PassQueen
a.k.a dumpsqueen.com

Nokia OSPF Routing Protocol Exam

Nokia 4A0-113

Version Demo

Total Demo Questions: 15

Total Premium Questions: 159

Buy Premium PDF

https://passqueen.com

support@passqueen.com

QUESTION NO: 1

What is the source IP address in an OSPF update?

A. 224.0.0.5

B. 224.0.0.6

C. The unicast IP address of the interface the update is sent on

D. The system address of the router.

ANSWER: C

Explanation:

The unicast IP address of the interface the update is sent on is correct. An OSPF Link State Update is carried in an IP packet, and the source field of that IP header identifies the sending router's local outgoing interface address. This is the address that is routable and valid on the connected network segment over which the OSPF packet is transmitted. OSPF's Router ID, often called the system address in Nokia contexts, is a separate logical identifier used within OSPF protocol processing and LSAs; it is not automatically the source address of every transmitted OSPF packet.

On broadcast and non-broadcast multi-access networks, OSPF commonly sends updates to an OSPF multicast destination, but multicast addresses are destination addresses only. The sending router still places its transmitting interface's unicast address in the IP source field. RFC 2328 specifies that OSPF packets are sent directly over IP and describes the IP destination handling for OSPF traffic; the IP header source remains the local sending interface. See RFC 2328: OSPF Version 2, especially the protocol packet and IP encapsulation details, and Appendix A.3.5, Link State Update packet format.

QUESTION NO: 2

Graceful Restart is supported by what IGPs?

A. RIPv1 and RIPv2

B. OSPF

C. IS-IS

D. BGP

ANSWER: B C

Explanation:

Graceful Restart is supported by **OSPF** and **IS-IS**. In both link-state IGPs, graceful-restart procedures are intended to preserve forwarding continuity while a routing process or control plane restarts. Neighboring routers can retain adjacency and topology information for a negotiated restart interval, allowing the restarting router to recover its protocol state and resume normal operation without necessarily causing an immediate traffic interruption. OSPF graceful restart is standardized by RFC 3623, which defines the mechanisms used by restarting and helper routers to maintain forwarding during a restart. See RFC 3623: Graceful OSPF Restart. IS-IS graceful restart is standardized by RFC 5306 and similarly specifies restart signaling and helper behavior for maintaining forwarding-plane stability while IS-IS state is rebuilt. See RFC 5306: Restart Signaling for IS-IS. Therefore, the applicable IGP selections are OSPF and IS-IS.

QUESTION NO: 3

Which of the following statements apply to link state protocol behavior? Choose three answers.

A. Routers broadcast the entire route table to all neighbors.

B. Information about directly connected links is sent to all neighbors.

C. An adjacency database is maintained by each router.

D. The sequence number for an update is incremented as it is flooded from router to router.

E. The topological database is the same for all routers in a single area.

ANSWER: B C E

Explanation:

Link-state routing protocols such as OSPF distribute information describing router links, their states, and associated metrics rather than distributing a complete routing table. Accordingly, “Information about directly connected links is sent to all neighbors” describes the essential link-state advertisement and flooding behavior: routers originate link-state information for their local interfaces and propagate it through the applicable routing domain so routers can build a shared topology view.

“An adjacency database is maintained by each router” is also applicable. OSPF routers track neighboring routers and the state of their adjacencies, including which peers have formed the synchronized relationships used for reliable link-state database exchange and flooding.

Finally, “The topological database is the same for all routers in a single area” reflects the link-state model. After successful flooding and synchronization, routers within an OSPF area hold a consistent link-state database for that area. Each router independently runs the SPF calculation against this common topology information to derive its own routing table. The OSPF specification describes neighbor adjacency operation, link-state database synchronization, and area-scoped LSAs in RFC 2328; related protocol fundamentals are also summarized by the IETF RFC repository.

QUESTION NO: 4

Which of the following correctly lists the 5 OSPF packet types?

