PassQueen
a.k.a dumpsqueen.com

Alcatel-Lucent Interior Routing Protocols and High Availability

Alcatel-Lucent 4A0-101

Version Demo

Total Demo Questions: 44

Total Premium Questions: 442

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Volume A	81
Topic 2, Volume B	101
Topic 3, Volume C	98
Topic 4, Volume D	66
Topic 5, Volume E	76
Total	442

QUESTION NO: 1

What type of LSA is used to flood information about prefixes from one OSPF area to other attached areas?

A. Type 1

B. Type 3

C. Type 4

D. Type 7

ANSWER: B

Explanation:

Type 3 is correct because an OSPF Area Border Router (ABR) uses a Type 3 Summary LSA, also called an Inter-Area-Prefix LSA, to advertise reachable network prefixes from one area into another attached area. The ABR learns the intra-area topology and prefixes in each connected area, then originates summary advertisements so routers in other areas can calculate inter-area routes without receiving the complete link-state database for the originating area. This behavior supports OSPF's hierarchical area design: detailed topology information remains within its area, while inter-area prefix reachability is distributed through summary LSAs. An ABR can also perform route summarization when creating these advertisements, reducing the number of inter-area prefixes advertised and improving routing-table and LSDB scalability. In OSPFv2, this function is defined as the Type 3 Summary-LSA. In OSPFv3 terminology, the corresponding advertisement is called an Inter-Area-Prefix-LSA, but it serves the same fundamental purpose of communicating prefixes between areas. See the OSPFv2 specification in RFC 2328, Appendix A.4.3 and the OSPFv3 inter-area prefix definition in RFC 5340, Appendix A.4.3.

QUESTION NO: 2

How does an IS-IS router acknowledge the receipt of an IS-IS LSP on a point-to-point link?

A. They are explicitly acknowledged by sending a CSNP to the sender identifying the LSP.

B. They are implicitly acknowledged as a part of the link-state database synchronization and maintenance procedures.

C. They are explicitly acknowledged by sending a PSNP to the sender identifying the LSP.

D. They are implicitly acknowledged by the DIS multicast of CSNPs.

ANSWER: C

Explanation:

They are explicitly acknowledged by sending a PSNP to the sender identifying the LSP. On an IS-IS point-to-point adjacency, reliable flooding requires the receiving router to confirm that it has received a Link State PDU (LSP). It does this by transmitting a Partial Sequence Number PDU (PSNP) that includes the LSP identifier and the received sequence number, checksum, and remaining lifetime information. This provides an explicit acknowledgment to the neighbor and lets the sender remove the LSP from its retransmission processing for that adjacency. PSNPs can also identify LSPs that a router needs from its neighbor, making them part of the database synchronization and reliable LSP flooding mechanisms. The important point for a point-to-point link is that the acknowledgment is sent directly and explicitly as a PSNP, rather than depending on a designated router's periodic complete database summary. The IS-IS protocol specification describes PSNPs as serving the acknowledgment function on point-to-point links; its flooding procedures define the receipt and processing of these sequence-number PDUs. See RFC 1195, Partial Sequence Number PDUs and RFC 1195, Complete Sequence Number PDUs.

QUESTION NO: 3

Which LSA type is not found in point-to-point only OSPF networks?

A. Type 1 Router LSA

B. Type 2 Network LSA

C. Type 3 Summary LSA

D. Type 5 External LSA

ANSWER: B

Explanation:

Type 2 Network LSA is not found in an OSPF network composed only of point-to-point links. A Network LSA represents a transit multi-access network and is originated by the Designated Router for that network segment. Its purpose is to describe the routers attached to the shared network, allowing the OSPF link-state database to model that segment as a pseudonode. Point-to-point interfaces connect exactly two routers and do not elect a Designated Router or create a transit-network pseudonode. Instead, each router describes its point-to-point adjacency and interface information in its own Router LSA. Therefore, where every OSPF link is point-to-point, there is no eligible transit multi-access segment and no Designated Router that could originate a Type 2 Network LSA. This behavior follows the OSPF LSA definitions in RFC 2328, section 12.4.2, which specifies that Network LSAs are generated by a Designated Router for transit networks. Nokia SR OS OSPF documentation also identifies Network LSAs as describing multi-access networks and their attached routers: Nokia SR OS OSPF documentation.

QUESTION NO: 4

Which of the following are major features of OSPF? (Choose 2.)

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 routing information learned from other routing domains into OSPF. OSPF represents this imported information as external routes, allowing internal OSPF routers to calculate paths to destinations that are not natively advertised by the OSPF domain. This capability is essential when OSPF must interoperate with other routing protocols or with statically configured routes. The OSPF specification defines the handling of external routing information and the associated external link-state advertisements.

Traffic engineering extensions are also a major OSPF feature. OSPF traffic engineering distributes additional link attributes, such as available bandwidth and administrative constraints, throughout the routing domain. These attributes enable constraint-based path calculation, commonly for MPLS traffic-engineered tunnels, so paths can be selected using more than ordinary shortest-path cost. The base OSPF protocol and its traffic-engineering extensions therefore provide both topology awareness and the information needed for engineered path selection. See RFC 2328: OSPF Version 2 and RFC 3630: Traffic Engineering Extensions to OSPF Version 2.

QUESTION NO: 5

Click on the exhibit.

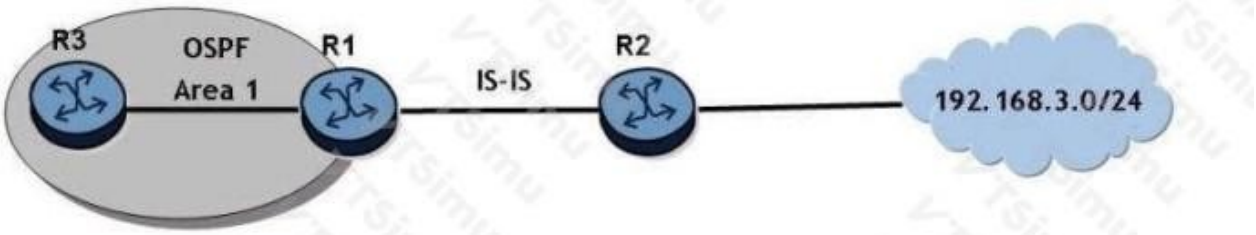


```
**A:AIRP_RO1>config>router>ospf# info
----------------------------------------------
            asbr
            export "isis-ospf"
            area 0.0.0.0
            exit
            area 0.0.0.1
                interface "toR3"
                    interface-type point-to-point
                    no shutdown
                exit
            exit
----------------------------------------------
*A:AIRP_RO1>config>router>ospf#
```

```
*A:AIRP_RO1# show router policy "isis-ospf"
    entry 10
        from
            protocol direct
        exit
        to
            protocol ospf
        exit
        action accept
        exit
    exit
*A:AIRP_RO1#
```

Router R1 learns the network 192.168.3.0/24 from IS-IS. Given the OSPF configuration shown, why isn't the 192.168.3.0/24 route in router R3's route table?

A. The route policy should be applied as an import policy.

B. The route policy should say "from protocol isis" instead of "from protocol direct".

C. The ASBR configuration should be removed.

D. The interface between routers R1 and R3 must be in OSPF area 0.

ANSWER: B

Explanation:

The route policy should say "from protocol isis" instead of "from protocol direct" because 192.168.3.0/24 was installed on R1 as an IS-IS-learned route. On SR OS, OSPF redistribution is controlled by the export policy under the OSPF configuration. Before R1 can advertise this destination into OSPF as an external route, the route must match a policy entry that accepts its source protocol. A match condition of `from protocol direct` selects only directly connected routes; it does not select routes learned through IS-IS. Consequently, the IS-IS route does not pass the export policy, R1 does not originate an OSPF external LSA for 192.168.3.0/24, and R3 has no OSPF route to install for that prefix. Changing the policy match to `from protocol isis` permits the intended IS-IS route to be exported by the OSPF ASBR, allowing downstream OSPF routers such as R3 to learn it, subject to normal OSPF reachability and area behavior. Nokia documents route-policy protocol matching and OSPF route export in its SR OS Route Policies guide and SR OS OSPF guide.

QUESTION NO: 6

Which of the following types of network topologies are supported by IS-IS on the Alcatel-Lucent 7750 SR? Choose two answers.

A. Non-broadcast

B. Broadcast

C. Point-to-multipoint

D. Point-to-point

ANSWER: B D

Explanation:

Broadcast and **Point-to-point** are the supported IS-IS interface network types on the 7750 SR. On a broadcast-capable LAN, such as an Ethernet segment, IS-IS uses LAN procedures: routers discover neighbors through IS-IS hello packets and elect a designated intermediate system for each level where applicable. The designated system originates the pseudonode representation used to describe the shared LAN topology in link-state advertisements. This allows multiple IS-IS routers attached to the same broadcast segment to maintain an efficient and consistent link-state database.

A point-to-point interface is used for a direct adjacency between exactly two IS-IS routers. The routers exchange point-to-point hello packets, establish adjacencies at the configured IS-IS levels, and advertise the direct link into the IS-IS topology. This model is appropriate for router-to-router links and does not require the LAN designated-system behavior used on shared broadcast media. These two interface models align with the IS-IS LAN and point-to-point operational procedures defined by the protocol and implemented in SR OS. See the IS-IS protocol specification in RFC 1195 and Nokia's SR OS IS-IS documentation.

QUESTION NO: 7

Click on the exhibit.

