Cisco

(300-620)

Implementing Cisco Application Centric Infrastructure (DCACI)

Total: **311 Questions**

Link:

Question: 1

An engineer is implementing a Cisco ACI data center network that includes Cisco Nexus 2000 Series 10G fabric extenders. Which physical topology is supported?

A.

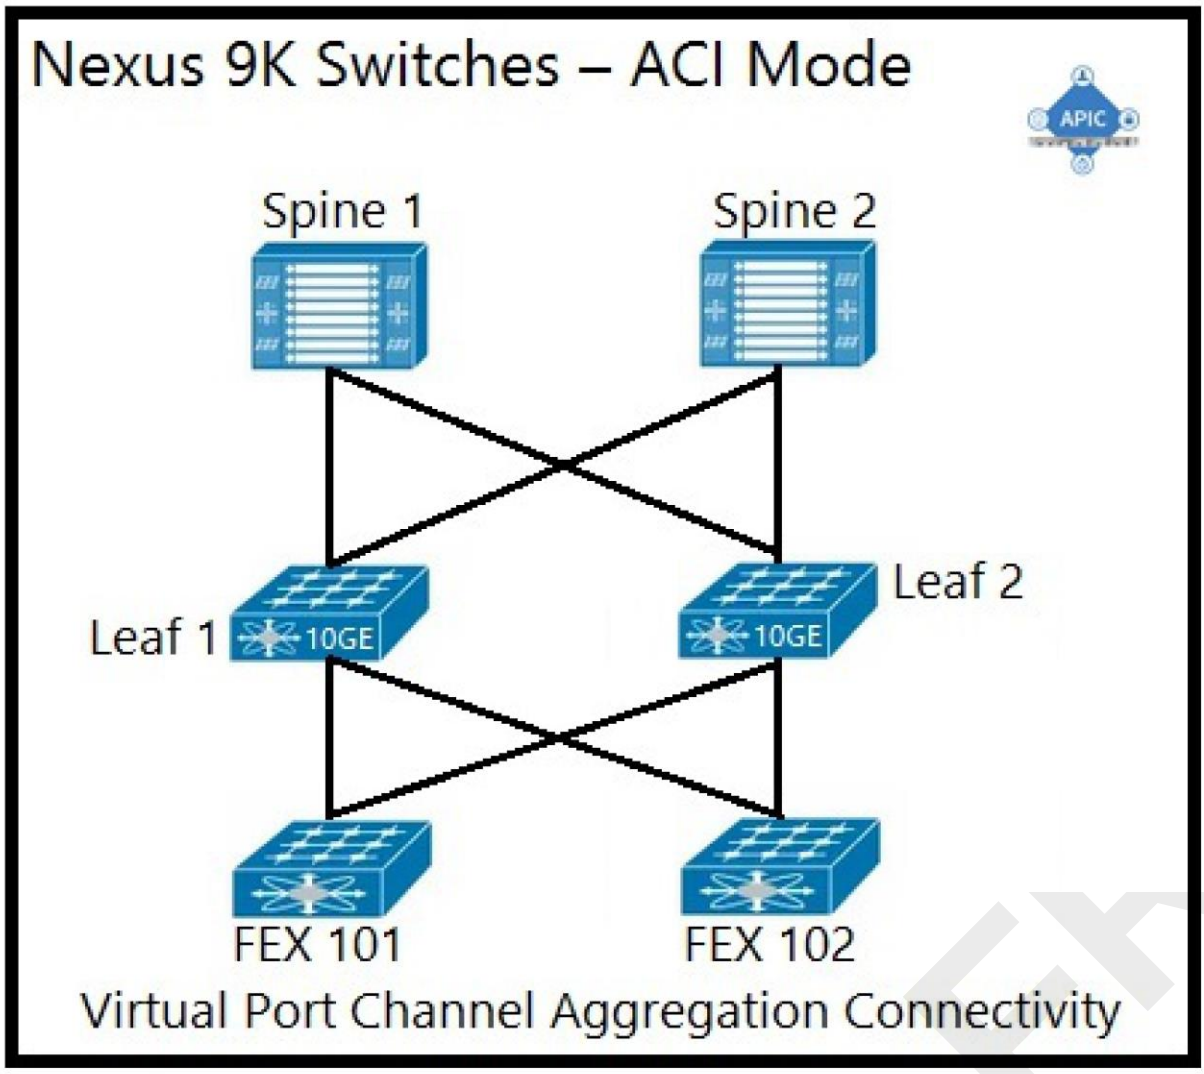

B.

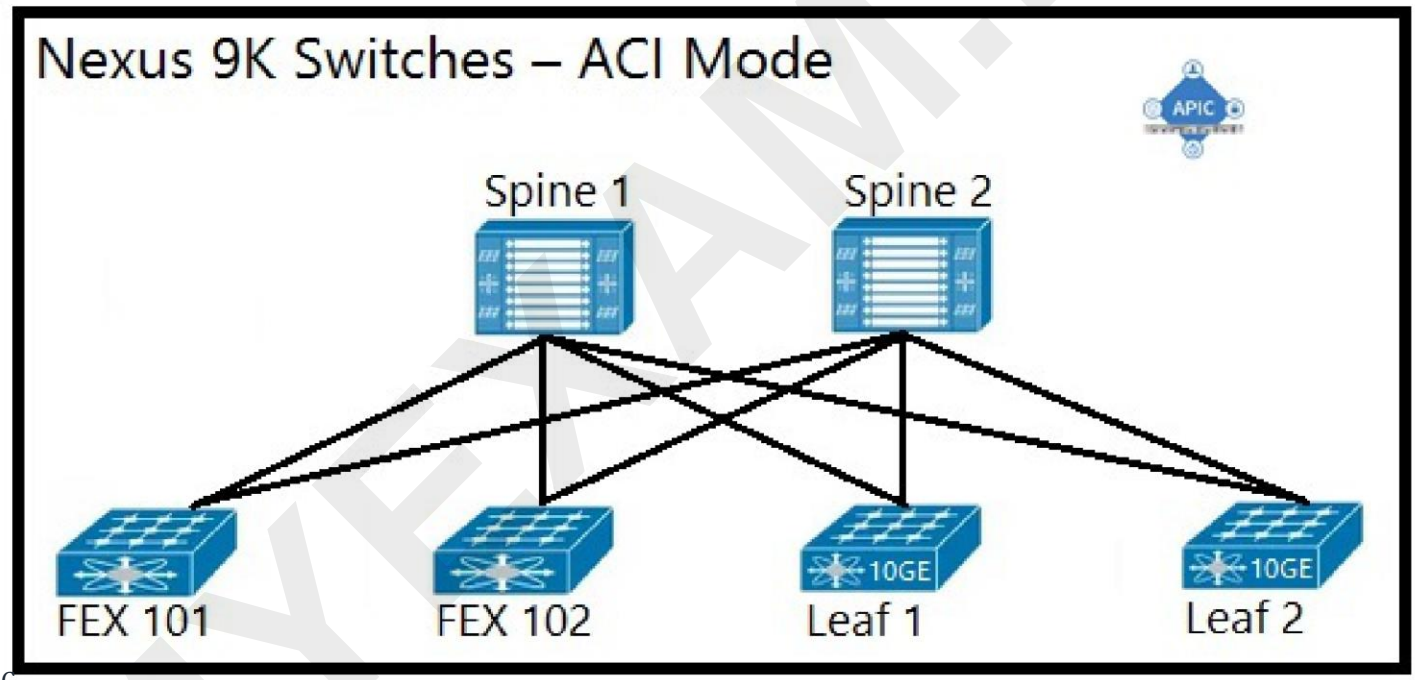

C.

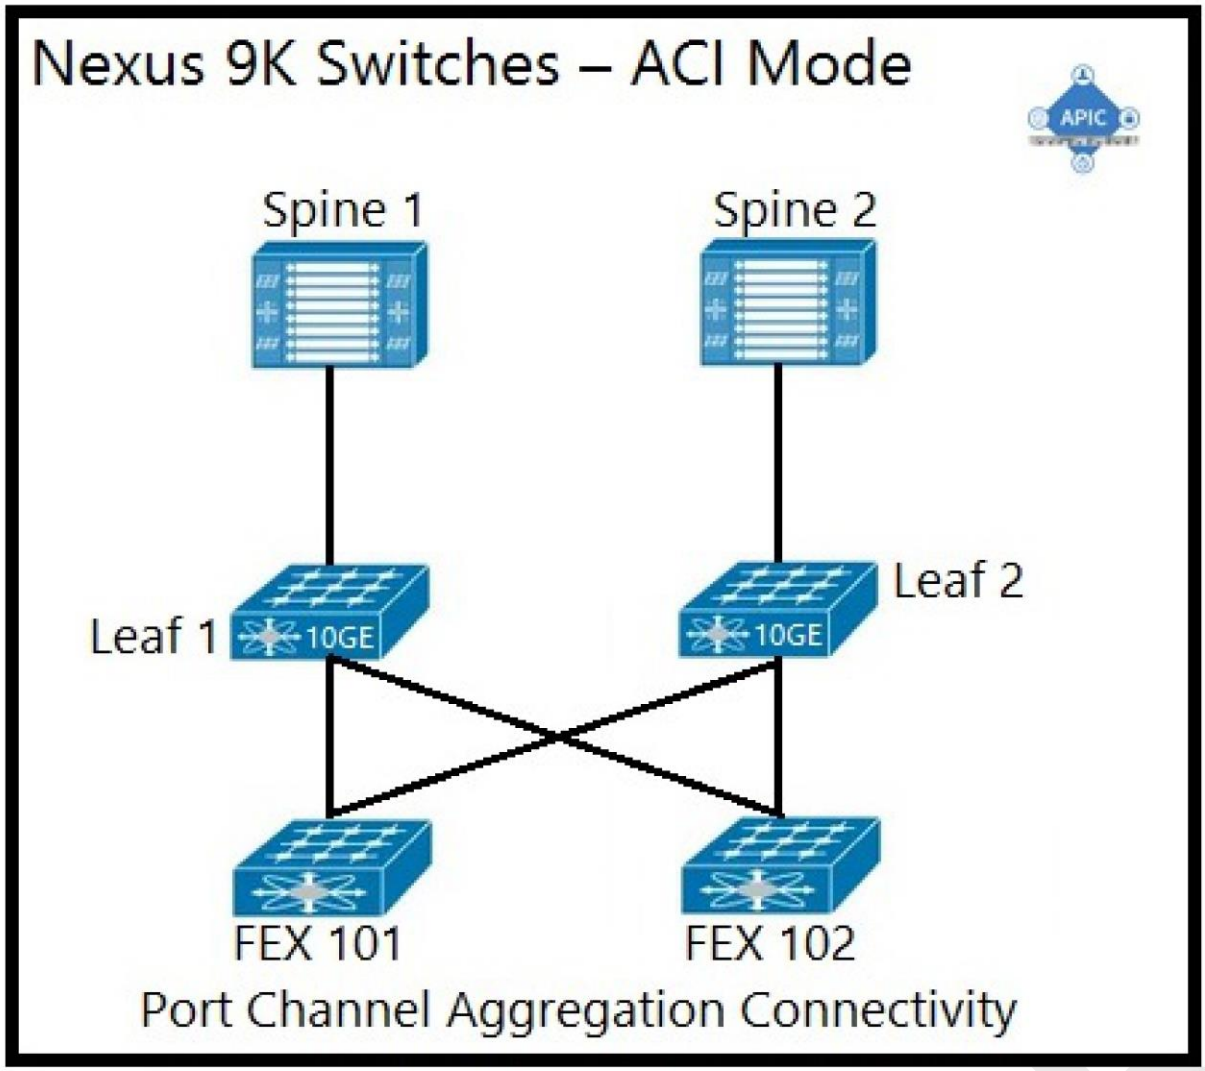
Nexus 9K Switches – ACI Mode
APIC
Spine 1
Spine 2
Leaf 1
10GE
Leaf 2
10GE
FEX 101
FEX 102
Port Channel Aggregation Connectivity

D.

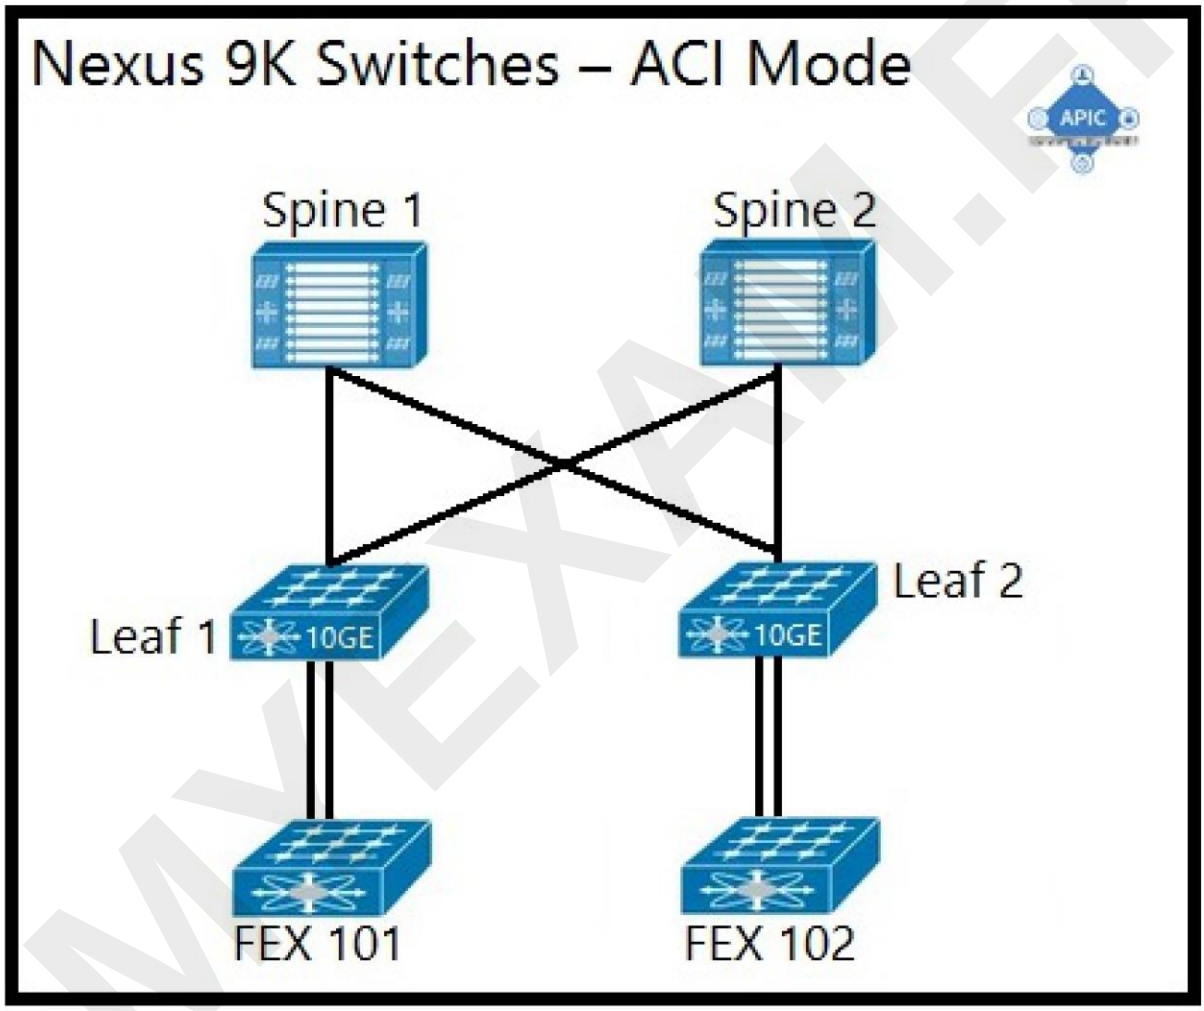
Nexus 9K Switches – ACI Mode
APIC
Spine 1
Spine 2
Leaf 1
10GE
Leaf 2
10GE
FEX 101
FEX 102

Answer: D

Explanation:

FEX (N2K Switches) must be connected to one parent switch (N9K) .

D is correct.

Question: 2

An ACI administrator notices a change in the behavior of the fabric. Which action must be taken to determine if a human intervention introduced the change?

A.Inspect event records in the APIC UI to see all actions performed by users.
B.Inspect /var/log/audit_messages on the APIC to see a record of all user actions.
C.Inspect audit logs in the APIC UI to see all user events.
D.Inspect the output of show command history in the APIC CLI.

Answer: C

Explanation:

Option C, "Inspect audit logs in the APIC UI to see all user events," is the correct answer because Cisco ACI's Application Policy Infrastructure Controller (APIC) maintains comprehensive audit logs to track all user-initiated changes within the fabric. These logs record specific actions, the user who performed them, and the timestamps, providing a clear timeline of events. This information is crucial for identifying if a human intervention, rather than an automated process or other system malfunction, caused a change in fabric behavior. Option A, while referring to events, doesn't specifically target user actions in a way that audit logs do. Option B, accessing the file system directly, is not the recommended or standardized method for reviewing user actions, and lacks the structured interface that the audit logs provide. Option D, the CLI command history, is limited to commands executed directly in the CLI, not the more comprehensive record of user actions across the system. Audit logs are a cornerstone of security and troubleshooting in enterprise systems, and in the APIC, they offer a centralized and readily accessible record of user actions, facilitating efficient root cause analysis. Examining these logs allows administrators to quickly pinpoint any human intervention that may have led to the observed change.

Here are some authoritative links for further research:

1. **Cisco ACI Fundamentals Guide:** While a specific guide on audit logs may not exist, the foundational guides on ACI will reference the logging and monitoring capabilities of the APIC, which covers audit logging. Search on cisco.com for "Cisco ACI Fundamentals".
2. **Cisco ACI Documentation:** Search the specific documentation for your APIC version on cisco.com for more details on audit log management. Keywords include "audit log", "user activity", "monitoring", "APIC logs".
3. **Cisco Live Presentations:** Search Cisco Live archives on ciscolive.com using keywords such as "ACI monitoring," or "ACI Troubleshooting" for relevant presentations that cover audit logs within APIC.
4. **Cisco ACI Configuration Guides:** Look for configuration guides on ciscolive.com which often have sections relating to monitoring and audit logging.

Question: 3

An engineer is creating a configuration import policy that must terminate if the imported configuration is

incompatible with the existing system. Which import mode achieves this result?

A.merge
B.atomic
C.best effort
D.replace

Answer: B

Explanation:

The correct answer is **B. atomic**. Here's why:

Cisco ACI import policies offer different modes to handle how imported configurations interact with the existing system. The "atomic" import mode is designed to be an all-or-nothing operation. If, during the import process, ACI detects any incompatibility between the imported configuration and the running configuration, the entire import operation is immediately terminated and rolled back, leaving the existing system unchanged.

This ensures that the system doesn't end up in an inconsistent or unstable state due to partial configuration application. This behavior is aligned with the principles of transactional integrity in database systems, guaranteeing consistency by either completing all changes or none at all.

The other options don't provide the desired behavior. "Merge" attempts to combine the imported configuration with the existing one, potentially overwriting existing parameters or creating inconsistencies. "Best effort" imports as much of the configuration as possible, ignoring incompatibilities and possibly leaving the system in an undesirable or incomplete state. "Replace" completely discards the existing configuration and replaces it with the imported one, which is not the intent of the question. The primary goal here is to detect and prevent incompatible configurations from being applied, a characteristic solely provided by the atomic import mode.

Therefore, for scenarios requiring strict adherence to compatibility and preventing partial configuration updates, the atomic import mode is the most suitable choice.

For further research, refer to the Cisco documentation on ACI import policies. Specifically, search for “Cisco ACI configuration import policies” or look at documentation detailing the 'atomic' import mode and its characteristics. Here is a relevant link that will assist in your research:

Cisco ACI Configuration Import and Export

Question: 4

Which components must be configured for the BGP Route Reflector policy to take effect?

A.spine fabric interface overrides and profiles
B.access policies and profiles
C.pod policy groups and profiles
D.leaf fabric interface overrides and profiles

Answer: C

Explanation:

The correct answer is **C. pod policy groups and profiles**. BGP Route Reflectors within Cisco ACI are configured at the pod level. Pod policies provide a container for configuring features that span multiple switches within a specific pod. To implement a BGP Route Reflector policy, you must define this policy within a pod policy group and associate it with the relevant pod profiles. This pod-level configuration ensures that the chosen spines will function as BGP route reflectors for all the leafs within that specific pod. Options A and

D, relating to fabric interfaces, are not directly used for implementing a BGP Route Reflector configuration.

Interface policies deal with the configuration of specific switch ports, not the overall BGP route reflection mechanism. Option B, access policies and profiles, primarily focus on configuring external network connectivity and do not govern the internal routing within the ACI fabric. The BGP Route Reflector functionality applies within the context of the pod, reflecting routes between leafs and providing internal reachability. In summary, the pod policy groups and profiles act as the central configuration point for enabling BGP route reflectors within Cisco ACI.

Authoritative Links:

Cisco ACI Fabric Routing:https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/l3-config/b-cisco-apic-layer-3-configuration-guide/m-bgp.html (Refer to the section on BGP Route Reflector configuration within this guide)
Cisco ACI Pods:https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/getting-started/b-cisco-apic-getting-started-guide/m-aci-fabric-setup.html (Provides an understanding of pod structure within ACI)

Question: 5

Which type of policy configures the suppression of faults that are generated from a port being down?

A.fault lifecycle assignment
B.event lifecycle assignment
C.fault severity assignment
D.event severity assignment

Answer: C

Explanation:

The correct answer is **C. fault severity assignment**. Fault severity assignment policies within Cisco ACI (Application Centric Infrastructure) allow administrators to modify the default severity levels of various faults.

This includes the ability to suppress or "silence" faults. When a port goes down, it triggers a fault. By configuring a fault severity assignment policy that maps this specific fault (associated with a port down event) to a severity of "cleared" or "ignored," the system effectively stops generating alerts for that condition. This is a direct way to control which fault events are actively reported and which are suppressed, providing a mechanism for focused event management. It doesn't change the fact that the fault occurred but prevents it from being an operational notification. Event lifecycle assignments are related to actions taken during the event occurrence, not the severity, while event severity assignments deal with changing the severity of events, which are different from faults. Fault lifecycle assignments are more focused on managing the overall progression of a fault rather than simply suppressing it. Hence, fault severity assignment is the specific mechanism used to suppress faults such as port downs.

Further Reading:

Cisco ACI Fault Management:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/fault-management/guide/b-cisco-apic-fault-management-guide/m-fault-policies.html
Cisco ACI Policy Model:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/white-paper-c11-730216.html

Question: 6

Which type of profile needs to be created to deploy an access port policy group?

A.attachable entity
B.Pod
C.module
D.leaf interface

Answer: A

Explanation:

The correct answer is **A. attachable entity profile (AEP)**. An AEP in Cisco ACI serves as a crucial intermediary, linking infrastructure policies to endpoints. When deploying an access port policy group, which defines configurations for ports connecting to devices like servers or virtual machines, you must associate it with an AEP. This association essentially tells the ACI fabric where to apply the policy.

An AEP defines a scope for policy deployment. It logically represents where endpoints connect to the fabric (e.g., server farms, network segments). It groups physical or virtual entities that share common policy requirements. When an endpoint, or device, is attached to a port configured with an access port policy group, the ACI fabric applies the configurations defined in that policy based on the AEP mapping.

Options B (Pod), C (Module), and D (Leaf Interface) are not directly used to associate access port policy groups. A Pod represents a collection of ACI switches, a module defines a specific piece of hardware, and a leaf interface is a physical port on a switch, although these are components of the ACI fabric, they don't define the policy attachment scope like an AEP does.

In summary, the AEP acts as the logical bridge that connects the port configuration policies defined in the access port policy group to the physical or virtual devices in the network. This link is critical to the policy deployment process in ACI.

Here are some authoritative links for further research:

1. **Cisco Application Centric Infrastructure Fundamentals:** https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/index.html (General information about Cisco ACI)
2. **Cisco ACI Policy Model White Paper:**https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-734447.html (Detailed explanation of the policy model, including AEPs)
3. **Cisco ACI Configuration Guide:** https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/configuration/cisco-aci-configuration-guide-52x/cisco-aci-configuration-guide-52x_chapter_01000.html (Specific configuration details, including AEP and policy group association)

Question: 7

A situation causes a fault to be raised on the APIC. The ACI administrator does not want that fault to be raised because it is not directly relevant to the environment. Which action should the administrator take to prevent the fault from appearing?

A.Under System -> Faults, right-click on the fault and select Acknowledge Fault so that acknowledged faults will immediately disappear.

B.Create a stats threshold policy with both rising and falling thresholds defined so that the critical severity threshold matches the squelched threshold.

C.Under System -> Faults, right-click on the fault and select Ignore Fault to create a fault severity assignment policy that hides the fault.

D.Create a new global health score policy that ignores specific faults as identified by their unique fault code.

Answer: C

Explanation:

The correct answer is **C. Under System -> Faults, right-click on the fault and select Ignore Fault to create a fault severity assignment policy that hides the fault.**

Here's why: Cisco ACI (Application Centric Infrastructure) uses a robust fault management system to alert administrators about issues. However, not all faults are relevant or require immediate action in every environment. Option C directly addresses this by utilizing the built-in fault severity assignment policy mechanism within the APIC (Application Policy Infrastructure Controller). By choosing "Ignore Fault," you're essentially creating a rule that suppresses the specific fault from being displayed, without necessarily resolving the underlying condition.

Option A is incorrect because acknowledging a fault doesn't prevent it from reappearing if the underlying condition persists; it merely marks it as reviewed. Option B, creating a stats threshold policy, deals with performance monitoring and alerts based on statistical data, not with suppressing specific, unwanted faults.

Option D, creating a global health score policy, modifies how the overall system health is calculated, but it does not suppress the appearance of individual faults in the fault list.

The "Ignore Fault" action within the APIC interface is designed explicitly for scenarios where an administrator wants to filter out noise and focus on more critical issues. It achieves this by adjusting the display of the fault based on a configured severity assignment policy, effectively preventing it from appearing in the fault list but the fault will continue to be generated by the system. This provides a method of dealing with known, non-critical alerts and simplifying the view for administrators.

