DUMPSARENA

Nokia Border Gateway Protocol

Alcatel-Lucent 4A0-102

Version Demo

Total Demo Questions: 27

Total Premium Questions: 274

Buy Premium PDF

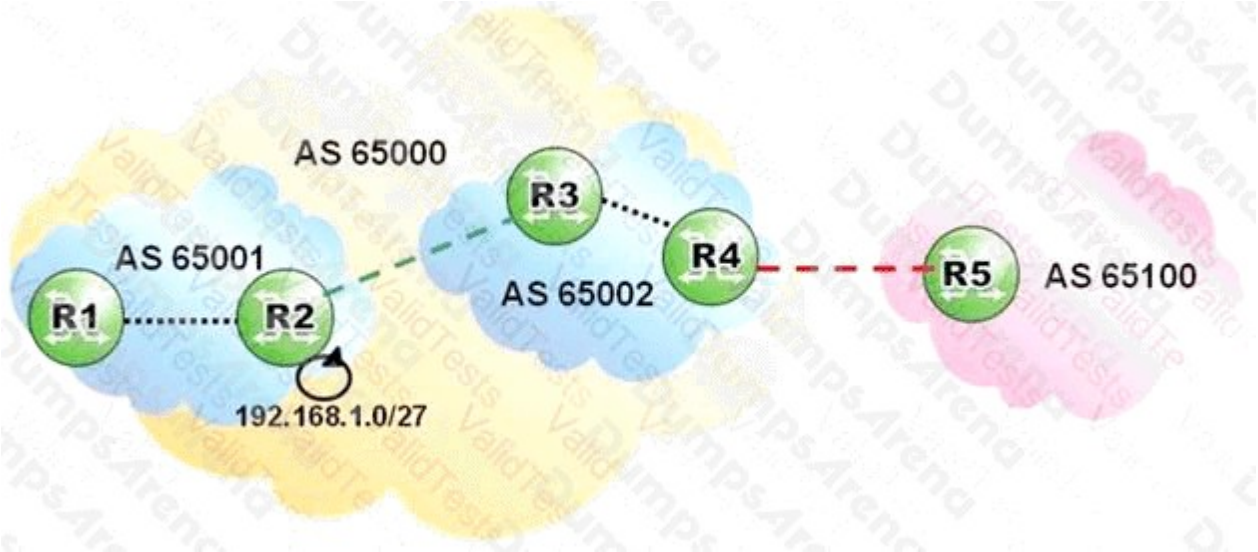
https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

Click the exhibit.



Router R2 is advertising prefix192.168.1.0/27 with community "no-advertise" towards AS 65002.

Which routers receive an update for that prefix?

A. Router R3 only.

B. Routers R3 and R4 only.

C. Routers R3, R4 and R5 only.

D. Router R2 does not advertise an update for the prefix.

ANSWER: A

Explanation:

Router R3 only. The well-known BGP community `NO_ADVERTISE` has the value `0xFFFFFF02` and controls onward propagation of a route. A BGP speaker may send a route carrying this community to its directly configured peer, so R2 can advertise 192.168.1.0/27 to R3, the neighboring router in AS 65002. R3 therefore receives the UPDATE and installs or evaluates the route according to its normal BGP policy and path-selection process. However, once R3 receives the route with `NO_ADVERTISE`, it must not advertise that route to any BGP peer—whether the peer is internal to AS 65002 or external to it. Consequently, the prefix does not proceed beyond R3 to the other routers shown in the topology. This behavior makes `NO_ADVERTISE` appropriate when a route must be made visible only to the immediate BGP neighbor while being explicitly prevented from further BGP distribution. The definition and mandatory propagation restriction for this well-known community are specified in RFC 1997, BGP Communities Attribute; the current BGP communities specification is also available in RFC 8092.

QUESTION NO: 2

An Alcatel-Lucent 7750 SR receives two updates for the same route with local preferences 200 and 300, respectively. What local preference is sent to eBGP peers?

A. 100

B. 200

C. 300

D. None

ANSWER: D

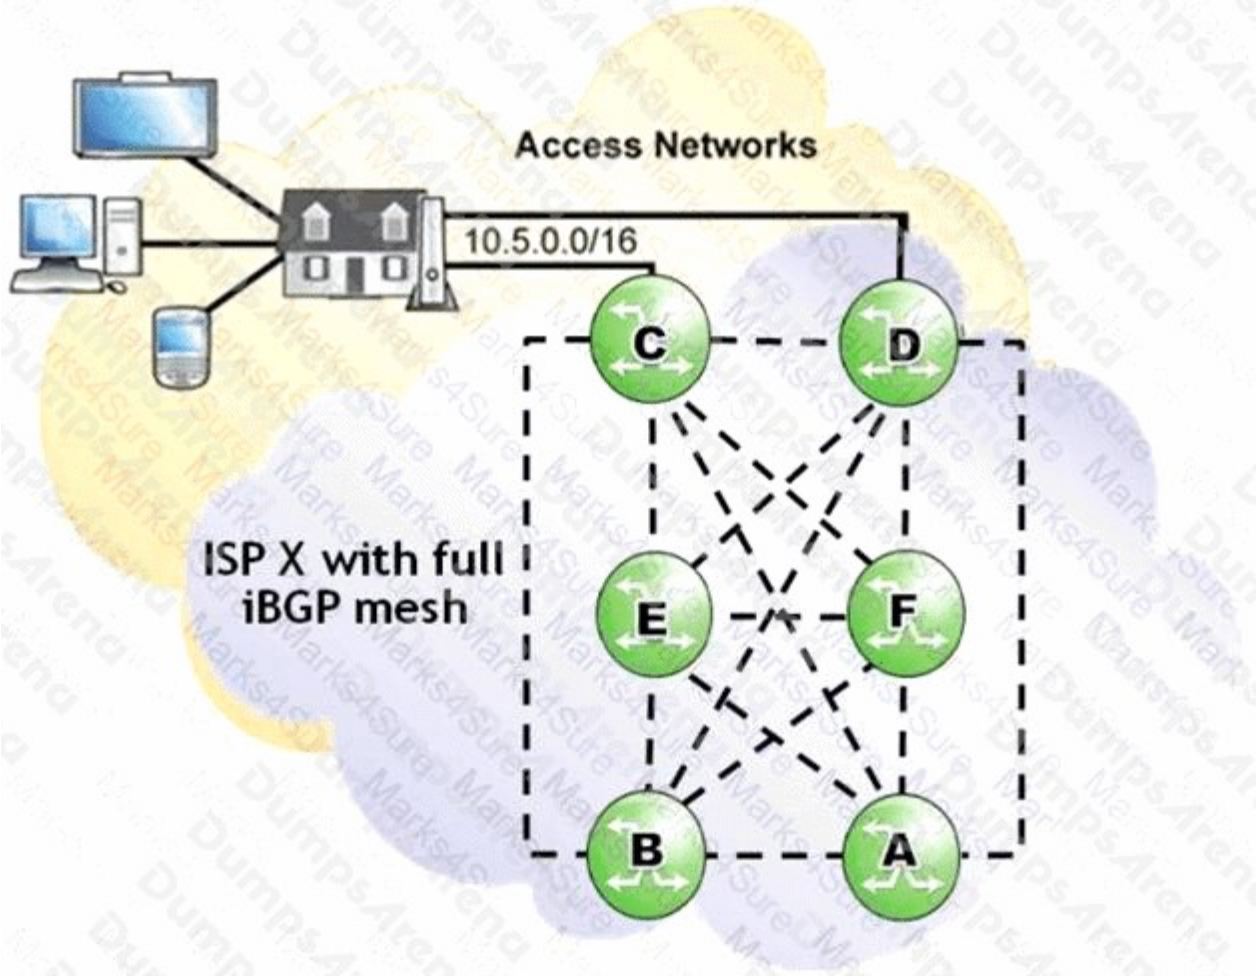
Explanation:

None is correct because LOCAL_PREF is an intra-autonomous-system BGP attribute and is not included when advertising routes to eBGP peers. On the 7750 SR, the locally selected path for this route will use the higher local-preference value, 300, as part of the BGP best-path decision process. That value influences which route the router prefers within its own AS and can be propagated to iBGP speakers so they apply a consistent outbound-routing policy. It does not, however, cross an AS boundary in an eBGP UPDATE.

Consequently, even though the route selected locally has a local preference of 300, the eBGP peer receives the route without any LOCAL_PREF attribute. The receiving external AS applies its own routing policy and its own local-preference value if appropriate. This behavior prevents one AS from imposing its internal exit preference on another AS. RFC 4271 explicitly states that a BGP speaker must not include the LOCAL_PREF attribute in UPDATE messages sent to external peers; see RFC 4271, section 5.1.5. The broader BGP path-attribute rules are also defined in RFC 4271, section 5.

QUESTION NO: 3

Click the exhibit.



Considering that both routers C and D are advertising the eBGP learned prefix 10.5.0.0/16 into

ISP X, which of the following best describes the route advertisement within ISP X?

A. Router C sends the update to routers E and F only.

B. Router C sends the update to routers D, E and F.

C. Router C sends the update to routers A, B, E and F.

D. Router C sends the update to routers A, B, D, E and F.

ANSWER: D

Explanation:

Router C sends the update to routers A, B, D, E and F. The prefix 10.5.0.0/16 was learned by router C through eBGP, so it is eligible for advertisement to router C's iBGP peers within ISP X. iBGP is used to distribute externally learned reachability throughout a single autonomous system while preserving the external AS-path information associated with the route. In the depicted full-mesh iBGP arrangement, router C has iBGP sessions with routers A, B, D, E, and F, and therefore advertises its selected eBGP-learned route to each of those peers. Router D also receives router C's advertisement even though it has independently learned and advertised reachability for the same destination prefix. Each receiving router can then apply the BGP path-selection process to choose its own best route for forwarding and subsequent eligible advertisement. The advertisement action from router C is determined by its iBGP peer relationships and the fact that the route was learned via eBGP, rather than by whether another border router has a competing path to the same prefix. This behavior follows the BGP UPDATE propagation rules defined in RFC 4271; Nokia's BGP documentation also describes iBGP's role in distributing external routes inside an AS: Nokia SR OS BGP documentation.

QUESTION NO: 4

What kind of attribute is the aggregator?

A. Well-known mandatory

B. Well-known discretionary

C. Optional transitive

D. Optionalnon-transitive

ANSWER: C

Explanation:

Optional transitive is correct. The AGGREGATOR path attribute is used when a BGP speaker creates an aggregated route. It identifies the router that performed the aggregation by carrying that router's autonomous system number and BGP identifier. This information helps preserve visibility into how the advertised aggregate was formed, particularly when route aggregation can obscure individual component routes.

RFC 4271 defines AGGREGATOR as an optional transitive attribute. "Optional" means that a BGP implementation is not required to support or include it on every route. "Transitive" means that, if a BGP speaker that does not recognize the attribute receives it, the speaker must still propagate it to other peers, while marking it as partially understood. That propagation behavior ensures the aggregation-origin information can remain available across autonomous-system boundaries even where intermediate speakers do not process its semantics. The attribute is specifically described in the route-aggregation section of the BGP specification. See RFC 4271, Section 5.1.7 and the aggregation procedures in RFC 4271, Section 9.2.2.2.

QUESTION NO: 5

What kind of BGP attribute is the Multi_Exit_Disc?

A. Optional transitive attribute.

B. Optional non-transitive attribute.

C. Well-known mandatory attribute.

D. Well-known discretionary attribute.

ANSWER: B

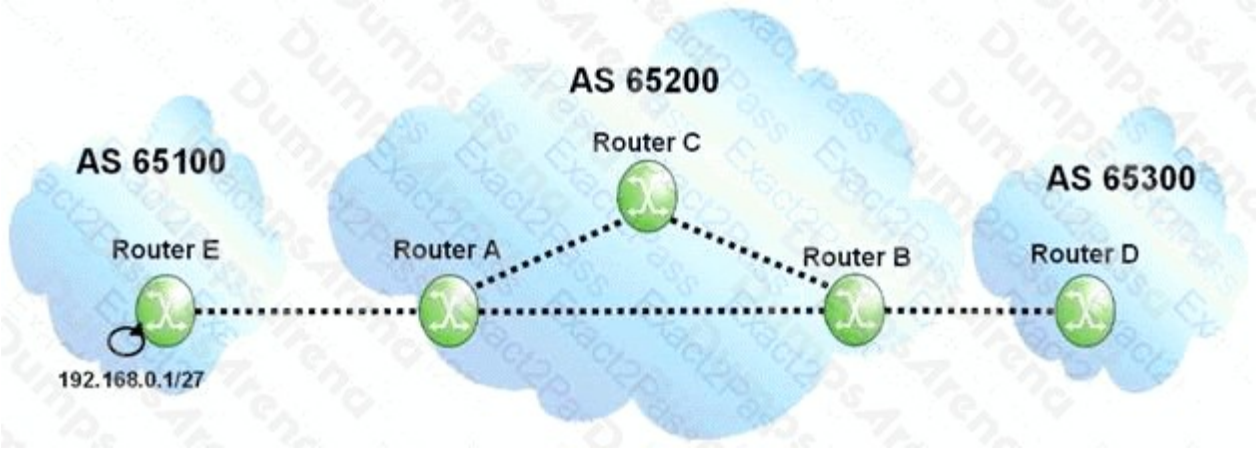
Explanation:

The Multi_Exit_Disc (MED) is an optional non-transitive BGP path attribute. MED is used to communicate to an external neighboring autonomous system the preferred entry point when multiple interconnection links are available. In general, a lower MED value is preferred, enabling an AS to suggest which of its links should be used by the neighboring AS for traffic entering that AS.

