PassQueen
a.k.a dumpsqueen.com

Aruba Certified Switching Professional Exam

HP HPE6-A73

Version Demo

Total Demo Questions: 17

Total Premium Questions: 173

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Planning and Designing	36
Topic 2, Implementing	108
Topic 3, Troubleshooting	29
Total	173

QUESTION NO: 1

Examine the partial output of the BGP routing table of an AOS-CX switch:

```
Switch# show bgp

<-output omitted->

   Network      Nexthop      Metric  LocPrf  Weight  Path
-------------------------------------------------------------
* e 1.0.0.0/8  192.168.1.5  0       100     0       100 ?
* e 1.0.0.0/8  192.168.2.5  0       100     0       200 100 i
* e 1.0.0.0/8  192.168.3.5  0       200     20      300 400 100 ?
* e 1.0.0.0/8  192.168.4.5  0       50      0       400 200 100 i
```

The switch is learning about four possible path to reach the 1.0.0.0/8 network. Based on this output, which next-hop route will the AOS-CX select to be placed in the IP routing table?

A. 192.168.1.5

B. 192.168.2.5

C. 192.168.3.5

D. 192.168.4.5

ANSWER: C

Explanation:

192.168.3.5 is the next hop selected for installation in the IP routing table because the corresponding BGP path is the best path in the displayed BGP table. In AOS-CX BGP output, the best-path indicator identifies the route that has won BGP's path-selection process for a given prefix. BGP evaluates eligible paths in a defined order, beginning with attributes such as Weight and Local Preference, then considering AS-path length, origin, MED where applicable, eBGP versus iBGP status, IGP cost to the next hop, and further tie-breakers as needed. The route through 192.168.3.5 has the preferred attribute combination shown in the output and is therefore the active BGP route for 1.0.0.0/8. Once its next hop is reachable, AOS-CX installs that selected BGP route into the routing information base; the remaining learned paths can be retained as BGP alternatives but are not installed as the active route for this prefix. This behavior follows the BGP best-path-selection rules described in RFC 4271 and Aruba's AOS-CX BGP documentation.

QUESTION NO: 2

What are best practices when implementing VSX on AOS-CX switches? (Choose two.)

A. The ISL lag should use the default MTU size.

B. Timers should be left at their default values.

C. The default system MAC addresses should be used.

D. The keepalive connection should use a direct layer-3 connection.

E. The ISL lag should use at least 10GbE links or faster.

ANSWER: B D

Explanation:

“Timers should be left at their default values” is a VSX best practice because the default keepalive and split-detection timer values are engineered for normal VSX operation. Altering them without a specific, validated design requirement can cause unnecessary sensitivity to transient delay or loss and complicate troubleshooting. Retaining the defaults provides predictable peer and keepalive behavior across the VSX pair.

“The keepalive connection should use a direct layer-3 connection” is also correct. The keepalive is a control-plane health mechanism used to determine whether the VSX peer remains reachable independently of the inter-switch link. A direct routed connection provides a simple, reliable, and independent path between the peers, helping the switches distinguish an inter-switch-link failure from a complete peer failure. This independence is important for split-brain protection and for stable VSX operation during link or topology failures. Aruba’s VSX documentation describes the keepalive as a Layer-3 connection and documents its role in VSX split detection and peer status. See the AOS-CX VSX documentation and the AOS-CX VSX overview.

QUESTION NO: 3

An AOS-CX switch uses RADIUS for 802.1X authentication and authorization.

Which authorization information can a RADIUS server return to dynamically apply access policy to an authenticated client? (Choose two.)

A. A VLAN assignment using RADIUS tunnel attributes

B. An Aruba user-role assignment

C. A DHCP address pool for the client

D. A DNS server address for the client operating system

E. A replacement IP address for the switch management interface

ANSWER: A B

Explanation:

RADIUS can return standard tunnel attributes to assign a client to a VLAN. For VLAN assignment, the RADIUS Access-Accept commonly includes attributes such as Tunnel-Type, Tunnel-Medium-Type, and Tunnel-Private-Group-ID. The switch uses the returned VLAN information when authorizing the client session.