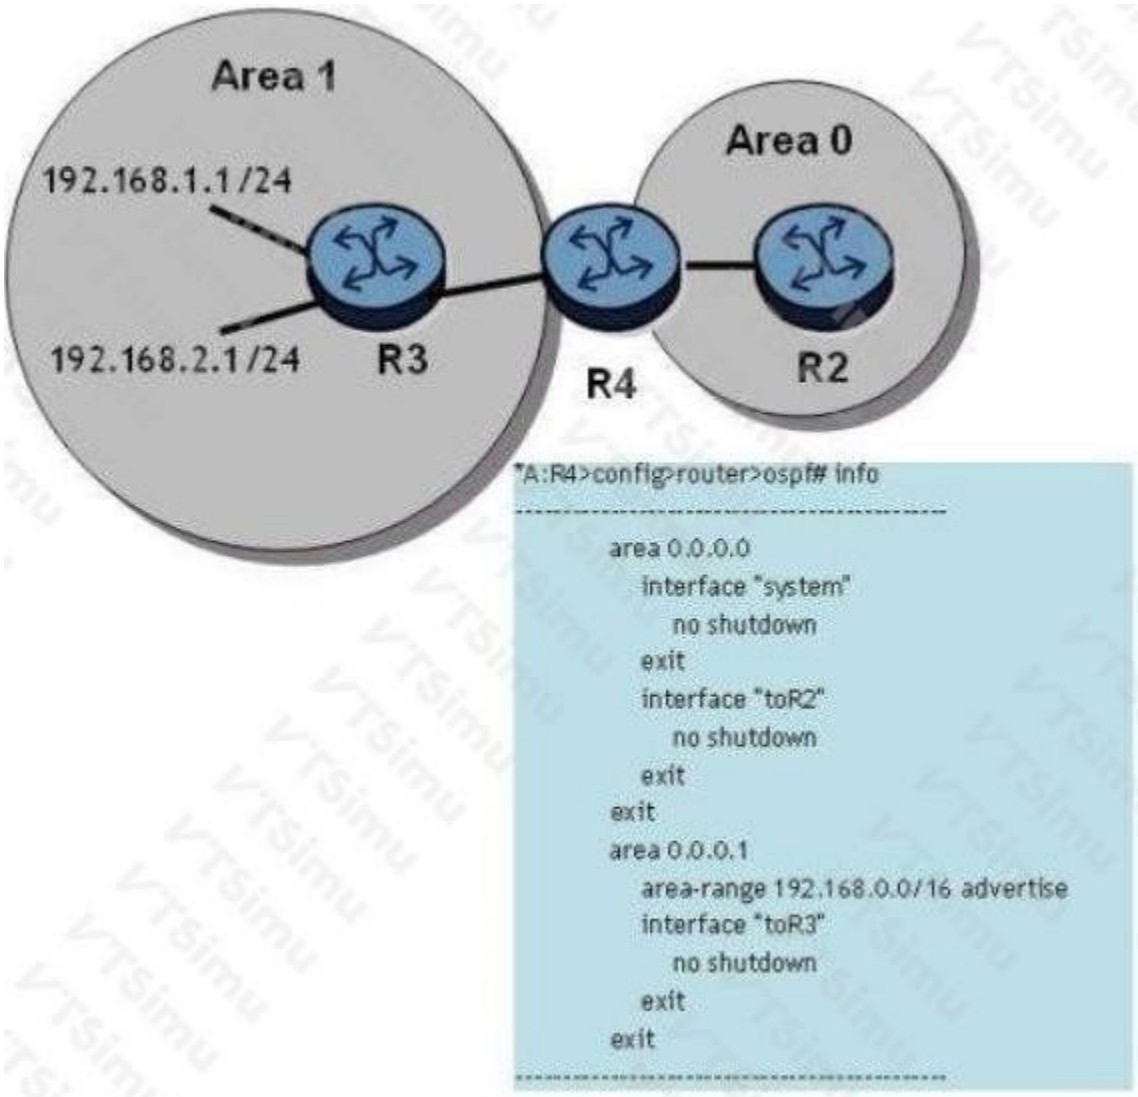

Which of the following would you expect to see in router R4's OSPF database?

A.

OSPF Link State Database (Type : All)

Type	Area Id	Link State Id	Adv Rtr Id	Age	Sequence	Cksum
Router	0.0.0.0	10.10.10.2	10.10.10.2	294	0x80000034	0x4adb
Router	0.0.0.0	10.10.10.4	10.10.10.4	81	0x80000033	0x8697
Network	0.0.0.0	10.2.4.4	10.10.10.4	653	0x80000025	0xdfdd
Summary	0.0.0.0	10.3.4.0	10.10.10.4	40	0x8000000a	0x737b
Summary	0.0.0.0	192.168.0.0	10.10.10.4	40	0x80000001	0xa6f8
Router	0.0.0.1	10.10.10.3	10.10.10.3	37	0x80000005	0xac12
Router	0.0.0.1	10.10.10.4	10.10.10.4	41	0x80000004	0x6d44
Network	0.0.0.1	10.3.4.4	10.10.10.4	46	0x80000001	0x4899
Summary	0.0.0.1	0.0.0.0	10.10.10.4	41	0x80000005	0x82b5

No. of LSAs: 9

B.

OSPF Link State Database (Type : All)

Type	Area Id	Link State Id	Adv Rtr Id	Age	Sequence	Cksum
Router	0.0.0.0	10.10.10.2	10.10.10.2	381	0x80000034	0x4adb
Router	0.0.0.0	10.10.10.4	10.10.10.4	168	0x80000033	0x8697
Network	0.0.0.0	10.2.4.4	10.10.10.4	741	0x80000025	0xdfdd
Summary	0.0.0.0	10.3.4.0	10.10.10.4	128	0x8000000a	0x737b
Summary	0.0.0.0	192.168.1.0	10.10.10.4	28	0x80000001	0x9b03
Summary	0.0.0.0	192.168.2.0	10.10.10.4	28	0x80000001	0x900d
Router	0.0.0.1	10.10.10.3	10.10.10.3	125	0x80000005	0xac12
Router	0.0.0.1	10.10.10.4	10.10.10.4	129	0x80000004	0x6d44
Network	0.0.0.1	10.3.4.4	10.10.10.4	134	0x80000001	0x4899
Summary	0.0.0.1	0.0.0.0	10.10.10.4	129	0x80000005	0x82b5

No. of LSAs: 10

C.

OSPF Link State Database (Type : All)

Type	Area Id	Link State Id	Adv Rtr Id	Age	Sequence	Cksum
Router	0.0.0.0	10.10.10.2	10.10.10.2	26	0x80000034	0x4adb
Router	0.0.0.0	10.10.10.4	10.10.10.4	157	0x80000031	0x8a95
Network	0.0.0.0	10.2.4.4	10.10.10.4	386	0x80000025	0xdfdd
Summary	0.0.0.0	10.3.4.0	10.10.10.4	119	0x80000008	0x7779
Summary	0.0.0.0	192.168.0.0	10.10.10.4	8	0x80000001	0xa6f8
Router	0.0.0.1	10.10.10.3	10.10.10.3	111	0x80000006	0x8c2f
Router	0.0.0.1	10.10.10.4	10.10.10.4	120	0x80000004	0x4f60
Network	0.0.0.1	10.3.4.4	10.10.10.4	125	0x80000001	0x2ab5
Summary	0.0.0.1	10.2.4.0	10.10.10.4	157	0x80000002	0x8540
Summary	0.0.0.1	10.10.10.2	10.10.10.4	157	0x80000002	0xcee6
Summary	0.0.0.1	10.10.10.4	10.10.10.4	157	0x80000002	0xce49

No. of LSAs: 11

D.

OSPF Link State Database (Type : All)

Type	Area Id	Link State Id	Adv Rtr Id	Age	Sequence	Cksum
Router	0.0.0.0	10.10.10.2	10.10.10.2	1421	0x80000033	0x4cda
Router	0.0.0.0	10.10.10.4	10.10.10.4	58	0x80000031	0x8a95
Network	0.0.0.0	10.2.4.4	10.10.10.4	287	0x80000025	0xdfdd
Summary	0.0.0.0	10.3.4.0	10.10.10.4	21	0x80000008	0x7779
Summary	0.0.0.0	192.168.1.0	10.10.10.4	21	0x80000001	0x9b03
Summary	0.0.0.0	192.168.2.0	10.10.10.4	21	0x80000001	0x900d
Router	0.0.0.1	10.10.10.3	10.10.10.3	13	0x80000006	0x8c2f
Router	0.0.0.1	10.10.10.4	10.10.10.4	22	0x80000004	0x4f60
Network	0.0.0.1	10.3.4.4	10.10.10.4	27	0x80000001	0x2ab5
Summary	0.0.0.1	10.2.4.0	10.10.10.4	58	0x80000002	0x8540
Summary	0.0.0.1	10.10.10.2	10.10.10.4	58	0x80000002	0xcee6
Summary	0.0.0.1	10.10.10.4	10.10.10.4	58	0x80000002	0xce49

No. of LSAs: 12

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: C

Explanation:

Option C is correct because an OSPF router's link-state database is populated from the Link-State Advertisements that are originated and flooded according to the topology and area boundaries shown in the exhibit. R4 maintains a synchronized database for its attached OSPF area and therefore contains the applicable router and network topology information represented in Option C. The database is not merely a list of routes selected for forwarding; it is the collection of received and self-originated LSAs from which the shortest-path-first calculation derives the routing table. Consequently, the expected database view must reflect the correct advertising routers, link-state identifiers, LSA types, and area-scoped topology information applicable to R4, as depicted by Option C.

OSPF defines the link-state database as an identical collection of area LSAs held by routers within an area, with flooding used to synchronize that information. The SPF algorithm then uses those LSAs to construct the area's shortest-path tree. This distinction between the LSDB and the installed routing table is fundamental when interpreting OSPF database output. See RFC 2328: OSPF Version 2 and Nokia's OSPF configuration and operation documentation.

