Nokia NRS II Composite Exam: OSPF version

Nokia 4A0-C04

Version Demo

Total Demo Questions: 42

Total Premium Questions: 427

Buy Premium PDF

https://passqueen.com

support@passqueen.com

QUESTION NO: 1

An OSPF router learns an external destination through a Type 5 LSA with an external metric type of E1. Which cost is used to calculate the route metric?

A. The external metric plus the internal OSPF cost to reach the ASBR.

B. Only the internal OSPF cost to reach the ASBR.

C. Only the external metric in the Type 5 LSA.

D. The highest interface cost between the router and the ASBR.

ANSWER: A

Explanation:

The external metric plus the internal OSPF cost to reach the ASBR is used for an E1 external route. A Type 1 external metric represents the complete cost of reaching the external destination. Therefore, each router adds its intra-domain path cost to the advertising ASBR to the external metric contained in the Type 5 LSA.

This differs from an E2 external route. For E2 routes, the external metric is normally used without adding the internal cost to the ASBR, except as a tie-breaker between otherwise equal E2 routes. OSPF external-route calculation is specified in RFC 2328, Section 16.4.

QUESTION NO: 2

Which of the following statements regarding an OSPF stub area are true? Choose three answers.

A. Type 5 AS-external LSAs are not flooded into the area.

B. An ASBR cannot be located within the area.

C. An ABR injects a default route into the area.

D. Type 7 NSSA-external LSAs are used for redistributed routes.

E. The area can be used as a virtual-link transit area.

ANSWER: A B C

Explanation:

A stub area blocks Type 5 AS-external LSAs and Type 4 Summary ASBR-LSAs from entering the area. Since routers in the stub area do not receive external destination information or ASBR reachability information, an ASBR cannot be located within a stub area.

An ABR attached to a stub area originates a default summary route into the area. This gives internal routers a path toward destinations outside the area without requiring the external LSAs themselves to be flooded into the area. All routers attached to the stub area must agree on the stub-area capability. OSPF stub-area operation is defined in RFC 2328, section 3.6.

QUESTION NO: 3

Which of the following are valid OSPF packet types? Choose all that apply.

A. Hello

B. Database Description

C. Link State Request

D. Link State Update

E. Link State Acknowledgment

F. Shortest Path First

ANSWER: A B C D E

Explanation:

OSPF defines five packet types: **Hello**, **Database Description**, **Link State Request**, **Link State Update**, and **Link State Acknowledgment**. Hello packets discover and maintain neighbors. Database Description, Link State Request, and Link State Update packets are used to synchronize the link-state database. Link State Acknowledgment packets provide reliable flooding acknowledgment.

LSAs are records carried within Link State Update packets; an LSA is not itself an OSPF packet type. The OSPF packet header and packet types are defined in [RFC 2328, Appendix A.3.2].

QUESTION NO: 4

Click on the exhibit.

```
*A:SRC_R3# oam lsp-ping "lsp-toR2"
LSP-PING lsp-toR2: 92 bytes MPLS payload
Seq=1, send from intf toR1, reply from 10.10.10.2
       udp-data-len=32 ttl=255 rtt=0.000ms rc=3 (EgressRtr)

---- LSP lsp-toR2 PING Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min = 0.000ms, avg = 0.000ms, max = 0.000ms, stddev = 0.000ms
*A:SRC_R3#
```

Which of the following is FALSE?

A. The lsp-ping tool is used to perform a unidirectional LSP test.

B. The router 10.10.10.2 is the egress router of the LSP.

C. The router can reach the far-end of the LSP through an RSVP-TE tunnel.

D. The router can reach the far-end of the LSP through an LDP tunnel.

ANSWER: D

Explanation:

The statement *"The router can reach the far-end of the LSP through an LDP tunnel."* is false. The exhibit identifies an RSVP-TE signaled LSP toward the far-end router, rather than an LDP-signaled transport LSP. RSVP-TE establishes MPLS label-switched paths using RSVP signaling and can apply explicit path and traffic-engineering constraints. Consequently, the forwarding and OAM result shown for this test applies to the RSVP-TE LSP that terminates at the indicated egress, not to an LDP tunnel.

LSP ping is an MPLS OAM mechanism used by an ingress LSR to send a test packet along a specified FEC/LSP toward its egress. It validates the forward MPLS data-plane path; the response is returned by the egress using an available return path. This behavior is standardized in [RFC 4379, Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures]. Nokia SR OS documentation describes MPLS, RSVP-TE, and associated operational tools within its [SR OS documentation library]. Therefore, the tunnel type displayed by the test is the decisive detail: it is RSVP-TE transport, so describing the observed far-end reachability as occurring through an LDP tunnel is inaccurate.

QUESTION NO: 5

Which of the following statements regarding OSPF Type 7 LSAs are true? Choose two answers.

A. They are originated by an ASBR in an NSSA.

B. They can be translated into Type 5 LSAs by an NSSA ABR.

C. They are flooded throughout every OSPF area.

D. They describe multi-access transit networks.

E. They are originated only by the Designated Router.

ANSWER: A B

Explanation:

Type 7 LSAs are used to advertise external routes that are redistributed by an ASBR within a Not-So-Stubby Area. They are flooded only within the NSSA, rather than as Type 5 LSAs throughout the OSPF domain.

An NSSA ABR can translate eligible Type 7 LSAs into Type 5 AS-external LSAs for advertisement into other OSPF areas. This permits external routes introduced inside the NSSA to become reachable from the rest of the OSPF domain while retaining NSSA external-routing restrictions. Type 7 LSA operation and translation are specified in RFC 3101.

QUESTION NO: 6

Which of the following OSPF area types are supported by the Alcatel- Lucent 7750 SR?

Choose two answers.

A. Not so stubby areas

B. Level1 area WC

C. Stub areas

D. Partially stubby areas

E. Level 2 areas

ANSWER: A C

Explanation:

The Alcatel-Lucent 7750 SR supports both **Not so stubby areas** and **Stub areas** as OSPF area types. A stub area is designed to reduce the amount of external routing information maintained by routers within that area. The area border router blocks Type 5 AS-external LSAs from entering the area and provides a default route so internal routers can still reach destinations outside the OSPF domain. This behavior is part of the OSPFv2 specification’s stub-area operation.

A not-so-stubby area (NSSA) provides a related capability for an area that needs to import external routes locally. An NSSA uses Type 7 LSAs for those locally redistributed routes; the area border router can translate eligible Type 7 LSAs into Type 5 AS-external LSAs for advertisement into the rest of the OSPF autonomous system. This allows controlled external-route injection while retaining the reduced external-LSA flooding model of a stub-like area. These are standardized OSPF mechanisms and are implemented in SR OS OSPF. See the OSPFv2 specification in RFC 2328 and NSSA operation in RFC 3101.

QUESTION NO: 7

If OSPF is used in a multi-area OSPF network, which of the following statements regarding route summarization is true?

A. Manual route summarization can only be done on autonomous system border routers.

B. Manual route summarization must be done on all the backbone routers to be effective. Manual route summarization is optional on the routers in non-backbone areas.

C. Manual route summarization is done on the ABRs. By default, automatic summarization is done; however, it may not be optimal for all networks.

D. Manual route summarization is done on the ABRs. By default, no route summarization is done and all routes are advertised to all areas.

ANSWER: D

Explanation:

“Manual route summarization is done on the ABRs. By default, no route summarization is done and all routes are advertised to all areas.” is correct. In a multi-area OSPF design, an Area Border Router (ABR) connects Area 0 to one or more non-backbone areas and originates inter-area summary LSAs. Without an explicitly configured range, the ABR advertises the individual intra-area prefixes that it learns from an attached area into the other areas. OSPF therefore does not perform automatic classful or automatic prefix summarization by default. An administrator can configure a summary range on the ABR so that multiple contiguous routes from an area are represented by one inter-area summary route when advertised outside that area. This can reduce the number of LSAs and routing-table entries in the rest of the OSPF domain, while containing topology detail within the originating area. The relevant OSPF behavior is defined in RFC 2328: ABRs originate summary-LSAs to describe routes available in their attached areas, and the advertised prefixes reflect configured summary ranges where applicable. Nokia SR OS OSPF documentation likewise describes route aggregation as an explicitly configured ABR function. See RFC 2328, OSPF Version 2 and Nokia SR OS OSPF documentation.

