Nokia 4A0-C03 Dumps

Shared by Hewitt on 24-11-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

This PE connects to an IPv4 core and is being configured to support 6VPE. Neighbor 10.10.10.2 is another PE that also supports 6VPE. Which of the following completes this configuration section?

Exhibit:

(gl)[/configure router "Base" bgp] A:admin@PE1# info group "multi-bgp" { family { ? ? ? } } neighbor "10.10.10.2" { group "multi-bgp" local-address 10.10.10.1 peer-as 64496 }

A)

family { ipv4 true }

B)

family { ipv6 true }

C)

family { ipv4 true ipv6 true }

D)

family { vpn-ipv6 true }

Options:

A- Option A
B- Option B
C- Option C
D- Option D

Answer:

D

Explanation:

IPv6 VPN Provider Edge (6VPE) enables service providers to deliver IPv6 Layer 3 VPN services across an existing IPv4 MPLS backbone.

With 6VPE, customer-facing Provider Edge routers maintain IPv6-capable VPRN instances, while the backbone can continue to use IPv4 routing and MPLS transport mechanisms.

The participating PE routers exchange IPv6 VPN routes using the MP-BGP VPN-IPv6 address family.

CE1 (IPv6) --- PE1 (10.10.10.1) --- P (IPv4/MPLS) --- PE2 (10.10.10.2) --- CE2 (IPv6).

MP-BGP VPN-IPv6 between PE1 and PE2; MPLS transport LSP across the IPv4/MPLS backbone. Core P routers do not require customer IPv6 routes.

A VPN-IPv6 route combines an IPv6 prefix with a Route Distinguisher and carries the associated VPN label and applicable Route Target communities.

The ordinary IPv6 unicast address family does not provide the VPN-specific route representation required for traditional 6VPE service signaling.

The appropriate Nokia MD-CLI configuration is:

group 'multi-bgp' { family { vpn-ipv6 true } }

This enables the exchange of VPN-IPv6 routing information between the participating PE routers.

The BGP session itself can operate over IPv4 TCP connectivity. The underlying provider core does not need native IPv6 customer routing to transport the encapsulated VPN traffic.

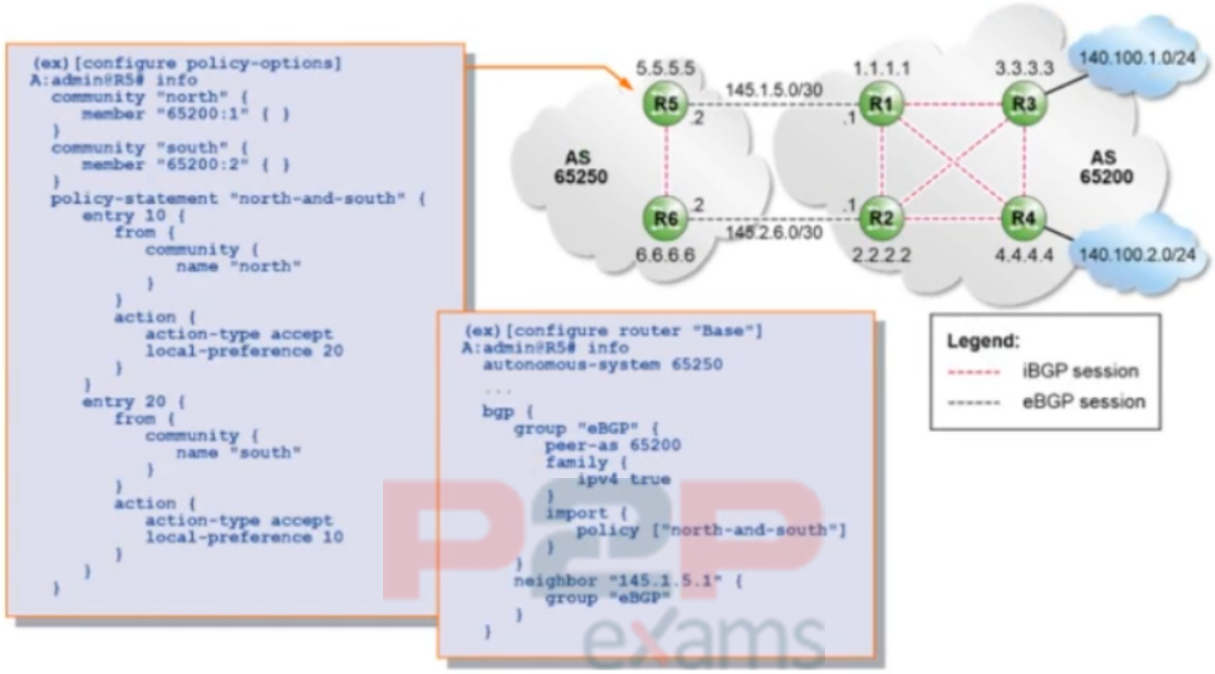
Therefore, option D is correct.

