PassQueen
a.k.a dumpsqueen.com

Nokia Border Gateway Protocol Fundamentals for Services

Nokia 4A0-114

Version Demo

Total Demo Questions: 7

Total Premium Questions: 73

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, BGP Fundamentals	20
Topic 2, BGP Route Selection	6
Topic 3, BGP Policy	5
Topic 4, BGP Route Reflection	15
Topic 5, BGP Confederations	2
Topic 6, BGP Multiprotocol Extensions	14
Topic 7, BGP and MPLS VPNs	9
Topic 8, Mix Questions	2
Total	73

QUESTION NO: 1

What is the purpose of an Ethernet Auto-Discovery (type 1) EVPN route?

A. It advertises EVPN Ethernet segment and service membership information.

B. It advertises the MAC address of an individual host.

C. It advertises an IP prefix reachable through an EVPN instance.

D. It advertises the default gateway MAC address for a subnet.

ANSWER: A

Explanation:

It advertises EVPN Ethernet segment and service membership information. is correct. An Ethernet Auto-Discovery route, also called an EAD route, is EVPN route type 1. It is used to advertise information associated with an Ethernet segment and an EVPN instance. In multihomed EVPN deployments, type 1 routes participate in procedures such as split-horizon operation and aliasing.

A type 1 route does not advertise an individual endpoint MAC address. Per-host MAC reachability is carried in an EVPN MAC/IP Advertisement route, which is route type 2. See RFC 7432, Section 7.1: Ethernet Auto-Discovery Route.

QUESTION NO: 2

Which of the following is considered a best practice for BGP address planning?

A. Advertising external networks into IGP.

B. Using system addresses for setting up iBGP sessions.

C. Exporting the networks associated with the external links of eBGP peers into BGP.

D. Optimizing BGP route table by exporting IGP routes into BGP.

ANSWER: B

Explanation:

Using system addresses for setting up iBGP sessions is a BGP address-planning best practice because a system address is a stable, router-level address that is not tied to the operational state of any one physical interface. On Nokia SR OS routers, the system interface is typically configured with a loopback-style address and is reachable through the IGP. This lets an iBGP session remain available when an individual link or interface fails, provided that the IGP can still find another path to the remote system address. It also provides a consistent endpoint for peer configuration, router identification, and operational troubleshooting, rather than requiring the BGP design to depend on transient link addresses. For this design to work correctly, the system addresses must be advertised and reachable in the IGP, and the BGP peer configuration must use the appropriate source address or local-address setting. This approach supports resilient, scalable iBGP connectivity, particularly in networks that use redundant links between routers. Nokia's SR OS routing documentation describes the system interface and BGP peer configuration concepts used for this design. See the Nokia SR OS BGP documentation and the Nokia SR OS system-interface documentation.

QUESTION NO: 3

A traditional (non EVPN) VPRN service is configured using MP-BGP and LDP transport tunnels. Which of the following statements about the VPN service label is TRUE?

A. The service label is advertised using LDP.

B. The service label is prepended to the customer IPv4 prefixes.

C. The service label must be manually configured for the VPRN on each PE router.

D. The service label is carried in an MP-BGP VPN-IPv4 route update.

ANSWER: D

Explanation:

The service label is carried in an MP-BGP VPN-IPv4 route update. In a traditional VPRN, MP-BGP distributes VPN-IPv4 reachability between PE routers. A VPN-IPv4 route consists of a route distinguisher together with the customer IPv4 prefix, and the BGP update includes an MPLS label allocated by the advertising egress PE for that VPN route. The ingress PE uses this advertised VPN/service label as the inner label when sending traffic toward the remote PE. On receipt at the remote PE, that label identifies the appropriate VPRN forwarding context and enables delivery toward the customer destination.

LDP has a separate role in this design: it establishes the transport label-switched path across the provider core. Consequently, packets normally carry a label stack, with the LDP-signaled transport label on the outside and the MP-BGP-advertised VPN/service label inside. This separation permits the provider core to transport packets without maintaining customer VPN routes, while PE routers use MP-BGP VPN route advertisements to associate customer prefixes with their service labels. This behavior follows the BGP/MPLS IP VPN model specified in RFC 4364; the encoding of labeled routes in BGP is defined by RFC 3107.