A. Hello, link state update, link state packet, link State acknowledgement, database description.

B. Link state update, link state refresh, link state acknowledgement, database description, notification.

C. Hello, link state refresh, link state update, link state acknowledgement, link state authentication.

D. Hello, link state update, link state request, link state acknowledgement, database description.

ANSWER: D

Explanation:

Hello, link state update, link state request, link state acknowledgement, database description. is correct because these are the five packet types defined by the OSPFv2 protocol. Hello packets discover neighboring OSPF routers and maintain neighbor adjacencies on an interface. Database Description packets summarize the contents of a router’s link-state database during adjacency formation, allowing neighbors to determine which link-state advertisements require synchronization. Link State Request packets then request specific missing or outdated advertisements. Link State Update packets carry one or more link-state advertisements to synchronize topology information and to flood changes through an OSPF area. Link State Acknowledgment packets explicitly acknowledge received link-state advertisements, providing reliable flooding over interfaces where acknowledgments are required. These packet types are identified in the OSPF packet header by packet-type values 1 through 5. The definitive protocol specification is RFC 2328, section 4.3, which names Hello, Database Description, Link State Request, Link State Update, and Link State Acknowledgment as the complete OSPFv2 packet-type set. Nokia SR OS documentation also describes OSPF neighbor establishment and LSDB synchronization using these packet exchanges. See RFC 2328, section 4.3 and Nokia SR OS OSPF documentation.

QUESTION NO: 5

Click the exhibit button.

A static route is created on router R1 "using the static-route 0.0.0.0/0 next-hop 10.3.3.2" command. What command can be used to test the static route on an Nokia 7750 SR?

A. ping static 10.5.5.5

B. ping static 0.0.0.0/0

C. ping 10.5.5.5

D. ping 0.0.0.0

ANSWER: C

Explanation:

`ping 10.5.5.5` is the appropriate test because it sends ICMP Echo traffic to an actual destination address that is expected to be reached through the configured default static route. When R1 performs the route lookup for 10.5.5.5, no more-specific route is presumed to exist, so the `0.0.0.0/0` route is selected. The router then forwards the packet to the configured next hop, 10.3.3.2. A successful reply confirms not only that the static route can be selected for a non-local destination, but also that next-hop reachability and the return path are functioning sufficiently for ICMP traffic. This is the practical operational use of a default route: it provides a forwarding path for destinations not matched by a more-specific prefix. Nokia SR OS includes ping as a network troubleshooting tool, and an ordinary destination-address ping performs the normal forwarding-table lookup needed to validate this behavior. See the Nokia SR OS documentation portal for SR OS operational and troubleshooting command documentation, and RFC 1812 section 5.2.4.3 for the default-route forwarding concept.

QUESTION NO: 6

Which of the following statements describe OSPF operation on a point-to-point interface? Choose three answers.

A. A DR and BDR election is not performed.

B. The two routers normally form a Full adjacency.

C. A Type 2 Network LSA is not originated for the link.

D. All OSPF packets must be sent using TCP.

E. The router with the highest priority becomes the DR.

ANSWER: A B C

Explanation:

On a point-to-point interface, OSPF does not elect a DR or BDR because exactly two routers are attached to the link. The two routers normally proceed beyond the 2-Way state and form a Full adjacency. Also, a Type 2 Network LSA is not originated for a point-to-point network because there is no multi-access transit network requiring representation by a designated router.

These interface behaviors are part of the OSPF interface and LSA definitions in RFC 2328, Section 8.1 and RFC 2328, Section 12.4.2.

QUESTION NO: 7

What causes an adjacency to change from down to two ways?