QUESTION NO: 8

Exhibit

```
ALA-1# configure router
ALA-1>config>router# router-id 138.120.54.73
ALA-1>config>router# ospf
ALA-1>config>router>ospf$ area 0.0.0.0
ALA-1>config>router>ospf>area$ interface ALA-2
ALA-1>config>router>ospf>area>if# interface-type broadcast
ALA-1>config>router>ospf>area>if# metric 100
ALA-1>config>router>ospf>area>if# priority 25
ALA-1>config>router>ospf>area>if# exit
```

```
•ALA-2# configure router
•ALA-2>config>router# router-id 138.120.54.72
•ALA-2>config>router# ospf
•ALA-2>config>router>ospf$ area 0.0.0.0
•ALA-2>config>router>ospf>area$ interface ALA-1
•ALA-2>config>router>ospf>area>if# interface-type point-to-
point
•ALA-2>config>router>ospf>area>if# metric 10
•ALA-2>config>router>ospf>area>if# priority 25
•ALA-2>config>router>ospf>area>if# exit
•ALA-2>config>router>ospf>area# interface system
```

Given the two configurations shown, identify the two incorrect statements below: (Choose two)

A. Both routers will generate OSPF Hello packets.

B. The system interface is missing in one configuration so OSPF will not operate correctly.

C. The interface type settings differ, but OSPF will still operate correctly.

D. Metric values differ, but this does not prevent OSPF operation.

ANSWER: B C

Explanation:

The two statements to select are “The system interface is missing in one configuration so OSPF will not operate correctly.” and “The interface type settings differ, but OSPF will still operate correctly.” OSPF operation on an SR OS router is enabled on the relevant IP interfaces under an OSPF area; the system interface does not have to be included as an OSPF-enabled interface for a peering over a different routed link to send and receive Hellos. The system interface commonly supplies a stable router ID when an explicit router ID is not configured, but its absence from the displayed OSPF interface configuration is not, by itself, proof that OSPF cannot operate. The interface network type, however, is an essential part of the operational design of an OSPF link. It controls neighbor handling and, where applicable, designated-router behavior and the expected adjacency model. Matching the OSPF interface type at both ends avoids inconsistent behavior and is required for a correctly operating adjacency in this scenario. OSPF Hello processing and adjacency establishment are defined in RFC 2328; Nokia’s SR OS OSPF configuration guidance is available in the Nokia SR OS OSPF documentation.

QUESTION NO: 9

Which of the following statements regarding an OSPF Not-So-Stubby Area (NSSA) are true? Choose three answers.

A. Type 5 AS-external LSAs are not flooded into an NSSA.

B. An ASBR in an NSSA advertises redistributed routes using Type 7 LSAs.

C. An NSSA ABR can translate eligible Type 7 LSAs into Type 5 LSAs.

D. All routers in an NSSA must be ASBRs.

E. Type 4 LSAs are flooded into the NSSA for every external route.

ANSWER: A B C

Explanation:

Type 5 AS-external LSAs are not flooded into an NSSA, **an ASBR in an NSSA advertises redistributed routes using Type 7 LSAs**, and **an NSSA ABR can translate eligible Type 7 LSAs into Type 5 LSAs** are correct. An NSSA provides a stub-like area boundary while still allowing routes external to OSPF to be injected from within the area.

External routes imported by an NSSA ASBR are advertised inside the NSSA as Type 7 NSSA-external LSAs. The elected translating ABR converts eligible Type 7 information into Type 5 AS-external LSAs for advertisement into the rest of the OSPF domain. NSSA operation and Type 7 to Type 5 translation are defined in RFC 3101.

QUESTION NO: 10

Which of the following statements describe OSPF passive-interface behavior? Choose two answers.

A. OSPF does not send Hello packets on a passive interface.

B. The connected network can still be advertised in the Router-LSA.

C. A passive interface forms a full adjacency without receiving Hello packets.

D. A passive interface suppresses all advertisement of its connected prefix.

E. A passive interface must be configured as an OSPF broadcast network.

ANSWER: A B

Explanation:

OSPF does not send Hello packets on a passive interface and **the connected network can still be advertised in the Router-LSA** are correct. Configuring an interface as passive prevents OSPF neighbor discovery and adjacency formation on that interface. This is useful on interfaces facing end hosts or networks where no OSPF peer should exist.

The interface can nevertheless remain part of the OSPF routing domain. Its associated connected prefix is advertised as a stub network in the router's Type 1 Router-LSA, allowing other OSPF routers to calculate reachability to that network. Passive-interface operation is a local implementation control that preserves prefix advertisement while suppressing protocol traffic and adjacency formation on the interface. The Router-LSA representation of stub networks is described in RFC 2328, Section 12.4.1. Nokia SR OS OSPF interface configuration is available in the Nokia SR OS Documentation Center.

QUESTION NO: 11

Which of the following statements describe the major features of OSPF? Choose two answers.

A. Fast reroute capability

B. Control traffic prioritization

C. Route redistribution

D. Traffic engineering extensions

E. Cut through forwarding

ANSWER: C D

Explanation:

Route redistribution is a major OSPF capability because an Autonomous System Boundary Router can inject routes learned from other routing domains or routing protocols into the OSPF domain. OSPF represents these redistributed destinations as external routes and advertises them using AS-external LSAs, allowing OSPF routers to calculate reachability to networks outside the OSPF autonomous system. This function is fundamental when OSPF must interoperate with static routing, BGP, RIP, or another OSPF process. The OSPFv2 specification defines the handling of external routing information and the role of AS boundary routers: RFC 2328.

Traffic engineering extensions are also a major OSPF feature in service-provider networks. OSPF traffic engineering distributes additional link attributes, such as available bandwidth, administrative groups, and TE metrics, through TE LSAs. MPLS RSVP-TE path computation can use this topology and attribute information to select paths that satisfy resource and policy constraints rather than relying only on the normal shortest-path calculation. The OSPF traffic-engineering extensions and their use of opaque LSAs are specified in RFC 3630. Together, these capabilities enable OSPF to exchange both external reachability and enhanced topology information needed by advanced MPLS TE deployments.

QUESTION NO: 12

What are some of the characteristics of Nokia's implementation of non-stop routing?

(Choose two)

A. No protocol extensions required

B. Only supported by OSPF and IS-IS

C. Transparent to routing neighbors

D. Uses Graceful Restart to inter-operate with other vendors

ANSWER: A C

Explanation:

Nokia non-stop routing is designed to preserve the control-plane continuity of a router during a control-processor switchover. The active and standby routing processors synchronize routing protocol state so that, if the active processor fails, the standby processor can take over with the information needed to continue protocol operation. Consequently, *No protocol extensions required* is a characteristic of Nokia NSR: its operation does not require protocol-specific extensions or special NSR awareness from adjacent routers. This enables NSR to be used without changing the protocol behavior expected by standard routing peers.

Transparent to routing neighbors is also correct because neighboring routers should not need to re-establish adjacencies merely due to the local control-plane switchover. By retaining synchronized protocol and adjacency state, NSR aims to make the event invisible to peers and avoid routing disruption. Nokia distinguishes this approach from graceful restart mechanisms, which involve protocol behavior and neighbor support to preserve forwarding while a router restarts. NSR instead relies on redundant control-plane operation and state synchronization within the local system. Nokia's routing documentation discusses NSR as a high-availability feature and describes its relationship to routing-protocol continuity: Nokia SR OS System Management Guide. Additional SR OS routing and high-availability documentation is available from the Nokia SR OS Documentation Center.

QUESTION NO: 13

An interface is configured on a router for OSPF but the command "show router ospf neighbor" does not show any neighbors. Which of the following is NOT a possible cause?

