HP.HPE7-A06.v2026-01-06.q32

Exam Code:	HPE7-A06
Exam Name:	HPE Campus Access Switching Expert Written Exam
Certification Provider:	HP
Free Question Number:	32
Version:	v2026-01-06
# of views:	147
# of Questions views:	320
https://www.freepdfdumps.com/HP.HPE7-A06.v2026-01-06.q32.html	

NEW QUESTION: 1
An administrator is monitoringthird-party WLAN transmitters m HPE Aruba Networking Central and some of them are classified as rogue and suspected rogue How aresuspected rogues classified when using the default classification method for the rule "Suspected AP On-Prem" in HPE Aruba Networking Central?
A. signal level = '-65 dbM- AND WLAN classification ="On-Prem"
B. signal level = "-55 dbM" AND WLAN classification ="Interfering"
C. signal level ="-50 dbM" AND WLAN classification = "Interfering"
D. signal level = "-50 dbM" ANDWLAN classification = "On Wire"
Answer: (SHOW ANSWER)
The question asks how suspected rogue APs are classified using the default classification method for the
"Suspected AP On-Prem" rule in HPE Aruba Networking Central.
* Analysis of Options:
* Option A:Correct. Suspected rogues are classified with a signal level of -65 dBm (indicating proximity) and WLAN classification of "On-Prem" (indicating they are on the premises).
* Option B:Incorrect. A signal level of -55 dBm is too strong, and "Interfering" is not specific to on-premises rogues.
* Option C:Incorrect. A signal level of -50 dBm is even stronger, and "Interfering" is incorrect.
* Option D:Incorrect. "On Wire" classification applies to wired rogue detection, not wireless on-premises APs.
* Why Option A is Correct:In HPE Aruba Networking Central, the "Suspected AP On-Prem" rule identifies rogue APs based on their signal strength and location. A signal level of -65 dBm indicates the AP is close enough to be on the premises, and the "On-Prem" classification confirms it's detected within the managed network's environment. This default rule helps identify potential security threats by flagging unauthorized APs with moderate to strong signals, distinguishing them from interfering or distant APs, as per Aruba's wireless security framework.

* Relevance to Certification Objectives:
* WLAN (9%):Designing and troubleshooting RF attributes and wireless security functions.
* Security (10%):Troubleshooting and identifying rogue APs in customer networks.
* Troubleshooting (10%):Analyzing wireless issues using Aruba Central tools.
References:
HPE Aruba Networking Central User Guide: Rogue AP Detection and Classification.
HPE7-A06Study Guide: Covers wireless security and rogue AP management.
HPE Aruba Networking Technical Documentation: Wireless Security and Rogue Detection Best Practices.

NEW QUESTION: 2
Exhibit.
The customer has VSX clusters intwo locations interconnected over an MC-LAG interface.
If active-gateway configuration uses the same virtual IP address and vMAC on each of the VSX nodes, what must you take into consideration0
A. Transit traffic will Increase over the VSX interconnect MC-LAG.
B. Each ARP request will result in four responses.
C. The configuration would end up in an async setup.
D. Outbound traffic will be load-balancedover all VSX members for each session.
Answer: C (LEAVE A REPLY)
The scenario describes two separate VSX clusters interconnected via MC-LAG, where both clusters are configured to use theexact samevirtual IP address and virtual MAC address for their respective Active Gateway SVIs.
* Active Gateway Scope & Conflict:Active Gateway provides a highly available default gatewaywithina single VSX cluster (L2 domain). The vIP/vMAC combination should be unique within its L2 broadcast domain.
* Interconnecting Clusters with Same vIP/vMAC:When two VSX clusters using the identical Active Gateway vIP/vMAC are interconnected at Layer 2 (even via MC-LAG), this creates a situation where the same active L2 (vMAC) and L3 (vIP) address exists in multiple places within the extended broadcast domain.
* Consequences:This leads to MAC address conflicts and L3 ambiguity. ARP resolution becomes unreliable, potentially causing ARP tables to flap on connected devices. Traffic forwarding becomes unpredictable, as packets destined for the vIP/vMAC might be delivered to the "wrong" cluster. This unstable and unpredictable state is sometimes referred to as an asymmetric or "async" setup.
* Analysis of Options:
* A: ISL traffic might change, but it's a symptom, not the root problem.
* B: Multiple ARP replies would occur, contributing to the confusion.
* C: The configuration results in an "async setup," accurately describing the unstable state caused by duplicate active L2/L3 addresses across the interconnected L2 domain.

* D: Load-balancing happens within a cluster; this setup causes conflict, not predictable load balancing across clusters.
* Conclusion:Reusing the same Active Gateway vIP and vMAC across interconnected VSX clusters is not a valid design and leads to an unstable, asymmetric ("async") environment due to address duplication within the extended L2 domain. Option C best describes this problematic outcome.
References:Aruba VSX Design and Best Practices Guides (Active Gateway uniqueness, Interconnecting VSX clusters). This relates to "Network Resiliency and virtualization" (8%), "Routing" (16%), and
"Troubleshooting" (10%) objectives.

NEW QUESTION: 3
Exhibit.
Acme Corp has VM workload running from ToR-1. and has noticed performancedegradation They suspect ToK-1 uplinks are periodically overutilized. List valid reasons whyToR-1 uplinks 3re being overutilized based on the diagram. (Select two.)
A. Core-2 has been incorrectly configured as the root bridge
B. The VLAN to instance mapping is not the same on all switches.
C. Core-1 and Core-2 are not running the same firmware
D. Thecustomer has used the default MSTP region configuration
E. ToR-1uplinks and downlinks are both running spanning-tree port-type admin-network.
Answer: (SHOW ANSWER)
The question involves Acme Corp experiencing performance degradation due to overutilized uplinks from ToR-1 to Core-1 and Core-2, with a diagram (not provided) indicating a potential MSTP (Multiple Spanning Tree Protocol) issue. The task is to identify valid reasons for uplink overutilization.
* Analysis of Options:
* Option A:Incorrect. Incorrect root bridge configuration (e.g., Core-2 as root) may cause suboptimal paths but is not directly linked to uplink overutilization without further context.
* Option B:Correct. Inconsistent VLAN-to-instance mappings across switches can cause MSTP to block unexpected ports, funneling traffic through fewer uplinks and causing overutilization.
* Option C:Incorrect. Firmware mismatches may cause compatibility issues but are unlikely to directly cause uplink overutilization.
* Option D:Correct. Using the default MSTP region configuration (e.g., default region name and revision) across switches can lead to all switches forming a single MSTP region, potentially causing suboptimal topology and uplink overuse.
* Option E:Incorrect. Running MSTP with admin-network port-type on uplinks and downlinks is not a standard cause of overutilization; it's a specific port role.
* Why B and D are Correct:MSTP relies on consistent region configurations (region name, revision number, VLAN-to-instance mappings) to create efficient topologies. If VLAN-to-instance mappings differ (Option B), switches treat each other as separate regions, leading to blocked

ports and traffic concentration on fewer uplinks, causing overutilization. Similarly, using the default MSTP region configuration (Option D) without customizing the region name or revision can result in all switches forming a single region with suboptimal spanning tree instances, potentially overloading specific uplinks. Both issues disrupt MSTP's ability to balance traffic across redundant paths, aligning with HPE Aruba Networking's MSTP troubleshooting scenarios.
* Relevance to Certification Objectives:
* Network Resiliency and Virtualization (8%):Troubleshooting MSTP for redundancy and fault tolerance.
* Switching (19%):Diagnosing Layer 2 issues, including MSTP misconfigurations.
* Performance Optimization (6%):Remediating uplink utilization issues.
References:
HPE Aruba Networking AOS-CX Configuration Guide: MSTP Configuration, detailing region and VLAN mapping.
HPE7-A06Study Guide: Covers MSTP troubleshooting and optimization.
HPE Aruba Networking Technical Documentation: MSTP Best Practices and Troubleshooting.

NEW QUESTION: 4
The customer is experiencing periodic uplink congestion between campus-1's AGG-1 and core. This has boon negativelyaffecting voice communications.The VOIP phones edge mark their packets with DSCP EF. The uplink from AGG-1 to core is LAG1.
The customer has already configured the following class and policy on AGG-1:
Based on this policy, which scrip), when deployed on AGG-1. will improve the reliable forwarding of voice trafficbetween AGG-1 and its uplink to the core?
A.
B.
C.
D.
Answer: (SHOW ANSWER)
The problem describes uplink congestion affecting VoIP traffic (marked with DSCP EF, value 46) on AGG-
1's LAG1 uplink. The existing configuration classifies this traffic into voip_class and applies voip_policy inbound, setting local-priority 6. To improve reliable forwarding during congestion, VoIP traffic needs strict priority queuing on the egress interface (LAG1).
* Analysis of Options:
* Option A applies a QoS schedule profile globally but doesn't modify the policy's local-priority or apply the schedule profile specifically to the congested LAG.
* Option B modifies voip_policy to set local-priority 7 (mapping DSCP 46 traffic to queue 7) and applies the 8qDwrStrict schedule profile to the egress interface lag 1. In the 8qDwrStrict profile, queue 7 is configured for strict priority, ensuring voice traffic gets precedence over lower-priority traffic during congestion. This aligns with best practices for QoS for VoIP.