QUESTION NO: 8

An Alcatel-Lucent 7750 SR has a port with the MAC address of 00;03:FA:AC:77: AF. The port is bound to an interface. If IPv6 is then enabled on that interface .what will the link-local address be?

A. FE80::203:FAFF:FEAC:77AF/64

B. 2000::203:FAFFFEAC:77AF/64

C. FE80::3FAFF:FEAC:77AF/64

D. FC00::203:FAFF:FEAC:77AF/64

E. FC00::3:FAFF:FEAC:77AF/64

ANSWER: A

Explanation:

FE80::203:FAFF:FEAC:77AF/64 is the link-local address formed from this interface MAC address. IPv6 link-local unicast addresses use the FE80::/64 prefix and commonly derive the 64-bit interface identifier from the interface's 48-bit MAC address through the modified EUI-64 procedure. Starting with 00:03:FA:AC:77:AF, the procedure inserts FF:FE between the first three and last three MAC octets, producing 00:03:FA:FF:FE:AC:77:AF. It then toggles the universal/local bit in the first octet: 00 becomes 02. The resulting interface identifier is therefore 0203:FAFF:FEAC:77AF, conventionally written without leading zeroes as 203:FAFF:FEAC:77AF. Combining this identifier with the FE80::/64 link-local prefix gives FE80::203:FAFF:FEAC:77AF/64. Link-local addresses are automatically applicable to the local Layer-2 segment and are fundamental to IPv6 neighbor discovery and routing-protocol adjacencies. The modified EUI-64 construction and the FE80::/64 link-local prefix are specified in RFC 4291, section 2.5.1; IPv6 address-autoconfiguration behavior is further described in RFC 4862.

QUESTION NO: 9

Which of the following best describes a unicast IPv6 address?

A. Fixed total of 128 bits, typically with 64 bits used for routing and 64 bits used as an interface identifier.

B. Fixed total of 128 bits, where 32 bits are used for the global prefix, 32 bits are used for the subnet and 64 bits are used as an interface identifier.

C. Variable total length, where at least 48 bits are used for the global prefix, 48 bits are used for the subnet and 64 bits are used as an interface identifier.

D. Variable total length, where at least 64 bits are used for routing and 64 bits are used as an interface identifier

ANSWER: A

Explanation:

An IPv6 unicast address is always 128 bits long. In the conventional IPv6 addressing model used for normal unicast interfaces, the address is divided into a 64-bit routing prefix and a 64-bit interface identifier. The routing prefix identifies the network and supports hierarchical routing, including provider and site/subnet allocation. The interface identifier identifies a specific interface on that subnet. This structure is particularly important for standard IPv6 LAN subnets, which normally use a /64 prefix; it supports Stateless Address Autoconfiguration and enables hosts to form or receive an interface identifier within the remaining 64 bits. Therefore, the description stating a fixed 128-bit total with 64 bits for routing and 64 bits for the interface identifier is the best description in this context. IPv6 architecture defines addresses as 128-bit values and describes

the common prefix/interface-identifier structure for unicast addressing. The IETF also documents why the 64-bit interface-identifier boundary is the established and operationally important convention for IPv6 networks. See RFC 4291, IPv6 Addressing Architecture and RFC 7421, Analysis of the 64-bit Boundary in IPv6 Addressing.

QUESTION NO: 10

What is the primary purpose of variable length subnet masks (VLSM)?

A. To reduce the size of each routing update.

B. To allow interoperability between routing protocols.

C. To allow different subnet masks for different subnets.

D. To reduce the size of the routing table.

ANSWER: C

Explanation:

To allow different subnet masks for different subnets is correct because Variable Length Subnet Masking (VLSM) lets an organization divide an allocated IP address block into subnets of unequal sizes. Rather than applying one uniform mask to every subnet, an administrator can assign a longer prefix where few hosts are needed and a shorter prefix where more hosts are needed. This enables more efficient use of IPv4 address space and supports hierarchical addressing plans that match real network requirements, such as point-to-point links, small branch LANs, and larger campus segments.

VLSM relies on classless addressing information: routes are identified by both the network prefix and its prefix length. In routing environments, classless protocols such as RIPv2, OSPF, and IS-IS can advertise or carry the relevant mask/prefix-length information, allowing routers to distinguish subnets of differing sizes within the same major address block. The fundamental purpose, however, is the ability to use different subnet masks where needed, improving address-allocation flexibility and reducing unnecessary consumption of addresses. See the IETF's classless IPv4 routing specification in RFC 1812 and Cisco's overview of VLSM and classless routing.

QUESTION NO: 12

Click the exhibit button.

Exhibit 4.6.g

```
*A:R3>config>router>ospf# show router route-table
Route Table (Router: Base)
===============================================================================
Dest Prefix                       Type   Proto   Age       Pref
    Next Hop[Interface Name]                                Metric
0.0.0.0/0                         Remote OSPF    00h00m23s  10
    10.1.3.1                                    101
10.1.3.0/27                        Local  Local   05d10h00m  0
    toR1                                        0
10.10.10.3/32                      Local  Local   0144d08h   0
    system                                      0
No. of Routes: 3

*A:R2# show router route-table 192.168.0.0/16 longer
Route Table (Router: Base)
===============================================================================
Dest Prefix                       Type   Proto   Age       Pref
    Next Hop[Interface Name]                                Metric
192.168.1.0/24                     Local  Local   01h09m41s  0
    loop1                                       0
192.168.2.0/24                     Local  Local   01h09m41s  0
    loop2                                       0
192.168.3.0/24                     Local  Local   01h09m40s  0
    loop3                                       0
192.168.4.0/24                     Local  Local   01h09m40s  0
    loop4                                       0
No. of Routes: 4
```

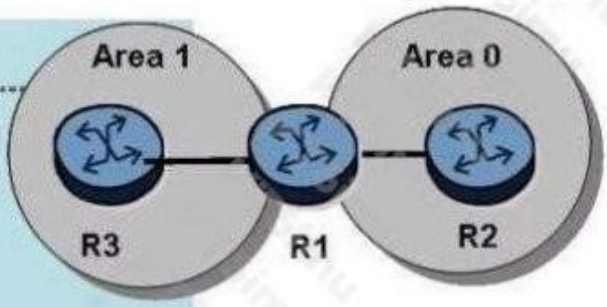

Given the topology and the show commands, and assuming that router R2 advertises all of its loopbacks into OSPF, what is the correct router R1 configuration?

A.
```
*A:R1>config>router>ospf# info
----------------------------------------------
        area 0.0.0.0
            interface "toR2"
                interface-type point-to-point
            exit
        exit
        area 0.0.0.1
            nssa
              originate-default-route
              no summaries
            exit
            interface "toR3"
                interface-type point-to-point
            exit
        exit
----------------------------------------------
*A:R1>config>router>ospf#
```

B.
```
*A:R1>config>router>ospf# info
----------------------------------------------
        area 0.0.0.0
            interface "toR2"
                interface-type point-to-point
            exit
        exit
        area 0.0.0.1
            nssa
              originate-default-route
            exit
            interface "toR3"
                interface-type point-to-point
            exit
        exit
----------------------------------------------
*A:R1>config>router>ospf#
```

C.
```
*A:R1>config>router>ospf# info
----------------------------------------------
        area 0.0.0.0
            interface "toR2"
                interface-type point-to-point
            exit
        exit
        area 0.0.0.1
            area-range 192.168.0.0/16 advertise
            interface "toR3"
                interface-type point-to-point
            exit
        exit
----------------------------------------------
*A:R1>config>router>ospf#
```