A. When a link state update is received in response to a link state request.

B. When a router receives a Hello packet that contains its own router ID in the neighbor list from a neighbor.

C. When a router receives a database description packet from a neighbor.

D. When a link state acknowledgement is received in response to a link state update.

ANSWER: B

Explanation:

When a router receives a Hello packet that contains its own router ID in the neighbor list from a neighbor is correct because this is OSPF's bidirectional-neighbor confirmation mechanism. Initially, a router that has not heard from a neighbor is in the Down state. Receipt of a valid Hello identifies the neighbor and moves the relationship to Init. The local router then verifies two-way communication when a subsequent Hello received from that neighbor lists the local router's own Router ID in the Hello packet's Neighbor field. This proves that the neighbor has received the local router's Hello packets, so the neighbor state advances to 2-Way.

The 2-Way state is the point at which OSPF has confirmed bidirectional connectivity. On point-to-point links, this normally leads to database synchronization and formation of a full adjacency. On multi-access networks, routers may remain 2-Way with neighbors that are neither the designated router nor backup designated router, while maintaining full adjacencies with the appropriate election participants. RFC 2328 defines this exact transition under the Hello-received event and describes the Neighbor ID check that establishes two-way communication. See [RFC 2328, Section 10.5] and [Nokia OSPF documentation].

QUESTION NO: 8

Click the exhibit button.

In the topology shown, router R1 is an ASBR configured to export external routes to OSP

F.

Assuming that there are no stub networks, which of the following statements regarding type 4 LSA generation is true?

A. Router R1 generates a type 4 LSA that is flooded to areas 0, 1, and 2.

B. Router R3 generates a type 4 LSA that is flooded to areas 0, 1, and 2.

C. Router R3 generates a type 4 LSA that is flooded to areas 0 and 2.

D. Router R3 generates a type 4 LSA that is flooded to area 0, and router R6 generates a type 4 LSA that is flooded to area 2.

ANSWER: D

Explanation:

Router R3 generates a type 4 LSA that is flooded to area 0, and router R6 generates a type 4 LSA that is flooded to area 2. This follows the OSPF inter-area mechanism used to make an autonomous-system boundary router reachable from areas other than the one containing it. R1 is the ASBR and originates the external-route information as Type 5 LSAs. However, routers outside R1's attached area also require the intra-AS reachability information needed to calculate a route to the ASBR itself. An ABR supplies that information through a Type 4 ASBR-summary LSA.

R3 is the ABR connecting R1's area to the backbone, so it originates the Type 4 LSA into area 0, advertising reachability to the ASBR. R6 is the ABR between the backbone and area 2. After receiving the ASBR-summary information through the backbone, it originates the corresponding Type 4 LSA into area 2. Thus, each ABR originates Type 4 information into its attached area as required; a Type 4 LSA is not simply flooded unchanged through all areas. This behavior is specified in [RFC 2328, OSPF Version 2], particularly the ASBR-summary LSA definition and ABR processing rules. Nokia's OSPF implementation documentation also describes ABR handling of inter-area OSPF information: [Nokia SR OS OSPF documentation].

QUESTION NO: 9

What are the ways the RID can be created on an Nokia 7750 SR router? (Choose three)

A. From the last 32 bits of the chassis MAC address

B. From the first 32 bits of the chassis MAC address

C. The loopback IP address

D. By using the command "config router ospf router-id X.X.X.X"

E. From the system IP address

F. From the highest interface IP address

ANSWER: A D E

Explanation:

On a Nokia 7750 SR, the OSPF router ID is a 32-bit value used to uniquely identify the router within an OSPF routing domain. It can be explicitly assigned under the OSPF configuration using `config router ospf router-id X.X.X.X`. This is the preferred operational method when a stable, deliberately planned identifier is required, because the configured value does not depend on automatically selected platform addressing information.

When an explicit OSPF router ID is not configured, SR OS uses the system IP address as the router ID. The system IP address is configured on the system interface and provides a stable, router-wide address suitable for control-plane protocols. If neither an explicit OSPF router ID nor a system IP address is available, SR OS derives the router ID from the last 32 bits of the chassis MAC address. This provides a deterministic fallback that can be represented in dotted-decimal IPv4 format.