A RADIUS server can also return an Aruba user-role attribute that identifies the role to apply to the authenticated client. The role can reference access policy such as ACLs, VLAN controls, QoS settings, or other supported session attributes. The switch enforces the authorization result after successful authentication.

See the RFC 2865 RADIUS specification and the AOS-CX Security Guide.

QUESTION NO: 4

A company has just purchased AOS-CX switches. The company has a free and open-source AAA solution. The company wants to implement access control on the Ethernet ports of the AOS-CX switches.

Which security features can the company implement given the equipment that they are using?

A. Port-based tunneling

B. Device fingerprinting

C. Local user roles

D. Downloadable user roles

ANSWER: C

Explanation:

Local user roles is the appropriate feature for this environment. AOS-CX supports port-access authentication, including IEEE 802.1X and MAC-based authentication, using an external RADIUS-based AAA service. A free, open-source RADIUS implementation can authenticate endpoint users or devices and return authorization information to the switch. The switch can then apply a user role that has been defined locally in its configuration.

A local user role provides the enforcement policy directly on the AOS-CX switch. Administrators can define the role's permitted traffic and network access behavior, such as access-control entries, VLAN assignment, and other session controls, then use that role as part of the port-access authorization design. This makes the approach suitable where the organization already has standards-based AAA but does not have a proprietary policy-management platform. The RADIUS server supplies identity and authorization results, while the switch retains and enforces the actual role definition. Aruba's AOS-CX security documentation describes port-access authentication and role-based authorization capabilities; see the AOS-CX Security Guide. For background on using RADIUS as the authentication and authorization component, see the FreeRADIUS documentation.

QUESTION NO: 5

A switch will apply a device profile to a port based on which pieces of information? (Select two.)

A. IP header

B. MAC address

C. LLDP

D. User role

E. 802.1Q

ANSWER: B C

Explanation:

Device profiles use endpoint-identification information learned on the switch port to recognize a connected device and apply the associated profile automatically. **MAC address** is a valid matching input because the switch learns the source MAC address of an endpoint on the port. A profile can use that Layer 2 identity, including a manufacturer-oriented MAC match where appropriate, to classify the connected device consistently.

LLDP is also a valid matching input. LLDP advertisements provide device-discovery attributes from an LLDP-capable endpoint, such as a system name, system description, capabilities, and other advertised identity details. These attributes enable the switch to identify device types such as phones or other managed infrastructure endpoints even when a simple MAC match is not sufficiently specific. After a device matches the configured profile criteria, the switch can automatically apply the profile's port-level settings, helping standardize endpoint deployment and reduce manual interface configuration. Aruba documents device-profile matching and LLDP-based endpoint discovery in its AOS-CX Security Guide: Device Profiles and the AOS-CX CLI Guide.

QUESTION NO: 6

Which statements are correct about OSPF link-state advertisements in a normal OSPF area? (Choose two.)

A. An ABR originates Type 3 summary LSAs to advertise inter-area prefixes.

B. An ASBR originates Type 5 LSAs to advertise external routes.

C. A DR originates Type 3 summary LSAs for all networks in its segment.

D. An internal router originates Type 5 LSAs for its connected interfaces.

E. An ABR uses Type 2 LSAs to advertise routes between areas.

ANSWER: A B

Explanation:

An ABR originates Type 3 summary LSAs to advertise inter-area prefixes is correct. Area border routers use Type 3 summary LSAs to describe networks from one area to routers in another area.

An ASBR originates Type 5 LSAs to advertise external routes is also correct for normal OSPF areas. When an autonomous system boundary router redistributes routes from another routing domain into OSPF, it advertises those external destinations by using Type 5 external LSAs. Type 5 LSAs are not flooded into stub areas or NSSAs. Aruba documents OSPF operation in the AOS-CX OSPF Guide; OSPF LSA behavior is defined in RFC 2328.