* Option C also sets local-priority 7 and applies the schedule profile to lag 1, but the profile itself configures queue 7 with DWRR (Deficit Weighted Round Robin) instead of strict priority, which is less suitable for delay-sensitive voice traffic.
* Option D applies a schedule profile globally and uses DWRR for queue 7.
* Conclusion:Option B is the correct solution because it maps the DSCP EF traffic to the highest local priority (7) and applies a QoS schedule profile to the specific congested uplink (lag 1) that treats queue
7 with strict priority. This ensures voice traffic is prioritized reliably.
References:AOS-CX QoS Guide (specifically sections on Classification, Queuing, Scheduling Profiles, Strict Priority vs. DWRR, applying policies to interfaces/LAGs), DSCP to Queue mapping concepts. This relates to the "Performance Optimization" (6%) and "Connectivity" (9%) objectives.

NEW QUESTION: 5
Review the diagram and existing configuration of RouterA above. Which configuration changes are necessary to permit load balancing between RouterA and RouterB? (Selecttwo) Exhibit.
A.
B.
C.
D.
E.
Answer: (SHOW ANSWER)
Analyze Topology and Existing Configuration:
* RouterA (AS 64500) peers with RouterB (AS 64512) using eBGP.
* Peering is configured between loopback interfaces (RouterA Lo0 10.3.0.3 to RouterB Lo0 10.255.0.12).
* Two parallel physical links connect the routers (10.255.102.0/30 and 10.255.102.4/30).
* RouterA has two static routes pointing to RouterB's loopback (10.255.0.12/32), one via each physical link's next hop (10.255.102.1 and 10.255.102.5). This provides reachability to the BGP peer address over both paths.
* RouterA's BGP config activates the neighbor 10.255.0.12 for IPv4 unicast but is missing key commands for stable loopback peering and load balancing.
Goal:Permit load balancing for traffic exchanged via BGP between RouterA and RouterB. This requires BGP ECMP (Equal Cost Multi-Path).
Requirements for eBGP ECMP over Loopbacks:
* Stable Peering:Peering must use loopback addresses. This requires:
* update-source loopback <id>: To source BGP TCP packets from the loopback IP.
* ebgp-multihop <ttl>: Because loopbacks are not directly connected (TTL > 1 needed).
* ECMP Enabled:BGP must be configured to allow multiple paths in the routing table. This requires:
* maximum-paths <n> (or maximum-paths ebgp <n>): To allow more than the default 1 path.

* Equal Paths:BGP must see multiple paths to thesameprefix learnedfrom RouterBthat are considered equal based on BGP path selection attributes (Weight, Local_Pref, AS_Path, Origin, MED, etc.). Since routes are learned from the same neighbor IP (RouterB's loopback), these attributes will likely be identical for routes learned via this peering. RouterA already has equal static routestothe BGP next hop (10.255.0.12).

NEW QUESTION: 6
Place the recommended troubleshooting steps in order.
Answer:
Explanation:
The correct order is:
* identify
* analyze
* hypothesize
* validate
* implement
* verify
This question requires arranging standard troubleshooting steps into a logical sequence. A systematic approach is crucial for effective network troubleshooting.
* identify:The first step is always to clearly identify and define the problem. What are the symptoms?
Who is affected? What is the scope? When did it start? Understanding the problem precisely is essential before proceeding.
* analyze:Once the problem is identified, gather relevant data and analyze the situation. This involves checking logs, looking at configurations, examining network topology diagrams, checking status commands, and potentially capturing packets. This analysis helps build context around the identified issue.
* hypothesize:Based on the identification and analysis, form a hypothesis (or multiple hypotheses) about the probable cause of the problem. This involves using technical knowledge and experience to theorize what might be wrong.
* validate:Test the hypothesis to determine if it's correct. This step involves performing specific tests or checks designed to confirm or refute the theory. For example, if the hypothesis is a bad cable, test the cable. If it's a routing issue, check the routing table and perform trace routes. This step validates the cause before implementing a fix.
* implement:Once the cause has been validated, implement the solution. This could involve replacing hardware, correcting configuration, clearing states, etc.
* verify:After implementing the solution, verify that the original problem is resolved. It's also critical to check that the fix hasn't introduced any new issues. Monitor the system to ensure stability.
References:Standard Network Troubleshooting Methodologies (e.g., CompTIA Network+, Cisco troubleshooting models), ITIL Problem Management processes. This directly relates to the

"Troubleshooting" (10%) objective, which emphasizes performing advanced troubleshooting and remediation.

NEW QUESTION: 7
Which is a best practice for configuringGBP?
A. Configure GBP classes to have a destination role that is different from theassociated user role.
B. Use static user roles (SUR) to configure GBP
C. Configure GBP classes to have a destination role that is the same as the associated user rote.
D. Use downloadable user roles (DUR) to configure GBP.
Answer: (SHOW ANSWER)
The question asks for a best practice when configuring Group-Based Policy (GBP). GBP simplifies policy management by assigning users/devices to roles and defining policies between these roles, often leveraging dynamic assignment from an authentication server.
* GBP Concepts:Policies are typically defined based on source and destination roles. Roles can be assigned statically on the switch or dynamically via an authentication server like ClearPass.
* Analysis of Options:
* A & C: Policies define interactionsbetweenroles (source role to destination role). These roles can be the same (intra-role policy) or different (inter-role policy). Neither option represents a singular
"best practice" for all configurations.
* B: Using Static User Roles (SUR) is possible but less flexible and scalable than dynamic assignment for large or complex environments.
* D: Using Downloadable User Roles (DUR) is generally considered a best practice. DUR allows roles and associated policies (including GBP attributes like GPID) to be centrally defined on an authentication server (e.g., ClearPass) and dynamically assigned to users/devices uponsuccessful authentication. This provides scalability, consistency, and easier management.
* Conclusion:Leveraging Downloadable User Roles (DUR) from a central authentication server like ClearPass is a best practice for implementing scalable and manageable Group-Based Policies.
References:Aruba Dynamic Segmentation concepts, Group-Based Policy (GBP) documentation, Aruba ClearPass integration guides. This relates to "Security" (10%) and "Authentication/Authorization" (9%) objectives.

NEW QUESTION: 8
Which issue may becausing the new door locks on the APs to notwork?
A. AT power to the AP is too much.
B. BT power to the AP is too much.
C. AF power to the AP is not enough.
D. AT power to the AP is notenough.
Answer: (SHOW ANSWER)

New PoE-powered door locks, connected via the PoE passthrough port on Aruba APs, are not working. We need to find the likely cause related to PoE power.
* PoE Passthrough:An AP feature where the AP, powered by PoE from a switch, provides PoE power out to another device connected to one of its Ethernet ports.
* Power Budget:The AP must receive enough power from the switch via its PoE input (e.g., 802.3af,
802.3at, 802.3bt) to power itselfandmeet the power demand of the downstream device (the door lock).
* PoE Standards Power (Approx. Available to Device):
* 802.3af (PoE): ~13 Watts
* 802.3at (PoE+): ~25.5 Watts
* 802.3bt (PoE++): 51W (Type 3) or 71W (Type 4)
* Analysis:Modern APs (especially Wi-Fi 6/6E) can consume significant power themselves (>15W or
>25W under load). Standard 802.3af PoE (supplying only ~13W) is often insufficient to power both a modern AP and a downstream PoE device like a door lock. The AP will power up, but won't enable PoE output if its input power budget is insufficient.
* Analysis of Options:
* A, B: Too much power (AT/BT) isn't the issue; devices only draw what they need.
* C: AF power (~13W) received by the AP is very likelynot enoughto power both the AP and the door lock.
* D: AT power (~25.5W)mightbe insufficient if the combined load of the AP and lock exceeds this, but AF being insufficient (C) is a more common limitation.
* Conclusion:Insufficient input power to the AP is the most common reason for PoE passthrough failure.
802.3af (PoE) power is often inadequate.
References:IEEE 802.3 PoE standards (af/at/bt), Aruba Access Point datasheets (PoE requirements, passthrough capabilities/budgets). This relates to "WLAN" (9%) and "Connectivity" (9%) objectives.