These selection methods and the OSPF configuration hierarchy are documented in the Nokia SR OS router configuration documentation and command reference. See Nokia SR OS OSPF configuration documentation and Nokia SR OS system configuration documentation.

QUESTION NO: 10

Which of the following statements regarding an OSPF Area Border Router are true? Choose three answers.

A. It has interfaces in more than one OSPF area.

B. It maintains a separate link-state database for each attached area.

C. It can originate Type 3 Summary LSAs.

D. It must redistribute routes from another routing protocol.

E. It is elected on every broadcast network.

ANSWER: A B C

Explanation:

An Area Border Router (ABR) has interfaces assigned to more than one OSPF area. It maintains a separate link-state database for each attached area and performs SPF calculations using the topology information for those areas. An ABR also distributes inter-area reachability by originating Type 3 Summary LSAs. When an ASBR is located in another area, an ABR can originate a Type 4 Summary LSA describing how to reach that ASBR.

An ABR is not required to redistribute routes from another routing protocol. That is the function of an Autonomous System Boundary Router (ASBR). OSPF ABR behavior and summary-LSA processing are specified in RFC 2328, Section 3.3.

QUESTION NO: 11

Choose the Distance Vector Protocols from the list below:

A. OSPF

B. RIPv1

C. RIPv2

D. IS-IS

ANSWER: B C

Explanation:

RIPv1 and RIPv2 are distance-vector interior gateway protocols. They implement the Routing Information Protocol algorithm, derived from the Bellman-Ford approach, in which routers periodically advertise routing information to directly connected neighbors. Route selection is primarily based on hop count, and RIP defines a maximum usable metric of 15 hops; a metric of 16 represents an unreachable destination. RIPv1 is the original classful version of RIP, while RIPv2 retains distance-vector route calculation but adds classless network support, including subnet-mask information in route advertisements, as well as capabilities such as authentication and multicast updates. Therefore, both RIPv1 and RIPv2 belong in the distance-vector category. This classification is important when comparing routing protocols because distance-vector protocols exchange reachability and metric information with neighbors rather than distributing a complete link-state database and independently running shortest-path calculations across the topology. The IETF RIP version 2 specification explicitly describes the protocol's route entries, metrics, update processing, and hop-count limits. See [RFC 2453: RIP Version 2]. For the underlying RIP version 1 behavior and packet format, see [RFC 1058: Routing Information Protocol].

QUESTION NO: 12

Click the exhibit button.