Its optional classification means that BGP implementations are not required to support or recognize it in the same way as well-known attributes. Its non-transitive classification means that, if a BGP speaker does not recognize the attribute, it must not pass it on to other BGP peers. This scope keeps the metric meaningful primarily between directly connected autonomous systems rather than allowing it to propagate broadly through the Internet. RFC 4271 defines the MULTI_EXIT_DISC attribute with the optional flag set and the transitive flag clear. See RFC 4271, Section 5.1.4 and the BGP-4 specification.

QUESTION NO: 6

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. The route received by router B from router A is an internally learned BGP route within the autonomous system shown in the exhibit. Standard iBGP propagation rules prevent a BGP speaker from advertising a route learned from one iBGP peer to another iBGP peer unless route reflection or confederation behavior has specifically been configured. This rule avoids routing-information loops in a full-mesh iBGP design. Router B can, however, advertise the selected route to its external BGP neighbor, router D. When advertising externally, BGP also applies its normal outbound processing, including adding the local autonomous-system number to the AS_PATH, so the advertisement can be safely propagated into the neighboring autonomous system. Therefore, once router B accepts the update, router D is the eligible recipient identified by the topology and BGP advertisement rules. This behavior follows the base BGP specification's rules for route advertisement and AS-path loop prevention; route-reflector exceptions are defined separately and are not implied by the diagram. See RFC 4271, A Border Gateway Protocol 4 (BGP-4) and RFC 4456, BGP Route Reflection.

QUESTION NO: 7

Assuming that "client1" and "client2" are directly-connected networks, what is the result of executing the following BGP export policy?

```
entry 10
    from
        protocol direct
    exit
    action next-entry
        community add "East"
    exit
exit
entry 20
    from
        prefix-list "client1"
    exit
    action accept
        community add "North"
    exit
exit
entry 30
    from
        prefix-list "client2"
      exit
      action accept
          community add "South"
      exit
  exit
  default-action reject
```

A. "client2" routes will be only tagged with communities 'East" and "South".

B. "client2" routes will be tagged with community "East".

C. "client2" routes will only be tagged with community "South".

D. "client2" routes will be rejected.

ANSWER: A

Explanation:

"client2" routes will be only tagged with communities 'East" and "South". A directly connected client2 prefix satisfies the policy's matching conditions and is exported after both applicable community operations have been performed. The resulting BGP advertisement therefore carries the two configured community values, East and South. Communities are optional transitive BGP path attributes used to mark routes so that subsequent routing policy—locally or at another BGP speaker—can classify and handle those routes according to the administrative intent. Adding these communities does not change the fact that the route is a directly connected prefix; it annotates the route as it is exported. This behavior follows the policy-processing model in which matching policy entries apply their configured route modifications to the route being evaluated before the export decision is completed. BGP communities and their propagation as path attributes are defined in RFC 1997, while the base BGP route advertisement and path-attribute framework is specified in RFC 4271.

QUESTION NO: 8

What isthe result of configuring the following policy statement as a BGP import policy on theNokia7750 SR?

A. All BGP routes are accepted.

B. All routes matching prefix lists "List-1" and "List-2" are rejected.

C. BGP routes matching prefix lists "List-1" or "List-2" are rejected.

D. BGP routes matching prefix lists "List-1" and "List-2" are rejected.

ANSWER: C

Explanation:

BGP routes matching prefix lists "List-1" or "List-2" are rejected. In an SR OS routing policy, a policy entry can use multiple prefix-list names as values for the prefix-list match criterion. The criterion is satisfied when the route's prefix matches any one of the listed prefix lists; the values are therefore evaluated as alternatives for that criterion. When the matching policy entry specifies a reject action and the policy is applied in the BGP import direction, any received BGP route that matches either named prefix list is rejected before it is installed or considered for further BGP route processing. This is useful when a single import filter must exclude several independently maintained groups of prefixes, such as private address space, customer-specific aggregates, or otherwise prohibited networks. The import policy is evaluated for inbound BGP updates, so the result controls which advertised routes are accepted from the peer. Nokia's SR OS routing-policy documentation describes policy entries, match criteria, actions, and their use by routing protocols, including BGP. See the Nokia SR OS routing-policy documentation and the Nokia SR OS BGP documentation.

QUESTION NO: 9

Assuming that "client1" and "client2" are directly-connected networks, what is the result ofexecuting the following BGP export policy?

A. Only entry 10 is executed for "client1" and "client2".

B. Entries 10 and 20 are executed for "client1", and entries 10 and 30 are executed for "client2".

C. Entries 10 and 20 are executed for "client1", andentries 10, 20 and 30 are executed for"client2".

D. Entry 10 is executed for directly-connected routes, entry 20 for "client1" and entry 30 for "client2".

E. The result cannot be determined because the BGP export policy is not provided.

ANSWER: E

Explanation:

The correct conclusion is that the result cannot be determined because the BGP export-policy configuration referred to by the question is not present. Whether a directly connected route for *client1* or *client2* reaches a particular policy entry depends on the policy's ordered entries, match criteria, and terminating or non-terminating actions. In SR OS routing policy, entries are evaluated in sequence, but an action can accept or reject a route immediately, or it can explicitly continue evaluation to a subsequent entry. Consequently, knowing only that the two networks are directly connected does not establish which entries are evaluated. The omitted policy is essential to determine whether an entry matches on route protocol, a prefix list containing either client network, or another route attribute, and whether evaluation stops after that match. A reliable answer therefore requires the complete policy statement, including every entry and its action. Nokia's routing-policy documentation describes policy evaluation and entry actions, while the BGP documentation explains how export policies control routes advertised by BGP: Nokia SR OS Routing Policies and Nokia SR OS BGP documentation.

QUESTION NO: 10

Which of the following regarding BGP confederations is FALSE?

A. A confederation is a collection of Autonomous Systems advertised as a single AS number to BGPspeakers outside the confederation.