A. An OSPF database description packet has not been received from the neighbor.

B. The interface facing the neighbor is down.

C. The interface is included in a different area than the neighbor router.

D. The router has not received a valid OSPF Hello packet from its neighbor.

ANSWER: A

Explanation:

An OSPF database description packet has not been received from the neighbor. is the correct answer because Database Description (DBD) packets are exchanged only after OSPF routers have already discovered one another through Hello packets and have created a neighbor-table entry. A valid received Hello packet first causes the router to identify the remote router ID and add it to the OSPF neighbor data structure. The neighbor is consequently visible in `show router ospf neighbor`, even while the adjacency is progressing through states such as ExStart or Exchange.

DBD packets are part of the later database-synchronization phase. They summarize the link-state advertisements known by each router so that the peers can determine which LSAs must be requested and synchronized. Therefore, a failure to receive a DBD packet can prevent an adjacency from reaching the Full state, but it does not account for the complete absence of a neighbor entry in the operational neighbor display. Nokia's OSPF documentation describes Hello processing and the subsequent adjacency/database-exchange procedures, while the OSPF base specification defines DBD exchange after neighbor discovery. See Nokia SR OS OSPF documentation and RFC 2328: OSPF Version 2.

QUESTION NO: 14

From the list below select the statements that describe the behavior of link state information aging. (Choose 2)

A. The age for newly created link state information is set to 0 for OSPF and 1200 for IS-IS

B. The age for newly created link state information is set to 1200 for OSPF and 0 for IS-IS

C. The age is incremented or decremented as it is flooded from router to router and as it is held in the topological database.

D. The age is incremented or decremented only as it is held in the topological database and not as it is flooded router to router.

ANSWER: A C

Explanation:

"The age for newly created link state information is set to 0 for OSPF and 1200 for IS-IS" is correct because the two link-state protocols represent aging in opposite directions. An OSPF router originates an LSA with an LS age of zero. The age increases while the LSA is retained and propagated, ultimately reaching the MaxAge value of 3,600 seconds unless it is refreshed or flushed earlier. IS-IS uses a remaining-lifetime model in its Link State PDUs. A newly generated IS-IS LSP is normally originated with the default maximum lifetime of 1,200 seconds, and that remaining lifetime counts down toward zero.

"The age is incremented or decremented as it is flooded from router to router and as it is held in the topological database" is also correct. OSPF age is incremented as time passes, including during flooding and while stored in the link-state database. IS-IS LSP lifetime is decremented under the same general conditions, reflecting the time remaining before expiration. This aging behavior ensures stale topology information is eventually removed if it is not refreshed by its originating router. Nokia's routing documentation describes the OSPF LS age and IS-IS LSP lifetime behavior; protocol specifications are available in RFC 2328 and RFC 1195.

QUESTION NO: 15

Which of the following prefix-lists is the most specific match for prefix 172.31.2.1/24?

A. Prefix 172.16.0.0/12 longer.

B. Prefix 172.31.0.0/22 longer.

C. Prefix 172.31.0.0/23 longer.

D. Prefix 172.31.0.0/24 exact.

ANSWER: B

Explanation:

Prefix 172.31.0.0/22 longer. is the most specific matching prefix-list entry. A /24 prefix is represented by its network address, so 172.31.2.1/24 belongs to the network 172.31.2.0/24. The aggregate 172.31.0.0/22 covers the address range from 172.31.0.0 through 172.31.3.255, encompassing four contiguous /24 networks: 172.31.0.0/24, 172.31.1.0/24, 172.31.2.0/24, and 172.31.3.0/24. Therefore, 172.31.2.0/24 is contained within that /22 block.

In Nokia SR OS routing policy terminology, the *longer* match type requires the evaluated route to be a subnet of the configured prefix and to have a prefix length greater than the configured prefix length. The evaluated /24 is longer than /22 and is within the /22 address range, satisfying both requirements. Among matching aggregate entries, /22 is the longest configured covering prefix and consequently provides the most specific match. This behavior is important when route policies use prefix matching to control the acceptance, advertisement, or modification of OSPF routes. Nokia documents route-policy prefix matching behavior in its SR OS Route Policies documentation; CIDR containment principles are also defined in RFC 4632.

QUESTION NO: 16

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:

OSPF defines exactly five packet types: Hello, Database Description, Link State Request, Link State Update, and Link State Acknowledgment. Therefore, *Hello, link state update, link state request, link state acknowledgement, database description.* correctly lists the complete set. Hello packets discover neighboring OSPF routers and maintain neighbor relationships. Database Description packets are exchanged while adjacent routers synchronize, summarizing the contents of their link-state databases. A router uses Link State Request packets to ask a neighbor for specific link-state advertisements that it needs. Link State Update packets carry one or more link-state advertisements, allowing routers to distribute topology information. Link State Acknowledgment packets reliably acknowledge received link-state advertisements, supporting OSPF's dependable flooding process. These packet types are assigned protocol values 1 through 5 in the OSPF header and form the basis for neighbor establishment, database synchronization, and ongoing topology propagation. See the OSPF version 2 packet-type definition in RFC 2328, Appendix A.3.2 and Nokia's OSPF documentation.

QUESTION NO: 17

Click the exhibit button.

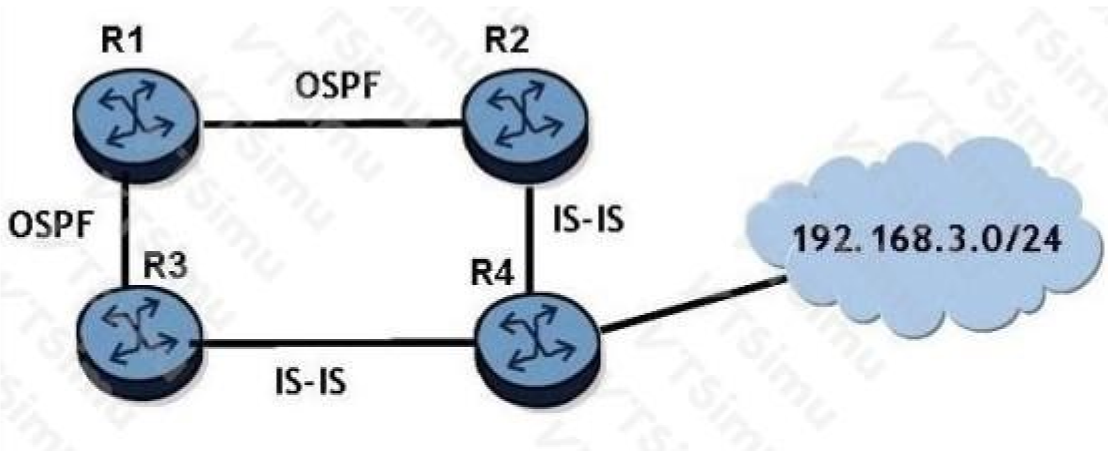

If router R2 redistributes the IS-IS route to 192.168.3.0/24 into OSPF, router R3 will receive two routes to 192.168.3.0/24. What will be the preference of these two routes? Assume that all IS-IS routers are L1/L2 capable and are in the same area. Choose two answers.

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:

R3 learns the prefix natively through IS-IS and also learns a second copy after R2 redistributes that prefix into OSPF. Because every IS-IS router is Level 1/Level 2 capable and belongs to the same IS-IS area, the native IS-IS advertisement is a Level 1 internal route. Therefore, the native copy at R3 has the IS-IS Level 1 internal preference. In SR OS route selection, this is the route class used for a destination learned internally within the local Level 1 area.

When R2 imports that same IS-IS destination into OSPF, the route does not become an OSPF intra-area or inter-area route. Redistribution causes it to be advertised as an OSPF autonomous-system external route, so R3's second route has the OSPF external preference. The two route entries are consequently an IS-IS Level 1 internal route and an OSPF external route. With Nokia's normal default preference behavior, the IS-IS Level 1 internal route is preferred over the OSPF external route for installation, while the OSPF external route remains a learned alternative. IS-IS Level 1 operation is defined in [RFC 1195], and OSPF external-route advertisement is specified in [RFC 2328].

QUESTION NO: 18

Which utility on the Nokia 7750 SR can be used to investigate an MPLS path?