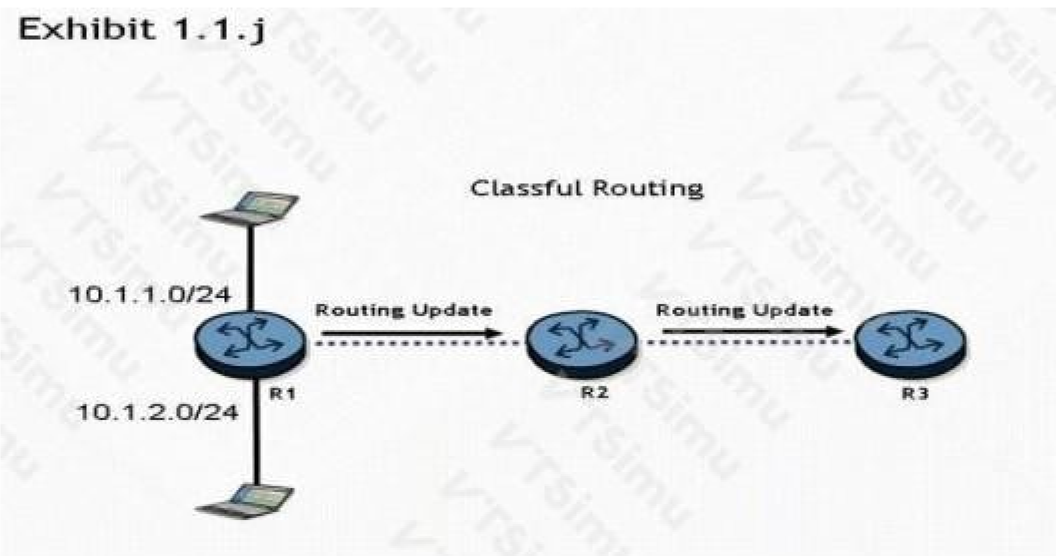

Routers R1 „ R2, and R3 are running a classful routing protocol between them. Assuming that router R1 advertises all directly connected networks, how will these networks be represented in router R3's routing table?

A. Router R3's routing table can only contain one of the routes, which will result in route flapping.

B. Router R3's routing table will have one entry for 10.1.1.0/24 and one entry for 10.1.2.0/24.

C. The networks will be represented with one entry of 10.0.0.0/8 in router R3's routing table.

D. The networks will be represented with one entry of 10.0.0.0/24 in router R3's route table.

ANSWER: C

Explanation:

The networks will be represented with one entry of 10.0.0.0/8 in router R3's routing table. A classful routing protocol carries only the network number for a route and does not include subnet-mask information in its route advertisements. Therefore, when R1 advertises its directly connected 10.1.1.0/24 and 10.1.2.0/24 subnets through the classful-routing domain, those subnets are interpreted and advertised using their classful major-network boundary. Since addresses beginning with 10 are Class A addresses, their classful network is 10.0.0.0 with an 8-bit prefix. R3 consequently installs the summarized classful destination 10.0.0.0/8 rather than retaining separate /24 subnet routes. This behavior is characteristic of protocols such as RIPv1, whose route entries contain an IP address but no subnet mask. RFC 1058 describes the original RIP route-entry format and its address-family/IP-address fields: RFC 1058: Routing Information Protocol. The classful IPv4 address ranges and default masks, including the Class A /8 boundary used here, are documented in RFC 791: Internet Protocol.

QUESTION NO: 13

Click the exhibit button.

Exhibit 4.3.h

•OSPF Version : 2	•OSPF Version : 2
•Router Id : 13.0.0.2	•Router Id : 13.0.0.1
•Area Id : 0.0.0.0	•Area Id : 0.0.0.0
•Checksum : 7c0e	•Checksum : 865e
•Authentication : Null	•Authentication : Null
•Authentication Key: 00 00 00 00 00 00 00 00	•Authentication Key: 00 00 00 00 00 00 00 00
•Packet Type : DB_DESC	•Packet Type : DB_DESC
•Packet Length : 32	•Packet Length : 32
•	•
•Interface MTU : 9000	•Interface MTU : 9000
•Options : 000042	•Options : 000042
•Flags : 7	•Flags : 7
•Sequence Num : 77793	•Sequence Num : 75667

The OSPF packets shown have been exchanged between two neighbors. What state is the OSPF adjacency in?

A. The adjacency state is two way because both routers have received a database description packet.

B. The adjacency is in the exchange state. The next step is to exchange link state request and link state update packets.

C. The adjacency is in the ex start state. The routers have exchanged database description packets, but still have to negotiate the starting sequence number.

D. The adjacency state is full. The routers are exchanging database description packets to notify each other that all link state advertisements have been received.

ANSWER: C

Explanation:

The adjacency is in the ex start state. In OSPF, ExStart is the initial phase of database synchronization after the neighbors have progressed beyond two-way communication and have decided to form a full adjacency. The routers use Database Description (DD) packets during this phase to establish which router will be the master and which will be the slave, and to negotiate the initial DD sequence number. The master controls the DD packet sequence, allowing both routers to synchronize the subsequent description exchange reliably.