B. Each member AS must either maintain a full mesh of iBGP sessions, or use route reflection.

C. Confederations can be used to subdivide ASs that have a large number of BGP speakers into smaller domains.

D. A full mesh of intra-confederation eBGP sessions is required between the member ASs.

ANSWER: D

Explanation:

A full mesh of intra-confederation eBGP sessions is required between the member ASs. is the false statement. A BGP confederation divides one public autonomous system into multiple member autonomous systems for internal scaling purposes. BGP sessions between member ASs are treated as intra-confederation eBGP sessions, but the confederation design does not impose a requirement that every member AS peer directly with every other member AS. Member AS connectivity can instead follow the network topology, with selected intra-confederation eBGP peerings carrying reachable routes across the confederation. This permits a scalable design while preserving the external presentation of the confederation as one autonomous system. To external BGP neighbors, the confederation identifier is the relevant public AS number; member-AS numbers are internal implementation details and are removed from the AS path when routes leave the confederation. The behavior and AS-path handling for BGP confederations are specified by the IETF in RFC 5065. Nokia SR OS BGP documentation also describes confederations as an approach for partitioning a large AS into smaller member ASs: Nokia SR OS BGP documentation.

QUESTION NO: 11

Which regular expression will match the AS Path of a route originating in remote AS 65100?

A. ".+ 65100"

B. "65100.+"

C. ".* 65100.*"

D. "65100"

ANSWER: A

Explanation:

`".+ 65100"` is the appropriate AS-path regular expression for identifying a route whose originating AS is remote AS 65100. AS_PATH is ordered from the most recently added, or adjacent, AS at the left toward the origin AS at the right. As an advertisement crosses each autonomous system, that system prepends its own AS number to the path; consequently, the originating AS remains the last AS in a normal AS_SEQUENCE path.

In the expression, `.+` represents one or more preceding characters in the received AS path, including the preceding AS-number portion, and the literal space before `65100` preserves the AS-path token boundary. Ending the expression with `65100` therefore selects paths whose final, origin-side AS is 65100, while requiring that the route be received through at least one intervening AS, as indicated by “remote.” This is useful in SR OS routing policy when matching BGP route attributes for import or export decisions. The AS_PATH ordering and AS-prepending behavior are defined in RFC 4271, A Border Gateway Protocol 4 (BGP-4); Nokia SR OS BGP policy and AS-path matching behavior is documented in the Nokia SR OS BGP documentation.

QUESTION NO: 12

What is the effect of configuring a maximum-prefix limit for a BGP neighbor?

A. It limits the number of prefixes accepted from the neighbor.

B. It limits the AS Path length accepted from the neighbor.

C. It prevents the neighbor from advertising default routes.

D. It limits the number of TCP connections used by BGP.

ANSWER: A

Explanation:

It limits the number of prefixes accepted from the neighbor. This is correct. A maximum-prefix limit protects the local BGP speaker from an unexpectedly large number of route advertisements received from a peer. The configured threshold is checked against prefixes received for the applicable address family.

If the neighbor exceeds the configured limit, the router takes the action defined by the implementation and configuration, which can include terminating the BGP session. This protection helps contain the impact of route leaks, peer misconfiguration, or other events that could consume excessive control-plane and routing-table resources. A warning threshold can also be used to alert an operator before the maximum is reached. See RFC 4271 and Nokia SR OS documentation.

QUESTION NO: 13

Which of the following prefixlists include prefixes 192.168.128.0/17 and 192.168.64.0/18?

A. Prefix list 192.168.0.0/16 through 18.

B. Prefix 192.168.0.0/17 longer.

C. Prefix list 192.168.0.0/17 through 18.

D. Prefix 192.160.0.0/13 longer.

ANSWER: A

Explanation:

Prefix list 192.168.0.0/16 through 18. is correct because its base aggregate, 192.168.0.0/16, covers the complete address interval from 192.168.0.0 through 192.168.255.255. Both stated networks fall inside that aggregate: 192.168.128.0/17 represents the upper half of the /16 address space, while 192.168.64.0/18 is contained in its lower half. The “through 18” portion specifies the permitted prefix-length range, inclusive, from /16 to /18. Consequently, a /17 and a /18 are both eligible as long as their network addresses are within 192.168.0.0/16. This reflects the standard route-filter or prefix-list matching model: matching requires both containment within the configured prefix and a prefix length within the configured minimum-to-maximum range. CIDR prefix boundaries and containment behavior are defined in RFC 4632. Nokia’s routing documentation is available through the Nokia Documentation Center, including SR OS policy and routing-filter material that uses this prefix-and-mask-range approach.

QUESTION NO: 14

Click the exhibit.



Which of the following is a valid BGP configuration for router R2?

A.
```
R2>config>router>bgp# info
----------------------------------------
        group "Conf-eBGP"
            peer-as 65541
            neighbor 10.16.10.1
            exit
        exit
        group "Member-AS"
            neighbor 10.16.10.6
                peer-as 65542
            exit
        exit
        no shutdown
```

B.
```
R2>config>router>bgp# info
----------------------------------------
        group "Conf-eBGP"
            peer-as 65541
            neighbor 10.16.10.1
            exit
        exit
        group "Member-AS"
            neighbor 10.16.10.6
                peer-as 65540
            exit
        exit
        no shutdown
```

C.
```
R2>config>router>bgp# info
----------------------------------------
        group "Conf-eBGP"
            peer-as 65540
            neighbor 10.16.10.1
            exit
        exit
        group "Member-AS"
            neighbor 10.16.10.6
                peer-as 65542
            exit
        exit
        no shutdown
```