A. Lsp-trace

B. Idp-trace

C. Sdp-ping

D. Svc-ping

E. Ping

ANSWER: A

Explanation:

Lsp-trace is the Nokia SR OS MPLS Operations, Administration, and Maintenance utility used to investigate an MPLS label-switched path. It traces the forwarding path followed by MPLS packets and identifies the transit nodes along that LSP. This provides operational visibility into the path actually used through the MPLS network, which is especially useful when validating LSP reachability, checking the path selected by signaling and routing, or isolating a forwarding issue at a particular hop. The command operates using MPLS OAM mechanisms: probes are sent with successively increasing time-to-live values, and responding LSRs reveal information that permits the operator to map the LSP hop by hop. On SR OS, the utility name is written `lsp-trace`; the initial lowercase character is the letter "l," not the letter "i." Nokia's SR OS documentation provides the platform command and MPLS-OAM context at [Nokia SR OS documentation]. The underlying MPLS LSP ping and traceroute procedures are standardized in [RFC 8029], which describes how echo requests and TTL expiration responses support LSP path discovery and fault isolation.

QUESTION NO: 19

What is the source IP address in an OSPF update?

A. 224.0.0.5

B. 224.0.0.6

C. The IP address assigned to the outgoing interface on which the update is sent.

D. The system address of the router

ANSWER: C

Explanation:

The IP address assigned to the outgoing interface is the source address of an OSPF Update packet. OSPF runs directly over IP, using protocol number 89, and the outer IPv4 header identifies the router interface that transmitted the packet. Consequently, a router originating an update on a particular OSPF-enabled link uses that link's configured interface address as the packet's source. This lets neighboring routers associate received protocol traffic with the directly connected adjacency and provides a reachable address for protocol exchanges where unicast delivery is used.

The destination used for an update can vary according to the OSPF network type and router role, but that does not alter the source-address rule. For example, updates on a broadcast network are commonly sent to an OSPF multicast destination, while the transmitting router still places its own egress-interface IP address in the source field. RFC 2328 explicitly specifies that OSPF packets have an IP source address equal to the address of the interface through which the packet is sent. See RFC 2328, OSPF Version 2 and Nokia’s OSPF configuration and operation documentation.

QUESTION NO: 20

What causes an adjacency to change from "Init" to "Two Way"?

A. When a Link State Update is received in response to a Link State Request

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

C. When OSPF neighbors exchange Database Descriptor packets

D. When a Link State Acknowledgement is received in response to a Link State Update

ANSWER: B

Explanation:

An OSPF neighbor transitions from the Init state to the Two-Way state when it receives a Hello packet from that neighbor and finds its own Router ID in the Hello packet’s Neighbor field. This is the bidirectional communication check: the receiving router already knows about the neighbor because it has received Hellos, and seeing its own Router ID listed confirms that the neighbor has also received its Hellos. OSPF therefore recognizes that two-way connectivity exists between the routers.

Hello packets are the mechanism used for neighbor discovery, parameter verification, maintenance of neighbor relationships, and detection of bidirectional communication. On multi-access networks, reaching Two-Way does not necessarily mean that every pair of routers will form a full adjacency; routers may remain in Two-Way depending on their designated-router role. However, the event that specifically drives the Init-to-Two-Way transition remains the appearance of the local Router ID in the neighbor’s received Hello packet. This behavior is defined in the OSPF neighbor state machine in RFC 2328, Section 10.5 and is also described in Nokia’s OSPF documentation.

QUESTION NO: 21

Which of the following statements regarding OSPF neighbor states are true? Choose three answers.

A. Link State Request packets are sent in the Loading state.

B. Database synchronization is complete in the Full state.

C. DROther routers normally remain in 2-Way state with each other on a broadcast network.

D. The Exchange state occurs only on point-to-point links.

E. A router originates Type 5 LSAs when a neighbor enters Init state.

ANSWER: A B C

Explanation:

In the Loading state, a router sends Link State Request packets for LSAs that are missing from its database or that require more recent instances. It remains in Loading until the requested LSAs have been received. A neighbor reaches the Full state when the databases have been synchronized.

On a broadcast network, routers that are neither the DR nor the BDR normally remain in the 2-Way state with each other. They form full adjacencies with the DR and BDR, reducing the number of required database exchanges on the shared segment. The OSPF neighbor state machine is described in RFC 2328, section 10.3.

QUESTION NO: 22

What is the default preference value for a static route in the Nokia 7750 SR router?

A. 0

B. 5

C. 10

D. 15

ANSWER: B

Explanation:

The default preference value for a static route on a Nokia 7750 SR running SR OS is **5**. Route preference is an administrative value used by the router to select between routes to the same destination that have been learned from different sources. A lower preference is more desirable, so the default value of 5 gives a configured static route a high level of trust in the route-selection process. This behavior is useful because static routes are explicitly provisioned by the operator and are commonly used for fixed reachability, default routing, or controlled backup designs.

In SR OS, the preference can be configured per static route when a design requires a different selection outcome. If multiple candidate routes have the same preference, the router proceeds with the normal route-selection process, including metric-related evaluation where applicable. The default static-route preference of 5 is therefore an important baseline when comparing a static route with dynamically learned routing information in the same routing table. Nokia's SR OS routing documentation describes static-route configuration and route preference behavior in the Static Routes documentation and the Route Table documentation.

QUESTION NO: 23

Click the exhibit button.

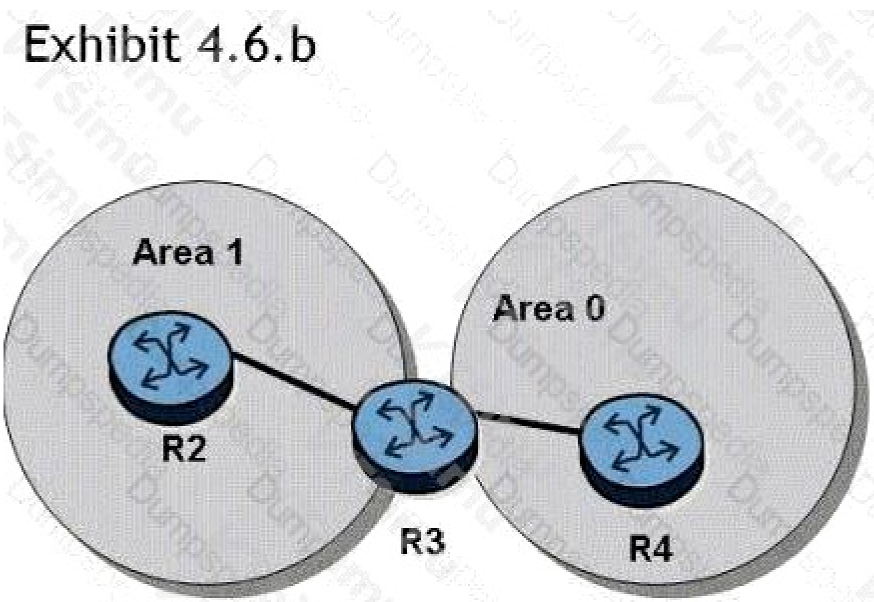

Which of the following would correctly configure routers R2, R3, and R4?