QUESTION NO: 7

Examine the following ACL rule policies:

- Permit traffic from 10.2.2.1 through 10.2.2.30 to anywhere
- Permit traffic from 10.2.2.40 through 10.2.2.55 to anywhere
- Deny all others

Based on this policy, place the following ACL rule statements in the correct order to accomplish the above filtering policy.

A. deny ip 10.2.2.31 255.255.255.255 any permit ip 10.2.2.40 255.255.255.248 any permit ip 10.2.2.48 255.255.255.248 any deny ip 10.2.2.32 255.255.255.224 any permit ip 10.2.2.0 255.255.255.192 any

B. permit ip 10.2.2.40 255.255.255.248 any permit ip 10.2.2.48 255.255.255.248 any permit ip 10.2.2.0 255.255.255.192 any deny ip 10.2.2.31 255.255.255.255 any deny ip 10.2.2.32 255.255.255.224 any

C. deny ip 10.2.2.31 255.255.255.255 any deny ip 10.2.2.32 255.255.255.224 any permit ip 10.2.2.40 255.255.255.248 any permit ip 10.2.2.48 255.255.255.248 any permit ip 10.2.2.0 255.255.255.192 any

D. deny ip 10.2.2.31 255.255.255.255 any permit ip 10.2.2.40 255.255.255.248 any deny ip 10.2.2.32 255.255.255.224 any permit ip 10.2.2.48 255.255.255.248 any permit ip 10.2.2.0 255.255.255.192 any

ANSWER: A

Explanation:

The correct ordered rule set is **deny ip 10.2.2.31 255.255.255.255 any permit ip 10.2.2.40 255.255.255.248 any permit ip 10.2.2.48 255.255.255.248 any deny ip 10.2.2.32 255.255.255.224 any permit ip 10.2.2.0 255.255.255.192 any**. ACL entries are evaluated sequentially, and processing stops at the first matching entry. This ordering therefore uses the more specific entries before the broader range permits.

The host-specific deny for 10.2.2.31 is evaluated first, preserving the required endpoint exclusion before the later permit covering the 10.2.2.0–10.2.2.63 range. The two /29 permit entries then allow 10.2.2.40–10.2.2.47 and 10.2.2.48–10.2.2.55, respectively. They must precede the deny for 10.2.2.32–10.2.2.63 so that these two intended permitted blocks match their permit entries first. The deny entry blocks the remaining addresses in that range, including 10.2.2.32–10.2.2.39 and 10.2.2.56–10.2.2.63. Finally, the /26 permit allows 10.2.2.0–10.2.2.30 because 10.2.2.31 has already matched its explicit deny. Traffic not matching a permit is denied by the ACL's implicit deny behavior. Aruba documents ordered ACL evaluation and first-match processing in its AOS-CX Security Guide and AOS-S Access Security Guide.

QUESTION NO: 8

Examine the following AOS-CX switch configuration:

```
Switch(config-addgroup-ip)# object-group ip address servers
Switch(config-addgroup-ip)# 10.1.0.100
Switch(config-addgroup-ip)# 10.1.1.100
Switch(config-addgroup-ip)# exit
```

Which access control entries would allow web traffic to the web servers 10.1.0.100 and 10.1.1.100?

A. permit tcp servers eq 80

B. permit tcp any 10.1.0.100 0.0.1.0 eq 80

C. permit tcp any 10.1.0.100/10.1.1.100 eq 80

D. permit tcp any 10.1.0.100/255.255.254.255 eq 80

ANSWER: B

Explanation:

`permit tcp any 10.1.0.100 0.0.1.0 eq 80` is correct because it defines a TCP ACL rule with an unrestricted source (`any`) and destination matching both specified web-server addresses. In an AOS-CX ACL, the value following the destination IP address is a wildcard mask, in which zero bits must match and one bits may vary. The wildcard mask `0.0.1.0` permits variation only in the third octet's least-significant bit. Starting with `10.1.0.100`, this matches `10.1.0.100` and `10.1.1.100`, while retaining an exact match for the first, second, and fourth octets. The `eq 80` operator matches TCP destination port 80, the standard HTTP service port, so clients can establish web connections to either server. This is a compact single ACE that precisely covers the two hosts without requiring separate entries. Aruba's AOS-CX documentation describes ACL matching and ACE configuration in its Access Control Lists documentation; protocol port assignments are maintained by the IANA Service Name and Port Number Registry.