D.
```
*A:R1>config>router>ospf# info
----------------------------------------------
        area 0.0.0.0
            area-range 192.168.0.0/16 advertise
            interface "toR2"
                interface-type point-to-point
exit
            exit
        area 0.0.0.1
            interface "toR3"
                interface-type point-to-point
            exit
        exit
----------------------------------------------
*A:R1>config>router>ospf#
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: A

Explanation:

Option A is the correct router R1 configuration. The intended configuration must match the OSPF information shown in the exhibits and establish R1's required OSPF participation with R2 so that the loopback prefixes originated by R2 can be learned and installed by R1. In SR OS, OSPF is enabled under a routing-instance OSPF context, with interfaces assigned to the appropriate OSPF area. The interface addressing, area membership, and any required passive-interface treatment must align with the topology and operational output. Once adjacency is established on the transit link, R2's advertised loopback prefixes are carried in OSPF link-state advertisements, processed by SPF, and become available in R1's OSPF routing table according to the area and route type indicated by the displayed commands. This behavior follows the OSPF link-state database exchange and shortest-path calculation defined in RFC 2328. Nokia's SR OS documentation provides the platform-specific OSPF configuration hierarchy, interface association, and route-advertisement behavior in its SR OS documentation library. Therefore, the configuration represented by Option A is the one that satisfies the exhibit-derived OSPF requirements on R1.

QUESTION NO: 13

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 protocols represent age in opposite directions. An OSPF router originates an LSA with an LS age of zero. The age then increases over time, and OSPF also accounts for transmission delay when an LSA is flooded between routers. This lets routers determine when an LSA has reached its maximum age and must be flushed from the link-state database. RFC 2328 specifies both the initial LS age and the handling of LS age during flooding; see RFC 2328, section 12.1.1.

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 increases, whereas IS-IS advertises a remaining lifetime that begins at 1200 seconds and counts downward while the LSP is retained. Thus, both protocols age link-state information during its database lifetime, with flooding behavior also preserving the aging semantics needed to prevent stale information from persisting. IS-IS lifetime and purge behavior are defined in RFC 1195, section 7.3.14.

QUESTION NO: 14

Click on the exhibit.

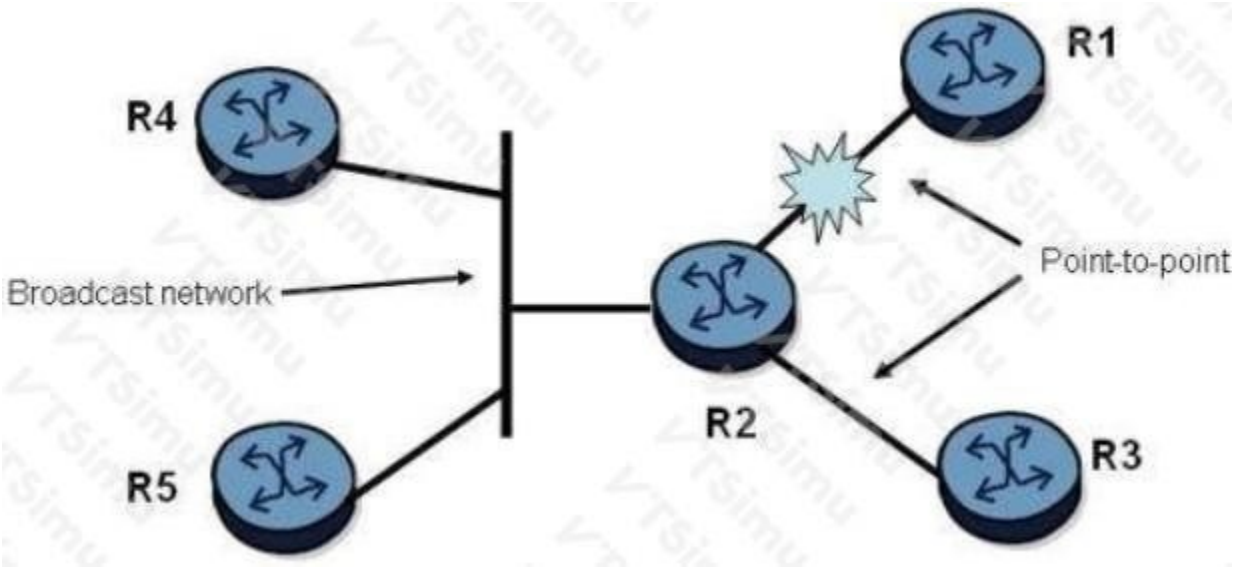

OSPF router R2 has both point-to-point and broadcast interfaces and it is not a DR for any interfaces. Router R2 detects the link failure between itself and router R1.

Which of the following is correct?

A. Router R2 sends an update to 224.0.0.5 on its broadcast and point-to-point interfaces. The DR does not send any updates.

B. Router R2 sends an update to 244.0.0.5 on its point-to-point interfaces. The DR sends an update to 224.0.0.6 on the broadcast network.

C. Router R2 sends an update to 224.0.0.5 out its point-to-point interfaces and an update to 224.0.0.6 out its broadcast interface. The DR sends updates to 224.0.0.5 on the broadcast network.

D. Router R2 sends an update to 224.0.0.6 out its point-to-point interfaces and an update to 224.0.0.5 out its broadcast interface. The DR sends an update to 224.0.0.6 on the broadcast networks.

ANSWER: C

Explanation:

Router R2 sends an update to 224.0.0.5 out its point-to-point interfaces and an update to 224.0.0.6 out its broadcast interface. The DR sends updates to 224.0.0.5 on the broadcast network. This follows OSPF's reliable flooding procedure after R2 detects the failed adjacency and originates the relevant changed link-state information. On a point-to-point link, there is no designated-router role, so R2 floods the Link State Update to the AllSPFRouters multicast address, 224.0.0.5. On a broadcast multi-access segment, because R2 is neither the DR nor the BDR, it sends the update to the AllDRouters address, 224.0.0.6, so that the DR can distribute it efficiently across the segment. The DR then refloods the update to AllSPFRouters, 224.0.0.5, allowing every OSPF router on that broadcast network to receive the topology change and run SPF as needed. This multicast behavior reduces duplicate flooding on shared media while retaining reliable synchronization of the link-state database. The OSPF flooding rules are defined in RFC 2328, section 13; the multicast address assignments are listed by the IANA IPv4 Multicast Address Space Registry.

QUESTION NO: 15

Which of the following statements regarding IS-IS Hello packets are true? (Choose two.)

A. Different types of Hello packets are exchanged on point-to-point networks than on broadcast networks.

B. L1/L2 Hello packets are exchanged on broadcast links.

C. Different L1 and L2 Hello packets are exchanged on broadcast links.

D. By default, Hello packets are sent every 5 seconds, and the timeout to declare a neighbor down is 15 seconds.

ANSWER: A C

Explanation:

"Different types of Hello packets are exchanged on point-to-point networks than on broadcast networks" is correct because IS-IS uses a Point-to-Point IS-to-IS Hello PDU on point-to-point circuits, while LAN circuits use LAN IS-to-IS Hello PDUs. The packet formats serve the differing adjacency and designated-router functions required by those network types. On a LAN, the hello exchange includes information used to maintain the shared-medium adjacency and support designated intermediate-system operation.

"Different L1 and L2 Hello packets are exchanged on broadcast links" is also correct. IS-IS maintains Level 1 and Level 2 operation independently, and a broadcast LAN uses distinct Level 1 LAN Hello and Level 2 LAN Hello PDUs. This permits routers to form and maintain adjacencies appropriate to each routing level on the same LAN interface. The original IS-IS protocol specification defines separate PDU types for Level 1 LAN hello, Level 2 LAN hello, and point-to-point hello processing. Nokia router documentation likewise distinguishes LAN and point-to-point IS-IS interface behavior and the associated hello processing. See RFC 1195 and Nokia SR OS IS-IS documentation.

QUESTION NO: 16

Which component of the Alcatel-Lucent 7750 SR is in charge of performing the longest prefix match lookup on packets that arrive on the physical interfaces?

A. CPM

B. MDA

C. Switch Fabric

D. IOM

ANSWER: D

Explanation:

IOM is correct because the Input/Output Module is the line-card subsystem that performs ingress packet processing for traffic received through its physical interfaces and associated MDAs. As part of this hardware-based ingress forwarding process, the IOM examines the packet headers and uses the forwarding information programmed into its lookup resources to identify the best route. For IP forwarding, that selection follows the longest-prefix-match rule: when more than one route could match a destination address, the route with the most-specific prefix is selected. The resulting forwarding decision determines the packet's required next-hop and egress handling before the packet is sent through the system's switching path. The routing protocols and control-plane functions provide and maintain the route information, but the high-speed per-packet lookup is performed in the forwarding hardware on the IOM, allowing the 7750 SR to forward traffic at line rate. Nokia's SR OS documentation describes the 7750 SR hardware architecture and the role of IOMs in packet forwarding; see the Nokia SR OS documentation portal and the 7750 SR Architecture Guide.

QUESTION NO: 17

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 an IGP, graceful restart is designed to preserve forwarding continuity while a router's control plane restarts. The restarting router signals its restart capability to adjacent routers, and its neighbors retain the relevant adjacencies and routing information for a negotiated or configured interval. This allows traffic to continue using the existing forwarding state while the restarting router rebuilds its protocol databases and re-establishes adjacencies, reducing packet loss and network disruption caused by a control-plane restart. OSPF implements this function through the Graceful Restart extension specified in RFC 3623, including helper behavior by neighboring routers. IS-IS graceful restart is specified in RFC 5306 and similarly enables neighboring systems to assist a restarting router while it recovers its link-state information. These are the link-state interior routing protocols for which the question's graceful-restart feature applies. See the IETF specifications for OSPF Graceful Restart and IS-IS Restart Signaling.

QUESTION NO: 18

Which of the following OSPF area types are supported by the Nokia 7750 SR? (Choose two.)

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 Nokia 7750 SR supports both **Not so stubby areas** and **Stub areas** as OSPF area configurations. A stub area is used to reduce the link-state database and routing information maintained by routers within the area. It does not carry external OSPF routes as Type 5 LSAs; instead, an Area Border Router can provide a default route so internal routers retain reachability to destinations outside the area. This is useful where an area has a single practical exit toward the rest of the OSPF domain.

A not-so-stubby area (NSSA) provides the stub-area behavior while also allowing an Autonomous System Boundary Router located inside the area to inject external routes. Those external routes are initially advertised as Type 7 LSAs and are translated by an Area Border Router for distribution beyond the NSSA when appropriate. Nokia SR OS documents the OSPF support and area configuration model for the 7750 SR in its OSPF documentation. The NSSA behavior and Type 7 external-LSA mechanism are standardized in RFC 3101.

QUESTION NO: 19

Click on the exhibit.

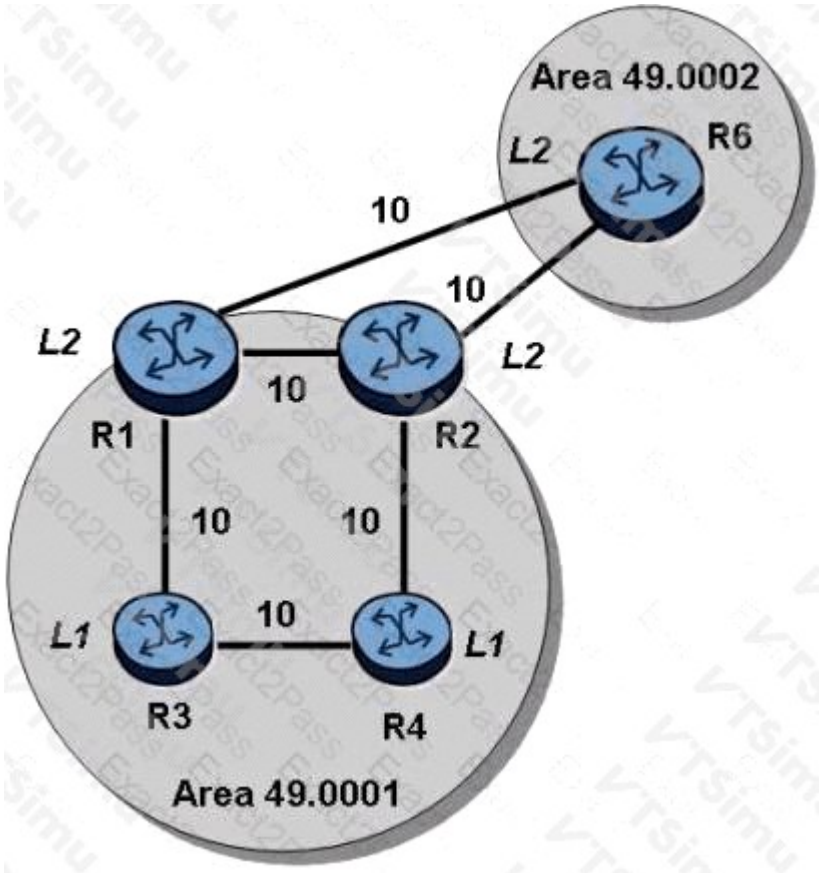

Examine the physical topology of the IS-IS network, the metrics of the links and the levels of the routers. All routers have a system address included in IS-IS.

Which of the following describes the route that router R4 will use to reach the system address of router R6?

A. Router R4 will have a route to router R6 ' s system address with router R2 as the next-hop.

B. Router R4 will have a default route with router R2 as the next-hop.

C. Router R4 will have a default route with router R3 as the next-hop.

D. Router R4 will not have a route that can be used to reach the system address of router R6.

ANSWER: D

Explanation:

Router R4 will not have a route that can be used to reach the system address of router R6. IS-IS route availability is governed by both the configured routing level and the area hierarchy represented in the link-state databases. A Level 1 router calculates routes only from the Level 1 information for its own area. It can use an attached Level 1–Level 2 router to leave the area only when that router advertises the appropriate Level 1 default route. A Level 2 topology does not, by itself, make every remote system address reachable from a Level 1 router. In the illustrated topology, the required Level 1 reachability or usable default path from R4 toward R6's system address is absent. Consequently, R4 cannot install a forwarding entry that resolves R6's system address, regardless of the physical links and their metrics. Metrics are considered only after a destination is present in the relevant IS-IS level's shortest-path calculation; they cannot create reachability across an unavailable level or area advertisement. This behavior follows the two-level IS-IS hierarchy defined in RFC 1195 and described in Nokia's IS-IS configuration documentation.

QUESTION NO: 20

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. An NSSA, or Not-So-Stubby Area, provides a controlled way to connect an OSPF area to an external routing domain while retaining the operational characteristics of a stub-like area. Within the NSSA, an Autonomous System Boundary Router can inject external routes as Type-7 LSAs. The 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 behavior is particularly useful where redistribution is needed at a branch, customer, or edge area, but unrestricted external-route flooding is undesirable. OSPFv3 retains this NSSA capability for IPv6 routing and specifies the associated Type-7 LSA handling and translation behavior. The OSPFv3 base specification explicitly describes NSSA support and references the NSSA functional model for the applicable area behavior. See the IETF's RFC 5340, OSPF for IPv6, especially its discussion of NSSAs, and RFC 3101, The OSPF Not-So-Stubby Area Option for the NSSA mechanism and Type-7 external advertisements.

QUESTION NO: 21

Click the exhibit button.

Exhibit 6.1.c

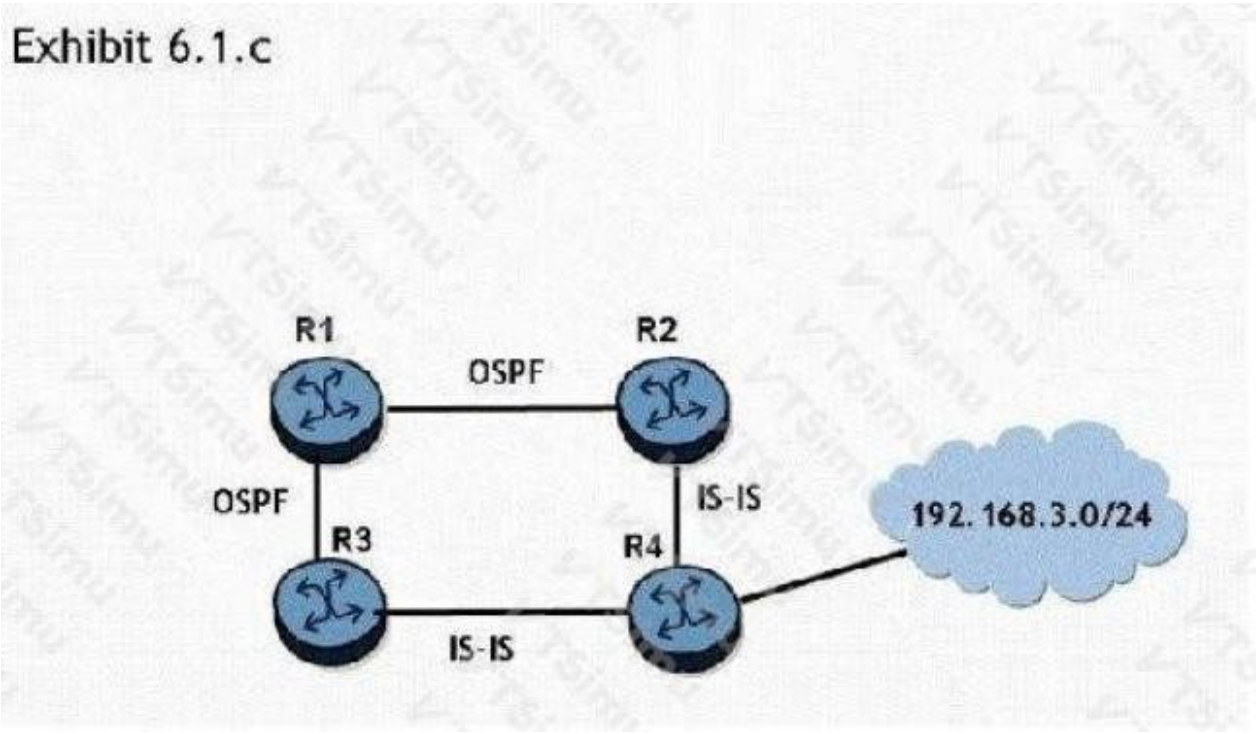

If router R2 redistributes the IS-IS route to 192.168.3.0/24 into OSPF, router R3 preference of these two routes? (Choose two.)