A.
```
*A:R3>config>router>ospf
        area 0.0.0.0
            interface "toR4"
        area 0.0.0.1
            interface "toR2
*A:R2>config>router>ospf
        area 0.0.0.0
            interface "toR3"
*A:R4>config>router>ospf
        area 0.0.0.0
            interface "toR3"
                nterface-type point-to-point
```

B.
```
*A:R3>config>router>ospf
        area 0.0.0.0
            interface "toR4"
                interface-type point-to-point
        area 0.0.0.1
            interface "toR2
*A:R2>config>router>ospf
        area 0.0.0.0
interface "toR3"
*A:R4>config>router>ospf
        area 0.0.0.0
            interface "toR3"
```

C.
```
*A:R3>config>router>ospf
        area 0.0.0.0
            interface "toR4"
                interface-type point-to-point
        area 0.0.0.1
            interface "toR2
                nterface-type point-to-point
*A:R2>config>router>ospf
        area 0.0.0.1
            interface "toR3"
                interface-type point-to-point
*A:R4>config>router>ospf
        area 0.0.0.0
            interface "toR3"
                interface-type point-to-point
```

```
D.    *A:R3>config>router>ospf
          area 0.0.0.0
            interface "toR4"
              interface-type point-to-point
          area 0.0.0.1
            interface "toR2
      *A:R2>config>router>ospf
          area 0.0.0.0
      interface "toR3" interface-type point-to-point
      *A:R4>config>router>ospf
          area 0.0.0.0
            interface "toR3"
              interface-type point-to-point
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: C

Explanation:

The configuration shown in *Option C* is the correct selection because it applies the OSPF parameters consistently to R2, R3, and R4 so that the routers can establish the required neighbor relationships and advertise the intended topology. OSPF operation depends on matching the fundamental interface and area settings required by the displayed design; after adjacency formation, each router originates the relevant link-state information and computes routes from a synchronized link-state database. This produces the expected reachability through the OSPF domain while retaining the area and interface roles specified in the exhibit.

In Nokia SR OS, OSPF is configured within a router instance and applied to the relevant interfaces and areas. The configuration must reflect the network design at both ends of every OSPF link so that the protocol can exchange Hello packets, form adjacencies, and flood link-state advertisements correctly. These behaviors follow the OSPF Version 2 protocol defined in [RFC 2328]; Nokia's product documentation is available through the [Nokia Documentation Center]. Therefore, the configuration represented by *Option C* correctly implements the exhibited R2–R4 OSPF design.

QUESTION NO: 24

The Nokia 7750 SR supports which two of the OSPF area types below?

A. Not So Stubby Areas

B. Level 1 Areas

C. Stub Areas

D. Partially Stubby Areas

ANSWER: A C

Explanation:

The Nokia 7750 SR supports both **Not So Stubby Areas** and **Stub Areas** for OSPF. A stub area is designed to reduce routing-table size and LSA flooding within an area by preventing external routing information from being propagated into it. Instead, the area border router provides a default route so routers in the stub area can send traffic for destinations outside the area toward the ABR. This is appropriate for an area with a single practical exit path to the rest of the OSPF domain.

A Not So Stubby Area (NSSA) retains the stub-area approach of limiting externally originated LSAs from the backbone while allowing an autonomous system boundary router located inside the area to inject external routes. Those routes are advertised as Type 7 LSAs within the NSSA and can be translated by an ABR into Type 5 external LSAs for distribution outside the NSSA. The SR OS OSPF configuration documentation includes support and configuration for stub and NSSA area behavior. See Nokia's [SR OS OSPF documentation]() and [RFC 3101]() for NSSA operation.

QUESTION NO: 25

Click on the exhibit.