Question 2

Question Type: MultipleChoice

The BGP import policy north-and-south is applied on router R5. Router R1 advertises prefix 140.100.1.0/24 with its community set to 65200:1. Which statement about R5's processing of the received route update for prefix 140.100.1.0/24 is TRUE?

Exhibit:



[ex][configure policy-options] A:admin@R5# info community "north" { member "65200:1" { } } community "south" { member "65200:2" { } } policy-statement "north-and-south" { entry 10 { from { community { name "north" } } action { action-type accept local-preference 20 } } entry 20 { from { community { name "south" } } action { action-type accept local-preference 10 } } }

Options:

A- The prefix is rejected.

B- The prefix's local preference is set to 20.

C- The prefix's community value is removed.

D- The prefix's community value is modified to 65200:2.

Answer:

B

Explanation:

Nokia SR OS routing policies allow BGP route advertisements to be evaluated against configured match criteria and modified through associated policy actions.

The north-and-south import policy contains separate entries matching the BGP communities named north and south.

The route advertised by R1 contains community 65200:1, which matches the community definition named north.

Received BGP route: 140.100.1.0/24 (Community: 65200:1) Policy entry 10 (Matches community

'north') Set LOCAL_PREF = 20 Accept route (Community remains 65200:1).

Once the route matches the first policy entry, the configured action accepts the route and sets its BGP LOCAL_PREFERENCE attribute to 20.

The LOCAL_PREFERENCE attribute influences BGP best-path selection within the local autonomous system. When otherwise eligible routes are compared, higher local-preference values are preferred.

The configured policy entry does not contain any community deletion, addition, or replacement action. Therefore, the original community value remains unchanged.

The second policy entry matches community 65200:2 and does not apply to this advertisement.

Consequently, the advertised prefix is accepted with its local preference set to 20, making option B correct.

Question 3

Question Type: MultipleChoice

Both outputs are from router R4. The show service command shows that an SDP is down. What could be the cause?

Exhibit:

A:admin@R4# show service id 333 base Service Basic Information Service Id : 333 Vpn Id : 0 Service Type : VPLS <.. output omitted ..> MPLS ECMP : Disabled Ignore MTU Mismatch : Disabled Tunnel ELMI : Disabled Service Access & Destination Points Identifier Type AdmMTU OprMTU Adm Opr ---------- sdp:1:333 M(10.10.10.1) Mesh 0 1556 Up Down sdp:2:333 M(10.10.10.2) Mesh 0 1556 Up Up sdp:3:333 M(10.10.10.3) Mesh 0 1556 Up Up A:admin@R4# show router ldp bindings services service-id 333 <.. output omitted ..> LDP Service FEC 128 Bindings Type VCId SDPId IngLbl LMTU Peer SvcId EgrLbl RMTU ---------- ---------- V-Eth 333 1 524266U 1500 10.10.10.1:0 333 524266 1386 V-Eth 333 2 524283U 1500 10.10.10.2:0 333 524275 1500 V-Eth 333 3 524282U 1500 10.10.10.3:0 333 524273 1500 No. of VC Labels: 3

Options:

A- The mesh SDP on router R4 has been disabled.
B- The mesh SDP on router 10.10.10.1 has been disabled.
C- The service MTU on router 10.10.10.1 has been set to a different value.
D- The IngLbl and EgrLbl that router R4 has for router 10.10.10.1 are identical.

Answer:

C

Explanation:

Nokia SR OS performs Maximum Transmission Unit (MTU) compatibility checks when establishing traditional LDP-signaled Layer 2 pseudowires.

During pseudowire signaling, participating Provider Edge routers exchange service-label bindings and associated interface parameters, including the advertised MTU.

The displayed command output identifies an MTU mismatch between router R4 and its peer 10.10.10.1.

SDP binding 1:333

Value

Local MTU

1500

Remote MTU

1386

MTU difference

114 bytes

Ignore MTU Mismatch

Disabled

Administrative state

Up

Operational state

Down

Because MTU mismatch handling is disabled, the incompatible advertised MTU values prevent the affected pseudowire binding from becoming operational.

The other two pseudowires advertise matching local and remote MTU values of 1500 bytes and are operationally up.

