses to send out the Hello protocol is based on the media type. The default HelloInterval of broadcast, point-to-point, and point-to-multipoint networks is 10 seconds. On NBMA networks, the default is 30 seconds.

**Question 55**

**If you have eight routers on an Ethernet network and you establish adjacencies with only the DR and BDR, how many circuits do you have?**

**Question 56**

**What is the OSPF router ID, and where does an OSPF router receive its router ID?**

### Question 55 Answer

14. The formula for figuring out the number of circuits (adjacencies or connections) needed to establish adjacencies on the DR and BDR is  $2(n - 1)$ , where  $n$  is the number of routers in the network. So if you have eight routers in a network,  $2(8 - 1) = 14$  adjacencies.

### Question 56 Answer

For OSPF to initialize, it must be able to define a router ID for the entire OSPF process. A router can receive its router ID from several sources. The most common and stable source is the IP address set on the loopback interface. The loopback interface is a logical interface that never goes down. If no loopback address is defined, an OSPF-enabled router selects the numerically highest IP address on all its interfaces as its router ID.

**Question 57**

An OSPF-enabled router has the following IP addresses configured on its interfaces:

Ethernet 0192.168.9.5

Serial 0172.16.3.1

Ethernet 1192.168.24.1

What is the router ID of the OSPF-enabled router?

**Question 58**

What are link-state advertisements?

### Question 57 Answer

The router ID is 192.168.24.1 because it is the numerically highest IP address on all interfaces on the router. If the router had a loopback address configured, it would choose the loopback address as the router ID (even if the loopback IP address was numerically lower than other IP addresses configured on the router).

### Question 58 Answer

Link-state advertisements (LSAs) are what OSPF-speaking routers send out all interfaces, describing the state of the router's links. LSAs are also packets that OSPF uses to advertise changes in the condition of a specific link to other OSPF routers.

**Question 59**

**How many LSAs are there in OSPF?**

Determining IP  
Routes

**Question 60**

**What are five reasons why you would use OSPF instead of RIP?**

Determining IP  
Routes

### Question 59 Answer

Six different and distinct link-state packet formats are used in OSPF—each for a different purpose. The ICND exam will test you on only two LSA types—Type 1 and Type 2.

Type 1 LSAs are router LSAs and are generated by each router for each area to which it belongs. These LSAs describe the states of the router's links to the area and are flooded within a single area.

Type 2 LSAs are network LSAs and are generated by the DR and BDR. They describe the set of routers attached to a particular network. They are flooded within a single area.

### Question 60 Answer

The reasons why you would use OSPF instead of RIP are as follows:

- Support for VLSMs
- Fast convergence
- No reachability limitations
- More efficient use of bandwidth
- Path selection is based on bandwidth rather than hops

### **Question 61**

**How do you enable OSPF on a Cisco router?**

### **Question 62**

**What command enables a loopback interface on a Cisco router?**



## Question 61 Answer

The **router ospf** *process-id* command enables the OSPF process, and the **network** *address wildcard-mask area area-id* command assigns networks to a specific OSPF area. Notice that you must specify the wildcard mask instead of the subnet mask. For example, the following commands enable OSPF process 10 and advertise the network 192.168.10/24 in area 0:

```
RouterA(config)#router ospf 10
RouterA(config-router)#network 192.168.10.0 0.0.0.255 area 0
```

### NOTE

The process ID is locally significant to the router, because all OSPF routes are assumed to belong to the same OSPF domain.

## Question 62 Answer

RouterA(config)#**interface loopback** *number* configures a loopback interface. The *number* option specifies the loopback interface number you are creating.

### NOTE

Remember that a loopback interface is a logical interface in the router. If you have a loopback interface created on an OSPF-enabled router, the router ID is the loopback address.

### Question 63

**What IOS command displays the OSPF neighbor information on a per-interface basis?**

### Question 63 Answer

The `show ip ospf neighbor` command displays OSPF neighbor information on a per-interface basis.