QUESTION NO: 4

Which BGP community value prevents a route from being advertised to any eBGP peer?

A. NO_ADVERTISE

B. NO_EXPORT

C. LOCAL_AS

D. INTERNET

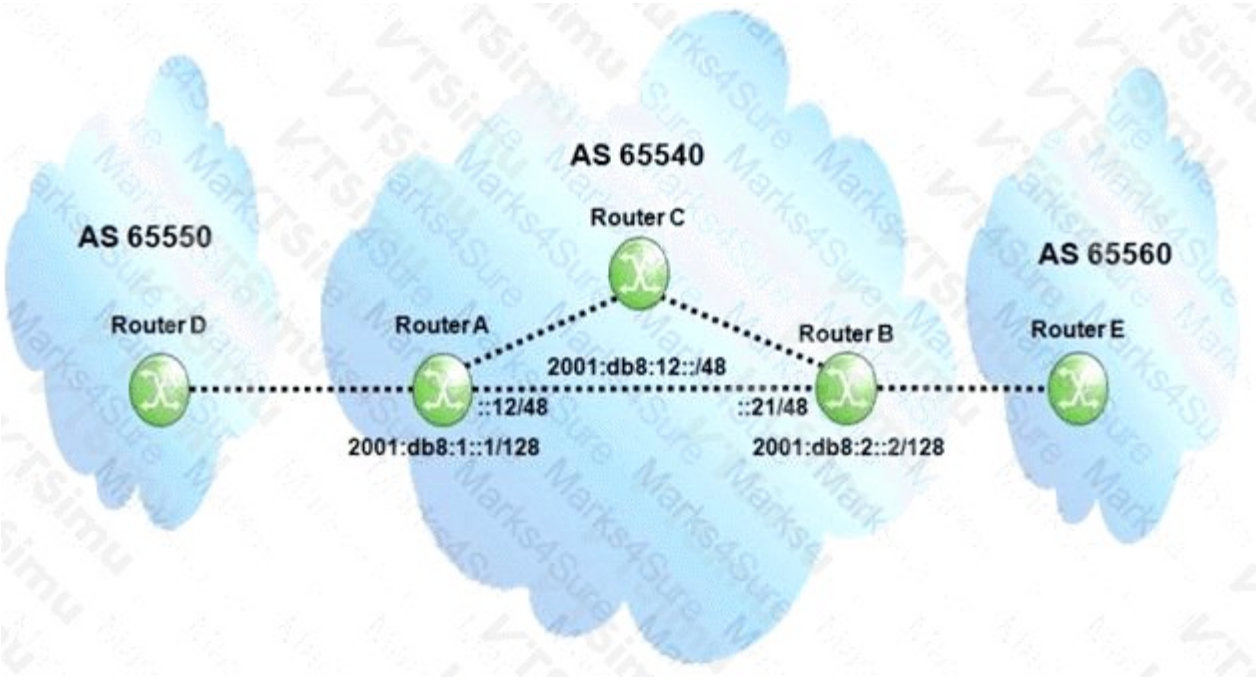
ANSWER: B

Explanation:

NO_EXPORT is correct. The well-known NO_EXPORT community instructs a BGP speaker not to advertise the route outside the local autonomous system. The route can still be advertised between iBGP peers within the AS and may be used internally, subject to normal BGP policy. This community is useful where a route must be distributed throughout an AS but must not be propagated to external autonomous systems. The NO_ADVERTISE community is more restrictive because it prevents advertisement to any BGP peer. See RFC 1997, Well-known Communities.

QUESTION NO: 5

Click the exhibit.



Assume all router-IDs are properly configured, and global IPv6 addresses are used to establish eBGP sessions. Which of the following configurations is required on router A so that router B can exchange traffic with router D?