NEW QUESTION: 9
The clientwouldlike to automate the process of troubleshooting issues to have better visibility. Which solution would you recommend for your client?
A. HPE Aruba Networking F3bric Compose
B. HPE Aruba Networking Switch Multi-Edit Software
C. Automate processes with scripting like Python.
D. AIOps integrated into HPE Aruba Networking Central
Answer: (SHOW ANSWER)
The client wants to automate troubleshooting processes and gain better visibility into their network. We need to identify the recommended Aruba solution.
* Analysis of Options:

* A. HPE Aruba Networking Fabric Composer: A tool primarily for data center fabric provisioning and management, not general campus troubleshooting automation.
* B. HPE Aruba Networking Switch Multi-Edit Software: Likely refers to configuration management features (e.g., in Central or NetEdit) for applying changes to multiple switches, not primarily focused on automated troubleshooting or visibility.
* C. Automate processes with scripting like Python: AOS-CX supports on-box scripting (NAE) and REST APIs, enabling custom automation for monitoring and troubleshooting. While powerful, it requires development effort.
* D. AIOps integrated into HPE Aruba Networking Central: Aruba Central's AIOps capabilities are specifically designed to enhance visibility and automate aspects of troubleshooting. It uses AI /ML to analyze network data, detect anomalies, provide insights into potential issues, correlate events, and offer prescriptive recommendations, directly addressing the client's need for better visibility and automated assistance with troubleshooting.
* Conclusion:While custom scripting (C) allows automation, Aruba Central AIOps (D) is the platform- integrated solution specifically marketed and designed by HPE Aruba Networking to provide enhanced visibility and automated insights fortroubleshooting campus networks. It is the most direct and recommended solution among the options for achieving these goals within the Aruba ecosystem.
References:Aruba Central documentation (AIOps features), AOS-CX NAE and REST API documentation.
This relates to "Troubleshooting" (10%) and "Performance Optimization" (6%) objectives.

NEW QUESTION: 10
Exhibit.
A.
B.
C.
D.
Answer: C (LEAVE A REPLY)
The question involves configuring an OSPF virtual link to extend area 0 across a non-backbone area, based on an exhibit (not provided) and four configuration options (A to D). Since the exhibit is unavailable, I will assume a typical scenario where a virtual link is needed to connect two area 0 segments through a transit area (e.g., area 1).
* Analysis of Options (Assumed Context):A virtual link is configured using the area <transit-area> virtual-link <router-id> command in the OSPF process. The correct option likely includes:
* Option A:Incorrect syntax or incorrect router ID/area for the virtual link.
* Option B:Incorrect configuration, possibly missing the virtual link or using wrong parameters.
* Option C:Correct. Likely includes the proper command, e.g., area 1 virtual-link 2.2.2.2, where area 1 is the transit area and 2.2.2.2 is the router ID of the remote ABR.
* Option D:Incorrect, possibly configuring an unnecessary or incorrect virtual link.

* Why Option C is Correct:OSPF requires all areas to connect to the backbone area (area 0). If two area
0 segments are separated by a non-backbone area (e.g., area 1), a virtual link is configured between the Area Border Routers (ABRs) to logically extend area 0 through the transit area. The command area
<transit-area> virtual-link <remote-router-id> is used, specifying the transit area and the router ID of the remote ABR. Option C is assumed to provide the correct syntax and parameters based on standard OSPF virtual link configurations, ensuring area 0 connectivity and proper route advertisement.
* Relevance to Certification Objectives:
* Routing (16%):Designing and troubleshooting OSPF topologies, including virtual links.
* Troubleshooting (10%):Resolving OSPF area connectivity issues.
References:
HPE Aruba Networking AOS-CX Configuration Guide: OSPF Configuration, detailing virtual link setup.
HPE7-A06Study Guide: Covers OSPF advanced configurations like virtual links.
RFC 2328: OSPF Version 2, explaining virtual link functionality.

NEW QUESTION: 11
Youare configuring an HPE Aruba NetworkingGateway Ouster with AOS-10. What is true about 802.1 X functionality incombination with gateways? (Select two.)
A. Users onL3-oonnoctod gateways need to perform a full authentication after re-association on theAP.
B. The UDG remains fixed on L2-connected gateways but not on 1.3-connected gateways.
C. Regardless of using gateways, the CoA message is always sent to the APs.
D. The gateways areused as a RADIUS proxy, while the AP is theauthenticator.
E. The gateways actasRADIUS Proxy only in Tunnel and Bridged Mode.
Answer: (SHOW ANSWER)
This question asks about 802.1X functionality in an AOS-10 environment involving Gateway Clusters.
* AOS-10 Gateway/802.1X Architecture:
* Authenticator:The Access Point (AP) typically acts as the 802.1X authenticator, handling EAPoL frames with the client.
* RADIUS Proxy:The Gateway Cluster (specifically the cluster leader or UDG anchor) often acts as a RADIUS proxy, forwarding RADIUS messages between the APs and the central RADIUS server (e.g., ClearPass). This simplifies RADIUS configuration as the server only needs to know about the gateway cluster.
* CoA:Change of Authorization messages from the RADIUS server are typically sent to the device acting as the RADIUS client, which is the Gateway Cluster when operating in proxy mode.
* Mobility (L2 vs L3):Roaming behavior and User Designated Gateway (UDG) assignment can differ based on whether clients maintain their IP address (L2 mobility) or potentially require new

IP information (L3 mobility). L2-connected gateway deployments generally allow for more seamless UDG persistence compared to L3-connected deployments where the client might roam across subnet boundaries managed by different gateways.
* Re-authentication:Seamless roaming mechanisms aim to minimize full re-authentications during roaming events.
* Analysis of Options:
* A: Full re-authentication after re-association on L3-connected gateways might occur in some scenarios but contradicts the goal of seamless roaming.
* B: States the UDG remains fixed on L2-connected but not on L3-connected gateways. This aligns with the architectural differences in handling mobility across L2 vs L3 boundaries within a cluster.
* C: Incorrect. CoA is generally sent to the RADIUS client/proxy (the Gateway Cluster), not always directly to the APs.
* D: Correct. Gateways commonly act as a RADIUS proxy, while the AP remains the authenticator handling EAPoL with the client.
* E: Incorrect. The RADIUS proxy function is not limited to only Tunnel and Bridged modes.
* Conclusion:Options B and D accurately describe common characteristics of 802.1X operation within an AOS-10 Gateway Cluster architecture.
References:Aruba AOS-10 documentation (Gateway Clusters, User-Based Tunneling, 802.1X/RADIUS interaction, L2/L3 Mobility). This relates to "Authentication/Authorization" (9%), "Connectivity" (9%), and
"WLAN" (9%) objectives.

NEW QUESTION: 12
You are configuring VSX active gateway on CX 8360 campus aggregation switches when the switch prompt returns the following error: "No more than IB VMACs can be configured." What should be done to address this issue?
A. Limit the number of SVIs with active-gatewayto 16.
B. Change the switch profile to "Leal" to increase the number of supported vMACs.
C. Change the aggregation switch to a higher-end model, such as a CX 8400.
D. As MAC addresses are link-local, use the same VMAC across SVIs.
Answer: D (LEAVE A REPLY)
The error "No more than 16 vMACs can be configured" occurs when trying to configure active-gateway on multiple SVIs on a CX 8360 VSX pair. This indicates a platform limit on the number of unique virtual MAC addresses has been reached.
* Active Gateway vMACs:Each SVI configured with Active Gateway requires a virtual MAC address (vMAC). While AOS-CX can auto-generate these, doing so consumes entries from a limited hardware pool (e.g., 16 on this platform/version).
* Best Practice & Solution:The recommended best practice to conserve these limited vMAC resources is to manually specify andreuse the same virtual MAC addressacross all SVIs configured with Active Gateway on that specific VSX pair. Since MAC addresses are Layer 2

local, using the same vMAC on different SVIs (different L3 subnets) does not cause conflicts within the VSX pair's operation.
* Analysis of Options:
* A: Limiting the number of SVIs using Active Gateway is a workaround, not a solution.
* B: Changing switch profiles doesn't typically alter hardware vMAC limits.
* C: Changing to a higher-end switch model might increase limits but is not the first or standard solution.
* D: Reusing the same VMAC across SVIs (active-gateway ip <vip> mac <SAME_VMAC>) avoids consuming a new vMAC entry for each SVI, thus staying within the platform limit. This is the standard, recommended solution.
* Conclusion:The correct approach to address the vMAC limit error is to explicitly configure the same virtual MAC address for all SVIs using the Active Gateway feature on the VSX pair.
References:AOS-CX VSX Guide (Active Gateway Configuration, Best Practices, vMAC considerations).
This relates to "Network Resiliency and virtualization" (8%) and "Routing" (16%).