The DD packet exchange shown is therefore evidence of the database-synchronization process beginning, rather than proof that the routers have already completed LSDB synchronization. Once the master/slave relationship and initial sequence number have been agreed, the adjacency advances to Exchange. In that later state, the routers continue sending DD packets that summarize the LSAs in their respective link-state databases. The ExStart state and its master/slave and sequence-number negotiation are defined in RFC 2328, Section 10.6. The DD packet fields and their role in adjacency database synchronization are described in RFC 2328, Section 10.8.

QUESTION NO: 14

Click the exhibit button.

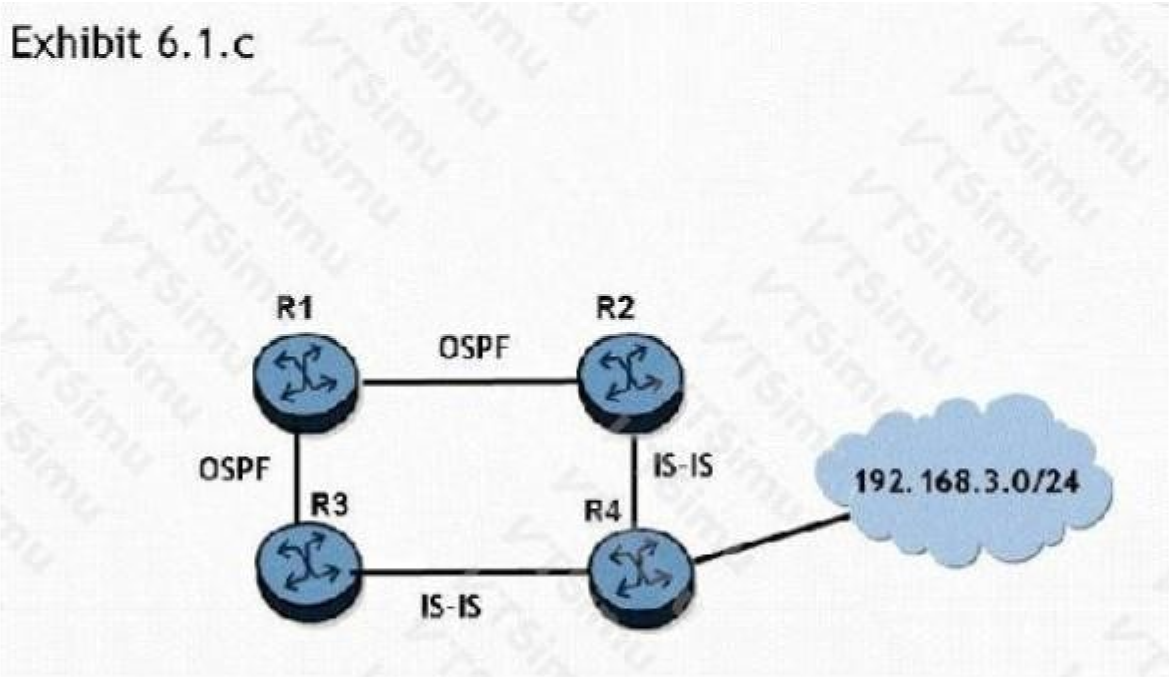

If router R2 redistributes the IS-IS route to 192.168.3.0/24 into OSPF, router R3 preference of these two routes? Choose two answers.

A. The OSPF internal preference

B. The IS-IS Level 1 internal preference

C. The IS-IS Level 2 internal preference

D. The OSPF external preference

E. The IS-IS Level 2 external preference

F. The IS-IS Level 1 external preference

ANSWER: B D

Explanation:

Router R3 learns the 192.168.3.0/24 prefix through two route sources. The directly learned IS-IS advertisement is an IS-IS Level 1 internal route, so it has the default IS-IS Level 1 internal preference. On Nokia SR OS, the default preference for an IS-IS Level 1 internal route is 15. This is the preference category applicable to the native route learned within the Level 1 area.