```
A:MPLS_R1# show router mpls lsp "toR6-8" path

===============================================================================
MPLS LSP toR6-8 Path
===============================================================================
-------------------------------------------------------------------------------
LSP Name    : toR6-8                          To          : 10.10.10.8
Adm State   : Up                              Oper State  : Up
-------------------------------------------------------------------------------
Path Name                        Next Hop     Type       Out I/F   Adm   Opr
-------------------------------------------------------------------------------
toR6-1                           10.1.2.2     Primary    1/1/2     Up    Up
toR6-2                           10.1.3.3     Standby    1/1/3     Up    Up
toR6-3                           10.1.2.2     Standby    1/1/2     Up    Up
toR6-4                           n/a          Secondary  n/a       Up    Dwn
toR6-5                           n/a          Secondary  n/a       Up    Dwn
toR6-6                           n/a          Secondary  n/a       Up    Dwn
toR6-7                           n/a          Secondary  n/a       Up    Dwn
===============================================================================
```

If the primary path fails, which backup LSP path will the head-end router pick first?

A. The Standby with the highest next-hop IP.

B. The first Standby in the list.

C. The Standby with the longest uptime.

D. The Standby with the lowest next-hop IP.

ANSWER: C

Explanation:

The Standby with the longest uptime is selected first. On Nokia SR OS RSVP-TE LSPs, standby paths are maintained as backup candidates while the primary path is operational. When more than one eligible standby path exists at the same selection level, the head-end uses the standby path that has been up for the longest time. This behavior favors the candidate that has remained operational and signaled for the greatest duration, rather than basing selection on the numerical value of its next-hop address or merely on its displayed order.

The uptime criterion is relevant to the operational standby-path information shown in the exhibit: it identifies which candidate has been continuously established longest and therefore has precedence when the active primary path can no longer be used. After the standby path takes over, RSVP-TE continues to maintain the LSP according to its configured path-selection and resiliency behavior. Nokia's MPLS documentation describes RSVP-TE LSP path operation and the distinction between primary, secondary, and standby path types; see the [Nokia SR OS MPLS Guide]() and the [MPLS overview]().

QUESTION NO: 26

Click the exhibit.

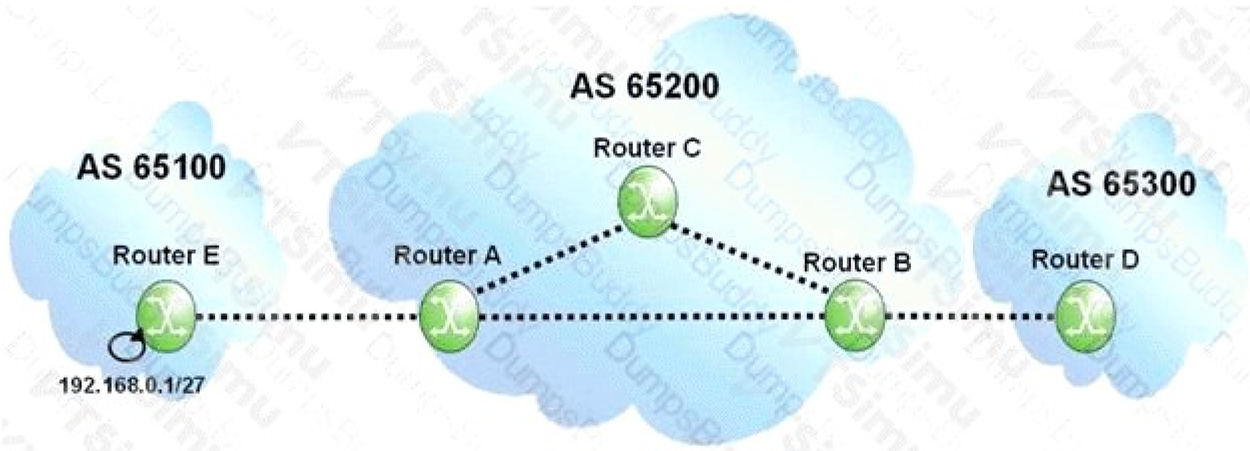

When router A receives the BGP update for the 192.168.0.1/27 prefix, it propagates it to routers B and C. What action is taken by router B when it receives the update?

A. Router B propagates the BGP update to routers D and C.

B. Router B propagates the BGP update to router D.

C. Router B detects duplicates in its BGP database and discards the update.

D. Router B propagates the update to routers A, C and D.

ANSWER: B

Explanation:

Router B propagates the BGP update to router D. Router B receives the prefix from router A over an iBGP relationship. A fundamental iBGP propagation rule is that a route learned from one iBGP peer is not advertised to another iBGP peer by default. This split-horizon behavior prevents routing-information loops within an autonomous system when a full iBGP mesh is used. The route remains eligible for advertisement from B to an eBGP neighbor, however. Therefore, B advertises the 192.168.0.1/27 route to router D through its external BGP session. When sending an eBGP advertisement, B applies the normal outbound BGP processing, including adding its own autonomous-system number to the AS path as appropriate. This behavior lets the route be distributed beyond the local AS while preserving iBGP's loop-prevention rule inside it. The base iBGP advertisement requirement is defined in RFC 4271, section 5.1.2; Nokia's BGP configuration and operational guidance is available in the Nokia SR OS BGP documentation.

QUESTION NO: 27

Identify two attributes of Distance Vector Protocols: (Choose two)

A. Use split horizon by default.

B. Send information about directly connected links to neighbors at periodic intervals.

C. Use holddown timers when topologies change to allow for convergence.

D. Maintain multiple databases: neighbor, topology and routing.

ANSWER: A C

Explanation:

"Use split horizon by default." and "Use holddown timers when topologies change to allow for convergence." are characteristic distance-vector routing behaviors. Distance-vector protocols make routing decisions from information learned from adjacent routers, rather than constructing a complete link-state topology map. Because routers can advertise routes back toward the neighbor from which they were learned, they require loop-prevention and convergence controls. Split horizon is a standard distance-vector loop-avoidance technique: route information learned through an interface is not advertised back through that same interface. This reduces the chance that two routers will incorrectly reinforce one another's routes after a failure.

Hold-down timers are also a classic distance-vector stability mechanism. When a route becomes unreachable or its metric deteriorates, a hold-down period limits acceptance of potentially stale or inferior updates. This gives failure information time to propagate through the routing domain and helps prevent transient routing loops and route flapping while the network converges. These mechanisms are commonly associated with RIP-style distance-vector operation. RFC 2453 describes split horizon as a technique used with RIP and discusses route timeout behavior: RFC 2453—RIP Version 2. Nokia's routing documentation also identifies split horizon as a route-advertisement loop-prevention behavior: Nokia SR OS Routing Protocols documentation.

QUESTION NO: 28

In an OSPF Hello packet what fields must match all neighbor routers on the segment?

(Choose 3)

A. Area ID

B. Hello and Dead Intervals

C. Stub flag

D. DR and BDR addresses

E. The list of neighbors

ANSWER: A B C

Explanation:

OSPF routers can form a neighbor relationship on a shared segment only when key Hello-packet parameters are compatible. **Area ID** must be identical because the interface on each router must participate in the same OSPF area. The area identifier is carried in the common OSPF packet header and prevents routers assigned to different areas from becoming neighbors on that link.

Hello and Dead Intervals must also agree. These timers establish how frequently Hellos are sent and how long a router waits without receiving a Hello before declaring its neighbor unavailable. Matching timer values ensures that both routers use a consistent liveness-detection process.

The **Stub flag**, represented through the Hello packet's OSPF options, must be consistent for routers within a stub area. This indicates whether the routers agree on the area's external-routing capability. Consistent area-type capability is required so that routers participating in that area maintain the same operational view of its routing behavior. These compatibility checks occur before OSPF establishes a full adjacency and exchanges topology information. See the OSPF Hello-packet and neighbor requirements in RFC 2328 and Nokia's OSPF documentation.

QUESTION NO: 29

What will happen if an update is received with an unrecognized optional non-transitive attribute?

A. The router will accept the update and send a notification.

B. The router will not accept the update and send a notification.

C. The router will accept the update and ignore the attribute.

D. The router will accept the update and propagate the attribute.

ANSWER: C

Explanation:

The router will accept the update and ignore the attribute. BGP path attributes contain flag bits that tell a receiving speaker both whether an attribute is optional and whether it is transitive. An optional attribute is not required to be recognized for the UPDATE message to be processed. When the attribute is also non-transitive, a BGP speaker that does not recognize it silently ignores it for its own path-selection processing and does not carry it forward in advertisements to other peers. Therefore, receipt of this attribute does not make the UPDATE malformed and does not require a NOTIFICATION message. This behavior lets vendors introduce optional, locally relevant BGP extensions without requiring every BGP implementation in the network to understand the extension. The transitivity bit is significant because it defines the handling of an unrecognized optional attribute at an intermediate BGP speaker: non-transitive information stops at that speaker rather than being retained for onward propagation. RFC 4271 defines these path-attribute categories and states that an unrecognized optional non-transitive attribute must be quietly ignored. See RFC 4271, Section 5: Path Attributes and RFC 4271, Section 5.1: Attribute Flags.

QUESTION NO: 30

An Nokia 7750 SR receives a route via an IS-IS LSP with internal reach ability, and receives the same route via an OSPF type 5 LSA. If al protocol preferences are default, which route will be installed in the route table?

A. The route learned via IS-IS will be installed.

B. The route learned via the OSPF type 5 LSA will be installed.

C. Both routes will be installed, if ECMP is enabled.

D. Both routes will be installed, regardless of whether ECMP is enabled.

ANSWER: A

Explanation:

The route learned via IS-IS will be installed. On SR OS, route selection first compares administrative preference when multiple routing protocols provide paths to the identical destination prefix. A lower preference value is more preferred. An IS-IS route advertised as internal reachability has the default IS-IS internal route preference, while a route carried in an OSPF Type 5 LSA is an OSPF autonomous-system external route and uses the default OSPF external preference. The default IS-IS internal preference is lower than the default OSPF external preference, so the IS-IS route is selected as the active route and entered in the route table. This decision is made before route metric comparison across these protocols; each protocol’s own metric is relevant only after the route-preference selection process has identified eligible routes of the preferred type. Equal-cost multipath does not alter this outcome because the competing routes do not have equal administrative preference. Nokia’s SR OS routing documentation describes preference as the mechanism used to choose between routes learned from different protocols, and its IS-IS documentation covers the internal reachability information advertised in IS-IS LSPs. See the Nokia SR OS routing protocols documentation and the Nokia SR OS IS-IS documentation.

QUESTION NO: 31

Which of the following conditions will prevent an OSPF adjacency from reaching the full state? Choose three answers.

A. MTU mismatch

B. Incorrect subnet mask

C. System interface not included in OSPF

D. Area ID not the same

E. Different metric set on each end of the link

F. Router ID not defined

ANSWER: A B D

Explanation:

MTU mismatch, **Incorrect subnet mask**, and **Area ID not the same** are OSPF neighbor-formation parameters that must be compatible for an adjacency to become fully operational. OSPF includes the interface MTU in Database Description packets. When a router receives a Database Description packet advertising an MTU larger than the receiving interface can support, database synchronization cannot proceed normally; the adjacency commonly remains in the ExStart or Exchange state rather than reaching Full.

On broadcast and NBMA network types, the subnet mask carried in OSPF Hello packets must agree with the receiving interface's mask. This validates that neighbors are attached to the same IP network before the neighbor relationship is accepted. The OSPF area identifier is also carried in Hello packets and must match the area configured on the receiving interface. Neighbor routers must belong to the same OSPF area on their shared link, so matching area IDs are essential before the routers can establish and synchronize an adjacency.

These requirements are defined in the OSPF neighbor and database-exchange procedures in RFC 2328, particularly the Hello-packet processing and Database Description packet rules. Nokia SR OS OSPF configuration guidance is available in the Nokia SR OS OSPF documentation.

QUESTION NO: 32

Which type of OSPF LSA has the following characteristics: It is flooded only within the area it originates from and can be originated by any OSPF router within the area (including non-

DR routers).

A. Type 1 Router LSA

B. Type 2 Network LSA

C. Type 3 Summary LSA

D. Type 4 ASBR LSA

ANSWER: A

Explanation:

Type 1 Router LSA is correct because every OSPF router originates one Router-LSA for each area to which it belongs. The LSA describes the router's links, interface-related metrics, and neighboring relationships within that specific area, allowing all routers in the area to construct the same link-state database and run the SPF calculation. Its flooding scope is area-local: it is flooded throughout the originating area but is not propagated unchanged into other OSPF areas. This behavior is fundamental to OSPF's hierarchical area design, which limits the size and scope of topology information outside an area.

The authority to originate a Router-LSA belongs to each participating router, not only to a designated router. Therefore, routers on point-to-point links, point-to-multipoint networks, and non-designated routers attached to broadcast or NBMA networks can originate their own Router-LSAs. The designated-router role affects network representation on multi-access segments, but it does not remove an individual router's responsibility to advertise its own reachable links in a Router-LSA. RFC 2328 defines Router-LSAs and their area flooding scope in Section 12.4.1, with the general LSA flooding procedure described in Section 13.

QUESTION NO: 33

Click the exhibit button.

Exhibit 6.1.c

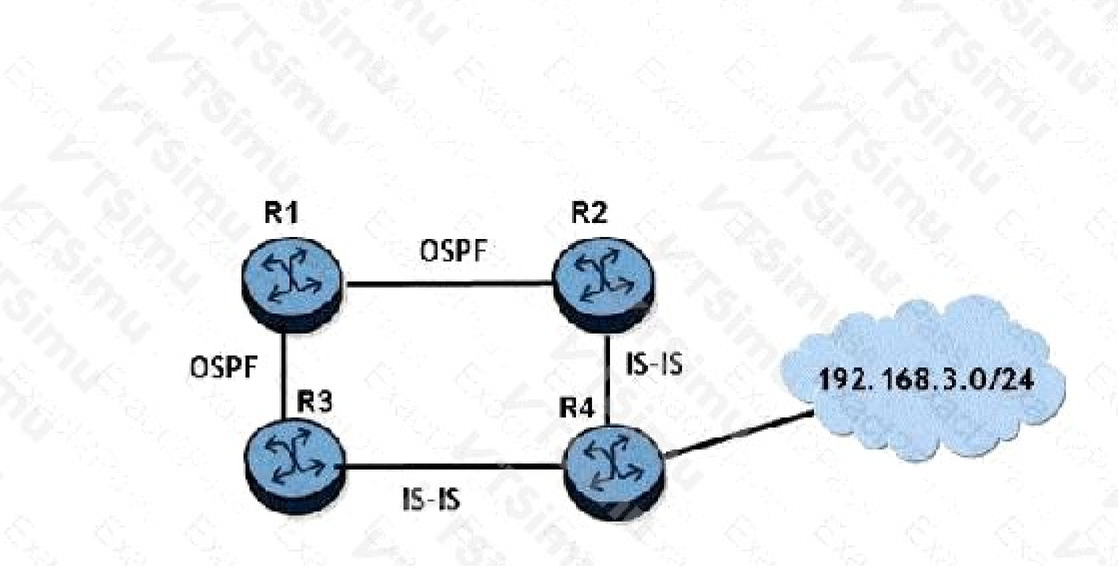

If router R2 redistributes the IS-IS route to 192.168.3.0/24 into OSPF, router R3 preference of these two routes? Choose two answers.

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:

R3 learns the original 192.168.3.0/24 route through its IS-IS Level 1 area, so that route is installed and evaluated as an **IS-IS Level 1 internal** route. In SR OS, route selection uses the route preference associated with the protocol and route type before considering metric values. The original intra-level IS-IS route therefore retains the IS-IS Level 1 internal route classification at R3.

When R2 redistributes that same IS-IS destination into OSPF, the redistributed advertisement is not an OSPF intra-area or inter-area route. Redistribution produces an OSPF autonomous-system external route, so R3 classifies the second learned path as an **OSPF external** route. This remains true even though the destination originally came from IS-IS: redistribution creates a new OSPF route advertisement and the receiving router evaluates it according to the OSPF external route type and its configured/default route preference. Nokia SR OS documentation describes protocol route preferences and route-table selection behavior in its router configuration material. OSPF defines redistributed autonomous-system destinations as external routes in RFC 2328; see also the Nokia SR OS documentation portal.

QUESTION NO: 34

Which of the following are similarities between OSPF and IS-IS? Choose two answers.

A. Metric is based on hop count.

B. Slow convergence after link failures.

C. Support hierarchy through the use of areas.

D. Support for CIDR.

E. Link state updates sent using layer 3 multicast.

ANSWER: C D

Explanation:

Both OSPF and IS-IS are link-state interior gateway protocols designed to scale within an autonomous system by organizing the network hierarchically. “Support hierarchy through the use of areas.” is correct because OSPF divides a routing domain into areas, with Area 0 providing the backbone for inter-area routing. IS-IS likewise supports a two-level hierarchy: Level 1 routers operate within an area, while Level 2 routers provide connectivity between areas. Although the terminology and operational details differ, both protocols use this hierarchy to limit flooding scope and reduce the size of the link-state database and shortest-path calculations.

“Support for CIDR.” is also correct. OSPF advertises IP prefixes with subnet masks in its link-state advertisements, allowing classless routing and variable-length subnet masks. Integrated IS-IS similarly carries IP reachability information with prefix masks, allowing classless prefixes to be propagated. This enables both protocols to support modern address planning, including route summarization at hierarchy boundaries and prefixes of arbitrary valid lengths rather than relying on classful network boundaries. The OSPF specification describes area-based operation and subnet-mask encoding in LSAs, while the Integrated IS-IS specification defines the hierarchical Level 1/Level 2 model and IP reachability information. See RFC 2328 and RFC 1195.

QUESTION NO: 35

A router receives an OSPF database description packet from a neighbor that references an LSA, which is already in the router's LSDB. The sequence number in the DBD packet is lower than the sequence number of the LSA in the LSDB.

What action does the router take?

A. The router sends a link-state update to its neighbor.

B. The router sends a link-state acknowledgement to its neighbor.

C. The router updates the information in its LSDB with the information received from its neighbor.

D. The router sends a link-state request to its neighbor.

E. The router takes no action.

ANSWER: A

Explanation:

The router sends a link-state update to its neighbor. A Database Description packet carries LSA headers rather than the complete LSA contents, allowing adjacent routers to compare their link-state databases during synchronization. Here, the router recognizes that its own LSDB contains a more recent instance of the LSA than the instance advertised in the neighbor’s Database Description packet, based on the LSA sequence number. The router must therefore provide the neighbor with its newer, authoritative copy so both databases can converge. It does this by transmitting the complete newer LSA in a Link State Update packet. This behavior avoids requesting information the router already knows is older and ensures the neighbor can install the latest LSA instance. OSPF determines LSA recency using the sequence number as part of its LSA-instance comparison rules; when an LSA in the local database is more recent than the received Database Description header, the local router sends that newer LSA to the neighbor. This is specified in the OSPF database-exchange processing rules in RFC 2328, section 10.8, with LSA recency comparison defined in RFC 2328, section 13.1.

QUESTION NO: 36

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 running SR OS, the OSPF router ID can be explicitly configured with the `config router ospf router-id X.X.X.X` command. Explicit configuration is the preferred operational practice because it provides a stable, predictable, unique 32-bit OSPF identifier that does not change as interfaces or addressing are modified. If an OSPF router ID is not explicitly configured, SR OS can derive it from the router's system IP address. The system IP address is intended to be a stable, router-wide address and is therefore an appropriate automatic source for an OSPF router ID. If neither an explicit OSPF router ID nor a system IP address is available, SR OS derives an identifier from the last 32 bits of the chassis MAC address. This provides a locally generated fallback value so that OSPF has a router ID available. These selection methods provide a clear hierarchy: configured OSPF router ID first, system IP address next, and chassis-MAC-derived value as the final fallback. Nokia's OSPF configuration documentation describes the router-ID configuration and behavior; see the Nokia SR OS OSPF documentation and the Nokia 7750 SR OSPF configuration guide.

QUESTION NO: 37

Which of the following about the components of a distributed service on the Nokia 7750 SR is TRUE?