NEW QUESTION: 13
A client would like to usetheHPE Aruba Networking Switch MultiEdit Software function in HPE Aruba Networking Central. Which option is available?
A. Usetemplates and apply them to selected switches.
B. Apply a configuration to an interface range for selected switches.
C. Run the same NAE scripts for selected switches.
D. Use CLI scripts and apply them to selected switches.
Answer: (SHOW ANSWER)
The question involves a client wanting to use the HPE Aruba Networking Switch Multi-Edit Software function in HPE Aruba Networking Central to manage multiple switches. The task is to identify the available option.
* Analysis of Options:
* Option A (Use templates and apply them to selected switches):Incorrect. Templates are used for configuration management in Central but are not part of the Multi-Edit Software function.
* Option B (Apply a configuration to an interface range for selected switches):Incorrect. Multi- Edit focuses on CLI scripting, not specifically interface range configurations.
* Option C (Run the same NAE scripts for selected switches):Incorrect. Network Analytics Engine (NAE) scripts are for monitoring, not configuration via Multi-Edit.
* Option D:Correct. Multi-Edit Software in Central allows administrators to apply CLI scripts to multiple selected switches for configuration changes.
* Why Option D is Correct:HPE Aruba Networking Central's Multi-Edit Software feature enables administrators to create and apply CLI scripts to multiple AOS-CX switches simultaneously, streamlining configuration tasks. This is particularly useful for bulk changes, such as VLAN configurations or policy updates, across selected switches. The feature supports direct CLI input

or script uploads, ensuring consistent application of commands, as per HPE Aruba Networking's management tools. This aligns with the client's need for efficient multi-switch management.
* Relevance to Certification Objectives:
* Connectivity (9%):Developing configurations for multiple devices based on customer requirements.
* Troubleshooting (10%):Applying consistent configurations to resolve network issues.
* Network Stack (4%):Analyzing solutions for network management automation.
References:
HPE Aruba Networking Central User Guide: Multi-Edit Software Feature, detailing CLI script application.
HPE7-A06Study Guide: Covers network management tools in Central.
HPE Aruba Networking Technical Documentation: Multi-Edit Software Best Practices.

NEW QUESTION: 14
A customer has configured eBGP peering using local AS 65000 with two routers from a CX 6300 VSF stack with thefollowing switch ports:
[ports connecting to router-1 10.10.10.2]
The LAGs are connected lo third-party L2 switches, which are used as a transit network for the remote eBGP routers. To optimise the possible BGP peering issues. The AOS-CX switch Is configured with theglobal settings:
What needs to be done on the AOS_CX switch to enable the bidirectional forwarding with the eBGP peers?
A. Option A
B. Option B
C. Option C
D. Option D
Answer: B (LEAVE A REPLY)
The goal is to enable Bidirectional Forwarding Detection (BFD) for eBGP neighbors 10.10.10.2 and
10.10.20.2 on the AOS-CX VSF stack (AS 65000). Global BFD settings are already configured. We need the specific commands to link BFD state to the BGP neighbor relationship.
* BFD for BGP Configuration:Requires enabling the fall-over bfd parameter for the specific neighbor within the router bgp <asn> configuration hierarchy.
* Analyzing the Options (New Image):
* Option 1 (Top):
router bgp 65000
address-family ipv4 unicast
neighbor 10.10.10.2 fall-over bfd
neighbor 10.10.20.2 fall-over bfd
This enables BFD specifically within the ipv4 unicast address family context for both neighbors. This is a valid configuration location.

* Option 2 (Second):
router bgp 65000
neighbor 10.10.10.2 fall-over bfd
neighbor 10.10.20.2 fall-over bfd
This enables BFD directly under the main neighbor <ip> configuration lines within router bgp 65000. This typically applies BFD to all address families configured for that neighbor relationship (including IPv4 unicast). This is also a valid and common configuration location.
* Option 3 (Third):
int 1/1/1-1/1/2, 2/1/1-2/1/2
fall-over-bfd
Incorrect. Applies BFD configuration under an interface range context, which is not how BFD is linked to BGP sessions.
* Option 4 (Bottom):
interface lag1-2
fall-over bfd
Incorrect. Applies BFD configuration under an interface LAG range context, which is not how BFD is linked to BGP sessions.
* Comparing Valid Options (1 vs 2):Both Option 1 and Option 2 correctly use the fall-over bfd command under router bgp. Option 1 provides per-address-family granularity, while Option 2 applies it to the neighbor generally. Without a specific requirement to enable BFDonlyfor IPv4, applying it at the neighbor level (Option 2) is often simpler and sufficient. Both achieve the goal for the required IPv4 peering. In many documentation examples, the configuration is shown at the neighbor level unless per- AF control is explicitly needed.
* Conclusion:Both Option 1 and Option 2 show valid configuration methods. Option 2 is arguably slightly more common/general when BFD is desired for the overall neighbor relationship.
References:AOS-CX BFD Guide, AOS-CX BGP Guide (neighbor commands, fall-over bfd option).
This relates to "Routing" (16%) and "Network Resiliency and virtualization" (8%) objectives.

NEW QUESTION: 15
Exhibit.
In the given example AGG-SW1 and AGG-SW2 use CX 8325 in VSX and Edge-1 withCX 6200F. You want toavcwl sub-optimal path.ng and ISL traffic for the VSX and upstream routers R1 and R2.
What is the HPE Aruba Networkingrecommended solution for me SVIs on the VSX switches connected to R1 and R2?
A. Configure the VSX SVI using the active-forwarding.
B. Configure the VSX SVI using the uncast IP.
C. Configure the VSX SVI using the VRRP virtual-ip.
D. Configure the VSX SVI using the active-gateway.
Answer: A (LEAVE A REPLY)

The scenario involves a VSX pair (AGG-SW1/SW2) connected upstream to routers R1/R2. The goal is to configure the SVIs on the VSX switches facing these upstream routers optimally to avoid suboptimal L3 paths and unnecessary traffic over the VSX Inter-Switch Link (ISL).
* VSX L3 Interface Options:
* Active Gateway:Primarily designed for downstream SVIs to provide a redundant default gateway to clients/access switches. Not typically used for upstream routed interfaces.
* Active Forwarding:Specifically designed for upstream routed interfaces (physical or SVIs) on a VSX pair. It allows both VSX members to actively route traffic arriving on that interface locally, without needing to forward L3 traffic across the ISL. This ensures optimal routing and utilizes both members effectively.
* Unicast IP (Standard IP):Without specific VSX features, standard routing applies. This could lead to suboptimal paths if, for example, return traffic prefers one VSX switch, but the optimal path requires crossing the ISL.
* VRRP:Can be run between VSX members but adds complexity and is generally superseded by Active Gateway (downstream) or Active Forwarding (upstream) in VSX designs.
* Analysis of Options:
* A. Configure active-forwarding: This enables local L3 forwarding on both VSX members for the upstream SVI, preventing unnecessary ISL traversal for routed traffic. This is the recommended best practice.
* B. Configure unicast IP: Standard configuration, potentially leading to suboptimal paths/ISL usage.
* C. Configure VRRP virtual-ip: Not the recommended approach for upstream links in VSX.
* D. Configure active-gateway: Incorrect, Active Gateway is for downstream SVIs.
* Conclusion:Using active-forwarding on the SVIs facing the upstream routers (R1/R2) is the HPE Aruba Networking recommended solution to ensure optimal routing and minimize L3 traffic across the ISL.
References:AOS-CX VSX Guide (Active Forwarding feature description and use cases). This relates to
"Network Resiliency and virtualization" (8%) and "Routing" (16%) objectives.

NEW QUESTION: 16
Ever since a recent firewall change at your WAN/Internet edge, the 8GP state in your VSX pair has not returned to Established. What should you check to restore BGP functionality at the site?
A. Restart the routing service so thatBGP auto-discovers its neighbors.
B. Confirm that appropriate TCP ports are still allowed.
C. Restart NAT service for the BGP interface.
D. Confirm that BGP Peer AS has not changed.
Answer: B (LEAVE A REPLY)
The BGP state on a VSX pair is stuck (not 'Established') after a recent firewall change at the WAN/Internet edge, where the BGP peering likely occurs.