The administrative state of SDP binding 1:333 is Up, eliminating option A.

Identical ingress and egress MPLS label values are valid because labels have local significance within their respective label spaces.

The appropriate corrective action is to investigate and align the service MTU configurations at both endpoints.

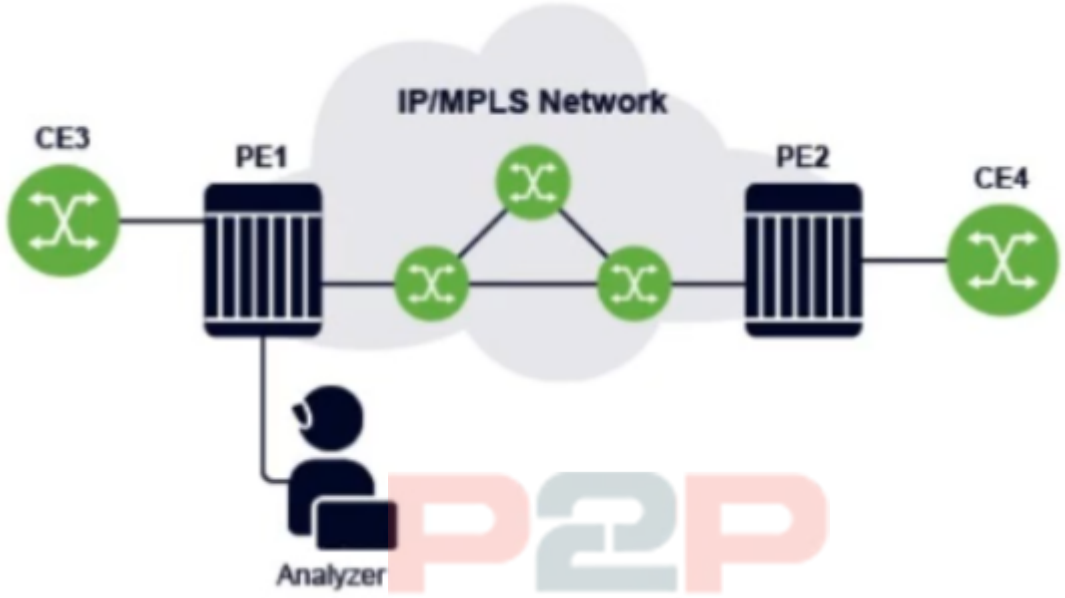
Therefore, option C is correct.

Question 4

Question Type: MultipleChoice

An operator at PE1 wants to analyze traffic between PE2 and CE4. On a Nokia 7750 SR, which of the following is required as part of the configuration?

Exhibit:



Options:

A- PE2 needs a mirror source with the remote source specified.
B- PE1 needs a mirror destination with the remote source specified.
C- PE1 needs a mirror source with the remote source specified.
D- PE1 and PE2 need a mirror source specified.

Answer:

B

Explanation:

Nokia SR OS supports remote traffic mirroring, enabling network operators to capture and analyze traffic traversing a remote router without requiring the analyzer to be physically connected to that router.

In the illustrated topology, the traffic to be monitored is exchanged between PE2 and CE4. However, the network analyzer is connected to PE1.

PE2 therefore operates as the mirror-source router, while PE1 operates as the final mirror-destination router.

Router

Required configuration

PE2

Mirror source identifying the monitored traffic and mirror service with remote transport

PE1

Mirror destination with remote-source configuration

PE1

Analyzer-facing destination SAP

IP/MPLS network

Operational transport connectivity

The mirror source on PE2 identifies the port, SAP, or supported traffic-selection criteria associated with the CE4 connection.

Copies of the selected packets are transported across the provider network using an appropriate SDP.

PE1 must configure its mirror destination to accept mirrored traffic originating from the authorized remote source. The replicated packets are subsequently delivered to the analyzer-facing SAP.

The original customer traffic continues to be forwarded normally.

Therefore, option B correctly identifies the required configuration on PE1.

Question 5

Question Type: MultipleChoice

How does advertising MP-BGP labeled IP prefixes enhance the scalability of VPN control-plane signaling?

Options:

A- They carry VPN-related information across regions, so there is no need to store per-VPN information on intermediate routers.
B- They help signal end-to-end inter-region transport tunnels, so there is no need to do per-VPN configuration on intermediate routers.
C- They help signal the intra-region transport tunnels, so LDP and RSVP are no longer needed.
D- They carry both transport and service labels, so only one MP-BGP address family is needed on the PE routers.