QUESTION NO: 9

An administrator is implementing a multi-area OSPF network. The network contains a backbone (area 0) and two other areas (1 and 2) connected to ABRs in the backbone The network has one routing switch connected to a service provider located in area 2 Which network design would minimize the number of routes in the routing switches' link state databases (LSDBs) while still allowing full connectivity?

A. Area 0: NormalArea 1: Totally stubby Area 2: Totally stubby

B. Area 0: NormalArea 1: Totally not-so-stubby Area 2: Totally stubby

C. Area 0: NormalArea 1: Totally stubbyArea 2: Totally not-so-stubby

D. Area 0: Not-so-stubbyArea 1: Totally not-so-stubby Area 2: Totally not-so-stubby

ANSWER: C

Explanation:

Area 0: Normal

Area 1: Totally stubby

Area 2: Totally not-so-stubby

is the appropriate design because the service-provider connection in area 2 requires that area to support an autonomous system boundary router and external route redistribution. A totally not-so-stubby area permits the required external information to enter the area as type 7 LSAs, which the area border router translates for distribution into the normal

backbone. This preserves reachability between the provider, area 2, the backbone, and area 1. It also minimizes the area-2 database by suppressing unnecessary interarea summary LSAs and replacing them with a default route where appropriate.

Area 1 has no stated requirement to redistribute external routes or host an autonomous system boundary router. Configuring it as totally stubby therefore provides the greatest LSDB reduction for that area: the ABR blocks external and interarea detail and injects a default route for destinations outside the area. The backbone must remain a normal OSPF area because it carries interarea routing information and participates in the distribution of routes between the nonbackbone areas. Aruba documents the behavior and restrictions of stub and NSSA area types in its OSPF area types documentation; the NSSA external-LSA behavior is standardized in RFC 3101.

QUESTION NO: 10

What would prevent two OSPF routers from forming an adjacency? (Select two.)

A. Different priorities

B. Different area types

C. Different MTU sizes

D. Different IP addresses

E. Different router IDs

ANSWER: B C

Explanation:

Different area types can prevent an OSPF adjacency because neighbors must agree on the area characteristics advertised in OSPF Hello packets. In particular, routers in a stub area must have matching external-routing capability settings. An area-type mismatch means the peers do not have compatible expectations for handling external LSAs, so they do not establish a usable neighbor relationship for that area. OSPF neighbors also need matching area membership on the shared link; area configuration is therefore a fundamental adjacency requirement.

Different MTU sizes can prevent the routers from reaching a fully established adjacency. After Hellos are exchanged and the routers begin database synchronization, OSPF includes the interface MTU in Database Description packets. If a router receives a Database Description packet that advertises an MTU larger than its local interface MTU, it rejects the packet. The neighbor relationship commonly remains stuck in the ExStart or Exchange state rather than reaching Full. Matching the Layer 3 interface MTU on both ends of an OSPF link is consequently a standard deployment requirement unless an intentional MTU-ignore configuration is used.

See the Cisco OSPF adjacency troubleshooting guide and RFC 2328, OSPF Version 2 for the OSPF Hello and database-exchange behavior.

QUESTION NO: 11

When comparing PIM-DM and PIM-SM, which multicast components are only found with PIM-SM in multicast routing? (Choose two.)