D.
```
R2>config>router>bgp# info
----------------------------------------
        group "Conf-BGP"
            peer-as 65540
            neighbor 10.16.10.1
            exit
            neighbor 10.16.10.6
            exit
        exit
        no shutdown
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: A

Explanation:

The configuration represented by *Option A* is the valid choice for R2. A usable SR OS BGP configuration must be created under the router BGP context, establish a neighbor relationship using the peer address depicted in the topology, and use the autonomous-system information appropriate to that peer relationship. The BGP session is associated with a group and neighbor on SR OS, with the remote autonomous system configured through the neighbor or group peer-AS setting. This ensures that R2 can correctly identify the relationship and perform the BGP OPEN-message AS validation required before the session can establish.

When configuring BGP on an Alcatel-Lucent/Nokia SR OS router, the local BGP AS identifies R2's own routing domain, while the peer AS identifies the remote router's domain. The neighbor address must also be reachable through the routing table, typically using the directly connected interface address or a loopback address together with the required reachability. The valid configuration therefore combines the proper BGP hierarchy, local AS, neighbor definition, and peer-AS relationship required by the exhibit. Nokia's SR OS BGP configuration and operational guidance is available in the SR OS BGP documentation and the Nokia BGP overview.

QUESTION NO: 15

Which of the following statements regarding BGP Add-Paths is FALSE?

A. Add-Paths allows non-best paths to be advertised to a peer.

B. Add-Paths capabilities are exchanged between peers after the BGP session is established.

C. A BGP speaker must advertise the capability to receive multiple paths from its peer.

D. A BGP speaker must advertise the capability to send multiple paths to its peer.

ANSWER: B

Explanation:

"Add-Paths capabilities are exchanged between peers after the BGP session is established." is false because Add-Paths support is negotiated as part of the BGP capability exchange in the OPEN messages, which occurs while the BGP session is being established. Each peer includes its supported capabilities in its OPEN message; only after both OPEN messages have been processed and the peers reach the Established state can UPDATE messages carry routes using the negotiated Add-Paths behavior. The Add-Paths Capability identifies whether a speaker can send multiple paths, receive multiple paths, or both, for a given address family. This advance negotiation is essential because the peers need an agreed interpretation of the Path Identifier field before any multipath route advertisements are exchanged. RFC 7911 defines Add-Paths and specifies that the capability is carried in the BGP OPEN message, not exchanged after establishment. This behavior is also consistent with normal BGP capability negotiation defined for OPEN-message processing. See RFC 7911: Advertisement of Multiple Paths in BGP and RFC 5492: Capabilities Advertisement with BGP-4.

QUESTION NO: 16

Assuming that "client1" and "client2" are directly-connected networks, what is the result of executing the following BGP export policy?

A. Only entry 10 is executed for "client1" and "client2".

B. Entries 10 and 20 are executed for "client1", and entries 10 and 30 are executed for "client2".

C. Entries 10 and 20 are executed for "client1", and entries 10, 20 and 30 are executed for"client2".

D. Entry 10 is executed for directly-connected routes, entry 20 for "client1" and entry 30 for "client2".

ANSWER: A

Explanation:

Only entry 10 is executed for "client1" and "client2". Both prefixes are learned as directly connected routes, so they satisfy the route-source match evaluated in entry 10. In SR OS routing policy processing, an entry whose match criteria are satisfied performs its configured action. Where that action is terminal, such as an accept action, policy evaluation stops immediately for that route; subsequent entries are not evaluated. Consequently, each directly connected client route reaches entry 10, matches it, and is accepted for BGP export at that point. The more specific later entries therefore are not reached during processing of either route. This behavior is important when designing BGP export policies: broad route-source matches with a terminating action take precedence over later policy terms, so policy entries must be ordered from the intended highest-precedence match to lower-precedence matches. Nokia's SR OS routing-policy documentation describes ordered policy-entry evaluation and the effect of policy actions, while its BGP documentation describes application of export policies to advertised routes. See Nokia SR OS policy statements documentation and Nokia SR OS BGP documentation.

QUESTION NO: 17

AnNokia7750 SR has the BGP configuration shown below. Assuming the router has established BGP sessions to routers R1 and R2, which of the following is TRUE?

A. Both neighbors R1and R2should be part of AS 65540.

B. Neighbor R1should be part of AS 65550, while neighbor R2 should be part of AS 65540.

C. Neighbor R1should be part of AS 65550, while neighbor R2 should be part of AS 65560.

D. Both neighbors R1and R2 should be part of AS 65560.

ANSWER: C

Explanation:

Neighbor R1 should be part of AS 65550, while neighbor R2 should be part of AS 65560. In Nokia SR OS BGP configuration, the autonomous system expected for a BGP peer is specified by the neighbor or peer-group `peer-as` setting. A successfully established BGP session means that the remote router has presented the ASN expected by the local router during BGP OPEN-message processing. Therefore, where the configuration associates R1 with peer AS 65550 and R2 with peer AS 65560, those are the ASNs to which the respective neighbors belong. These are valid four-octet ASN values; BGP supports four-octet AS numbers so that values above 65535 can be carried and checked as part of normal session establishment and AS-path handling. The peer ASN is a fundamental part of defining whether a session is internal BGP or external BGP and determines the expected remote speaker identity for the configured neighbor. Nokia's SR OS BGP documentation describes BGP neighbor and peer-group configuration, including peer-AS behavior, at Nokia SR OS BGP documentation. Four-octet ASN support is standardized in RFC 6793.

QUESTION NO: 18

An Alcatel-Lucent 7750 SR has the BGP configuration shown below. Assuming the router has established BGP sessions to routers R1 and R2, which of the following is TRUE?