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 has two independently learned routes for 192.168.3.0/24 after redistribution at R2. The route learned directly through the Level 1 IS-IS domain is classified as an IS-IS Level 1 internal route. In SR OS, this route type has the default preference value of 15. R3 also receives an OSPF route resulting from redistribution of that IS-IS route by R2. Because the prefix originated outside OSPF and was injected into OSPF through redistribution, R3 classifies it as an OSPF external route. The default SR OS preference for an OSPF external route is 150. These are the two route-source classifications and associated preferences that apply to the same destination at R3. Preference is an administrative route-selection attribute: SR OS installs the route with the lower preference when otherwise considering competing routes for the same prefix. Consequently, the native IS-IS Level 1 internal route is preferred for forwarding, while the OSPF external route remains the alternative route learned through redistribution. Nokia's SR OS routing documentation describes route preference as the mechanism used to select among routes learned from different protocols; OSPF defines redistributed routes as AS-external routes. See Nokia SR OS route preference documentation and RFC 2328, OSPF Version 2.

QUESTION NO: 22

If routers R1 and R2 have a Level 1 and Level 2 IS-IS adjacency, which of the following statements is true? Choose two answers.

A. Routers R1 and R2 are in different areas.

B. Routers R1 and R2 are in the same area.

C. Routers R1 and R2 will send a Level 1 Hello PDU and Level 2 Hello PDU, if the interface type between them is point-to-point.

D. Routers R1 and R2 will send a Level 1 Hello PDU and Level 2 Hello PDU, if the interface type between them is broadcast.

E. Routers R1 and R2 will send a Level 1/2 Hello PDU if the interface type between them is broadcast.

ANSWER: B D

Explanation:

“Routers R1 and R2 are in the same area.” is correct because an IS-IS Level 1 adjacency can be formed only by systems sharing the same area address. Since the routers have both a Level 1 and a Level 2 adjacency, the Level 1 relationship establishes that their area membership matches. Level 2 operation provides interarea connectivity, but it does not remove the same-area requirement that applies to the concurrent Level 1 adjacency.

“Routers R1 and R2 will send a Level 1 Hello PDU and Level 2 Hello PDU, if the interface type between them is broadcast.” is also correct. On a broadcast LAN, IS-IS uses LAN Hello PDUs (IIHs), whose PDU type identifies whether the hello is for Level 1 or Level 2. A router operating at both levels therefore sends the appropriate Level 1 LAN IIH and Level 2 LAN IIH on that broadcast interface, allowing the corresponding adjacencies to be maintained. This behavior follows the IS-IS PDU definitions and LAN adjacency procedures in the Integrated IS-IS specification. See RFC 1195 and the IS-IS protocol specification, RFC 1142.

QUESTION NO: 23

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:

The E bit in an OSPF Type 1 Router-LSA identifies the originating router as an Autonomous System Boundary Router (ASBR). An ASBR redistributes routes from another routing domain into OSPF, so its Router-LSAs advertise that capability in every applicable attached area. A router can operate as an ASBR in the Backbone Area, which is the OSPF transit core and carries the inter-area topology required for domain-wide routing. It can likewise operate as an ASBR in a Standard Area, where normal OSPF external-route processing is supported. A Not So Stubby Area also permits an ASBR: the ASBR injects redistributed reachability into that area using Type 7 NSSA external LSAs, with an NSSA area border router translating eligible Type 7 information to Type 5 external LSAs for distribution beyond the NSSA. Therefore, the E-bit setting is valid for the listed Backbone Area, Standard Area, and Not So Stubby Area cases. RFC 2328 defines the Router-LSA E-bit as the ASBR indication in the Router-LSA flags; RFC 3101 defines the NSSA behavior that enables ASBR-originated external information within an NSSA. See RFC 2328 and RFC 3101.

QUESTION NO: 24

A static route is created using the command "static-route 0.0.0.0/0 next-hop 3.3.3.1". What command could be used to test the static route on an Nokia 7750 SR?

A. Ping static 3.3.3.1

B. Ping static 0.0.0.0/0

C. Ping 3.3.3.1

D. Ping 0.0.0.0

ANSWER: C

Explanation:

`Ping 3.3.3.1` is the appropriate command because 3.3.3.1 is the configured next-hop address for the static default route. On a Nokia 7750 SR, the operational `ping` command sends ICMP echo requests to a specified IPv4 destination and reports whether the destination can be reached. Successful replies from 3.3.3.1 confirm IP reachability to the next-hop router, which is necessary for the static route using that next hop to become usable for forwarding traffic. This is a practical first verification after configuring a route such as `static-route 0.0.0.0/0 next-hop 3.3.3.1`. The route can additionally be inspected in the route table to verify its active state, but pinging the configured next-hop directly tests the connectivity on which the static route depends. Nokia's 7750 SR command documentation describes `ping` as an IP connectivity diagnostic command and documents static-route next-hop configuration under router routing commands. See the Nokia 7750 SR Router Configuration Guide and the Nokia 7750 SR CLI Command Reference.

QUESTION NO: 25

Click the exhibit button.

What is the command to configure a static route on the 7750 service router R1 to reach the network behind router R2?

A. configure router static-route 139.120.121.0/24 next-hop 138.120.199.2

B. configure router static-route 139.120.121.0/24 next-hop 138.120.199.1

C. configure router static-route 138.120.199.2/24 next-hop 138.120.199.1

D. configure router static-route 139.120.121.0/24 next-hop 139.120.121.2

ANSWER: A

Explanation:

`configure router static-route 139.120.121.0/24 next-hop 138.120.199.2` is correct because a static route must identify the remote destination prefix and a reachable next-hop address on the directly connected transit network. The destination network behind R2 is `139.120.121.0/24`, while `138.120.199.2` is R2's address on the shared link between R1 and R2. R1 can therefore forward packets for the remote prefix to R2, which is the router responsible for onward delivery to that network. In SR OS, the `static-route` command is configured in the router context with a prefix in CIDR notation, and the `next-hop` parameter specifies the IP address to which traffic matching that prefix is sent. This creates a route with an unambiguous destination and a resolvable forwarding next hop, allowing R1 to install and use the route as long as the next hop remains reachable. Nokia's SR OS routing documentation describes static routes as manually configured routes that associate a destination prefix with a next-hop gateway or outgoing interface. See the Nokia SR OS static routes documentation and the Nokia SR OS Router Configuration Guide.

QUESTION NO: 26

Click the exhibit button.

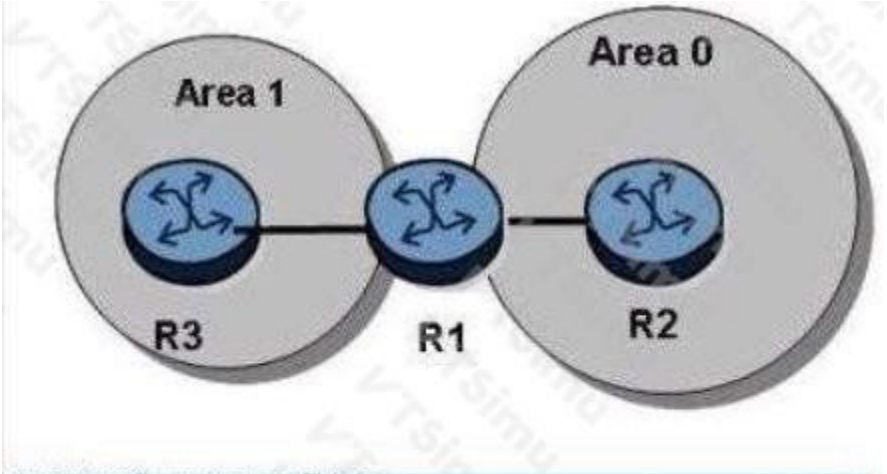