* BGP and Firewalls:BGP establishes sessions usingTCP port 179. Firewalls located between BGP peers must explicitly permit TCP port 179 traffic bidirectionally for the peering to establish and maintain. Firewall changes are a frequent cause of broken BGP sessions.
* Troubleshooting Steps After Firewall Change:The most logical first step is to verify that the firewall change did not inadvertently block TCP port 179 between the configured BGP neighbor IP addresses.
* Analysis of Options:
* A: Restarting routing service is disruptive and not the first step.
* B: Confirming that appropriate TCP ports (specifically 179) are still allowed through the firewall directly addresses the most probable cause related to the firewall change event.
* C: Restarting NAT service is likely irrelevant unless NAT is incorrectly configured for BGP peers.
* D: Confirming the peer AS is a basic configuration check but less likely related to thefirewall changeevent than port blocking.
* Conclusion:Given the problem occurred immediately following a firewall change, verifying that the firewall still permits TCP port 179 between the BGP peers is the most direct and likely troubleshooting step.
References:BGP protocol specifications (RFC 4271), Firewall management principles, Network troubleshooting methodology. This relates to "Routing" (16%), "Security" (10%), and "Troubleshooting" (10%) objectives.

NEW QUESTION: 17
Network administrators are reporting that switches arc taking a very long time to execute commands. Based on the configuration below, what is the mostlikelycause ofthe issue?
A. Too many administrators ace logged in.
B. Authentication fail-through is enabled.
C. The primary TACACS+ server is unreachable.
D. A Denial of Service attack on the data plane.
Answer: C (LEAVE A REPLY)
The issue is that switches are taking a very long time to execute commands. The question points towards the AAA configuration as the context (though the specific configuration is missing).
* AAA and Command Latency:When AAA servers (like TACACS+ or RADIUS) are used for authentication, authorization, or accounting, the switch must communicate with these servers.

* Impact of Unreachable Servers:If the primary AAA server configured on the switch becomes unreachable (due to network issues, server downtime, or firewall rules), the switch will attempt to connect, wait for a configured timeout period (often several seconds), and only then potentially try a secondary server or fall back to local credentials (if configured). This connection attempt and timeout period occurring before command execution (if command authorization is enabled) or during login introduces significant delays.
* Analysis of Options:
* A: Too many administrators might strain resources, but AAA timeouts cause more predictable, long delays per action.
* B: Authentication fail-through only comes into playafterthe primary server times out. The timeout itself causes the delay.
* C: An unreachable primary TACACS+ (or RADIUS) server is a classic cause of slow logins and command execution delays due to connection timeouts.
* D: A DoS attack might cause general slowness but isn't specifically linked to the AAA configuration context provided.
* Conclusion:The most likely cause, given the context of AAA configuration and the symptom of slow command execution, is that the primary configured AAA server (like TACACS+) is unreachable, causing the switch to wait for timeouts.
References:AOS-CX Security Guide (AAA, TACACS+, RADIUS), general network troubleshooting for AAA latency. This relates to "Authentication/Authorization" (9%) and "Troubleshooting" (10%) objectives.

NEW QUESTION: 18
A senior engineer from the network operations team has reported an intermittent problem where some PoE- powered devices are randomly losing power. During your investigation, you found that port 1 of the Acc-1 switch is currently presenting the behavior shown in the CLI output for the Acc-1.
What is a probablecausa lot poor 1/1/1 is denying PoE?
A. This switch does no! support PoE class 4.
B. switch PoE powerbudget exceeded
C. PoE port priority sot to low
D. PoE was manually disabled for port 1/1/1.
Answer: B (LEAVE A REPLY)
The question involves intermittent PoE-powered device power loss on port 1/1/1 of an AOS-CX switch (Acc-
1), with CLI output (not provided) indicating a PoE issue. The task is to identify a probable cause.
* Analysis of Options:
* Option A:Incorrect. AOS-CX switches typically support PoE Class 4 (802.3at, 30W), sufficient for most devices.
* Option B:Correct. If the switch's PoE power budget is exceeded, it may deny power to port 1/1 /1, causing intermittent device failures.

* Option C:Incorrect. Low PoE port priority may deprioritize the port but is less likely to cause complete power loss compared to budget issues.
* Option D:Incorrect. Manual disabling of PoE would cause consistent power loss, not intermittent issues.
* Why Option B is Correct:AOS-CX switches have a finite PoE power budget (e.g., 370W or 740W, depending on the model and power supply). If the total power demand from connected devices exceeds this budget, the switch denies power to some ports, often intermittently as devices cycle or negotiate power. For port 1/1/1, this could manifest as random power loss for the connected device. The CLI output likely shows a "power denied" status (e.g., via show power-over-ethernet brief). Checking the PoE budget (show power-over-ethernet) and upgrading power supplies or prioritizing critical ports resolves the issue, aligning with HPE Aruba Networking's PoE troubleshooting guidelines.
* Relevance to Certification Objectives:
* Connectivity (9%):Troubleshooting PoE deployment issues.
* Troubleshooting (10%):Diagnosing power-related issues in campus networks.
* Switching (19%):Implementing PoE configurations for Layer 2 devices.
References:
HPE Aruba Networking AOS-CX Configuration Guide: PoE Configuration and Troubleshooting.
HPE7-A06Study Guide: Covers PoE management and diagnostics.
HPE Aruba Networking Technical Documentation: PoE Budget Troubleshooting.

NEW QUESTION: 19
When trying to add a now access switch to the network, theswitch port at the aggregation switch is automatically disabled.
What needs to be done to fix this issue?
A. Disable spanning tree bpdu-tilter al the interface level.
B. Disable spanning tree root-guard at the interface level.
C. Disable spanning tree loop-guard at the interface level.
D. Disable spanning tree bpdu-guard at the interface level.
Answer: D (LEAVE A REPLY)
The issue involves a new access switch's port being automatically disabled when connected to an aggregation switch, likely due to a Spanning Tree Protocol (STP) protection mechanism.
* Analysis of Options:
* Option A (Disable bpdu-filter):BPDU filtering prevents BPDUs from being sent or processed, which could cause loops, not resolve the issue.
* Option B (Disable root-guard):Root guard prevents a port from becoming the root bridge but does not cause port disablement in this context.
* Option C (Disable loop-guard):Loop guard prevents alternate ports from becoming designated but is unrelated to port disablement.
* Option D:Correct. Disabling BPDU guard on the aggregation switch's interface prevents it from disabling the port when it receives BPDUs from the new access switch.

* Why Option D is Correct:BPDU guard is an STP feature that disables a port if it receives BPDUs, assuming an unauthorized device is connected. When a new access switch isconnected, it sends BPDUs as part of normal STP operation, triggering BPDU guard on the aggregation switch and disabling the port. Disabling BPDU guard on the aggregation switch's interface (e.g., no spanning-tree bpdu-guard) allows the access switch to participate in STP without being disabled, resolving the issue while maintaining network stability.
* Relevance to Certification Objectives:
* Network Resiliency and Virtualization (8%):Involves troubleshooting STP mechanisms for fault tolerance.
* Troubleshooting (10%):Includes diagnosing and remediating STP-related issues in campus networks.
* Switching (19%):Covers Layer 2 technologies like STP and its protection features.
References:
HPE Aruba Networking AOS-CX Configuration Guide: Spanning Tree Configuration, detailing BPDU guard.
HPE7-A06Study Guide: Covers STP troubleshooting and protection mechanisms.
HPE Aruba Networking Technical Documentation: STP Best Practices, explaining BPDU guard behavior.

NEW QUESTION: 20
Identify the required configuration steps to enable DHCP EndpointProfiling with HPE Aruba Networking ClearPass. (Not all will be used)
Answer:
Explanation:
To enable DHCP Endpoint Profiling, the switch needs to forward relevant DHCP packets from the client to the ClearPass server. The most common method is configuring the switch to act as a DHCP relay agent for the ClearPass server IP address on the client VLAN's Switched Virtual Interface (SVI).
In scenarios where port access control (like 802.1X or MAC Auth) is enabled, clients might need to send DHCP requestsbeforethey are fully authenticated. To allow this while maintaining security, a pre- authentication role with limited access (specifically allowing DHCP) can be applied to the port initially.
The logical sequence based on the provided steps, assuming a pre-authentication workflow is intended, is:
* Create the role:Define the pre-authentication role container and associate it with the appropriate initial VLAN if needed.
* Permit DHCP in the role:Apply an Access Control List (ACL) or policy to this role that permits the necessary DHCP traffic (UDP ports 67 and 68). The step provided only mentions UDP 67, which allows the client's initial Discover/Request packets towards the server/relay. (A complete solution requires allowing return traffic on UDP 68 as well).

* Apply the role:Configure the client-facing physical interface to use this pre-authentication role before the final role is assigned post-authentication.
* Configure DHCP Relay:Configure the ip helper-address <clearpass_ip> command on the client's VLAN SVI. This instructs the switch to forward the DHCP packets it receives from clients in that VLAN to the ClearPass server (in addition to forwarding them to the actual DHCP server). ClearPass receives these packets and extracts information for profiling.
This sequence ensures that even before full authentication, DHCP is permitted, and the necessary packets are relayed to ClearPass for profiling.
References:AOS-CX Security Guide (Port Access, Roles, AAA), AOS-CX IP Helper / DHCP Relay Guide, ClearPass Deployment Guides (Endpoint Profiling using DHCP). This relates to "Authentication
/Authorization" (9%), "Security" (10%), and "Switching" (19%).