```
group "BGP Peers"
    peer-as 65540
    neighbor 10.16.10.1
        description "R1"
        peer-as 65550
    exit
    neighbor 10.16.10.2
        shutdown
        description "R2"
        peer-as 65560
    exit
exit
```

A. Both neighbors R1and R2 should be part of AS 65540.

B. Neighbor R1should be part of AS 65550, while neighbor R2 should be part of AS 65540.

C. Neighbor R1should be part of AS 65550, while neighbor R2 should be part of AS 65560.

D. Both neighbors R1and R2 should be part of AS 65560.

ANSWER: C

Explanation:

Neighbor R1 should be part of AS 65550, while neighbor R2 should be part of AS 65560. In SR OS BGP configuration, a peer AS configured at the BGP group level supplies the default remote AS for neighbors in that group. Therefore, R1 inherits the group peer-AS value of 65550. The configuration for R2 includes its own neighbor-level peer-AS value of 65560. A peer-AS configured directly under a neighbor takes precedence over the peer-AS inherited from the group for that particular BGP session. Consequently, the router can establish an external BGP session to R1 in AS 65550 and a separate external BGP session to R2 in AS 65560, even though both neighbors are associated with the same BGP group. This hierarchy allows common BGP policy and operational parameters to be placed at group scope while permitting an individual neighbor to use a distinct remote autonomous system. Nokia's SR OS BGP documentation describes BGP groups, neighbor configuration, and the peer-AS association used to identify the remote AS: Nokia SR OS BGP documentation. Additional SR OS routing configuration documentation is available through the Nokia SR OS documentation portal.

QUESTION NO: 19

What is the function of the Cluster_List attribute?

A. It is used for loop detection within the same cluster.

B. It is used for loop detection between clusters.

C. It is used for loop detection in a confederation.

D. It is used to identify the route reflectors within an AS.

ANSWER: B

Explanation:

It is used for loop detection between clusters. In a route-reflector deployment, a cluster is identified by a cluster ID, and a route reflector adds its local cluster ID to the `CLUSTER_LIST` when reflecting a route. The attribute therefore records the

sequence of route-reflector clusters through which the route has passed. Before reflecting or accepting a route, a route reflector examines this list; if its own cluster ID is already present, the route has returned to a cluster it previously traversed and is discarded. This prevents reflected BGP updates from circulating through interconnected route-reflector clusters and forming persistent control-plane loops.

The mechanism is defined as part of BGP route reflection. It is particularly important where route reflectors exchange routes with other route reflectors, since ordinary AS-path loop prevention alone does not identify every reflection path within the same autonomous system. The `CLUSTER_LIST` provides the route-reflection-specific path history required for that check. RFC 4456 defines the attribute as a sequence of cluster IDs representing the reflection path and specifies the local-cluster-ID loop check: RFC 4456, BGP Route Reflection. Additional BGP attribute and route-reflector background is available in Nokia's SR OS routing documentation: Nokia SR OS BGP documentation.

QUESTION NO: 20

Which of the following is TRUE?

A. There can only be one community associated with a route.

B. Community is an optional non-transitive attribute.

C. Community is an optional transitive attribute.

D. Community is a well-known discretionary attribute.

ANSWER: C

Explanation:

Community is an optional transitive attribute. BGP communities provide a scalable way to attach routing-policy information to a route. A community is encoded as a set of one or more four-octet community values, allowing a route to carry multiple community tags simultaneously. Routers can use these values in routing policy to match routes and apply actions such as changing preference, controlling route advertisement, or adding and removing community values.

The attribute is optional because a BGP UPDATE does not have to contain it. It is transitive because a BGP speaker that recognizes and accepts the Communities attribute propagates it to eligible BGP peers, subject to configured export policy. This preserves policy-signaling information across autonomous systems when it is intended to travel with the route. The original BGP Communities specification defines the attribute as optional transitive and describes its format and propagation behavior. Nokia SR OS routing policy also supports matching and manipulating BGP community values as route attributes. See RFC 1997: BGP Communities Attribute and Nokia SR OS BGP documentation.

QUESTION NO: 21

Click the exhibit.



Assuming that none of the routers within AS 65200 is configured with next-hop-self", what will the BGP update for the 192.168.0.1/27 prefix contain when it arrives at router B?

A. AS Path of 65200 65100, Next Hop of router A.

B. AS Path of 65200 65100, Next Hop of router E.

C. AS Path of 65100, Next Hop of router E.

D. AS Path of 65100, Next Hop of router A.

ANSWER: C

Explanation:

AS Path of 65100, Next Hop of router E. is correct. Router E advertises the 192.168.0.1/27 prefix from AS 65100 to the AS 65200 edge router. When a route is received through eBGP, the receiving router prepends neither its own AS on receipt nor changes the originating external AS path; therefore, the route has an AS path containing 65100. When that route is subsequently advertised within AS 65200 through iBGP, iBGP preserves the AS_PATH attribute. AS 65200 is not inserted into the path during internal propagation because doing so would incorrectly represent an external AS traversal and could trigger loop-prevention behavior within the same autonomous system. iBGP also preserves the NEXT_HOP attribute by default. Since no router in AS 65200 uses next-hop-self, the next hop learned from the external advertisement remains router E when the update reaches router B. Router B must therefore be able to resolve/reach router E's next-hop address through the IGP or another installed route. This behavior follows the BGP attribute handling rules in RFC 4271, section 5.1.3, and is consistent with Nokia's SR OS documentation library for BGP operation.

QUESTION NO: 22

Click the exhibit.

Assume all router-IDs are properly configured, and IPv6 link-local addresses are used to establish all eBGP sessions. Given the following iBGP configuration on router A, what is the Next Hop forprefix 2001:DB8:A:401::1/128 on router B?

A. The IPv6 link-local address of the router D interface towards router A.

B. The IPv6 link-local address of the router A interface towards router D.

C. The IPv6 global address of router A.

D. The IPv6 globaladdress of router D.

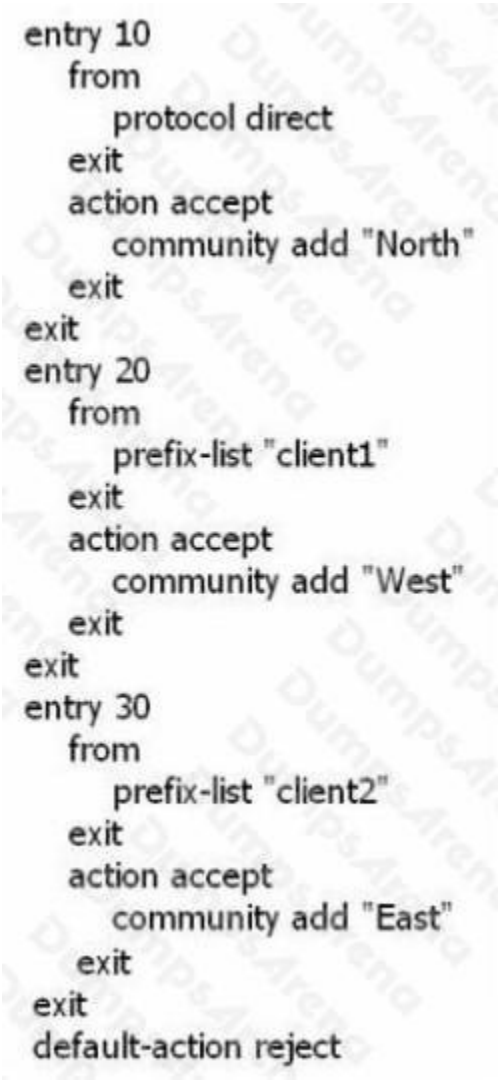
ANSWER: C

Explanation:

The IPv6 global address of router A is the next hop installed and advertised to router B for prefix 2001:DB8:A:401::1/128. Router A learns the route through an eBGP session established using IPv6 link-local addressing, but the iBGP policy/configuration shown causes A to apply next-hop self when it advertises the route internally. In IPv6 BGP, a usable next hop for an internally distributed route must be meaningful and reachable in the receiving router's routing context. A link-local address is scoped to a single interface and link, so it cannot serve as a generally reachable transit next hop for router B when the original external peer is connected to A. By replacing the externally learned next hop with its own globally routable IPv6 address, router A gives B a next hop that B can resolve through the IGP and use to forward traffic toward the external destination. This is the practical purpose of next-hop-self on an iBGP advertisement at the edge of the autonomous system. Nokia SR OS documents BGP next-hop processing and the use of next-hop-self in its BGP routing-protocol documentation. The underlying IPv6 BGP next-hop encoding, including global and link-local next-hop considerations, is specified in RFC 2545.

QUESTION NO: 23

Assuming that "client1" and "client2" are directly-connected networks, what is the result of executing the following BGP export policy?



A. "client1" routes will be tagged with communities "North" and "West".

B. "client1" routes will be tagged with community "West".

C. "client1" routes will be tagged with community "North".

D. "client1" routes will be rejected.

ANSWER: C

Explanation:

"client1" routes will be tagged with community "North". The route for client1 is a directly connected route, but it first matches the policy entry specifically intended for the client1 prefix. That matching entry adds the North community and accepts the route for BGP export. In SR OS routing policy evaluation, an accept action is terminal for that policy evaluation: once the route is accepted, subsequent entries are not evaluated. Consequently, the later condition that can match directly connected routes does not add the West community to client1's route. The exported BGP update therefore carries the North community attribute for the client1 prefix. BGP communities are optional transitive path attributes used to label routes so that later routing policy can identify and process them consistently. See Nokia's SR OS BGP documentation and RFC 1997, BGP Communities Attribute.

QUESTION NO: 24

Which of the following is TRUE if a BGP session state is Connect?

A. BGP is waiting for the TCP connection to be completed.

B. BGP has established a TCP connection, and an Open message has been sent to that neighbor.

C. The BGP session has been interrupted, and a TCP connection is being re-established.

D. BGPhas failed to establish a TCP connection to that neighbor, and ConnectRetry has been reset to zero.

ANSWER: A

Explanation:

“BGP is waiting for the TCP connection to be completed.” is correct because the Connect state is the stage in the BGP finite-state machine in which the router is waiting for the underlying TCP transport connection to complete. A BGP peer relationship depends on TCP, normally using destination port 179, before BGP can exchange its protocol messages. If the TCP connection succeeds while the session is in Connect, BGP initializes the connection resources, sends its Open message, sets the Hold Timer to a large value, and transitions to the OpenSent state. The OpenSent state—not Connect—is therefore the state associated with having sent an Open message. The precise behavior is defined by the BGP finite-state-machine rules in RFC 4271, which specifies that a successful TCP connection-completion event in Connect causes initialization and transmission of the Open message. This transport-first progression is important operationally: a router can be in Connect even though no BGP capability negotiation, authentication result, or route exchange has occurred yet. See the IETF’s BGP Connect-state definition and the related BGP finite-state machine.

QUESTION NO: 25

The configuration of an Alcatel-Lucent 7750 SR is given below. Router Rl has established BGP sessions with routers R2 and R3. Which of the following is TRUE?