Answer:

B

Explanation:

Multiprotocol BGP Labeled Unicast (BGP-LU) enhances the scalability of MPLS VPN services by enabling labeled transport connectivity across multiple routing regions or autonomous systems.

In conventional inter-region VPN architectures, each region can maintain its own underlying IGP and MPLS transport infrastructure. BGP-LU distributes labeled IP prefixes representing reachable transport destinations, including remote Provider Edge (PE) router system addresses.

PE1 --- RBR1 --- RBR2 --- PE2: end-to-end BGP-LU transport across Regions 1, 2 and 3; MP-BGP VPN service signaling between PE endpoints.

No per-VPN state required on transit border routers.

Regional border routers use this labeled reachability information to connect individual transport segments into an end-to-end MPLS forwarding path.

In an Inter-AS Model C deployment, intermediate border routers do not need to maintain customer-specific VPN routing information. This separation significantly improves control-plane scalability.

The VPN service information and associated service labels are distributed separately through the VPN-IPv4 and VPN-IPv6 address families.

LDP, RSVP-TE, or segment routing can still provide the underlying intra-region transport connectivity.

Therefore, option B correctly identifies the scalability advantage of using MP-BGP labeled prefixes for inter-region transport signaling.

Question 6

Question Type: MultipleChoice

Which statement about the function of the Route Target (RT) is TRUE?

Options:

A- RT is carried in each MP-BGP route update as an extended community to indicate the service to which the route is relevant.
B- IP prefixes are concatenated with the RT to form unique VPN-IPv4 and VPN-IPv6 prefixes.
C- VPRN PE routers encapsulate customer packets with RT to identify the service tunnel used.
D- VPRN PE routers encapsulate customer packets with RT to identify the transport tunnel used.

Answer:

A

Explanation:

A Route Target is a BGP extended community used to control the distribution and import of VPN routing information between participating Provider Edge routers.

In a traditional Nokia VPRN service, the exporting PE associates eligible customer routes with the appropriate Route Target communities.

Receiving PE routers evaluate these communities against their configured VRF import policies to determine which VPN instances are eligible to import the received route.

A Route Distinguisher (RD) performs a different function. It is combined with the customer IP prefix to construct a unique VPN-IPv4 or VPN-IPv6 route representation.

Control-plane components: Route Distinguisher + Customer IP prefix Unique VPN prefix. Route Target controls which VPNs import the advertised route.

Data-plane components: Transport label + Service label + Customer packet.

Route Targets are control-plane attributes, not MPLS packet encapsulation fields.

Neither the Route Target nor the Route Distinguisher serves as an MPLS forwarding label.

Customer traffic is transported using the appropriate VPN service label and backbone transport label.

Therefore, option A correctly identifies the purpose of the Route Target in MP-BGP VPN route advertisements.

Technical clarification: Route Targets are carried in applicable VPN route advertisements, not necessarily in every MP-BGP UPDATE across all supported address families.

Question 7

Question Type: MultipleChoice

Which statement below about BGP peering sessions is FALSE?

Options:

A- BGP peers need IP reachability between them.
B- iBGP peers can be non-adjacent.
C- eBGP peers have to be adjacent to each other.
D- BGP peers can be configured with different AS numbers.

Answer:

C

Explanation:

Border Gateway Protocol establishes routing relationships between devices using TCP connections. The fundamental transport requirement is IP reachability between the configured BGP peering addresses.

Internal BGP (iBGP) sessions are established between routers belonging to the same autonomous system. These routers do not need to be directly connected because their peering addresses can be reached through an underlying routing protocol.

External BGP (eBGP) sessions are established between routers belonging to different autonomous systems.

Although conventional eBGP sessions frequently use directly connected interfaces, Nokia SR OS also supports multihop eBGP peering.

Conventional single-hop eBGP: AS 65001 --- eBGP --- AS 65002.

Multihop eBGP: AS 65001 --- R1 --- R2 --- AS 65002; multihop TCP/BGP session between the AS endpoints.

Multihop eBGP permits routers separated by intermediate routing devices to establish an external BGP relationship, provided the necessary reachability and appropriate TTL configuration exist.

Therefore, physical adjacency is not an unconditional requirement for eBGP.

Option C is false because it excludes supported multihop BGP operation.

P2P exams

Question 8

Question Type: MultipleChoice

Which of these statements about BGP is FALSE?

Options:

A- BGP uses Hello messages to dynamically discover BGP peers.
B- To establish a BGP session, a TCP session must first be established between the two routers.
C- BGP uses the OPEN message to exchange capability parameters between routers.
D- BGP was originally designed as an exterior routing protocol.