When R2 redistributes that IS-IS prefix into OSPF, the redistributed prefix is learned by R3 as an OSPF external route. Redistribution creates an external OSPF route rather than an OSPF internal route, and SR OS applies the OSPF external preference category to it. The default OSPF external preference is 150. Thus, the two route types and preference categories present at R3 are the IS-IS Level 1 internal preference and the OSPF external preference. Because SR OS selects the route with the numerically lowest preference, the Level 1 internal route is the preferred route for installation, while the external OSPF route remains an alternative learned route. Nokia documents route preference as the primary protocol-level route-selection value and documents separate defaults for IS-IS and OSPF route classes. See the [Nokia SR OS route-policy documentation]() and the [Nokia SR OS OSPF documentation]().

QUESTION NO: 15

Click on the exhibit.

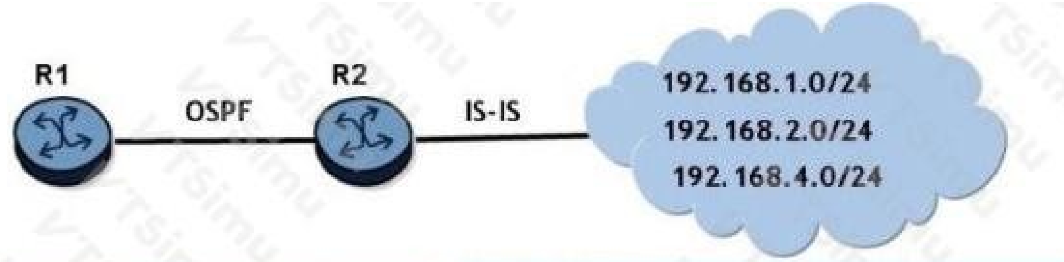


```
*A:AIRP_R02>config>router>ospf# info
----------------------------------------------
            asbr
            export "isis-ospf"
            area 0.0.0.0
                interface "system"
                    no shutdown
                exit
                interface "toR1"
                    interface-type point-to-point
                    no shutdown
                exit
            exit
----------------------------------------------
*A:AIRP_R02>config>router>ospf#
```

```
*A:AIRP_R02>config>router>policy-options# info
----------------------------------------------
            prefix-list "no-redist"
                prefix 192.168.2.0/24 exact
            exit
            prefix-list "redist-16"
                prefix 192.168.0.0/16 longer
            exit
            policy-statement "isis-ospf"
                entry 10
                    from
                        protocol isis
                        prefix-list "redist-16"
                    exit
                    action accept
                    exit
                exit
                entry 20
                    from
                        protocol isis
                        prefix-list "no-redist"
                    exit
                    action reject
                exit
            exit
----------------------------------------------
*A:AIRP_R02>config>router>policy-options#
```

Router R2 learns the networks shown from IS-IS. Router R2 is configured to redistribute routes from IS-IS to OSPF as shown. Assume the OSPF and IS-IS adjacencies are operational.

How many OSPF Type 5 LSAs will router R1 receive from router R2?

A. 0

B. 1

C. 2

D. 3

E. 4

ANSWER: D

Explanation:

3 is correct. R2 acts as an ASBR when it redistributes IS-IS routes into OSPF. The redistribution policy shown in the exhibit determines which IS-IS-learned prefixes are eligible for export into OSPF. Three of the learned routes match the policy criteria and are therefore injected as OSPF external routes. For each eligible external prefix, R2 originates a separate Type 5 AS-external LSA, which is flooded throughout the normal OSPF flooding scope. Because the OSPF adjacency between R1 and R2 is operational, R1 receives those three Type 5 LSAs from R2 and installs or evaluates the corresponding external destinations according to OSPF route selection rules. A Type 5 LSA describes a single external destination prefix, rather than combining multiple redistributed prefixes into one advertisement. Nokia SR OS routing policy processing controls route import and export decisions, including redistribution between routing protocols; OSPF external-route handling is described in Nokia's [Route Policies documentation] and [OSPF documentation].