A. IGMP querier

B. Rendezvous point

C. Bootstrap router

D. Shortest path tree

E. Designated router

ANSWER: B C

Explanation:

PIM Sparse Mode depends on a shared-tree rendezvous mechanism for initial multicast distribution. A **Rendezvous point** is the central router to which sources register and toward which receivers' routers initially build (*,G) shared-tree state. This enables PIM-SM to operate efficiently when receivers are sparsely distributed, because multicast traffic is not flooded throughout the routed domain before receiver interest is known. The PIM-SM specification defines the RP's role in source registration, shared-tree construction, and subsequent transition to source-specific shortest-path state where appropriate.

A **Bootstrap router** is a PIM-SM control-plane component used by the Bootstrap Router mechanism to advertise and distribute Rendezvous Point information throughout the PIM domain. Candidate RPs supply RP information, and the elected BSR originates Bootstrap messages so routers can learn the RP set and perform deterministic RP selection for multicast groups. Thus, the Bootstrap Router mechanism supports scalable, distributed RP discovery in PIM-SM deployments. These PIM-SM-specific RP-discovery and shared-tree functions are documented in RFC 4601, Protocol Independent Multicast—Sparse Mode; the dense-mode behavior is separately specified in RFC 3973, Protocol Independent Multicast—Dense Mode.

QUESTION NO: 12

An administrator configures a local mirroring session on an AOS-CX switch to capture traffic for analysis.

Which statements are correct regarding a local mirroring session? (Choose two.)

A. An interface can be configured as a source of mirrored traffic.

B. A VLAN can be configured as a source of mirrored traffic.

C. The mirror destination interface continues normal switching operation while receiving mirrored traffic.

D. A local mirror session requires GRE encapsulation.

E. The packet analyzer must be configured as an LACP peer to receive mirrored traffic.

ANSWER: A B

Explanation:

An interface can be configured as a source of mirrored traffic is correct. The switch can copy traffic observed on a specified source interface to the configured monitoring destination.

A VLAN can be configured as a source of mirrored traffic is also correct. VLAN-based mirroring enables the administrator to observe traffic associated with the selected VLAN rather than limiting the capture to one physical interface. The destination interface is connected to the packet analyzer and should not be used as a normal forwarding interface while it is assigned as a mirror destination. Aruba documents traffic mirroring in the AOS-CX Diagnostics and Supportability Guide.

QUESTION NO: 13

Which protocols are used by NetEdit to interact with third-party devices? (Choose two.)

A. telnet

B. SNMP

C. SSH

D. Restful API

E. CDP

ANSWER: B C

Explanation:

NetEdit uses **SNMP** and **SSH** when working with supported third-party devices. SNMP provides the management and monitoring mechanism needed to retrieve device information, such as inventory and operational details, through standard management information bases. This enables NetEdit to identify and collect relevant state from multivendor infrastructure without requiring a proprietary management interface.

SSH provides the secure remote command-line channel used for device interaction where configuration commands or CLI-based operational commands must be sent to the third-party platform. Using SSH protects management traffic through encryption and supports authenticated administrative access, which aligns with standard network-management security practice. Together, SNMP for information retrieval and SSH for secure CLI access give NetEdit the interoperability mechanisms required to include supported non-Aruba devices in its management workflows. Aruba's NetEdit documentation describes its multivendor capabilities and device-management workflows; see the Aruba NetEdit documentation portal. For the protocol standards underlying these functions, consult RFC 3411, SNMP Management Framework.