For further research, you can refer to the official Cisco ACI documentation, specifically sections on Fault Management and Policies:

Cisco ACI Fault Management
Cisco ACI Fault Policies

Question: 8

A RADIUS user resolves its role via the Cisco AV Pair. What object does the Cisco AV Pair resolve to?

A.tenant
B.security domain
C.primary Cisco APIC
D.managed object class

Answer: D

Explanation:

The correct answer is **D. managed object class**. Here's why:

In Cisco ACI, user roles and permissions are mapped to managed objects, which represent configurable entities within the system. When a RADIUS user authenticates, the Cisco AV Pair (Attribute-Value Pair) sent by the RADIUS server carries information about the user's authorization level. Cisco ACI uses this information to resolve the user's permissions. Specifically, the AV Pair doesn't directly assign a user to a tenant, security domain, or APIC. Instead, it maps the user to a specific managed object class within the ACI object model. This class defines the permissions the user should have access to. Examples of managed object classes include

polUser, polUserRole, and other related classes that control user access and privileges in ACI. The specific managed object class used will depend on the desired role or permission level for the user. Therefore, the Cisco AV Pair essentially acts as a key, allowing ACI to find the appropriate managed object class which then governs the user's access and permissions.

For further research, consult the Cisco ACI documentation on role-based access control and external authentication:

Cisco ACI Security Guide:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/security/b_Cisco_APIC_Security_Guide.ht
(Focus on sections related to Role-Based Access Control and external authentication.)
Cisco ACI Object Model Documentation:https://developer.cisco.com/docs/aci/#!aci-object-model (Explore the specific managed object classes involved in user authentication and permissions)
Cisco ACI REST API User Guide:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/restapi/b_Cisco_APIC_REST_API_User
(Review the authentication and authorization sections and how managed objects are referenced.)

Question: 9

DRAG DROP -
An engineer is configuring a VRF for a tenant named Cisco. Drag and drop the child objects on the left onto the correct containers on the right for this configuration.
Select and Place:

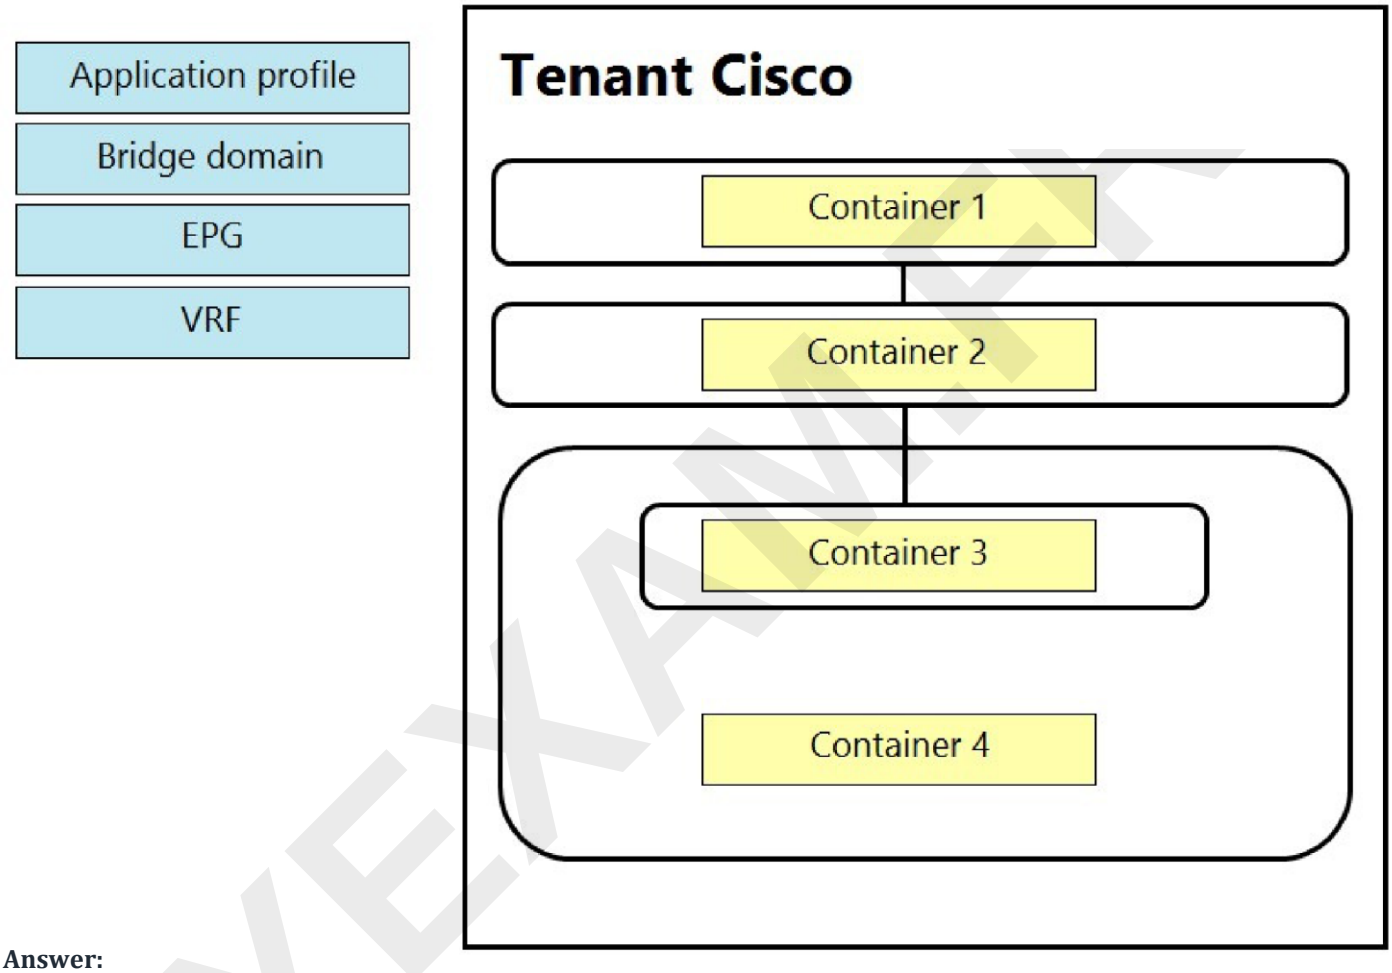

Answer:

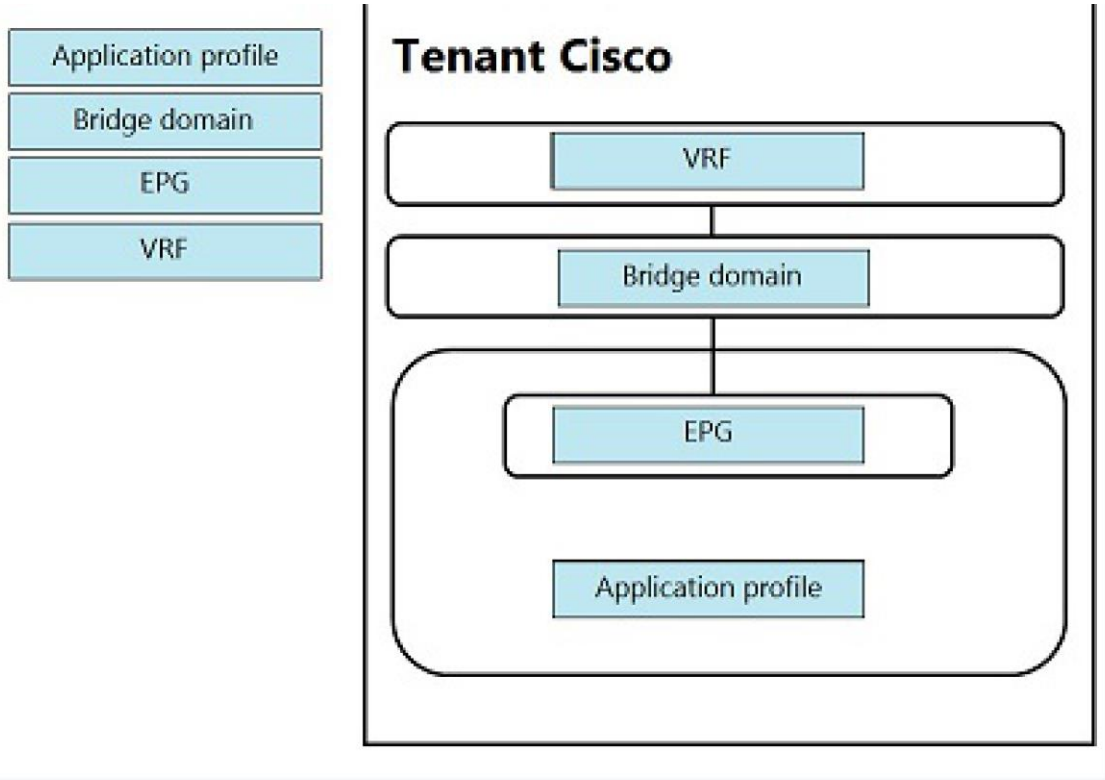

Explanation:

VRF>Bridge domain>EPG>Application Profile.

Question: 10

Which feature dynamically assigns or modifies the EPG association of virtual machines based on their attributes?

A.vzAny contracts
B.standard contracts
C.application EPGs
D.uSeg EPGs

Answer: D

Explanation:

uSeg EPGs (micro-segmented EPGs), also known as user-defined segmentation EPGs, are specifically designed to dynamically classify and associate virtual machines (VMs) to different EPGs based on their attributes. These attributes, which can include VM names, operating systems, custom tags, or any other configurable parameter, are used as selectors in the policy rules that govern EPG membership. Unlike standard or application EPGs, uSeg EPGs don't require predefined network segments or static IP
assignments. VMs are dynamically placed in the appropriate uSeg EPG as their attributes change, enabling granular security and policy enforcement. vzAny contracts, on the other hand, focus on providing
communication between different EPGs and are not concerned with EPG membership itself. Standard contracts define specific communication patterns between EPGs. Application EPGs represent applications as a whole and don't offer the dynamic, attribute-based VM membership feature inherent in uSeg EPGs. Therefore, only uSeg EPGs can dynamically assign or modify EPG associations based on VM attributes. This dynamic classification is essential for cloud environments where VM attributes and workloads can frequently change, ensuring ongoing compliance and security.

Here are links for further research:

Cisco Application Centric Infrastructure (ACI) Fundamentals:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/index.html - This link provides a general overview of Cisco ACI.

Cisco ACI Policy Model:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/policy/guide/b_Cisco_APIC_Policy_Model- This is the official documentation detailing the ACI policy model which includes uSeg EPGs.

Cisco ACI Microsegmentation White Paper:https://www.cisco.com/c/dam/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739104.pdf - A white paper specifically focusing on microsegmentation using uSeg EPGs in ACI.

Question: 11

Which feature allows firewall ACLs to be configured automatically when new endpoints are attached to an EPG?

A.ARP gleaning
B.dynamic endpoint attach
C.hardware proxy
D.network-stitching

Answer: B

Explanation:

Dynamic endpoint attachment, or dynamic EPG learning, in Cisco ACI facilitates the automatic configuration of firewall Access Control Lists (ACLs) based on endpoint IP and MAC addresses. When a new endpoint joins an EPG (Endpoint Group), ACI dynamically learns the endpoint's attributes and updates the policies relevant to that EPG. This eliminates manual ACL configurations and allows for seamless scaling and management of endpoint connections. This automatic learning process triggers updates to contracts and their filters, including firewall rules associated with the EPG. Essentially, when a new endpoint is attached, ACI recognizes it as a member of that EPG and dynamically modifies the policies. ARP gleaning, while related to endpoint discovery, doesn't directly drive ACL configuration changes. Hardware proxy and network-stitching, on the other hand, are separate concepts for optimizing resource usage and bridging different networks
respectively. Therefore, dynamic endpoint attachment directly relates to automation of ACL configurations when new endpoints connect to EPGs.

For further research on Dynamic EPG Learning and Cisco ACI, please refer to the official Cisco documentation:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/configuration/guide/b-cisco-apic-configuration-guide/b-cisco-apic-configuration-guide_chapter_0101.html#concept_A3633C457D7A4F14AC042FD6F52F5B9D

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/programming/guide/b-cisco-apic-programming-guide/b-cisco-apic-programming-guide_chapter_011.html#task_7C4105173C164954A3E049B4D74DF834

Question: 12

An engineer is implementing Cisco ACI at a large platform-as-a-service provider using APIC controllers, 9396PX leaf switches, and 9336PQ spine switches. The leaf switch ports are configured as IEEE 802.1p ports. Where does the traffic exit from the EPG in IEEE 802.1p mode in this configuration?

A.from leaf ports tagged as VLAN 0
B.from leaf ports untagged
C.from leaf ports tagged as VLAN 4094
D.from leaf ports tagged as VLAN 1

Answer: A

Explanation:

The correct answer is **A. from leaf ports tagged as VLAN 0**.

Here's why: When an ACI leaf port is configured for IEEE 802.1p mode, it doesn't use traditional VLAN tagging based on an assigned VLAN ID. Instead, it leverages the 802.1p priority field within the Ethernet frame. This priority field is used to indicate quality of service (QoS) requirements. However, ACI still needs a VLAN internally to track the traffic within the fabric. In this scenario, the system utilizes VLAN 0 internally to identify EPG traffic that's been tagged using the 802.1p priority field, and this tag is preserved when traffic exits an EPG via a leaf port configured in 802.1p mode. Because the VLAN tagging used in ACI is internal, the VLAN 0 tag is not what goes out onto the port.

Although the 802.1p priority field is what is directly related to QoS classification, ACI needs a VLAN to internally manage traffic on it's fabric, therefore when an EPG is configured for 802.1p mode, it uses VLAN 0 internally, and keeps the 802.1p tag intact as the traffic leaves the fabric, tagged with VLAN 0.

Options B, C, and D are incorrect because they don't reflect how ACI handles 802.1p traffic. Option B is incorrect because 802.1p uses a field in the Ethernet frame, so there are not untagged frames. Option C (VLAN 4094) is not a default or reserved VLAN for 802.1p tagging, and option D (VLAN 1) is normally the default VLAN, but it's not used to identify 802.1p EPG traffic in this context.

Authoritative Links for Further Research:

Cisco ACI Fundamentals:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/index.html - Provides a comprehensive overview of Cisco ACI concepts.

Cisco ACI Configuration Guides:https://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/products-configuration-examples-list.html -Offers detailed configuration examples and explanations for various ACI features, including EPGs and VLANs. **Cisco ACI Layer 2 Bridging:**https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/l2-configuration/cisco-apic-layer-2-config-guide-52x/m-aci-l2-bridging.html - More information about Layer 2 functionality in ACI

Question: 13

How is an EPG extended outside of the ACI fabric?

A.Create an external bridged network that is assigned to a leaf port.
B.Create an external routed network that is assigned to an EPG.
C.Enable unicast routing within an EPG.
D.Statically assign a VLAN ID to a leaf port in an EPG.

Answer: D

Explanation:

The correct answer is **D. Statically assign a VLAN ID to a leaf port in an EPG.** Here's why:

In Cisco ACI, an Endpoint Group (EPG) represents a collection of network endpoints that share similar policy

requirements. To extend an EPG outside the ACI fabric, typically to connect with legacy networks or external resources, you need to bridge the ACI's internal VXLAN-based network with the external world, which often uses traditional VLANs. This is achieved by statically associating a VLAN ID with a specific port on a leaf switch that belongs to the EPG.

Option A is incorrect because an external bridged network is a construct for interconnecting different subnets within ACI, not for reaching outside the fabric. Option B is wrong because external routed networks are used to interconnect L3 domains via a router, not for directly connecting external VLAN segments to an EPG. Option C is incorrect as unicast routing is an internal ACI function, not a mechanism for external connectivity with VLANs.

Statically assigning a VLAN ID to a leaf port essentially creates a "trunk" or a connection where a specific VLAN from the external network is mapped to the EPG on that port. Traffic tagged with that VLAN ID coming into that leaf port is considered to belong to that specific EPG, and vice-versa for traffic leaving. This is the most direct way to integrate external Layer 2 VLANs with an EPG in ACI.

For further research, see the Cisco documentation on EPGs and external connectivity within ACI:

Cisco ACI Fundamentals Guide:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/fundamentals/b-aci-fundamentals/m-epgs-and-bd.html
Cisco ACI External Routed Networks:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/getting-started/b_ACI_Getting_Started_Guide_4_x/b_ACI_Getting_Started_Guide_4_x_chapter_0110.html

Question: 14

DRAG DROP -
Drag and drop the Cisco ACI filter entry options from the left onto the correct categories on the right indicating what are required or optional parameters.
Select and Place:

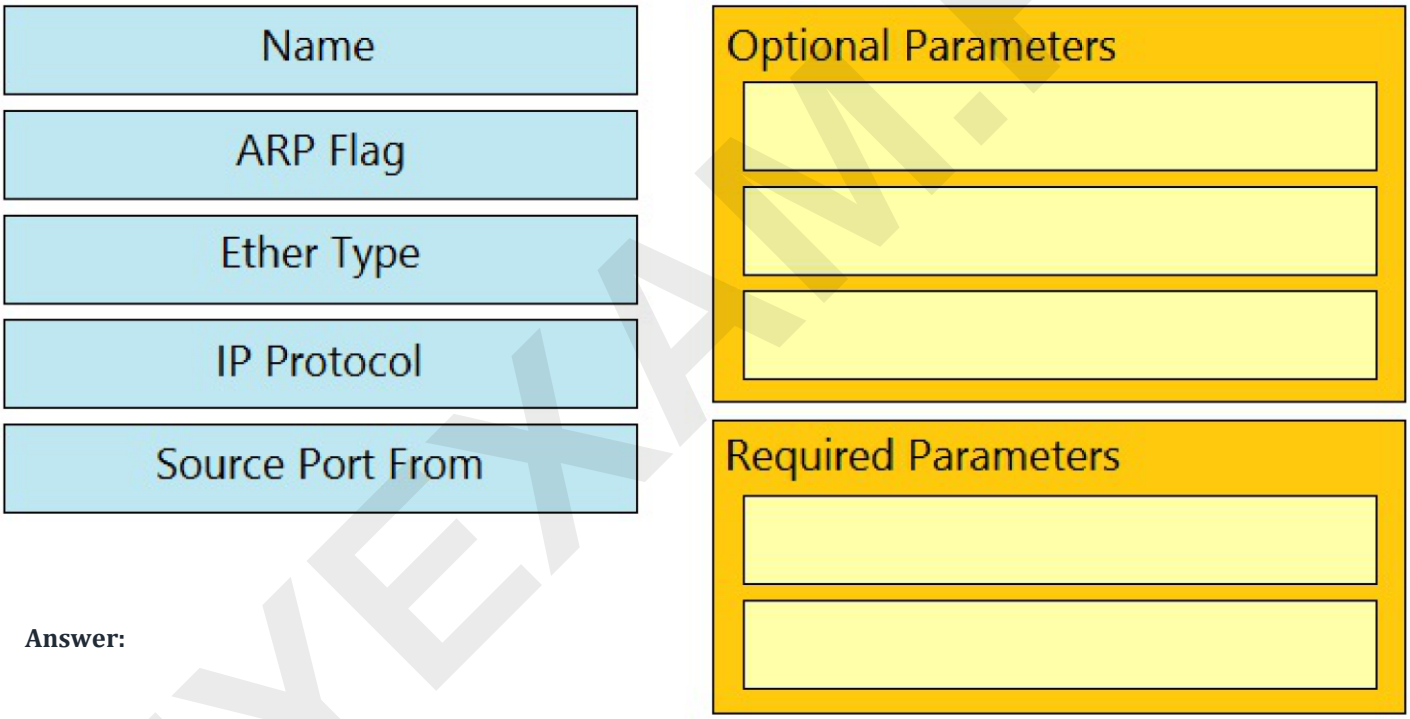

Answer:

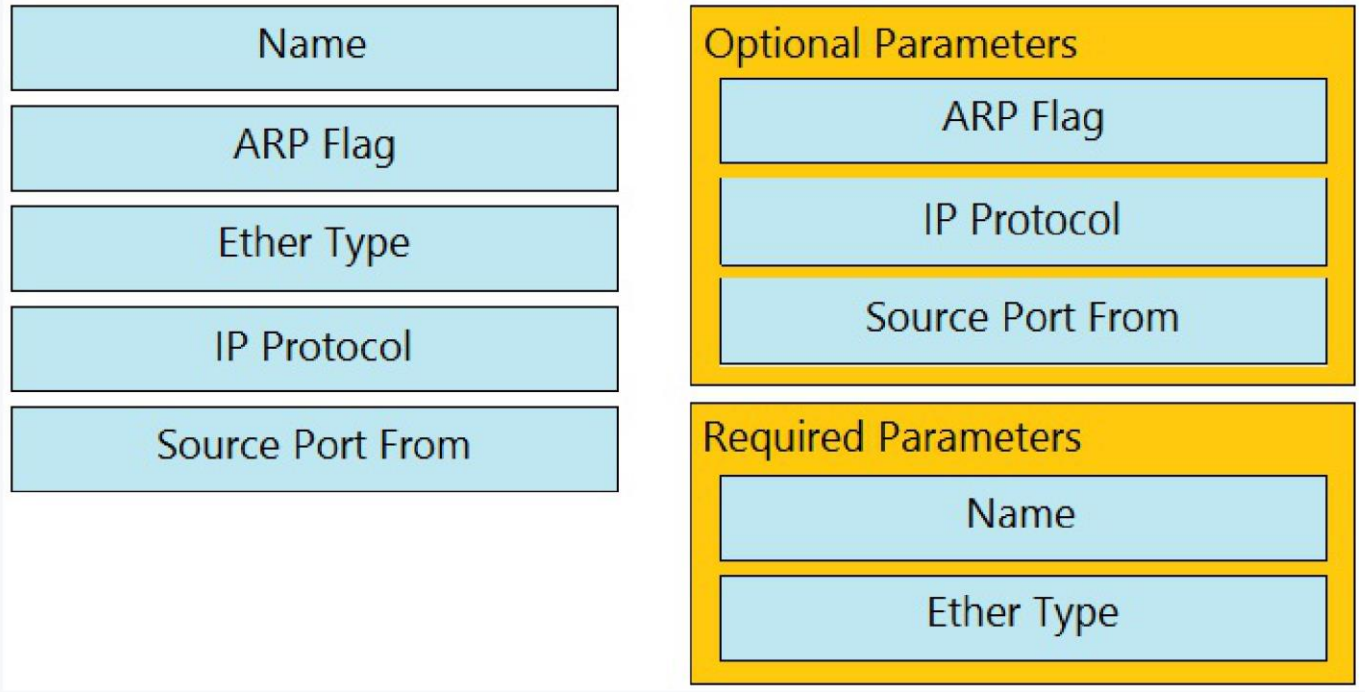

Explanation:

Reference:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/Operating_ACI/guide/b_Cisco_Operating_ACI/ b_Cisco_Operating_ACI_chapter_01000.html

Question: 15

Where is the COOP database located?

A.leaf
B.spine
C.APIC
D.endpoint

Answer: B

Explanation:

The correct answer is **B. spine**. The COOP (Council of Oracle Protocol) database, a crucial component in Cisco ACI (Application Centric Infrastructure), is housed on the spine switches. COOP is essentially a distributed database responsible for mapping endpoint identifiers (like MAC addresses and IP addresses) to their respective locations within the ACI fabric. Leaf switches learn about endpoints connected to them and register this information with the spine switches via the COOP protocol. The spines act as the central directory for endpoint location information. This architecture enables efficient forwarding and policy enforcement throughout the fabric. When a leaf switch needs to send traffic to a particular endpoint, it queries the spine's COOP database to determine the destination's location. This distributed, spine-centric approach to endpoint mapping ensures scalability and eliminates the need for each switch to hold a complete mapping of all endpoints. The APIC (Application Policy Infrastructure Controller) manages the configuration and policies within the ACI fabric, but it is not directly involved in the operational storage of the COOP database itself.

Endpoints, of course, are the network devices being managed within the ACI infrastructure, not components of the infrastructure, and therefore do not host the COOP database. Thus, the spine layer is the designated location for the COOP database storage and processing.

Authoritative Links for Further Research:

Cisco ACI Fundamentals:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/index.html (This is Cisco's general ACI page, you may need to navigate to specific documentation from here).

Understanding Cisco ACI Fabric Forwarding:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/fundamentals/b_ACI-Fundamentals/b_ACI-Fundamentals_chapter_010.html (Check this chapter on Forwarding).

Cisco ACI COOP protocol:https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b_kb-aci-coop-protocol.html (Cisco Knowledge base article on COOP protocol).

Question: 16

Which description regarding the initial APIC cluster discovery process is true?

A.The APIC uses an internal IP address from a pool to communicate with the nodes.
B.Every switch is assigned a unique AV by the APIC.
C.The APIC discovers the IP address of the other APIC controllers by using Cisco Discovery Protocol.
D.The ACI fabric is discovered starting with the spine switches.

Answer: A

Explanation:

The correct answer is **A: The APIC uses an internal IP address from a pool to communicate with the nodes.**

Here's why:

During the initial ACI fabric discovery, APICs (Application Policy Infrastructure Controllers) don't rely on Cisco Discovery Protocol (CDP) for inter-APIC communication or switch discovery. Instead, APICs form a cluster amongst themselves by electing a leader which facilitates the initial setup of the system. They establish communication using a private internal IP address space, typically from the 10.0.0.0/16 subnet, which is exclusively for control plane communication within the ACI fabric. This internal address space is not routable and should not conflict with any other network address space. This ensures a secure and isolated
communication channel for control plane operations. When an APIC is added to the cluster, it discovers its peers by utilizing multicast and a defined cluster membership protocol. This internal IP addressing scheme allows APICs to establish communication regardless of the external addressing or other network
configuration. APIC nodes leverage the out-of-band management network initially for bootstrapping. This mechanism enables the APIC to then establish internal communication within the ACI fabric. After cluster formation, a leader APIC emerges. Each switch in the fabric is assigned a unique node ID, not an AV, and these switches are discovered via the Fabric Discovery Protocol (FDP) which the APICs use. The ACI fabric discovery process begins with discovering the leaf switches first, rather than spine switches.

Therefore, option A accurately reflects the internal IP address usage during the APIC cluster discovery.

Authoritative Links:

Cisco ACI Fundamentals:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/fundamentals/b-cisco-aci-fundamentals.html
Cisco ACI Fabric Discovery Process:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/getting-started/b_APIC_Getting_Started_Guide/b_APIC_Getting_Started_Guide_chapter_0100.html#task_2838E19F8A48

Question: 17

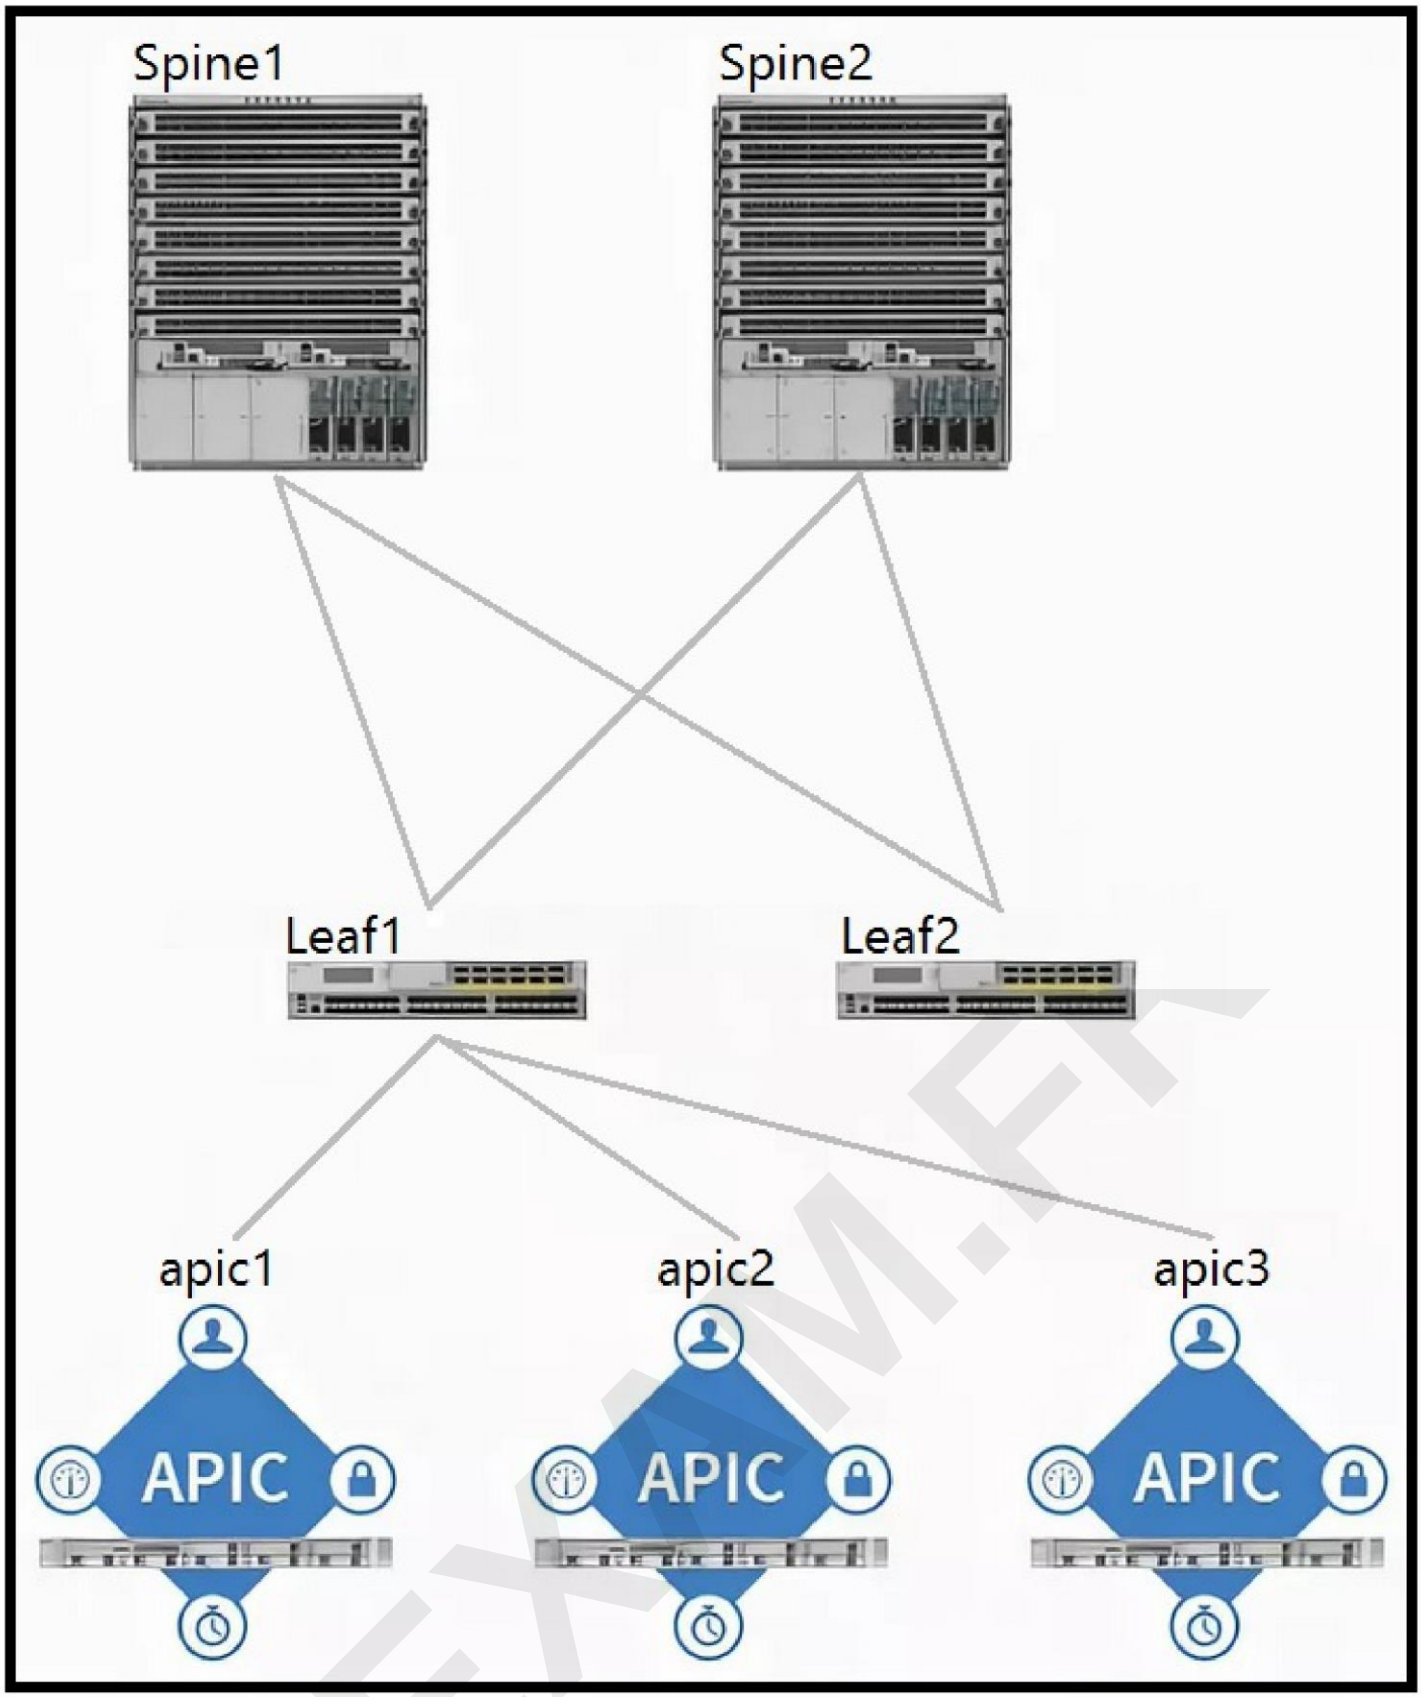

Refer to the exhibit. Which two components should be configured as route reflectors in the ACI fabric? (Choose two.)

A.Spine1
B.apic1
C.Spine2
D.Leaf1
E.Leaf2
F.apic2

Answer: AC

Explanation:

A.Spine1.

C.Spine2.

Question: 18

When creating a subnet within a bridge domain, which configuration option is used to specify the network visibility of the subnet?

A.limit IP learning to subnet
B.scope
C.gateway IP
D.subnet control

Answer: B

Explanation:

The correct answer is **B. scope**. In Cisco ACI, when configuring a subnet within a bridge domain, the "scope" parameter determines the network visibility of that subnet. This parameter dictates how the subnet is advertised and handled within the ACI fabric. The "scope" option allows you to define whether a subnet is private to the bridge domain, advertised to external networks, or both. Options include: "private to VRF," meaning the subnet is only visible within the associated Virtual Routing and Forwarding (VRF) instance; "shared between VRFs," making it visible across VRFs; or "advertised externally," which enables the subnet to be reachable beyond the ACI fabric through configured external Layer 3 connections.

Option A, "limit IP learning to subnet," does not directly control network visibility. Instead, it relates to how the fabric learns MAC and IP addresses associated with the subnet. Option C, "gateway IP," specifies the IP address of the default gateway for devices in that subnet but has no bearing on network visibility. Finally, "subnet control" (Option D) typically refers to granular access control lists applied to the subnet rather than its overall visibility scope. Therefore, only the "scope" option precisely defines the intended reach and accessibility of the created subnet. Configuring scope is crucial for network segmentation, security, and routing within the ACI environment. It dictates which endpoints can communicate with each other and how traffic is routed across the fabric and external connections.

For further reading and clarification, refer to Cisco's official documentation on ACI configuration, specifically regarding bridge domains and subnets:

Cisco ACI Bridge Domain Configuration Guide
Cisco ACI Fundamentals

Question: 19

What does a bridge domain represent?

A.Layer 3 cloud
B.Layer 2 forwarding construct
C.tenant

D.physical domain

Answer: B

Explanation:

A bridge domain in Cisco ACI represents a Layer 2 forwarding construct. It's crucial for understanding how network traffic is handled within an ACI fabric. Bridge domains essentially function as broadcast domains, similar to VLANs in traditional networks. They define the boundaries of Layer 2 flooding and learning, enabling communication between endpoints within the same logical segment. Unlike Layer 3, where routing decisions are made, bridge domains focus on MAC address learning and forwarding. Endpoints within a bridge domain can directly communicate using their MAC addresses, without involving routing. A single bridge domain can contain multiple subnets, but these subnets will share the same Layer 2 domain. The association of endpoints to a specific bridge domain is managed via endpoint groups (EPGs). These EPGs allow for the grouping of similar endpoints regardless of their physical location. In contrast, a physical domain refers to the physical infrastructure, a tenant is a container for managing ACI configurations, and a cloud, in the context of ACI, is usually where external connectivity resides, typically via Layer 3. Therefore, option B, Layer 2 forwarding construct, accurately describes the function and purpose of a bridge domain within the Cisco ACI framework.Here are some authoritative links for further research:

Cisco ACI Fundamentals:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/index.html
Cisco ACI Bridge Domain Concepts:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/configuration/guide/b_ACI_Config_Gu
Cisco ACI Documentation:https://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/products-technical-reference-list.html

Question: 20

Which table holds IP address, MAC address and VXLAN/VLAN information on a Cisco ACI leaf?

A.endpoint
B.adjacency
C.RIB
D.ARP

Answer: A

Explanation:

The correct answer is **A. endpoint**.

The endpoint table within a Cisco Application Centric Infrastructure (ACI) leaf switch is the central repository for mapping endpoint identifiers to their associated network attributes. This crucial table stores vital information including an endpoint's IP address, its corresponding MAC address, and the VLAN or VXLAN (Virtual Extensible LAN) identifier within which the endpoint resides. The leaf switch utilizes this table to make forwarding decisions, ensuring that traffic reaches the correct destination within the ACI fabric. Without the endpoint table, the leaf wouldn't be able to know how to deliver packets to connected devices, including physical servers and virtual machines. The endpoint database is built through a combination of data-plane learning (by observing source MAC addresses) and control-plane mechanisms like ARP (Address Resolution Protocol). In essence, the endpoint table acts as a dynamic directory for all connected endpoints on a leaf switch.

The adjacency table, option B, primarily concerns neighbor discovery and protocol adjacency relationships,

not the association of IP addresses and MAC addresses to VXLAN/VLANs. Option C, RIB (Routing Information Base), focuses on routing protocols, subnet information, and forwarding paths but doesn't hold the per-endpoint details. ARP, option D, is used for resolving IP addresses to MAC addresses, but does not
persistently store the complete endpoint information needed for the mapping to VLAN/VXLANs in the same way the endpoint table does. The endpoint table is specific to the ACI architecture and essential for the fabric's functionality.

For more in-depth information on Cisco ACI's endpoint table, you can refer to the following authoritative links:

Cisco ACI Fundamentals White Paper:https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-732200.html (Look for sections on Endpoint Learning and the Endpoint Table).

Cisco ACI Documentation:https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/whitepaper/aci-epg-contract.html (Focus on the endpoint learning and forwarding mechanisms.)
Cisco Live Sessions: Search for "Cisco ACI Endpoint Learning" on the Cisco Live On-Demand Library.

Question: 21

Which two types of interfaces are supported on border leaf switches to connect to an external router? (Choose two.)

A.subinterface with VXLAN tagging
B.subinterface with 802.1Q tagging
C.FEX host interface
D.out of band interface
E.Switch Virtual Interface

Answer: BE

Explanation:

The correct answer, B (subinterface with 802.1Q tagging) and E (Switch Virtual Interface), pertains to how border leaf switches in a Cisco Application Centric Infrastructure (ACI) fabric connect to external routers.
Border leaf switches act as the gateway between the ACI fabric and the external network. They need a mechanism to segregate traffic from different tenants or applications as it traverses the boundary.

802.1Q tagging, option B, is a standard protocol for VLAN trunking. Subinterfaces on a physical interface, each associated with a unique VLAN ID, allow multiple logical networks to share the same physical link. This is crucial for separating traffic from various endpoint groups (EPGs) in the ACI fabric as they exit toward the external network. The external router then uses these tags to route traffic accordingly to the correct destination networks.

A Switch Virtual Interface (SVI), option E, is a virtual interface representing a VLAN. In the ACI context, an SVI provides an L3 interface on the leaf switch for the specific subnet associated with a bridge domain (BD). The SVI serves as the default gateway for endpoints within that BD. When routing external traffic, the SVI is the interface used as the next hop. When configuring external routed connections, an SVI is created on the border leaf for the L3 outside network. This allows for Layer 3 routing between the ACI fabric and external routers.

Option A, “subinterface with VXLAN tagging,” is incorrect because VXLAN is primarily used for internal ACI fabric overlays, not for communication with external routers which would use 802.1q. Option C, “FEX host interface,” is incorrect as Fabric Extender (FEX) interfaces are primarily used for connecting access ports to end devices, not for external routing. Option D, "out of band interface", is used for switch management and not for data plane forwarding. Therefore, 802.1Q tagging on subinterfaces and SVIs are fundamental for establishing connections with external routers in Cisco ACI, enabling the border leaf switch to effectively

route traffic between the ACI fabric and the external networks.

Authoritative Links:

Cisco ACI Fundamentals:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/index.html
Cisco ACI Layer 3 Out:https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/l3-configuration/cisco-apic-layer3-configuration-5x.html
802.1Q Protocol:https://ieeexplore.ieee.org/document/927870

Question: 22

Which Cisco APIC configuration prevents a remote network that is not configured on the bridge domain from being learned by the fabric?

A.enable Limit IP Learning to Subnet
B.enable Unicast Routing
C.enable IP Data-plane Learning
D.enable ARP Flooding to BD

Answer: A

Explanation:

The correct answer is **A. enable Limit IP Learning to Subnet**.

Limiting IP learning to the subnet within a bridge domain (BD) is crucial for controlling which IP addresses are learned and advertised by the Cisco Application Centric Infrastructure (ACI) fabric. By enabling this feature, the fabric only learns IP addresses that fall within the configured subnets of the bridge domain. This prevents the learning and propagation of routes from remote networks that are not part of the defined bridge domain, thereby ensuring network segmentation and preventing unintended routing loops or reachability issues.

Options B, C, and D do not provide the same level of control over IP learning. Unicast routing (B) enables routing, it doesn't restrict IP learning. IP data-plane learning (C) allows the fabric to learn IPs based on data traffic but does not limit the learning to specific subnets. ARP flooding to BD (D) concerns only the ARP protocol handling within the bridge domain. Limiting IP learning to the subnet provides a more granular control over the learned IP address space, adhering to the principle of least privilege in network design. This practice ensures that the fabric only learns and advertises routes that are directly relevant to the specific bridge domain and its associated endpoints. Therefore, "enable Limit IP Learning to Subnet" directly prevents learning of IP addresses that are not within the configured subnets. This maintains controlled routing and avoids issues related to address space overlapping.

For further research, refer to the Cisco documentation on bridge domain configurations within ACI: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/basic-config/b_Cisco_APIC_Basic_Configuration_Guide/b_Cisco_APIC_Basic_Configuration_Guide_chapter_0111.html and search for "Limit IP Learning to Subnet" and specifically under the "Configure a Bridge Domain" subsection. You might also want to search on terms such as ACI L2 forwarding for more detail on how the BD works.

Question: 23

An engineer needs to deploy a leaf access port policy group in ACI Fabric to support the following requirements:☞ Control the amount of application data flowing into the system
☞ Allow the newly connected device to auto-negotiate link speed with the leaf switch

Which two ACI policies must be configured to achieve these requirements? (Choose two.)

A.link level policy
B.L2 interface policy
C.slow drain policy
D.ingress data plane policing policy
E.ingress control plane policing policy

Answer: AD

Explanation:

Let's break down why options A (link level policy) and D (ingress data plane policing policy) are the correct choices for this ACI deployment scenario.

Link Level Policy (A): This policy directly addresses the requirement of enabling auto-negotiation of link speed. A link level policy defines parameters such as speed and duplex for the physical interface. When configured for auto-negotiation, the connected device and the ACI leaf switch automatically negotiate the optimal link speed, ensuring compatibility and efficient data transfer. Without this policy, manual speed
configuration would be needed, potentially causing connectivity issues. Cisco's documentation details how link level policies control these physical interface parameters:
https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/policy/b_Cisco_APIC_Policy_Model/m_creating_link_level_po

Ingress Data Plane Policing Policy (D): This policy fulfills the requirement of controlling the amount of application data entering the system. Data plane policing policies limit the traffic flow rate (bandwidth) entering the ingress interface. By configuring a policing policy, we can prevent a single device or application from overwhelming the network. This mechanism is essential for maintaining network stability and ensuring that critical applications have adequate bandwidth. It is crucial for security and resource management. You can find more information on ACI policing policies here:
https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/configuration/2x/b-aci-config-guide-2x/b-aci-config-guide-2x_chapter_0101.html#concept_86B3B6E164674A24A9F5702B26D67055

Why other options are incorrect:

L2 interface policy (B): While L2 interface policies configure various Layer 2 parameters, they do not directly address rate limiting or auto-negotiation in the same way link-level and policing policies do.

Slow drain policy (C): A slow drain policy is used to prevent buffer overflow issues with slow-receiving endpoints and not for general rate-limiting or auto-negotiation.

Ingress control plane policing policy (E): Control plane policing protects against malicious traffic targeting the ACI control plane itself rather than the application data traffic.

In conclusion, the combination of a link level policy for auto-negotiation and an ingress data plane policing policy for traffic control effectively meets the outlined requirements for deploying the leaf access port policy group within the ACI fabric.

Question: 24

A customer migrates a legacy environment to Cisco ACI. A Layer 2 trunk is configured to interconnect the two environments. The customer also builds ACI fabric in an application-centric mode. Which feature should be enabled in the bridge domain to reduce instability during the migration?

A.Set Multi-Destination Flooding to Flood in BD.
B.Enable Flood in Encapsulation.
C.Set Multi-Destination Flooding to Flood in Encapsulation.

D.Disable Endpoint Dataplane Learning

Answer: C

Explanation:

The correct answer is **C. Set Multi-Destination Flooding to Flood in Encapsulation.** Here's a detailed justification:

During a migration from a legacy environment to Cisco ACI, especially when a Layer 2 trunk connects the two, a crucial challenge is handling broadcast, unknown unicast, and multicast (BUM) traffic. Initially, endpoints might exist in both environments, leading to potential instability if not handled properly. ACI learns endpoints based on source MAC addresses; however, during the migration phase, endpoints may move between ACI and the legacy network, potentially causing MAC address flapping.

By setting **Multi-Destination Flooding to Flood in Encapsulation**, the ACI fabric is instructed to flood BUM traffic out all ports associated with the bridge domain's encapsulation VLAN, including the trunk link to the legacy environment. This flooding mechanism allows legacy devices to receive broadcasts and discover endpoints located in the ACI fabric, even if ACI hasn't yet learned their location. This reduces potential MAC flapping and aids smoother endpoint mobility during the migration as the legacy network slowly relinquishes ownership of certain MAC addresses.

Options A, B, and D are not the best fit for reducing instability during this migration phase:

A. Set Multi-Destination Flooding to Flood in BD: While this option does involve flooding BUM traffic, it floods across the entire bridge domain. This behavior is less desirable during migration because flooding is not limited to the specific encapsulation used for interconnecting with the legacy environment, leading to potentially wider and unnecessary impact during migration.

B. Enable Flood in Encapsulation: This option is too broad. Enabling flood in encapsulation means flooding of all encapsulation traffic which is not practical as it floods traffic that would be unnecessary.

D. Disable Endpoint Dataplane Learning: Disabling learning would severely hinder ACI's ability to effectively manage endpoint information and would create operational and troubleshooting difficulties. This is contrary to the goal of a smooth transition.

Therefore, setting **Multi-Destination Flooding to Flood in Encapsulation** provides a targeted way to ensure legacy devices can discover endpoints in the ACI fabric and prevents potential instability during the migration of endpoints between the legacy network and ACI.

Authoritative Links for further research:

Cisco ACI Bridge Domain Configuration:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/configuration/guide/b_ACI_Config_Gu
Cisco ACI Multi-Destination Flooding:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/configuration/guide/b_ACI_Config_Gu
(Specifically, review the section relating to "Multi-Destination Flooding")

Question: 25

New ESXi hosts are procured in a data center compute expansion project. An engineer must update the configuration on the Cisco APIC controllers to support the addition of the new servers to the existing VMM domain. Which action should be taken to support this change?

A.Create a range of internal VLANs in the associated VLAN pool.

B.Set the encapsulation mode as VXLAN.

C.Enable infrastructure VLAN in the associated AEP.

D.Map the leaf interface selector to the AEP that is associated with the VMM domain.

Answer: D

Explanation:

The correct answer is **D. Map the leaf interface selector to the AEP that is associated with the VMM domain.**

Here's why:

In Cisco ACI, integrating ESXi hosts into a VMM domain involves associating physical infrastructure (leaf interfaces) with a logical representation of the environment (the VMM domain). An Attachable Entity Profile (AEP) acts as a bridge, defining policies and settings applicable to endpoints connected to specific interfaces. The VMM domain, in turn, utilizes the AEP to interact with the connected infrastructure. When new ESXi hosts are added, their connected interfaces on the ACI fabric (leaf switches) must be explicitly linked to the existing AEP associated with the VMM domain.

Mapping the leaf interface selector to the AEP establishes this connection, informing the ACI fabric which interfaces belong to the specific VMM domain. This action ensures that the required VLANs, policies, and configurations tied to the VMM domain are extended to the new hosts. Without this mapping, the new ESXi hosts cannot participate in the VMM domain and receive necessary network connectivity.

Options A, B, and C are not correct in this scenario:

A. Create a range of internal VLANs in the associated VLAN pool: While VLAN pools are essential for segmentation, adding new VLANs isn't necessary when expanding existing infrastructure and you're using the same policies; the existing VLAN pool should suffice.

B. Set the encapsulation mode as VXLAN: VXLAN is not mandatory for all VMM integrations; it depends on the specific requirements. The encapsulation mode setting is usually a one-time decision for a given environment and generally does not need modification in this situation.

C. Enable infrastructure VLAN in the associated AEP: Infrastructure VLANs are essential for the management plane within ACI, but this setting isn't directly tied to the integration of new hypervisor hosts into an existing VMM domain.

In essence, the leaf interface mapping is the key operation to integrate new ESXi servers into an existing ACI-controlled VMware environment by attaching the physical infrastructure to the logical VMM domain.

Authoritative links for further research:

Cisco ACI Fundamentals:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/fundamentals/b_ACI_Fundamentals.html (This is a high-level overview; look for the specific topics mentioned below.)
Attachable Access Entity Profile (AEP): Search within the above link, for example search for "Attaching Entity Profile AEP" .
Cisco ACI VMM Integration: Search within the Cisco documentation for 'VMM integration', for example:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/vmware/b_Cisco_APIC_VMware_Integration_Guide/b_Cisco_APIC_VMware_Integration_Guide_chapter_010.html

Question: 26

DRAG DROP -

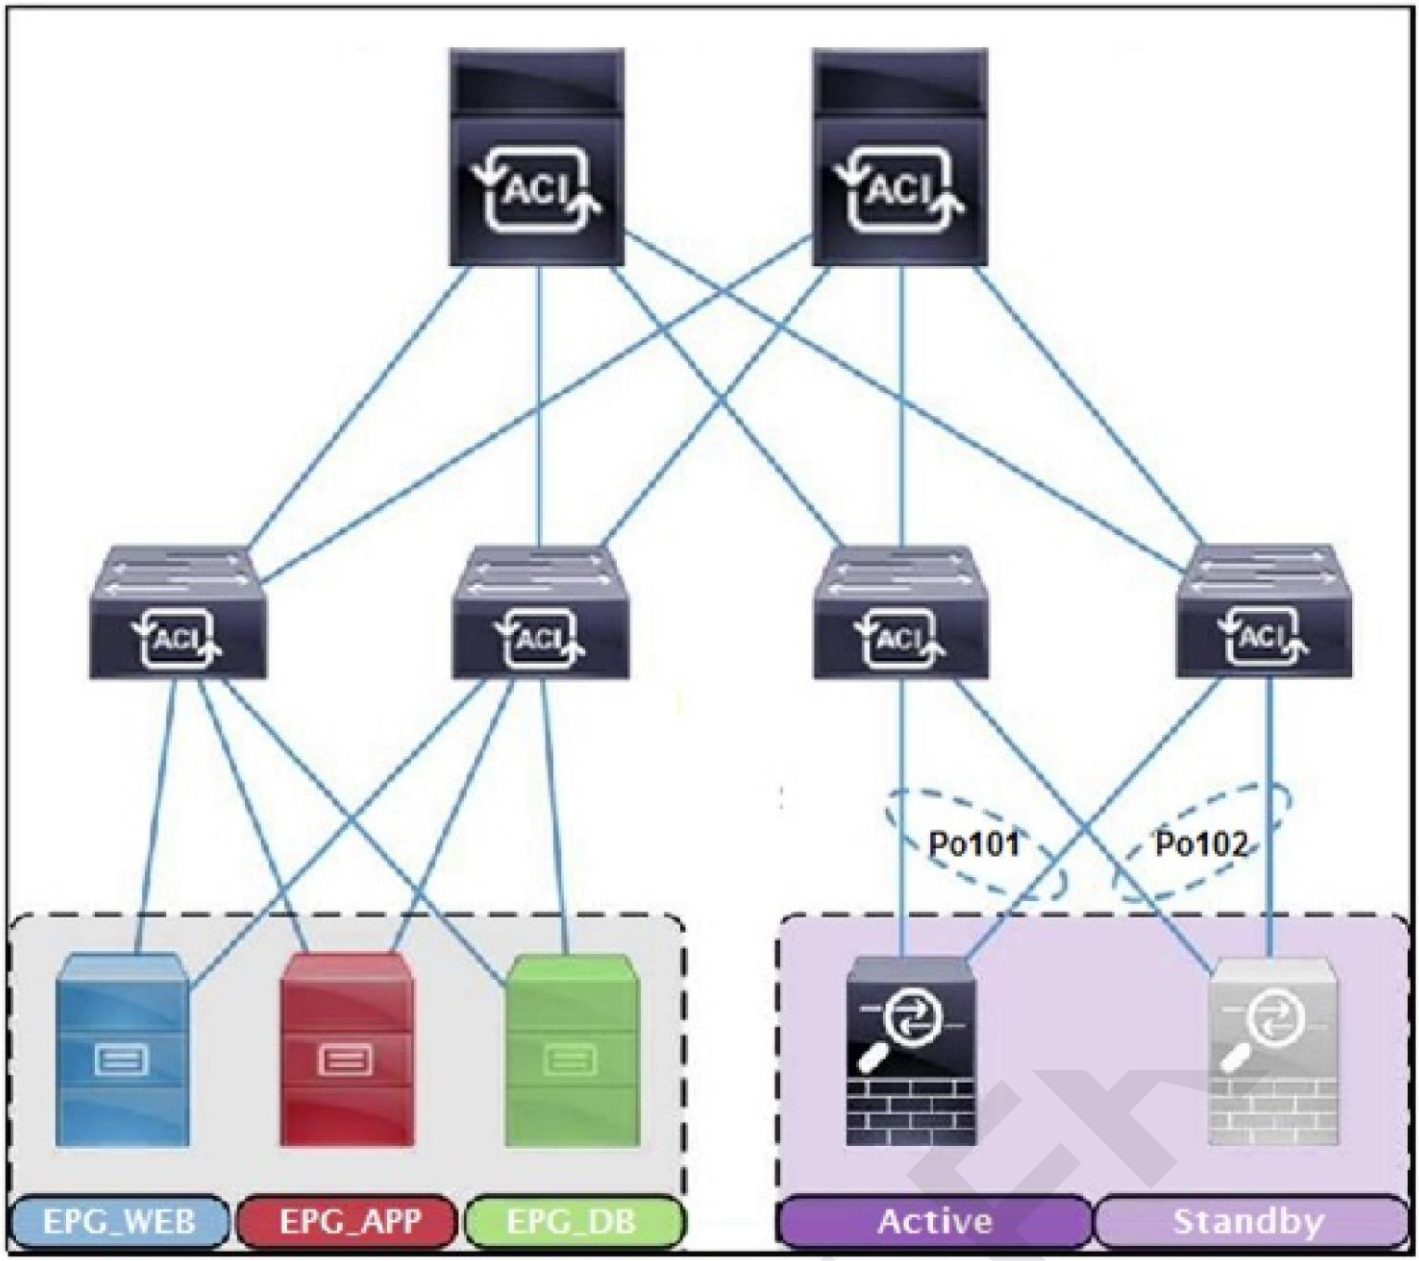

Refer to the exhibit. A Cisco ACI fabric is newly deployed, and the security team requires more visibility of all inter EPG traffic flows. All traffic in a VRF must be forwarded to an existing firewall pair. During failover, the standby firewall must continue to use the same IP and MAC as the primary firewall. Drag and drop the steps from the left into the implementation order on the right to configure the service graph that meets the requirements. (Not all steps are used.)
Select and Place:

Answer Area

Apply a service graph template and select vzAny EPG as the consumer and provider.	Step 1
Select a redirect policy with the Layer 3 destination.	Step 2
Create a Layer 4 to Layer 7 service graph template.	Step 3
Select a redirect policy with enabled anycast and the Layer 3 destination.	Step 4
Select the same cluster interface under Consumer Connector and Provider Connector.	Step 5
Create a service bridge domain and a Layer 4 to Layer 7 device with one cluster interface.	Step 6
Select the existing contract with custom IP EtherType filter.	

Answer:

Answer Area

Apply a service graph template and select vzAny EPG as the consumer and provider.	Create a service bridge domain and a Layer 4 to Layer 7 device with one cluster interface.
Select a redirect policy with the Layer 3 destination.	Create a Layer 4 to Layer 7 service graph template.
Create a Layer 4 to Layer 7 service graph template.	Apply a service graph template and select vzAny EPG as the consumer and provider.
Select a redirect policy with enabled anycast and the Layer 3 destination.	Select the existing contract with custom IP EtherType filter.
Select the same cluster interface under Consumer Connector and Provider Connector.	Select a redirect policy with enabled anycast and the Layer 3 destination.
Create a service bridge domain and a Layer 4 to Layer 7 device with one cluster interface.	Select the same cluster interface under Consumer Connector and Provider Connector.
Select the existing contract with custom IP EtherType filter.	

Explanation:

1.Create a service bridge domain and a Layer 4 to Layer 7 device with on cluster interface.

2.Create a Layer 4 to Layer 7 service graph template .

3.Apply a service a graph template and select vzAny EPG as the consumer and provider.

4.Select the existing contract with customer IP Ether Type filter .

5.Select a redirect policy with enabled any cast and the Layer 3 destination .

6.Select the same cluster interface under Consumer Connector and Provider connector .

Question: 27

An engineer is extending an EPG out of the ACI fabric using static path binding. Which statement about the endpoints is true?

A.Endpoints must connect directly to the ACI leaf port.
B.External endpoints are in a different bridge domain than the endpoints in the fabric.
C.Endpoint learning encompasses the MAC address only.
D.External endpoints are in the same EPG as the directly attached endpoints.

Answer: D

Explanation:

Here's a detailed justification for why option D is the correct answer:

When extending an Endpoint Group (EPG) outside the ACI fabric using static path binding, the fundamental concept is that you're essentially creating a logical extension of that EPG. This means that the external endpoints, despite not being directly connected to the ACI leaf, are still considered members of the same EPG. They inherit the policy and configurations defined for that EPG. Option D accurately reflects this by stating that external endpoints are in the same EPG as directly attached endpoints.

Option A is incorrect because static path binding allows external devices to connect through an intermediary device. They are not required to directly connect to the leaf. Option B is incorrect because a core feature of EPG extension is that all endpoints, regardless of attachment location, exist within the same bridge domain. They share layer 2 connectivity. Option C is inaccurate because endpoint learning in ACI encompasses more than just MAC addresses, often including IP addresses (especially when contracts are using IP based filters).

Static path binding in ACI allows for scenarios where external endpoints might be connected to a router or switch outside the fabric, and the ACI fabric uses configuration to treat those endpoints as part of the same logical group. This ensures consistent policy enforcement, security, and management across all endpoints within that EPG, whether physically inside or outside the ACI fabric. This is a core element of how ACI facilitates a consistent policy plane across extended networks.

Here are some links for further research on ACI EPGs and static path binding:

Cisco Application Centric Infrastructure Fundamentals:
https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/fundamentals/cisco-aci-fundamentals-5x-chapter_010.html
Cisco ACI Static Binding:https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b-aci-static-binding-kb.html
Cisco ACI Endpoint Groups:https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/configuration/l2-configuration/cisco-aci-l2-configuration-5x-chapter_0100.html

Question: 28

Which setting prevents the learning of Endpoint IP addresses whose subnet does not match the bridge domain subnet?

A."Limit IP learning to network" setting within the bridge domain.
B."Limit IP learning to subnet" setting within the EPG.
C."Limit IP learning to network" setting within the EPG.
D."Limit IP learning to subnet" setting within the bridge domain.

Answer: D

Explanation:

The correct answer is **D. "Limit IP learning to subnet" setting within the bridge domain.**

Here's why: Cisco ACI's bridge domains define a Layer 2 forwarding domain. They dictate how traffic is forwarded within a subnet. When an endpoint (a server, VM, etc.) sends traffic, ACI learns the endpoint's IP and MAC address and associates them with the port where the traffic entered. The "Limit IP learning to subnet" setting, configured within the bridge domain, directly impacts how this IP learning occurs.

By enabling "Limit IP learning to subnet," the bridge domain restricts the learning of endpoint IPs to only those that fall within the defined subnet for that bridge domain. If an endpoint sends traffic with an IP outside the configured subnet, its IP address will not be learned. This is crucial for maintaining subnet isolation and preventing accidental communication between different subnets within the same bridge domain. This prevents issues like traffic leakage and unintended connectivity. This setting is a safeguard against misconfigurations or potentially rogue devices. Options A and C concern network level limit of IP learning and they are not specific to subnet based isolation. Option B concerns EPG specific restrictions and not Bridge domain. EPGs are containers for endpoints that share the same policy profile, which is different from the IP address learning function of the bridge domain. Disabling this setting would allow IP addresses from other subnets to be learned on this bridge domain. This may lead to unpredictable connectivity issues, because traffic intended for other subnets might inadvertently be routed within the original bridge domain.

Authoritative Links:

Cisco ACI Bridge Domain Configuration:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/configuration/guide/b_Cisco_APIC_Co bridge-domains.html - Specifically, search for the "Limit IP Learning to Subnet" section for detailed explanation.

Cisco ACI Fundamentals:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/index.html - To gain an overall understanding of Cisco ACI architecture.

Question: 29

Which endpoint learning operation is completed on the egress leaf switch when traffic is received from an L3Out?

A.The source MAC and IP address of the traffic is learned as a local endpoint.
B.The source MAC address of the traffic is learned as a remote endpoint.
C.No source MAC or IP address of the traffic is learned as a remote endpoint.
D.The source IP address of the traffic is learned as a remote endpoint.

Answer: C

Explanation:

The correct answer is **C. No source MAC or IP address of the traffic is learned as a remote endpoint.** Here's why:

When traffic enters an ACI fabric from an L3Out (external Layer 3 network), the egress leaf switch handles the packet leaving the ACI fabric. Endpoint learning within ACI primarily focuses on tracking endpoints within the fabric (connected to leaf ports). When dealing with L3Out traffic, the source is external and often doesn't have visibility within the ACI fabric's endpoint database. The ACI fabric only learns the IP addresses when an IP-based forwarding is being used. The endpoint information such as the MAC addresses is not stored in the ACI fabric database when the traffic enters the ACI fabric through the L3Out. In essence, traffic from an L3Out is treated as a routed packet at the edge of the fabric, and the ACI fabric is not designed to learn and track external endpoints in its internal database. Source endpoints reachable via L3Out are considered part of the external routed network, and ACI does not need to track them like internal endpoints. The egress leaf primarily performs routing based on the destination, using the routing table. The ACI fabric doesn't learn about the specific MAC address of external endpoints connected to an L3Out because it's not within its scope of endpoint tracking. This differs from internal traffic where endpoints are explicitly learned and tracked for efficient local forwarding within the fabric. Therefore, neither source MAC nor IP address of traffic originating from an L3Out is recorded as a remote endpoint within the fabric's endpoint database.

Further research:

Cisco ACI Fundamentals:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/index.html (This provides general information on ACI architecture)
Cisco ACI L3Out Documentation: Search Cisco documentation for "ACI L3Out" for specific details on external connectivity behavior.
ACI Endpoint Learning: Refer to ACI documentation sections on how endpoints are learned within the fabric.

Question: 30

Refer to the exhibit. Which two objects are created as a result of the configuration? (Choose two.)

A.application profile
B.attachable AEP
C.bridge domain
D.endpoint group
E.VRF

Answer: CE

Explanation:

C.bridge domain.

E.VRF.

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/rest_cfg/2_1_x/b_Cisco_APIC_REST_API_Configuration_Guide/b_Cisco_APIC_REST_API_Configuration_Guide_chapter_01110.html

Question: 31

What must be enabled in the bridge domain to have the endpoint table learn the IP addresses of endpoints?

A.L2 unknown unicast: flood
B.GARP based detection
C.unicast routing
D.subnet scope

Answer: C

Explanation:

The correct answer is C, unicast routing. In Cisco ACI, bridge domains primarily function at Layer 2, handling MAC address learning and forwarding. However, to enable the endpoint table to associate IP addresses with MAC addresses, a bridge domain must have unicast routing enabled. This feature essentially turns the bridge domain into a routed Layer 3 entity, allowing it to perform IP address learning. Without unicast routing, the endpoint table only tracks MAC addresses. Specifically, when unicast routing is enabled, the bridge domain becomes associated with an IP subnet. As endpoints communicate, their IP and MAC addresses are learned and added to the endpoint table. Options A, L2 unknown unicast flooding and B, GARP based detection, are related to MAC address learning, not IP address learning. Option D, subnet scope, is related but does not directly enable IP address learning without unicast routing enabled. Unicast routing transforms the bridge domain's role and is crucial for IP-based forwarding and policy enforcement in ACI environments.

For further information, you can refer to the following Cisco documentation:

Cisco ACI Fabric L2 Bridge Domains
Cisco ACI Configuration Guide
Cisco ACI White Paper

Question: 32

An engineer is extending EPG connectivity to an external network. The external network houses the Layer 3 gateway and other end hosts. Which ACI bridge domain configuration should be used?

A.Forwarding: Custom L2 Unknown Unicast: Hardware Proxy L3 Unknown Multicast Flooding: Flood Multi Destination Flooding: Flood in BD ARP Flooding: Enabled

B.Forwarding: Custom L2 Unknown Unicast: Flood L3 Unknown Multicast Flooding: Flood Multi Destination Flooding: Flood in BD ARP Flooding: Enabled

C.Forwarding: Custom L2 Unknown Unicast: Hardware Proxy L3 Unknown Multicast Flooding: Flood Multi Destination Flooding: Flood in BD ARP Flooding: Disabled

D.Forwarding: Custom L2 Unknown Unicast: Flood L3 Unknown Multicast Flooding: Flood Multi Destination Flooding: Flood in BD ARP Flooding: Disabled

Answer: B

Explanation:

The correct answer is B because it provides the necessary flooding configuration for optimal communication between the ACI fabric and the external network. When extending EPG connectivity to an external Layer 3 network, the ACI bridge domain needs to handle unknown unicast, multicast, and ARP traffic appropriately.

Option B's "L2 Unknown Unicast: Flood" ensures that when a destination MAC address is unknown within the bridge domain, the traffic is flooded across all ports in that domain. This is essential for initial learning and reaching hosts that are not yet known by the ACI fabric. "L3 Unknown Multicast Flooding: Flood" enables the bridge domain to forward multicast traffic when it doesn't have specific multicast routes. "Multi Destination Flooding: Flood in BD" ensures traffic destined to multiple unknown end-points within the bridge domain is properly flooded. Finally, "ARP Flooding: Enabled" is crucial for the ACI fabric to discover IP addresses of hosts within the external network and learn their MAC addresses to avoid having to flood for every packet.

This enables both ACI and external devices to learn each other's MAC addresses, therefore enabling communication.

Option A's "L2 Unknown Unicast: Hardware Proxy" would require the spine switches to act as proxies for unknown unicast traffic, which is inefficient for external connectivity where it's expected the bridge domain to learn them. Option C and D incorrectly disable "ARP Flooding", which is required for address discovery within the external Layer 3 network. If ARP flooding is disabled, devices on the external network and the ACI network would not discover each other, which would cause communication problems.

In summary, option B's configuration is standard and necessary for proper operation when connecting to an external Layer 3 network because it allows the ACI bridge domain to learn and forward traffic for communication with devices on the external network.

Further Research:

Cisco ACI Bridge Domains:https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/l2-configuration/cisco-apic-layer2-configuration-guide-5x/m-aci-bd.html
Cisco ACI L3 Out:https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/l3-configuration/cisco-apic-layer3-configuration-guide-5x/m-l3-out.html

Question: 33

An engineer configured a bridge domain with the hardware-proxy option for Layer 2 unknown unicast traffic. Which statement is true about this configuration?

A.The leaf switch drops the Layer 2 unknown unicast packet if it is unable to find the MAC address in the local forwarding tables.

B.The Layer 2 unknown hardware proxy lacks support of the topology change notification.

C.The leaf switch forwards the Layers 2 unknown unicast packets to all other leaf switches if it is unable to find the MAC address in its local forwarding tables.

D.The spine switch drops the Layer 2 unknown unicast packet if it is unable to find the MAC address in the proxy database.

Answer: D

Explanation:

Here's a detailed justification for why option D is the correct answer, along with supporting concepts and links:

Justification:

When a bridge domain in Cisco ACI is configured with the "hardware-proxy" option for Layer 2 unknown unicast traffic, the behavior of forwarding such packets changes significantly. Instead of leaf switches directly flooding these packets to all other leaves, they are first sent to the spine switches. The spine switches then maintain a "proxy database" that stores MAC addresses learned throughout the ACI fabric. This database acts as a centralized lookup for these packets.

Option D correctly states that "the spine switch drops the Layer 2 unknown unicast packet if it is unable to find the MAC address in the proxy database." This accurately reflects the behavior of the hardware-proxy feature. When a leaf switch receives an unknown unicast frame, it sends it to the spine. If the spine doesn't have an entry for the destination MAC in its proxy database, it means the destination is not known within the fabric, and thus the spine drops the packet to avoid uncontrolled flooding. This centralized decision-making reduces unnecessary flooding across the ACI fabric, enhancing efficiency and scalability.

Options A, B, and C are incorrect because:

Option A: Leaf switches do not drop the packet if the MAC isn't in their local table when hardware proxy is enabled; instead, they forward the packet to the spine switches.

Option B: The hardware proxy does support topology change notifications. The spines maintain the proxy database, which is updated with the latest changes in the fabric.

Option C: As mentioned, the leaf switches send the traffic towards the spine rather than directly flooding to all leaves.

Cloud Computing Concepts:

The hardware proxy mechanism in ACI embodies several important cloud computing concepts:

Centralized Control: The spine switches manage the proxy database and make forwarding decisions based on this centralized information. This aligns with the principle of centralized management for efficiency and consistency.

Scalability: By reducing the uncontrolled flooding of unknown unicast traffic, this approach improves the scalability of the network. It prevents unnecessary load on leaf switches.

Efficient Resource Utilization: The spine switches act as a gatekeeper for unknown unicast traffic, thus reducing the strain on leaf switches. This contributes to efficient resource utilization.

Control Plane Optimization: Centralizing the MAC address learning and forwarding decision in the spines optimizes the control plane operations and results in a more scalable and efficient network.

Authoritative Links:

Cisco ACI Layer 2 Forwarding: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b_Cisco_ACI_Layer_2_Forwarding.htm **Cisco ACI Fundamentals:**https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/fundamentals/b_ACI_Fundamentals_2_x.html (Look for sections on Layer 2 and proxy behavior)
Cisco ACI White Paper:https://www.cisco.com/c/dam/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-733060.pdf (This whitepaper provides a comprehensive overview of ACI architecture and functions)

By using the hardware proxy, the system moves to a more controlled and scalable approach to handle unknown traffic, enhancing network efficiency. Therefore, **D** is the correct answer.

Question: 34

An engineer configured Layer 2 extension from the ACI fabric and changed the Layer 2 unknown unicast policy from Flood to Hardware Proxy. How does this change affect the flooding of the L2 unknown unicast traffic?

A.It is forwarded to one of the spines to perform as a spine proxy.

B.It is flooded within the whole fabric.

C.It is dropped by the leaf when the destination endpoint is not present in the endpoint table.

D.It is forwarded to one of the APICs to perform as a proxy.

Answer: A

Explanation:

Okay, let's break down why the answer is A.

When Layer 2 unknown unicast policy in Cisco ACI is changed from "Flood" to "Hardware Proxy", the behavior for handling unknown unicast traffic shifts significantly. Previously, with "Flood", the traffic would be sent to all leaf ports within the bridge domain (essentially mimicking a traditional switch flooding behavior). However, with "Hardware Proxy", the leaf switch encountering an unknown destination MAC address no longer floods.

Instead, it sends this traffic to a designated spine switch. This spine switch then acts as a "hardware proxy," holding a more comprehensive view of MAC addresses across the entire ACI fabric. The spine uses its endpoint table to determine the destination location. If the destination is known, the spine proxy forwards the traffic; if not, the spine may either drop the traffic or forward it based on the remaining configuration. This approach significantly reduces unnecessary flooding within the fabric and optimizes traffic flow. Options B and C are therefore incorrect: B describes the behavior of the "flood" policy, not "hardware proxy" and C explains what would happen if the traffic is dropped due to the destination MAC not being in the EP table, and doesn't explain how the proxying works. Option D is also incorrect as it is a spine switch, not an APIC, that performs the proxy role.

In summary, the change to "Hardware Proxy" centralizes the lookup process at the spine, enabling more intelligent forwarding decisions and reducing broadcast traffic within the ACI environment. This ultimately improves performance, efficiency, and scalability.

Authoritative Links for Further Research:

Cisco ACI Layer 2 Forwarding:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b-aci-layer-2-forwarding.html
Cisco ACI White Paper - Understanding Layer 2 forwarding:
https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-736500.html
Cisco ACI L2 Unknown Unicast Policy Configuration Guide: (Search on Cisco website)

Question: 35

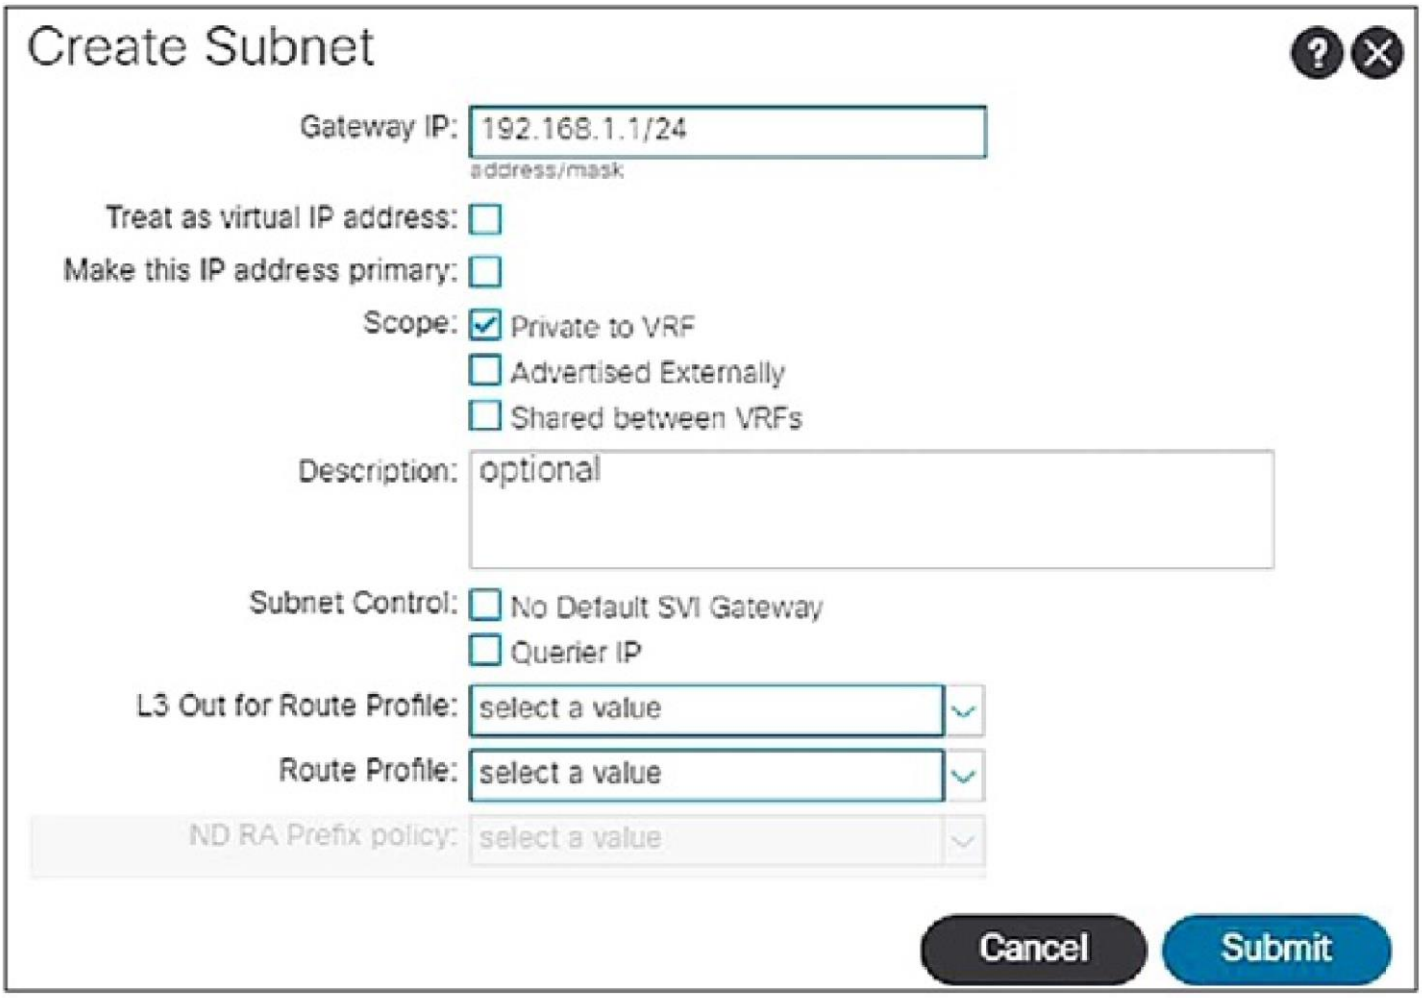

When the subnet is configured on a bridge domain, on which physical devices is the gateway IP address configured?

A.all leaf switches and all spine nodes
B.only spine switches where the bridge domain of the tenant is present C.only leaf switches where the bridge domain of the tenant is present D.all border leaf nodes where the bridge domain of the tenant is present

Answer: C

Explanation:

only leaf switches where the bridge domain of the tenant is present.

Reference:

http://www.netdesignarena.com/index.php/2016/06/16/aci-tenant-building-blocks-forwarding-logic/

Question: 36

Which method does the Cisco ACI fabric use to load-balance multidestination traffic?

A.forwarding tag trees
B.PIM routing
C.spanning trees
D.shortest-path trees

Answer: A

Explanation:

The correct answer is **A. forwarding tag trees**. Cisco ACI utilizes a unique method of load balancing called forwarding tag trees (F-Tags) for multidestination traffic. Unlike traditional Layer 2 or Layer 3 forwarding mechanisms that rely on MAC addresses or IP addresses, ACI assigns a unique F-Tag to each endpoint. When a packet destined for multiple endpoints needs to be sent, the ACI fabric computes a forwarding tag tree based on the destinations. Each leaf switch in the ACI fabric holds a portion of the forwarding information for the F-Tag and only delivers the traffic to the endpoints it's directly connected to or its downstream leaf switches. This approach allows for efficient multicast and broadcast forwarding, eliminating the need for traditional spanning tree protocol (STP) and minimizing traffic flooding. Spanning trees (C) are intentionally avoided in ACI due to their inherent inefficiencies and potential for blocking paths; shortest-path trees (D) while used in routing, aren't the primary method for multidestination within ACI. PIM routing (B) is also not the primary forwarding method, although it can be used for specific multicast scenarios on top of the fabric's forwarding mechanisms. F-Tag implementation ensures optimal path usage and eliminates broadcast storms within the ACI fabric.

Authoritative Links for further research:

Cisco ACI Multi-Destination Forwarding:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/configuration/guide/b-cisco-apic-configuration-guide/m-multi-destination-forwarding.html
Cisco ACI Fabric Forwarding:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/white-paper-c11-733445.html (Look for sections on forwarding)

Question: 37

What happens to the traffic flow when the Cisco ACI fabric has a stale endpoint entry for the destination endpoint?

A.The leaf switch does not learn the source endpoint through data plane learning.
B.The leaf switch drops the traffic that is destined to the endpoint.
C.The leaf switch floods the traffic to the endpoint throughout the fabric.
D.The leaf switch sends the traffic to the wrong destination leaf.

Answer: D

Explanation:

Okay, let's break down why option D is the correct answer and why the others aren't, focusing on how Cisco ACI handles endpoint learning and traffic forwarding.

In Cisco ACI, when a leaf switch receives traffic destined for an endpoint, it first consults its local endpoint table (also known as an endpoint database or EPG endpoint table). This table maps endpoint identifiers (like MAC addresses or IP addresses) to their associated leaf switch location. If the table has a stale entry, it means the recorded location of the destination endpoint is outdated.

When a stale entry exists, the leaf switch incorrectly uses this outdated information. Specifically, it will forward the traffic based on the stale location data. This means the traffic will be sent to the wrong leaf switch, as per answer option D. It's vital to understand that ACI uses a centralized control plane (APIC) to maintain the truth for endpoints, and individual leaves use this cached information.

Option A is incorrect because ACI leaf switches do learn source endpoints, regardless of destination endpoint table issues. They'll learn the source through data plane learning. Option B is incorrect because traffic is not

dropped in this scenario, it is misdirected, rather than dropped. Option C is incorrect because ACI doesn't flood traffic unless there's an unknown destination endpoint. The leaf switch thinks it knows where the destination is, even if it is wrong.

Therefore, a stale entry in the endpoint table leads to misforwarding rather than dropping, not learning, or flooding. The key issue is that the cached endpoint information on the leaf doesn't reflect the actual current location, leading to incorrect delivery.

Authoritative Links for Further Research:

1. **Cisco ACI Fundamentals:**https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/index.html This link takes you to the Cisco ACI main page, which provides a lot of resources.
2. **Cisco ACI Endpoint Learning:** While there isn't one specific document covering stale endpoint entries explicitly, many Cisco documents about ACI delve into the mechanisms for endpoint learning and management. Refer to their documentation on the ACI APIC and leaf switch interaction related to endpoint discovery, specifically look for "endpoint database", "endpoint group (EPG)" and "leaf switch forwarding."

In summary, a stale endpoint entry in a Cisco ACI fabric leads to traffic being forwarded to the incorrect leaf due to outdated location information. The traffic is not dropped or flooded, and the source endpoint learning process is not impeded by this. The system relies on a central truth from the APIC, and inconsistency leads to misdirection, and the leaf will use its cached information in the absence of the central truth.

Question: 38

Which action sets Layer 2 loop migration in an ACI Fabric with a Layer 2 Out configured?

A.Enable MCP on the ACI fabric.

B.Disable STP in the external network.

C.Disable STP on the ACI fabric.

D.Enable STP on the ACI fabric.

Answer: A

Explanation:

Here's a detailed justification for why the correct answer is A: Enabling MCP (Mis-Cabling Protocol) is the action that sets Layer 2 loop mitigation in an ACI fabric with a Layer 2 Out configured. An ACI fabric, by design, does not run traditional Spanning Tree Protocol (STP) internally. Instead, it employs a loop-free architecture. When connecting an external Layer 2 network to an ACI fabric using a Layer 2 Out, the potential for loops arises at the boundary where the STP-based external network meets the ACI's non-STP
environment. MCP is specifically designed to handle this situation. If loops occur due to misconfiguration or other issues in the external network, MCP detects these loops through the exchange of special control packets. Upon detection, MCP will automatically block specific ports at the ACI boundary to mitigate the loop and prevent network disruption. This is crucial for maintaining network stability when integrating with traditional Layer 2 infrastructure. Disabling STP, whether on the ACI fabric (which doesn't run it) or the external network, would exacerbate the loop risk instead of mitigating it. Enabling STP on the ACI fabric is not an option, as it's not how ACI handles loops internally. Therefore, MCP is the dedicated mechanism for L2 loop protection in ACI when integrating with external networks via L2 outs.

Here are some authoritative links for further research on ACI and MCP:

Cisco's official documentation on MCP:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b_Cisco_ACI_Mis-Cabling_Protocol.html
Cisco ACI Fundamentals:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/index.html
Cisco Live presentation on ACI L2 Out: (Search Cisco Live archives for relevant sessions)

Question: 39

An engineer is implementing a connection that represents an external bridged network. Which two configurations are used? (Choose two.)

A.Layer 2 remote fabric
B.Layer 2 outside
C.Layers 2 internal
D.Static path binding
E.VXLAN outside

Answer: BD

Explanation:

Here's a detailed justification for why options B (Layer 2 outside) and D (Static path binding) are the correct configurations for implementing a connection representing an external bridged network in Cisco ACI:

When integrating external Layer 2 networks with ACI, the ACI fabric needs to be aware of these external networks to allow communication. Option B, "Layer 2 outside," directly refers to the configuration required to define these external networks within the ACI fabric. An "outside" network signifies a network that resides outside the ACI domain, typically connected to a Leaf switch. By defining the Layer 2 outside connection, the fabric recognizes that traffic with a specific VLAN or VLAN range is associated with an external bridge domain. This configuration is crucial for the fabric to appropriately forward traffic coming from and going to the external Layer 2 network.

Option D, "Static path binding," is then needed to connect this "Layer 2 outside" to a physical interface on an ACI leaf switch. Since this is a bridged (Layer 2) connection, we need to specify which port or interface the external network is accessible from. Static path binding is a mechanism within ACI that explicitly maps an External Bridged Network to a specific path on a leaf switch (i.e. a physical interface). Without the static path binding, ACI would not know on which interface to receive or transmit traffic related to this external Layer 2 network. ACI dynamically handles internal paths using endpoints and policy models, but when connecting to the outside via Layer 2, static path bindings are required.

Options A, C, and E are incorrect. Option A, "Layer 2 remote fabric," describes a network that connects to an ACI fabric through another fabric and is not the correct option. Option C, "Layers 2 internal," refers to internal layer 2 bridge domains within the ACI fabric and therefore it is not suitable for an external bridged network.
Option E, “VXLAN outside”, is used to encapsulate Layer 2 traffic inside VXLAN to extend L2 domains, and this method is not used in this context where the network is natively bridged.

In summary, creating a Layer 2 outside definition and binding it to a physical interface using static path binding forms a complete configuration that enables the ACI fabric to connect with an external bridged network.

Authoritative Links:

1. **Cisco ACI Fundamentals:**https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/index.html (This provides a broad overview and

understanding of ACI)

2. **Cisco ACI Layer 2 Outside Connectivity:** [Specific document or configuration guide on Cisco ACI Layer 2 external connectivity, would need to be specific to a Cisco resource](The Cisco website would have specific documentation on connecting to external Layer 2 networks within ACI. search for keywords “ACI Layer 2 external connectivity configuration” or “ACI bridged network configuration” or “external bridged domain ACI configuration” in Cisco documentation)

Question: 40

Which two actions extend a Layer 2 domain beyond the ACI fabric? (Choose two.)

A.extending the routed domain out of the ACI fabric
B.creating a single homed Layer 3 Out
C.creating an external physical network
D.extending the bridge domain out of the ACI fabric
E.extending the EPG out of the ACI fabric

Answer: DE

Explanation:

The question asks about methods to extend a Layer 2 domain beyond the Cisco ACI fabric. Layer 2 domains operate at the data link layer, dealing with MAC addresses and local network segments.

Option **D**, "extending the bridge domain out of the ACI fabric," is correct because bridge domains are the core constructs representing Layer 2 networks within ACI. By connecting a bridge domain to external Layer 2 devices through a Layer 2 Out, you extend that domain beyond the ACI fabric. This maintains the same broadcast domain across the ACI fabric and external networks.

Option **E**, "extending the EPG out of the ACI fabric," is also correct. An EPG (Endpoint Group) is a logical grouping of endpoints with common policy requirements. When an EPG is associated with a Layer 2 external segment, endpoints connected to that external segment are treated as part of the same EPG, even though they reside outside the ACI fabric. This effectively stretches the EPG and its associated Layer 2 domain to external networks.

Option **A**, "extending the routed domain out of the ACI fabric," is incorrect. Routed domains operate at Layer 3, using IP addresses for routing, and don't extend a Layer 2 domain. Option **B**, "creating a single homed Layer 3 Out," is incorrect as Layer 3 Outs connect the ACI fabric to routed networks and are not related to extending a Layer 2 domain. Option **C**, "creating an external physical network," is incorrect because it merely defines the external infrastructure, without inherently extending the ACI Layer 2 domain.

In essence, extending a Layer 2 domain involves allowing devices outside the ACI fabric to be part of the same broadcast domain as devices inside the fabric, which is achieved through bridge domain and EPG extensions to external Layer 2 networks.

Authoritative Links:

Cisco ACI Fundamentals:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/index.html
Cisco ACI Layer 2 and 3 Networking:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/basic/GUI/b_Cisco_APIC_Basic_Configuration_Guide/b_Cisco_APIC_Basic_Configuration_Guide_chapter_011.htm

Question: 41

When Cisco ACI connects to an outside Layers 2 network, where does the ACI fabric flood the STP BPDU frame?

A.within the bridge domain
B.within the APIC
C.within the access encap VLAN
D.between all the spine and leaf switches

Answer: C

Explanation:

The correct answer is **C. within the access encap VLAN**. When a Cisco ACI fabric connects to an external Layer 2 network, it needs a mechanism to handle bridging and prevent loops. Spanning Tree Protocol (STP) is often used for this. However, ACI operates with a policy-based model, and traditional STP is not directly implemented within the ACI fabric's internal structure. Instead, the ACI fabric handles STP BPDUs in a specific way when interacting with external Layer 2 networks. When an STP BPDU enters the ACI fabric from an external Layer 2 network, the leaf switch receiving that BPDU treats it like any other Layer 2 frame associated with an access encap VLAN configured for that external connection. The fabric then floods the BPDU within the access encap VLAN only. It's crucial to note that ACI isolates broadcast domains by design and does not flood STP BPDUs across the entire fabric, which could cause issues in the ACI fabric. ACI keeps the broadcast domain limited to a single bridge domain or access encap VLAN, ensuring that STP is confined to that domain and prevents any possible loop within the fabric itself. The access encap VLAN essentially maps an external VLAN to a logical construct within the ACI fabric. Therefore, the STP BPDU will propagate within that specific VLAN, maintaining network segmentation and preventing widespread flooding within the fabric. The access encapsulation is key to correctly identifying and processing traffic coming from the external L2 Network. ACI uses policies for endpoint learning, and BPDUs are not used for that, meaning they don't go beyond the access encap VLAN.

Here are some links for further research:

Cisco ACI Fundamentals:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/index.html
Cisco ACI Layer 2 Bridging:https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/l2-configuration/cisco-apic-layer-2-configuration-guide-52x/m_l2-external-connectivity-52x.html
ACI and STP Interoperability:https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/l2-configuration/cisco-apic-layer-2-configuration-guide-52x/m_l2-external-connectivity-52x.html#topic_88753E9548B5429E8117FD2C7C03E63F

Question: 42

On which two interface types should a user configure storm control to protect against broadcast traffic? (Choose two.)

A.APIC facing interfaces
B.port channel on a single leaf switch
C.all interfaces on the leaf switches in the fabric
D.endpoint-facing trunk interface
E.fabric uplink interfaces on the leaf switches

Answer: BD

Explanation:

Let's analyze why options B and D are the correct choices for configuring storm control in a Cisco ACI environment.

B. port channel on a single leaf switch: Port channels, also known as link aggregations, combine multiple physical links into a single logical link. This increases bandwidth and redundancy. However, a broadcast storm on one member link of a port channel could potentially propagate to all member links, impacting multiple devices connected via that port channel. Therefore, configuring storm control on a port channel mitigates this risk by preventing excessive broadcast traffic from flooding the aggregated link.

D. endpoint-facing trunk interface: Endpoint-facing trunk interfaces connect to servers, hypervisors, and other endpoints. These interfaces often carry traffic from multiple VLANs. A broadcast storm originating from an endpoint can flood all VLANs on the trunk, potentially causing severe network disruptions and consuming valuable resources. Applying storm control on these interfaces prevents endpoint-generated broadcast storms from disrupting other endpoints and impacting the fabric.

Options A, C, and E are less relevant for storm control in this context:

A. APIC facing interfaces: The APIC is the controller of the ACI fabric, and its interfaces primarily handle control-plane traffic, not data-plane traffic vulnerable to broadcast storms.

C. all interfaces on the leaf switches in the fabric: While technically possible, applying storm control on all interfaces is not efficient and might disrupt normal traffic flow. The best practice is to apply it where broadcast storms are most likely to originate and propagate - endpoint facing interfaces and port channels.

E. fabric uplink interfaces on the leaf switches: These interfaces connect to the spine switches and handle inter-fabric traffic. Broadcast storms are less likely to propagate from an endpoint across the fabric via the uplinks but originate from the endpoints themselves, making endpoint facing interfaces a more relevant focus.

In essence, storm control should be strategically applied at the edges of the network, where broadcast traffic from endpoints or redundant connections are most likely to cause congestion or disruption.

Authoritative Links:

Cisco ACI Best Practices Guide : https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/white-paper-c11-737762.html
Cisco Nexus 9000 Storm Control Configuration : https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/configuration/guide/b_Cisco_Nexus_9OS_Interfaces_Configuration_Guide/b_Cisco_Nexus_9000_Series_NX-OS_Interfaces_Configuration_Guide_chapter_0101.html

Question: 43

Which two dynamic routing protocols are supported when using Cisco ACI to connect to an external Layer 3 network? (Choose two.)

A.iBGP
B.VXLAN
C.IS-IS
D.RIPv2
E.eBGP

Answer: AE

Explanation:

Cisco ACI supports external Layer 3 network connectivity primarily through Border Gateway Protocol (BGP), specifically both internal BGP (iBGP) and external BGP (eBGP). iBGP is typically used for routing within the same Autonomous System (AS), enabling communication between ACI fabric leaf switches and external routers that belong to the same AS. eBGP is used for routing between different AS's, allowing ACI to connect to networks controlled by separate organizations or entities. While other routing protocols exist, the ACI design philosophy and implementation highly favor BGP for its scalability and policy-based routing
capabilities essential for a data center environment. IS-IS, RIPv2, and VXLAN are not supported as routing protocols in this specific context for external Layer 3 connectivity. Instead, VXLAN serves as the primary tunneling mechanism within the ACI fabric, not for routing to external networks. RIPv2 is outdated and not designed for larger-scale data center environments. IS-IS is more frequently seen in Service Provider networks and less commonly integrated in the context of ACI external connectivity. Therefore, the supported and standard protocols for Cisco ACI external Layer 3 routing are iBGP and eBGP, making choices A and E the correct answer.Further Research:Cisco ACI Layer 3 External ConnectivityCisco ACI External Routed Connections - YouTube

Question: 44

What must be configured to redistribute externally learned OSPF routes within the ACI fabric?

A.Route Control Profile
B.BGP Route Reflector
C.BGP Inter-leak Route Map
D.PIM Sparse Mode

Answer: B

Explanation:

The correct answer is **B. BGP Route Reflector**. To redistribute externally learned OSPF routes within the ACI fabric, a Border Gateway Protocol (BGP) route reflector is required. ACI's internal routing protocol is MP-BGP, and it doesn't directly understand or process OSPF routes. When OSPF routes are learned at the ACI border leaf switches, they must be converted into BGP routes before being propagated within the ACI fabric. A route reflector simplifies this process by eliminating the need for full mesh BGP peering between all spine switches.

The route reflector receives external routes via BGP, which are then advertised to all other iBGP speakers in the fabric. Option A, Route Control Profile, is used for filtering routes but not for redistribution. Option C, BGP Inter-leak Route Map, is used for leaking routes between different address families, which isn't the case here.

Option D, PIM Sparse Mode, is unrelated as it's a multicast protocol, not for redistributing unicast routes.

Therefore, only a BGP route reflector handles the required conversion and redistribution of OSPF routes within the ACI fabric.

Here are some authoritative links for further research:

1. **Cisco Application Centric Infrastructure Fundamentals:** https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/index.html
2. **Cisco ACI Routing Concepts:**https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/routing/cisco-aci-routing-deployment-guide-52x.html
3. **BGP Route Reflector Configuration:**https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13753-rreflector.html

Question: 45

Regarding the MTU value of MP-BGP EVPN control plane packets in Cisco ACI, which statement about communication between spine nodes in different sites is true?

A.By default, spine nodes generate 9000-bytes packets to exchange endpoints routing information. As a result, the Inter-Site network should be able to carry 9000-bytes packets.

B.By default, spine nodes generate 1500-bytes packets to exchange endpoints routing information. As a result, the Inter-Site network should be able to carry 1800-bytes packets.

C.By default, spine nodes generate 1500-bytes packets to exchange endpoints routing information. As a result, the Inter-Site network should be able to carry 1500-bytes packets.

D.By default, spine nodes generate 9000-bytes packets to exchange endpoints routing information. As a result, the Inter-Site network should be able to carry 9100-bytes packets.

Answer: D

Explanation:

The correct answer is D. Let's break down why:

In Cisco ACI Multi-Pod deployments, MP-BGP EVPN is used for the control plane, enabling endpoint reachability information exchange between spine nodes located in different ACI sites (pods). These MP-BGP EVPN packets encapsulate endpoint routing information, which can become quite large, especially when many endpoints are present.

By default, Cisco ACI spine nodes generate MP-BGP EVPN packets with a 9000-byte Maximum Transmission Unit (MTU). This large MTU is chosen to optimize the transmission of larger routing updates, reducing overhead and improving efficiency.

However, it is not sufficient for the Inter-Site network to merely support 9000-byte packets. An additional layer of encapsulation is added for inter-site transport. This encapsulation typically involves VXLAN headers and possibly others, adding a few bytes to the overall packet size. Due to this overhead, the Inter-Site network must support an MTU slightly larger than 9000 bytes to prevent packet fragmentation, which can lead to performance degradation. Generally a buffer of 100 bytes is added to the maximum packet size. Therefore, 9100 is necessary. While the added header size depends on the underlying technology, the important takeaway is that some bytes must be added to the base MTU for effective operation. Options A, B, and C fail to consider this crucial point.

Options A and C incorrectly state that the base MTU is 1500 bytes, which is the standard Ethernet MTU, not the MTU used for MP-BGP EVPN packets within ACI Multi-Pod. Option B is inaccurate on the standard MTU and the inter-site MTU calculation. The need for an inter-site network MTU to accommodate additional encapsulation explains why choice D is the correct answer.

For further research, refer to official Cisco documentation on ACI Multi-Pod and inter-site connectivity. Specific topics to investigate would be EVPN control plane, MTU considerations, and inter-site forwarding for Multi-Pod.

Authoritative Links:

Cisco ACI Multi-Pod White Paper: (Search for "Cisco ACI Multi-Pod Deployment Guide" on Cisco's website.) -Look for sections discussing inter-site networking and MTU.

Cisco ACI Configuration Guides: (Search for "Cisco APIC Configuration Guide" on Cisco's website.) - Look for specifics on configuring Multi-Pod and related MTU settings.

Cisco Live Presentations: Search for Cisco Live presentations on ACI Mult-pod which often contain detailed discussions and best practices.

Question: 46

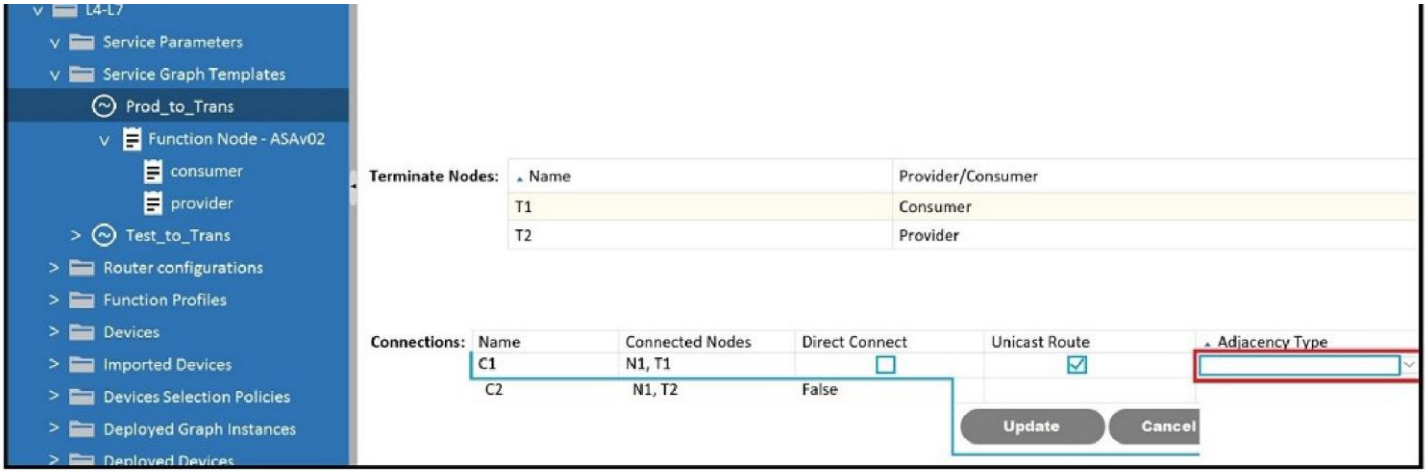

Refer to the exhibit. Which Adjacency Type value should be set when the client endpoint and the service node interface are in a different subnet?

A.Routed
B.Unicast
C.L3Out
D.L3

Answer: D

Explanation:

D. L3 or L2 available for configuration only.

Question: 47

Which endpoint learning operation is completed on the ingress leaf switch when traffic is received from a Layer 3 Out?

A.The source MAC address of the traffic is learned as a local endpoint.
B.The source MAC address of the traffic is learned as a remote endpoint.
C.The source IP address of the traffic is learned as a remote endpoint.
D.The source IP address of the traffic is learned as a local endpoint.

Answer: A

Explanation:

The correct answer is A. When traffic enters the ACI fabric via a Layer 3 Out (L3Out), the ingress leaf switch performs endpoint learning. This process involves examining the incoming frame's source MAC address. Since the traffic originated from outside the ACI fabric, the source MAC address represents an external device. The ingress leaf switch learns this MAC address as a local endpoint. This means it creates an entry in its local endpoint table, associating the source MAC address with the specific ingress interface and VLAN the traffic arrived on. This learned endpoint is considered local because it's directly connected to the fabric from the perspective of the ingress leaf. The IP address is not learned at this stage of the learning process; IP endpoint learning occurs when the destination IP is inside the fabric. In essence, the switch now knows how to reach this MAC address if it sees traffic destined for it within the ACI fabric. This is fundamental to how ACI efficiently forwards packets within and out of the fabric based on MAC address learning.

Authoritative sources for further research:

Cisco ACI Fundamentals:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/index.html (General overview of ACI)
Cisco ACI Endpoint Learning: (Refer to official documentation for detailed information on endpoint learning, specifically around ingress switches and layer 3 out scenarios). Search for "Cisco ACI endpoint learning" on Cisco's official documentation portal for the most accurate details and best practices.

Question: 48

An engineer must connect Cisco ACI fabric using Layer 2 with external third-party switches. The third-party switches are configured using 802.1s protocol. Which two constructs are required to complete the task? (Choose two.)

A.spanning tree policy for mapping MST Instances to VLANs
B.MCP policy with PDU per VLAN enabled
C.MCP instance policy with administrative slate disabled
D.dedicated EPG for native VLAN
E.static binding of native VLAN in all existing EPGs

Answer: AD

Explanation:

Here's a detailed justification for why options A and D are the correct constructs when connecting a Cisco ACI fabric to external switches using Layer 2 and the 802.1s protocol (Multiple Spanning Tree Protocol - MSTP):

A. Spanning tree policy for mapping MST Instances to VLANs: When connecting to a third-party switch using MSTP, the ACI fabric needs to be configured to understand the MSTP instances and their corresponding VLANs. ACI doesn't automatically infer this mapping. A spanning tree policy allows you to explicitly map MST instances to VLANs within the ACI fabric. This ensures the ACI fabric participates correctly in the MSTP domain and avoids loops. Without this mapping, the ACI fabric wouldn't correctly process BPDU messages from the external switch or form a consistent spanning tree.

D. Dedicated EPG for native VLAN: When connecting switches that utilize the native VLAN concept on trunks, a dedicated EPG within the ACI fabric is essential. This is because ACI EPGs by default are associated with a tagged VLAN. By creating a dedicated EPG for the native VLAN, you explicitly configure the ACI fabric to handle untagged traffic appropriately. If native VLAN traffic is not handled by a specific EPG and you use the default VLAN (untagged), the ACI fabric might misinterpret traffic or experience other unexpected behavior. This ensures the native VLAN traffic from the third-party switch is properly categorized within the ACI fabric and receives the intended policy treatment.

Options B, C and E are not correct:

B. MCP policy with PDU per VLAN enabled: MCP (Misconfiguration Protocol) is used for loop detection and mitigation within an ACI fabric. While important for ACI itself, it is not directly related to connecting to external MSTP domains. Enabling PDU per VLAN wouldn't resolve issues related to VLAN mappings or the native VLAN.
C. MCP instance policy with administrative slate disabled: Similarly to B, MCP configurations within the ACI fabric have no effect on external connectivity using MSTP. Disabling the administrative state does not solve the issue of VLAN mapping with the external device
E. Static binding of native VLAN in all existing EPGs: Static binding of the native VLAN in all existing EPGs is not an efficient or proper solution. Instead of using a dedicated EPG, static binding adds the untagged VLAN to all other existing EPGs, creating inefficiencies and making it harder to implement targeted security or QoS

policies.

Authoritative Links:

Cisco ACI Layer 2 Integration:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/l2-configuration/b-cisco-apic-layer-2-configuration-guide/m-layer-2-external-connectivity.html (Look for sections about "External Layer 2 Connectivity" and "Spanning Tree Configuration")
Cisco ACI Spanning Tree Protocol:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/l2-configuration/b-cisco-apic-layer-2-configuration-guide/m-layer-2-bridging.html#task_44087693A0C147B598426E949B895B2C (Focus on the section pertaining to MST and mapping)
Cisco ACI EPGs:https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/l2-configuration/b-cisco-apic-layer-2-configuration-guide/m-aci-epg-overview.html (Understand EPG behavior and VLAN associations)

In summary, proper MST instance mapping using a spanning tree policy and the creation of a dedicated EPG for native VLAN traffic are critical for a successful Layer 2 integration between a Cisco ACI fabric and external switches using the 802.1s protocol.

Question: 49

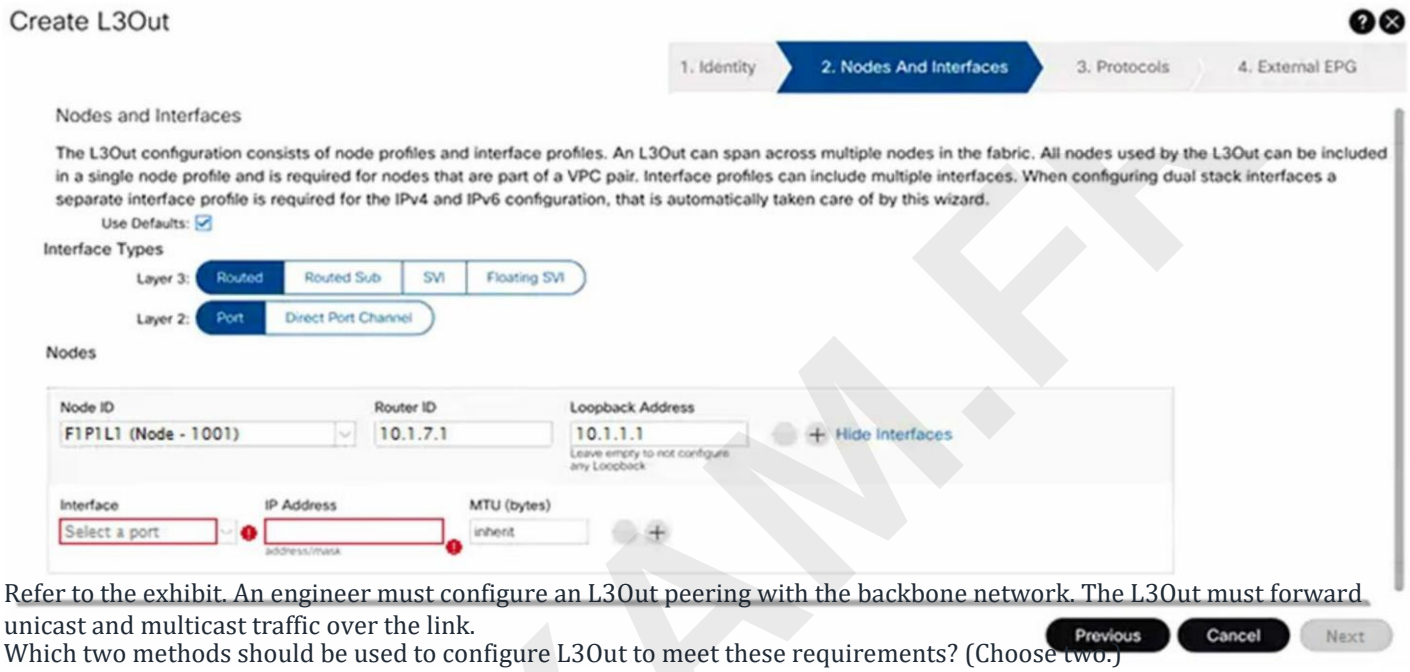

Refer to the exhibit. An engineer must configure an L3Out peering with the backbone network. The L3Out must forward unicast and multicast traffic over the link.
Which two methods should be used to configure L3Out to meet these requirements? (Choose two.)

A.Layer 3 routed port
B.VPC with SVI
C.port channel with SVI
D.Layer 3 routed subinterface
E.Layer 3 floating SVI

Answer: AD

Explanation:

A and D. "PIM is supported on Layer 3 Out routed interfaces and routed subinterfaces including Layer 3 port-channel interfaces. PIM is not supported on Layer 3 Out SVI interfaces."

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/L3_config/b_Cisco_APIC_Layer_3_Configuration_Guide/b_Cisco_APIC_Layer_3_Configuration_Guide_chapter_

01111.html#id_21570

Question: 50

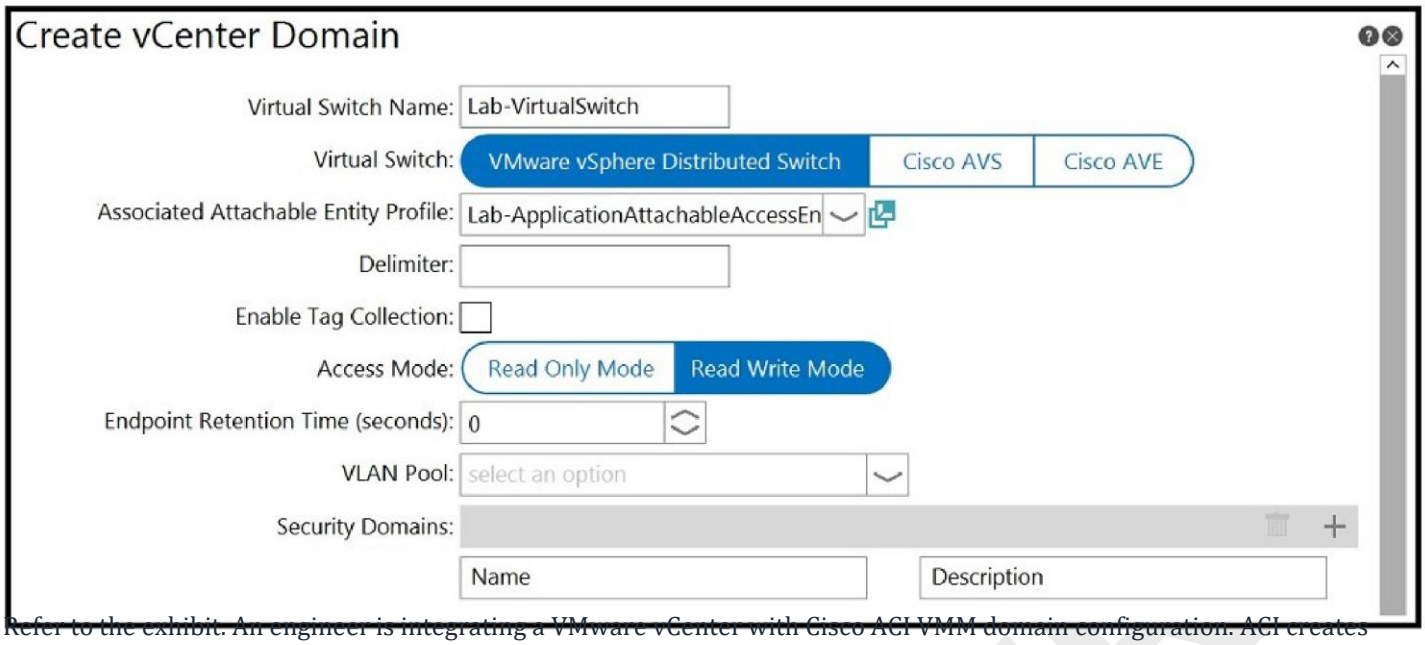

Refer to the exhibit. An engineer is integrating a VMware vCenter with Cisco ACI VMM domain configuration. ACI creates port-group names with the format of
"Tenant | Application | EPG". Which configuration option is used to generate port groups with names formatted as "Tenant=Application=EPG"?

A.enable tag collection
B.security domains
C.delimiter
D.virtual switch name

Answer: C

Explanation:

Answer is C: DelimiterStep 8 Configuring the delimiter during VMM domain creation, perform the following actions:On the menu bar, choose VM NETWORKING > InventoryIn the Navigation pane, right-click VMware and click Create vCenter Domain.In the Create vCenter Domain dialog box, enter a Name.Optional: In the Delimiter field, enter one of the following: |, ~, !, @, ^, +, or =. If you do not enter a symbol, the system default | delimiter will appear in the VMware PortGroup name.

Question: 51

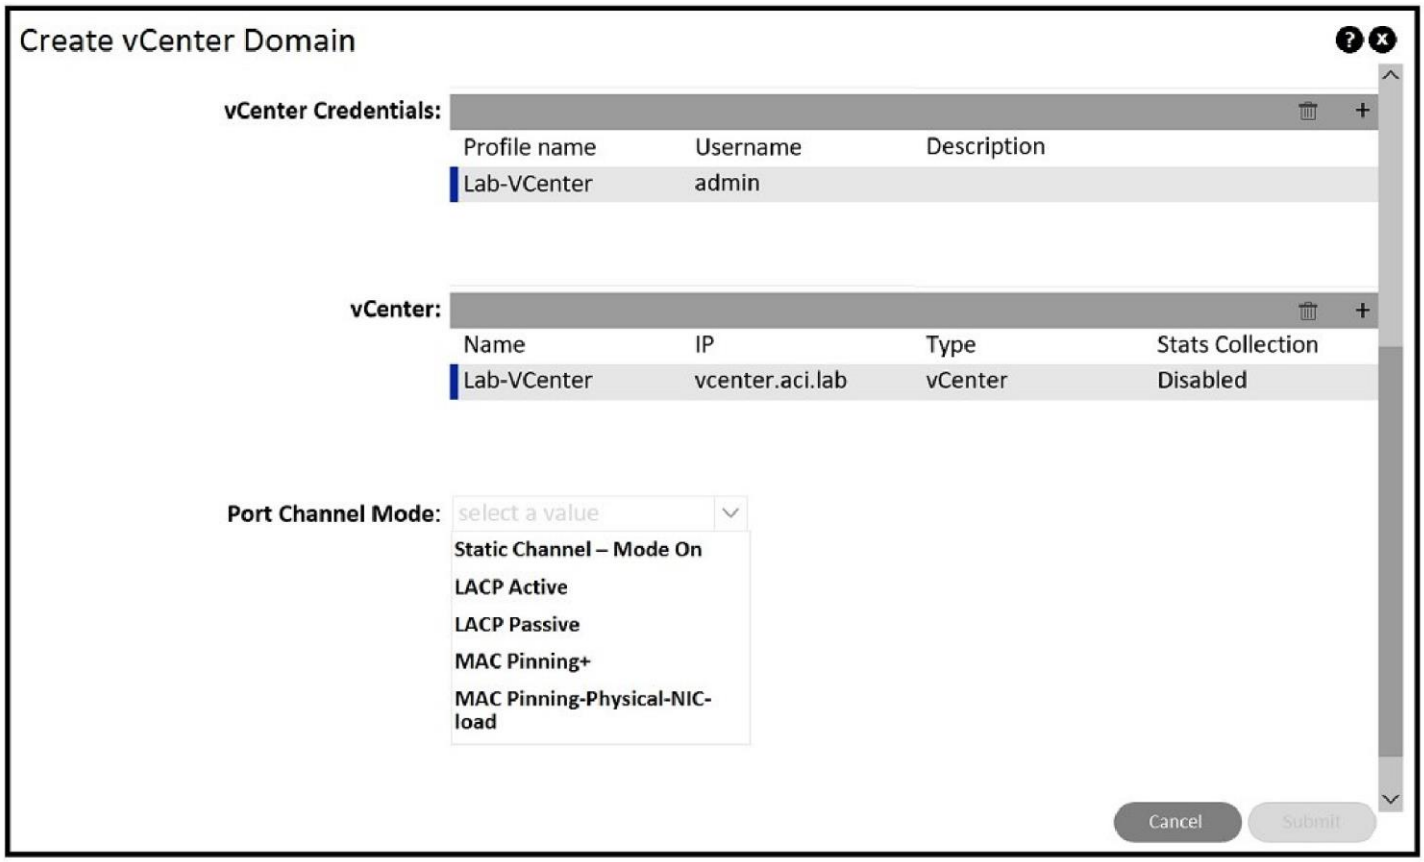

Refer to the exhibit. An engineer is implementing Cisco ACI "" VMware vCenter integration for a blade server that lacks support of bonding. Which port channel mode results in "route based on originating virtual port" on the VMware VDS?

A.Static Channel "" Mode On
B.MAC Pinning-Physical-NIC-load
C.LACP Passive
D.MAC Pinning+
E.LACP Active

Answer: D

Explanation:

MAC Pinning+.

Reference:

https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/6x/virtualization/cisco-aci-virtualization-guide-60x/ACI-Virtualization-Guide-60x-aci-with-vmware-vds.pdf

Question: 52

When configuring Cisco ACI VMM domain integration with VMware vCenter, which object is created in vCenter?

A.datacenter
B.VMware vSphere Standard vSwitch
C.VMware vSphere Distributed Switch
D.cluster

Answer: C

Explanation:

When integrating Cisco ACI with VMware vCenter, a crucial step is establishing network connectivity for virtual machines. This connectivity is facilitated through a VMware vSphere Distributed Switch (VDS). ACI utilizes the VDS as its external virtual network interface for managing virtual machine traffic. The VDS provides a centralized management plane for networking across multiple ESXi hosts, aligning with the ACI's centralized policy management paradigm. Cisco ACI's VMM domain configures and deploys the VDS in vCenter. The ACI fabric does not create or manage datacenters or clusters in vCenter; these are pre-existing elements of the virtualized infrastructure. While standard vSwitches can exist within vSphere, the ACI VMM integration specifically relies on VDS for features such as link aggregation (LACP) support, Port Group management, and advanced policy enforcement that can be centrally managed by the ACI fabric. Therefore, out of the listed options, the VDS is the object created in vCenter during the ACI VMM integration process.

Relevant Links:

1. **Cisco ACI Virtual Machine Manager (VMM) Integration:** https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/vmm-integration/cisco-aci-vmm-integration-51x.html
2. **VMware vSphere Networking:**https://docs.vmware.com/en/VMware-vSphere/index.html (Refer to sections on Distributed Switches)

Question: 53

DRAG DROP -
Drag and drop the Cisco ACI Layer 4 to Layer 7 service insertion terms on the left to the correct descriptions on the right.
Select and Place:

Answer:

concrete interfaces	ensures reachability between L3 domains
service graph	rendered with local resources that are available in the fabric
device cluster	contains an active-standby pair of firewalls or load balancers
VRF stitching	encapsulations programmed based on their association with logical interfaces

Explanation:

concrete interfaces	encapsulations programmed based on their association with logical interfaces
service graph	rendered with local resources that are available in the fabric
device cluster	contains an active-standby pair of firewalls or load balancers
VRF stitching	ensures reachability between L3 domains

concrete interfaces: encapsulations programmed based on their association with logical interfaces.

service graph: rendered with local resources that are available in the fabric.

device cluster: contains an active-standby pair of firewalls or load balancers.

VRF stitching: ensures reachability between L3 domains.

Question: 54

An engineer has set the VMM resolution immediacy to pre-provision in a Cisco ACI environment. No Cisco Discovery Protocol neighborship has been formed between the hypervisors and the ACI fabric leaf nodes. How does this affect the download policies to the leaf switches?

A.No policies are downloaded because LLDP is the only supported discovery protocol.
B.Policies are downloaded when the hypervisor host is connected to the VMM VDS.
C.Policies are downloaded to the ACI leaf switch regardless of Cisco Discovery Protocol neighborship.
D.No policies are downloaded because there is no discovery protocol neighborship.

Answer: C

Explanation:

The correct answer is C. With VMM resolution immediacy set to "pre-provision" in Cisco ACI, policies are pushed to the leaf switches regardless of the immediate discovery of hypervisors via CDP or LLDP. This pre-provisioning approach decouples policy deployment from real-time discovery events. The system anticipates the connection of hypervisors and configures the leaf ports accordingly in advance. It is designed to ensure that when a hypervisor eventually connects, the necessary policies are already in place. Options A and D are incorrect because the pre-provision setting overrides the need for a discovery protocol adjacency. The absence of CDP/LLDP neighborship only affects real-time discovery of entities, not the policy distribution when pre-provisioning is active. Option B is misleading; while connection to the VMM VDS is the ultimate goal, the pre-provision mode ensures policies are present even before that connection is fully established. This mode is beneficial in situations where there might be latency or delays in VMM integration. Therefore, the policies are distributed proactively based on the defined configuration instead of reactively based on a discovery event. This strategy is especially useful in large-scale deployments where timely application of policies is essential.**Authoritative Links for Further Research:**

Cisco ACI VMM Integration Guide: (Search on Cisco's official documentation portal using keywords like "Cisco ACI VMM Integration Pre-provision" to find the most relevant official Cisco guide)
Cisco ACI Fundamentals: (Search on Cisco's official documentation portal using keywords like "Cisco ACI Policy Deployment" to find related concepts in ACI Fundamentals)

Question: 55

In the context of VMM, which protocol between ACI leaf and compute hosts ensures that the policies are pushed to the leaf switches for immediate and on demand resolution immediacy?

A.VXLAN
B.LLDP
C.ISIS
D.STP

Answer: B

Explanation:

The correct answer is **B. LLDP (Link Layer Discovery Protocol)**. While VXLAN, ISIS, and STP play crucial roles in ACI fabric operation, they don't directly handle the immediate, on-demand policy push between the leaf switches and compute hosts concerning Virtual Machine Manager (VMM) integration.

Here's a detailed justification:

LLDP is a layer 2 discovery protocol that facilitates the exchange of device information between directly connected network devices. In the ACI context, when a virtual machine (VM) on a compute host is created or moved (via VMM), the host informs its directly connected leaf switch of the VM's presence and relevant attributes, including its network association. The leaf switch leverages LLDP to discover the compute host and its VM. This allows ACI to immediately identify the VM, enabling it to program the appropriate policies onto the leaf switch ports.

The immediacy is critical in dynamically changing environments managed by VMMs. Without this direct and immediate discovery mechanism, the ACI fabric would struggle to keep up with VM churn, potentially causing network connectivity disruptions or policy misapplication.

VXLAN is used for network virtualization within the ACI fabric, creating a stretched layer 2 network between ACI endpoints. However, it is primarily used within the fabric and not as an adjacency protocol between hypervisors and leaf switches.

ISIS is a routing protocol utilized within the ACI fabric for internal routing and loop prevention, not directly involved in interaction with compute hosts. STP is a spanning tree protocol, which is disabled within the ACI fabric.

Therefore, LLDP is the crucial protocol for the immediate on-demand resolution of policies in response to VMM events, creating a dynamic and automated environment.

Authoritative links for further research:

Cisco ACI Fundamentals:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/white-paper-c11-732984.html
Cisco Application Policy Infrastructure Controller (APIC) Basic Configuration Guide: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/basic-config/b_Cisco_APIC_Basic_Configuration_Guide/b_Cisco_APIC_Basic_Configuration_Guide_chapter_011.html **LLDP Overview:**https://en.wikipedia.org/wiki/Link_Layer_Discovery_Protocol

Question: 56

Which two components are essential parts of a Cisco ACI Virtual Machine Manager (VMM) domain policy configuration? (Choose two.)

A.Layer 3 outside interface association
B.EPG static port binding
C.VMM domain profile
D.EPG association
E.IP address pool association

Answer: CD

Explanation:

The correct answer is C and D. A Cisco ACI Virtual Machine Manager (VMM) domain policy configuration necessitates two crucial elements: a VMM domain profile and EPG (Endpoint Group) association.

A **VMM domain profile (C)** defines the specific characteristics of the hypervisor environment. It establishes communication parameters between the Cisco APIC (Application Policy Infrastructure Controller) and the hypervisor, such as vCenter or Hyper-V servers, including credentials and connection details. This profile enables the APIC to interact with the virtual infrastructure. Without a VMM domain profile, the APIC has no

context of the virtualization layer to implement ACI policies.

The **EPG association (D)** establishes the connection between the virtualized endpoints (VMs) and the ACI fabric. It dictates to which EPG the traffic originating from a VM or a specific vNIC (virtual network interface card) is assigned. This association maps VMs to their proper network security and access control policies defined by the EPG. Without this mapping, traffic from the VMs cannot be managed by the ACI fabric.

Options A, B, and E are not direct requirements for a VMM domain policy configuration. Layer 3 outside interface association is relevant for external connectivity, EPG static port binding relates to physical endpoints, and IP address pool association is for IP management within the ACI fabric, but they aren't primary aspects of integrating virtual environments using a VMM domain. The core of VMM integration hinges on communicating with the virtualized environment (VMM profile) and ensuring that VMs are correctly classified within ACI (EPG association).

Cisco ACI VMM Domain Configuration Guide:
https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/configuration/l3-infrastructure/cisco-aci-fabric-l3-infrastructure-config-5x/cisco-aci-fabric-l3-infrastructure-config-5x_chapter_010011.html
Cisco ACI Virtualization Guide:https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/virtualization/cisco-aci-virtualization-5x/cisco-aci-virtualization-5x_chapter_01.html

Question: 57

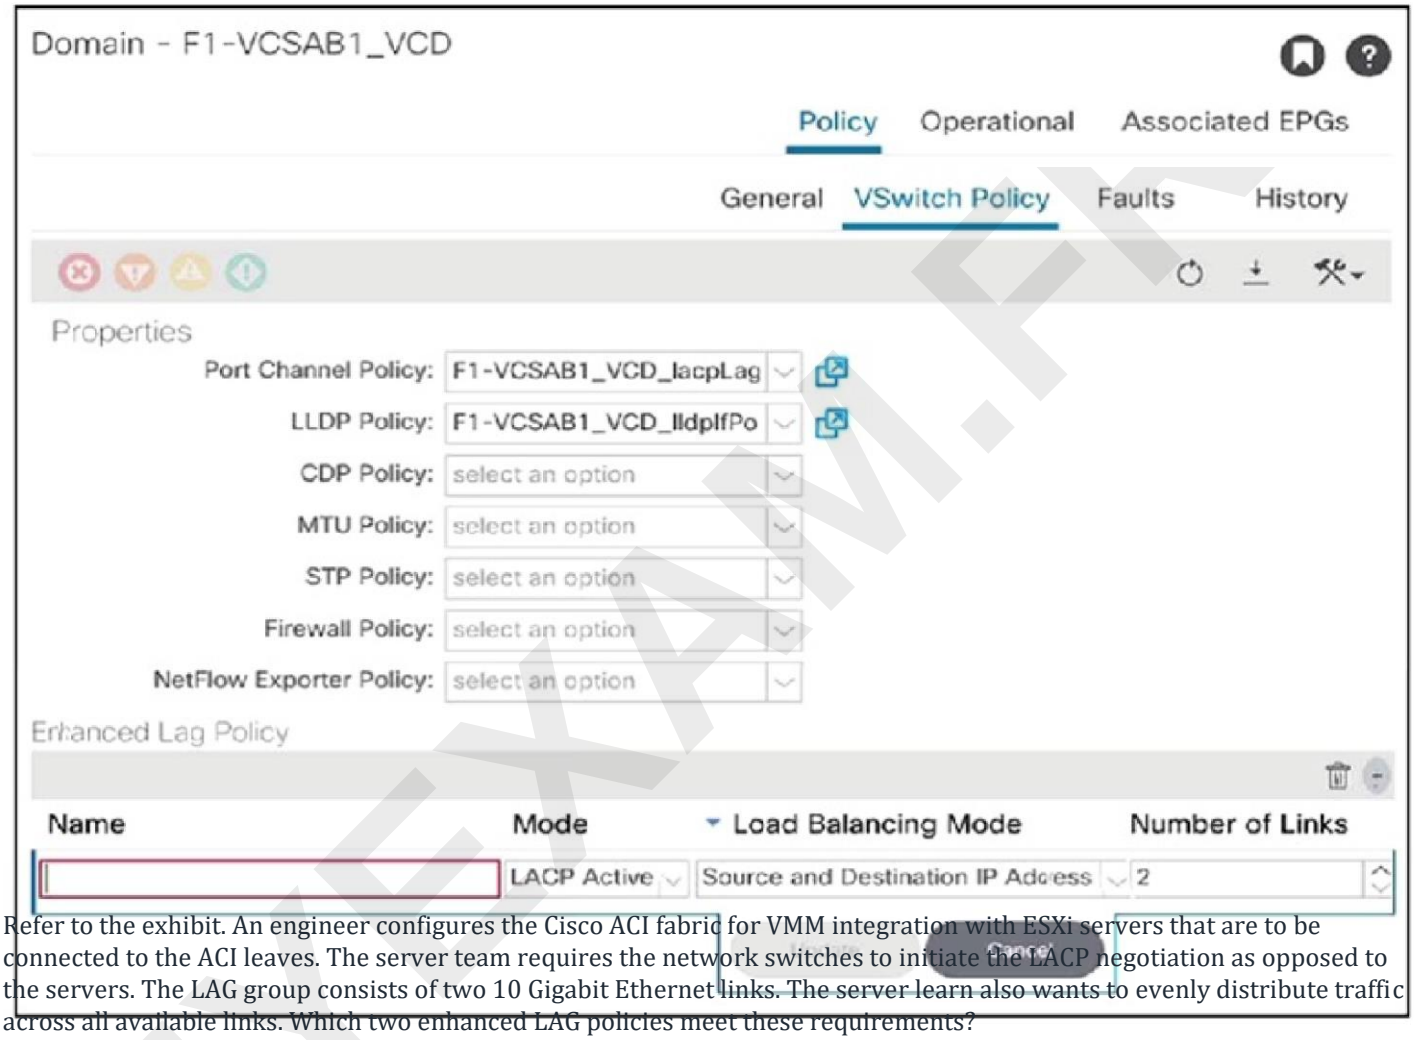

Refer to the exhibit. An engineer configures the Cisco ACI fabric for VMM integration with ESXi servers that are to be connected to the ACI leaves. The server team requires the network switches to initiate the LACP negotiation as opposed to the servers. The LAG group consists of two 10 Gigabit Ethernet links. The server learn also wants to evenly distribute traffic across all available links. Which two enhanced LAG policies meet these requirements?
(Choose two.)

A.LACP Mode: LACP Standby
B.LB Mode: Destination IP Address and TCP/UDP Port
C.LB Mode: Source and Destination MAC Address
D.LB Mode: Source IP Address and TCP/UDP Port
E.LACP Mode: LACP Active

Answer: BD

Explanation:

The answer is B and D. LACP Mode: LACP Standby - In LACP available only two modes - Active and PassiveLACP Mode: Active its standard LACP modeLB Mode: Source and Destination MAC Address its a standard LACP balancing modeAnd question says: Which two enhanced LAG policies meet these
requirements? (Choose two.) LB Mode: Destination IP Address and TCP/UDP PortLB Mode: Source IP Address and TCP/UDP PortThis two modes are LACP enhanced balancing modes.

Question: 58

Which tenant is used when configuring in-band management IP addresses for Cisco APICs, leaf nodes, and spine nodes?

A.default
B.infra
C.common
D.mgmt

Answer: D

Explanation:

The correct answer is **D. mgmt**. Cisco ACI utilizes distinct tenants for different administrative and operational purposes. The mgmt (management) tenant is specifically designed for managing the ACI fabric itself, including the infrastructure components. It provides a dedicated space for configuring out-of-band and, importantly, in-band management IP addresses used by the APICs, leaf nodes, and spine nodes for communication and monitoring. This segregation of management traffic within its own tenant ensures proper isolation and security for the critical management functions of the ACI fabric. The default tenant is for user-created applications and services. The infra tenant houses internal infrastructure components not directly manageable by end users. The common tenant is used to configure resources and policies that apply to multiple tenants but not specifically management functions. Using the mgmt tenant for in-band management aligns with the best practices of network segmentation, allowing for granular control over access and policy enforcement of management plane traffic, reducing the potential for accidental interference with application traffic.
Therefore, the mgmt tenant is the designated tenant for configuring the necessary in-band management IP addresses.

Authoritative Links:

Cisco ACI Fundamentals:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/fundamentals/b-cisco-aci-fundamentals.html (Refer to the sections on tenants)
Cisco ACI Management Configuration Guide:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b-aci-kb/b-aci-kb_chapter_0100.html (Search for management tenant information and in-band management details)

Question: 59

What represents the unique identifier of an ACI object?

A.universal resource identifier (URI)
B.application programming interface
C.management information tree
D.distinguished name

Answer: D

Explanation:

The correct answer, Distinguished Name (DN), is fundamental to how Cisco ACI uniquely identifies objects within its management information tree (MIT). A DN is a hierarchical naming convention that specifies an object's location within the ACI fabric's structured database. Think of it like a fully qualified path to a file on a computer, except instead of files, it's for ACI objects like tenants, application profiles, bridge domains, and more. Each component of the DN represents a level in the hierarchy, starting from the root and progressing to the specific object. This ensures that each object, regardless of its type, has a globally unique identifier within the ACI environment. While a URI might be used to address an object over an API, it doesn't define the object's inherent identity within the system's management structure like a DN does. An application programming interface (API) is a method of interacting with the system, not an identifier. The management information tree (MIT) is the overall data structure, not the identifier itself. The DN's hierarchical and unambiguous structure is critical for efficient management, policy enforcement, and troubleshooting within the ACI fabric. This identification system allows the system to consistently manage, identify, and locate resources across the entire infrastructure.

Relevant Links:

Cisco ACI Fundamentals:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/index.html (General overview of ACI concepts)
Cisco ACI Object Model:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/management/b-cisco-apic-mgmt-config-guide/b-cisco-apic-mgmt-config-guide_chapter_010.html#task_560430 (Explanation of ACI's managed objects, including DNs)

Question: 60

Which new construct must a user create when configuring in-band management?

A.VLAN pool
B.management contract
C.management tenant
D.bridge domain

Answer: A

Explanation:

The correct answer is **A. VLAN pool**.

When configuring in-band management within a Cisco Application Centric Infrastructure (ACI) fabric, a VLAN

pool is a necessary construct. In-band management implies that the management traffic flows within the same data plane network used by application traffic. This contrasts with out-of-band management, which uses a separate, dedicated network.

To enable in-band management, you need to allocate a specific range of VLAN IDs for this purpose. This is achieved by creating a VLAN pool. The pool acts as a container for a set of VLANs that can be used for management. The ACI fabric then utilizes this pool to assign VLANs to endpoints, allowing devices like the APIC (Application Policy Infrastructure Controller) and switches to communicate for management tasks.

Options B, C, and D are not directly involved in defining the VLANs for in-band management. A management contract (B) controls communication between management endpoints. A management tenant (C) is a logical grouping, and bridge domain (D) is related to network segmentation and forwarding but doesn't define the fundamental VLAN assignment. While these are related concepts and can be used in conjunction with in-band management, the VLAN pool is the crucial initial construct for specifying the management VLAN range. Without a defined VLAN pool, in-band management cannot be established.

In summary, the VLAN pool forms the foundation for providing the necessary VLANs for in-band management within ACI, making it the correct answer.

For further research, refer to Cisco documentation:

Cisco ACI Fundamentals: https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/white-paper-c11-732208.html
Cisco ACI Configuration Guide: (Refer to specific sections on management and VLAN pools in the appropriate release documentation) https://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/products-installation-and-configuration-guides-list.html

Question: 61

What must be configured to allow SNMP traffic on the APIC controller?

A.out-of-band management interface
B.contract under tenant mgmt
C.SNMP relay policy
D.out-of-band bridge domain

Answer: B

Explanation:

The correct answer is **B. contract under tenant mgmt**. To enable SNMP traffic on the APIC controller, you must define a contract under the 'mgmt' tenant. The Application Policy Infrastructure Controller (APIC) in Cisco ACI uses a policy-based model where access control is managed through contracts. These contracts dictate the communication allowed between endpoints. The 'mgmt' tenant specifically manages the infrastructure components, including the APIC itself. By creating a contract within the 'mgmt' tenant that permits SNMP traffic (typically UDP port 161), you explicitly allow external SNMP managers to communicate with the APIC. This method leverages the fundamental ACI concept of security via whitelisting, where only explicitly permitted traffic is allowed. An out-of-band management interface (A) is necessary for initial access and management but doesn't control traffic flow after configuration. An SNMP relay policy (C) is used to forward SNMP traps but not to allow initial SNMP communication to the APIC. Finally, an out-of-band bridge domain (D) facilitates management network access but doesn't directly enable SNMP communication through contracts. Only the creation of a policy, via a contract, within the mgmt tenant will allow traffic to and from the APIC controller itself.

Authoritative Links:

Cisco ACI Fundamentals: https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/index.html
Cisco ACI Management Guide: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/management/cisco-apic-management-guide-401.html (Specifically, search for 'SNMP' to find relevant configuration details within the management context)
Cisco ACI Contracts: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b_Cisco_APIC_Best_Practices/b_Cisco (Explains the general use of contracts in ACI)

Question: 62

Which type of port is used for in-band management within ACI fabric?

A.spine switch port
B.APIC console port
C.leaf access port
D.management port

Answer: C

Explanation:

The correct answer is C, leaf access port. In Cisco's Application Centric Infrastructure (ACI), in-band management refers to managing the fabric using the data plane, which means leveraging the existing network infrastructure for management traffic. This contrasts with out-of-band management, which utilizes a separate, dedicated network. Leaf access ports are specifically designed to connect endpoints (servers, virtual machines, etc.) to the ACI fabric, and these endpoints often require in-band management. APIC (Application Policy Infrastructure Controller) console ports (B) are used for initial setup and out-of-band access, while spine switch ports (A) primarily interconnect leaf switches and do not directly connect to managed endpoints. Management ports (D) usually provide out-of-band access for the APIC itself, not for the general management of endpoints within the fabric. Therefore, leaf access ports, due to their connection to endpoints requiring management, are the designated points for in-band management within an ACI fabric. Endpoints connected to leaf access ports can be reached via the same data plane used for normal application traffic, thereby enabling in-band control. The ACI fabric leverages policies and contracts to ensure secure and controlled in-band management.

Further research:

1. **Cisco ACI Fundamentals:**https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/white-paper-c11-731704.html (This white paper provides a comprehensive overview of ACI architecture, including the role of leaf and spine switches and in-band management concepts.)
2. **Cisco ACI Fundamentals Documentation:** https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/Cisco_APIC_Fundamentals.ht (This documentation covers basic concepts and components of ACI including network access.)

Question: 63

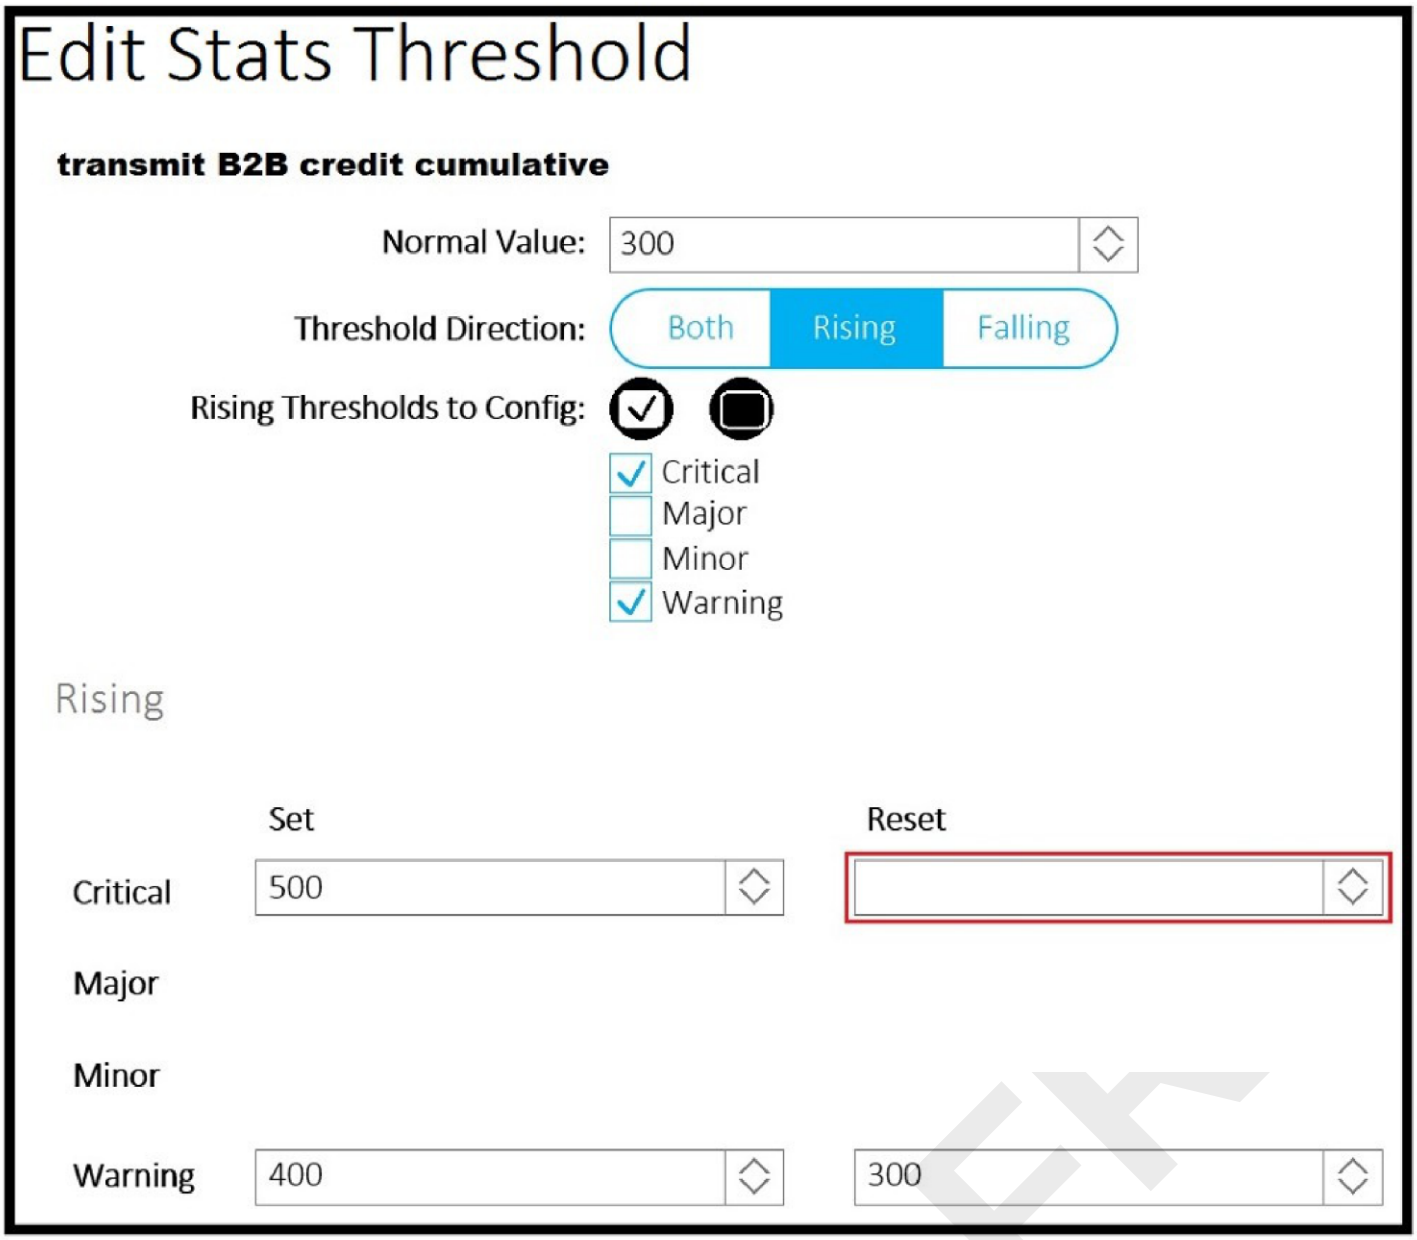

Refer to the exhibit. A client reports that the ACI domain connectivity to the fiber channel storage is experiencing a B2B credit oversubscription. The environment has a SYSLOG server for state collection messages. Which value should be chosen to clear the critical fault?

A.300
B.410
C.350
D.510

Answer: B

Explanation:

Correct answer is B:410.

Question: 64

Which statement about ACI syslog is true?

A.Notifications for different scopes of syslog objects can be sent only to one destination. B.Syslog messages are sent to the destination through the spine.

C.All syslog messages are sent to the destination through APIC.

D.Switches send syslog messages directly to the destinations.

Answer: D

Explanation:

Here's a detailed justification for why option D is the correct answer regarding ACI syslog behavior:

In Cisco ACI (Application Centric Infrastructure), syslog messages are generated by various components, including the switches (leaf and spine nodes), and the Application Policy Infrastructure Controller (APIC).

However, when it comes to forwarding these syslog messages to external destinations, the leaf and spine switches send the messages directly. They don't rely on the APIC as a transit point. This direct forwarding mechanism is designed to minimize latency and overhead. Option D correctly reflects this fundamental aspect of ACI's syslog implementation, stating "Switches send syslog messages directly to the destinations."

Option A is incorrect as ACI allows for multiple syslog destinations and different levels (scopes) of objects can be configured to use specific destinations. Option B is incorrect because while spines are part of the fabric, the syslog messages are sent directly from the source switches to the designated destinations without traversing the spine as an intermediate hop for syslog purposes. Option C is also incorrect because APIC is involved in configuring syslog policies, it doesn't act as a central proxy for forwarding syslog messages generated by switches. Instead, it's the switches that directly send the messages to the configured servers.

The distributed nature of syslog forwarding ensures efficiency and avoids overloading the central APIC controller. The switches use the management network out-of-band for sending syslog data. This approach provides a reliable mechanism for real-time monitoring and troubleshooting across the ACI fabric, avoiding any performance impact of sending high-volume logging through APIC itself.

Authoritative Links:

1. **Cisco ACI Logging and Monitoring Guide:** https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/system-management/latest/cisco-aci-system-management-guide/b_Cisco_ACI_System_Management_Guide_latest_chapter_01101.html (Focus on the syslog section).

2. **Cisco ACI Best Practices:**https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/white-paper-listing.html (Search for relevant content on monitoring and logging).

Question: 65

A data center administrator is upgrading an ACI fabric. There are 3 APIC controllers in the fabric and all the servers are dual-homed to pairs of leaf switches configured in VPC mode. How should the fabric be upgraded to minimize possible traffic impact during the upgrade?

A.1. Create two maintenance groups for the APIC controllers: VPC left and VPC right. 2. Upgrade the first group of controllers. 3. Upgrade the second group of controllers. 4. Upgrade the leaf switches.

B.1. Create two maintenance groups for APIC controllers: VPC left and VPC right. 2. Upgrade the leaf switches. 3. Upgrade the first group of controllers. 4. Upgrade the second group of controllers.

C.1. Create two maintenance groups for the leaf switches: VPC left and VPC right. 2. Upgrade the APIC controllers. 3. Upgrade the first group of leaf switches. 4. Upgrade the second group of leaf switches.

D.1. Create two maintenance groups for the leaf switches: VPC left and VPC right. 2. Upgrade the first group of switches. 3. Upgrade the second group of switches. 4. Upgrade the APIC controllers.

Answer: C

Explanation:

The correct upgrade procedure for an ACI fabric with dual-homed servers and multiple APIC controllers, aiming for minimal traffic disruption, involves prioritizing the APIC controllers' stability before touching the data plane. Option C is the correct approach because it ensures the control plane remains healthy during the leaf switch upgrades. First, creating two maintenance groups based on VPC pairs (VPC left and VPC right) allows for a controlled, staged upgrade of the leaf switches. This approach ensures that traffic can still flow through the other peer of each VPC while one side undergoes maintenance. Next, upgrading the APIC controllers before upgrading the leaf switches is the crucial step. This order maintains control plane stability during the subsequent leaf switch upgrades. Only after ensuring the APIC controllers are operational and synchronized should the leaf switches be upgraded in a staggered fashion, respecting the VPC pairs.

Upgrading VPC pairs separately prevents total isolation of any servers connected to the fabric. Upgrading both peers simultaneously would cause a temporary connectivity loss for those servers, which is why option C is the optimal approach. Options A, B, and D incorrectly prioritize upgrading the APIC controllers at different steps or after the leaf switches which introduces the risk of control plane issues before the data plane is stable. Upgrading the fabric in such an order would not guarantee a stable control plane during the leaf switch upgrades and could result in unexpected behavior. Option C follows the best practices for staged upgrades.

For further research:

Cisco ACI Upgrade Guide:https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/system-management/latest/cisco-apic-system-management-guide/b-cisco-apic-system-management-guide-latest/m-upgrading-the-cisco-apic-software.html
Cisco ACI Best Practices:https://www.cisco.com/c/en/us/td/docs/dcn/aci/architecture/white-paper/cisco-aci-best-practices/cisco-aci-best-practices.html

Question: 66

Which protocol does ACI use to securely sane the configuration in a remote location?

A.SCP
B.HTTPS
C.TFTP
D.FTP

Answer: A

Explanation:

The correct answer is A, SCP (Secure Copy Protocol). ACI (Application Centric Infrastructure) employs SCP to securely back up and restore its configuration to a remote location. SCP provides secure file transfer over an encrypted SSH connection, ensuring data integrity and confidentiality during transmission. Unlike TFTP and FTP, which lack built-in security features, SCP protects the configuration data from eavesdropping and tampering. HTTPS, while secure for web traffic, is not the protocol ACI uses for configuration backups. SCP's secure connection and reliable data transfer make it ideal for managing sensitive configuration data in a datacenter environment. Cisco specifically implements SCP for this purpose within the ACI framework, adhering to best practices for data protection. This aligns with the general security principles of cloud infrastructure, where encryption of data in transit is essential. The ACI documentation and best practices emphasize utilizing secure protocols like SCP when dealing with configuration backups, underlining its importance in maintaining a secure and reliable environment.

Authoritative links for further research:

Cisco ACI Backup and Restore
Secure Copy Protocol (SCP)
Understanding Secure Shell (SSH) and Related Protocols

Question: 67

Which two protocols support accessing backup files on a remote location from the APIC? (Choose two.)

A.TFTP
B.FTP
C.SFTP
D.SMB
E.HTTPS

Answer: BC

Explanation:

The correct answer is B. FTP and C. SFTP. Cisco APIC allows for exporting configuration backups to remote locations. The protocols used for these transfers need to be reliable and capable of handling file transfers. FTP (File Transfer Protocol) is a standard network protocol for transferring files between a client and a server.

While FTP is widely used, it transmits data in plain text, making it inherently insecure. SFTP (Secure File Transfer Protocol), on the other hand, is a secure file transfer protocol. It uses SSH for encryption, providing a secure channel for data transmission, and thus protects sensitive data during transfer. These protocols are suitable for transferring configuration files that can be used in disaster recovery or regular backups. TFTP (Trivial File Transfer Protocol) is not suitable for backups because it lacks security features and reliable delivery mechanisms. SMB (Server Message Block) is typically used for shared file access on a network and is not typically used for backup processes initiated by the APIC. HTTPS (Hypertext Transfer Protocol Secure) is mainly for secure web browsing and communication, not for directly transferring files in the context of APIC backups. Thus, FTP and SFTP offer suitable options for APIC backups to remote locations, with SFTP being the preferred option due to its enhanced security.

Further research:

Cisco APIC Configuration Guide:https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/config/GUI/b_APIC_GUI_Configuration_Guide/b_APIC_GUI_Configuration_Guide_chapter_01010.html#concept_7
FTP:https://en.wikipedia.org/wiki/File_Transfer_Protocol
SFTP:https://en.wikipedia.org/wiki/SSH_File_Transfer_Protocol

Question: 68

Which attribute should be configured for each user to enable RADIUS for external authentication in Cisco ACI?

A.cisco-security domain
B.cisco-auth-features
C.cisco-aci-role
D.cisco-av-pair

Answer: D

Explanation:

The correct attribute to configure for each user to enable RADIUS for external authentication in Cisco ACI is **D. cisco-av-pair**. This is because the cisco-av-pair attribute allows administrators to pass specific Cisco ACI authorization parameters within the RADIUS authentication process. RADIUS, a widely used networking protocol, primarily handles authentication and authorization using Attribute-Value pairs (AV-pairs).

In the context of Cisco ACI, cisco-av-pair is the designated mechanism to communicate a user's assigned role and security domain information from the RADIUS server back to the ACI fabric. This allows the ACI fabric to correctly determine the user's privileges and access level based on the information received from the RADIUS server. Options A, B, and C are incorrect. cisco-security-domain is a security-related attribute, but it's not the attribute specifically used to pass authorization data within RADIUS. cisco-auth-features is not a recognized attribute within the Cisco ACI RADIUS context, and cisco-aci-role is the actual role assigned to a user but not the attribute used in the RADIUS context. Through cisco-av-pair, an administrator defines the specific role and security domain the user should have upon successful login. This ensures consistent and secure access control based on externally authenticated identities. Without properly configuring cisco-av-pair within the RADIUS server, users cannot be assigned the appropriate ACI roles and security domains during
authentication. This effectively renders external authentication through RADIUS ineffective for controlled access.

Further Reading:

Cisco ACI Security Configuration Guide
Cisco ACI External Authentication
RFC 2865 - Remote Authentication Dial In User Service (RADIUS) (for general RADIUS understanding)

Question: 69

```
aaa authentication login fallback
  realm radius
  group radius-1

aaa authentication login console
  realm radius
  group radius-1

aaa authentication login default
  realm radius
  group radius-1

aaa banner 'WELCOME TO ACI'
aaa group radius radius-1
  server 10.1.1.1 priority 0
  server 10.2.2.2 priority 1

aaa user default-role-no-login
```

Refer to the exhibit. Which action should be taken to ensure authentication if the RADIUS servers are unavailable?

A.Adjust the priority of server 10.1.1.1 to 1.
B.Assign the user to the default role.
C.Set the default login realm to LDAP.
D.Set the fallback login to local.

Answer: D

Explanation:

Set the fallback login to local.

Question: 70

A network engineer demonstrates Cisco ACI to a customer. One of the test cases is to validate a disaster recovery event by resetting the ACI fabric to factory and then restoring the fabric to the state it was in before the event. Which setting must be enabled on ACI to export all configuration parameters that are necessary to meet these requirements?

A.enabled AES encryption
B.generated a tech-support file
C.encrypted export destination
D.enabled JSON format export

Answer: A

Explanation:

The correct answer is **A. enabled AES encryption**. Here's why:

Cisco ACI's backup and restore process is critical for disaster recovery scenarios, allowing a fabric to be quickly rebuilt after a catastrophic failure. To fully capture and restore the state of an ACI fabric, including all configuration parameters, encryption is essential to protect the sensitive data in the backup file. When exporting configurations for backup, **AES encryption** is mandatory to protect sensitive information like passwords, user credentials, and other confidential configurations that are present in the configuration data. Simply having a file (option B) or using a JSON format (option D) doesn't inherently guarantee the security and completeness needed for a full restore, especially if the export location is not secure. Without encryption, the backup file itself might be vulnerable to unauthorized access, creating additional security risks. Encrypted export destination (option C) is relevant to where the file is stored rather than how it is created for the purpose of a full fabric restore. Therefore, **enabling AES encryption is the necessary security measure to safeguard the configuration data during export, ensuring it can be securely restored** to the ACI fabric during a disaster recovery event and is critical for successful restoration of a fabric to its previous state. This ensures the security and usability of the exported data when returning to the previous configuration. The lack of encryption compromises the integrity of the sensitive information, preventing a safe and secure restore.

For further research on Cisco ACI backup and restore, consult the following resources:

Cisco Application Centric Infrastructure Fundamentals:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/fundamentals/b_Cisco_ACI_Fundamental (Specifically, look for sections on backup/restore and data protection.)
Cisco ACI Configuration and Backup Best Practices:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b_ACI_Config_and_Backup.html (This document is a Cisco Knowledge Base article focusing on best practices.)
Cisco ACI APIC Configuration Export/Import:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/configuration/guide/b-aci-cfg/b-aci-cfg_chapter_011.html (This provides guidance in the APIC configuration guide.)

Question: 71

An engineer wants to filter the System Faults page and view only the active faults that are present in the Cisco ACI fabric. Which two lifecycle stages must be selected for filtering? (Choose two.)

A.Raised
B.Retaining
C.Soaking, Clearing
D.Raised, Clearing
E.Soaking

Answer: AE

Explanation:

The correct answer is A. Raised and E. Soaking. Cisco ACI uses a lifecycle model for faults to track their progression. When a fault is initially detected, it is in the 'Raised' state, indicating an active issue. Selecting 'Raised' will filter the faults page to show only these currently occurring problems. The 'Soaking' state represents a situation where a fault has been cleared but the system is monitoring to ensure it does not reoccur. This stage is still relevant for actively monitored faults. By selecting both 'Raised' and 'Soaking', the engineer can view the faults currently active and those that have recently been cleared. Options B 'Retaining' and C 'Soaking, Clearing' are not appropriate for filtering active faults because 'Retaining' typically means the fault is historical and no longer active, while the 'Soaking, Clearing' option is not a valid lifecycle stage

combination. Option D 'Raised, Clearing' is partially correct by including 'Raised' but incorrectly includes 'Clearing' which represents resolved faults. Therefore, selecting 'Raised' and 'Soaking' provides the necessary view of the active fault lifecycle.

Authoritative Links:

Cisco ACI Fault Management:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/fault-management/b-cisco-aci-fault-management/b-cisco-aci-fault-management_chapter_01.html (This Cisco document provides an overview of fault management in ACI, which includes lifecycle stages)
Cisco ACI Troubleshooting:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/troubleshooting/b-cisco-aci-troubleshooting/b-cisco-aci-troubleshooting_chapter_010.html (This document covers troubleshooting in Cisco ACI and may contain references to fault states)

Question: 72

An engineer must limit management access to the Cisco ACI fabric that originates from a single subnet where the NOC operates. Access should be limited to SSH and HTTPS only. Where should the policy be configured on the Cisco APIC to meet the requirements?

A.policy in the management tenant
B.ACL on the console interface
C.ACL on the management interface of the APIC
D.policy on the management VLAN

Answer: A

Explanation:

The correct answer is **A. policy in the management tenant.**

Here's why: Cisco ACI uses a policy-driven approach for network management. Access control is handled through policies defined within the ACI fabric itself, not directly on device interfaces like traditional networking. The "management tenant" is a dedicated tenant within the ACI fabric specifically designed for managing the ACI infrastructure, including the APIC controllers. Therefore, any access control rules for management traffic must be configured within this tenant.

Specifically, policies configured in the management tenant allow for fine-grained control over access to the APIC controllers. These policies can be applied to incoming connections based on source IP address, protocol (SSH, HTTPS), and user roles. By defining a policy that permits only SSH and HTTPS access from the NOC's subnet within the management tenant, the engineer effectively secures the APIC fabric management interface according to the requirements.

Options B, C, and D are incorrect: An ACL on the console interface is not used for network management access, but rather for local physical access to the device. An ACL on the management interface of the APIC implies configuring access directly on the device's interface, bypassing the ACI's policy-based control mechanism. Similarly, configuring policies on the management VLAN is incorrect because ACI uses the concept of tenants to segment resources, not VLANs, so access rules must be defined within the context of the management tenant.

The correct and secure way to control access within ACI is via ACI policies within the management tenant. This ensures the policy is centrally managed, consistent, and applied appropriately within the ACI framework.

Authoritative links for further research:

Cisco ACI Fundamentals:https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/fundamentals/b-aci-fundamentals.html
Cisco ACI Security Guide:https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/security/b-cisco-apic-security-guide.html
Cisco ACI Management Tenant Configuration:https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/gui/l3/b-aci-gui-l3-configuration/b-aci-gui-l3-configuration_chapter_01000.html

Question: 73

In the context of ACI Multi-Site, when is the information of an endpoint (MAC/IP) that belongs to site 1 advertised to site 2 using the EVPN control plane?

A.Endpoint information is not exchanged across sites unless COOP protocol is used.

B.Endpoint information is not exchanged across sites unless a policy is configured to allow communication across sites.

C.Endpoint information is exchanged across sites as soon as the endpoint is discovered in one site.

D.Endpoint information is exchanged across sites when the endpoints are discovered in both sites.

Answer: B

Explanation:

The correct answer is B: "Endpoint information is not exchanged across sites unless a policy is configured to allow communication across sites." In Cisco ACI Multi-Site, endpoint information (MAC/IP address mappings) is not automatically shared between sites via the EVPN control plane. This behavior is intentional to maintain isolation and control. By default, each ACI site operates as an independent fabric with its own endpoint database. To enable cross-site communication, administrators must explicitly define a contract that allows the sharing of specific endpoint information between sites. This is achieved by configuring inter-site policies that govern the propagation of endpoint reachability information using the Multi-Site Orchestrator (MSO).

Without such explicit policies, the EVPN control plane within each site remains isolated, preventing the leakage of endpoint data and maintaining site autonomy. This design ensures security and scalability by allowing administrators fine-grained control over the exchange of endpoint information, enabling targeted communication between specific applications or tenants across sites. Option C is incorrect because information isn't automatically exchanged; explicit configuration is needed. Option D is incorrect as cross-site endpoint discovery is not the trigger; rather, it's the policy that dictates information exchange. Option A is misleading; while COOP plays a role in some ACI architectures, it's not the primary factor governing EVPN cross-site endpoint sharing in Multi-Site scenarios.

Authoritative Links for Further Research:

Cisco ACI Multi-Site White Paper:https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-737528.html (Focus on Multi-Site Design and Policy)
Cisco ACI Multi-Site Configuration Guide:
https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/multisite/1x/configuration/cisco-aci-multisite-configuration-guide/m-overview.html (Explore Inter-site Communication & Policy)
Cisco ACI Multi-Site FAQ:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/faq-q-a-multisite.html (Review FAQs relating to cross-site communication).

Question: 74

Which statement regarding ACI Multi-Pod and TEP pool is true?

A.The IP addresses used in the IPN network can overlap TEP pool of the APIC.

B.A different TEP pool must be assigned to each Pod.

C.The Pod1 TEP pool must be split and a portion of the TEP pool allocated to each Pod.

D.The same TEP pool is used in all Pods.

Answer: B

Explanation:

Here's a detailed justification for why option B is correct, regarding ACI Multi-Pod and TEP pools:

Option B, "A different TEP pool must be assigned to each Pod," is the correct statement. In a Cisco ACI Multi-Pod deployment, each Pod is a physically separate ACI fabric that is interconnected through an Inter-Pod Network (IPN). To ensure proper communication between endpoints within different pods, a unique Tunnel Endpoint (TEP) pool is required for each Pod. TEP pools are used to encapsulate VXLAN packets for forwarding between nodes in the ACI fabric. Using the same TEP pool across multiple pods would create addressing conflicts and prevent correct routing. Each pod needs its own unique address space for both intra-pod and inter-pod communication. By assigning a unique TEP pool to each pod, we maintain segmentation, prevent overlap, and ensure proper functioning of the VXLAN tunnels. This segmentation ensures that traffic within each pod is isolated and routed correctly through the IPN when necessary. This design mirrors the logical separation of ACI fabrics when using Multi-Pod. The TEP pool is fundamentally an IP address allocation that allows communication between elements within the ACI fabric. Because of this, each separate fabric requires a separate and distinct allocation.

Option A is incorrect because the IPN network IP addresses must not overlap with TEP pools, as this would cause routing conflicts. Option C is incorrect as TEP pools are not divided, but assigned separately and completely to each Pod. Finally, option D is incorrect because using the same TEP pool in all pods would defeat the purpose of multi-pod design, as discussed earlier.

Authoritative Links for Further Research:

1. **Cisco ACI Multi-Pod White Paper:**https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-737855.html This document provides a detailed overview of ACI Multi-Pod architecture, including TEP pool usage.
2. **Cisco ACI Fabric L3 Out and Multi-Pod Deployment Guide:** https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/l3-config/cisco-aci-fabric-layer-3-external-networks-and-multi-pod-deployment-guide-5x.html This guide offers detailed explanations and configuration examples for ACI Multi-Pod, including the management of TEP pools.
3. **Cisco ACI Configuration Guide:** (Search for "Multi-Pod" and "TEP Pool") https://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/products-installation-and-configuration-guides-list.html This official documentation provides the most accurate information and is a great source of information to gain further insight.

Question: 75

Which two statements regarding ACI Multi-Site are true? (Choose two.)

A.The Multi-Site orchestrator must be directly attached to one ACI leaf.

B.Routers in the Inter-Site network must run OSPF, DHCP relay, and MP-BGP.

C.ACI Multi-Site is a solution that supports a dedicated APIC cluster per site.

D.ACI Multi-Site is a solution that allows one APIC cluster to manage multiple ACI sites.

E.The Inter-Site network routers should run OSPF to establish peering with the spines.

Answer: CE

Explanation:

Here's a detailed justification for why options C and E are correct regarding Cisco ACI Multi-Site, while A, B, and D are incorrect:

Option C is Correct: ACI Multi-Site is designed around the concept of having a dedicated Application Policy Infrastructure Controller (APIC) cluster for each individual ACI site. This architecture ensures that each site maintains its own independent control plane and fault domain. This enhances resilience and prevents a single point of failure from affecting multiple sites. Each APIC cluster manages its local ACI fabric, providing local control and policy enforcement.

Option E is Correct: The Inter-Site Network (ISN), responsible for connecting ACI sites, typically uses routers that establish peering with the ACI spine switches using Border Gateway Protocol (BGP). This BGP peering is crucial for exchanging routing information and enabling communication between endpoints across different ACI sites. OSPF, a routing protocol, is not mandatory for establishing peering with the spines in an ACI Multi-Site scenario.

Option A is Incorrect: The Multi-Site orchestrator, which is typically the Multi-Site Orchestrator (MSO), is not directly connected to a specific ACI leaf. It operates as a centralized management platform that interacts with the APIC clusters in each ACI site via the out-of-band network.

Option B is Incorrect: The ISN routers do not require OSPF or DHCP relay. While BGP is essential for peering and route exchange between sites, DHCP relay is generally not required on the ISN routers. MP-BGP (Multiprotocol BGP) is the typical protocol, particularly when exchanging VPNv4 or VPNv6 routes.

Option D is Incorrect: ACI Multi-Site does not allow one single APIC cluster to manage multiple ACI sites. Each ACI site operates with its own independent APIC cluster, which is key to the distributed control plane and fault isolation design principle. The MSO coordinates between APIC clusters.

In essence, ACI Multi-Site utilizes a distributed architecture with per-site APIC clusters to maintain control plane independence and fault tolerance, while the ISN utilizes BGP for inter-site routing. This approach ensures that a failure in one site does not propagate to other sites.

Authoritative Links for Further Research:

Cisco ACI Multi-Site Design Guide:
https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/multisite/2x/config/cisco-aci-multi-site-design-guide-2x.html
Cisco ACI Multi-Site White Paper:https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-741871.html

Question: 76

What are two requirements for the IPN network when implementing a Multi-Pod ACI fabric? (Choose two.)

A.EIGRP routing
B.PIM ASM multicast routing
C.BGP routing
D.VLAN ID 4
E.OSPF routing

Answer: DE

Explanation:

Here's a detailed justification for why the correct answers are D (VLAN ID 4) and E (OSPF routing) when implementing a Multi-Pod ACI fabric, focusing on the IPN (Inter-Pod Network) requirements:

The IPN serves as the crucial connection between different ACI pods in a Multi-Pod deployment. It's not an extension of the ACI fabric but a separate L3 network enabling communication across pods. This requires specific configurations for seamless operation.

Firstly, **OSPF routing (E)** is mandatory on the IPN. ACI leverages OSPF for internal communication, particularly for the Multipod Control Plane. This dynamic routing protocol ensures reachability of the spine nodes across pods, which is critical for Multi-Pod operation. The IPN routers must be configured with OSPF and be part of the ACI OSPF process. This facilitates the exchange of routing information for communication between different pods.

Secondly, **VLAN ID 4 (D)** is the predefined VLAN used by the ACI fabric specifically for Multi-Pod communication. All inter-pod traffic and control-plane communication will utilize this VLAN. This means that your IPN infrastructure needs to be configured to allow traffic tagged with VLAN 4 to pass through correctly.

This ensures that the internal fabric communication protocols can function correctly across the multiple PODs.

While options like EIGRP, PIM, and BGP are routing and multicast protocols, they aren't specifically required for IPN within ACI Multi-Pod configuration. BGP is only needed for external routing. PIM isn't needed internally within ACI and is only used for external multicast. EIGRP is not used by Cisco ACI for routing protocol needs. VLAN 4 is not an arbitrary value that you can change, it is the standard internal VLAN used for Multipod communications.

In summary, the IPN facilitates seamless communication between ACI pods, and it must have VLAN 4 transport along with OSPF routing for this to occur. Without these foundational elements, the Multi-Pod architecture won't function correctly.

Here are some authoritative links for further research:

1. **Cisco ACI Multi-Pod White Paper**: This document provides in-depth details about the design and implementation of Multi-Pod ACI, which highlights the requirement for VLAN 4 and OSPF on the IPN.https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-737305.html

2. **Cisco ACI Multi-Pod Configuration Guide**: This guide provides technical guidance on configuring Multi-Pod in ACI, focusing on the IPN requirements.https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/configuration/l3-multisite/cisco-apic-layer-3-multisite-config-guide-52x/m-multi-pod.html

These sources provide further insights and details that will support the correctness of the answer.

Question: 77

A Solutions Architect is asked to design two data centers based on Cisco ACI technology that can extend L2/L3, VXLAN, and network policy across locations. ACI
Multi-Pod has been selected. Which two requirements must be considered in this design? (Choose two.)

A.ACI underlay protocols, i.e. COOP, IS-IS and MP-BGP, spans across pods. Create QoS policies to make sure those protocols have higher priority.

B.A single APIC Cluster is required in a Multi-Pod design. It is important to place the APIC Controllers in

different locations in order to maximize redundancy and reliability.

C.ACI Multi-Pod requires an IP Network supporting PIM-Bidir.

D.ACI Multi-Pod does not support Firewall Clusters across Pods. Firewall Clusters should always be local.

E.Multi-Pod requires multiple APIC Controller Clusters, one per pod. Make sure those clusters can communicate to each other through a highly available connection.

Answer: BC

Explanation:

Here's a detailed justification for why options B and C are the correct requirements for a Cisco ACI Multi-Pod design, and why the other options are incorrect:

Correct Options:

B. A single APIC Cluster is required in a Multi-Pod design. It is important to place the APIC Controllers in different locations in order to maximize redundancy and reliability.

ACI Multi-Pod architecture is managed by a single, unified Application Policy Infrastructure Controller (APIC) cluster. This single cluster is the central point of management for the entire multi-pod fabric, encompassing all pods.

For high availability and resilience, it is strongly recommended to deploy the APIC controllers in separate physical locations. This provides tolerance against site failures. If one physical location housing an APIC becomes unavailable, the other APICs can still manage the fabric across all pods, ensuring uninterrupted operations.

Authoritative Link: Cisco ACI Multi-Pod White Paper, section on "Single APIC Cluster Deployment" [search on cisco.com or google for "Cisco ACI Multi-Pod White Paper"]

C. ACI Multi-Pod requires an IP Network supporting PIM-Bidir.

ACI Multi-Pod utilizes multicast for control plane communication across inter-pod networks (IPN).

Specifically, Protocol Independent Multicast - Bidirectional (PIM-Bidir) is required. PIM-Bidir is necessary for the VXLAN control plane to learn about endpoints in different pods. This allows for seamless L2/L3 extension and communication across pods.

The IPN is essential to provide the connectivity needed between the PODs for communication. PIM-Bidir provides for efficient multicast which enables the control plane to manage the various pods of the single fabric.

Authoritative Link: Cisco ACI Multi-Pod Configuration Guide, section on "IPN Requirements" [search on cisco.com or google for "Cisco ACI Multi-Pod Configuration Guide"]

Incorrect Options:

A. ACI underlay protocols, i.e. COOP, IS-IS and MP-BGP, spans across pods. Create QoS policies to make sure those protocols have higher priority.

While COOP, IS-IS and MP-BGP are core to the ACI fabric's operation, they do not span across pods in a Multi-Pod design. The IPN provides an overlay between the PODs. Each pod runs its own instance of these protocols internally. QoS priority is generally implemented for control plane traffic within the Pod itself but not across the IPN (inter-pod network).

D. ACI Multi-Pod does not support Firewall Clusters across Pods. Firewall Clusters should always be local.

While it is true that firewall clusters within a pod are often preferred due to latency and management, ACI Multi-Pod does support stretched firewall services across pods. This is done through service graph chaining and using appropriate load-balancing methods. However, there can be potential design considerations.

E. Multi-Pod requires multiple APIC Controller Clusters, one per pod. Make sure those clusters can communicate to each other through a highly available connection.

As explained in Option B, Multi-Pod requires a single APIC cluster, not separate clusters per pod. Each pod does not have its own APIC cluster.

In summary, a single APIC cluster manages the entire ACI Multi-Pod fabric, requiring its resilient deployment. The IPN that connects the PODs requires PIM-Bidir to support multicast for control plane communications.

Question: 78

An engineer configures a Multi-Pod system with the default getaway residing outside of the ACI fabric for a bridge domain. Which setting should be configured to support this requirement?

A.disable Limit IP Learning to Subnet
B.disable IP Data-plane Learning
C.disable Unicast Routing
D.disable Advertise Host Routes

Answer: C

Explanation:

The correct answer is **C. disable Unicast Routing**.

Here's the justification:

In a Cisco ACI Multi-Pod environment, when the default gateway for a bridge domain resides outside the ACI fabric, the ACI fabric needs a mechanism to forward traffic destined for that gateway. By default, ACI assumes that all routing is handled within the fabric. When the default gateway is external, ACI needs to be told not to perform routing decisions for this specific subnet inside the fabric. Disabling unicast routing for the bridge domain achieves this. This allows the ACI leaf nodes to simply forward traffic for that subnet based on Layer 2 information (MAC addresses) to the external router acting as the gateway.

Specifically:

A. disable Limit IP Learning to Subnet: This option affects how the ACI fabric learns IP addresses within the subnet, not how routing is performed. It doesn't address the issue of forwarding traffic to an external gateway.

B. disable IP Data-plane Learning: Disabling IP data-plane learning prevents the ACI fabric from dynamically learning IP addresses via data plane traffic, but it doesn't stop the fabric from attempting to route that traffic within the fabric.

D. disable Advertise Host Routes: This setting relates to how ACI advertises routes to external routing devices, not how it handles traffic within the fabric destined for an external gateway.

Disabling Unicast Routing means that the ACI fabric will not attempt to route packets based on IP address to devices inside this bridge domain. Instead, if the destination IP is not found in its internal routing table the traffic will be forwarded as L2 traffic to the Layer 2 interface. By having the default gateway outside of the fabric the traffic will be sent to the gateway device for routing.

This approach ensures proper traffic forwarding when a default gateway is located outside of the ACI fabric, allowing ACI to maintain its core function while integrating with external routing infrastructures.

Authoritative Links for Further Research:

Cisco ACI Multi-Pod Design:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/white-paper-c11-738182.html - This white paper provides a deep dive into ACI Multi-Pod architecture and related configurations.

Cisco ACI Bridge Domain Configuration:
https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/configuration/l2-configuration/cisco-apic-layer-2-configuration-52x.html - This document details the various bridge domain configurations, including options for Layer 2 forwarding and disabling unicast routing.

Question: 79

What do Pods use to allow Pod-to-Pod communication in a Cisco ACI Multi-Pod environment?

A.over Layer 3 directly connected back-to-back spines
B.over Layer 3 Out connectivity via border leafs
C.over Layer 3 IPN connectivity via border leafs
D.over Layer 3 IPN connectivity via spines

Answer: D

Explanation:

In a Cisco ACI Multi-Pod environment, Pods, which are logically separated ACI fabrics, need a mechanism to communicate with each other. Direct Layer 3 connections between spines (option A) are not the standard approach and lack the scalability required for inter-pod communication. Layer 3 Out connectivity via border leafs (option B) is used for external network access, not internal pod-to-pod communication.

The correct approach utilizes the IPN (Inter-Pod Network) for connectivity between Pods. This network is typically a Layer 3 infrastructure connecting the border leafs of different Pods. Therefore, pod-to-pod traffic doesn't flow directly via the spines (option A), or using the external L3Out directly (option B). Border leafs, which connect each Pod to the IPN, are crucial in facilitating this inter-pod communication (option C).

However, the traffic is not directly via those border leafs, but is sent over the **Layer 3 IPN connectivity via spines** (option D) that are connected to border leafs. This means the border leafs within each pod connect to spines within that pod which then transit through the IPN to other pods. This method ensures routed communication, scalability, and flexibility in connecting multiple ACI pods. The IPN is configured to route traffic based on IP subnets that are assigned to different ACI pods. Spines facilitate this routing process. The IPN provides a transit network between the pods. Option D is the standard design as recommended by Cisco.

Further research can be done on the Cisco website for detailed configuration, design, and deployment information for ACI multi-pod environments:

Cisco ACI Multi-Pod White PaperCisco ACI Multi-Pod Configuration Guide

Question: 80

An engineer must advertise a selection of external networks learned from a BGP neighbor into the ACI fabric. Which L3Out subnet configuration option creates an inbound route map for route filtering?

A.External Subnets for the External EPG
B.Shared Route Control Subnet
C.Import Route Control Subnet
D.Shared Security Import Subnet

Answer: C

Explanation:

The correct answer is **C. Import Route Control Subnet**.

Here's why: In Cisco ACI, L3Outs facilitate external network connectivity. To control which routes learned from external BGP neighbors are advertised into the ACI fabric, you use route maps. These route maps are applied at specific points within the L3Out configuration. The "Import Route Control Subnet" option specifically creates a subnet that, when associated with a route map, filters routes inbound to the ACI fabric. When a route is imported into the ACI fabric, the import route control subnet will filter what is allowed to be ingested. The route map associated with this subnet examines incoming BGP prefixes and either permits or denies them based on the defined criteria. The other options do not fulfill the need to filter inbound routes: "External Subnets for the External EPG" defines the subnets used for the external EPG, "Shared Route Control Subnet" is used for shared services, and "Shared Security Import Subnet" is more concerned with security policies than route filtering. Therefore, configuring the “Import Route Control Subnet” alongside an appropriate route map gives the network engineer fine-grained control over which routes enter the ACI fabric from external BGP peers.

Authoritative Resources:

Cisco ACI Layer 3 Outside Guide:https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/6x/l3-configuration/cisco-aci-layer-3-outside-configuration-guide-60x/cisco-aci-layer-3-outside-configuration-guide-60x_chapter_010.html
Cisco ACI Fundamentals:https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/6x/fundamentals/cisco-aci-fundamentals-60x/cisco-aci-fundamentals-60x_chapter_010.html