```
R1>config>service# info
----------------------------------------------
        customer 1 create
            description "Default customer"
        exit
        customer 100 create
            description "Cust 100"
        exit
        sdp 2 mpls create
            far-end 10.10.10.2
            ldp
            keep-alive
                shutdown
            exit
            no shutdown
        exit
        epipe 12 customer 100 create
            sap 1/1/1:30 create
            exit
            spoke-sdp 2:44 create
                no shutdown
            exit
            no shutdown
        exit
----------------------------------------------
```

A. A SAP can only be used by one service on a router.

B. A dedicated SDP needs to be created for each service.

C. A service is created on each site before the customer is created.

D. SAPs must be created before a service is created.

ANSWER: A

Explanation:

A SAP can only be used by one service on a router. This is correct because a Service Access Point (SAP) is the service-specific attachment between a customer-facing physical port or logical interface and an SR OS service. When a SAP is configured, its encapsulation value identifies the traffic presented to that service, such as an untagged interface, a VLAN tag, or another supported encapsulation. SR OS associates that SAP with one service context, allowing the router to apply that service's forwarding, QoS, and customer policy treatment consistently to frames received at the access interface.

This one-to-one service association is an important part of service separation on the 7750 SR. Multiple services may use the same physical port when distinct encapsulations are used, for example separate VLANs, but each resulting SAP is individually bound to its own service. The SAP is therefore not a reusable attachment object that can simultaneously belong to multiple services on the same router. Nokia's SR OS service documentation describes SAPs as the access-side service attachment and covers their configuration under the relevant service context. See the [Nokia SR OS documentation portal] and the [7750 SR Services Overview documentation].

QUESTION NO: 38

Click the exhibit.

If router A originates a BGP route for prefix 192.168.0.1/27, what will the update contain when it reaches router B?

A. AS Path of 65200, Next Hop of router A, Origin of IGP.

B. AS Path of 65200, Next Hop of router A, Origin of incomplete.

C. Null AS Path, Next Hop of router A, Origin of IGP.

D. Null AS Path, Next Hop of router A, Origin of incomplete.

ANSWER: C

Explanation:

Null AS Path, Next Hop of router A, Origin of IGP. is correct. Router A originates the prefix directly into BGP using a BGP network origination, which creates a locally originated route. A locally originated BGP route has an empty AS_PATH because the route has not traversed any autonomous-system boundary. The ORIGIN attribute is set to IGP for a route introduced with the BGP *network* command; this name is historical and does not require that the prefix itself was learned through an IGP.

The exhibit's peering context is iBGP, so advertising the route from router A to router B does not add the local AS number to AS_PATH. Router A is also the BGP speaker that originates and advertises the NLRI, so its address is carried as the next-hop value shown for the received update. These behaviors follow the BGP path-attribute rules: AS_PATH is modified when advertising to an external peer, whereas an iBGP advertisement does not prepend the local AS. The origin code for routes injected through a BGP network statement is IGP. See [RFC 4271, A Border Gateway Protocol 4 (BGP-4)](), especially the AS_PATH, ORIGIN, and NEXT_HOP attribute definitions, and Nokia's [SR OS BGP documentation]().

QUESTION NO: 39

Click the exhibit button.

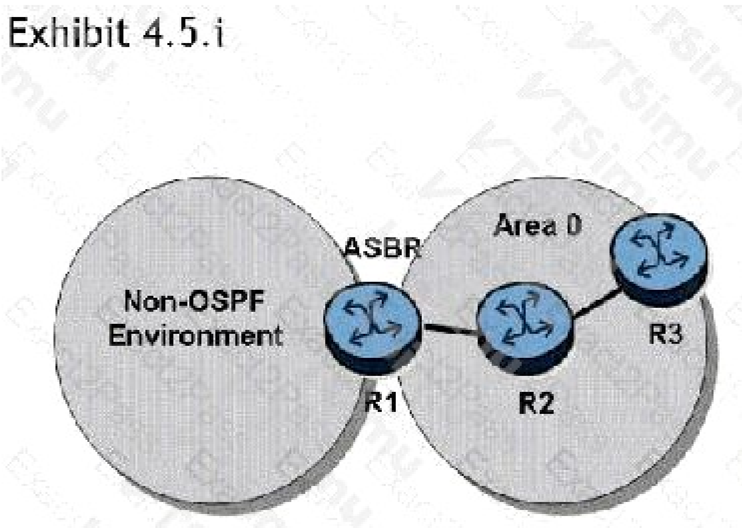

In the topology shown, router R1 is an ASBR configured to export external routes to OSPF. How many type 4 LSAs will be present in the network?

A. One.

B. One for each of the routers in area 0

C. One for each of the external routes exported by router R1.

D. Type 4 LSAs are not generated in this network topology.

ANSWER: D

Explanation:

Type 4 LSAs are not generated in this network topology. A Type 4 Summary ASBR LSA is an inter-area LSA generated by an ABR only when routers in another OSPF area need reachability information for an ASBR. Its purpose is to advertise the ASBR's router ID and the cost to reach that ASBR across an area boundary, allowing routers receiving Type 5 external LSAs to calculate a complete route toward the external destination.

In the shown topology, no separate-area advertisement of R1's ASBR reachability is required. Therefore, although R1 originates Type 5 AS-External LSAs for the redistributed routes, there is no requirement for an ABR to originate a Type 4 LSA. Type 4 LSAs describe reachability to an ASBR, not individual external prefixes, and their existence depends on inter-area topology rather than on the number of redistributed routes.

This behavior is defined in the OSPF version 2 specification's treatment of ASBR-summary LSAs and inter-area routing: [RFC 2328, OSPF Version 2](). Nokia's OSPF configuration and protocol documentation is available through the [Nokia SR OS OSPF documentation]().

QUESTION NO: 40

Which of the following are characteristics of an OSPF Area Border Router (ABR)? Choose three answers.

A. Maintains a separate link-state database for each attached area.

B. Originates summary LSAs for inter-area reachability.

C. Provides connectivity between a non-backbone area and the backbone area.

D. Must redistribute routes from another routing protocol.

E. Originates all Type 5 AS-external LSAs in the domain.

ANSWER: A B C

Explanation:

An ABR **maintains a separate link-state database for each attached area**, **originates summary LSAs for inter-area reachability**, and **provides connectivity between a non-backbone area and the backbone area**. OSPF uses areas to limit the scope of detailed topology flooding. An ABR participates in more than one area and performs the inter-area route calculations and advertisement functions needed to exchange reachability between those areas.

For IPv4 OSPFv2, an ABR originates Type 3 Summary-LSAs to advertise network destinations from one area into another. The ABR is distinct from an ASBR: an ASBR redistributes external routing information, while an ABR provides inter-area connectivity and summarization. The ABR functions and inter-area route processing are described in RFC 2328, Section 3.7 and Section 16.2.

QUESTION NO: 41

If a Router originates a type 1 router LSA with the E bit set what types of areas can it belong to? Select all that apply.

A. Not So Stubby Area

B. Stub Area

C. Backbone Area

D. Standard Area

ANSWER: A C D

Explanation:

In an OSPF Type 1 Router-LSA, the E bit indicates that the originating router is an Autonomous System Boundary Router (ASBR), meaning that it injects external routing information into the OSPF domain. An ASBR may operate in a standard area and may also be located in the backbone area, since both support normal external-route operation and flooding of Type 5 AS-external LSAs.

An ASBR can additionally belong to a Not-So-Stubby Area. NSSA functionality was specifically defined to allow external routes to be introduced within the NSSA. The ASBR originates Type 7 NSSA-external LSAs in that area; an Area Border Router translates eligible Type 7 LSAs into Type 5 LSAs for distribution into the rest of the OSPF domain. The E bit therefore remains a meaningful indication that the router has ASBR capability even though the NSSA uses Type 7, rather than Type 5, LSAs for its locally introduced external routes.

These area behaviors are defined by the OSPF Version 2 base specification and the NSSA extension specification. See RFC 2328 and RFC 3101.

QUESTION NO: 42

Which of the following concerning OSPFv3 is TRUE?

A. OSPFv3 uses MD5 Authentication.

B. Router LSAs carry IPv6 prefix information.

C. NSSA is supported in OSPFv3.

D. OSPFv3 uses a 64-bit router ID.

ANSWER: C

Explanation:

NSSA is supported in OSPFv3. A Not-So-Stubby Area provides the controlled external-route injection behavior of an NSSA while retaining the area restrictions associated with stub-style designs. In OSPFv3, an autonomous-system-boundary router within the NSSA can originate Type 7 AS-external LSAs for external IPv6 reachability. An area border router then translates eligible Type 7 LSAs into Type 5 AS-external LSAs for advertisement into the rest of the OSPF routing domain. This allows routes learned from another routing domain, connected networks, or redistribution policies to be introduced without making the NSSA a normal area that receives Type 5 external LSAs from elsewhere. OSPFv3 preserves this NSSA capability while adapting OSPF operation to IPv6, including the OSPFv3 LSA formats and IPv6 addressing model. RFC 5340 explicitly specifies support for NSSAs in OSPFv3 and identifies the applicable Type 7 external-LSA behavior. See RFC 5340, Not-So-Stubby Areas and the IETF’s OSPF for IPv6 specification.