A.
```
group "IPv6_iBGP"
    family ipv6
    next-hop-self
    peer-as 65540
    neighbor 2001:DB8:2::2
    exit
exit
```

B.
```
group "IPv6_iBGP"
    next-hop-self
    type internal
    peer-as 65540
    neighbor 2001:DB8:2::2
    exit
exit
```

C.
```
group "IPv6_iBGP"
    family ipv6
    peer-as 65540
    neighbor 2001:DB8:2::2
    exit
exit
```

D.
```
group "IPv6_iBGP"
    family ipv6
    next-hop-self
    type internal
    neighbor 2001:DB8:2::2
    exit
exit
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: D

Explanation:

The required configuration is the one that builds an IPv6 iBGP peering from router A to router B and applies `next-hop-self`. Router A and router B belong to the same autonomous system, so their BGP session must be configured as an internal peer relationship. Because the peering address is IPv6, the IPv6 address family must also be enabled for the BGP group so that IPv6 NLRI can be exchanged.

The key requirement for actual forwarding is `next-hop-self`. When router A advertises to router B a route that it learned from router D through eBGP, BGP would otherwise normally retain the externally learned next-hop. Router B may not have a route to that external IPv6 next-hop. Applying `next-hop-self` causes router A to advertise itself as the next hop, making router A the reachable forwarding point for traffic that router B sends toward destinations learned from router D. Thus, the complete configuration needs the IPv6 family, an internal peer type, router B's IPv6 neighbor address, and next-hop-self.

Nokia's SR OS BGP documentation describes BGP group and neighbor configuration, including address-family activation and iBGP peer settings: Nokia SR OS BGP documentation. Nokia's routing documentation is also available through the SR OS documentation portal.

QUESTION NO: 6

Which of the following regarding route reflectors is TRUE?

A. Only the route reflector is configured with a cluster id.

B. There can only be one route reflector in a cluster.

C. Clients should have iBGP sessions with each other to maintain the mesh.

D. There can only be one level of route reflectors in the topology.

ANSWER: A

Explanation:

Only the route reflector is configured with a cluster id. A cluster ID identifies the route-reflection cluster and is attached to routes reflected by the route reflector in the non-transitive CLUSTER_LIST attribute. When an RR receives an update whose CLUSTER_LIST already contains its own cluster ID, it recognizes that the route has traversed the cluster before and discards it. This supplies loop prevention for reflected iBGP routes, which do not use the normal iBGP AS-path modification behavior to detect every possible reflection loop.

When redundancy is deployed, multiple route reflectors may serve the same client group. Those route reflectors are configured with the same cluster ID so that they are treated as members of one reflection cluster and can consistently detect routes previously reflected within it. The RR is therefore the device that requires the cluster-id configuration; client routers participate as iBGP clients and process reflected routes, but do not operate as reflectors for that cluster. The cluster ID is commonly derived from an IPv4 address, although it is an identifier rather than a reachable address requirement. RFC 4456 defines route reflection, CLUSTER_LIST handling, and cluster loop detection: RFC 4456. Nokia SR OS BGP configuration documentation also covers route-reflector cluster configuration: Nokia SR OS BGP documentation.

QUESTION NO: 7

Which information is carried in the MP_REACH_NLRI attribute?

A. Reachable NLRI and the next-hop information for that NLRI.

B. Withdrawn NLRI and the negotiated Hold Time.

C. The BGP Identifier and the local AS number.

D. The CLUSTER_LIST and ORIGINATOR_ID attributes.

ANSWER: A

Explanation:

Reachable NLRI and the next-hop information for that NLRI. is correct. MP_REACH_NLRI is used by Multiprotocol BGP to advertise reachable destinations for an address family. It contains the AFI, SAFI, next-hop information, an optional SNPA field, and the reachable NLRI. The encoding allows the next hop to be represented appropriately for the address family being advertised.

Withdrawn reachability is carried in the MP_UNREACH_NLRI attribute. Standard IPv4-unicast reachable NLRI can instead be carried in the conventional UPDATE message NLRI field. See RFC 4760, Section 3: MP_REACH_NLRI.