QUESTION NO: 14

Examine the VSX-related configuration of the core layer AOS-CX switch:

```
ICX-Tx-Core1(config)# vrf KA
ICX-Tx-Core1(config)# interface 1/1/45
ICX-Tx-Core1(config-if-1/1/45)# no shutdown
ICX-Tx-Core1(config-if-1/1/45)# vrf attach KA
ICX-Tx-Core1(config-if-1/1/45)# ip address 192.168.0.0/31
ICX-Tx-Core1(config-if-1/1/45)# exit
ICX-Tx-Core1(config)# interface lag 256
ICX-Tx-Core1(config-if)# no shutdown
ICX-Tx-Core1(config-if)# no routing
ICX-Tx-Core1(config-if)# vlan trunk native 1
ICX-Tx-Core1(config-if)# vlan trunk allowed all
ICX-Tx-Core1(config-if)# lacp mode active
ICX-Tx-Core1(config-if)# exit
ICX-Tx-Core1(config)# interface 1/1/46-1/1/47
ICX-Tx-Core1(config-if-<1/1/46-1/1/47>)# mtu 9198
ICX-Tx-Core1(config-if-<1/1/46-1/1/47>)# exit
ICX-Tx-Core1(config)# vsx
ICX-Tx-Core1(config-vsx)# inter-switch-link lag 256
ICX-Tx-Core1(config-vsx)# role primary
ICX-Tx-Core1(config-vsx)# vsx-sync vsx-global
ICX-Tx-Core1(config-vsx)# exit
ICX-Tx-Core1(config)# vsx
ICX-Tx-Core1(config-vsx)# keepalive peer 192.168.0.1 source 192.168.0.0 vrf KA
ICX-Tx-Core1(config-vsx)# exit
ICX-Tx-Core1(config)# interface lag 1 multi-chassis
ICX-Tx-Core1(config-lag-if)# no routing
ICX-Tx-Core1(config-lag-if)# vlan access 1
ICX-Tx-Core1(config-lag-if)# lacp mode active
ICX-Tx-Core1(config-lag-if)# exit
ICX-Tx-Core1(config)# int 1/1/1
ICX-Tx-Core1(config-if)# description access1
ICX-Tx-Core1(config-if)# lag 1
ICX-Tx-Core1(config-if)# no shutdown
ICX-Tx-Core1(config-if)# exit
```

A network administrator is troubleshooting a connectivity issue involving the VSX LAG (link aggregation) between the core and access layer switch, during HW replacement of one of the core switches.

Which configuration should the administrator add to the core switch to fix this issue?

A. ICX-Tx-Core1(config)# vsx
ICX-Tx-Core1(config-vsx)# system-mac 02:01:00:00:01:00

B. ICX-Tx-Core1(config)# interface lag 1 multi-chassis ICX-Tx-Core1(config-if-lag-if)# mtu 9198

C. ICX-Tx-Core1(config)# interface 1/1/46-1/1/47
ICX-Tx-Core1(config-if-vlan)# active-gateway ip 10.1.11.1 mac 02:02:00:00:01:00

D. ICX-Tx-Core1(config)# interface 1/1/45
ICX-Tx-Core1(config-if-vlan)# active-gateway ip 192.168.0.0 mac 02:02:00:00:01:00

ANSWER: A

Explanation:

`ICX-Tx-Core1(config)# vsx`
`ICX-Tx-Core1(config-vsx)# system-mac 02:01:00:00:01:00` is the required configuration because a VSX pair needs a stable, shared system MAC address for its multi-chassis link-aggregation identity. The downstream access switch sees the two core members as one logical LACP partner for the VSX LAG. When a core switch is replaced, its factory-assigned base MAC address can differ from that of the failed chassis. If the VSX system MAC is not explicitly retained, the logical LACP system identity presented by the VSX pair can change, which can interrupt LAG negotiation and forwarding while the access switch updates its LACP state.

Configuring the original, administratively defined `system-mac` under the VSX context preserves the pair's logical identity across hardware replacement. The same system MAC must be configured consistently on both VSX peers, using the value planned for that VSX pair. This allows the access-layer switch to continue recognizing the VSX LAG as connected to the same LACP system and restores stable aggregation behavior. Aruba documents the VSX configuration model and commands in the AOS-CX VSX Guide and provides additional VSX deployment guidance in its VSX Configuration Best Practices.

QUESTION NO: 15

An administrator of a company has concerns about upgrading the access layer switches. The users rely heavily on wireless and VoIP telephony. Which is the best recommendation to ensure a short downtime for the users during upgrading the access layer switches?