NEW QUESTION: 21
You are configuring an SSID that is using PSK as a security mechanism. Why should you use WPA3- Personal with WPA3 Transition Mode disabled?
A. WPA3-Porsonal with Transition Modedisabled is optional tor 6 GHz-enabled networks as there is a built-in tailback to 6 GHz mode with WPA2
B. WPAS-Personalwith Transition Mods disabled is mandatory for 5 GHz-enabled networks.
C. WPA3-Porsonal with Transition Mode disabled should be used to prevent legacy clients from connecting to thenetwork.
D. WPA3-Personal with Transition Mode disabled is mandatory for 6 GHz-enabled networks.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 22
Refer to the exhibit.
Acme Corp has VM workload running downstream of ToR-1 and has noticed performance degradation. They suspect ToR-1 uplinks are periodically over utilized. A partner has suggested you migrate your legacy 1U Coie-1 and Cote-2 to the CX 6400 series.
Which aspects of this platform would solve the customer's problem, while focusing on implementing HPE Aruba Networking best practices? (Select two.)
A. The CX 6400 series supports multiple active forwardingpathways from ToR-1 based on multi-region design.
B. The proposed new core's VSF capabilityallows multiple active forwarding pathways from ToR-1 based while eliminating the need for STP.
C. The proposed solutions backplane stacking permits the directly connected ESXI hosts to load balance using active LACP.
D. MC-LAG permits Core-1 and Core-2 to present the edge 602.3ad device as a common system ID"
Answer: B,D (LEAVE A REPLY)

The question involves a customer experiencing performance degradation due to periodic overutilization of ToR-1 uplinks to legacy Core-1 and Core-2 switches. The proposed solution is to migrate to CX 6400 series switches, and the task is to identify which aspects of the CX 6400 platform address the issue while adhering to HPE Aruba Networking best practices.
* Analysis of Options:
* Option A:Incorrect. The CX 6400 does not support "multi-region design" as a feature for active forwarding pathways.
* Option B:Correct. Virtual Switching Framework (VSF) on the CX 6400 allows multiple active forwarding pathways by creating a single logical switch from multiple physical switches, eliminating the need for STP in the core and reducing uplink congestion.
* Option C:Incorrect. Backplane stacking does not directly enable ESXi hosts to load balance using active LACP; this is unrelated to uplink utilization.
* Option D:Correct. Multi-Chassis Link Aggregation (MC-LAG) allows Core-1 and Core-2 to form a single logical 802.3ad (LACP) device, enabling active-active uplinks from ToR-1 and load balancing traffic to prevent overutilization.
* Why B and D are Correct:The performance degradation is caused by uplink overutilization, likely due to STP blocking redundant paths or inefficient load balancing. The CX 6400's VSF capability combines multiple switches into a single logical device, allowing all uplinks from ToR-1 to be active without relying on STP, which often blocks redundant paths. MC-LAG further enhances this by presenting Core-1 and Core-2 as a single LACP system, enabling ToR-1 to use all uplinks actively via LACP load balancing. These features align with HPE Aruba Networking best practices for high-availability and performance in campus core deployments.
* Relevance to Certification Objectives:
* Network Resiliency and Virtualization (8%):Designing and troubleshooting VSF and MC- LAG for resiliency and redundancy.
* Performance Optimization (6%):Analyzing and remediating uplink utilization issues.
* Connectivity (9%):Applying advanced networking architectures like VSF and MC-LAG.
References:
HPE Aruba Networking AOS-CX Configuration Guide: VSF and MC-LAG Configuration, detailing active forwarding and load balancing.
HPE7-A06Study Guide: Covers core switch resiliency and performance optimization.
HPE Aruba Networking Technical Documentation: CX 6400 Series Deployment Best Practices.

NEW QUESTION: 23
You want to use OSPF to advertise a only .\16 summary route for the SVIsbelow to a neighbor In the same area (area 0).
Which configuration will achieve this?
A.
B.
C.
D.

E.

Answer: E (LEAVE A REPLY)

The goal is to configure OSPF on a router so that it advertisesonlya 10.1.0.0/16 summary route for the specific SVIs (VLAN 11, 12, 13, assumed to be within the 10.1.x.x range) to its OSPF neighborswithin the same area (Area 0).

* OSPF Intra-Area Behavior:A fundamental principle of OSPF (link-state protocols) is that all routers within the same area must have an identical Link State Database (LSDB) for that area. This means all routers learn about all the specific networks (Type-1 Router LSAs, Type-2 Network LSAs) within their area. OSPFv2 doesnotsupport summarizing routes in a way that hides specific network LSAs from other routerswithin the same area. Summarization occurs only at area boundaries (by ABRs using Type-
3 Summary LSAs via the area range command) or for external routes redistributed into OSPF (by ASBRs using Type-5 External LSAs via the summary-address command).

* Analysis of Options:

* A) area 0 range 10.1.0.0/16:This command is used on an Area Border Router (ABR) to summarize routes originatingfromArea 0 when advertising theminto another area(e.g., the backbone). It does not affect LSA floodingwithinArea 0. It also includes redistribute connected, which is unrelated here.

* B) summary-address 10.1.0.0/16:This command is used on an Autonomous System Boundary Router (ASBR) to summarizeexternalroutes being redistributed into OSPF. It is not used for summarizing internal OSPF routes like SVIs defined within an OSPF area.

* C) & D) summary-address 10.1.0.0/16:Same issue as B; incorrect command for summarizing internal OSPF routes.

* E) area 0 range 10.1.0.0/16:Similar to A, this uses the area range command. It correctly shows the SVIs configured for OSPF Area 0 first. However, like A, this command performs inter-area summarization on an ABR and does not suppress the specific LSAswithinArea 0.

* Conclusion:The question asks for something that OSPFv2 cannot do: advertiseonlya summary route within the same area while suppressing specifics. Therefore, none of the configurations will achieve the exactstated outcome. However, if the question is flawed and intends to ask which configuration uses the correct command structure for summarizinginternalOSPF routes (even if only effective between areas), then the area range command is the relevant one. Both A and E use this command. Option E is slightly better structured as it shows the interfaces being added to OSPF Area 0 first. Assuming this is the intended direction despite the impossibility of the specific request, E is the most plausible choice among the given options.

References:RFC 2328 (OSPFv2), OSPF Configuration Guides for AOS-CX (explaining area range for ABRs and summary-address for ASBRs). This relates to the "Routing" (16%) objective.

NEW QUESTION: 24

Match the AOS-CX switch BGP keepalive and holddown timersto the default.

Answer:

Explanation:

The question requires matching the default BGP keepalive and hold-down timers on AOS-CX switches to their respective values.
* Analysis of Options:
* Keepalive Timer:The keepalive timer determines how often BGP keepalive messages are sent to maintain a session. The default value on AOS-CX switches is 60 seconds.
* Hold-down Timer:The hold-down timer specifies the maximum time a BGP session can remain active without receiving a keepalive or updatemessage before it is considered down. The default value on AOS-CX switches is 180 seconds.
* Why This Mapping is Correct:Per BGP standards (RFC 4271) and HPE Aruba Networking AOS-CX documentation, the default BGP keepalive timer is 60 seconds, and the hold-down timer is 180 seconds (three times the keepalive interval). These timers ensure BGP sessions remain stable while allowing timely detection of peer failures. The AOS-CX implementation adheres to these defaults unless explicitly configured otherwise.
* Relevance to Certification Objectives:
* Routing (16%):Involves designing and troubleshooting BGP routing topologies, including timer configurations.
* Troubleshooting (10%):Includes diagnosing BGP session issues related to timers.
References:
HPE Aruba Networking AOS-CX Configuration Guide: BGP Configuration, detailing default timer values.
HPE7-A06Study Guide: Covers BGP session management and timers.
RFC 4271: A Border Gateway Protocol 4 (BGP-4), specifying default keepalive and hold-down timers.

NEW QUESTION: 25
A customer wants to deploy IoT security devices that are PoE-powered. Dueto its criticality. it is required that those devices remain active even during a switch software upgrade. What is a valid solution to meet customer requirements?
A. a VSX pairof switches for redundancy
B. power-over-ether net quick-poe
C. power-over-ethernet always-on
D. power-over-ethernel priority
Answer: (SHOW ANSWER)
The question involves a customer deploying PoE-powered IoT security devices (e.g., door locks) that must remain active during an AOS-CX switch software upgrade. The task is to identify a valid solution.
* Analysis of Options:
* Option A:Incorrect. A VSX pair provides redundancy but does not guarantee PoE continuity during a single switch's upgrade.
* Option B:Incorrect. quick-poe reduces PoE startup time but does not ensure power during upgrades.

* Option C:Correct. power-over-ethernet always-on ensures PoE remains active during software upgrades, meeting the requirement.
* Option D:Incorrect. PoE priority adjusts power allocation but does not guarantee continuity during upgrades.
* Why Option C is Correct:The power-over-ethernet always-on feature on AOS-CX switches ensures that PoE power delivery continues uninterrupted during software upgrades or reboots, critical for devices like IoT security door locks that require constant power. This feature prevents power cycling on PoE ports, maintaining device operation. For example, enabling it on relevant ports (e.g., interface 1/1/1 power-over-ethernet always-on) ensures compliance with the customer's requirement, as per HPE Aruba Networking's PoE high-availability guidelines.
* Relevance to Certification Objectives:
* Connectivity (9%):Configuring PoE for critical device deployment.
* Network Resiliency and Virtualization (8%):Ensuring device uptime during maintenance.
* Troubleshooting (10%):Resolving PoE continuity issues.
References:
HPE Aruba Networking AOS-CX Configuration Guide: PoE Always-On Feature.
HPE7-A06Study Guide: Covers PoE configuration for high-availability devices.
HPE Aruba Networking Technical Documentation: PoE Best Practices for IoT.

NEW QUESTION: 26
Match the customer requirement with the relevant commands.
Answer:
Explanation:
* Aggregate links across multiple switches -->
vsx
role primary
inter-switch-link lag 256
keepalive peer 192.168.0.1 source 192.168.0.0 vrf KA
(Snippet 4)
* Establish redundant links between the aggregation and core layers --> router ospf 1 maximum-paths 2 (Snippet 2)
* Extend layer 2 across multiple sites -->
interface vxlan 1
no shutdown
source ip 10.1.0.4
(Snippet 1)
* Identify individual layer 2 segments in an overlay -->
vni 11
vtep-peer 10.1.0.5
vlan 11
(Snippet 3)

Comprehensive Detailed Explanation along with All References available from related to the HPE Campus Access Switching Expert certification objectives at end of each question below:
* Aggregate links across multiple switches:This requirement describes Multi-Chassis Link Aggregation (MC-LAG), where a device forms a LAG to two separate upstream switches that act as a logical pair. In AOS-CX, VSX (Virtual Switching Extension) enables this functionality. Snippet 4 shows commands related to setting up VSX (vsx, role primary, inter-switch-link, keepalive), which is the foundation for MC-LAG.
References:AOS-CX VSX Guide.Relates to "Network Resiliency and virtualization" (8%), "Switching" (19%).
Establish redundant links between the aggregation and core layers:This often involves Layer 3 routing protocols utilizing multiple paths. Snippet 2 (router ospf 1, maximum-paths 2) configures OSPF to use up to two Equal Cost Multi-Paths (ECMP). If redundant links between aggregation and core result in equal OSPF costs, this command enables load sharing and redundancy at Layer 3.
References:AOS-CX IP Routing Guide (OSPF, ECMP). Relates to "Routing" (16%), "Network Resiliency and virtualization" (8%).
Extend layer 2 across multiple sites:VXLAN (Virtual Extensible LAN) is the standard overlay technology for extending Layer 2 segments over an underlying Layer 3 network, enabling L2 adjacency across different physical locations (sites, racks, pods). Snippet 1 shows the basic configuration of a VXLAN tunnel interface (interface vxlan 1, source ip), which is the core component for VXLAN tunneling.
References:AOS-CX VXLAN Guide.Relates to "Switching" (19%), "Connectivity" (9%).
Identify individual layer 2 segments in an overlay:Within a VXLAN overlay, each separate Layer 2 broadcast domain (typically corresponding to a VLAN) is identified by a unique VXLAN Network Identifier (VNI). This VNI tags the encapsulated traffic. Snippet 3 shows the configuration associating VNI 11 with the local VLAN 11 (vni 11, vlan 11). The vtep-peer command is relevant when using EVPN as the control plane.
This configuration directly maps an L2 segment (VLAN 11) to its identifier (VNI 11) within the overlay.
References:AOS-CX EVPN Guide, AOS-CX VXLAN Guide.Relates to "Switching" (19%), "Connectivity" (9%).

NEW QUESTION: 27
Match the network technology to the customer requirement.
Answer:
* Establish redundant links between the aggregation and core layers:When using Layer 3 routing between network layers (like Aggregation and Core),ECMP (Equal Cost Multi-Path)allows the routing protocol (e.g., OSPF, BGP) to utilize multiple links simultaneously if they have the same routing cost. This provides both redundancy (if one link fails, traffic uses the others) and load sharing across the links.

References:AOS-CX IP Routing Guide (OSPF, BGP, ECMP). Relates to "Routing" (16%), "Network Resiliency and virtualization" (8%).
Extend layer 2 across multiple sites:VXLAN (Virtual Extensible LAN)is the overlay technology specifically designed for this purpose. It encapsulates Layer 2 Ethernet frames within UDP packets, allowing them to be tunneled across an underlying Layer 3 network infrastructure, effectively stretching Layer 2 domains (VLANs) between physically separate locations.
References:AOS-CX VXLAN Guide.Relates to "Switching" (19%), "Connectivity" (9%).
Identify individual layer 2 segments in an overlay:Inside the VXLAN header, theVNI (VXLAN Network Identifier)serves as the segment identifier. Each unique Layer 2 segment (like a specific VLAN being extended) is mapped to a unique 24-bit VNI, allowing the overlay network to differentiate between traffic belonging to different L2 domains, even when tunneled between the same VTEPs (VXLAN Tunnel Endpoints).
References:AOS-CX VXLAN Guide, RFC 7348 (VXLAN).Relates to "Switching" (19%), "Connectivity" (9%).
Minimize configuration steps to establish tunnels between sites:While VXLAN provides the data plane encapsulation,EVPN (Ethernet VPN)acts as the modern control plane for VXLAN overlays. Using MP-BGP extensions, EVPN dynamically discovers VTEPs and advertises MAC address and IP reachability information. This significantly reduces configuration complexity compared to older static VXLAN or flood- and-learn methods, as VTEP peer relationships and endpoint learning are automated by the control plane, thus minimizing manual steps to establish connectivity.
References:AOS-CX EVPN Guide.Relates to "Routing" (16%), "Switching" (19%), "Connectivity" (9%).

NEW QUESTION: 28
You are configuring an SSID thatisusing 802.IX as a security mechanism. Whatis the reason tor using WPA3- Ertorpnse (CCM-128) when deploying Wi-Fi 6 networks?
A. WPA3-Entorprise(CCM-128) is also called WPA3-Entorpnso 192-bit mode. It is WPA3 only and enforces specific EAP certificate ciphers.
B. WPA3-Enterprise (CCM-128) is also called WPA3-Ente'prise Transition Mode. It will allow WPA2 clients to conned.
C. WPA3-Enterprise(CCM-128) is also called WPA3-Eniarpnse Compatibility Mode. It will allow WPA2 clients to connect.
D. WPA3-Enterprise(CCM-128) is also called WPA3-EnterpriseOnly Mode. There is no support tor WPA2 clients.
Answer: (SHOW ANSWER)
The question asks for the reason for using WPA3-Enterprise (CCM-128) when deploying Wi-Fi 6 networks.
* WPA3-Enterprise Modes:
* CCM-128:Uses AES-CCMP-128 (same cipher as WPA2). Its main purpose is to provide a transition path from WPA2 to WPA3. It allows both WPA3-capable and WPA2-only clients to

connect to the same SSID. It enforces Protected Management Frames (PMF, 802.11w) when possible(required for WPA3, optional for WPA2). It's often called "Transition Mode" or "Compatibility Mode".
* GCMP-256:Uses stronger AES-GCMP-256. It operates in "WPA3-Only Mode" and doesnot allow WPA2 clients.
* Wi-Fi 6 (802.11ax) & WPA3:Wi-Fi 6 certification requires support for WPA3.
* Analysis of Options:
* A: Incorrectly calls CCM-128 "192-bit mode" and "WPA3 only".
* B: Correctly calls CCM-128 "Transition Mode" and states it allows WPA2 clients.
* C: Correctly calls CCM-128 "Compatibility Mode" and states it allows WPA2 clients. "Compatibility Mode" and "Transition Mode" are used interchangeably for this WPA3 mode.
* D: Incorrectly calls CCM-128 "Only Mode" and states no WPA2 support.
* Conclusion:Both Option B and Option C accurately describe WPA3-Enterprise (CCM-128). It is designed as a transition/compatibility mode to allow environments to adopt WPA3 features (like mandatory PMF for capable clients) while still supporting legacy WPA2 clients on the same network during the migration period. Selecting either B or C would be functionally correct based on common terminology.
References:Wi-Fi Alliance WPA3 specifications, Aruba WPA3 deployment guides, 802.11ax standard information. This relates to the "WLAN" (9%) and "Security" (10%) objectives.

NEW QUESTION: 29
When using the cable diagnostic feature on an AOS-CX switch to test a 1000BaseT connection, whatthe accuracy of 'distance to fault"?
A. +/- 10m
B. +/- 2m
C. +/-6m
D. +/-1m
Answer: D (LEAVE A REPLY)
The question asks about the accuracy of the 'distance to fault' measurement provided by the cable diagnostic feature (using Time Domain Reflectometry - TDR) on an AOS-CX switch for a 1000BaseT connection.
* TDR Accuracy:TDR works by sending a signal down the cable and measuring the time it takes for reflections to return, which indicates faults like opens or shorts. The accuracy depends on the quality of the TDR circuitry, the calibration,and the cable characteristics. Network equipment vendors typically specify the expected accuracy.
* AOS-CX Specification:According to HPE Aruba Networking documentation for AOS-CX switches, the accuracy of the TDR-based cable diagnostics for distance to fault on copper cabling is typically specified as +/- 1 meter.
* Analysis of Options:
* A: +/- 10m - Too inaccurate.
* B: +/- 2m - Less accurate than specified.

* C: +/- 6m - Too inaccurate.
* D: +/- 1m - Matches the documented accuracy for AOS-CX TDR.
References:AOS-CX Fundamentals Guide, AOS-CX CLI Reference Guide (under diag cable-diagnostic command description or general troubleshooting sections). This relates to the "Troubleshooting" (10%) objective.

NEW QUESTION: 30
Match the BGP connection states to the conditions that could have caused that state.
Answer:
Explanation:
The last keepalive is less than 3 times the negotiated holddown timer. -->established The router has not received a response. The neighbor might be unreachable. -->active The router is waiting for an initial response from the neighbor. -->connect The router starts listening for a connection. -->idle This question requires matching specific BGP connection states from the BGP Finite State Machine (FSM) to descriptions of the router's activity or condition in those states.
* Idle:This is the starting state. The BGP process is administratively up but is not actively trying to connect. It refuses all incoming BGP connection attempts but listens for a start event (like configuration or operator initiation) or potentially listens for incoming connections if configured for passive peering.
* Matches:"The router starts listening for a connection." (This describes the passive aspect of the Idle state before active attempts begin).
* Connect:In this state, BGP is actively trying to establish a TCP connection with the peer. It has initiated the TCP three-way handshake and is waiting for it to complete, or it is waiting for a remote peer to initiate the TCP connection.
* Matches:"The router is waiting for an initial response from the neighbor." (Specifically, waiting for the TCP handshake to complete).
* Active:If the TCP connection attempt in the Connect state fails (e.g., timeout), the router transitions to the Active state. In this state, it will periodically retry establishing the TCP connection while also listening for an incoming connection from the peer. This state indicates repeated failures to establish TCP connectivity.
* Matches:"The router has not received a response. The neighbor might be unreachable." (This reflects the condition in the Active state where connection attempts fail, suggesting the neighbor is unreachable at the TCP level).
* Established:This is the final, operational state where the TCP connection is up, BGP session parameters have been successfully negotiated via OPEN messages, and KEEPALIVE messages are being exchanged. Routing information (UPDATEs) can be exchanged. The condition described implies the session is healthy and timers are being maintained.
* Matches:"The last keepalive is less than 3 times the negotiated holddown timer." (While phrased slightly unusually, this indicates the holddown timer hasnotexpired because keepalives are being received within the expected window (Holddown Timer = ~3 * Keepalive Interval). This confirms the session is alive, which is true in the Established state).

References:RFC 4271 (BGP4 Specification - Section 8, Finite State Machine), BGP configuration and troubleshooting guides for AOS-CX. This relates to the "Routing" (16%) and "Troubleshooting" (10%) objectives.

NEW QUESTION: 31
With the configuration oftwo CX 8325 switches in the VSX cluster, how would you prepare a link-aggregation for a 7000 gateway for a zero-touch provision to support protocol-based port redundancy?
A.
B.
C.
D.
Answer: B (LEAVE A REPLY)
The goal is to configure a Link Aggregation Group (LAG) on a VSX cluster (pair of CX 8325 switches) that connects to an Aruba 7000 series gateway undergoing Zero Touch Provisioning (ZTP). The LAG needs to support "protocol-based port redundancy" (LACP) and allow connectivity during ZTP.
* VSX Requirement:Since the LAG connects to two separate physical switches operating as a VSX pair, the LAG must be configured as a Multi-Chassis LAG (MC-LAG) on the switches. This allows the gateway to form a single LAG across both upstream devices. The command multi-chassis under the interface lag <id> context enables this.
* Protocol Redundancy Requirement:"Protocol-based port redundancy" indicates that Link Aggregation Control Protocol (LACP) should be used to dynamically negotiate and manage the LAG bundle between the switches and the gateway. The command lacp mode active enables LACP in active negotiation mode.
* ZTP Requirement:During ZTP, the gateway might not have its full configuration, including LACP settings, enabled immediately. To ensure the gateway can establish basic IP connectivity for ZTP (e.g., reach Activate/Central via DHCP/DNS), the switch ports should allow traffic even if LACP negotiation hasn't completed. The lacp fallback feature enables this, allowing individual LAG member ports to become active if LACP PDUs are not received from the peer.
* Analyzing the Options:
* A)Configures lacp mode active and lacp fallback butlacksthe multi-chassis command required for VSX.
* B)Correctly configures the LAG as multi-chassis, enables lacp mode active, and enables lacp fallback. This meets all requirements.
* C)Configures multi-chassis but uses potentially older or less standard syntax lacp enable and lacp fail-over instead of lacp mode active and lacp fallback.
* D)Lacks the multi-chassis command and uses potentially older/less standard syntax.
* Conclusion:Option B provides the complete and correct configuration using standard AOS-CX syntax to create an MC-LAG on the VSX pair with LACP enabled for redundancy and LACP fallback enabled to support gateway connectivity during ZTP.

References:AOS-CX VSX Guide (MC-LAG configuration), AOS-CX Link Aggregation Guide (LACP, LACP Fallback commands and usage), ArubaGateway ZTP documentation. This relates to "Network Resiliency and virtualization" (8%), "Switching" (19%), and "Connectivity" (9%) objectives.

NEW QUESTION: 32
The user's device is failing 802.1 Xwith EAP-TLS authentication. We know that theclient-side certificate is valid. What is the likely cause of this issue? (Select two.)
A. The user's device is not configured to use the correct gateway.
B. There is a problem with the ACL applied to the switch port
C. There Is an EAP-type mismatch.
D. The user's device is using the wrong MAC address
E. The NAD is not able to communicate with DNS servers.
Answer: (SHOW ANSWER)
The user's device fails 802.1X EAP-TLS authentication, but the client-side certificate is known to be valid.
We need two likely causes.
* EAP-TLS Process:Involves mutual certificate validation and TLS handshake between client and RADIUS server (proxied by NAD).
* Causes (Client Cert OK):
* Server Certificate Issues: Client doesn't trust server cert (Untrusted CA, name mismatch, expired).
* EAP Type Mismatch:Client supplicant configured for different EAP type than RADIUS server policy.
* RADIUS Server Issues:Policy misconfiguration, user not found, internal errors.
* NAD <-> RADIUS Communication Failure:Switch cannot reach RADIUS server (IP connectivity, firewall, routing), incorrect shared secret.
* Client Supplicant Misconfiguration:Incorrect identity, settings other than the certificate itself.
* Network packet loss.
* Analysis of Options (Select Two):
* A: Wrong gateway affects L3 post-authentication.
* B: ACL blocking EAPoL/RADIUS is possible but less common than config errors.

* C:EAP-type mismatch:A very common configuration error leading to failure.
* D: Wrong MAC address is irrelevant for EAP-TLS failure itself.
* E: NAD not able to communicate with DNS servers: DNS isn't directly involved in EAP-TLS. However, if interpreted more broadly asNAD not able to communicate with the RADIUS server(due to IP routing, firewall, or incorrect server address), this is a very common cause of failure.
* Conclusion:An EAP-type mismatch (C) is a prime suspect when basic certificate validity is assumed.
Failure of the Network Access Device (NAD - the switch) to communicate with the RADIUS server (E, interpreted broadly as RADIUS reachability) is another major category of failure causes.
References:EAP-TLS (RFC 5216), 802.1X Troubleshooting Guides, ClearPass Documentation.
This relates to "Troubleshooting" (10%), "Security" (10%), and "Authentication/Authorization" (9%).