P2P exams

Answer:

A

Explanation:

Border Gateway Protocol is a path-vector routing protocol originally designed to exchange reachability information between autonomous systems.

Unlike link-state routing protocols such as IS-IS and OSPF, conventional BGP does not use periodic Hello messages to discover neighboring routers automatically.

BGP peers are ordinarily identified through administrative configuration or supported dynamic-neighbor mechanisms. A TCP connection is subsequently established using port 179.

Once the TCP connection becomes operational, the routers exchange BGP OPEN messages containing information required to establish the BGP relationship.

The fundamental BGP message types include:

Message type

Function

OPEN

Establishes BGP peering and exchanges capabilities

UPDATE

Advertises and withdraws routing information

KEEPALIVE

Maintains BGP session liveness

NOTIFICATION

Reports protocol errors and terminates sessions

ROUTE-REFRESH

Requests route readvertisement

BGP KEEPALIVE messages are exchanged after the initial TCP and BGP establishment procedures. They are not neighbor-discovery Hello messages.

Although modern Nokia SR OS supports dynamic BGP neighbor configurations, the base BGP protocol does not implement an IS-IS-style or OSPF-style Hello discovery mechanism.

Therefore, option A is false.

Question 9

Question Type: MultipleChoice

Which of these statements concerning VPRN VRF tables is FALSE?

Options:

A- A PE uses the same VRF table for all the VPRN services on it.

B- A PE uses a VRF table to maintain forwarding information.

C- VRF information is advertised between PEs using MP-BGP.
D- A PE uses a VRF table to forward traffic to CEs.

Answer:

A

Explanation:

A Virtual Private Routed Network on Nokia SR OS uses a dedicated Virtual Routing and Forwarding instance to maintain customer-specific Layer 3 routing information.

Each conventional VPRN service is associated with its own routing context. This separation allows multiple customer VPNs to operate on the same physical provider edge router without mixing their routing and forwarding information.

For example, two different VPRN customers may both use the private IPv4 prefix 192.168.1.0/24. Separate VRF instances allow the provider edge router to maintain independent next-hop information for each customer.

Routing instance

Destination

Next hop

VPRN 100

192.168.1.0/24

Customer A

VPRN 200

192.168.1.0/24

Customer B

MP-BGP distributes eligible VPN routes between participating PE routers. Route Distinguishers provide VPN-prefix uniqueness, while Route Targets control route import and export.

When a packet reaches the terminating PE, the service label identifies the appropriate VPN forwarding treatment. The relevant customer routing and forwarding information determines how the packet is delivered toward its destination CE.

Using a single VRF for every independent VPRN would eliminate the required routing isolation.

Therefore, option A is false.

Question 10

Question Type: MultipleChoice

What is a benefit of VPRN when compared with VPLS?

Options:

A- Service providers can apply QoS, ingress/egress traffic shaping and policing on a per service basis.
B- Different customers with overlapping address space can be accommodated.
C- There is a clear demarcation of customer and service provider networks.
D- Customers can connect to the same service using different subnets.

Answer:

D

Explanation:

Virtual Private Routed Network (VPRN) and Virtual Private LAN Service (VPLS) provide different service architectures within Nokia SR OS.

VPLS provides Layer 2 Ethernet bridging across geographically distributed customer sites. It forwards traffic according to MAC addresses and emulates a shared Ethernet broadcast domain.

VPRN provides Layer 3 routed connectivity through dedicated Virtual Routing and Forwarding instances. The provider edge routers exchange customer routes and perform IP forwarding between customer subnetworks.

Service

Site A

Forwarding

Site B

VPLS --- Layer 2 bridging

192.168.1.0/24

L2 bridge

192.168.1.0/24

VPRN --- Layer 3 routing

192.168.1.0/24

L3 routing

192.168.2.0/24

Illustration of direct subnet connectivity through provider-managed VPN routing.

Both service types support QoS and traffic-management functionality. Separate VPLS instances also allow different customers to use overlapping IP address spaces because their Ethernet forwarding contexts are isolated.

The principal distinguishing advantage is that VPRN provides native IP routing between different customer subnetworks without requiring customer-managed routing across the provider's Layer 2 service.

Therefore, option D is the most appropriate answer.

Technical clarification: Multiple IP subnets can coexist within a VPLS broadcast domain. However, VPLS does not inherently provide Layer 3 routing between those subnets, whereas VPRN does.

To Get Premium Files for 4A0-C03 Visit
https://www.p2pexams.com/products/4a0-c03

For More Free Questions Visit
https://www.p2pexams.com/nokia/pdf/4a0-c03