```
R1>config>router>bgp# info
----------------------------------------
        group "BGP-Peers"
            peer-as 65540
            neighbor 10.10.10.2
                description "R2"
            exit
            neighbor 10.10.10.3
                description "R3"
                peer-as 65550
            exit
        exit
```

A. Router R1forms an iBGP session with router R2, and an eBGP session with router R3.

B. Router R1forms iBGP sessions with routers R2 and R3.

C. Router R1forms eBGP sessions with routers R2 and R3.

D. The types of BGP sessions with routers R2 and R3 cannot be determined.

ANSWER: D

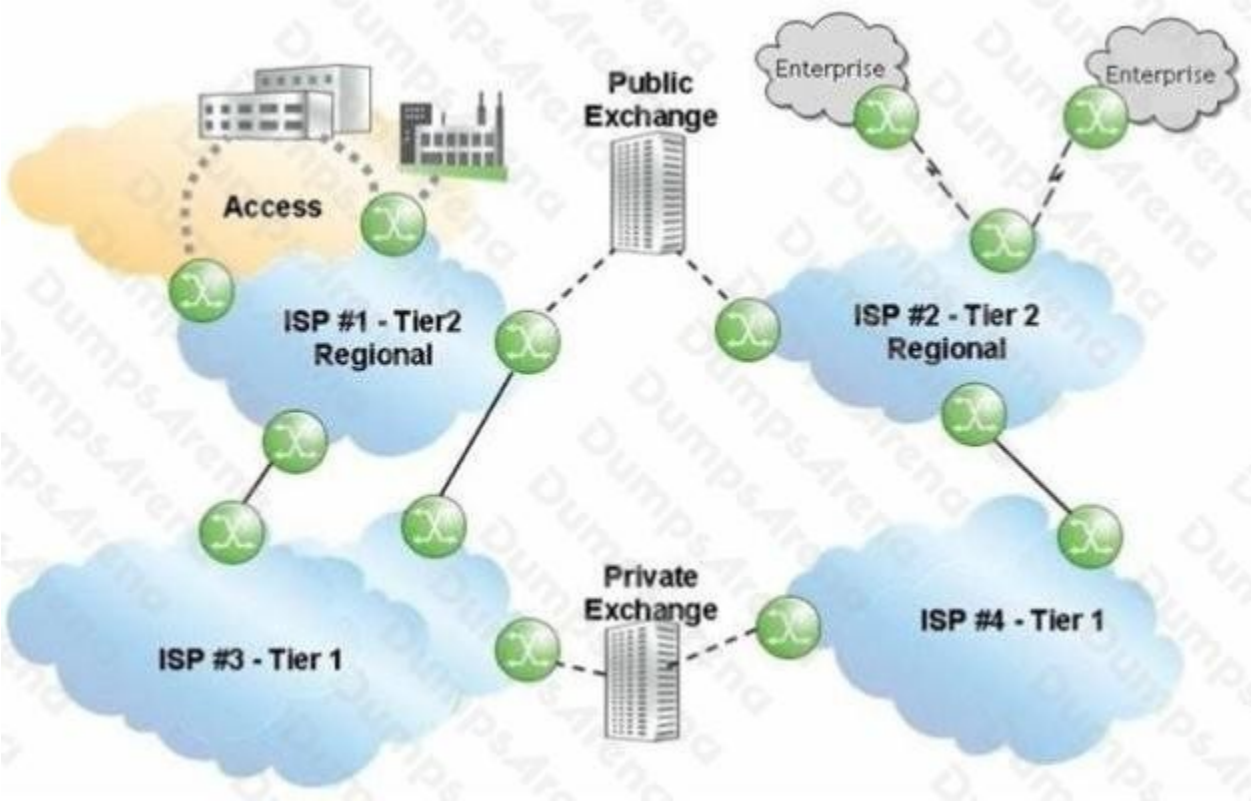
Explanation:

The types of BGP sessions with routers R2 and R3 cannot be determined. A BGP session is classified by comparing the local speaker’s autonomous system number with the autonomous system number configured for, or advertised by, the peer. When the AS numbers are the same, the relationship is internal BGP (iBGP); when they differ, it is external BGP (eBGP).

The fact that BGP sessions are established confirms that the peers can communicate and that their session parameters are compatible, but it does not, by itself, identify whether each neighbor belongs to the same AS as R1. The displayed configuration must explicitly provide the relevant local AS and peer AS information, or an explicit internal/external group type, before either session can be classified. This distinction matters because iBGP and eBGP have different route-propagation and attribute-handling behavior. The underlying BGP protocol defines an internal peer as one in the same autonomous system and an external peer as one in a different autonomous system; see RFC 4271, Section 3. Nokia's SR OS documentation is available from the Nokia SR OS documentation portal.

QUESTION NO: 26

Click the exhibit.



Given that ISP 1 and ISP 2 are Tier 2 providers, and that ISP 3 and ISP 4 are Tier 1 providers, what is the most likely relationship between the ISPs?

A. ISP 1 and ISP 2 have a public transit relationship, while ISP 3 and ISP 4 have a private transit relationship.

B. ISP 1 and ISP 2 have a private transit relationship, while ISP 3 and ISP 4 have a public transit relationship.

C. ISP 1 and ISP 2 have a private peering relationship, while ISP 3 and ISP 4 have a public peering relationship.

D. ISP 1 and ISP 2 have a public peering relationship, while ISP 3 and ISP 4 have a private peering relationship.

ANSWER: D

Explanation:

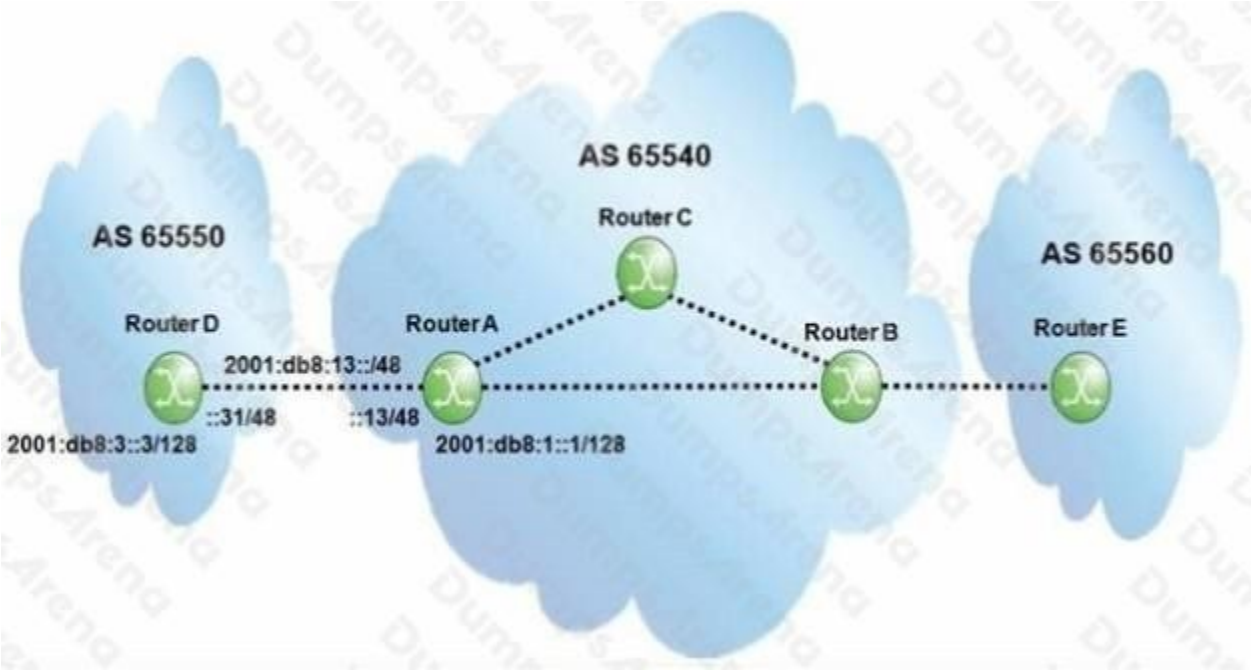
ISP 1 and ISP 2 have a public peering relationship, while ISP 3 and ISP 4 have a private peering relationship. The exhibit distinguishes the two interconnection models: ISP 1 and ISP 2 connect through a shared public exchange infrastructure, which is characteristic of public peering at an Internet exchange point. Public peering enables multiple independently operated networks to establish BGP sessions using the common exchange fabric rather than requiring a dedicated physical circuit between every pair of participants.

The Tier 1 providers are shown with a direct, dedicated interconnection. This is private peering: two networks establish a bilateral physical connection, commonly used where traffic volume, performance requirements, or routing-policy needs justify dedicated capacity. Tier designation provides useful context because Tier 1 networks commonly maintain settlement-free

bilateral interconnections with other major networks, while Tier 2 networks may use exchange points to peer efficiently with many networks. Peering is a bilateral exchange of traffic between networks; it is distinct from transit, where one provider carries traffic to destinations beyond the customer network. These concepts and the operational distinctions between public and private interconnection are described in RFC 7454, BGP Operations and Security and Nokia's BGP documentation.

QUESTION NO: 27

Click the exhibit.



Assume all router-IDs are properly configured. Which of the following configurations is required on the Alcatel-Lucent 7750 SR router A to establish an IPv6 BGP session to router D?

A.
```
group "IPv6_eBGP"
    family ipv6
    type external
    peer-as 65550
    neighbor 2001:DB8:3::3
    exit
exit
```

B.
```
group "IPv6_eBGP"
    type external
    peer-as 65550
    neighbor 2001:DB8:3::3
    exit
exit
```

C.
```
group "IPv6_eBGP"
    family ipv6
    peer-as 65550
    neighbor 2001:DB8:13::31
    exit
exit
```

D.
```
group "IPv6_eBGP"
    type external
    peer-as 65550
    neighbor 2001:DB8:13::31
    exit
exit
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: C

Explanation:

The configuration shown in *Option C* is required because an IPv6 BGP peering must be configured with an IPv6-capable BGP neighbor relationship and the IPv6 address family enabled for that peer. BGP still uses router IDs as 32-bit values, but those router IDs do not determine the transport address used to establish the TCP session. For an IPv6 session, router A must identify router D using the appropriate IPv6 peer address and activate the IPv6-unicast family so that IPv6 network-layer reachability information can be exchanged after the session is established.

On SR OS, BGP configuration is organized around the BGP group and neighbor, with address-family controls determining which NLRI families are negotiated and advertised. The required configuration therefore combines the correct IPv6 neighbor

definition with IPv6-family activation in the applicable BGP context. This follows the multiprotocol BGP mechanism, in which IPv6 routing information is carried using the multiprotocol extensions rather than the original IPv4-only BGP update format. See Nokia's SR OS BGP documentation and RFC 2545, Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing.