DUMPSARENA

Aruba Certified Switching Expert Written Exam

HP HPE6-A69

Version Demo

Total Demo Questions: 10

Total Premium Questions: 103

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Network Design	12
Topic 2, Network Implementation	50
Topic 3, Network Troubleshooting	19
Topic 4, Network Security	22
Total	103

QUESTION NO: 1

The customer is currently planning the migration of their current switches from the existing Aruba 5400R to ArubaOS-CX 6400. Which statements are correct about the mitigation of rogue DHCP servers with the selected product? (Select two)

A. With Aruba CX 6400 configure dhcp-snooping for the selected VLANs with authorized servers

B. DHCP snooping can be enabled on ArubaOS-CX switches, once it's disabled on ArubaOS switches for same VLANs

C. DHCP snooping is supported on both IPv4 and IPv6

D. DHCP snooping needs to be enabled per VRF basis on Aruba CX 6400

ANSWER: A B

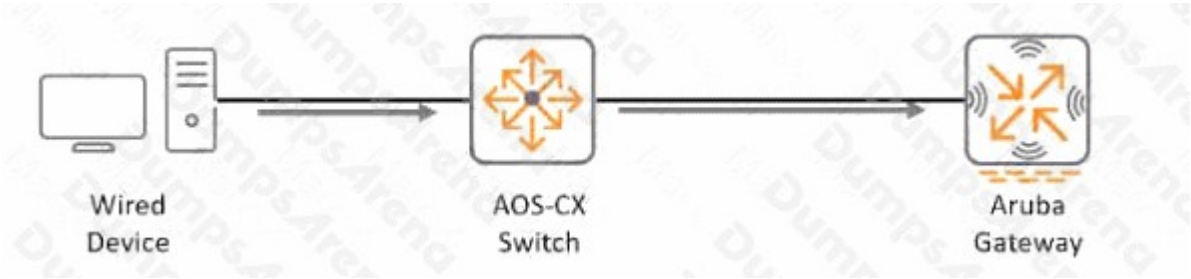
Explanation:

With Aruba CX 6400 configure dhcp-snooping for the selected VLANs with authorized servers is correct because DHCP snooping is a Layer 2 protection mechanism used to prevent untrusted access ports from supplying illegitimate DHCP server responses. On ArubaOS-CX, the feature is enabled for the relevant VLANs, while interfaces toward the legitimate DHCP infrastructure are configured as trusted. This permits valid DHCP offer and acknowledgement traffic to reach clients while enforcing inspection and filtering at the access layer.

DHCP snooping can be enabled on ArubaOS-CX switches, once it's disabled on ArubaOS switches for same VLANs is also correct in the context of a controlled migration. DHCP-snooping policy must be transitioned with the VLAN and its DHCP forwarding path so that enforcement is applied consistently by the active switching platform. Disabling the former ArubaOS-Switch enforcement for VLANs moved to the CX 6400 and enabling the corresponding CX policy supports a clean cutover and avoids retaining an obsolete enforcement point in the migrated path. ArubaOS-CX DHCP snooping configuration and trusted-interface behavior are documented in the ArubaOS-CX DHCP Snooping documentation. Platform capabilities for the chassis are available in the Aruba CX 6400 Switch Series datasheet.

QUESTION NO: 2

With the given topology, the customer Has ArubaOS-CX 6300 switches and Aruba Gateway in use.



What is required for the client traffic to be tunneled as per best practice between the connected switch port and the Aruba Gateway ' ' (Select two.)

A. IP Protocol 6 should not be blocked on me datapath

B. IP Protocol 47 should not be blocked in the data-path

C. The ArubaOS-CX switch and Aruba gateway should have an end-to-end MacSec connection

D. The ArubaOS-CX switch and Aruba gateway should be EBGP peers.

E. Change the default MTU on the data-path between the switch and gateway

ANSWER: B E

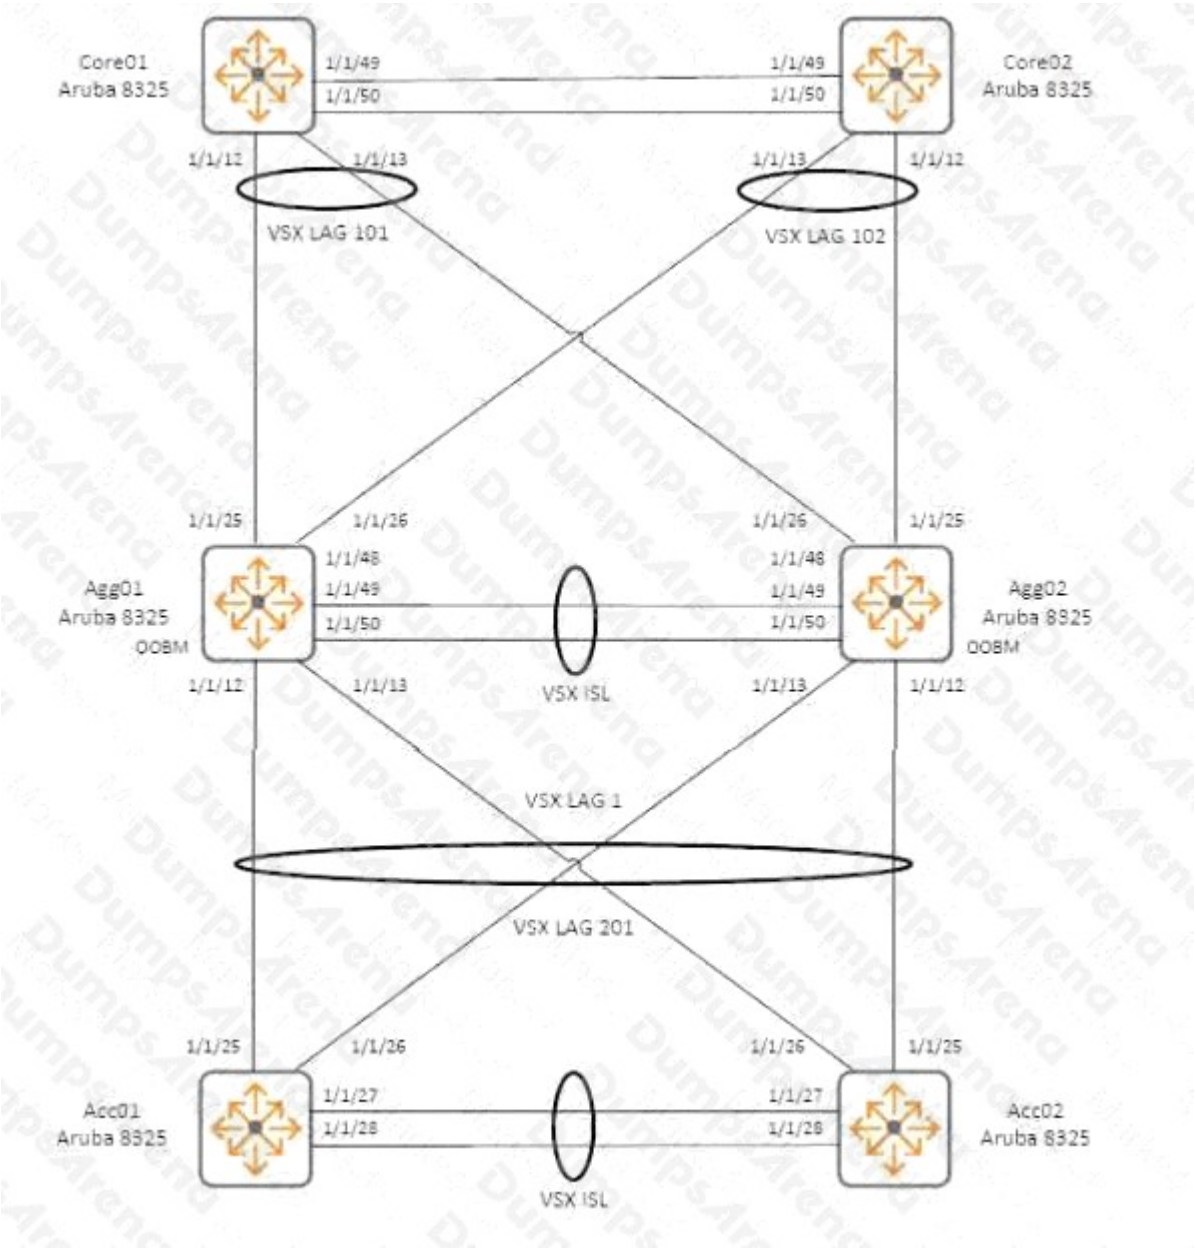
Explanation:

IP Protocol 47 should not be blocked in the data-path because ArubaOS-CX tunneled-node forwarding uses Generic Routing Encapsulation (GRE) to carry client traffic between the access switch and the Aruba Gateway. GRE is identified by IP protocol number 47, rather than a TCP or UDP port. Every routed hop, firewall, access-control policy, and security device on the path must therefore permit GRE traffic between the switch tunnel source and the gateway tunnel endpoint. This is essential for the switch to encapsulate traffic received on tunneled-node ports and for the gateway to decapsulate and apply the appropriate user-role policy.

Change the default MTU on the data-path between the switch and gateway because GRE adds encapsulation overhead to each tunneled frame. The underlay must provide sufficient MTU headroom for the original client packet plus the GRE and outer IP headers. Configuring an appropriately increased MTU end to end avoids fragmentation and prevents oversized packets from being dropped, which is particularly important for reliable client connectivity and higher-MTU application traffic. Aruba's tunneled-node documentation and switch security guidance describe the GRE-based tunnel model and the associated underlay requirements: ArubaOS-CX Tunneled Node documentation and ArubaOS-CX Security Guide.

QUESTION NO: 3

Refer to the exhibit.



You want to protect the aggregation layer if the VSXISL falls. Where should you place a VSX keepalive link?

A. On VSX LAG 1

B. On a dedicated link created using port 1 ' 1/48 of each aggregation switch

C. On the OOBM ports of both aggregation switches

D. On VSX LAG 101

ANSWER: C

Explanation:

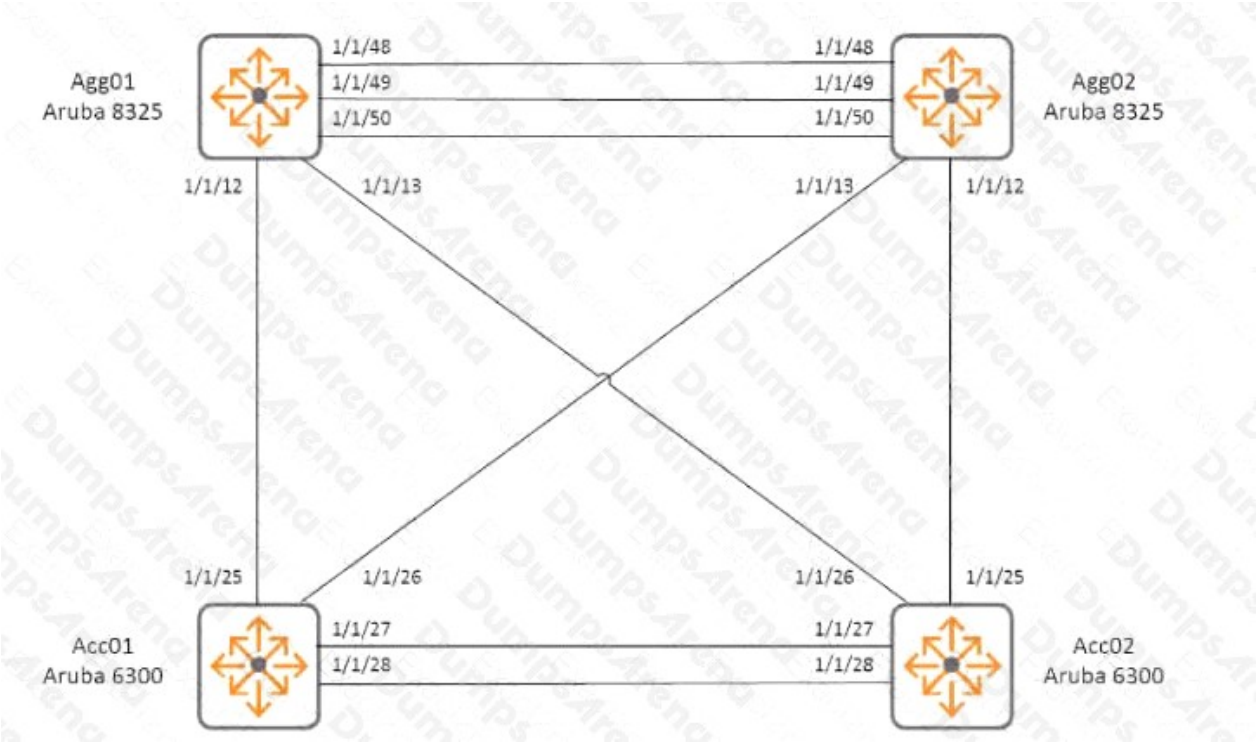
On the OOBM ports of both aggregation switches is correct because the VSX keepalive must use a path that is independent of the VSX inter-switch link and normal production forwarding topology. Its purpose is to allow each VSX peer to determine whether its counterpart is still operational when the inter-switch link is unavailable. A keepalive carried through the dedicated out-of-band management interfaces remains separate from the VSX inter-switch link, VSX LAGs, and data-plane links. This separation is essential for detecting an actual peer or connectivity failure and for helping prevent both peers from incorrectly operating independently after an inter-switch-link failure.

Aruba recommends using the management network for the VSX keepalive where possible. The out-of-band management interfaces provide a dedicated Layer 3 reachability path between the aggregation switches and avoid dependence on links whose status may be affected by the same fault domain as the VSX inter-switch link. The keepalive does not carry user traffic; it is a peer-health mechanism used alongside the inter-switch link, rather than as a replacement for it. Aruba's VSX documentation describes the keepalive role and its independence requirement: AOS-CX VSX Guide and VSX Keepalive documentation.

QUESTION NO: 4

(Scenarios may contain multiple errors which may or may not impact the solution)

Refer to the Exhibit.



An engineer has attempted to configure two pairs of switches m the referenced configuration it Is required to implement

Multi-Chassis Link Aggregation for each pair of switches

The ports of the Aruba 8325 switches used for Agg01 and Agg02 are populated as follows

```
1/1/12 10G SFP+ LC SR 300m MMF Transceiver
1/1/13 10G SFP+ LC SR 300m MMF Transceiver
1/1/48 1GBaseT 100m Cat5e Transceiver
1/1/49 40G QSFP+ 15m Active Optical Cable
1/1/50 40G QSFP+ 15m Active Optical Cable
```

The configuration includes:

```
!
!Version ArubaOS-CX GL.10.04.2000
!export-password: default
hostname Agg01
profile L3-agg
no usb
vrf KA
ntp server 10.77.77.77
ntp vrf mgmt
interface mgmt
   no shutdown$
   ip static 10.177.177.70/24
   default-gateway 10.177.177.128
system interface-group 2 speed 10g
system interface-group 4 speed 10g
interface lag 1
   no shutdown
   no routing
   vlan trunk native 1
   vlan trunk allowed 700-701
   lacp mode active
interface lag 2
interface 1/1/50
  no shutdown
  mtu 9198
  lag 256
vsx
  system-mac 02:01:00:00:20:00
  inter-switch-link lag 256
  role primary
  keepalive peer 192.168.20.2 source 192.168.20.1 vrf KA
  linkup-delay-timer 600
  vsx-sync aaa acl-log-timer bfd-global bgp copp-policy dhcp-relay dhcp-server dhcp-
snooping dns icmp-tcp lldp loop-protect-global mac-lockout mclag-interfaces neighbor ospf
qos-global route-map sflow-global   snmp
ssh stp-global time vsx-global
ip dns server-address 10.25.110.250 vrf mgmt
https-server rest access-mode read-write
https-server vrf mgmt
```

There is an error message stating " mismatched group speed " What should you add to the configuration to correct the error?

A. " system Interface-group 1 speed I0g "

B. " allow-unsupported-transceiver

C. " speed 10-full " under interfaces 1/1/12 and 1/1/13

D. " speed 40-full " under interfaces 1/1/49 and 1/1/50

ANSWER: C

Explanation:

Adding **" speed 10-full " under interfaces 1/1/12 and 1/1/13** explicitly places those physical links at the same 10-Gb/s full-duplex operating speed. A multi-chassis LAG presents member links from both VSX peers as one logical aggregation to the connected device. For that aggregation to operate consistently, corresponding member links must have matching operational characteristics, including link speed. A speed mismatch can arise when port defaults, transceiver behavior, or negotiation

cause the affected interfaces to operate at different speeds than the intended LAG group. Explicitly configuring 10-Gb/s full duplex on the two listed interfaces removes that ambiguity and establishes the common group speed required by the MC-LAG design. This is particularly important for links using 10-Gb/s-capable interfaces where automatic negotiation or inherited defaults might not produce the intended operational state. Aruba's AOS-CX documentation describes link aggregation and LACP configuration requirements, while Aruba's 8325 documentation provides the platform-specific interface capabilities and configuration guidance. See the Aruba AOS-CX technical documentation and the Aruba 8325 Switch Series data sheet.

QUESTION NO: 5

The example ACL is used with Aruba CX 6400 switches to permit traffic to web servers configured on interface 1/1/1 inbound with IP 10.254.4.1/24.

```
access-list ip ACL_web_server
  10 permit tcp any 10.254.0.0/22 eq 443
  20 permit tcp any 10.254.0.0/22 eq 80
  30 permit udp any any eq 67
  40 permit udp any any eq 68
  50 permit udp any any eq 53
  60 permit icmp any 10.254.0.0/22
```

How does this ACL affect the HTTPS traffic to destination 10.254.4.7 from a client 10.253.1.5 via interface 1/1/2 with IP 10 254.3.1/24?

A. The ACL would be valid only if bound to 1/1/2 outbound.

B. The implicit rule allows the traffic

C. The implicit rules deny the traffic.

D. The traffic is allowed when ACL is Pound to inbound traffic.

E. The ACL does not affect this traffic because it enters interface 1/1/2, not interface 1/1/1 inbound.

ANSWER: E

Explanation:

The HTTPS request from 10.253.1.5 to 10.254.4.7 enters the switch through interface 1/1/2 and is then routed toward the server subnet through interface 1/1/1. An ACL associated with interface 1/1/1 in the inbound direction is evaluated only for packets arriving *from* the network connected to 1/1/1. The client-to-server request does not arrive on that interface, so it is not subject to that inbound ACL at all. Consequently, this ACL does not decide whether the stated HTTPS request is permitted or denied; forwarding is determined by the applicable routing and any ACLs applied at the packet's actual ingress or egress processing point. To filter this client-to-server flow using an interface ACL, the policy should be placed inbound on 1/1/2, where the request enters the switch, or in an appropriate outbound direction on the server-facing path if that direction is supported for the intended ACL type and platform. Aruba documents that interface ACL direction determines whether packets are matched as they enter or leave the interface. See the Aruba AOS-CX Access Control Lists guide and the ip access-group command reference.

QUESTION NO: 6

Which statements are true with regards to Strict Priority scheduling? (Select two.)

A. Strict Priority scheduling protects against starvation of low priority queues.

B. Under strict Priority, the port will never forward a packet in a lower priority queue if a higher priority queue includes even one packet.

C. Strict Priority must be used with egress queue shaping to prevent starvation.

D. Strict Priority works best when the network uses the higher priority queues for the majority of applications

ANSWER: B C

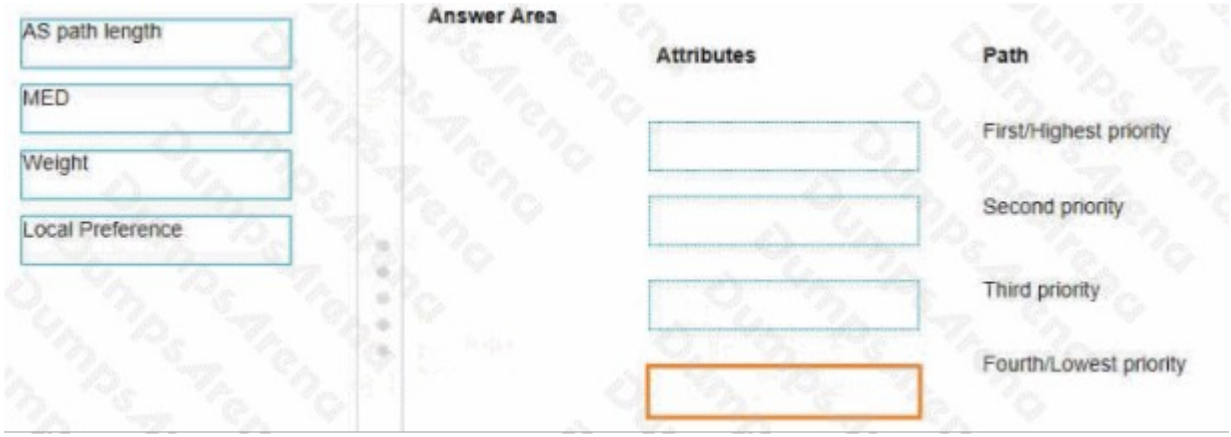
Explanation:

“Under strict Priority, the port will never forward a packet in a lower priority queue if a higher priority queue includes even one packet” describes the defining behavior of strict-priority queue scheduling. The scheduler always services the highest-priority nonempty queue first. As long as traffic remains queued at that higher priority, lower-priority queues do not receive transmission opportunities. This behavior is useful for delay- and jitter-sensitive traffic when the high-priority class is tightly controlled.

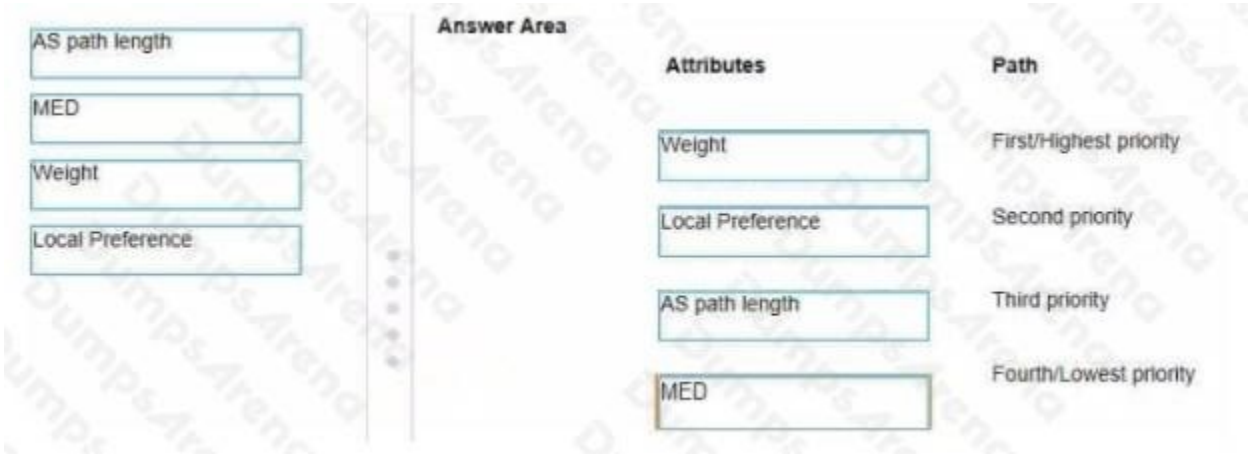
“Strict Priority must be used with egress queue shaping to prevent starvation” reflects the required design safeguard when strict priority is deployed. Because continuous high-priority traffic can consume all available egress bandwidth, shaping the prioritized traffic constrains its rate and leaves capacity in which lower-priority queues can be serviced. Egress queue shaping therefore makes strict-priority treatment suitable for production use by protecting the service availability of lower-priority traffic while retaining expedited treatment for the intended class. Aruba’s QoS documentation describes strict scheduling and queue-rate controls in its QoS guidance; see the Aruba AOS-CX queueing documentation and the Aruba AOS-CX shaping documentation.

QUESTION NO: 7 - (DRAG DROP)

Match the BGP path selection criteria attributes (Each option may be used once, more than once, or not at all)



ANSWER:



Explanation:

BGP selects a single best route by evaluating path attributes in a defined sequence, so the relative order of these four attributes is important. **Weight** is the first/highest-priority attribute in this group. It is a local routing-policy value and is evaluated before the standard propagated BGP attributes, allowing the local router to strongly prefer one learned path over another. The path with the higher Weight is preferred.

Next, BGP evaluates **Local Preference**. Local Preference communicates path preference throughout the local autonomous system, making it the next decision point after Weight. A higher Local Preference indicates the preferred exit path for traffic leaving the autonomous system.

When Weight and Local Preference do not produce a unique path, BGP then compares **AS path length**. The route that has traversed fewer autonomous systems is preferred, which generally identifies a shorter interdomain route. This places AS path length after Local Preference but before MED in the requested ordering.

Finally, **MED** is the fourth/lowest-priority item among the listed attributes. MED is used to signal a preferred ingress point into a neighboring autonomous system, and a lower metric is preferred when MED comparison is applicable. Aruba documents BGP route selection and its ordered best-path criteria in the AOS-CX BGP routing documentation. The correct priority sequence is therefore Weight, Local Preference, AS path length, then MED.

QUESTION NO: 8

The customer is implementing new ArubaOS-CX 6300 switches into an environment that has third-party switches and uses centralized authentication from Microsoft Active Directory and TACACS* servers

What would be a valid configuration requirement for ArubaOS-CX switches for centralized user authentication?

A. Importing a TACACS server certificate for TLS authentication into the switch

B. Defining a TACACS server group for authentication in the switch configuration

C. Enabling the active directory VSA in the switch for radius authentication

D. Using a PAM extension for LDAP authentication with Active Directory Server.

ANSWER: B

Explanation:

Defining a TACACS server group for authentication in the switch configuration is a valid requirement because ArubaOS-CX uses TACACS+ server definitions and groups as part of its centralized administrative-access authentication configuration. The switch must know the address of each TACACS+ server, the shared secret used to protect TACACS+ exchanges, and the server group to which those servers belong. That group can then be referenced in an authentication method list for management login access, normally with a local method retained as a fallback to prevent loss of administrative access if centralized services are unavailable. In this design, the TACACS+ servers can in turn use Microsoft Active Directory as their identity store, allowing the organization to maintain user accounts centrally while enforcing device-administration authorization and accounting policies through TACACS+. Server groups also support resiliency by allowing more than one TACACS+ server to be configured and contacted according to the defined group behavior. Aruba's AOS-CX Security Guide documents TACACS+ server and authentication configuration for the 6300 platform: AOS-CX 6300 and 6400 Security Guide. The TACACS+ protocol's role in centralized device administration is also specified in RFC 8907.

QUESTION NO: 9

An administrator configures an IPv4 static route on an ArubaOS-CX switch using the reject route type.

What happens when a packet matches this route?

A. The switch forwards the packet to the CPU for normal routing.

B. The switch silently discards the packet without notifying the sender.

C. The switch discards the packet and returns an ICMP destination-unreachable message.

D. The switch floods the packet on every interface in the matching VRF.

ANSWER: C

Explanation:

A reject static route prevents forwarding for the matching destination prefix and causes the switch to discard the packet while returning an ICMP destination-unreachable message to the sender. This differs from a blackhole route, which silently discards matching packets. Reject routes are useful when the administrator wants to explicitly inform the sending host that the destination is unreachable.

See the AOS-CX Static Routes documentation.

QUESTION NO: 10

A customer requires centralized time synchronization for ArubaOS-CX switches so that event logs from multiple devices can be correlated accurately.

Which statements about NTP are true? (Select two.)

A. It synchronizes the switch clock with a configured time source.

B. Multiple NTP servers can be configured for redundancy.

C. It distributes the running configuration to managed switches.

D. It replaces TACACS+ for administrator authentication.

E. It automatically creates routing adjacencies.

ANSWER: A B

Explanation:

NTP synchronizes the switch clock to a configured time source, allowing event timestamps to be correlated across switches and management systems. Multiple NTP servers can be configured to provide redundancy if one time source becomes unavailable.

NTP does not distribute switch configurations and does not provide user authentication for administrative access. See the ArubaOS-CX NTP documentation.