```
*A:R3>config>router>ospf# info
----------------------------------------------
        area 0.0.0.1
            stub
            exit
            interface "system"
            exit
            interface "toR1"
                interface-type point-to-point
                authentication-type password
                authentication-key
"i/XVK6X/L7YM44ZaNqVrJKWb5q244dGbfRvF4eS5KH." hash2
            exit
        exit
----------------------------------------------
*A:R3>config>router>ospf#
```

```
*A:R1>config>router>ospf# info
----------------------------------------------
        area 0.0.0.0
            interface "toR2"
                interface-type point-to-point
            exit
        exit
        area 0.0.0.1
            stub
            exit
            interface "system"
            exit
            interface "toR3"
                interface-type point-to-point
                authentication-type password
                message-digest-key 1 md5
"d7oaaU5pMrRRJ7Xl3w7bZxF9h6rw3DOoYBUOKD.mtFs" hash2
            exit
        exit
----------------------------------------------
*A:R1>config>router>ospf#
```

The OSPF adjacency between routers R3 and R1 is down. What is the problem?

A. Area 1 of router R1 should not be configured as a stub.

B. Router R3 has the incorrect authentication key type.

C. Router R1 has the incorrect authentication key type.

D. There is not enough information to tell.

ANSWER: C

Explanation:

Router R1 has the incorrect authentication key type. OSPF neighbors form an adjacency only after their Hello packets successfully pass the protocol's compatibility checks. Authentication is part of that validation: the authentication method and associated authentication parameters used on the shared OSPF interface must be compatible at both ends. The exhibit identifies that R1 is configured with a different authentication key type from the one used by R3 on their common link. Consequently, R1 and R3 cannot authenticate each other's OSPF Hellos, so they do not progress to a valid neighbor relationship and the adjacency remains down.

On SR OS, OSPF interface authentication is configured as an interface-level attribute, and matching authentication configuration is required between adjacent routers. Correcting R1 so that its key type matches R3's configuration allows authenticated OSPF Hellos to be accepted and permits adjacency establishment, assuming the remaining normal OSPF neighbor parameters are aligned. Nokia's OSPF configuration guidance is available in the SR OS OSPF documentation; the protocol authentication behavior is also defined in RFC 2328.

QUESTION NO: 27

What command will show the OSPF neighbors, and their status on the Nokia 7750 SR?

A. Show ospf neighbors

B. Show router ospf neighbors

C. Show router ospf adjacency

D. Show ospf adjacency

ANSWER: B

Explanation:

`Show router ospf neighbors` is the operational CLI command used to display OSPF neighbor information on a Nokia 7750 SR. It is entered from the operational environment and queries the OSPF protocol instance under the router context. The resulting output identifies discovered OSPF peers and provides the adjacency state for each, such as whether the relationship is progressing through database synchronization or has reached the Full state. This is a primary verification command when checking that OSPF adjacencies have formed as expected and when diagnosing routing-protocol connectivity. The display also supplies useful peer details, including the neighbor router ID, interface association, priority, and timing-related values, enabling an operator to correlate an adjacency with its local interface and network segment. Nokia's SR OS OSPF documentation describes the neighbor operational display as part of OSPF monitoring and troubleshooting, while the SR OS command reference documents the router OSPF show-command hierarchy. See the Nokia SR OS OSPF configuration guide and the Nokia SR OS command reference.

QUESTION NO: 28

Which of the LSA types stay within an OSPF area and are not exported outside of the area?

(Choose two.)

A. Router LSA

B. Network LSA

C. Summary Network LSA

D. Summary Router LSA

ANSWER: A B

Explanation:

Router LSA and Network LSA have area-local flooding scope in OSPF. Every router originates one Router LSA for each area to which it is attached. This LSA describes the router's links, interface costs, and neighboring relationships as seen within that specific area, allowing all routers in the area to construct the same area topology database and run the shortest-path-first calculation.

Network LSA is also restricted to its originating area. On a multi-access network that has a designated router, the designated router originates this LSA to describe the routers attached to that transit network. Together, Router LSAs and Network LSAs provide the topology information needed for intra-area routing. They are flooded throughout the relevant area but are not carried unchanged across an area boundary. This containment supports OSPF's hierarchical design, keeping detailed topology knowledge local to each area while allowing inter-area reachability to be represented separately.

RFC 2328 defines Router LSAs as Type 1 and Network LSAs as Type 2, and specifies their area flooding scope. See RFC 2328, Router-LSAs and RFC 2328, Network-LSAs.

QUESTION NO: 29

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 and maintain a neighbor relationship only when fundamental Hello-packet parameters describe the same logical OSPF domain and compatible adjacency behavior. **Area ID** must match because an OSPF interface belongs to a specific area, and neighbors exchange topology information within that shared area context. **Hello and Dead Intervals** must also match. The Hello interval governs how frequently neighbor-maintenance packets are sent, while the Router Dead interval determines how long a router waits before declaring a silent neighbor unavailable; identical values ensure both routers use the same liveness expectations. The **Stub flag** must match as well. This is carried in the OSPF Hello options field as the external-routing-capability indication (the E-bit). Its value identifies whether the area supports external routing information, so a mismatch would mean the routers disagree about the area type and cannot establish an adjacency. These requirements are defined by the OSPF neighbor compatibility checks in the base OSPF specification and apply to Nokia/Alcatel-Lucent router implementations of OSPF. See RFC 2328, Receiving Hello Packets and RFC 2328, Hello Packet Format.

QUESTION NO: 30

Which of the LSA types stay within an OSPF area and are not exported outside of the area? (Choose all that apply)

A. Router LSA

B. Network LSA

C. Summary Network LSA

D. Summary Router LSA

ANSWER: A B

Explanation:

Router LSA and Network LSA are area-scoped OSPF link-state advertisements. A Router LSA is originated by every router and describes that router's links, interfaces, and OSPF adjacencies within the specific area. A Network LSA is originated by the designated router on a multi-access segment and describes the routers attached to that transit network in that same area. This detailed topology information is used by routers to construct the area's link-state database and run the shortest-path-first calculation for intra-area routing.

Because these advertisements represent local topology, they are flooded only to routers belonging to their area. An area border router maintains a separate link-state database for each attached area and does not propagate these detailed Router LSA and Network LSA records into another area. Instead, when inter-area reachability must be advertised, the area border router creates summarized inter-area information. This separation is central to OSPF's hierarchical design: it limits topology flooding and SPF scope while allowing areas to exchange reachable prefixes through ABRs. RFC 2328 identifies Router LSAs as Type 1 and Network LSAs as Type 2, both having area flooding scope; see RFC 2328, Router-LSAs and RFC 2328, Network-LSAs.

QUESTION NO: 31

What are some of the characteristics of Alcatel-Lucent'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:

Alcatel-Lucent's Non-Stop Routing (NSR) is designed to preserve control-plane routing operation when the active control module fails and the standby module takes over. **No protocol extensions required** is a key characteristic because NSR synchronizes applicable routing-protocol state between the active and standby control modules inside the router. The recovery mechanism is therefore implemented locally rather than requiring a special capability to be negotiated with every peer.

Transparent to routing neighbors is also correct. During a successful NSR switchover, the router retains the protocol and forwarding state needed to continue operating without forcing adjacent routing devices to rebuild their neighbor relationship solely because the control module changed. This minimizes routing disruption, avoids unnecessary reconvergence, and helps maintain forwarding continuity. The behavior is particularly valuable in high-availability deployments because peer routers can continue treating the local router as the same routing neighbor while the redundant control plane assumes responsibility.

Nokia's SR OS documentation describes high-availability and redundancy mechanisms, including NSR operation and the synchronization of protocol state between redundant control modules. See the Nokia SR OS documentation portal and the SR OS Router Configuration Guide NSR documentation.

QUESTION NO: 32

Click the exhibit button.

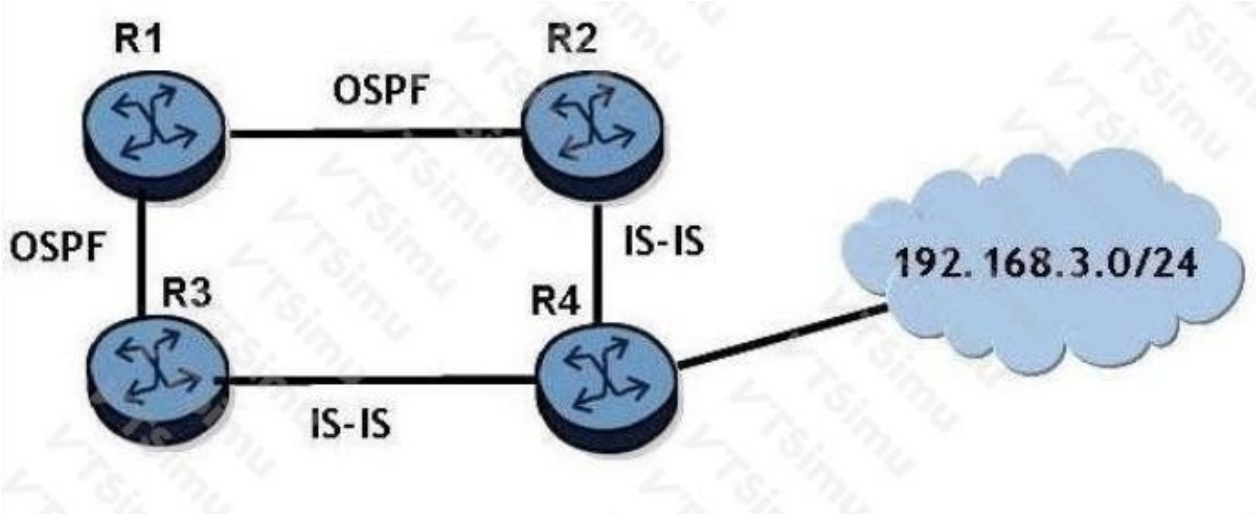

If router R2 redistributes the IS-IS route to 192.168.3.0/24 into OSPF, router R3 will receive two routes to 192.168.3.0/24. What will be the preference of these two routes? Assume that all IS-IS routers are L1/L2 capable and are in the same area. (Choose two.)

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 destination natively through IS-IS and also through the OSPF redistribution performed by R2. Because all IS-IS routers are Level 1/Level 2 capable and belong to the same IS-IS area, the native IS-IS advertisement for 192.168.3.0/24 is an intra-area Level 1 route at R3. Therefore, its route type uses the IS-IS Level 1 internal preference. This classification is based on the route's origin within the same IS-IS area, rather than on an inter-area or externally injected IS-IS route type.

The second copy reaches R3 through OSPF after R2 redistributes the IS-IS route. Redistribution makes the destination an external route from OSPF's perspective: it originated outside the OSPF domain and was injected by an autonomous-system boundary router. Consequently, this route uses the OSPF external preference. Nokia SR OS documents the protocol-specific route preference treatment for IS-IS and OSPF routes in its unicast-routing documentation, including the distinction between internal and external route sources. See the Nokia IS-IS documentation and Nokia OSPF documentation.

QUESTION NO: 33

What command will show the OSPF neighbors, and their status on the Alcatel-Lucent 7750 SR?

A. Show ospf neighbors

B. Show router ospf neighbor

C. Show router ospf adjacency

D. Show ospf adjacency

ANSWER: B

Explanation:

`show router ospf neighbor` is the SR OS operational command used to display OSPF neighbor information and the current state of each adjacency. It is entered from the operational CLI and queries the OSPF process configured under the router context. The resulting display identifies neighboring routers and includes useful adjacency details such as the neighbor router ID, interface, address, priority, dead-time information, and the OSPF neighbor state. This makes it the appropriate verification command when confirming that OSPF peering has formed and determining its present status, such as whether an adjacency has reached the expected Full state. The command can also be refined with supported context or interface qualifiers when troubleshooting a particular OSPF instance or link. Nokia's SR OS OSPF documentation describes the neighbor relationship and its state-machine behavior, while the router command reference documents the associated operational show commands. See the Nokia SR OS OSPF guide and the Nokia SR OS OSPF command reference.

QUESTION NO: 34

Which of the statements about the output in the exhibit is FALSE?

A. The system interface is configured with IPv6 address 2001DB8 1:100:1/128

B. All interfaces are configured for IPv4 and IPv6.

C. Interface " toR2 " is configured with a globally routed IPv6 address.

D. Interface " toR3 " is not configured for IPv6 with a global IPv6 address.

ANSWER: B

Explanation:

"All interfaces are configured for IPv4 and IPv6." is the false statement because the displayed interface information does not show a dual-stack configuration on every interface. In SR OS, IPv4 and IPv6 addressing are configured and displayed independently for an interface. The presence of an IPv4 address does not imply that an IPv6 address has also been configured, and an IPv6 link-local address alone is not equivalent to a configured global IPv6 address. The exhibit distinguishes the system interface and the routed interfaces by their address-family information, and it shows that at least one interface does not have the required IPv6 global-address configuration. Therefore, it is not accurate to generalize the output as showing IPv4 and IPv6 configured on all interfaces. IPv6 address syntax and prefix notation, including use of a /128 prefix for a host or system address, are defined in RFC 4291, IP Version 6 Addressing Architecture. Nokia SR OS interface configuration treats IPv4 and IPv6 as separate protocol-family settings, as described in the Nokia SR OS interface configuration documentation.

QUESTION NO: 35

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. This LSA describes the router's own links, interfaces, connected networks, link costs, and adjacent neighbors as they apply within that specific OSPF area. Consequently, its flooding scope is area-local: it is distributed throughout the originating area but is

not flooded unchanged into other areas. This behavior lets all routers in the area construct the same area topology database and independently run the SPF calculation.

A designated router role is not required to originate a Type 1 Router LSA. Any participating OSPF router, including a router on a broadcast segment that is not the designated router, creates and refreshes its own Router LSA. The OSPF specification defines Router-LSAs as being generated by each router and having area flooding scope. This directly matches both key parts of the question: origination by any router in the area and containment of flooding within that area. See the Router-LSA definition and flooding-scope details in RFC 2328, section 12.4.1 and the LSA scope table in RFC 2328, Appendix A.4.2.

QUESTION NO: 36

Which of the following LSA types stay within an OSPF area, and are not flooded outside of the area? (Choose two.)

A. Router LSA

B. Network LSA

C. Summary network LSA

D. Summary router LSA

E. Type 5 (AS external) LSAs

ANSWER: A B

Explanation:

Router LSA and Network LSA are the OSPF LSA types with area-local flooding scope. A Router LSA (Type 1) is originated by every router for each area to which it is attached. It describes that router's interfaces, links, neighbors, and associated costs as they apply within that particular area. Because this information represents the area's internal topology, it is flooded only to routers in the same area.

A Network LSA (Type 2) is originated by the designated router on a broadcast or non-broadcast multi-access network when that network has at least one fully adjacent router. It identifies the attached routers and represents the shared transit network in the area's link-state database. Like Router LSAs, Network LSAs are used by routers to construct the shortest-path tree for the local area and are not propagated across an area boundary.

This area scoping is a core part of OSPF hierarchy: each area maintains its own detailed topology database, while inter-area reachability is advertised separately by area border routers. RFC 2328 defines Type 1 and Type 2 LSAs as having area flooding scope and describes their origination and contents. See RFC 2328, Router-LSAs and RFC 2328, Network-LSAs.

QUESTION NO: 37

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

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:

OSPF neighbors must agree on key Hello-packet parameters before they can form an adjacency, and they must successfully synchronize databases before reaching Full. **MTU mismatch** can stop the database-description exchange: the receiving router compares the interface MTU advertised in Database Description packets and, when the values are incompatible, the neighbors can remain stuck during ExStart or Exchange rather than completing synchronization. **Incorrect subnet mask** prevents adjacency formation on network types whose Hello processing requires matching network masks, such as broadcast and non-broadcast multi-access networks. The mask is a Hello parameter used to ensure that neighbors share the same IP network. **Area ID not the same** also prevents an adjacency because an OSPF interface and its prospective neighbor must belong to the same area; received Hellos identifying another area are not accepted for neighbor establishment. These checks ensure that routers synchronize link-state databases only with valid peers that share the same OSPF network and area context. The OSPF Hello validation and Database Description MTU behavior are specified in RFC 2328; Nokia's OSPF configuration guidance is available in the Nokia SR OS OSPF documentation.

QUESTION NO: 38

Click on the exhibit.

```
Use this one for 5.2.a
-------------------------------------------------------------------------
???  LSA for Area 0.0.0.0
-------------------------------------------------------------------------
Area Id           : 0.0.0.0              Adv Router Id    : 10.10.10.4
Link State Id     : 0.0.0.0 (0)
LSA Type          : ???
Sequence No       : 0x8000002e           Checksum         : 0x9744
Age               : 40                   Length           : 56
Options           : --R--EV6
Flags             :                      Link Count       : 2
Link Type (1)     : Transit Network      DR Rtr ID (1)    : 10.10.10.4
I/F Index (1)     : 2                    DR I/F Index (1) : 2
Metric (1)        : 100
Link Type (2)     : P2P Link             Nbr Rtr ID (2)   : 10.10.10.3
I/F Index (2)     : 3                    Nbr I/F Index (2): 3
Metric (2)        : 50
```

What type of OSPFv3 LSA is shown?

A. Router LSA

B. Network LSA

C. Intra-Area Prefix LSA

D. Link LSA

ANSWER: A

Explanation:

Router LSA is correct. In OSPFv3, a Router-LSA is an area-scoped LSA originated by every router for each area to which it belongs. Its purpose is to describe the router's links and associated topology information within that area, allowing routers to construct the shortest-path-first topology database. The LSA header identifies its type as the Router-LSA function code, and the body contains a router-bit field followed by a count of described links and the individual link descriptions. This is the OSPFv3 equivalent of the fundamental per-router topology advertisement used for intra-area route calculation. OSPFv3 carries IPv6 addressing information separately from much of the topology information, but Router-LSAs remain central to advertising the area topology. RFC 5340 defines the Router-LSA format and specifies that each router originates one or more Router-LSAs for every attached area. See the IETF's OSPFv3 Router-LSA specification and the associated Router-LSA packet format.

QUESTION NO: 39

In an OSPF Hello packet, which of the following fields must match for all neighbor routers on the segment? Choose three answers.

A. Area ID

B. Hello and Dead intervals

C. Stub flag

D. Checksum values

E. The list of neighbors

ANSWER: A B C

Explanation:

OSPF routers use Hello packets to discover neighbors and verify that the routers attached to a common segment have compatible OSPF operational parameters before forming an adjacency. **Area ID** must match because an OSPF interface belongs to one specific area, and neighbors exchange topology information only within that shared area context. **Hello and Dead intervals** must also match. The Hello interval controls how often Hellos are sent, while the Dead interval specifies how long a router waits without receiving Hellos before declaring a neighbor down; matching timers ensures both routers use the same neighbor-liveness expectations. The **Stub flag** is carried in the OSPF Hello packet's Options field, specifically through the external-routing-capability indication used for stub-area compatibility. Routers on a segment must agree on this area capability so that they operate consistently in the same area type. RFC 2328 specifies that routers will not become neighbors unless key Hello parameters, including Area ID, HelloInterval, RouterDeadInterval, and Options compatibility, agree. See RFC 2328, Hello Packet Processing and RFC 2328, Hello Packet Format.

QUESTION NO: 40

Which of the following statements regarding distance vectors protocols are true? (Choose two.)

A. RIPv1 and RIPv2 are distance vector protocols; BGP is a path-vector protocol.

B. OSPF and IS-IS are distance vector protocols.

C. Routing tables are exchanged between neighbors; however, no routing table is transmitted beyond the immediate neighbor.

D. Distance vector protocols use the Dijkstra SPF algorithm.

E. Routers that participate in distance vector routing protocols maintain full knowledge of distant routers and how they interconnect

ANSWER: A C

Explanation:

RIPv1 and RIPv2 are classic distance-vector protocols: routers advertise reachability information to directly connected neighbors, and each router uses the received vector information to calculate routes through those neighbors. BGP should be described more precisely as a path-vector protocol. It distributes reachability information between BGP peers along with path attributes, most notably the AS_PATH, which records the autonomous systems traversed and supports loop prevention and policy-based route selection. Thus, the corrected statement accurately identifies the RIP protocols as distance-vector protocols while correctly classifying BGP's related but distinct path-vector approach.

The statement that routing information is exchanged only with immediate neighbors describes the fundamental hop-by-hop propagation model of distance-vector routing. A router learns routes from an adjacent router and may subsequently advertise its own selected routes to its own neighbors; it does not send a routing-table update directly through an intermediate neighbor to a nonadjacent router. RIP's specification describes updates being sent to neighboring gateways, and BGP's path-vector behavior is standardized separately in its protocol specification. See RFC 1058: Routing Information Protocol and RFC 4271: A Border Gateway Protocol 4 (BGP-4).

QUESTION NO: 41

What address is used when RIPv2 uses multicast to send its updates?

A. 224.0.0.5

B. 224.0.0.6

C. 224.0.0.9

D. 224.0.0.10

E. RIPv2 does not have support for multicast

ANSWER: C

Explanation:

RIPv2 sends routing updates to the IPv4 link-local multicast address **224.0.0.9**. This multicast group is specifically assigned to RIP version 2 routers, allowing updates to be delivered to RIP-speaking routers on the local network segment without requiring every host to process the packets. Because it is a link-local multicast address, routers do not forward these RIPv2 update packets beyond the originating subnet; RIP neighbors must therefore be directly connected at Layer 3 for normal multicast update exchange. RIPv2 uses UDP port 520 for its routing protocol messages. The use of multicast is one of the operational improvements of RIPv2 over RIPv1, which sends updates using the limited broadcast address 255.255.255.255. In addition to multicast delivery, RIPv2 supports classless routing information, including subnet masks, and can carry authentication information. RFC 2453, the RIPv2 specification, explicitly defines 224.0.0.9 as the multicast address used by RIP-2 routers. Nokia's routing protocol documentation also identifies RIP version 2 updates as UDP-based routing messages using this well-known multicast destination. See RFC 2453: RIP Version 2 and Nokia 7750 SR RIP documentation.

QUESTION NO: 42

Which of the following features are supported by IS-IS? Choose three answers.

A. PDU authentication.

B. The ability to customize the link cost metric.

C. Use of Bellman-Ford routing algorithm.

D. Classless routing.

E. Updates sent as layer 3 multicast.

ANSWER: A B D

Explanation:

IS-IS supports **PDU authentication.** Authentication information can be carried in an IS-IS authentication TLV, allowing routers to validate received routing protocol data units before accepting the routing information. This helps protect the IS-IS control plane and is standardized for both authentication mechanisms and their placement within IS-IS PDUs.

The ability to customize the link cost metric is also a fundamental IS-IS capability. Each IS-IS interface is assigned a metric that contributes to the path cost computed by the link-state shortest-path calculation. IS-IS supports both the original narrow metric style and extended/wide metrics, enabling a network to represent larger metric values and engineer path selection with more granularity.

Classless routing is supported through the encoding and advertisement of IP prefixes with explicit prefix lengths. This permits variable-length subnet masks and route summarization, rather than requiring classful network boundaries. In Integrated IS-IS, IP reachability information is distributed using TLVs that contain the advertised prefix and mask information; later extensions define extended IP reachability for wider metrics and modern deployments.

These capabilities are specified in RFC 1195, which defines Integrated IS-IS for IP, and in RFC 5305, which specifies IS-IS extensions for traffic engineering and extended IP reachability.

QUESTION NO: 43

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

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:

Support hierarchy through the use of areas. is a similarity because both protocols scale a link-state domain by dividing it into hierarchical portions. OSPF organizes routers and links into areas, with Area 0 serving as the backbone and inter-area routing performed through area border routers. IS-IS similarly supports a two-level hierarchy: Level 1 routing provides intra-area reachability, while Level 2 routing provides connectivity between areas. This hierarchical design limits the scope of detailed topology information and supports larger routed networks.

Support for CIDR. is also shared by both protocols. Each protocol can advertise IP prefixes with a prefix length, rather than being restricted to classful network boundaries. OSPF version 2 carries subnet masks in its IP routing information, allowing variable-length subnet masks and route aggregation. Integrated IS-IS carries IP reachability information using prefix and mask data, likewise enabling classless addressing and summarization. These capabilities allow either protocol to operate effectively in networks designed with VLSM and CIDR-based address plans. OSPF area operation and IP route information are specified in RFC 2328; Integrated IS-IS IP routing extensions are specified in RFC 1195.