A. Install the in-service software upgrade (ISSU) feature with clustering enabled

B. Install AOS-CX 6300 or 6400 switches with always-on POE

C. Implement VSF on the AOS-CX access switches

D. Implement VSX on the AOS-CX access switches

ANSWER: B

Explanation:

Install AOS-CX 6300 or 6400 switches with always-on POE is the best recommendation because it specifically addresses the outage most visible to users of access-layer wireless and voice services: loss of power to access points and IP phones during a switch software upgrade or reboot. Always-on PoE is designed to continue supplying Power over Ethernet to connected powered devices while the switch control plane is restarting. This allows compatible APs and phones to remain powered instead of performing a full power-cycle and boot sequence. As a result, the time required for wireless coverage and VoIP endpoints to become usable again is substantially reduced, which directly supports the requirement for short user downtime during maintenance.

For an access layer that serves these endpoint types, preserving endpoint power is particularly valuable because it avoids the additional registration, discovery, and initialization delay that follows a PoE interruption. Aruba identifies always-on PoE as a capability of the CX 6300 and CX 6400 switch families, intended to maintain PoE delivery across switch reboots and improve service continuity during upgrades. See the Aruba CX 6300 Switch Series and Aruba CX 6400 Switch Series product documentation.

QUESTION NO: 16

A network has two AOS-CX switches connected to two different service providers. The administrator is concerned about bandwidth consumption on the service provider links and learned that the service providers were using the company as a transit AS.

Which feature should the administrator implement to prevent this situation?

A. Configure route maps and apply them to BGP

B. Configure the two switches as route reflectors

C. Configure a classifier policy to disable MED

D. Configure bi-directional forwarding detection on both switches

ANSWER: A

Explanation:

Configure route maps and apply them to BGP is correct because route maps provide the policy-control mechanism needed to stop the enterprise AS from advertising one provider's routes to the other provider. The administrator should apply an outbound BGP route map to each eBGP neighbor, typically in conjunction with prefix lists or other matching criteria. The policy should permit advertisement of only prefixes that belong to the company and deny learned transit routes. As a result, each service provider can reach the company's advertised networks, but neither provider receives routes learned from the other provider through the company AS. This enforces the intended non-transit design and prevents traffic between the providers from being attracted across the company's service-provider links, conserving their bandwidth. Route maps can also be used to set BGP attributes where needed, but route filtering is the essential action in this scenario. Aruba AOS-CX documents BGP neighbor route-map configuration, including the ability to apply a route map in the outbound direction, at AOS-CX BGP neighbor route-map command reference. The underlying BGP route-advertisement behavior is defined in RFC 4271.

QUESTION NO: 17

An administrator has configured the following on an AOS-CX switch:

What is the correct ACL rule configuration that would allow traffic from anywhere to reach the web ports on the two specified servers?

A. access-list ip server 10 permit tcp any web-servers group web-ports

B. access-list ip server 10 permit tcp any object-group web-servers object-group web-ports

C. access-list ip server 10 permit tcp any group web-servers group web-ports

D. access-list ip server 10 permit tcp any web-servers web-ports

ANSWER: B

Explanation:

`access-list ip server 10 permit tcp any object-group web-servers object-group web-ports` is correct because AOS-CX ACL entries use the `object-group` keyword when an address object group or port object group is referenced in a rule. Here, `any` specifies every IPv4 source. The first `object-group web-servers` then supplies the destination addresses: the two servers defined in that address group. The second `object-group web-ports` supplies the TCP destination ports defined in the port group, allowing web traffic to those services on either server. The sequence number `10` places this permit entry at the intended position in the ACL named `server`. Object groups provide a maintainable way to reuse multiple hosts and service ports in a single ACL entry; updates to the group membership can be reflected without duplicating equivalent ACL rules for each server and port combination. Aruba AOS-CX ACL syntax documents object-group references in address and service positions, including the explicit `object-group` prefix. See the Aruba AOS-CX ACL command documentation and the AOS-CX object-group documentation.