ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

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)

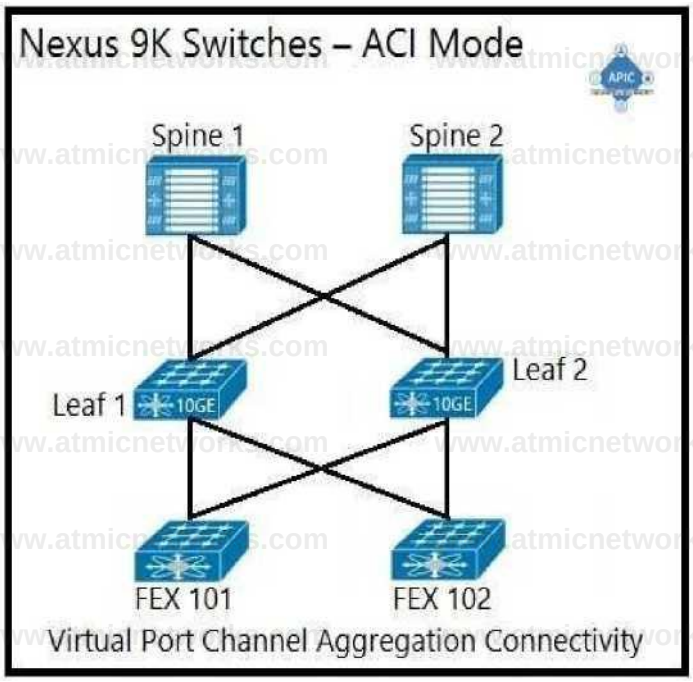

B)

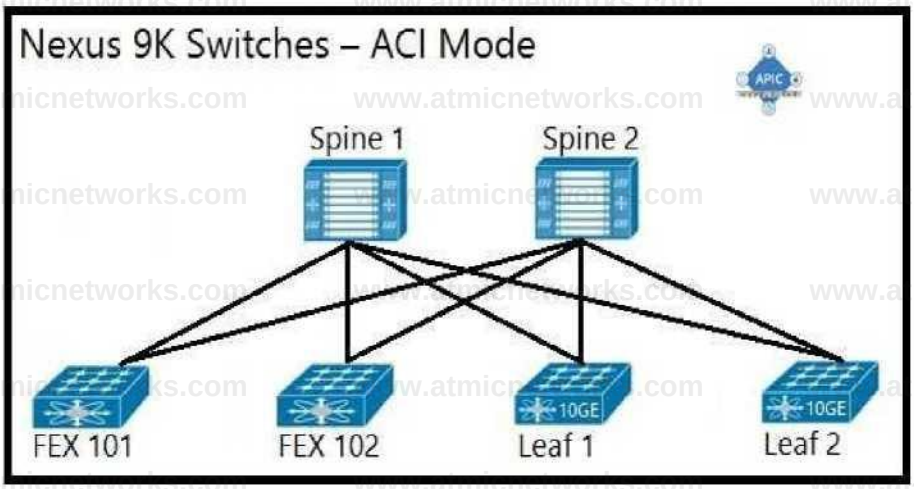
Nexus 9K Switches – ACI Mode
APIC
Spine 1
Spine 2
10GE
10GE
FEX 101
FEX 102
Leaf 1
Leaf 2

C)

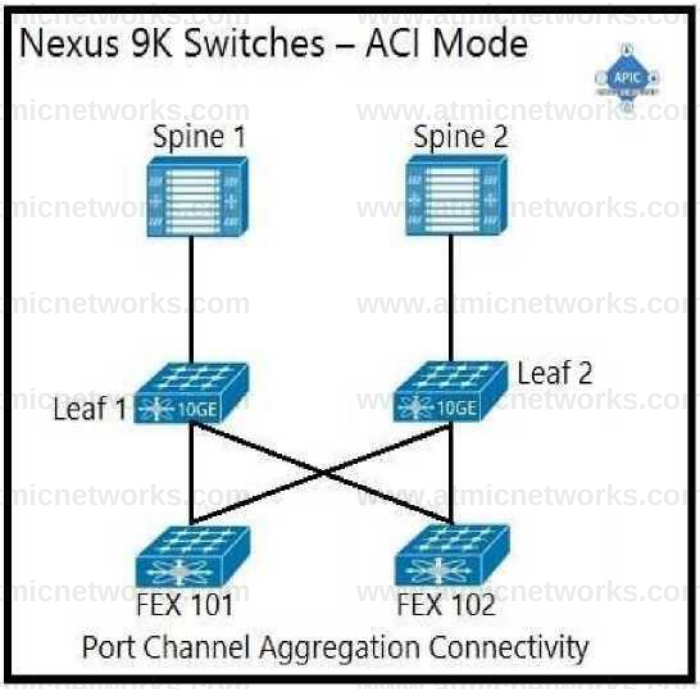
Nexus 9K Switches – ACI Mode
APIC
Spine 1
Spine 2
Leaf 1
10GE
Leaf 2
10GE
FEX 101
FEX 102
Port Channel Aggregation Connectivity

D)

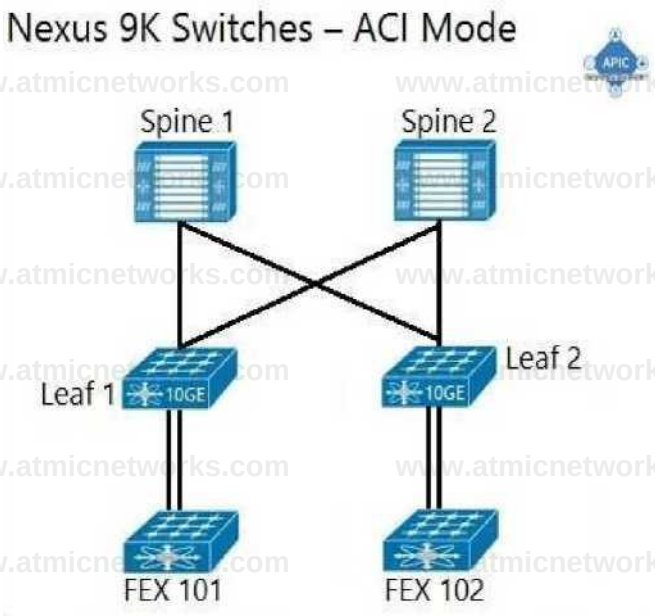

A. Option A

B. Option B

C. Option C

D. Option D

Answer: A

Explanation:

The supported physical topology for a Cisco ACI data center network that includes Cisco Nexus 2000 Series 10G fabric extenders is depicted in Option A. This topology illustrates a pair of spine switches connected to leaf switches, which are then connected to the fabric extenders. The Cisco Nexus 2000 Series Fabric Extenders act as remote line cards for a parent Cisco Nexus switch, extending the network fabric. The topology shown is a spine-leaf architecture, which is a scalable, high-bandwidth framework that is typical in ACI deployments.

Reference:

For a detailed understanding of the ACI fabric and supported topologies, you can refer to the Cisco Application Centric Infrastructure Design Guide, which provides comprehensive information on design

recommendations and deployment models: Cisco ACI Design Guide.

To learn more about the role of fabric extenders in the ACI architecture, the following resource offers insights into how they integrate with the ACI fabric: Understanding the ACI Fabric.

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:

To determine if a change in the behavior of the ACI fabric was due to human intervention, an ACI administrator should inspect the audit logs in the APIC UI. The audit logs provide a record of all user events, which includes actions performed by users, making it possible to trace any changes back to specific human activities1.

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 import mode that must be used to ensure that the import process terminates if the imported configuration is incompatible with the existing system is the 'atomic' mode. This mode ignores shards that contain objects that cannot be imported while proceeding with shards that can be imported. If the incoming configuration's version is incompatible with the existing system, the import process will terminate2.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/aci- fundamentals/ Cisco-ACI-Fundamentals-401/Cisco-ACI-Fundamentals-401_chapter_01011.html

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: D

Explanation:

For the BGP Route Reflector policy to take effect, the components that must be configured are the leaf fabric interface overrides and profiles. These configurations are necessary to establish the route reflector relationships within

the BGP protocol on the leaf switches

Reference: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/L3- configuration/Cisco-APIC-Layer-3-Networking-Configuration-Guide-401/Cisco-APIC-Layer-3-Networking-Configuration-Guide- 401_chapter_01.html

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 type of policy that configures the suppression of faults generated from a port being down is the 'fault severity assignment' policy. This policy allows administrators to change the severity of a fault or suppress it altogether, which can be useful for managing expected faults and reducing noise from non-critical events

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/faults/guide/b_APIC_Faults_Errors/b_IFC_Faults_Errors_chapter_01.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:

To deploy an access port policy group within Cisco ACI, an attachable entity profile (AEP) needs to be created. An AEP is a policy construct that groups together various domains and allows them to be attached to the fabric access layer

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_0110.html

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:

To prevent a fault from appearing that is not directly relevant to the environment, the ACI administrator should create a fault severity assignment policy that hides the fault. This can be done by selecting "Ignore Fault" under System -> Faults in the APIC UI2

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 Cisco AV Pair resolves to a managed object class. In the context of Cisco ACI, the AV Pair is used to define specific

attributes for RADIUS users, which then map to managed objects within the ACI fabric3.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/Security_config/b_Cisco_APIC_Security_Configuration_Guide/b_Cisco_APIC_Security_Guide_chapter_01011.html

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.

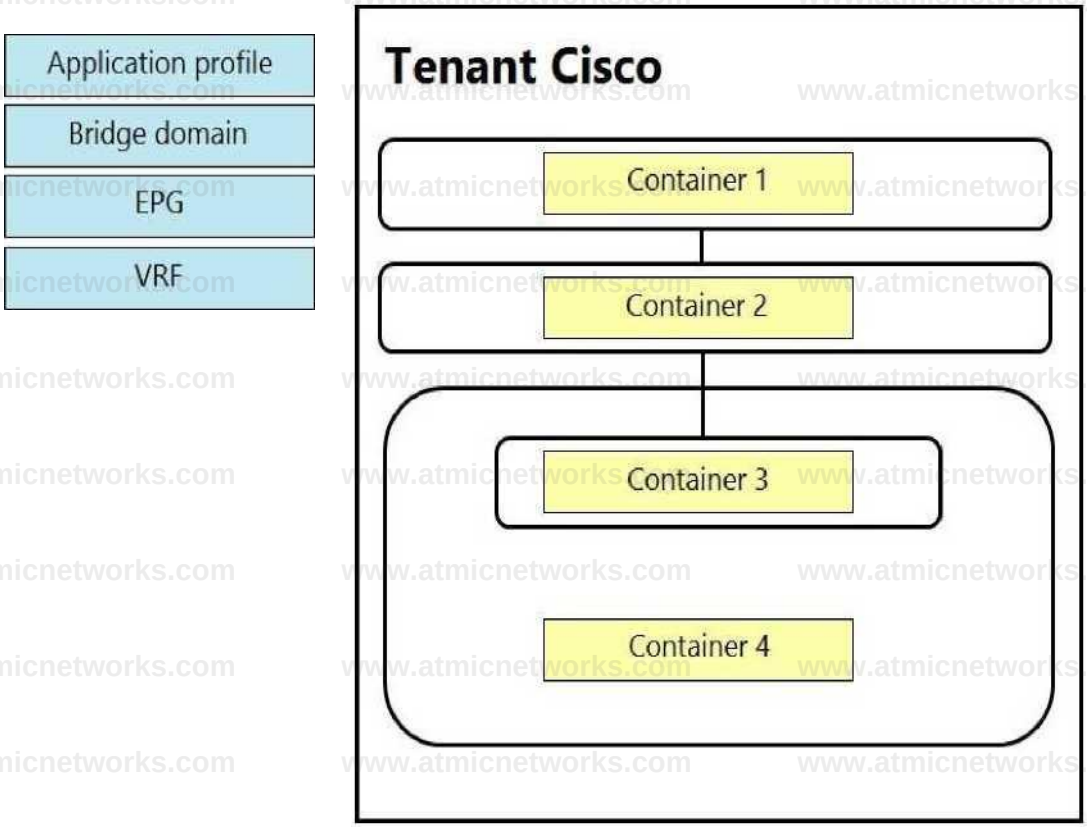

Answer:

Explanation:

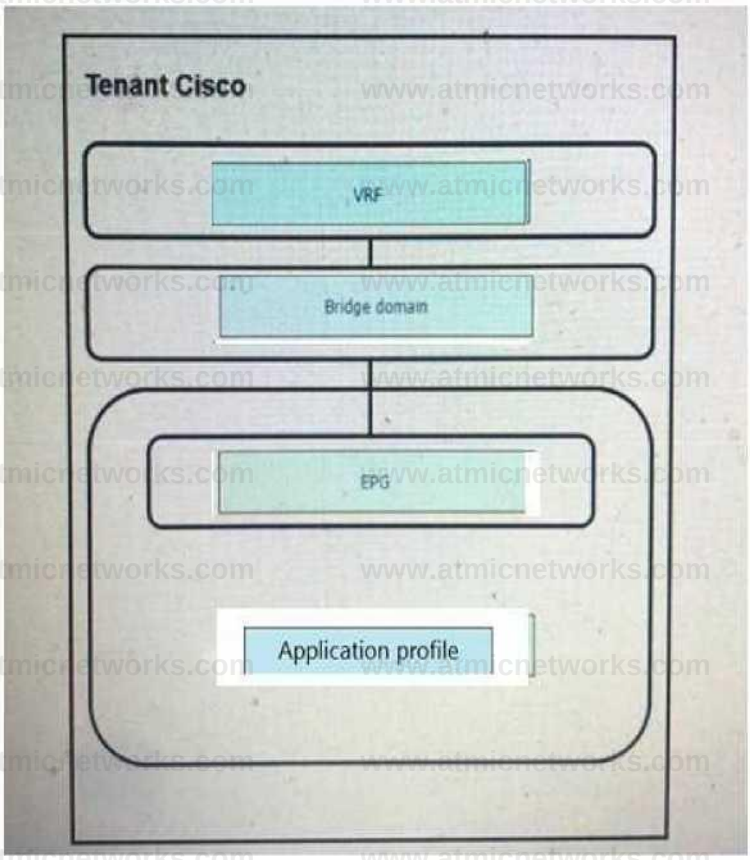

Question: 10

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:

The feature that dynamically assigns or modifies the EPG association of virtual machines based on their attributes is uSeg EPGs. uSeg EPGs allow for microsegmentation within the Cisco ACI environment, enabling the dynamic assignment of VMs to EPGs based on attributes such as VM name, tag, or other identifiers

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: A

Explanation:

The feature that allows firewall ACLs to be configured automatically when new endpoints are attached to an EPG is ARP gleaning. This feature helps in identifying the endpoints and applying the necessary ACLs to them as they are discovered

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: B

Explanation:

In a Cisco ACI configuration with IEEE 802.1p ports, the traffic exits from the EPG untagged from leaf ports. This means that when the port is configured in Access (802.1p) mode and the access VLAN is the only VLAN deployed on the port, then traffic will be untagged on egress56.

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:

An EPG is extended outside of the ACI fabric by statically assigning a VLAN ID to a leaf port in an EPG. This method maps the traffic received on the leaf port to the EPG, and the policy for this EPG is enforced

Reference: https://www.dclessons.com/l2-external-network-with-aci

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.

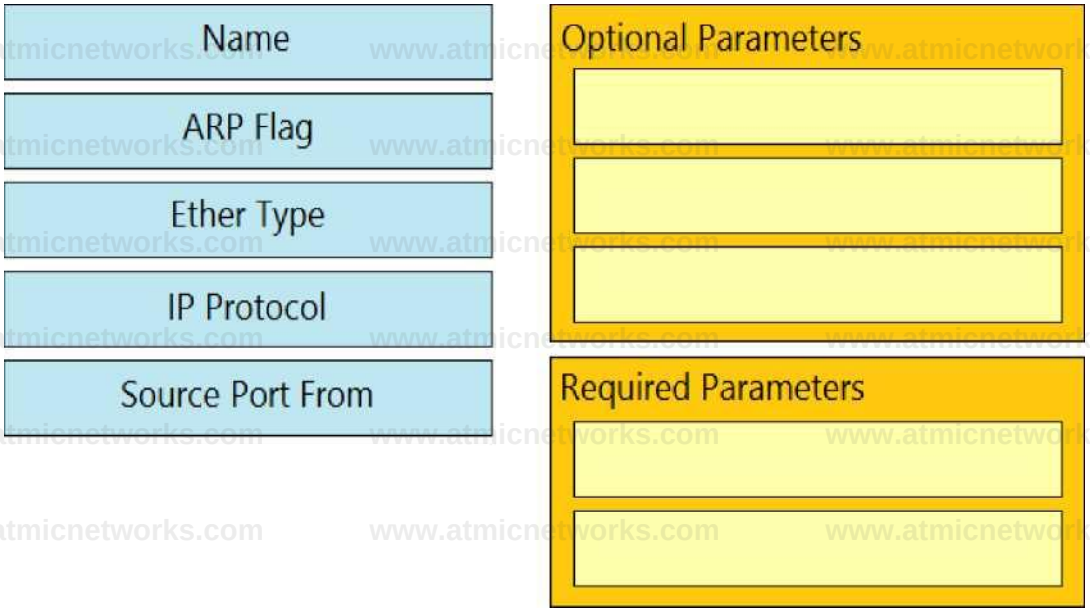

Answer:

Explanation:

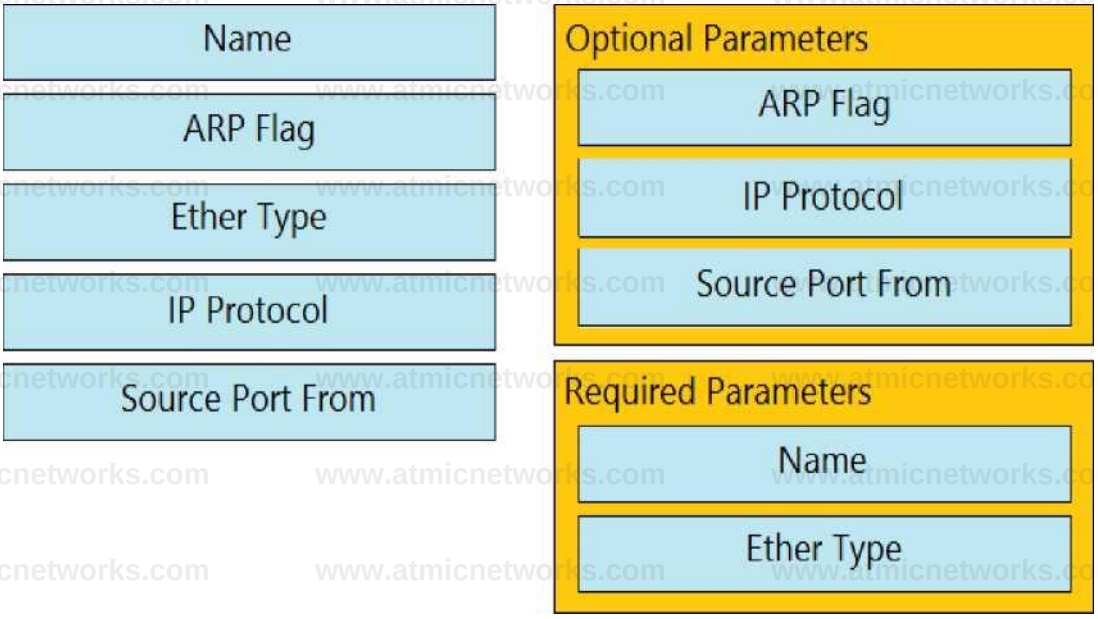

Reference: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/Operating_ACI/guide/ b_Cisco_Operating_ACI/b_Cisco_Operating_ACI_chapter_01000.html

Question: 15

Where is the COOP database located?

A. leaf

B. spine

C. APIC

D. endpoint

Answer: B

Explanation:

The COOP (Council Of Oracle Protocol) database is located on each spine switch within the Cisco ACI fabric.

The COOP database is responsible for maintaining a consistent copy of endpoint address and location

information across the fabric

Reference: https://www.cisco.com/c/en/us/solutions/collateral/data-center- virtualization/application-centric- infrastructure/white-paper-c11-739989.html

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 initial APIC cluster discovery process involves each APIC using an internal private IP address from a pool to communicate with the nodes and other APICs in the cluster. The APICs discover the IP addresses of other APIC controllers through an LLDP-based discovery process

Reference: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/aci- fundamentals/b_ACI-Fundamentals/b_ACI-Fundamentals_chapter_010011.html

Question: 17

Refer to the exhibit.

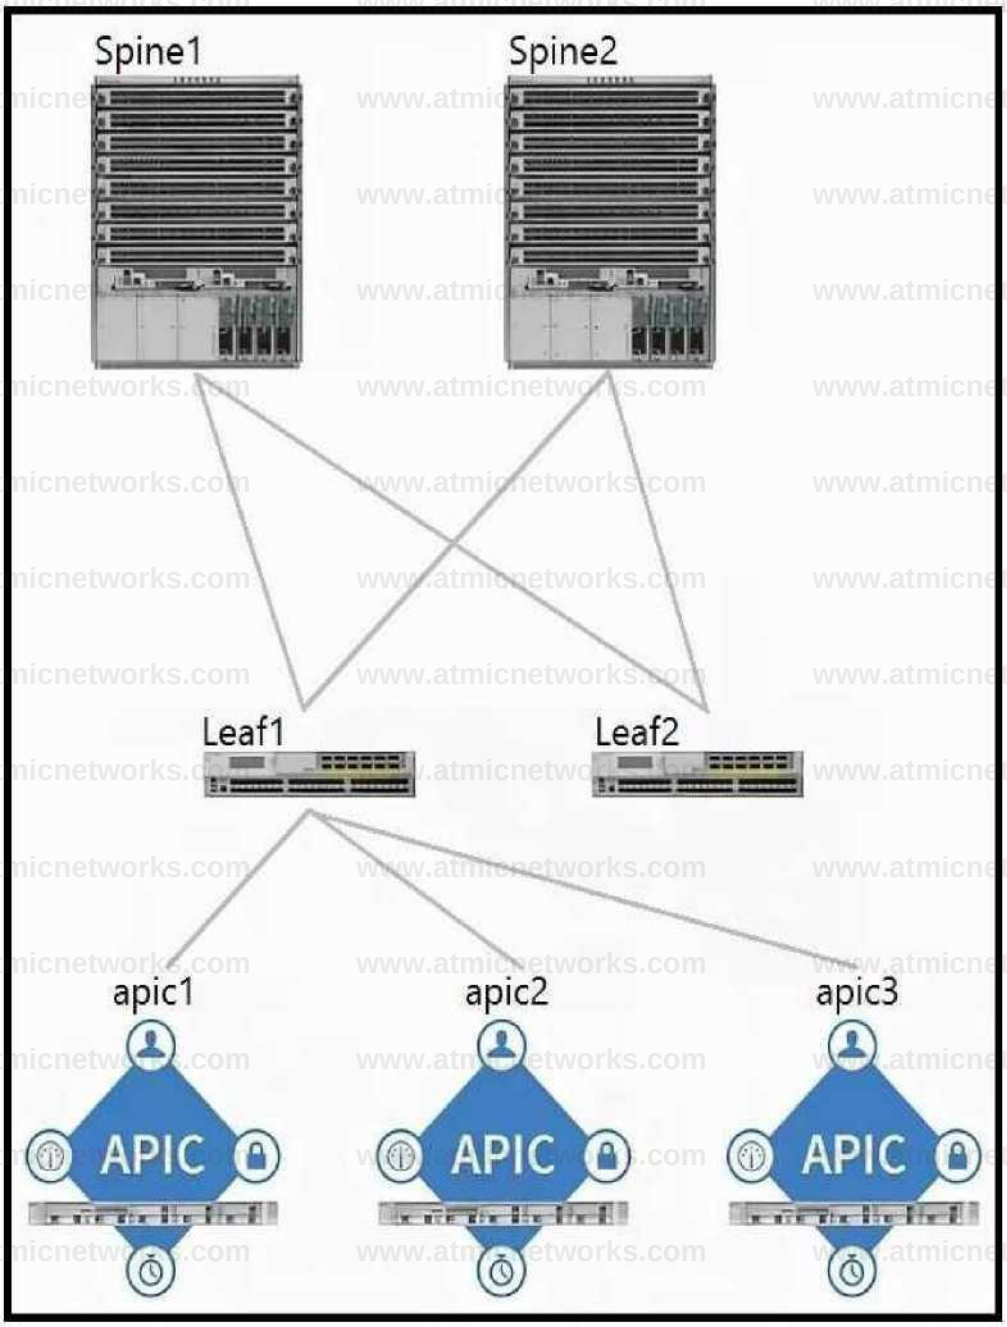

Which two components should be configured as route reflectors in the ACI fabric? (Choose two.)

A. Spine1

B. apic1

C. Spine2

D. Leaf1

E. Leaf2

F. apic2

Answer: A, C

Explanation:

In a Cisco ACI fabric, the spine switches are typically configured as route reflectors. This is because spine switches are central to the fabric's architecture and are responsible for interconnecting all the leaf switches, making them ideal for reflecting BGP routes within the fabric. Therefore, the components that should be configured as route reflectors are:

Spine1 (Option A)

Spine2 (Option C)

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:

When creating a subnet within a bridge domain in Cisco ACI, the configuration option used to specify the network visibility of the subnet is the scope12. The scope can be set to private, public, or shared, determining how the subnet is advertised and accessed within and outside of the ACI fabric12.

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 construct345. It defines a unique Layer 2 MAC address space and can include one or more subnets. Bridge domains are associated with a VRF and can contain multiple EPGs345.

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 table that holds IP address, MAC address, and VXLAN/VLAN information on a Cisco ACI leaf is the endpoint table6. This table is essential for the ACI fabric's ability to track and apply policies to endpoints.

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: B, E

Explanation:

On border leaf switches in Cisco ACI, the two types of interfaces supported to connect to an external router are subinterface with 802.1Q tagging and Switch Virtual Interface (SVI)7. These interfaces allow for the segmentation of traffic and the extension of Layer 2 and Layer 3 services to external networks7.

Question: 22

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:

When extending an EPG out of the ACI fabric using static path binding, it is true that external endpoints are in the same EPG as the directly attached endpoints8. This means that traffic received on a statically bound leaf port with a specific VLAN ID is mapped to the EPG, and the same policies applied to directly attached endpoints are enforced on the external endpoints8.

Question: 23

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 setting that prevents the learning of Endpoint IP addresses whose subnet does not match the bridge domain subnet is the "Limit IP learning to subnet" setting within the bridge domain1. This setting ensures that only IP addresses belonging to the subnets configured on the bridge domain are learned, preventing mis-

learning of IP addresses that may not belong to the fabric1.

Question: 24

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:

On the egress leaf switch, when traffic is received from an L3Out, no source MAC or IP address of the traffic is learned as a remote endpoint1. The Cisco ACI fabric does not learn the IP addresses from the data plane in an L3Out domain; instead, it uses ARP to resolve next-hop IP and MAC relationships to reach the prefixes behind external routers1.

Question: 25

Refer to the exhibit.

```
<fvTenant name="ACILab">
    <fvCtx name="pvnl"/>
    <fvBD name="bdl">
        <fvRsCtx tnFvCtxName="pvnl"/>
          <fvSubnet ip="10.1.100.1/24"/>
    </fvBD>
```

```
</fvTenant>
```

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: C, E

Explanation:

The XML configuration snippet provided in the exhibit results in the creation of two objects within the Cisco ACI environment:

Bridge Domain (Option C): This is indicated by the XML element <fvBD name="bd1">, which defines a bridge domain with the name “bd1”.

VRF (Option E): This is shown by the XML element <fvCtx name="pvn1"/>, which defines a Virtual Routing and Forwarding instance named “pvn1”.

These objects are essential for network segmentation and routing within the ACI fabric, where the bridge domain represents a Layer 2 broadcast domain and the VRF is used for Layer 3 routing separation.

Question: 26

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:

To have the endpoint table learn the IP addresses of endpoints in a bridge domain, unicast routing must be enabled1. This allows the bridge domain to learn the IP addresses of endpoints

through the data plane1.

Question: 27

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: A

Explanation:

When extending EPG connectivity to an external network that houses the Layer 3 gateway and other end hosts, the ACI bridge domain configuration should use Forwarding: Custom, L2 Unknown Unicast: Hardware Proxy, L3 Unknown Multicast Flooding: Flood, Multi Destination Flooding: Flood in BD, and ARP Flooding: Enabled23. This configuration ensures that unknown unicast traffic is handled efficiently while still allowing ARP flooding, which is necessary for the ACI fabric to learn about endpoints on the external network23.

Question: 28

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:

When a bridge domain is configured with the hardware-proxy option for Layer 2 unknown unicast traffic, if the spine switch is unable to find the MAC address in the proxy database, it will drop the Layer 2 unknown unicast packet4. This behavior is controlled by the hardware proxy option associated with a bridge domain4.

Question: 29

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:

Changing the Layer 2 unknown unicast policy from Flood to Hardware Proxy results in the traffic being forwarded to one of the spines to perform as a spine proxy4. If the spine proxy also does not know the address, the packet is discarded4.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/L2_config/b_Cisco_APIC_Layer_2_Configuration_Guide/b_Cisco_APIC_Layer_2_Configuration_Guide_chapter_010.html

Question: 30

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:

To set Layer 2 loop migration in an ACI Fabric with a Layer 2 Out configured, MCP (Misconfiguration Protocol) must be enabled on the ACI fabric56. MCP helps to prevent Layer 2 loops by detecting misconfigurations that could potentially cause loops56.

Question: 31

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: B, D

Explanation:

When implementing a connection that represents an external bridged network, the configurations used are Layer 2 outside (Option B) and Static path binding (Option D)12. Layer 2 outside is typically used to connect a bridged external network trunk switch to a leaf switch in the ACI fabric, while static path binding is used to assign specific ports on a leaf switch to an external bridged network12.

Question: 32

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: D, E

Explanation:

The two actions that extend a Layer 2 domain beyond the ACI fabric are extending the bridge domain out of the ACI fabric (Option D) and extending the EPG out of the ACI fabric (Option E)34. Extending the bridge domain involves creating a Layer 2 outside connection, while extending the EPG involves statically assigning a port with a VLAN ID to an EPG34.

Question: 33

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: A

Explanation:

When Cisco ACI connects to an outside Layer 2 network, the ACI fabric floods the STP BPDU frame within the bridge domain (Option A)5. This ensures that BPDUs are properly propagated within the relevant VLAN encapsulation associated with the bridge domain5.

Reference: https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric- infrastructure/white-paper-c07-732033.html#_Toc395143571

Question: 34

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: B, D

Explanation:

Storm control should be configured on port channel on a single leaf switch (Option B) and endpointfacing trunk interface (Option D) to protect against broadcast traffic6. These interfaces are where storm control policies are typically applied to prevent disruptions caused by traffic storms6.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/L2_config/b_Cisco_APIC_Layer_2_Configuration_Guide/b_Cisco_APIC_Layer_2_Configuration_Guide_chapter_01010.html

Question: 35

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: A, E

Explanation:

The two dynamic routing protocols supported when using Cisco ACI to connect to an external Layer 3 network are iBGP (Option A) and eBGP (Option E)78. These protocols are included in the routing protocol options for a Layer 3 external outside network configuration in ACI

Reference: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/ACI_Best_Practices/ b_ACI_Best_Practices/b_ACI_Best_Practices_chapter_010010.html

Question: 36

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: A

Explanation:

To redistribute externally learned OSPF routes within the ACI fabric, a Route Control Profile must be configured1. This profile allows for the control of route redistribution and the application of route maps to influence routing decisions within the fabric1.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/ACI_Best_Practices/ b_ACI_Best_Practices/b_ACI_Best_Practices_chapter_010010.html

Question: 37

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: A

Explanation:

The statement about the MTU value of MP-BGP EVPN control plane packets in Cisco ACI that is true for communication between spine nodes in different sites is that 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 packets2. This ensures that the control plane traffic, which is not VXLAN encapsulated, can be properly handled across the sites2.

Question: 38

Refer to the exhibit.

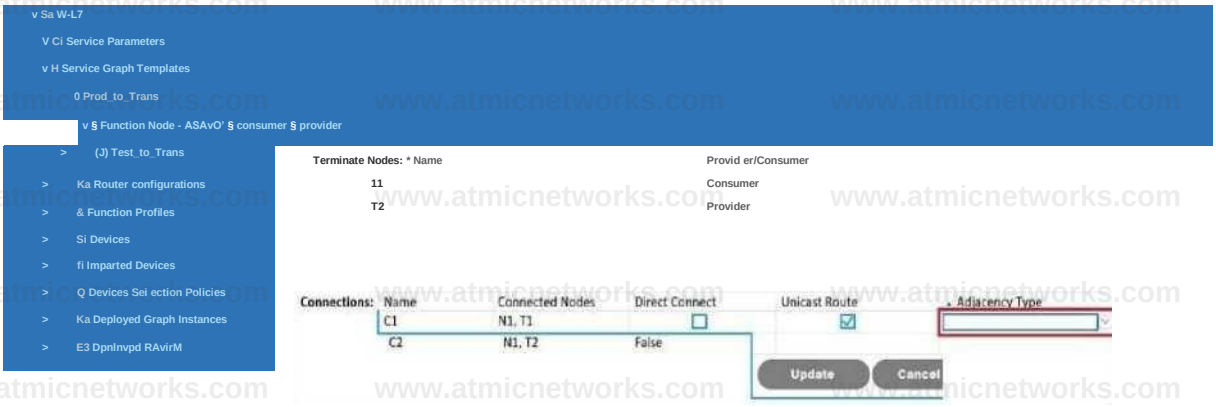

Which Adjacency Type value should be set when the client endpoint and the service node interface are in a different subnet?

A. Routed

B. Unicast

C. L3Out

D. L3

Answer: A

Explanation:

When the client endpoint and the service node interface are in different subnets, the Adjacency Type value should be set to Routed. This setting is used to indicate that routing is required between the client endpoint and the service node interface, as they are not in the same Layer 2 broadcast domain and need Layer 3 routing to communicate.

Question: 39

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:

In the Cisco ACI VMM domain configuration, to generate port groups with names formatted as "Tenant=Application=EPG", the configuration option used is delimiter. The delimiter is a character that separates the elements of the port-group name. By changing the delimiter setting to "=", the port-group names will be generated with the desired format.

Question: 40

Refer to the exhibit.

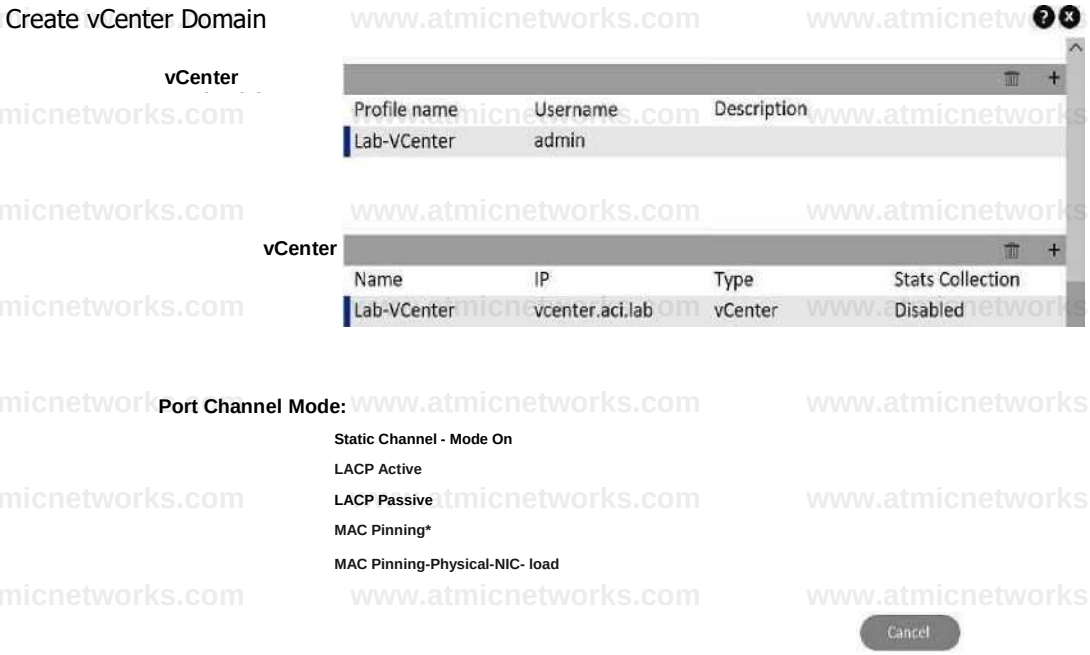

An engineer is implementing Cisco ACI – VMware vCenter integration for a blade server that lacks support of bonding. Which port channel mode results in "route based on originating virtual port" on the VMware VDS?

A. Static Channel – Mode On

B. MAC Pinning-Physical-NIC-load

C. LACP Passive

D. MAC Pinning+

E. LACP Active

Answer: B

Explanation:

For a blade server that lacks support of bonding, the port channel mode that results in "route based on originating virtual port" on the VMware VDS is MAC Pinning-Physical-NIC-load (Option B). This mode ensures that the virtual machine traffic is distributed across the available uplinks based on the MAC address of the originating virtual port1.

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: 41

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 configuring Cisco ACI VMM domain integration with VMware vCenter, the object that is created in vCenter is a VMware vSphere Distributed Switch (VDS)12. The VDS is a virtual switch that extends across all esxi hosts in the cluster and provides a centralized interface from which network configurations can be managed12. This switch allows network administrators to manage networking for multiple hosts and virtual machines from a single interface within vCenter12.

Question: 42

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.

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

Answer:

Explanation:

concrete interfaces	VRF stitching
service graph	service graph
device cluster	device cluster
VRF stitching	concrete interfaces

Question: 43

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:

When the VMM resolution immediacy is set to pre-provision in a Cisco ACI environment, policies are downloaded to the ACI leaf switch regardless of Cisco Discovery Protocol neighborship1. This setting ensures that the policies are in place on the leaf switches even before a VM controller is attached to the virtual switch1.

Question: 44

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:

In the context of VMM, the protocol between ACI leaf and compute hosts that ensures that the policies are pushed to the leaf switches for immediate and on-demand resolution immediacy is LLDP (Link Layer Discovery Protocol)2. LLDP is used for discovery and to ensure that the policies are applied as needed for the virtual machines2.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/aci-fundamentals/b_ACI-Fundamentals/b_ACI-Fundamentals_chapter_01011.html

Question: 45

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:

When configuring in-band management IP addresses for Cisco APICs, leaf nodes, and spine nodes, the tenant used is the mgmt tenant1.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b_KB_Configuring_Static_Management_Access.html#concept_CFF63FEBE947424291B0F10E6F23DA7D

Question: 46

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: A

Explanation:

The unique identifier of an ACI object is represented by the universal resource identifier (URI)2.

Reference: https://www.slideshare.net/CiscoDevNet/introduction-to-aci-apis

Question: 47

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: D

Explanation:

When configuring in-band management, the new construct that must be created is a bridge domain1.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b_KB_Configuring_Static_Management_Access.html

Question: 48

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:

To allow SNMP traffic on the APIC controller, a contract under tenant mgmt must be configured3.

Question: 49

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: A

Explanation:

The type of port used for in-band management within ACI fabric is the spine switch port4.

Question: 50

Refer to the exhibit.

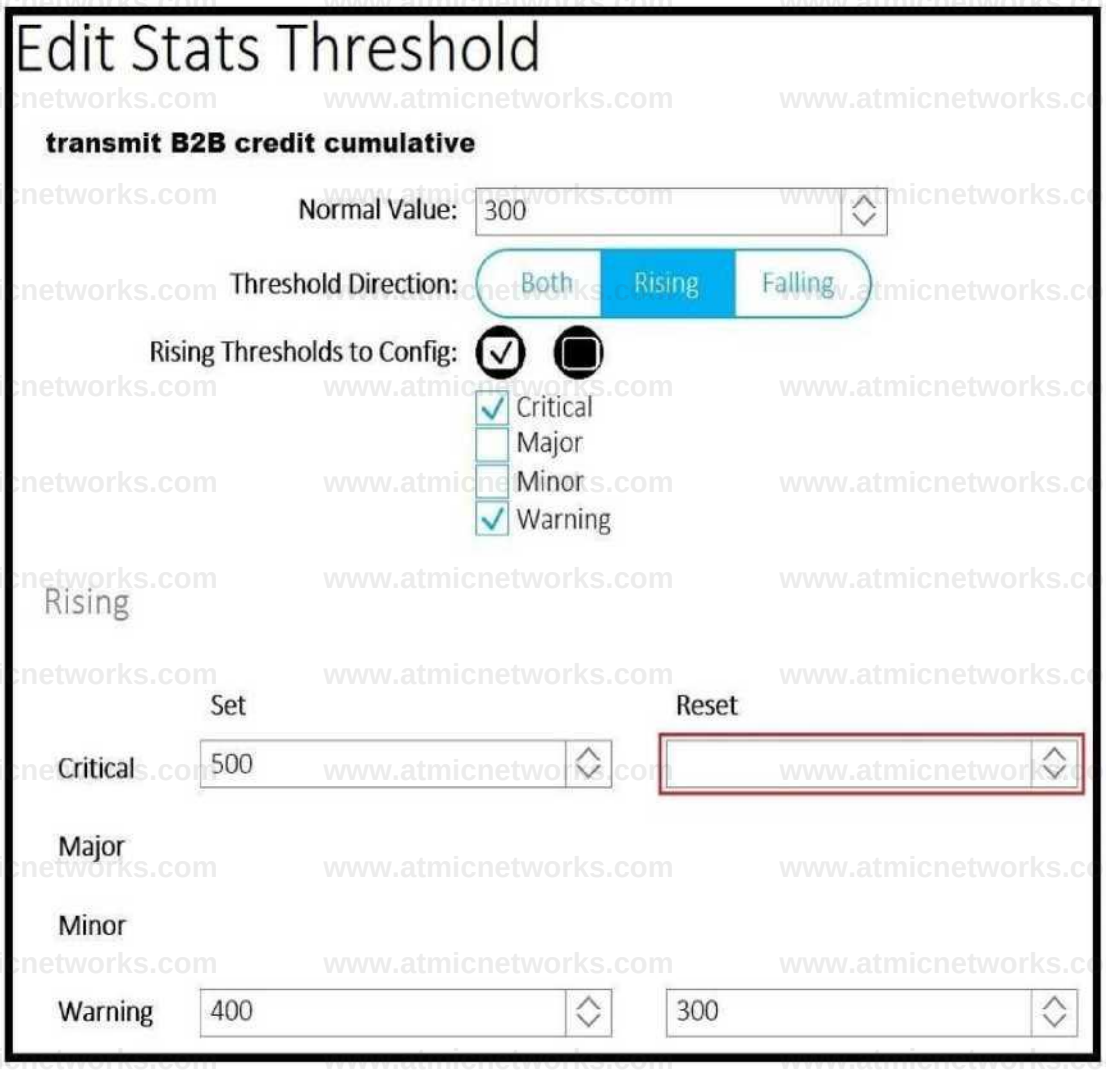

A client reports that the ACI domain connectivity to the fiber channel storage is experiencing a B2B credit oversubscription. The environment has a SYSLOG server for state collection messages. Which value should be chosen to clear the critical fault?

A. 300

B. 410

C. 350

D. 510

Answer: D

Explanation:

To clear the critical fault related to B2B credit oversubscription in the ACI domain connectivity to the fiber channel storage, the value that should be chosen is 5101. This value is set to address the issue of B2B credit oversubscription and clear the critical fault as reported by the client1.

Question: 51

Which statement about ACI syslog is true or Which statement describes the ACI syslog?

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: C

Explanation:

The correct statement about ACI syslog is that all syslog messages are sent to the destination through APIC1. This centralized approach allows for consistent logging and monitoring across the ACI fabric1.

Reference: https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2016/pdf/BRKACI-2303.pdf

Question: 52

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?

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

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

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

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

A. Option A

B. Option B

C. Option C

D. Option D

Answer: A

Explanation:

To minimize possible traffic impact during the upgrade of an ACI fabric with 3 APIC controllers and servers dual-homed to pairs of leaf switches configured in VPC mode, the fabric should be upgraded following Option A2. This option involves creating two maintenance groups for the APIC controllers, upgrading one group at a time, and then proceeding with the leaf switches2. This staged approach helps ensure that not all paths are affected simultaneously, thereby reducing the potential for traffic disruption2.

Question: 53

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:

ACI uses the SCP (Secure Copy Protocol) to securely save the configuration in a remote location. SCP is a network protocol that supports file transfers and relies on the Secure Shell (SSH) protocol to provide security for the transferred data.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b_KB_Using_Import_Export_to_Recover_Config_States.html

Question: 54

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: B, C

Explanation:

The two protocols that support accessing backup files on a remote location from the APIC are FTP (File Transfer Protocol) and SFTP (SSH File Transfer Protocol). Both FTP and SFTP allow for the transfer of files to and from a remote location, but SFTP provides an additional layer of security by utilizing SSH.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/troubleshooting/ b_APIC_Troubleshooting/b_APIC_Troubleshooting_appendix_010011.html

Question: 55

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 attribute that should be configured for each user to enable RADIUS for external authentication in Cisco ACI is cisco-av-pair. This attribute allows for the assignment of roles and permissions to users authenticated via RADIUS.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/Security_config/
b_Cisco_APIC_Security_Configuration_Guide/b_Cisco_APIC_Security_Guide_chapter_01011.html

Question: 56

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: C

Explanation:
In the context of ACI Multi-Site, the information of an endpoint (MAC/IP) that belongs to site 1 is advertised to site 2 using the EVPN control plane as soon as the endpoint is discovered in one site. This allows for immediate visibility of the endpoint across sites.

Reference: https://www.cisco.com/c/en/us/solutions/collateral/data-center- virtualization/application-centric- infrastructure/white-paper-c11- 739609.html#CiscoACIMultiSiteoverlaydataplane

Question: 57

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: D

Explanation:

The true statement regarding ACI Multi-Pod and TEP pool is that the same TEP pool is used in all Pods. This shared TEP pool ensures consistent addressing and simplifies the management of the overlay network across multiple Pods.

Reference: https://www.cisco.com/c/en/us/solutions/collateral/data-center- virtualization/application-centric- infrastructure/white-paper-c11-739714.html

Question: 58

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: C, D

Explanation:

The two true statements regarding ACI Multi-Site are:

ACI Multi-Site is a solution that supports a dedicated APIC cluster per site1.

ACI Multi-Site is a solution that allows one APIC cluster to manage multiple ACI sites1.

Question: 59

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: B, E

Explanation:

The two requirements for the IPN network when implementing a Multi-Pod ACI fabric are:

PIM ASM multicast routing2.

OSPF routing3.

Question: 60

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: B, C

Explanation:

When designing two data centers based on Cisco ACI technology that can extend L2/L3, VXLAN, and network policy across locations with ACI Multi-Pod, the two requirements that must be considered are:

A single APIC Cluster is required in a Multi-Pod design4.

ACI Multi-Pod requires an IP Network supporting PIM-Bidir2.

Question: 61

An engineer must limit management access to me 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. policy on the management VLAN

C. ACL on the management interface of the APIC

D. ACL on the console interface

Answer: A

Explanation:

To limit management access to the Cisco ACI fabric that originates from a single subnet where the NOC operates, the policy should be configured in the management tenant on the Cisco APIC5.

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_0111.html

Question: 62

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 spines

D. over Layer 3 IPN connectivity via border leafs

Answer: C

Explanation:

In a Cisco ACI Multi-Pod environment, Pods use over Layer 3 IPN connectivity via spines to allow Pod-to-Pod communication6.

https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-737855.html

Question: 63

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

A. VMM domain profile

B. EPG static port binding

C. Layer 3 outside interface association

D. IP address pool association

E. EPG association

Answer: A, E

Explanation:

The two essential components of a Cisco ACI Virtual Machine Manager (VMM) domain policy configuration are:

VMM domain profile1.

EPG association1.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/aci-fundamentals/b_ACI-Fundamentals/b_ACI-Fundamentals_chapter_01011.html#concept_74EFC437C0AA44A391676F70ACE59DF3

Virtual Machine Manager Domain Main Components

AOI fabric virtual machine manager (VMM) domains enable an administrator to configure connectivity policies for virtual machine controllers. The essential components of an ACI VMM domain policy include the following:

- **Virtual Machine Manager Domain Profile—**Groups VM controllers with similar networking policy requirements For example, VM controllers can share VLAN pools and application endpoint groups (EPGs). The APIC communicates with the controller to publish network configurations such as port groups that are then applied to the virtual workloads. The VMM domain profile includes the following essential components:
  - Credential-Associates a valid VM controller user credential with an APIC VMM domain.
  - Controller-Specifes how to connect to a VM controller that is part of a policy enforcement domain. For example, the controller specifies the connection to a VMware vCenter that is part a VMM domain.

  * A single VMM domain can contain multiple instances of VM controllers, but they must be from the same vendor Not® (for example, from VMware or from Microsoft.

- **EPG** Association-Endpoint groups regulate connectivity and visibility among the endpoints within the scope of the VMM domain policy. VMM domain EPGs behave as follows:
  - The APIC pushes these EPSs as port groups into the VM controller.
  - An EPG can span multiple VMM domains, and a VMM domain can contain multiple EPGs.
- **Attachable Entity Profile** Association-Associates a VMM domain with the physical network infrastructure. An attachable entity profile (AEP) is a network interface template that enables deploying VM controller policies on a large set of leaf switch ports. An AEP specifies which switches and ports are available, and how they are configured.
- **VLAN Pool Association-A** VLAN pool specifies the VLAN IDs or ranges used for VLAN encapsulation that the VMM domain consumes.

Question: 64

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. L2 interface policy

B. link level policy

C. slow drain policy

D. ingress control plane policing policy

E. ingress data plane policing policy

Answer: B, E

Explanation:

To deploy a leaf access port policy group in ACI Fabric that controls the amount of application data flowing into the system and allows auto-negotiation of link speed, the two ACI policies that must be configured are:

Link level policy2.

Ingress data plane policing policy2.

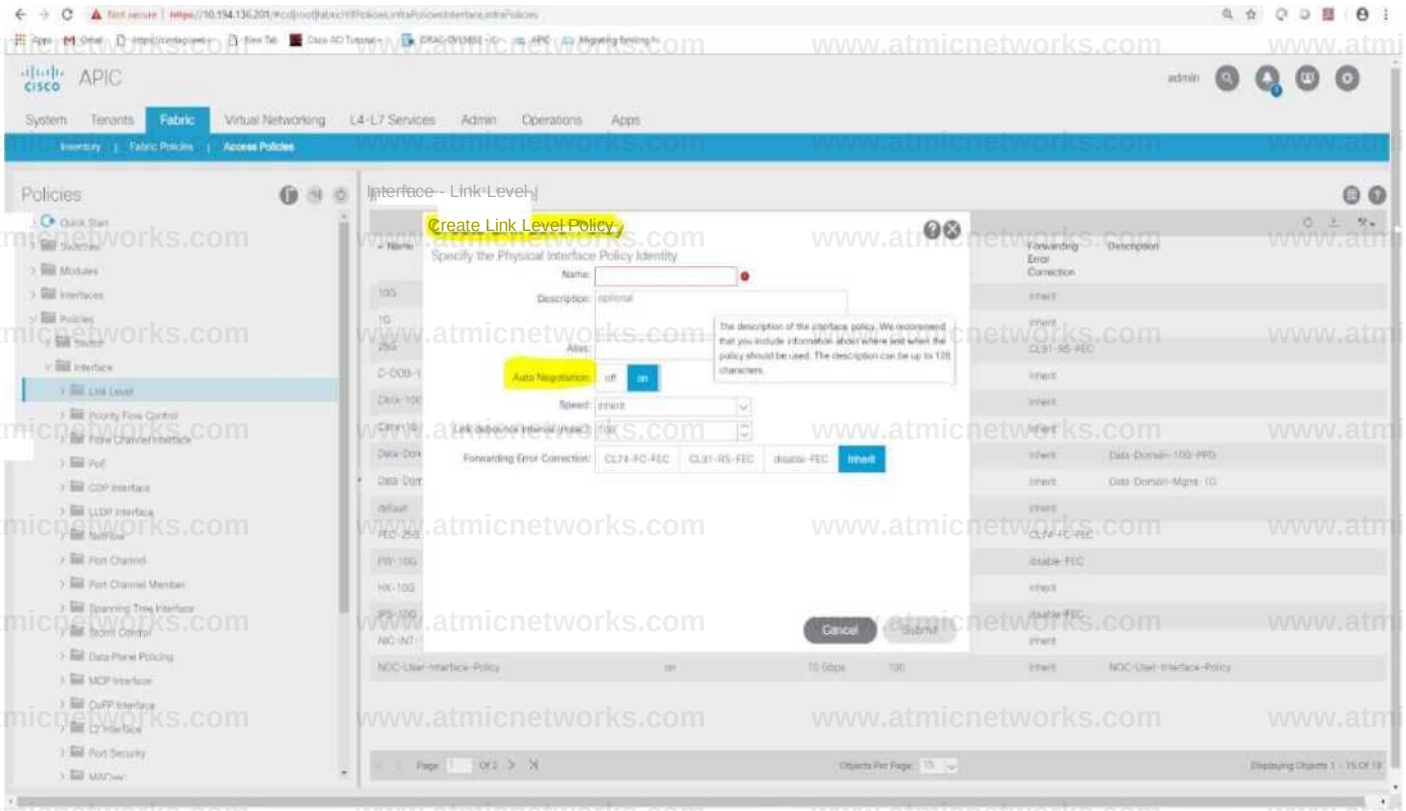

Slow Drain handles FCoE packets that are causing traffic congestion on ACI fabric. So, it is wrong.
Ingress control plane is wrong, because the request is for “application data flowing”.

L2 interface policy is concerned about QinQ and VLAN scope.

Question: 65

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

A. PIM routing

B. spanning trees

C. shortest-path trees

D. forwarding tag trees

Answer: D

Explanation:

The method that the Cisco ACI fabric uses to load-balance multidestination traffic is forwarding tag trees3. This method involves assigning a forwarding tag (FTAG) to the traffic when it is forwarded within the fabric, ensuring efficient load balancing of multi-destination traffic3.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/aci- fundamentals/b_ACI-Fundamentals/b_ACI-Fundamentals_chapter_010010.html

The ACI fabric uses Forwarding Tag (FTAG) trees to load balance multi-destination traffic. All multi-destination traffic is forwarded in the form of encapsulated IP multicast traffic within the fabric. The ingress leaf assigns an RAG to the traffic when forwarding it to the spine. The RAG is assigned in the packet as part of the destination multicast address. In the fabric, the traffic is forwarded along the specified RAG tree. Spine and any intermediate leaf switches forward traffic based on the RAG ID. One forwarding tree is built per RAG ID. Between any two nodes, only one link forwards per RAG. Because of the use of multiple RAGs, parallel links can be used with each RAG choosing a different link for forwarding. The larger the number of RAG trees in the fabric means the better the load balancing potential is. The ACI fabric supports up to 12 RAGs.

Question: 66

DRAG DROP

Refer to the exhibit.

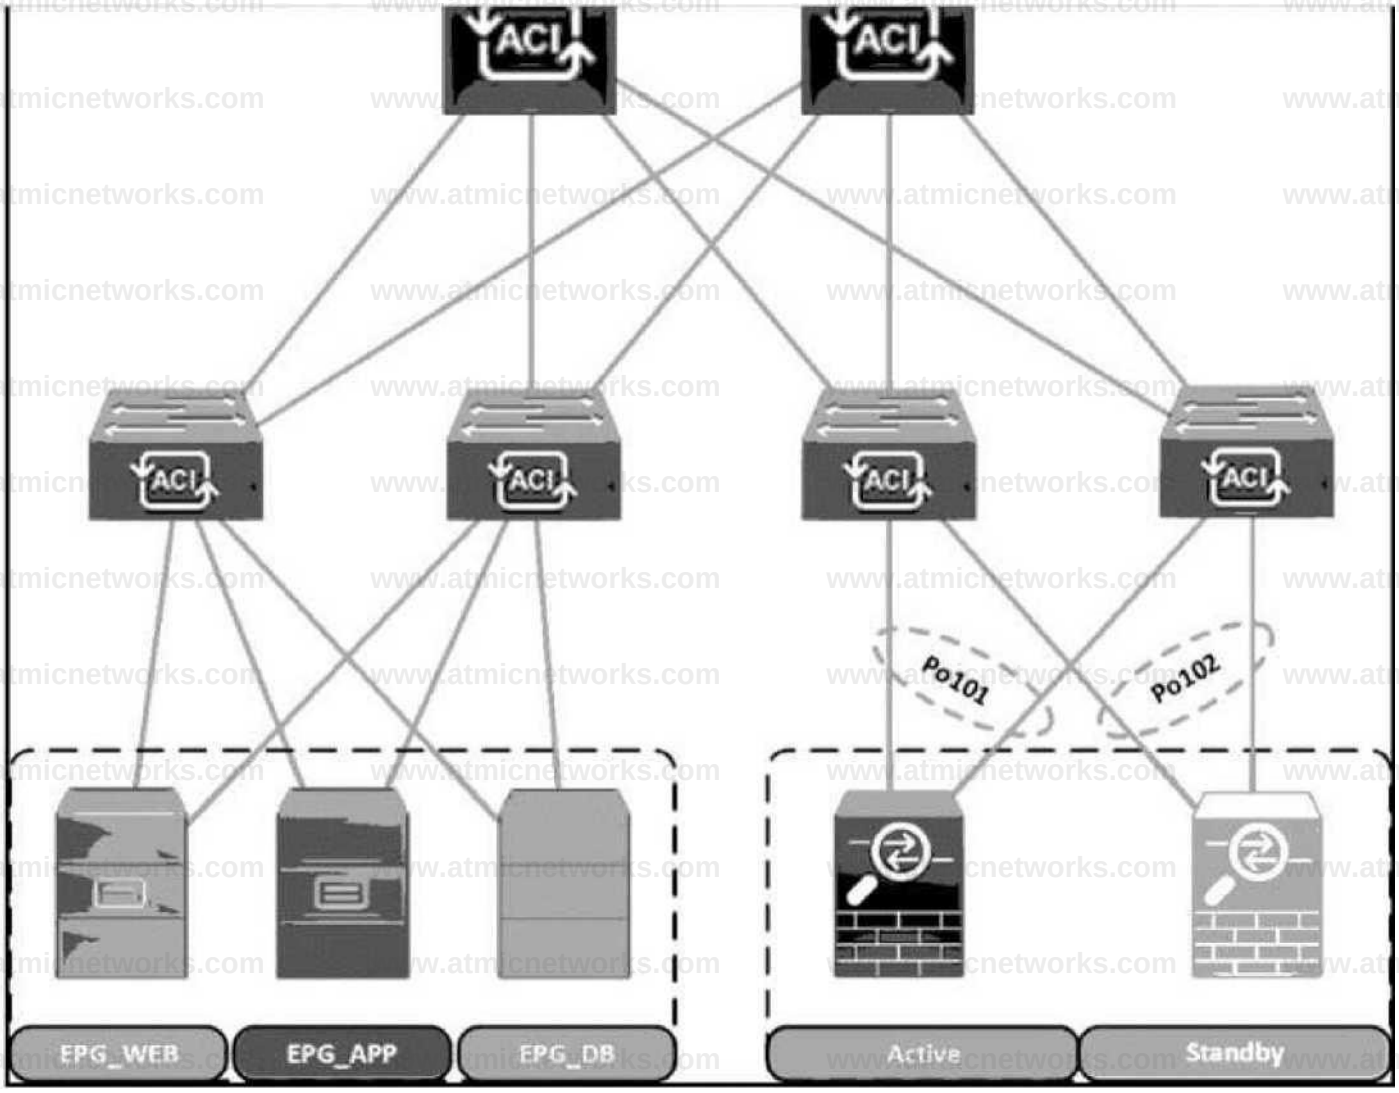

A Cisco ACI fabric is newly deployed, and the security team requires more visibility of all inter-EPG traffic flows. All traffic in a VRF must be forwarded to an existing firewall pair. During fallover, the standby firewall must continue to use the same IP and MAC as the primary firewall. Drag and drop 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.)

Apply o service graph template end 'elect vtAny 6XS as the consumes and provide'	step:
Select a redirect policy with the Layer J desBnatxsn	[1]1 Step?
Create a Layer 4 to Layr 7 service ytaph template.	Step 3
Select a tedrect policy with enabled aryuM and the Layer J destination.	Step 4

Select the same Ouster interface under Consume' Connector ar Provider CwxMxtw

Create a MS vice bridge domain end a Layw 4 to layer 7 device with one duster interface.

Select the exabng contract with custom P Ether Type fitter.

Answer:

Explanation:

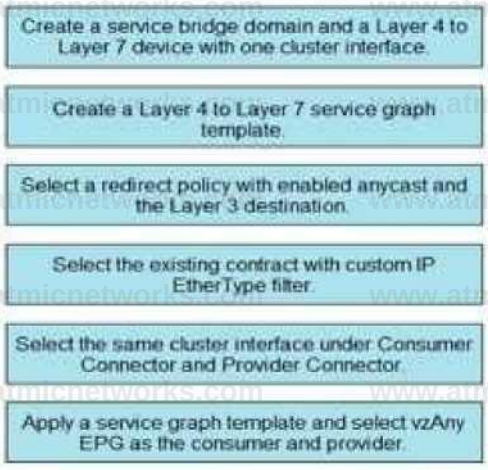

Question: 67

Refer to the exhibit.

```
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
```

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. Set the fallback login to local.

C. Assign the user to the default role.

D. Set the default login realm to LDAP

Answer: B

Explanation:

To ensure authentication if the RADIUS servers are unavailable, the action that should be taken is to set the fallback login to local. This setting allows users to authenticate using the local database on the Cisco APIC when external RADIUS servers cannot be reached.

Question: 68

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: A, D

Explanation:

To configure an L3Out peering with the backbone network that must forward both unicast and multicast traffic, the two

methods that should be used are:

Layer 3 routed port (Option A): This method involves configuring a physical port on the Cisco ACI leaf switch as a Layer 3 interface, which is suitable for routing unicast and multicast traffic.

Layer 3 routed subinterface (Option D): This method involves configuring a subinterface under a physical port, which allows for traffic segregation and routing over the same physical link, supporting both unicast and multicast traffic.

These methods ensure that the L3Out can handle the required traffic types and maintain proper routing with the backbone network.

https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application- centricinfrastructure/ guide-c07-743150.html#_L3Out_sStatic_rRoutes

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/L3- configuration/Cisco-APIC-Layer-3-Networking-Configuration-Guide-401/Cisco-APIC-Layer-3- Networking-Configuration-Guide-401_chapter_010010.html

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

Question: 69

Refer to the exhibit.

Domain - F1-VCSAB1_VCD

oC

Policy Operational Associated EPGs

General VSwitch Policy Faults History

0 ± ^

Properties

Port Channel Policy; Fl -VCSAB1_VCDJacplae

LLDP Policy F1-VCSAB1VCD ttopllPo 0

CDP Policy: select an option

MTU Policy: select an option

STP Policy: select an option

Firewall Policy. select «m option

Klei Dow Exporter Policy: select an option

Eri anced Lao Policy

HHHHBHHHHHHBI -

Name	Mode	' Load Balancing Mode	Number of Links
	LACP Active	Source and Destination IP Ado ess 2	

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: B, E

Explanation:

To meet the server team's requirements for VMM integration with ESXi servers, the two enhanced LAG policies that should be configured are:

LB Mode: Destination IP Address and TCP/UDP Port (Option B): This load balancing mode ensures that traffic is evenly distributed across all available links based on the destination IP address and the TCP/UDP port numbers, which can provide a balanced traffic flow for different sessions.

LACP Mode: LACP Active (Option E): Setting the LACP mode to active allows the network switches to initiate the LACP negotiation, which is necessary when the servers are not configured to initiate the LACP process.

These settings ensure that the LAG group consisting of two 10 Gigabit Ethernet links will have traffic evenly distributed across them and that the network switches will initiate the LACP negotiation as required by the server team.

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/aci_virtual_edge/configuration/2- x/Cisco-ACIVirtual-Edge-Configuration-Guide-202/Cisco-ACI-Virtual-Edge-Configuration-Guide-202_chapter_0100.html

Question: 70

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: A, E

Explanation:

To connect Cisco ACI fabric using Layer 2 with external third-party switches configured using 802.1s protocol, the two constructs required are:

Spanning tree policy for mapping MST Instances to VLANs (Option A): This policy will ensure that the Multiple Spanning Tree (MST) instances on the third-party switches are correctly mapped to the VLANs in the ACI fabric.

Static binding of native VLAN in all existing EPGs (Option E): This construct is necessary to maintain the native VLAN across all EPGs when integrating with third-party switches that use 802.1s protocol.

https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2019/pdf/BRKACI-3101.pdf

https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric- infrastructure/white-paper-c07-732033.html

Question: 71

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: A

Explanation:

When configuring a Multi-Pod system with the default gateway residing outside of the ACI fabric for a bridge domain, the setting that should be configured is to disable Limit IP Learning to Subnet (Option A). This allows the ACI fabric to learn IP addresses outside of the defined subnet, which is necessary when the default gateway is external.

Question: 72

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:

To support the addition of new ESXi hosts to the existing VMM domain, the action that should be taken is to map the leaf interface selector to the AEP that is associated with the VMM domain (Option D). This ensures that the correct policies are applied to the interfaces where the new ESXi hosts are connected.

Question: 73

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: A

Explanation:

To reduce instability during the migration of a legacy environment to Cisco ACI, the feature that should be enabled in the bridge domain is to Set Multi-Destination Flooding to Flood in BD (Option A). This setting allows unknown multicast, broadcast, and unknown unicast traffic to be flooded within the bridge domain, which can help prevent traffic loss during the migration process.

Question: 74

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: A, D

Explanation:

To filter the System Faults page and view only the active faults present in the Cisco ACI fabric, the two lifecycle stages that must be selected for filtering are:

Raised (Option A): This stage indicates that the fault is currently active and has been raised.

Raised, Clearing (Option D): This combination indicates that the fault is active but is in the process of being cleared.

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/faults/guide/b_APIC_Faults_Errors/b_IFC_Faults_Errors_chapter_01.html

Question: 75

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: C

Explanation:

When the Cisco ACI fabric has a stale endpoint entry for the destination endpoint, the traffic flow is affected in that the leaf switch floods the traffic to the endpoint throughout the fabric (Option C).

This flooding occurs because the leaf switch does not have the correct location of the endpoint in its endpoint table, leading it to flood the traffic in an attempt to reach the destination.

Reference:

https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2019/pdf/BRKACI-2641.pdf

Question: 76

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: D

Explanation:

To export all configuration parameters necessary for disaster recovery in Cisco ACI, the setting that must be enabled is enabled JSON format export1. This format allows for a complete and structured export of the configuration that can be used for recovery purposes1.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b_KB_Using_Import_Export_to_Recover_Config_States.html

Question: 77

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 Cisco APIC configuration that prevents a remote network that is not configured on the bridge domain from being learned by the fabric is to enable Limit IP Learning to Subnet2. This setting restricts IP address learning to only those subnets that are defined on the bridge domain2.

Reference:

https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric- infrastructure/white-paper-c11-739989.html

Question: 78

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 L3Out subnet configuration option that creates an inbound route map for route filtering is Import Route Control Subnet3. This option allows for the application of route maps to incoming routes from external networks, providing granular control over which routes are allowed into the ACI fabric3.

Question: 79

An engineer must set up a Cisco ACI fabric to send Syslog messages related to hardware events, such as chassis line card failures. The messages should be sent to a dedicated Syslog server. Where in the Cisco APIC should the policy be configured to meet this requirement?

A. uni/tn-common/monepg-default

B. uni/infra/monifra-default

C. uni/fabric/monfab-default

D. uni/fabric/moncommon

Answer: C

Explanation:

To set up a Cisco ACI fabric to send Syslog messages related to hardware events to a dedicated Syslog server, the policy should be configured at uni/fabric/monfab-default4. This location in the Cisco APIC allows for the configuration of Syslog policies that can be applied to the entire fabric4.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/aci-fundamentals/Cisco-ACI-Fundamentals-401/Cisco-ACI-Fundamentals-401_chapter_01100.html

Configuring Monitoring Policies

Administrators can create monitoring policies with the following four broad scopes:

- Fabric Wide: includes both fabric and access objects
- Access (also known as infrastructure): access ports, FEX, VM controllers, and so on
- Fabric: fabric ports, cards, chassis, fans, and so on
- Tenant: EPGs, application profiles, services, and so on

The APIC includes the following four classes of default monitoring policies:

- monCommonPol (uni/fabric/moncommon): applies to both fabric and access infrastructure hierarchies
- monFabricPol (uni/fabric/monfab-default): applies to fabric hierarchies
- moninfraPol (uni/infra/monifra-default): applies to the access infrastructure hierarchy
- monEPGPol (uni/tn-common/monepg-default): applies to tenant hierarchies

Question: 80

The existing network and ACI fabric have been connected to support workload migration. Servers will physically terminate at the Cisco ACI, but their gateway must stay in the existing network. The solution needs to adhere to

Cisco's best practices. The engineer started configuring the relevant Bridge Domain and needs to complete the configuration.

Which group of settings are required to meet these requirements?

A. L2 Unknown Unicast: Hardware Proxy

L3 Unknown Multicast Flooding: Flood

Multi Destination Flooding: Flood in BD

ARP Flooding: Enable

B. L2 Unknown Unicast: Flood

L3 Unknown Multicast Flooding: Flood

Multi Destination Flooding: Flood in BD

ARP Flooding: Enable

C. L2 Unknown Unicast: Flood

L3 Unknown Multicast Flooding: Optimize Flood

Multi Destination Flooding: Flood in BD

ARP Flooding: Disable

D. L2 Unknown Unicast: Hardware Proxy

E. Unknown Multicast Flooding: Optimize Flood

Multi Destination Flooding: Flood in BD

ARP Flooding: Disable

Answer: B

Explanation:

The group of settings required to meet the requirements for workload migration where servers will physically terminate at the Cisco ACI, but their gateway must stay in the existing network, is:

F. Unknown Unicast: Flood

G. Unknown Multicast Flooding: Flood

Multi Destination Flooding: Flood in BD

ARP Flooding: Enable5

These settings ensure that the bridge domain is configured to handle unknown unicast and multicast traffic appropriately, and ARP flooding is enabled to allow for the resolution of IP addresses when the gateway is outside the fabric5.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/migration_guides/migrating_existing_networks_to_aci.html

Default Gateway Migration Considerations

The default gateway used by the workloads to establish communication outside their IP subnet is initially maintained in the Brownfield network; this implies that the ACI fabric initially provides only Layer 2 services for devices part of EPG1, and the workloads already migrated to the ACI fabric send traffic to the Brownfield network when they need to communicate with devices external to their IP subnet (shown in the following diagram).

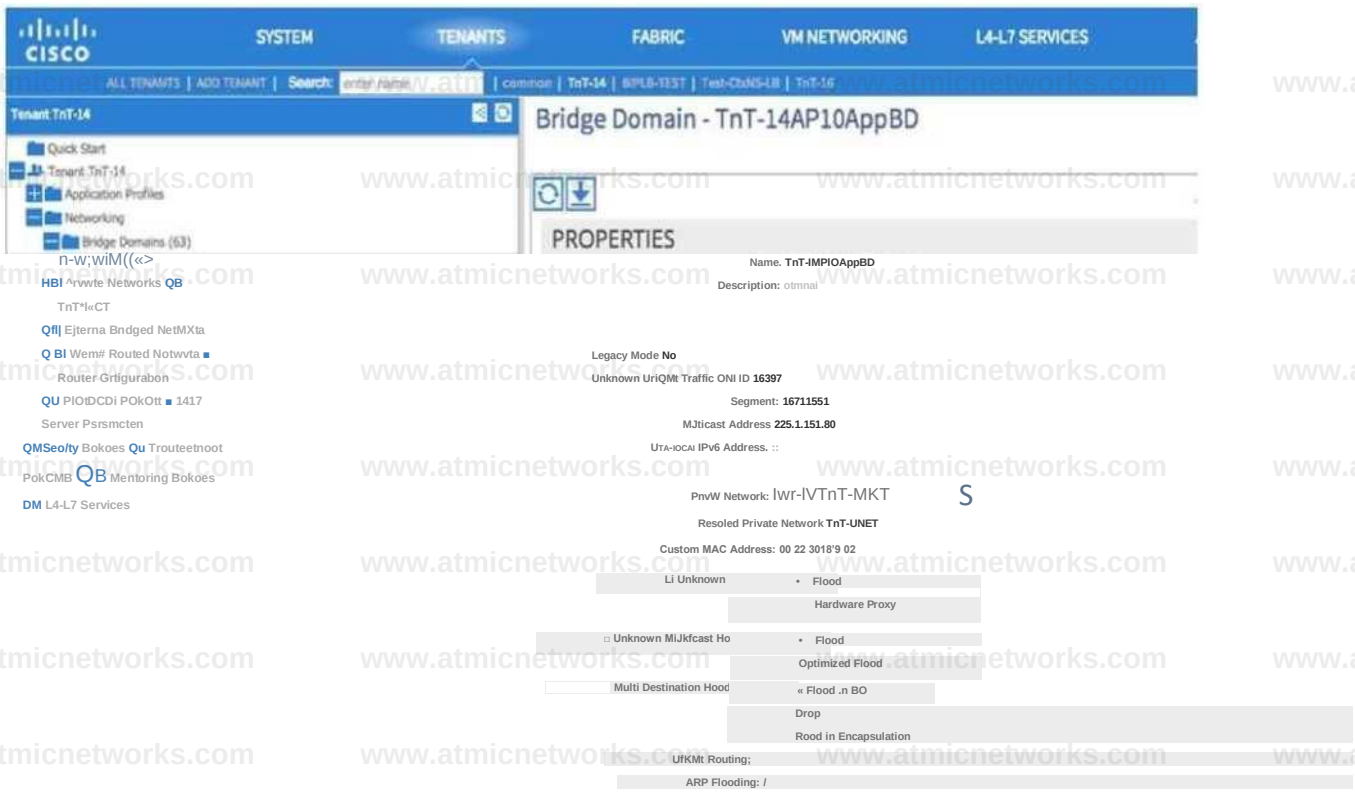

Question: 81

An engineer must implement management policy and data plane separation in the Cisco ACI fabric. Which ACI object must be created in Cisco APIC to accomplish this goal?

A. Application profile

B. Tenant

C. Contract

D. Bridge domain

Answer: D

Explanation:

To implement management policy and data plane separation in the Cisco ACI fabric, the ACI object that must be created in Cisco APIC is a Bridge domain

Question: 82

An engineer is implementing a Cisco ACI environment that consists of more than 20 servers. Two of the servers support only Cisco Discovery Protocol with no order link discovery protocol. The engineer wants the servers to be discovered automatically by the Cisco ACI fabric when connected. Which action must be taken to meet this requirement?

A. Create an override policy that enables Cisco Discovery Protocol after LLDP is enabled in the default policy group.

B. Configure a higher order interface policy that enables Cisco Discovery Protocol for the interface on the desired leaf switch.

C. Configure a lower order policy group that enables Cisco Discovery Protocol for the interface on the desired leaf switch.

D. Create an interface profile for the interface that disables LLDP on the desired switch that is referenced by the interface policy group.

Answer: A

Explanation:

To ensure that servers supporting only Cisco Discovery Protocol are discovered automatically by the Cisco ACI fabric when connected, the action that must be taken is to Create an override policy that enables Cisco Discovery Protocol after LLDP is enabled in the default policy group

Question: 83

An engineer wants to monitor all configuration changes, threshold crossing, and link-state transitions in a Cisco ACI fabric. Which action must be taken to receive the required messages?

A. Add Faults and Events to the monitor policy.

B. Add Session Logs and Audit Logs to the monitor policy.

C. Include Audit Logs and Events in the Syslog source policy.

D. Include Events and Session Logs in the Syslog source policy.

Answer: C

Explanation:

To monitor all configuration changes, threshold crossing, and link-state transitions in a Cisco ACI fabric, the action that must be taken is to Include Audit Logs and Events in the Syslog source policy

https://community.cisco.com/legacyfs/online/attachments/blog/technote-aci-syslog_external- latest.pdf

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/faults/guide/b_APIC_Faults_Errors/b_IFC_Faults_Errors_chapter_010.html

Question: 84

An organization has encountered many STP-related issues in the past due to failed hardware components. They are in the process of long-term migration to a newly deployed ACI fabric. Senior engineers are worried that spanning-tree loops in the existing network may be extended to the ACI fabric. Which feature must be enabled on the ACI leaf ports to protect the fabric from spanning-tree loops?

A. BPDU Guard

B. per-VLANMCP

C. Storm Control

D. BPDU Filter

Answer: A

Explanation:

o protect the ACI fabric from spanning-tree loops, the feature that must be enabled on the ACI leaf ports is BPDU Guard

Reference: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/aci-fundamentals/Cisco-ACI-Fundamentals-401/Cisco-ACI-Fundamentals-401_chapter_0101.html

Question: 85

A network engineer must design a method to allow the Cisco ACI to redirect traffic to the firewalls. Only traffic that matches specific L4-L7 policy rules should be redirected. The load must be distributed across multiple firewalls to scale the performance horizontally. Which action must be taken to meet these requirements?

A. Configure ACI Service Graph with Unidirectional PBR.

B. Implement ACI Service Graph with GIPo.

C. Implement ACI Service Graph Two Nodes with GIPo.

D. Configure ACI Service Graph with Symmetric PBR.

Answer: A

Explanation:

To design a method that allows the Cisco ACI to redirect traffic to the firewalls based on specific L4- L7 policy rules and distribute the load across multiple firewalls, the action that must be taken is to Configure ACI Service Graph with Unidirectional PBR (

https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739971.html

Question: 86

An engineer created two interface protocol policies called Pol_CDP40275332 and Pol_LLDP46783451. The policies must be used together in a single policy. Which ACI object must be used?

A. interface policy group

B. switch policy group

C. switch profile

D. interface profile

Answer: A

Explanation:

When two interface protocol policies need to be used together in a single policy, the ACI object that must be used is an interface policy group

Question: 87

What is the minimum number of APICs does Cisco recommend to deploy in a production cluster?

A. 1

B. 3

C. 4

D. 5

Answer: B

Explanation:

The minimum number of APICs that Cisco recommends to deploy in a production cluster is 3 (Option B). This recommendation is based on the need for high availability and redundancy in a production

environment.

Question: 88

Refer to the exhibit.

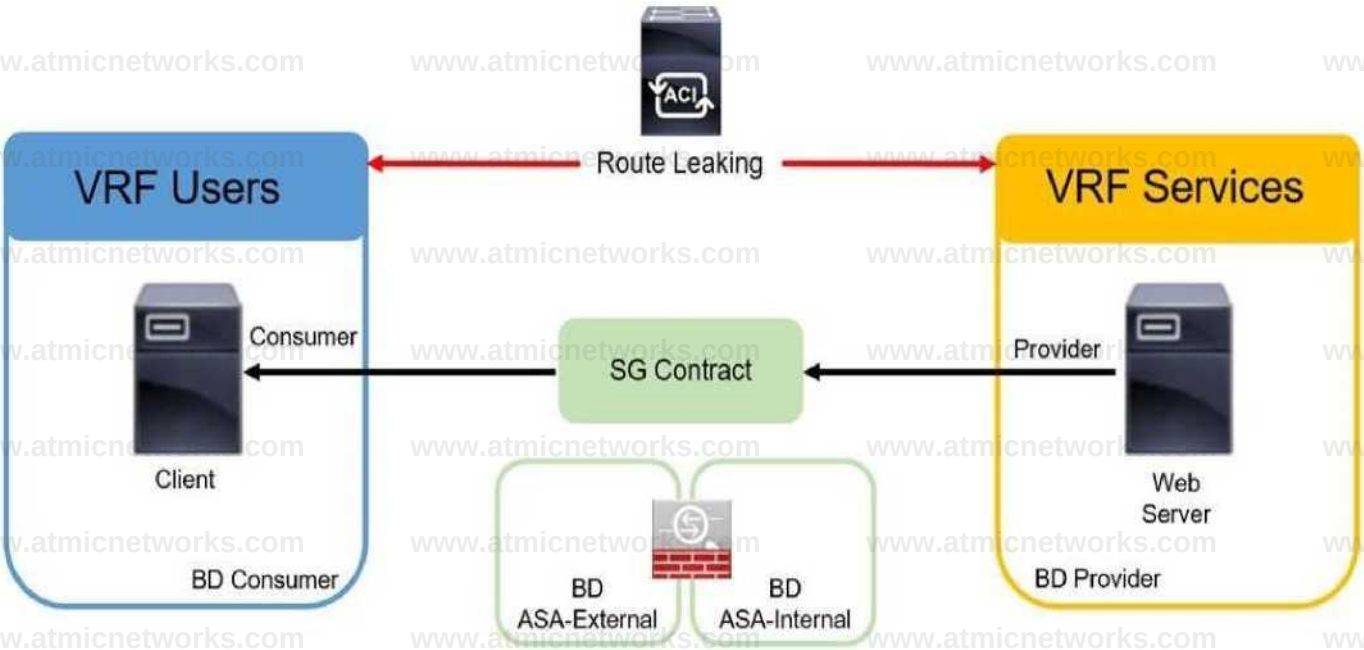

An engineer must implement the inter-tenant service graph. Which set of actions must be taken to accomplish this goal?

A. • Define the contract in the provider tenant and export it to the consumer tenant.
• Define the L4-L7 device, service graph template, and ASA bridge domains in the provider tenant.

B. • Define the contract in the provider tenant and export it to the consumer tenant.
• Define the L4-L7 device and service graph template in the provider tenant and the ASA bridge domains in the consumer tenant.

C. • Define the contract in the provider tenant and export it to the provider tenant.

• Define the L4-L7 device and service graph template in the provider tenant and the ASA bridge domains in the consumer tenant.

D. • Define the contract in the provider tenant and export it to the provider tenant.

• Define the L4-L7 device, service graph template, and ASA bridge domains in the consumer tenant.

Answer: B

Explanation:

To implement the inter-tenant service graph, the set of actions that must be taken are:

Define the contract in the provider tenant and export it to the consumer tenant (Option B).

Define the L4-L7 device and service graph template in the provider tenant and the ASA bridge domains in the consumer tenant (Option B).

This approach allows for the proper configuration and association of the service graph between the provider and consumer tenants, ensuring that the services are correctly applied to the traffic flowing between them.

Question: 89

Refer to the exhibit.

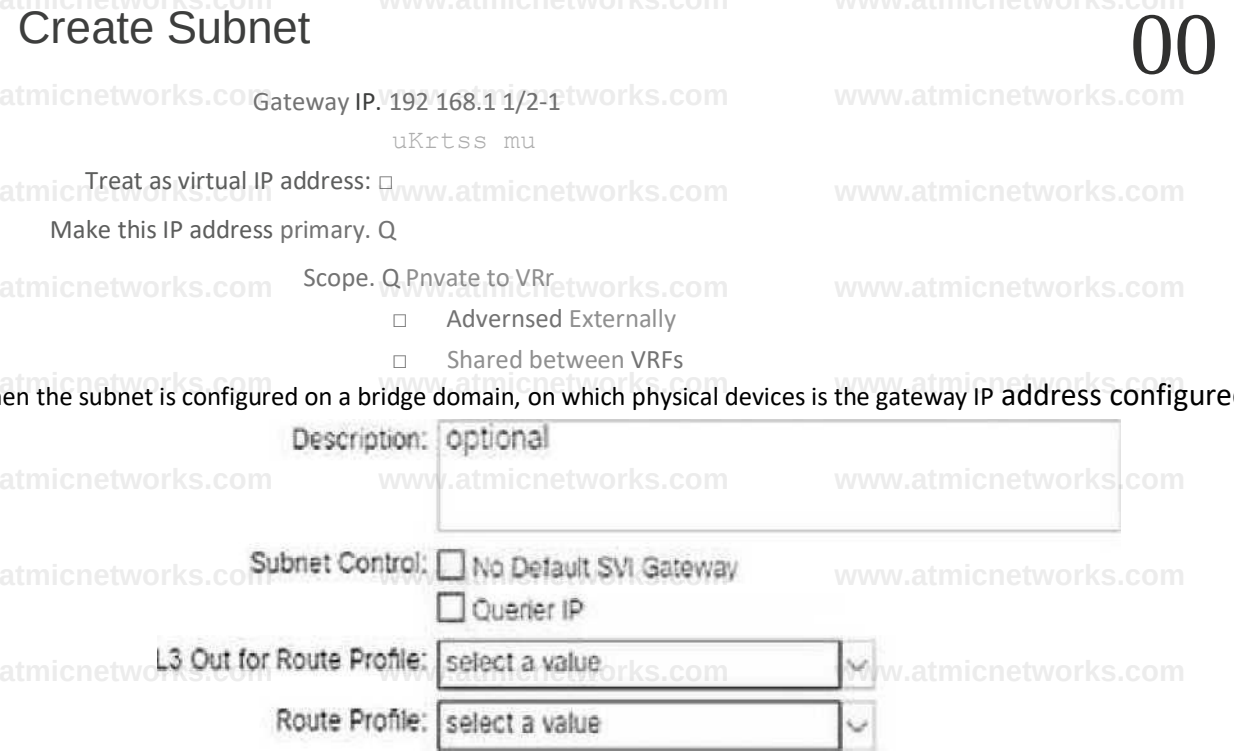

When the subnet is configured on a bridge domain, on which physical devices is the gateway IP address configured?

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:

In Cisco Application Centric Infrastructure (ACI), when a subnet is configured on a bridge domain, the gateway IP address is configured on the leaf switches where the bridge domain of the tenant is present. This configuration allows for localized routing within the fabric, providing efficient east-west traffic handling without requiring traffic to traverse up to the spine nodes unless necessary for northsouth routing or policy enforcement.

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

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_0111.html

From a practical perspective, each bridge domain will exist in a particular leaf if there is a connected endpoint that belongs to that endpoint group. Each bridge domain receives a VLAN ID in the leaf switches.

Question: 90

Refer to the exhibit.

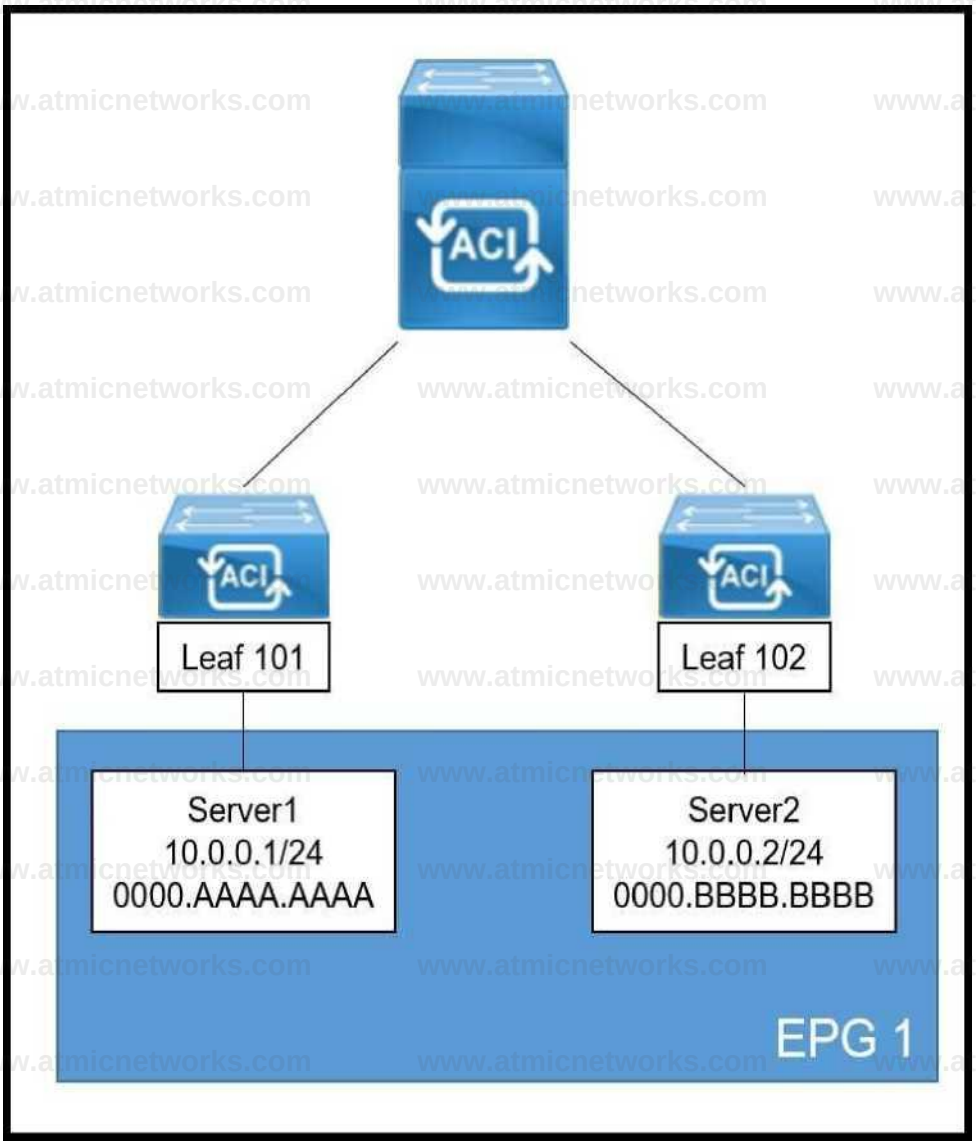

A systems engineer is implementing the Cisco ACI fabric. However, the Server2 information is missing from the Leaf 101 endpoint table and the COOP database of the spine. The requirement is for the bridge domain configuration to enforce the ACI fabric to forward the unicast packets generated by Server1 destined to Server2. Which action must be taken to meet these requirements?

A. Enable ARP Flooding

B. Set L2 Unknown Unicast to Flood

C. Set IP Data-Plane Learning to No

D. Enable Unicast Routing

Answer: B

Explanation:

In the scenario where Server2 information is missing from the Leaf 101 endpoint table and the COOP database of the spine, setting Layer 2 Unknown Unicast to Flood (Option B) is the correct action to meet the requirements. This setting will cause the ACI fabric to flood unicast packets destined for unknown destinations within the bridge domain. This ensures that packets from Server1 will reach Server2 even though its location isn't known by Leaf 101 or the spine's COOP database.https://www.cisco.com/c/en/us/solutions/collateral/data-center- virtualization/application-centric-infrastructure/white-paper-c11-739989.html

Silent hosts considerations

In the case of silent hosts, where an ACI leaf hasn't learned a local endpoint, ACI has some mechanisms to detect those silent hosts. Some of them are controlled by BD configurations. Following are explanations of each scenario with related BD configurations.

For (L2) switched traffic to an unknown MAC, the L2 Unknown Unicast option under the BD may need to be set to "Flood". This is because the ACI fabric with the L2 Unknown Unicast "Hardware-Proxy" configuration drops the L2 unicast packets on the spine in cases where the destination MAC has not been learned as an endpoint anywhere on the BD in ACI, and the COOP database doesn't have the information.

Question: 91

An engineer must allow multiple external networks to communicate with internal ACI subnets.

Which action should the engineer take to assign the prefix to the class ID of the external Endpoint Group?

A. Enable the Export Route Control Subnet for the External Endpoint Group flag.

B. Enable an L30ut with Shared Route Control Subnet.

C. Configure subnets with the External Subnets for External EPG flag enabled.

D. Configure subnets with the Import Route Control Subnet flag enabled.

Answer: C

Explanation:

To allow multiple external networks to communicate with internal ACI subnets and assign the prefix to the class ID of the

external Endpoint Group, the engineer should configure subnets with the External Subnets for External EPG flag enabled1. This configuration allows the specified subnets to be associated with an external EPG, which facilitates communication between internal tenants and external routed networks via L3Outs1.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/ACI_Best_Practices/b_ACI_Best_Practices/b_ACI_Best_Practices_chapter_01001.html

The external subnets for an external EPG are used to define the subnets that should be classified to the external EPG. This policy does not affect routing. It is similar to an Access Control List (ACL) that assigns a prefix to the class id (pcTag) of the external EPG.

Question: 92

An engineer must ensure that Cisco ACI flushes the appropriate endpoints when a topology change notification message is received in an MST domain. Which three steps are required to accomplish this goal? (Choose three.)

A. Enable the BPDU interface controls under the spanning tree interface policy.

B. Configure a new STP interface policy.

C. Bind the spanning tree policy to the switch policy group.

D. Associate the STP interface policy to the appropriate interface policy group.

E. Create a new region policy under the spanning tree policy.

F. Map VLAN range to MAT instance number.

Answer: A, C, D

Explanation:

To ensure that Cisco ACI flushes the appropriate endpoints when a topology change notification message is received in an MST domain, the three steps required are:

Enable the BPDU interface controls under the spanning tree interface policy (Option A)2.

Bind the spanning tree policy to the switch policy group (Option C)3.

Associate the STP interface policy to the appropriate interface policy group (Option D)3. These steps ensure that the ACI

fabric responds correctly to STP events, such as topology changes, by flushing the MAC address table and the Local Station Table for the affected VLAN on the leaf switch4.

Question: 93

A Cisco ACI bridge domain and VRF are configured with a default data-plane learning configuration. Which two endpoint attributes are programmed in the leaf switch when receiving traffic? (Choose two.)

A. Remote MAC. IP

B. Remote Subnet

C. Local IP, not MAC

D. Local MAC, IP

E. Local Subnet

F. Remote IP

Answer: D, E

Explanation:

In a Cisco ACI bridge domain and VRF configured with a default data-plane learning configuration, the two endpoint attributes that are programmed in the leaf switch when receiving traffic are:

Local MAC, IP (Option D)5.

Local Subnet (Option E)5. These attributes are learned by the ACI fabric through common network methods such as ARP, GARP, and ND, as well as through the data-plane. The local MAC and IP addresses are learned and used for traffic routing and bridging, while the local subnet information is used for traffic optimization and endpoint location tracking56.

Question: 94

Refer to the exhibit.

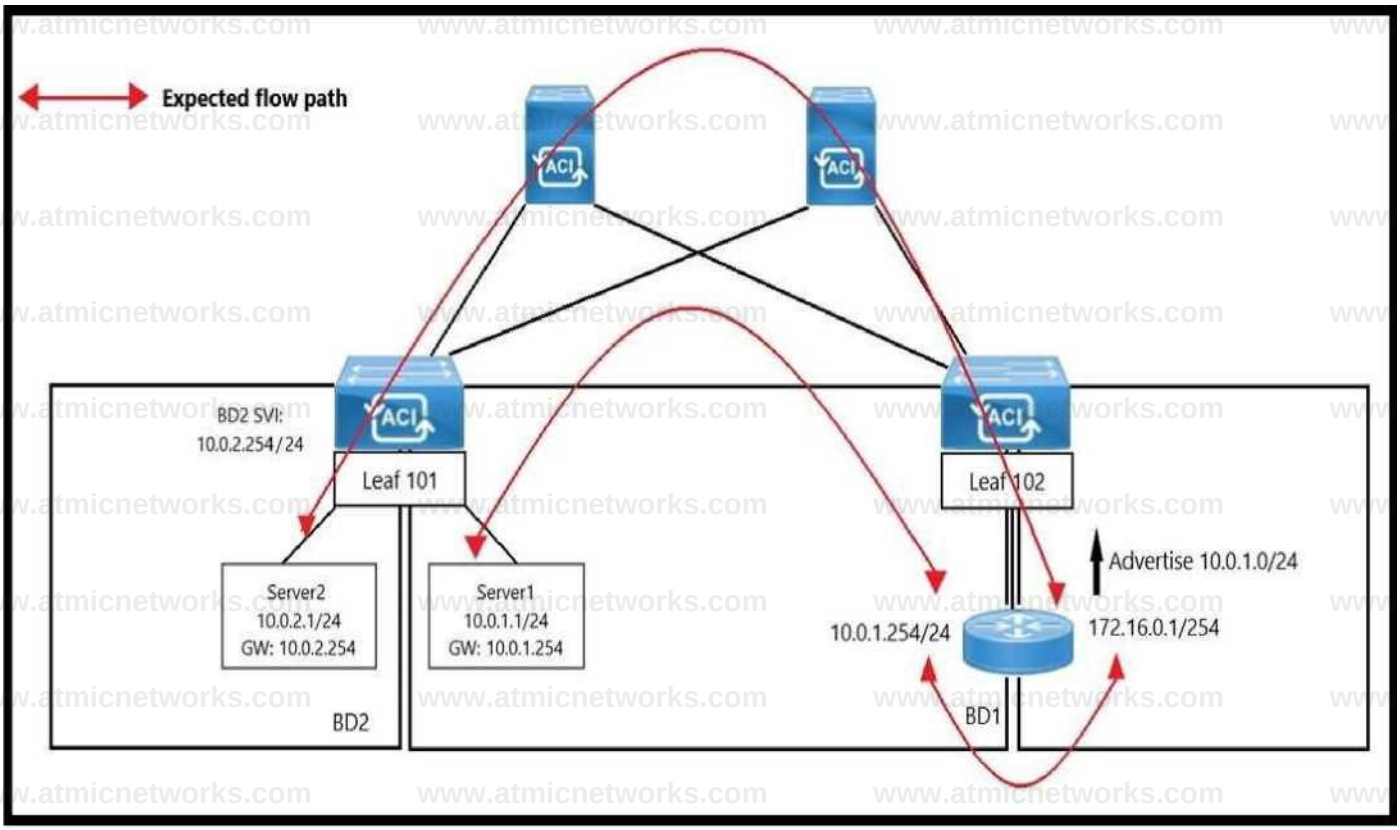

An engineer wants to initiate an ICMP ping from Server1 to Server2. The requirement is for the BD1 to enforce ICMP replies that follow the expected path. The packets must be prevented from taking the direct path from Leaf1 to Server1.

Which action must be taken on BD1 to meet these requirements?

A. Set L2 Unknown Unicast to Flood.

B. Set L2 Unknown Unicast to Hardware Proxy.

C. Disable Unicast Routing.

D. Enable ARP Flooding.

Answer: B

Explanation:

In the Cisco ACI environment, when an engineer wants to enforce a specific path for ICMP replies and prevent packets from taking a direct path, setting Layer 2 Unknown Unicast to Hardware Proxy on the bridge domain (BD1) is the appropriate action. This setting ensures that the unknown unicast traffic, which includes ICMP packets destined for an unknown MAC address, is forwarded to the spine proxy. The spine proxy then forwards the packets based on the endpoint table information, ensuring that the packets follow the expected path through the fabric and do not take any shortcuts or direct paths that bypass the intended routing.

Question: 95

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: C

Explanation:

When traffic is received from a Layer 3 Out (L3Out) on the ingress leaf switch in a Cisco ACI fabric, the source IP address of the traffic is learned as a remote endpoint. This is because the traffic is entering the ACI fabric from an external network, and the source IP is considered remote to the fabric.

Question: 96

An engineer must configure a group of servers with a contract that uses TCP port 80. The EGP that contains the web servers requires an external Layer 3 cloud to initiate communication. Which action must be taken to meet these requirements?

A. Configure the EGP as a provider and L3 out as consumer of the contract.

B. Configure OSPF to exchange routes between the L3 out and EGP.

C. Configure a taboo contract and apply it to the EPG.

D. Configure the EPG as a consumer and L3 out as a provider of the contract.

Answer: D

Explanation:

To configure a group of servers with a contract that uses TCP port 80 and requires an external Layer 3 cloud to initiate communication, the EPG that contains the web servers should be configured as a consumer of the contract, and the Layer 3 Out (L3Out) should be configured as the provider of the contract. This configuration establishes the direction of the contract and allows the external Layer 3 cloud to initiate communication with the web servers within the EPG.

Question: 97

The unicast routing feature is enabled on the bridge domain. Which two conditions enable the Cisco ACI leaf to learn a source IP as a local endpoint? (Choose two.)

A. Through Ethernet traffic received in a bridge domain.

B. IP traffic routed through an SVI.

C. Through VXLAN traffic received on the uplink.

D. IP traffic routed through a Layer 3 Out.

E. Through ARP received on an SVI.

Answer: A, E

Explanation:

With unicast routing enabled on the bridge domain, the two conditions that enable the Cisco ACI leaf to learn a source

IP as a local endpoint are:

Through Ethernet traffic received in a bridge domain (Option A): The leaf switch learns the source IP address as a local endpoint when it receives Ethernet traffic within the bridge domain.

Through ARP received on an SVI (Option E): The leaf switch also learns the source IP address as a local endpoint when it receives an ARP request or reply on a Switched Virtual Interface (SVI) associated with the bridge domain.

These conditions allow the ACI fabric to maintain an accurate endpoint database for efficient routing and forwarding of traffic within the fabric.

https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739989.html

A Cisco ACI leaf switch follows these steps to learn a local endpoint MAC address and IP address:

1. The Cisco ACIleaf receives a packet with a source MAC Address (MAC A) and source IP Address (IP A).
2. The Cisco ACIleaf learns MAC A as a local endpoint.
3. The Cisco ACIleaf learns IP A tied to MAC A if the packet is an ARP packet.
4. The Cisco ACIleaf learns IP A tied to MAC A if the packet is routed.

In Figure 5, the packet is Layer 3 traffic with the Cisco ACI bridge domain Switch Virtual Interface (SVI) as its default gateway. Therefore, both the MAC address and IP address (Src MAC S and Src IP 192.168.1.1 in the figure) are learned as a single local endpoint on LEAF1, and only IP address 192.168.1,1 is learned as a remote endpoint on LEAF2.

Question: 98

When does the Cisco ACI leaf learn a source IP or MAC as a remote endpoint?

A. When VXLAN traffic arrives on a leaf fabric port from the spine and outer source IP is in the Layer 3 Out EPG subnet range.

B. When VXLAN traffic arrives on a leaf fabric port from the spine and outer source IP is in the bridge domain subnets range.

C. When VXLAN traffic arrives on a leaf fabric port from the spine and inner source IP is in the Layer 3 Out EPG subnet range.

D. When VXLAN traffic arrives on a leaf fabric port from the spine and inner source IP is in the bridge domain subnets range.

Answer: D

Explanation:

In Cisco ACI, a leaf switch learns a source IP or MAC as a remote endpoint when VXLAN traffic arrives on a leaf fabric port from the spine, and the inner source IP is within the range of the bridge domain subnets associated with the leaf. This process is part of the ACI's COOP (Council Of Oracles Protocol),

which ensures that endpoint information is accurately disseminated across the fabric for efficient traffic forwarding.

Question: 99

DRAG DROP

An engineer must configure VMM domain integration on a Cisco UCS B-Series server that is connected to a Cisco ACI fabric. Drag and drop the products used to create VMM domain from the bottom into the sequence in which they should be

On the		interface, create a dynamic VLAN pool.
On the		interface, create a VMware vCenter domain.
On the		interface, create a vCenter/vShield controller.
On the		user interface, verify that the VMware vDS is created.

implemented at the top. Products are used more than once.

Answer:

Explanation:

APIC vCenter UCS Manager

On the	UCS Manager	Interface, create a dynamic VLAN pool.
On the	vCenter APIC	interface, create a VMware vCenter domain.
On the	vCenter	interface, create a vCenter/vShield controller.
On the		user interface, verify that the VMware vDS is created.

Question: 100

The company ESXi infrastructure is hosted on the Cisco UCS-B Blade Servers. The company decided to take advantage of ACI VMM integration to enable consistent enforcement of policies across virtual and physical workloads. The requirement is to prevent the packet loss between the distributed virtual switch and the ACI fabric. Which setting must be implemented on a vSwitch policy to accomplish this goal?

A. Static Channel

B. MAC Pinning

C. LACP

D. LLDP

Answer: B

Explanation:

To prevent packet loss between the distributed virtual switch and the ACI fabric in a company's ESXi infrastructure hosted on Cisco UCS-B Blade Servers, the vSwitch policy must implement MAC Pinning. This setting ensures that each virtual interface card (vNIC) on the UCS servers is pinned to an uplink Ethernet port, providing a stable and consistent path for traffic and preventing packet loss due to path changes.

Question: 101

An engineer is configuring ACI VMM domain integration with Cisco UCS-B Series. Which type of port channel policy must be configured in the vSwitch policy?

A. LACP Active

B. MAC Pinning

C. LACP Passive

D. MAC Pinning-Physical-NIC-load

Answer: D

Explanation:

When configuring ACI VMM domain integration with Cisco UCS-B Series, the type of port channel policy that must be configured in the vSwitch policy is MAC Pinning-Physical-NIC-load. This policy type allows for the distribution of traffic across physical NICs based on MAC address hashing, which is suitable for environments where dynamic port channels (LACP) are not used.

Reference: https://www.cisco.com/c/en/us/support/docs/cloud-systems-management/application- policy-infrastructure-controller-apic/118965-config-vmm-aci-ucs-00.html

Question: 102

In-band is currently configured and used to manage the Cisco ACI fabric. The requirement is for leaf and spine switches to use out-of-band management for NTP protocol. Which action accomplishes this goal?

A. Select Out-of-Band as Management EPG in the default DateTimePolicy.

B. Create an Override Policy with NTP Out-of-Band for leaf and spine switches.

C. Change the interface used for APIC external connectivity to ooband.

D. Add a new filter to the utilized Out-of-Band-Contract to allow NTP protocol.

Answer: B

Explanation:

To configure leaf and spine switches in a Cisco ACI fabric to use out-of-band management for NTP protocol, an Override Policy with NTP Out-of-Band must be created. This policy specifies that NTP traffic should be directed through the out-of-

band management network, separate from the in-band management traffic, ensuring dedicated management connectivity for time synchronization.

https://www.cisco.com/c/en/us/support/docs/cloud-systems-management/application-policy- infrastructure-controller-apic/200128-Configuring-NTP-in-ACI-Fabric-Solution.html

Question: 103

DRAG DROP

An engineer must configure RADIUS authentication with Cisco ACI for remote authentication with out-of-band management access. Drag and drop the RADIUS configuration steps from the left into the required implementation order on the right. Not all steps are used.

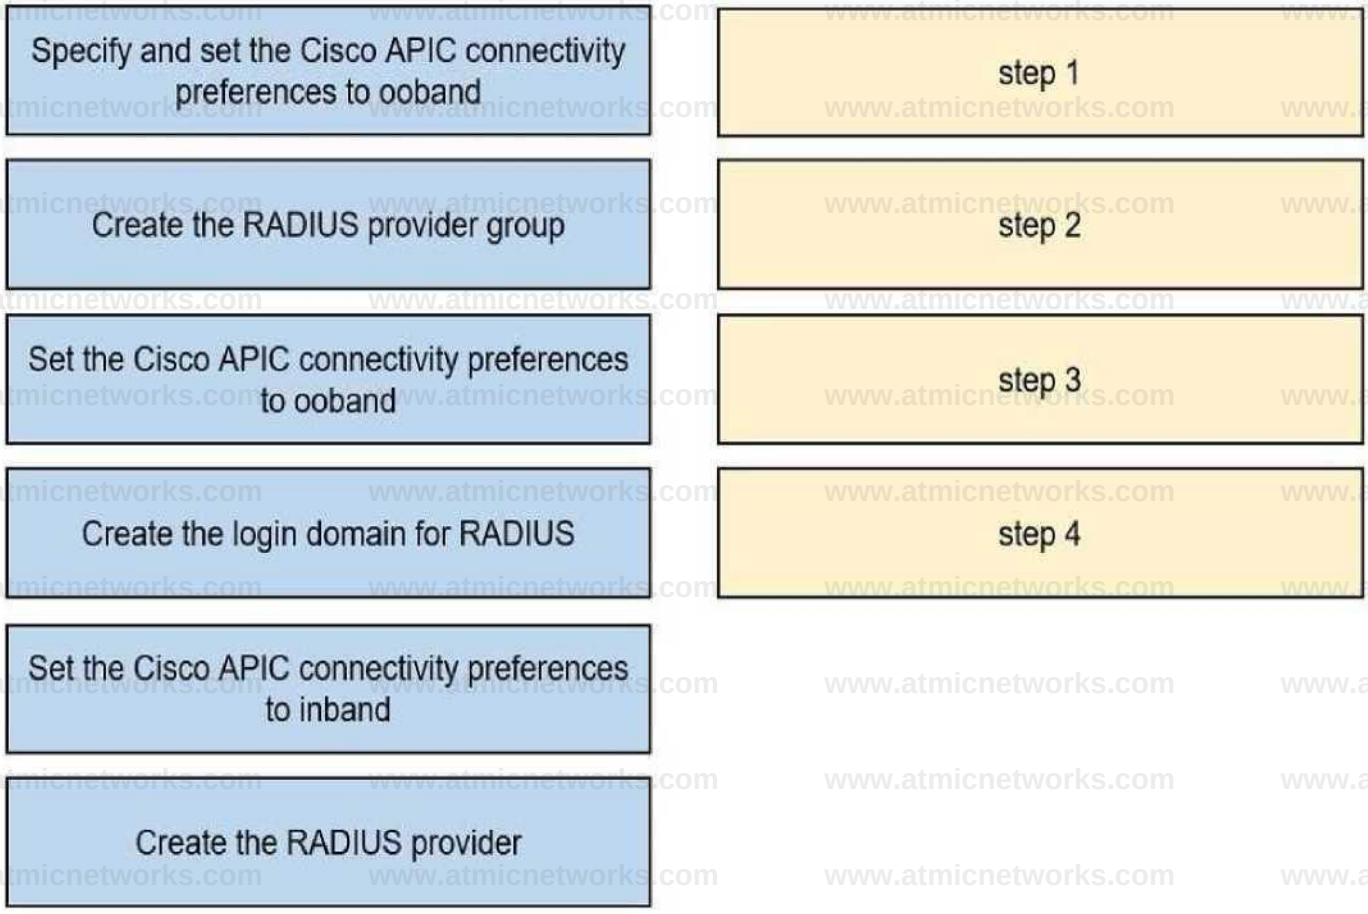

Answer:

Explanation:

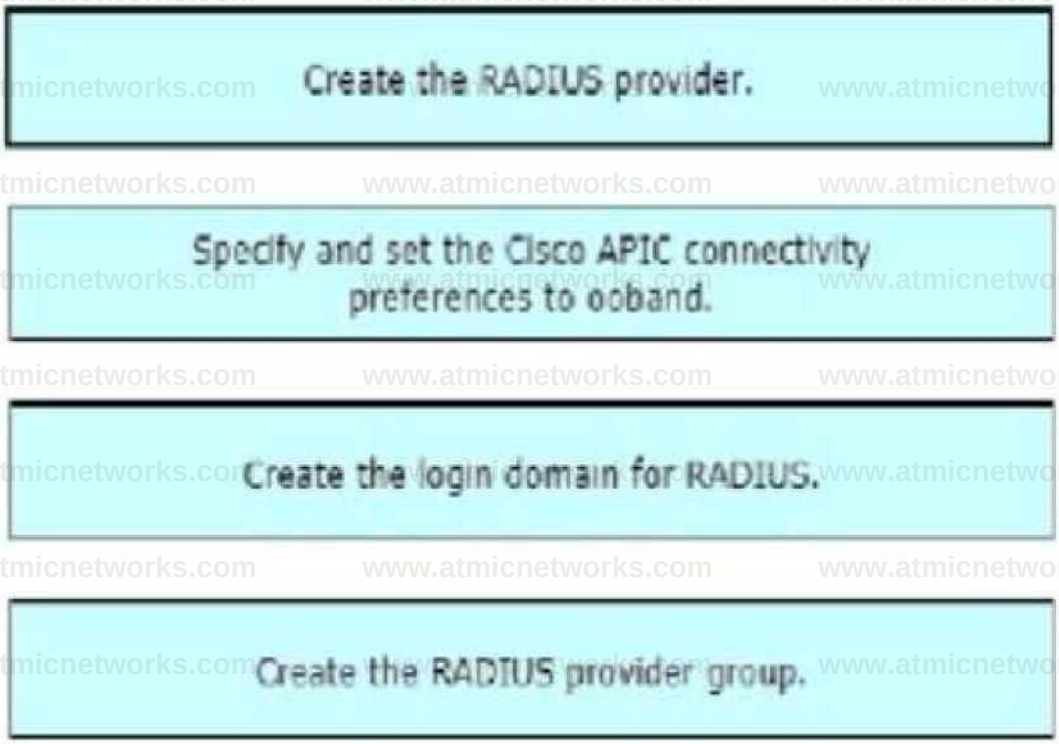

Question: 104

An administrator must migrate the vSphere Management VMkernel of all ESXi hosts in the production cluster from the standard default virtual switch to a VDS that is integrated with APIC in a VMM domain. Which action must be completed in this scenario?

A. The Management VMkernel EPG resolution must be set to Pre-Provosion.

B. The administrator must create an in-band VMM Management EPG before performing the migration.

C. The administrator must set the Management VMkernel BD resolution immediacy to On-Demand.

D. The VMkernel Management BD must be located under the Management Tenant.

Answer: A

Explanation:

When migrating the vSphere Management VMkernel of all ESXi hosts from the standard default virtual switch to a Virtual Distributed Switch (VDS) that is integrated with APIC in a VMM domain, it is essential to set the Management VMkernel EPG resolution to Pre-Provision. This action ensures that the necessary policies are in place on the ACI fabric before the migration occurs, allowing for a seamless transition and continuous management connectivity.

Question: 105

A customer implements RBAC on a Cisco APIC using a Windows RADIUS server that is configured with network control policies. The APIC is as follows:

Tenant = TenantX

Security Domain = Tenantx-SD

User = X

The customer requires User X to have access to TenantX only, without any extra privilege in the Cisco ACI fabric domain. Which Cisco AV pair must be implemented on the RADIUS server to meet these requirement?

A. shell:domains = TenantX-SD/fabric-admin/,common//read-all

B. shell:domains = TenantX-SD/tenant-admin

C. shell:domains = TenantX-SD/tenant-ext-admin/,common//read-all

D. shell:domains = TenantX-SD/tenant-admin/,common//read-all

Answer: B

Explanation:

To restrict User X to have access only to TenantX without any extra privileges in the Cisco ACI fabric domain, the Cisco AV pair that must be implemented on the RADIUS server is shell:domains = TenantX-SD/tenant-admin. This AV pair assigns User X the role of tenant admin within the security domain of TenantX, ensuring that the user has the necessary permissions to manage resources within TenantX and no additional privileges outside of this scope.

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

x/Security_config/b_Cisco_APIC_Security_Guide/b_Cisco_APIC_Security_Guide_chapter_0100.html

Configuring an AV Pair on the External Authentication Server

The numerical value within the parentheses in the attribute/value (AV) pair string is used as the UNIX user ID of the user who is togged tn using Secure Shell (SSH) or Telnet.

Procedure

Configure an AV pair on the external authentication server. The Cisco AV pair definition is as follows (Cisco supports AV pa

Example:

* shell:domains = domain/t/writeRolel|writeRole2]writeRole3/readRolel|readRole2ₓdomainB/writeRolel|writeRole2 vri
    * shell:domains = doirainA/writeRolel(writeRo lei writeRole3/readRolel|readRole2,donainB/writeRolel|unite

These are the boost regexes supported by APIC:

uid_regex("shell:domaiₙs\\s*[=:]\\s*((\\S+?/\\5'?/\\S*?)(ₓ\\5+?/\\S*?/\\S*?){0ₓ31})(\\(\\d+\\))$")I

regex("shell:domains\\s*[=]\\s*((\\S+?/\\S*?/\\5*?)(,\\S+?/\\S*?/\\S*?){0,31})$");

The following is an example:

shell:domains - coke/tenant-admin/read-all^peps i//read-all(16001)

Question: 106

A network engineer must backup the PRODUCTION tenant. The configuration backup should be stored on the APIC using a markup language and contain all secure information. Which export policy must be used to meet these requirement?

A)

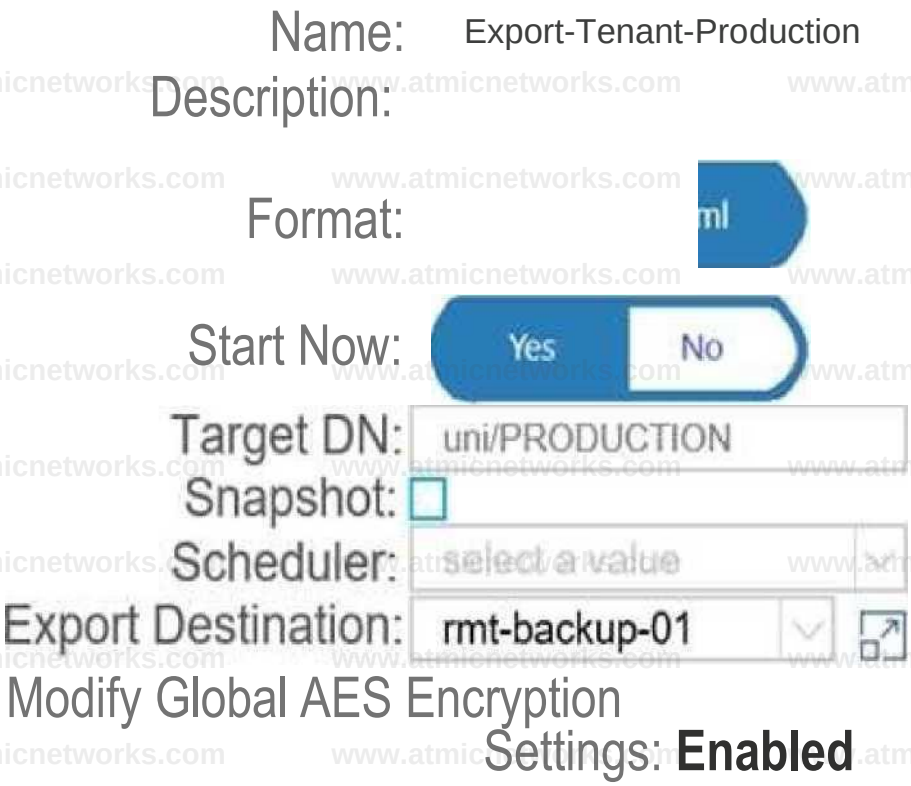
Name: Export-Tenant-Production
Description:
Format: ml
Start Now: Yes No
Target DN: uni/PRODUCTION
Snapshot:
Scheduler: select a value
Export Destination: rmt-backup-01
Modify Global AES Encryption Settings: Enabled

B)

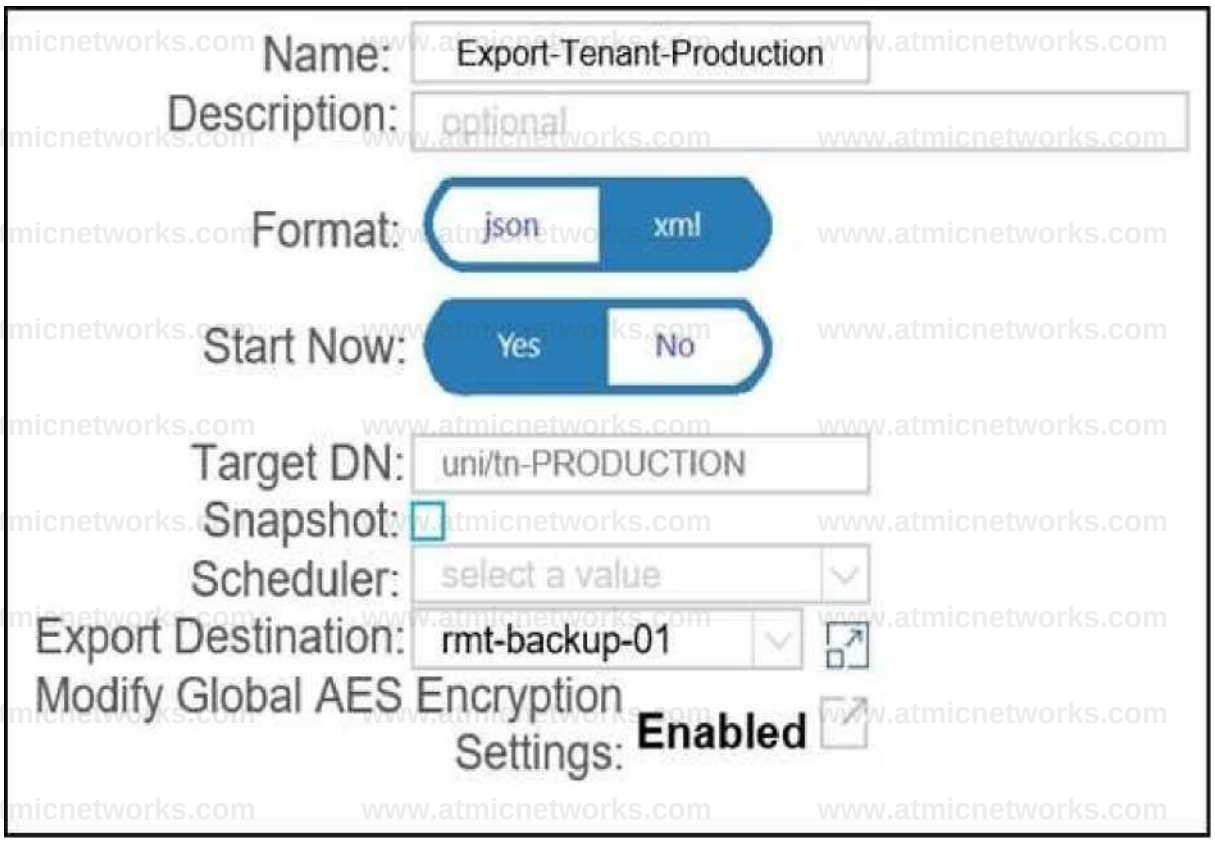
Name: Export-Tenant-Production
Description: optional
Format: json xml
Start Now: Yes No
Target DN: uni/tn-PRODUCTION
Snapshot:
Scheduler: select a value
Export Destination: rmt-backup-01
Modify Global AES Encryption Settings: Enabled

C)

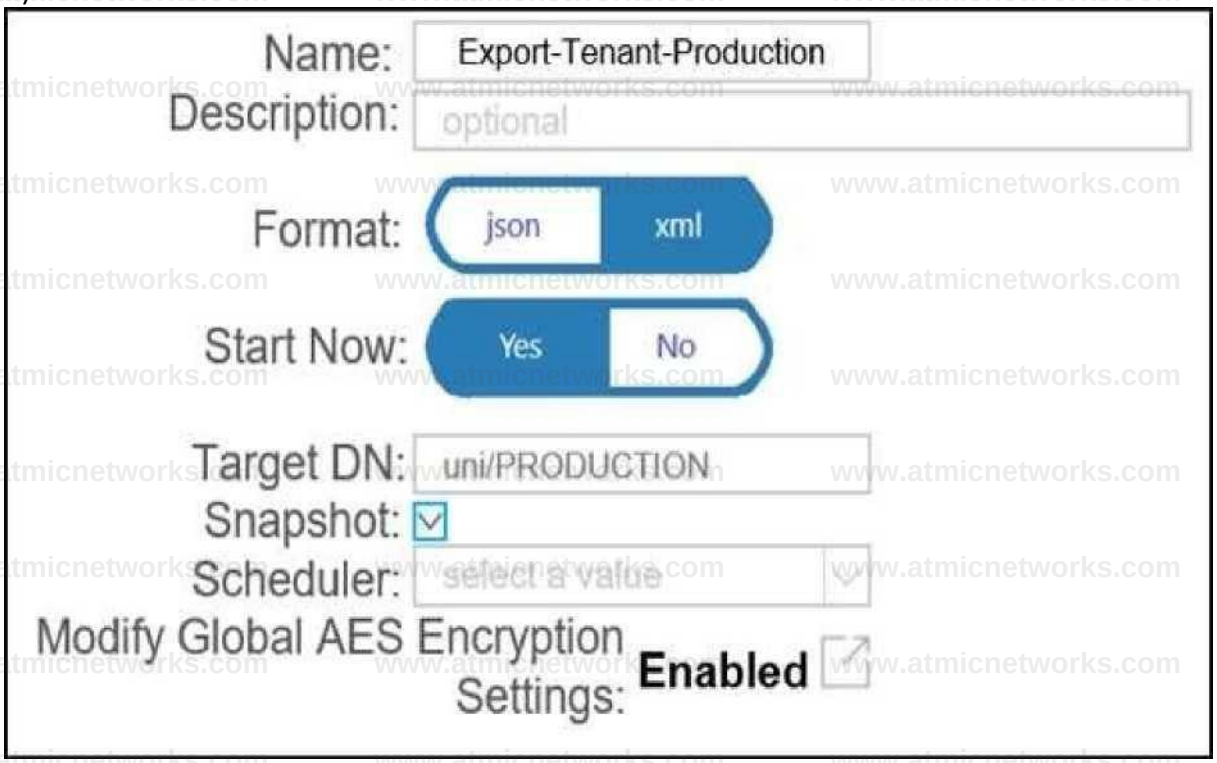

D)

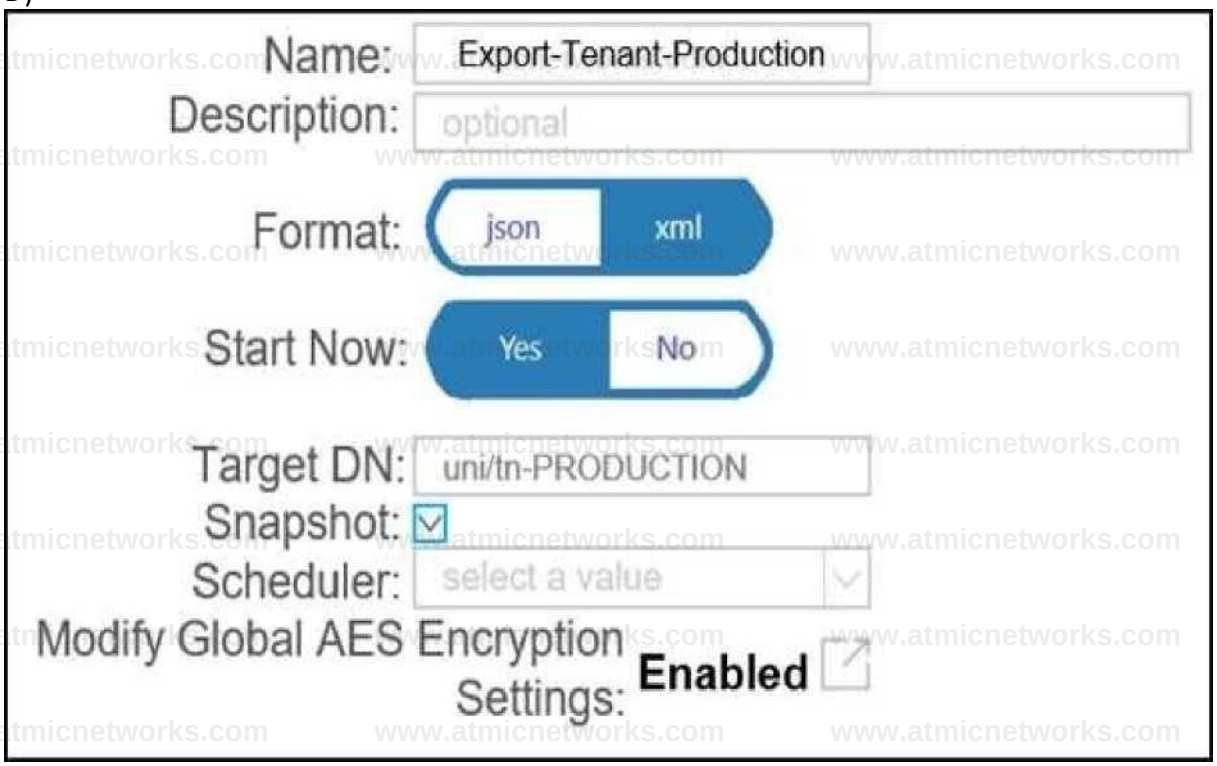

A. Option A

B. Option B

C. Option D

D. Option D

Answer: A

Explanation:

To backup the PRODUCTION tenant configuration on the APIC using a markup language and include all secure information, the network engineer must use an export policy that allows for exporting in a markup language format such as XML or JSON and includes secure attributes. In this case, the correct export policy is Option A, which shows an interface with "Export-Tenant-Production" where the format can be selected as 'xml' or 'json', and there is an option to modify global AES encryption settings, indicating that secure information can be included in the backup.

Reference:

Cisco Application Policy Infrastructure Controller (APIC) - Cisco APIC

Cisco Application Centric Infrastructure Best Practices Guide - Cisco ACI Design Guide

The image shows a graphical user interface for an export policy named "Export-Tenant-Production." The interface provides options to select the format of the backup (either 'json' or 'xml'), whether to start now with options 'Yes' or 'No', a field for Target DN with '/tn-PRODUCTION' filled in, a scheduler dropdown menu, and a checkbox for modifying global AES encryption settings which is enabled. This image is relevant because it depicts how to configure an export policy on Cisco's APIC to backup tenant configurations securely.

Question: 107

An engineer must create a backup of the Cisco ACI fabric for disaster recovery purposes. The backup must be transferred over a secure and encrypted transport. The backup file must contain all user and password related information. The engineer also wants to process and confirm the backup file validity by using a Python script. This requires the data structure to have a format similar to a Python dictionary. Which configuration set must be used to meet these requirements?

A. Under the Create Remote location settings, select Protocol: FTP

Under the Export policy, select

- Format: XML
- Modify Global AES Encryption Settings: Enabled

B. Under the Create Remote location settings, select Protocol: FTP

Under the Export policy, select

- Format: XML
- Modify Global AES Encryption Settings: Disabled

C. Under the Create Remote location settings, select Protocol: SCP

Under the Export policy, select

- Format: JSON
- Modify Global AES Encryption Settings: Disabled

D. Under the Create Remote location settings, select Protocol: SCP

Under the Export policy, select

- Format: JSON
- Modify Global AES Encryption Settings: Enabled

Answer: D

Explanation:

To create a backup of the Cisco ACI fabric for disaster recovery that is secure, encrypted, and in a format conducive to processing with a Python script, the engineer must select Secure Copy Protocol (SCP) for secure and encrypted transport. The backup file should be in JSON format, which is similar to a Python dictionary and can be easily processed by a Python script. Additionally, enabling Global AES Encryption ensures that all user and password-related information is included securely in the

backup1.

Question: 108

The Application team reports that a previously existing port group has disappeared from vCenter. An engineer confirms that the VM domain association for the EPG is no longer present. Which action determines which user is responsible for the change?

A. Check the EPG audit logs for the 'deletion' action and compare the affected object and user.

B. Evaluate the potential faults that are raised for that EPG.

C. Examine the health score and drill down to an object that affects the EPG combined score.

D. Inspect the server logs to see who was logging in to the APIC during the last few hours.

Answer: A

Explanation:

To determine which user is responsible for the disappearance of a previously existing port group from vCenter, the engineer should check the EPG audit logs for any 'deletion' actions. By comparing the affected object and the user who performed the action, the engineer can identify the responsible individual2.

Question: 109

An application team tells the Cisco ACI network administrator that it wants to monitor the statistics of the unicast and BUM traffic that are seen in a certain EPG. Which statement describes the collection statistics?

A. All EPGs in the Cisco ACI tenant object must be enabled for statistics to be collected.

B. Cisco ACI does not capture statistics at the EPG level. Only statistics that pass through ACI contracts can be monitored.

C. EPG statistics can be collected only for VMM domains. If a physical domain exists, statistics are not collected.

D. The collection of statistics is enabled on the EPG level by enabling the statistics for unicast and BUM traffic.

Answer: D

Explanation:

To monitor the statistics of unicast and BUM traffic seen in a certain EPG, the collection of statistics must be

enabled at the EPG level. This is done by enabling the statistics specifically for unicast and BUM traffic within the EPG settings3.

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_01011.html

EPG Level Statistics

The application owner would like to be able to monitor network-related information fortheir application, such as the aggregate amount of traffic to a specific tier As an example, we will monitor the amount of traffic to the web tier of a given application In this example, the default monitoring policies are appropriate, and they are simply extracting them from the system to be consumed externally. This information is useful in scenarios such as a new release being pushed, and to make sure that no traffic anomalies are created after the push.

Question: 110

Which routing protocol is supported between Cisco ACI spines and IPNs in a Cisco ACI Multi-Pod environment?

A. OSPF

B. ISIS

C. BGP

D. EIGRP

Answer: A

Explanation:

In a Cisco ACI Multi-Pod environment, the supported routing protocol between Cisco ACI spines and IPNs (Inter-Pod Network) is Open Shortest Path First (OSPF)4.

Question: 111

An engineer must deploy Cisco ACI across 10 geographically separated data centers. Which ACI site deployment feature enables the engineer to control which bridge domains contain Layer 2 flooding?

A. GOLF

B. Multi-Site

C. Multi-Pod

D. Stretched Fabric

Answer: B

Explanation:

The Cisco ACI site deployment feature that enables an engineer to control which bridge domains contain Layer 2 flooding across geographically separated data centers is Multi-Site. This feature allows for the extension of Layer 2 and Layer 3 connectivity between different locations with consistent policy enforcement5.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/aci_multi-site/sw/2x/fundamentals/Cisco-ACI-Multi-Site-Fundamentals-Guide-211/Cisco-ACI-Multi-Site- Fundamentals-Guide-211_chapter_011.html#id_51188

Question: 112

Which class of ACI object is presented in this output?

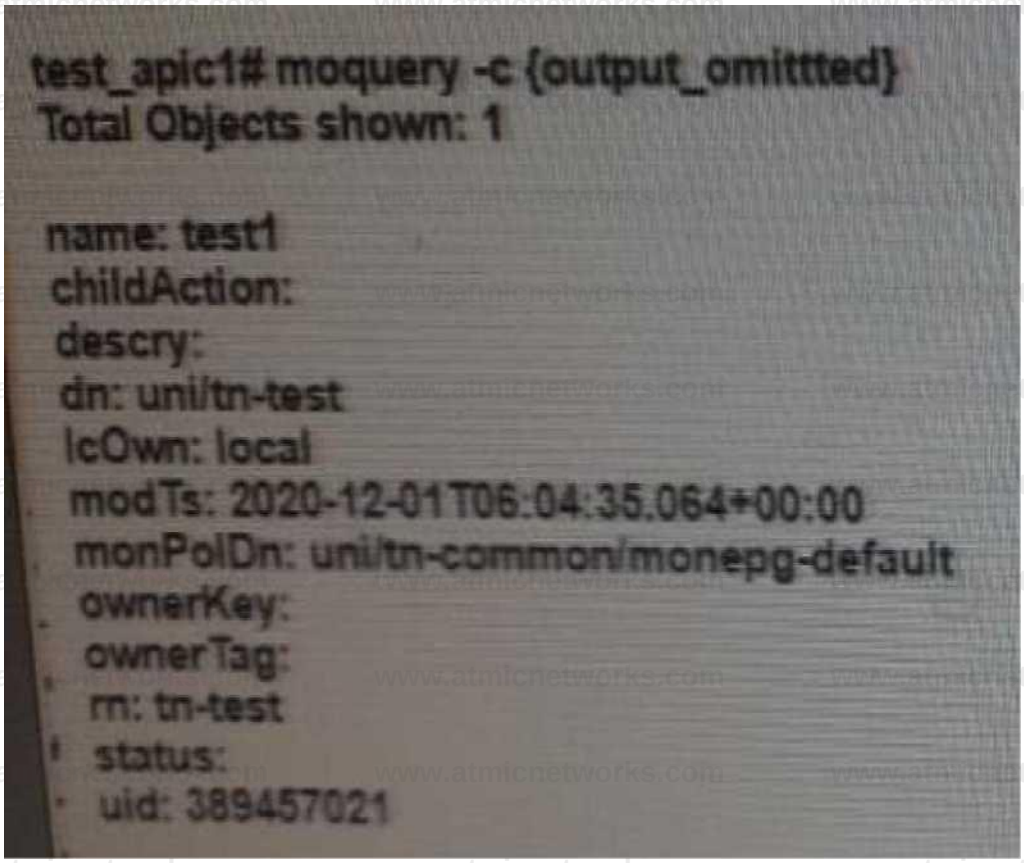

A. Contract

B. Bridge Domain

C. Tenant

D. Endpoint

Answer: C

Explanation:

In Cisco ACI, the object classes are used to categorize different types of managed objects within the ACI model. The output presented is indicative of a Tenant class object. Tenants in ACI are a top-level container for application policies, essentially providing a unit of isolation from other tenants. Each tenant can contain its own application policies, services, and network policies, and they are separate from other tenants to ensure multi-tenancy1.

Reference:

Cisco ACI Policy Model Guide1

Object Renaming in Cisco ACI - Cisco2

Application Centric Infrastructure (ACI) REST API Guide - Cisco DevNet

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/openstack/ACI- Installation-Guide-for-Red-Hat-Using-OSP13-Director/m-configuring-ironic-for-openstack.html

Question: 113

What is the effect of enabling the disable Remote EP learn feature?

A. It disables remote IP endpoint learning on all leaf nodes in the fabric.

B. It disables remote IP endpoint learning on leaf switches that do not have L3Outs.

C. It limits learning of compute leaf endpoints on border leaves.

D. It prevents border leaf switches from receiving routes through peering with external routers.

Answer: D

Explanation:

Enabling the “disable Remote EP learn” feature in Cisco ACI has the effect of preventing border leaf switches from receiving routes through peering with external routers. This feature is particularly useful when there is a mix of Gen1 and Gen2 hardware in the fabric, as it clears all the remote endpoints from the border leaf only. Traffic will still use hardware proxy if the endpoint is not learned on the border leaf, ensuring no operational impact12.

Reference:

ACI Fabric Endpoint Learning White Paper2

Cisco Community Discussion on Disable Remote EP Learn

https://unofficialaciguide.com/2018/11/29/aci-best-practice-configurations/

Question: 114

What are two descriptions of ACI multi-site? (Choose two.)

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

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

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

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

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

Answer: D, E

Explanation:

ACI Multi-Site architecture is designed to interconnect geographically dispersed data centers and extend Layer 2 and Layer 3 connectivity between those locations with consistent policy enforcement. It allows for either

a single APIC cluster to manage multiple ACI sites or supports a dedicated APIC cluster per site. This architecture ensures a scalable and flexible approach to managing multiple data center sites, providing a unified policy framework and operational model across the sites34.

Reference:

Cisco Multi-Site Deployment Guide for ACI Fabrics3

Cisco ACI Multi-Site Architecture White Paper4

Question: 115

An engineer must connect a new host to port 1 »'1 on Leaf 101. A Cisco ACI fabric has an MOP policy configured but experience excessive Layer 2 loops The engineer wants the Cisco ACI fabric to detect and prevent Layer 2 loops m the fabric Which set of actions accomplishes these goals'?

Enabte MCPgfOtai y
Associate Ihe MCP policy with an interface selector

Enable MCP globatty
Associate Ihe MCP pokey wth an interface policy group

EMe MCP oca^y
Asscwe tfe MCP policy WMF an interface policy group

Enable MCP locally
Associate the MCP policy ^ ar interface profifc

A. Option A

B. Option B

C. Option C

D. Option D

Answer: B

Explanation:

To detect and prevent Layer 2 loops in a Cisco ACI fabric, the engineer must configure the MisCabling Protocol (MCP) and related policies. MCP is designed to detect loops in Layer 2 network segments connected to ACI access ports and operates in conjunction with Spanning Tree Protocol (STP) running on external Layer 2 networks. The steps to accomplish these goals include:

Enable MCP Globally: Ensure that MCP is enabled globally on all access ports, virtual ports, and virtual port channels (VPCs) unless they are disabled at the individual port level1.

Configure MCP Transmit Frequency: Set the transmit frequency to a value that allows for quick loop detection. Starting with the 3.2(1) release, the Cisco ACI fabric provides faster loop detection with transmit frequencies from 100 milliseconds to 300 seconds1.

Set Action on Loop Detection: Decide how the MCP policies will act upon loop detection. Options include generating a syslog message or disabling the port upon detection of a loop1.

Implement Error Disabled Recovery Policy: Configure an error disabled recovery policy to automatically re-enable ports that were disabled due to loop detection after a configurable interval1.

By following these steps, the engineer can ensure that the Cisco ACI fabric will detect and prevent Layer 2 loops, thereby maintaining a stable and efficient network environment.

Reference:

Cisco APIC Online Help - Loop Detection1

ACI Layer 2 loop detection and Mitigation - Cisco Video Portal2

Question: 116

A network engineer is integrating a new Hyperflex storage duster into an existing Cisco ACI fabric The Hyperflex cluster must be managed by vCenter so a new vSphere Distributed switch must be created In addition the

hardware discovery must be performed by a vendor-neutral discovery protocol Which set of steps meets these requirements'?

Configure an Interface Policy group. select COP and apply 1 Io the dewed interfaces Enter (Jie '/Center IP and credentials in (he Create (Center Cento er dialog box In the Create VMware VMM domam dialog box, select Read-Only I,lode

Configure an Interface Policy group, select COP. and apply it to the designated interlaces
Create a VMwate VMM dew add il Io Ihe VIAN pool and assoctfie it Io the designated interlaces Select Read Only Mode In the Create VMwate VMM domain dialog box

Cordie an Interface Policy group, select LLOP. and apply il to the selected interfaces
Create a VLAN pool add it to the VMware VMM domain art include the appropriate interfaces
Enter Ihe /Center IP and credentials m the Create /Center Cceltoer dialog box

Configure a Switch Policy group select LLOP and apply it to the md uM interfaces
Set up a VMware VMM domain and apply it to the approprate interfaces
Enter the APtC management IP and credentials m the Create /Center Controller dialog box

A. Option A
B. Option B
C. Option C
D. Option D

Answer: B

Explanation:

To integrate a new Hyperflex storage cluster into an existing Cisco ACI fabric and manage it with vCenter, the following steps should be taken:

Create a new vSphere Distributed Switch (vDS): This is done within the vCenter interface. The vDS acts as a single virtual switch across all associated hosts in the data center, providing centralized management and monitoring of the network configuration1.

Configure the vDS for the Hyperflex cluster: This involves setting up the correct port groups, VLANs, and uplink profiles to ensure proper network segmentation and performance2.

Enable hardware discovery using a vendor-neutral protocol: This could involve using protocols like LLDP (Link Layer Discovery Protocol) or CDP (Cisco Discovery Protocol), which are supported by vCenter and can help in identifying and managing physical network devices3.

Integrate the Hyperflex cluster with ACI using the Application Policy Infrastructure Controller (APIC): This involves creating the necessary policies and profiles in ACI to manage the Hyperflex storage traffic effectively4.

Reference:

Cisco HyperFlex Systems Installation Guide for VMware ESXi1

Cisco HyperFlex Virtual Server Infrastructure 3.0 with Cisco ACI 3.2 and VMware vSphere 6.52

How to Set Up Networking with vSphere Distributed Switches - VMware Docs3

Tutorial: How to integrate HyperFlex and ACI with VMM (VMware) and UCSM4

Question: 117

A network engineer must configure a Cisco ACI system to detect network loops for untagged and tagged traffic The loop must be detected and slopped by disabling an interface within 4 seconds Which configuration must be used?


Admin State
Disabled
Enabled
Controls
Enable MCP PDU per VLAN
Key
Confirm Key
Loop Detect Multiplication Factor
2
t^h i#^^^ n^tui wwtatlikl
IK
I >WMilll*i<<i IwJ I
J


Admin State
Disabled
Enabled
Key
Confirm Key
Loop Detect Multiplication Factor
4
Loop Protection Action
Port Disable
Initial Delay (sec)
180
Transmission Frequency (sec)
1
(msec)
0


Admin State
Disabled
Enabled

A. Option A

B. Option B

C. Option C

Answer: B

Explanation:

To configure a Cisco ACI system to detect network loops for both untagged and tagged traffic and to stop the loop by disabling an interface within 4 seconds, the following configuration must be used:

Enable Mis-Cabling Protocol (MCP): MCP detects loops from external sources and will err-disable the interface on which Cisco ACI receives its own packet. This is crucial for preventing loops in the network1.

Configure MCP Transmit Frequency: Set the MCP transmit frequency to a value that allows for quick loop detection. The Cisco ACI fabric provides faster loop detection with transmit frequencies from 100 milliseconds to 300 seconds. To meet the requirement of detecting and stopping a loop within 4 seconds, you would set the transmit frequency to a value less than 4 seconds2.

Set Action on Loop Detection: Configure the MCP policies to identify loops and decide how to act upon them. You can set the policy to generate a syslog message or disable the port upon loop detection2.

Implement Error Disabled Recovery Policy: Configure an error disabled recovery policy to automatically re-enable ports that were disabled due to loop detection after a configurable interval2.

By implementing these configurations, the Cisco ACI system will be able to detect and stop network loops quickly and efficiently, ensuring network stability and preventing potential disruptions.

Reference:

Cisco APIC Online Help - Loop Detection2

Cisco ACI Best Practices Quick Summary

Question: 118

Refer to the exhibit.

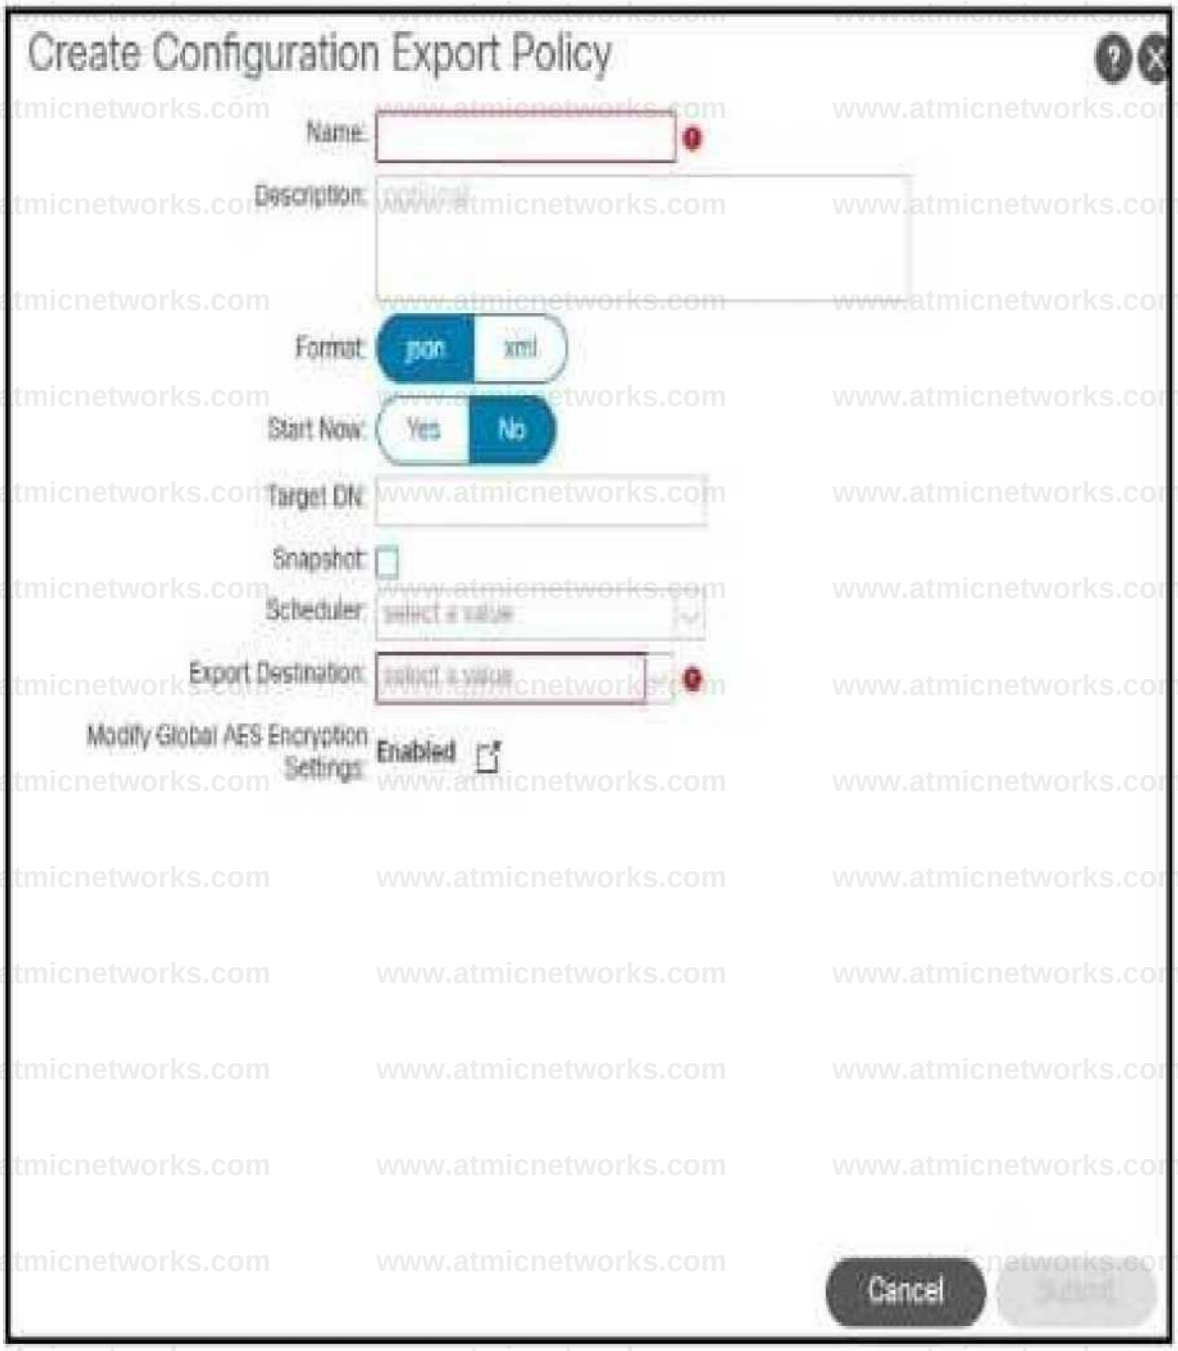

Refer to the exhibit A customer must back up the current Cisco ACl configuration securely to the remote location using encryption and authentication. The backup job must run once per day The customer s security policy mandates that any sensitive information including passwords, must not be exported from the device Which set of steps meets these requirements?

Expert destination using FTP protocol
Use XML format

Export destination using SOP protocol
Disable Global AES Encryption

Expert dest nation using SCP protocol
Use XML format

Export destination using FTP protocol
Disable Global AES Encryption

A. Option A

B. Option B

C. Option C

D. Option D

Answer: B

Explanation:

To securely back up the current Cisco ACI configuration to a remote location with encryption and authentication, and ensure that sensitive information is not exported, the following steps should be taken:

Create a Remote Location: Define a remote location where the backup will be stored. This involves specifying the hostname or IP address of the remote server and selecting the protocol (SCP/FTP/SFTP) to be used1.

Create a Configuration Export Policy: Set up an export policy that defines the format (XML/JSON) and the destination for the backup. This policy will also include the schedule for the backup to run once per day1.

Configure Global AES Encryption: To ensure that the backup is encrypted, configure the Global AES Encryption setting. This will allow for hashed secure properties, such as passwords and certificates, to be exported in an encrypted form. It's important to configure a passphrase that is at least 16 characters long for the encryption1.

Schedule the Backup Job: Use the scheduler to set up the backup job to run once per day. This ensures that the backup process is automated and adheres to the customer's security policy1.

Exclude Sensitive Information: To comply with the security policy that mandates sensitive information not be exported, ensure that the configuration export policy is set to exclude secure fields unless encryption is enabled1.

Reference:

Cisco Guide on Creating a Backup for Your APIC Cluster1

Cisco Community Discussion on Backing Up ACI Configuration2

Question: 119

What is MP-BGP used for in Cisco ACl fabric?

A. MP-BGP VPNv4 AF is used to propagate L3Out routes that are received from a border leaf to the fabric.

B. MP-BGP VPNv4 AF is used between spines in an ACI Multi-Pod fabric to propagate the endpoint

C. MP-BGP VPNv4 AF is used as protocol on L3Out between a border leaf and an external router

D. MP-BGP Layer 2 VPN EVPN AF is used to propagate L30ut routes that are received from a border leaf.

Answer: A

Explanation:

In Cisco ACI fabric, MP-BGP (Multi-Protocol Border Gateway Protocol) is utilized for several purposes, one of which includes the propagation of L3Out (Layer 3 Out) routes within the fabric. Specifically, MP-BGP with VPNv4 address family (AF) is used in the ACI infra VRF (overlay-1 VRF) to distribute external routes from a border leaf to other leaf switches1. This allows for the efficient dissemination of routing information and ensures that all relevant leaf switches are aware of the routes necessary to reach external networks.

Reference:

ACI Fabric L3Out White Paper - Cisco1

Configuring ACI Fabric BGP Route Reflectors - Cisco2

ACI Fabric - Underlying protocols - Cisco Community3

Cisco Application Centric Infrastructure Fundamentals, Release 5.3(x)4

Question: 120

An engineer must securely export Cisco APIC configuration snapshots to a secure, offsite location The exported configuration must be transferred using an encrypted tunnel and encoded with a platformagnostic data format that provides namespace support Which configuration set must be used?

Policy Dor Policy Protocol' TLS
Fimi XML

•'Policy Export Rhcy PJOLOMI TLS Fcr-at £$|

Policy Import Policy protocol SCP
FomaL JSD*J

Policy Export Policy Protocol- SEP
Fermat XML

A. Option A

B. Option B

C. Option C

D. Option D

Answer: B

Explanation:

To securely export Cisco APIC configuration snapshots to a secure, offsite location using an encrypted tunnel and a platform-agnostic data format that provides namespace support, the following configuration set must be used:

Choose a Platform-Agnostic Data Format: Export the configuration in a data format like JSON or XML, which are platform-agnostic and support namespaces1.

Use an Encrypted Tunnel: Transfer the configuration snapshot over a secure protocol such as SFTP or SCP, which provides an encrypted tunnel for the data transfer1.

Schedule the Export: Set up a scheduled export of the configuration in the Cisco APIC to occur at regular intervals, ensuring the process is automated1.

Enable Encryption: Make sure the configuration export policy includes encryption settings. Cisco APIC supports AES-256 encryption for securing exported configuration files2.

Verify Namespace Support: Ensure that the chosen data format and the method of export support the use of namespaces, which are essential for organizing elements and attributes in XML documents1.

By following these steps and choosing Option B, the engineer can meet the requirements for securely exporting Cisco APIC configuration snapshots to an offsite location.

Question: 121

A Cisco APIC is configured with RADIUS authentication as the default The network administrator must ensure that users can access the APIC GUI with a local account if the RADIUS server is unreachable. Which action must be taken to accomplish this goal?

A. Create an additional login domain that references local accounts

B. Enable the fallback check with the default authentication domain

C. Associate console authentication with the "RADIUS" realm.

D. Reference the local realm in the fallback domain

Answer: B

Explanation:

To ensure that users can access the Cisco APIC GUI with a local account if the RADIUS server is unreachable, the network administrator must enable the fallback check with the default authentication domain. This allows for local authentication as a backup method when RADIUS is not available. The fallback mechanism is crucial for maintaining access to the APIC GUI in case of RADIUS server issues1.

Reference:

Cisco Community Discussion on Auth Radius fallback to Local1

Cisco APIC Security Configuration Guide

Question: 122

A network engineer must allow secure access to the Cisco ACl out-of-band (OOB) management only from external subnets 10 0 0024 and 192.168 20 G'25. Which configuration set accomplishes this goal?

Cress a LSCut in the MOW 7 Lentil ir OOB VRF
Set Ow.31 Management N^ow Inside Profile as a consumer at (he 006 Kfliaet
Crtafe an Exltfra[1] EPG /nth Iwa 'iuinei eiltrie- with If- - cMcnel SUL-HEK

Creel? a PER Soviet gianh n HIE I^FT teharU
Ci idle a mataQHne 1 Profile w=lh Uis rEquh ed 000 EPC
Redirect all traffic going mtfi ACI managef^nt to Lhecxltrr^[1] linewall
Create MO subnet entraiJ witrthe OOP Sric^ ta^n ufo tin i&ired sulfate

Create an 008 contact mai allays tne required SKHS
Provide Uie carvia:1 from Bw 000 z[J]G.
Cars, me *e zon^aerb? ms IV-E ExTem*! MsnawmeRi frer.'.^fk-i^nrce Prof Jr
Create IM subnet mies[1] n ^ e Extern# Managenwt Work P'oJiie wn n the reared subnets

Create an EPS and 00 In the MGMT lenani ft MS VRF
WE 009 YR- lo j-rguid? ft? ^Sfl
$a A new EFG lo consume re 006 conrad

A. Option A

B. Option B

C. Option C

D. Option D

Answer: C

Explanation:

To enable the APIC in a Cisco ACI fabric using out-of-band management connectivity to access a routable host with an IP address of 192.168.11.2, you need to add a Fabric Access Policy that allows management connections. This involves configuring a contract that will be consumed and provided to your OOB devices. The contract will let the system know what traffic is allowed. In this case, you will use the default/common contract to permit any traffic. This is necessary because the APIC needs

to be able to send traffic to and receive traffic from the management network1.

Reference:

ACI: Configuring Out-of-Band (OOB) Access for Your Fabric - Cisco1

Troubleshoot ACI Management and Core Services - In-band and Out-of-band Management - Cisco2

Question: 123

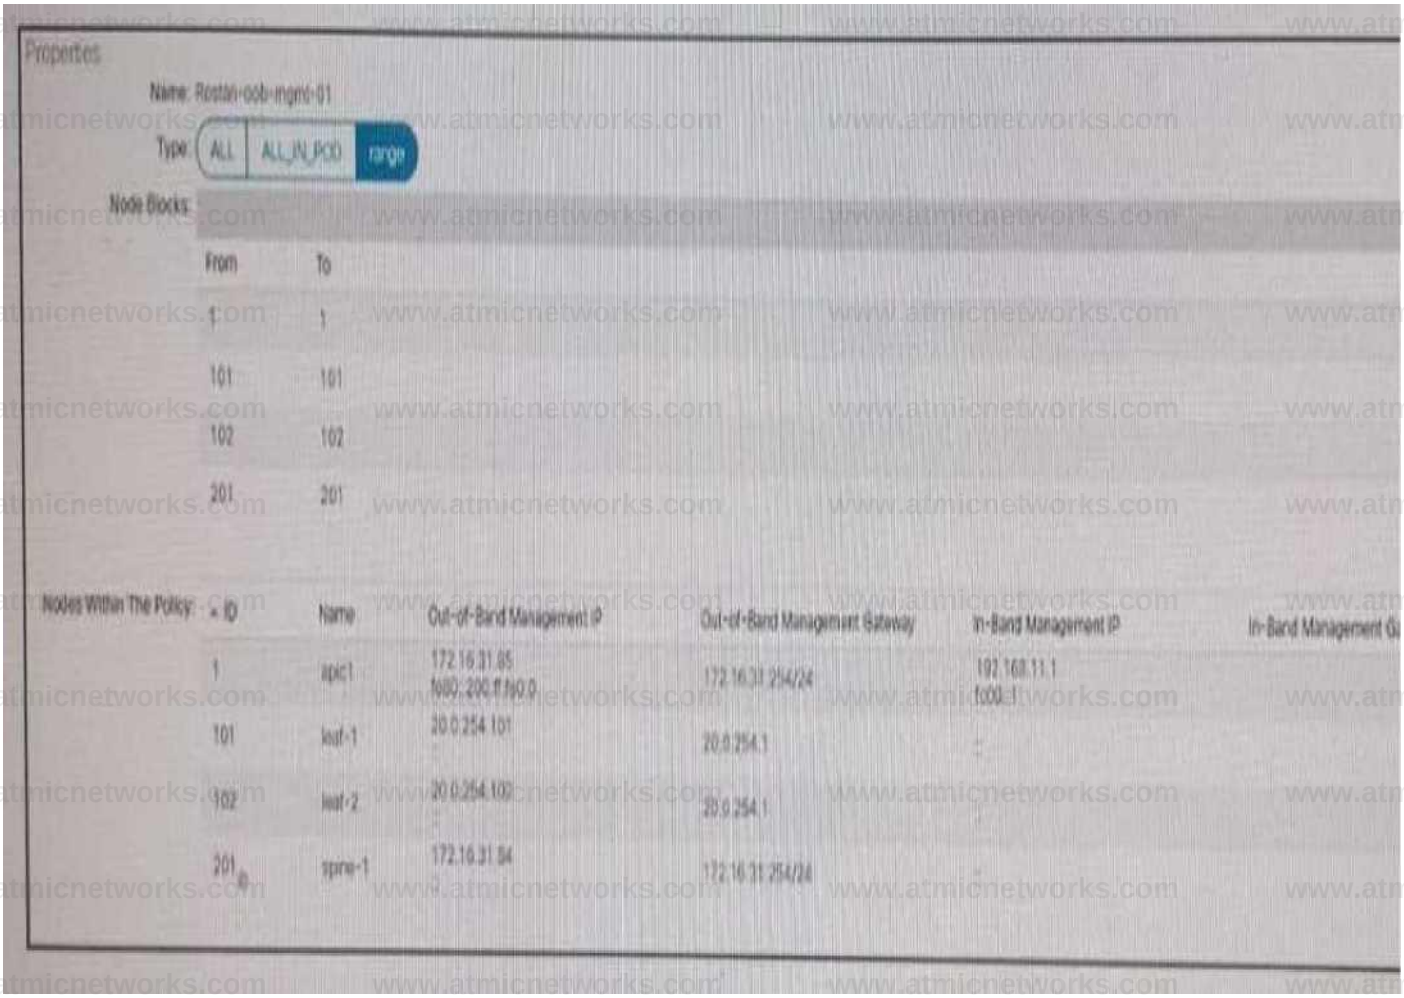

Refer to the exhibit A Cisco ACI fabric is using out-of-band management connectivity The APIC must access a routable host with an IP address of 192 168 11 2 Which action accomplishes this goal?

A. Change the switch APIC Connectivity Preference to in-band management

B. Remove the in-band management address from the APIC.

C. Add a Fabric Access Policy to allow management connections.

D. Modify the Pod Profile to use the default Management Access Policy

Answer: C

Explanation:

To enable the APIC in a Cisco ACI fabric using out-of-band management connectivity to access a routable host with an IP address of 192.168.11.2, you need to add a Fabric Access Policy that allows management connections. This involves configuring a contract that will be consumed and provided to your OOB devices. The contract will let the system know what traffic is allowed. In this case, you will use the default/common contract to permit any traffic. This is necessary because the APIC needs to be able to send traffic to and receive traffic from the management network1.

Reference:

ACI: Configuring Out-of-Band (OOB) Access for Your Fabric - Cisco1

Troubleshoot ACI Management and Core Services - In-band and Out-of-band Management - Cisco2

Question: 124

A bridge domain for an EPC called "Web Servers" must be created in the Cisco APIC. The configuration must meet these requirements:

Only traffic to known Mac addresses must be allowed to reduce noice.

The multicast traffic must be limited to the ports that are participating in multicast routing.

The endpoints within the bridge domain must be kept in the endpoint table for 20 minultes without any updates.

Which set of steps configures the bridge domain that satisfies the requirements?

A. Select the ARP Flooding checkbox.

Create an Endpoint Retention Policy with a Remote Endpoint Aging Interval of 20 minutes.

Set L3 Unknown Multicast Flooding to Optimized Flooding

B. Set L2 Unknown Unicast to Hardware Proxy.

Configure L3 Unknown Multicast Flooding to Optimized Flood.

Create an Endpoint Retention Policy with a Local Endpoint Aging interval of 1200 seconds.

C. Switch L2 Unknown Unicast to Flood.

Select the default Endpoint Retention Policy and set the Local Endpoint Aging to 20 minutes.

Set Multicast Destination Flooding to Flood in Encapsulation.

D. Multicast Destination Flooding should be set to Flood in BD.

Set L3 Unknown Multicast Flooding to Flood.

Select the default Endpoint Retention Policy with a Local Endpoint Aging Interval of 1200 seconds.

Answer: B

Explanation:

To configure a bridge domain for an EPC called "Web Servers" that meets the specified requirements, the following steps should be taken:

Set L2 Unknown Unicast to Hardware Proxy: This setting ensures that only traffic to known MAC addresses is allowed, which helps to reduce noise by preventing unknown unicast flooding within the bridge domain1.

Configure L3 Unknown Multicast Flooding to Optimized Flood: This setting limits multicast traffic to the ports that are participating in multicast routing, which optimizes the delivery of multicast packets1.

Create an Endpoint Retention Policy with a Local Endpoint Aging interval of 1200 seconds (20 minutes): This policy ensures that endpoints within the bridge domain are kept in the endpoint table for 20 minutes without any updates, maintaining the stability of the network1.

By following these steps, the bridge domain will be configured to satisfy the requirements for the EPC called "Web Servers" in the Cisco APIC environment.

Question: 125

The company's Cisco ACI fabric hosts multiple customer tenants. To meet a service level agreement, the company is constantly monitoring the Cisco ACI environment. Syslog is one of the methods used for monitoring. Only events related to leaf and spine environmental information without specific customer data should be logged. To which ACI object must the configuration be applied to meet these requirements?

A. access policy

B. infra tenant

C. switch profile

D. fabric policy

Answer: D

Explanation:

To ensure that only events related to leaf and spine environmental information are logged without including specific customer data, the configuration should be applied to the fabric policy. The fabric policy in Cisco ACI is used to define global settings that apply to the entire fabric, such as syslog settings, which can be configured to filter and forward only the desired types of system messages12.

Question: 126

A Cisco ACI is integrated with a VMware vSphere environment. The port groups must be created automatically in vSphere and propagated to hypervisors when created in the ACI environment.

Which action accomplishes this goal?

A. Associate the VMM domain with the EPGs that must be available in vCenter.

B. Assign the uplinks of the ESXi hosts to the vDS that the APIC created.

C. Configure contracts for the EPGs that are required on the ESXi hosts.

D. Create the port groups on the vCenter that reflect the EPG names in the APIC.

Answer: A

Explanation:

Associating the VMM domain with the EPGs (Endpoint Groups) that must be available in vCenter is the action that accomplishes the goal of creating port groups automatically in vSphere and propagating them to hypervisors when created in the ACI

environment. This integration allows the APIC to automatically create port groups in the VMware vCenter under the VDS (vSphere Distributed Switch), which provisions the network policy in the VMware vCenter

Question: 127

An engineer is troubleshooting fabric discovery in a newly deployed Cisco ACI fabric and analyzes this output:

An engineer Is troubleshooting fabric discovery In a newly deployed Cisco ACI fabric and analyzes this output:

```
LEAF 101 # show ip int brief vrf overlay-1 (...output truncated for brevity...) IP Interface Status for VRF"overlay•1"M
Interface          Address                 Interface Status
lo1023             10.233.44,32/32         protocol*up/link*upi'admin-up

LEAFIOlff show vlan extended
VLAN    Name              Encap             Ports

        infraidefault     vxlan-38802518, Eth1/1, Eth1/2. Eth1/47
                          vlan-3600
```

Which ACI fabric address is assigned to interface lo1023?

A. Dynamic tunnel endpoint

B. Physical tunnel endpoint

C. Fabric tunnel endpoint

D. VXLAN tunnel endpoint

Answer: C

Explanation:

In Cisco ACI, the interface lo1023 is assigned as a fabric tunnel endpoint (FTEP). This is a pervasive address found on every leaf and it's always the same. It is used mostly for AVS (Application Virtual Switch) when the infra VXLAN is extended out of the fabric1. The FTEP is crucial for the internal operation of the fabric, particularly for scenarios where the infrastructure VXLAN needs to be extended outside of the ACI fabric.

Reference:

Default IP interfaces on Fabric nodes - Cisco Community1

Question: 128

An ACI engineer is implementing a Layer 3 out inside the Cisco ACI fabric that must meet these requirements:

The data center core switch must be connected to one of the leaf switches with a single 1G link.

The routes must be exchanged using a link-state routing protocol that supports hierarchical network design.

The data center core switch interface must be using 802.1Q tagging, and each vlan will be configured with a dedicated IP address.

Which set of steps accomplishes these goals?

A. Set up the ElGRP Protocol policy with the selected Autonomous System number. Set up the Routed External Network object ana Node Profile, selecting ElGRP Create the Switch profile, selecting Portchannel and the appropriate interfaces Create the default network and associate it with the Routed Outside object.

B. Set up the BGP Protocol policy with the Autonomous System number of 0.

Configure an interface policy and an External Bridged Domain.

Create an External Bridged Network using the configured VLAN pool.

Build the Leaf profile, selecting Routed sub-interface and the appropriate VLAN.

C. Configure the OSPF Protocol policy with an area of 0.

Create Routed Outside object and Node Profile, selecting OSPF as the routing protocol. Build the Interface profile, selecting Routed Sub-interface and the appropriate VLAN. Configure the External Network object with a network of 0.0.0.0/0.

D. Set up the ElGRP Protocol policy with the selected Autonomous System number. Create the Routed Outside object and Node Profile selecting ElGRP Configure the Interface profile selecting Routed Interface and the appropriate interfaces. Create the External Network object with a network of 0.0.0.0/0.

Answer: C

Explanation:

To implement a Layer 3 out (L3Out) inside the Cisco ACI fabric that meets the specified requirements, the following steps should be taken:

Configure the OSPF Protocol policy with an area of 0: This step involves setting up OSPF, a link-state routing protocol that supports hierarchical network design, as the routing protocol for the L3Out1.

Create Routed Outside object and Node Profile, selecting OSPF: This step involves creating a Routed Outside object and associating it with a Node Profile that specifies OSPF as the routing protocol1.

Build the Interface profile, selecting Routed Sub-interface and the appropriate VLAN: This step involves creating an Interface profile for the connection to the data center core switch, using 802.1Q tagging for VLAN separation2.

Configure the External Network object with a network of 0.0.0.0/0: This step involves setting up the External Network object to advertise routes to the rest of the network1.

Question: 129

An engineer must advertise a bridge domain subnet out of the ACI fabric to an OSPF neighbor. Which two configuration steps are required? (Choose two.)

A. Configure Subnet scope to Advertised Externally

B. Add External Subnet for External EPG flag under External EPG.

C. Create Route Control Profile with the export direction under External EPG.

D. Add L30ut profile to the bridge domain using Associated L30uts section

E. Configure the Subnet under the EPG level.

Answer: A, D

Explanation:

To advertise a bridge domain subnet out of the ACI fabric to an OSPF neighbor, the following configuration steps are required:

Configure Subnet scope to Advertised Externally: This step involves setting the scope of the subnet within the bridge domain to be advertised externally, which allows the subnet to be propagated to external routing entities3.

Add L30ut profile to the bridge domain using Associated L30uts section: This step involves associating the bridge domain with an L3Out profile, which is necessary for the bridge domain subnet to be advertised to OSPF neighbors4.

Question: 130

An engineer must perform a Cisco ACI fabric upgrade that minimizes the impact on user traffic and allows only permitted users to perform an upgrade. Which two configuration steps should be taken to meet these requirements?

A. Divide Cisco APIC controllers into two or more maintenance groups.

B. Grant tenant-ext-admin access to a user who performs an upgrade

C. Combine all switches into an upgrade group.

D. Divide switches into two or more maintenance groups.

E. Grant the fabric administrator role to a user who performs an upgrade.

Answer: A, D

Explanation:

To perform a Cisco ACI fabric upgrade that minimizes the impact on user traffic and allows only permitted users to perform an upgrade, the following configuration steps should be taken:

Divide Cisco APIC controllers into two or more maintenance groups: This step involves organizing the APIC controllers into separate maintenance groups to allow for a phased upgrade process, reducing the impact on user traffic5.

Divide switches into two or more maintenance groups: This step involves organizing the switches into separate maintenance groups, which allows for a controlled and phased upgrade process, minimizing

the impact on the network during the upgrade5.

Reference: https://www.cisco.com/c/en/us/td/docs/dcn/whitepapers/cisco-application-centric- infrastructure-design-guide.html

Question: 131

An engineer is in the process of discovering a new Cisco ACI fabric consisting of two spines and four leaf switches. The discovery of leaf 1 has just been completed. Which two nodes are expected to be discovered next? (Choose two.)

A. spine 1

B. leaf 4

C. spine 2

D. leaf 3

E. leaf 2

Answer: A, C

Explanation:

In the process of discovering a new Cisco ACI fabric, after the discovery of leaf 1 has been completed, the next nodes expected to be discovered are the spine switches. This is because the spines are typically connected to the leaf switches and play a central role in the fabric's topology6.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/GSG/b_APIC_Getting_Started_Guide_Rel_2_x/b_APIC_Getting_Started_Guide_Rel_2_x_chapter_0 100.pdf

Question: 132

All workloads in VLAN 1001 have been migrated into EPG-1001. The requirement is to move the gateway address for VLAN 1001 from the core outside the Cisco ACI fabric into the Cisco ACI fabric.

The endpoints in EPG-1001 must route traffic to endpoints in other EPGs and minimize flooded traffic in the fabric. Which configuration set is needed on the bridge domain to meet these requirements?

A. Disable ARP Flood
Disable Limn Endpoint Learning

B. Enable Hardware Proxy Enable Unicast Routing

C. Disable Local IP Learning Limit Disable Unicast Routing

D. Enable Flood
Enable Unicast Routing

Answer: B

Explanation:

To move the gateway address for VLAN 1001 from the core outside the Cisco ACI fabric into the Cisco ACI fabric and ensure that endpoints in EPG-1001 route traffic to endpoints in other EPGs while minimizing flooded traffic in the fabric, the following configuration set is needed on the bridge domain:

Enable Hardware Proxy: This step involves enabling the hardware proxy feature, which allows the fabric to learn and maintain endpoint information more efficiently7.

Enable Unicast Routing: This step involves enabling unicast routing within the bridge domain, which allows for the routing of traffic between different EPGs and minimizes flooded traffic

Question: 133

Refer to the exhibit.

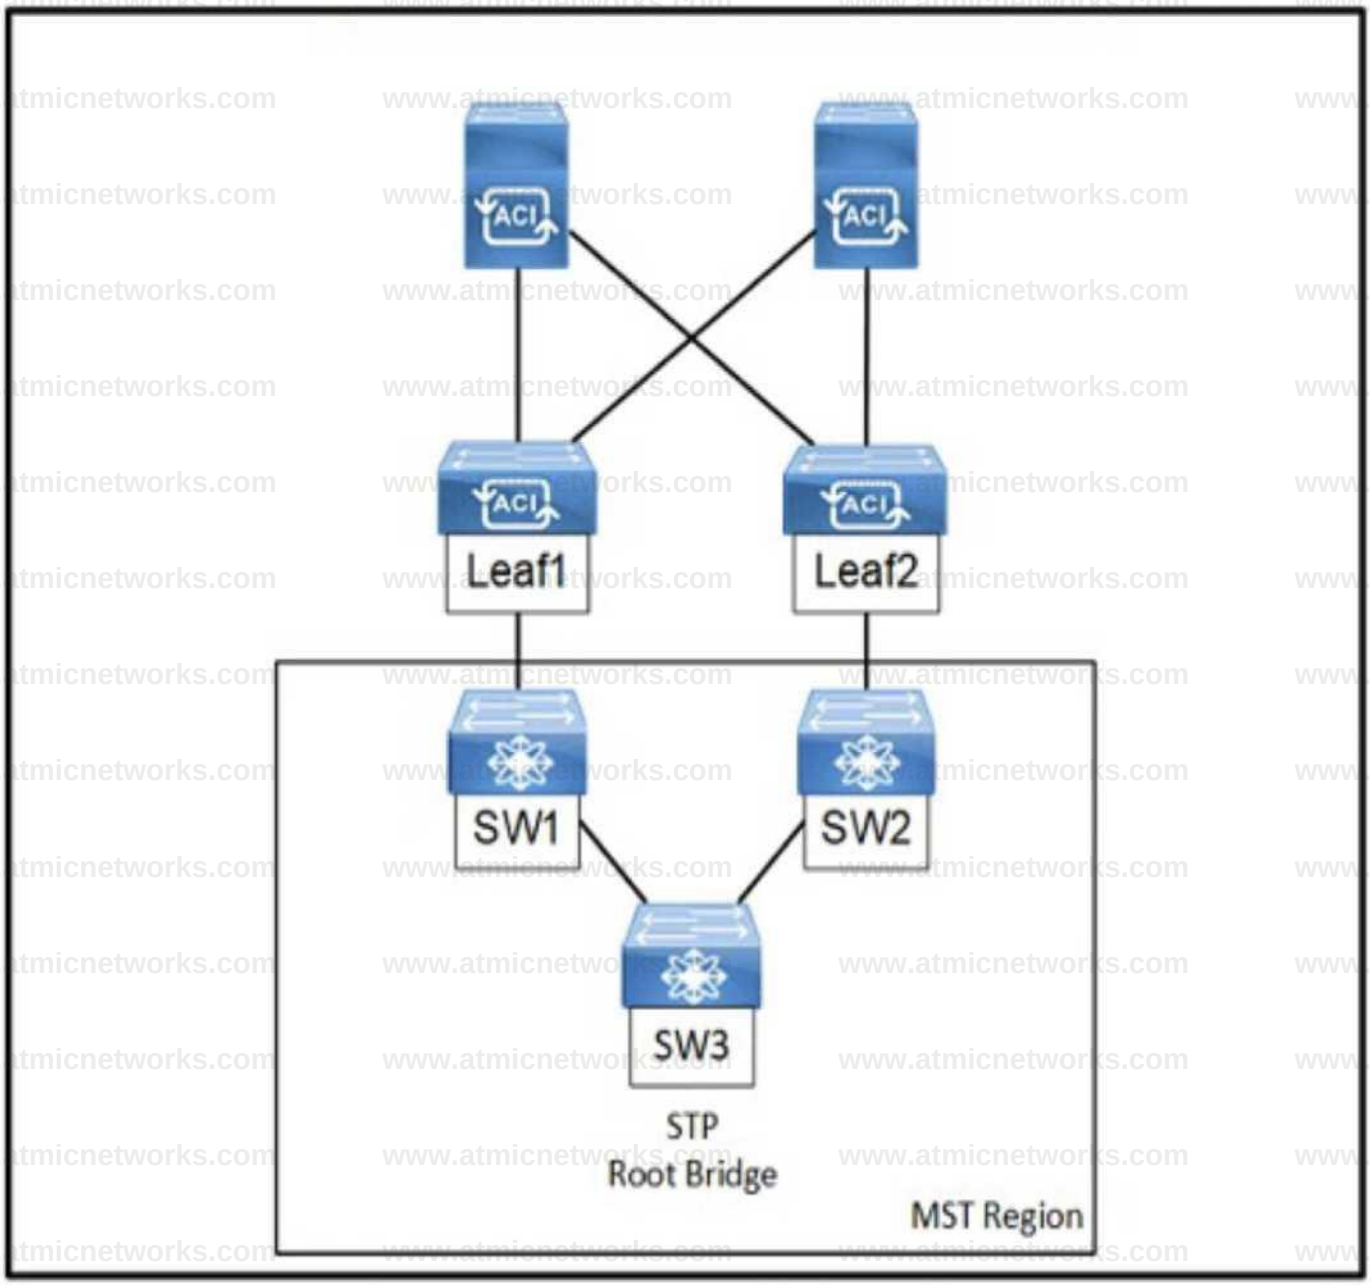

Refer to the exhibit, An engineer is deploying a Cisco ACI environment but experiences a STP loop between switch1 and switch2. Which configuration step is needed to break the STP loop?

A. Configure the STP instance to VLAN mapping under the switch STP policy.

B. Configure a Layer 2 external bridged network on the interfaces facing the MST switches.

C. Enable the native VLAN on the interfaces facing the MST switches using static pons in a dedicated EPG.

D. Enable BPDU filter under the STP interface policy on the Interfaces lacing the MST switches.

Answer: D

Explanation:

To break the STP loop between switch1 and switch2 in a Cisco ACI environment, enabling BPDU filtering on the interfaces facing the MST (Multiple Spanning Tree) switches is the necessary step. BPDU filtering prevents BPDUs from being sent or received through a port, which effectively stops the STP process on those ports and breaks the loop1.

When BPDU filtering is enabled on an interface, the switch does not send any BPDUs and ignores any BPDUs it receives on that interface. This can be used to prevent unwanted STP negotiations on ports where loops are not expected to happen or where STP is not desired1.

Reference:

Use of ACI Operation with L2 Switches and Spanning Tree Link Types - Cisco1

Question: 134

A customer creates Layer 3 connectivity to the outside network. However, only border leaf switches start receiving destination updates to other networks from the newly created L3Out. The updates must also be propagated to other Cisco ACI leaf switches. The L3Out is linked with the EPGs via a contract. Which action must be taken in the pod policy group to accomplish this goal?

A. Apply a BGP route reflector policy.

B. Enable a COOP policy.

C. Configure an IS-IS policy.

D. Implement an access management policy.

Answer: A

Explanation:

To ensure that destination updates from a newly created L3Out are propagated to all Cisco ACI leaf switches, a BGP route reflector policy should be applied. This policy allows the border leaf switches, which act as BGP route reflectors, to redistribute the learned routes to other leaf switches within the fabric1.

Question: 135

A network administrator configures AAA inside the Cisco ACI fabric. The authentication goes through the local users if the TACACS+ server is not reachable. If the Cisco APIC is out of the cluster, the access must be granted through the fallback domain. Which configuration set meets these requirements?

A. Ping Check: True

Default Authentication Realm: Local

Fallback Check: True

B. Ping Check: True

Default Authentication Realm: TACACS+

Fallback Check: False

C. Ping Check: False

Default Authentication Realm: Local

Fallback Check: False

D. Ping Check: False

Default Authentication Realm: TACACS+

Fallback Check: True

Answer: D

Explanation:

For AAA configuration within the Cisco ACI fabric, where authentication should fall back to local users if the TACACS+ server is unreachable, and access must be granted through the fallback domain if the Cisco APIC is out of the cluster, the configuration set that meets these requirements is to have the Default Authentication Realm set to TACACS+ with Fallback Check enabled. This ensures that when the TACACS+ server cannot be reached, the system will fall back to using local authentication

Question: 136

A Cisco ACI environment consists of multiple silent hosts that are often relocated between leaf switches. When the host is relocated, the bridge domain takes more than a few seconds to relearn the host's new location. The requirement is to minimize the relocation impact and make the ACI fabric relearn the new location of the host faster. Which action must be taken to meet these requirements?

A. Set Unicast Routing to Enabled.

B. Configure ARP Flooding to Enabled.

C. Set L2 Unknown Unicast to Hardware Proxy.

D. Configure IP Data-Plane Learning to No.

Answer: D

Explanation:

In a Cisco ACI environment with multiple silent hosts that are often relocated between leaf switches, configuring IP Data-Plane Learning to 'No' will minimize the relocation impact and make the ACI fabric relearn the new location of the host faster. Disabling IP Data-Plane Learning prevents the fabric from learning IP addresses from the data plane, which can speed up the process of relearning host locations when they move3.

Question: 137

Refer to the exhibit.

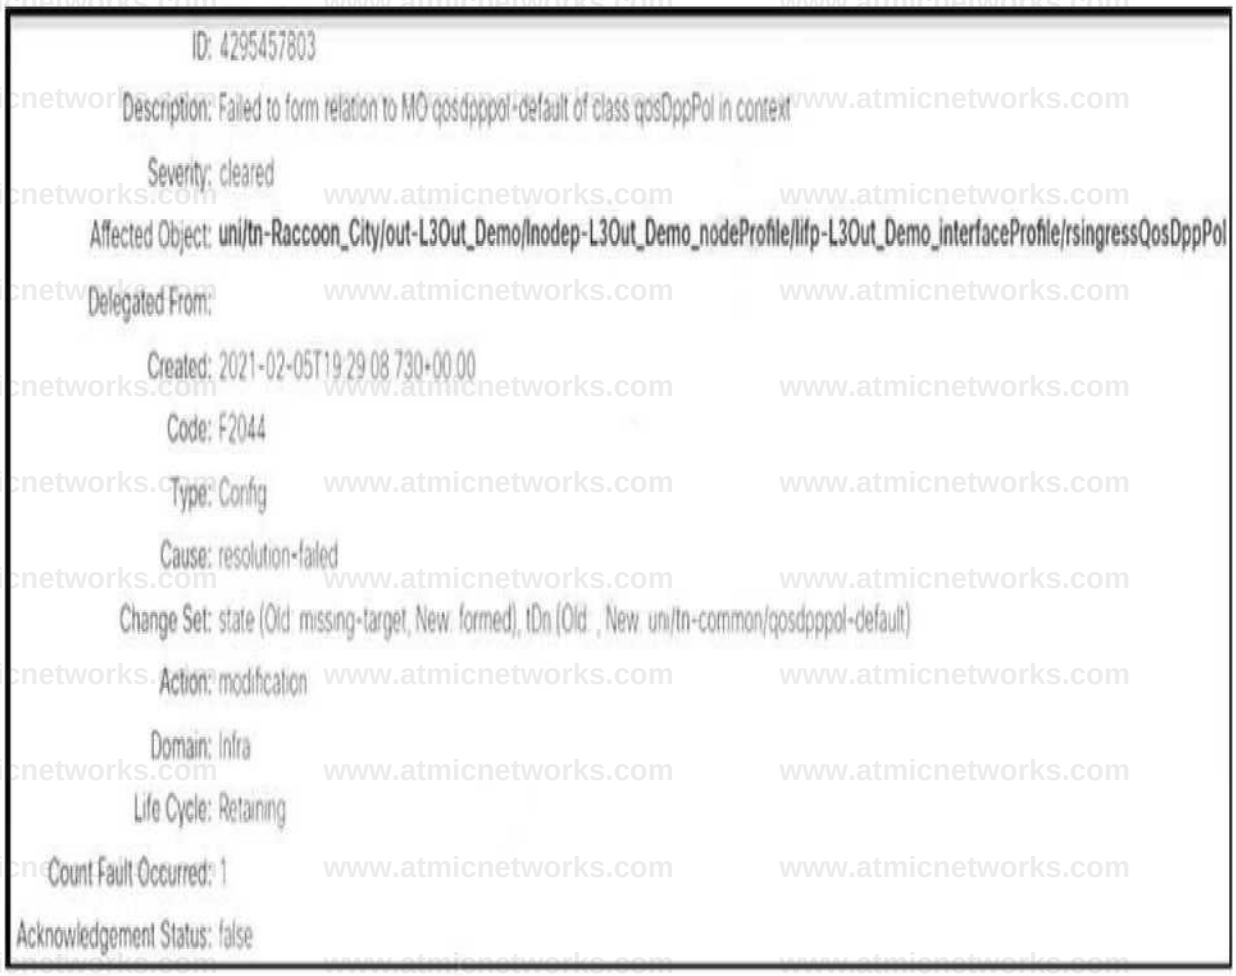

Refer to the exhibit. An engineer configures an L3Out but receives the error presented. Which action clears the fault?

A. Acknowledge the QoS-related error.

B. Associate a custom QoS class.

C. Create a custom QoS policy.

D. Set the QoS policy to Level 3.

Answer: C

Explanation:

When configuring an L3Out in Cisco ACI and encountering a QoS-related error, creating a custom QoS policy is often necessary to clear the fault. This involves defining a QoS policy that meets the specific requirements of the L3Out configuration and applying it to the relevant interfaces or globally within

the fabric. The custom QoS policy should be designed to prioritize traffic appropriately and ensure that the necessary QoS settings are in place for the L3Out connections1.

Reference:

Cisco ACI L3Out Configuration Examples1

Question: 138

A customer must upgrade the Cisco ACI fabric to use a feature from the new code release. However, there is no direct path from the current release to the desired one. Based on the Cisco APIC Upgrade/Downgrade Support Matrix, the administrator must go through one intermediate release.

Which set of steps must be taken to upgrade the fabric to the new release?

A. Upgrade the APICs to an interim release.
Upgrade the switches to an interim release.
Upgrade the APICs to the targeted release.
Upgrade the leaf and spine switches to the targeted release.

B. Upgrade the APICs to an interim release and then switches to an interim release.
When all switches are operational, upgrade leaf switches to the targeted release.
Upgrade the spine switches to the targeted release.
Upgrade the APICs to the targeted release.

C. Upgrade the APICs to an interim release.
Upgrade the leaf switches directly to the targeted release.
Upgrade the spine switches directly to the targeted release.
Upgrade the APICs to the targeted release.

D. Upgrade the APICs directly to the targeted release.
Upgrade the switches to an interim release.
When all switches are operational, upgrade the leaf switches to the targeted release.
Upgrade the spine switches to the targeted release.

Answer: A

Explanation:

When upgrading the Cisco ACI fabric through an intermediate release, the following steps should be taken:

Upgrade the APICs to an interim release: This is the first step to ensure that the APICs are running a stable version that is compatible with both the current and the targeted release1.

Upgrade the switches to an interim release: After the APICs are upgraded, the next step is to upgrade the leaf and spine switches to the same interim release. This ensures that the entire fabric is running on a compatible version before moving to the targeted release1.

Upgrade the APICs to the targeted release: Once the fabric is stable on the interim release, the APICs can be upgraded to the targeted release. It's important to verify that the APICs are fully operational before proceeding to upgrade the switches1.

Upgrade the leaf and spine switches to the targeted release: The final step is to upgrade the leaf and spine switches to the targeted release. This should be done after the APICs have been successfully upgraded and are stable on the targeted release1.

By following these steps, the customer can successfully upgrade the Cisco ACI fabric to the new code release that includes the desired feature.

Question: 139

Refer to the exhibit.

```
leaf-102# show interface brief
!snip
--------------------------------------------------------------------------------
Port-channel VLAN    Type Mode    Status Reason                    Speed     Protocol
Interface
--------------------------------------------------------------------------------
Po3          46      eth  trunk   down   mac-pinning              inherit(D lacp
Po11         --      eth  fabric  up     none                     10G(D)    none
Po12         0       eth  trunk   down   mcp-loop-err-disable     inherit(D none
```

Refer to the exhibit. Which two configuration steps are completed before this output is generated? (Choose two.)

A. MCP policy for the interface policy group for Port-channel 12 is enabled.

B. MCP Instance Policy default in the global access policies is enabled.

C. Error Disabled Recovery Policy for Loop Indication by MCP is set to True.

D. BPDU Guard is enabled for the interface policy group for Port-channel 12.

E. Spanning Tree Policy Region STP_4CAF232E48FF20 is added to the spanning-tree policy of the switch.

Answer: AB

Explanation:

Question: 140

The customer is looking for redundant interconnection of the existing network to the new ACI fabric. Unicast and multicast traffic must be routed between the two networks. Which L3Out implementation meets these requirements?

A.

B.

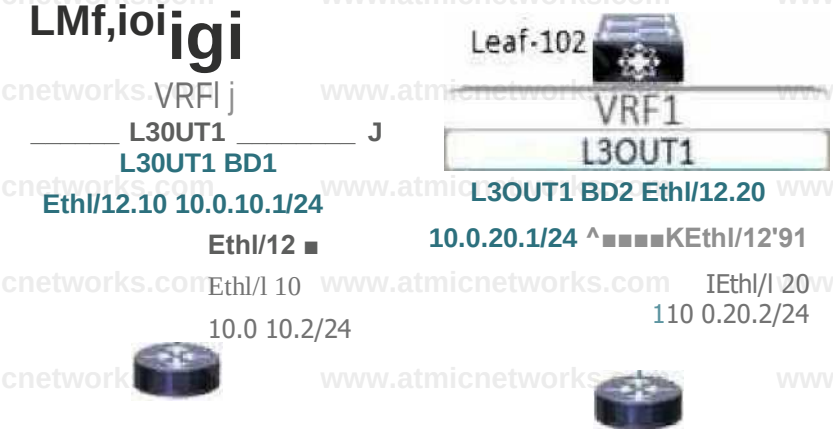

C.

D.

A. Option A

B. Option B

C. Option C

D. Option D

Answer: A

Explanation:

For a redundant interconnection that routes both unicast and multicast traffic, the L3Out (Layer 3 Outside) implementation in Cisco ACI should include the following:

Redundant Connections: Ensure there are at least two connections from the ACI fabric to the external network for redundancy. These connections should be to different leaf switches if possible.

Routing Protocols: Configure a routing protocol that supports both unicast and multicast routing. OSPF or EIGRP can be used for unicast, and PIM (Protocol Independent Multicast) should be configured for multicast.

Route Redistribution: Set up route redistribution between the ACI fabric and the external network to ensure that routes are shared across both networks.

Multicast Configuration: Enable multicast routing within the ACI fabric and configure the necessary multicast policies.

Contract and Policy Configuration: Define contracts and policies that allow the required unicast and multicast traffic to flow between the ACI fabric and the external network.

Reference:

Cisco documentation on configuring L3Outs in ACI

Cisco white paper on ACI Multicast Routing

Question: 141

A network engineer configures the Cisco ACI fabric to connect to vCenter with these requirements:

Port groups must be automatically created on the distributed virtual switch.

Port groups must use the VLAN allocation in the range between 20-30.

The deployment must optimize the CAM space on the leaf switches.

Which set of actions meets these criteria?

A. Create a dynamic VLAN pool with the VLAN range of 20-30.

Create a VMM domain and associate it with the VLAN pool.

Create the EPG and associate the domain.

Set the deployment immediacy to On Demand.

B. Create a dynamic VLAN pool with the VLAN range of 20-30.

Create a physical domain and associate it with the VLAN pool.

Create the EPG and associate the domain.

Set the deployment immediacy to On Demand.

C. Create a static VLAN pool with the VLAN range of 20-30.

Create a physical domain and associate it with the VLAN pool.

Create the EPG and associate the domain.

Set the deployment immediacy to Immediate.

D. Create a static VLAN pool with the VLAN range of 20-30.

Create a VMM domain and associate it with the VLAN pool.

Create the EPG and associate the domain.

Set the deployment immediacy to Immediate.

Answer: A

Explanation:

To meet the criteria for configuring the Cisco ACI fabric to connect to vCenter, the following actions should be taken:

Create a dynamic VLAN pool with the VLAN range of 20-30: This step involves creating a VLAN pool that dynamically assigns VLANs within the specified range to the port groups as they are created1.

Create a VMM domain and associate it with the VLAN pool: The VMM domain is a representation of the VMware vSphere Distributed Switch (VDS) in ACI, and associating it with the VLAN pool allows for the automatic creation of port groups on the VDS1.

Create the EPG and associate the domain: An Endpoint Group (EPG) is a logical entity to which workloads are attached. Associating the EPG with the VMM domain ensures that the port groups are automatically created in vCenter when the EPG is associated with a VMM domain1.

Set the deployment immediacy to On Demand: Setting the deployment immediacy to 'On Demand' optimizes the CAM space on the leaf switches by deploying policies to leaf nodes only when endpoints assigned to this EPG are connected, rather than immediately upon configuration12.

By following these steps, the network engineer can ensure that the Cisco ACI fabric is correctly configured to connect to vCenter, with port groups being automatically created on the distributed

virtual switch, using the VLAN allocation in the range between 20-30, and optimizing the CAM space on the leaf switches.

Question: 142

Refer to the exhibit.

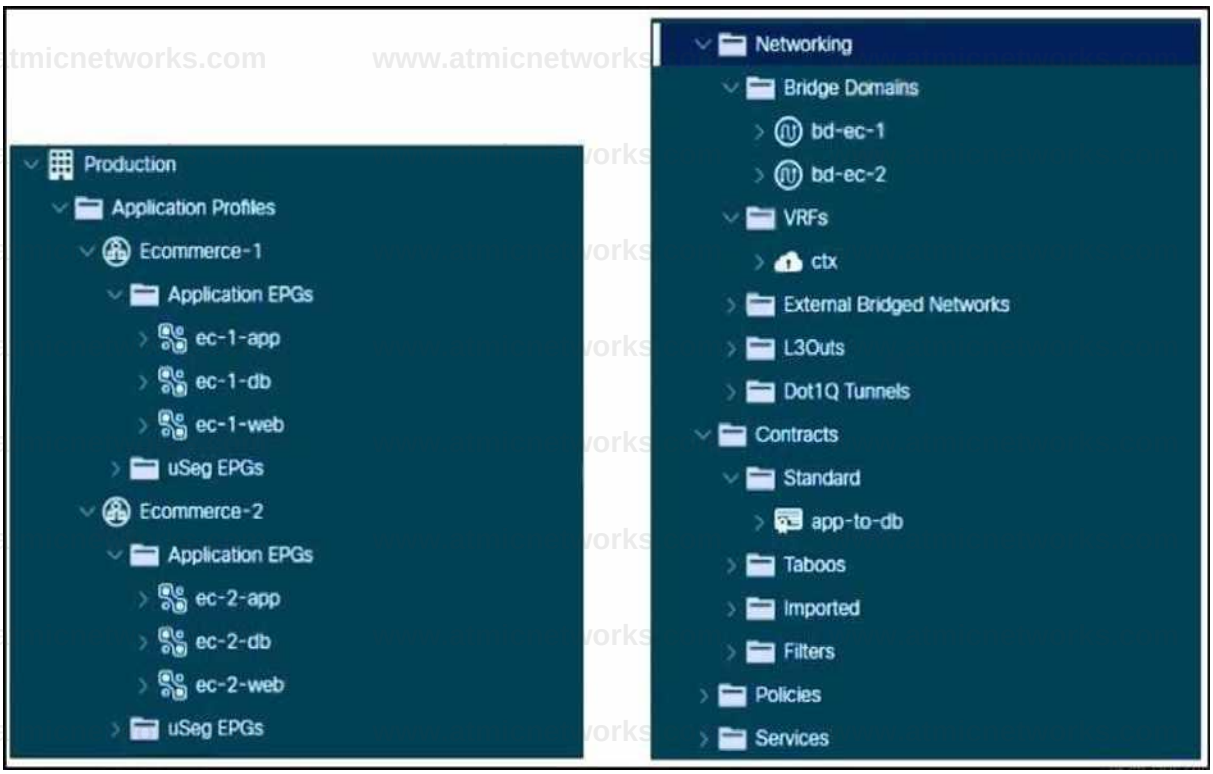

Refer to the exhibit. A Cisco ACI environment hosts two e-commerce applications. The default contract from a common tenant between different application tiers is used, and the applications work as expected. The customer wants to move to more specific contracts to prevent unwanted traffic between EPGs. A network administrator creates the app-to-db contract to meet this objective for the application and database tiers. The application EPGs must communicate only with their respective database EPGs. How should this contract be configured to meet this requirement?

A. Set the app-to-db scope to Global.

B. Set the app-to-db scope to Application Profile.

C. Implement the app-to-db scope as VRF.

D. Implement the app-to-db as a Taboo contract.

Answer: B

Explanation:

To ensure that application EPGs communicate only with their respective database EPGs, the app-to- db contract should be configured with a scope set to the Application Profile. This scope ensures that the contract is applied only within the specific application profile, allowing for granular control over the communication between the application and database tiers. By setting the scope to Application Profile, the contract will not apply to other application profiles, thus preventing unwanted traffic between EPGs of different applications1.

Reference:

Cisco ACI Contract Guide White Paper1

Question: 143

Refer to the exhibit.

Refer to the exhibit. An engineer configured subnets on the external EPG called L3OUT_CORE. The external endpoints in the 10.1.0.0/24 subnet can reach internal endpoints, but the external endpoints in the 172.16.1.0/24 subnet are unreachable. Which set of actions enables the connectivity?

A.

Delete both external EPG subnets.

Create the 0.0.0.0/1 subnet.

B.

Delete the external EPG subnet 0.0.0.128/1.

Create the 128.0.0.0/1 subnet.

C.

Delete both external EPG subnets.

Create the 0.0.0.0/0 subnet.

D.

Delete the external EPG subnet 0.0.0.0/0.

Create the 0.0.0.0/128 subnet.

Answer: C

Explanation:

To enable connectivity for the external endpoints in the 172.16.1.0/24 subnet, the following actions should be taken:

Delete both external EPG subnets: This action removes any specific subnet configurations that might be causing routing issues or conflicts within the external EPG1.

Create the 0.0.0.0/0 subnet: By configuring a 0.0.0.0/0 subnet, you are effectively creating a default route that allows all traffic to be routed, ensuring that endpoints in any subnet, including the 172.16.1.0/24 subnet, can reach internal endpoints1.

This configuration change addresses the issue where traffic from the 172.16.1.0/24 subnet is not reaching the internal endpoints, likely due to a lack of a default route or misconfigured subnets within the external EPG

Question: 144

An engineer deploys a two-pod Cisco ACI Multi-Pod environment. Why should no more than two Cisco APIC controllers be deployed in the same pod?

A. to enable equal capacity to scale in each pod

B. to avoid losing all replicas of a shard if a pod fails

C. to avoid hair-pinning traffic that is destined for the primary APIC controller between pods

D. to ensure that all nodes in all pods have local access to a controller

Answer: B

Explanation:

In a two-pod Cisco ACI Multi-Pod environment, it is recommended to deploy no more than two Cisco APIC controllers in the same pod to avoid losing all replicas of a shard if a pod fails. The APIC cluster is a distributed system where each APIC contains a replica of the configuration database, referred to as a shard. If all replicas of a shard were located in the same pod and that pod were to fail, it would result in the loss of the shard and potentially impact the entire fabric's operation. By distributing the APIC controllers across pods, the risk of losing all replicas due to a single point of failure is mitigated, enhancing the overall resiliency of the system1.

Reference:

ACI Multi-Pod White Paper - Cisco1

Question: 145

Refer to the exhibit.

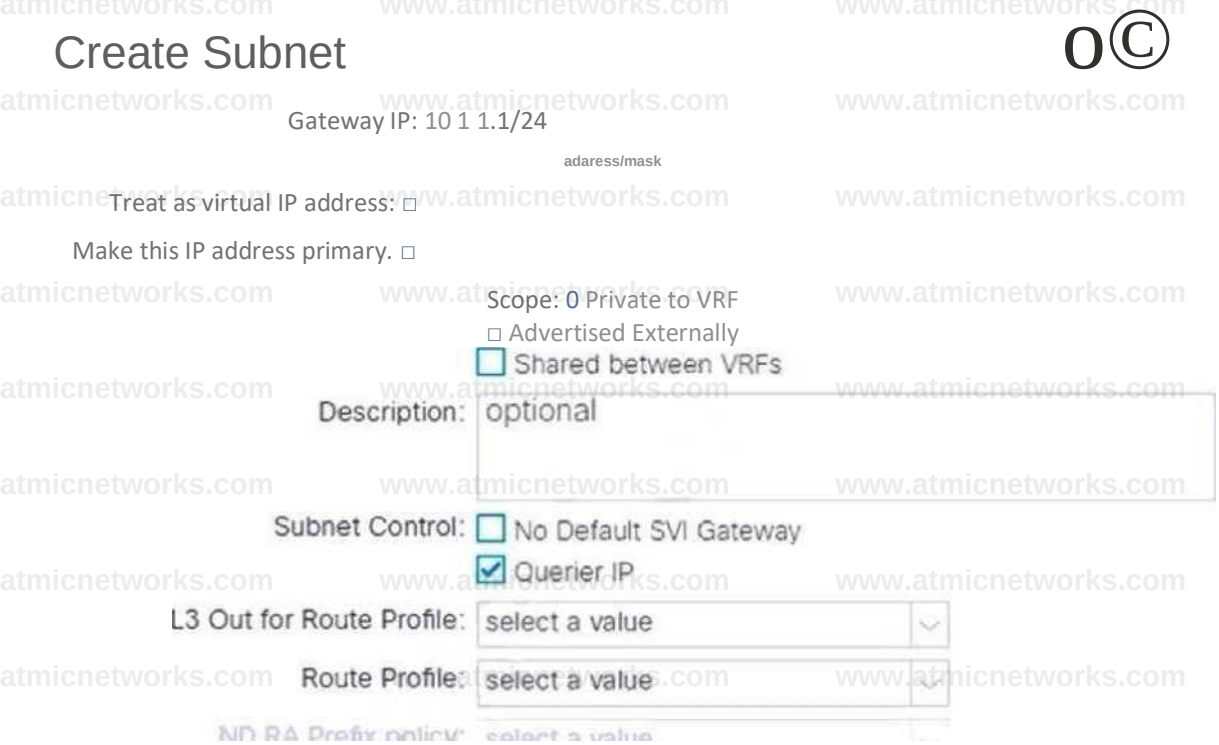

Refer to the exhibit. An engineer configures communication between the EPGs in different tenants. Which action should be taken to create the subnet?

A. Change Scope to Shared between VRFs.

B. Leave Scope set to Private to VRF.

C. Add the L3Out for Route Profile value.

D. Change Scope to Advertised Externally.

Answer: A

Explanation:

To configure communication between the EPGs in different tenants, the subnet scope should be set to “Shared between VRFs”. This allows the subnet to be shared across different VRFs, enabling communication between EPGs that are in different tenants but within the same VRF. By marking the subnet as shared, it becomes visible to other VRFs, which is necessary for inter-tenant

communication1.

Reference:

Learning ACI - Part 12: Inter-VRF and Inter-Tenant Communication

Question: 146

Refer to the exhibit.

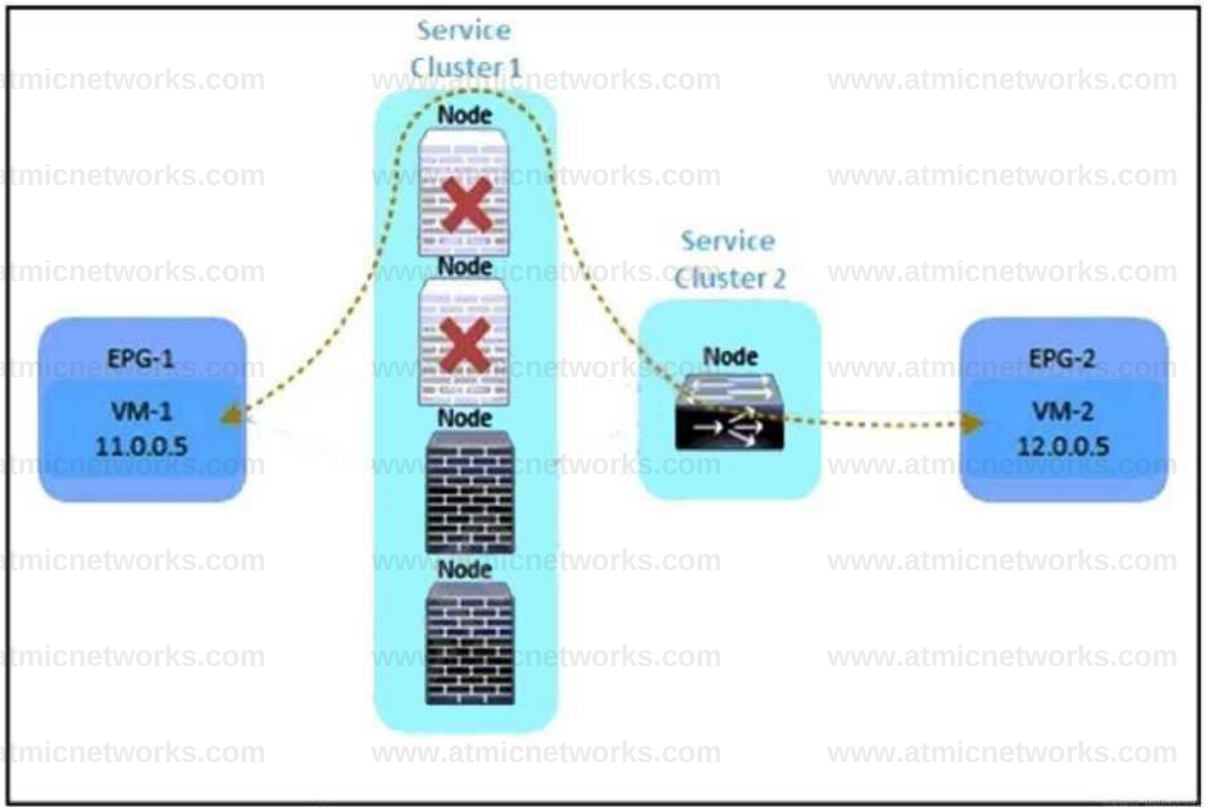

Refer to the exhibit. An engineer must divert the traffic between VM-1 and VM-2 by using a MultiNode service graph. The solution should prevent an insufficient number of available Layer 4 to Layer 7 devices in the first cluster. Which configuration set accomplishes this goal?

A. PBR node tracking

tracking threshold with action bypass

symmetric PBR

resilient hashing

B. PBR node tracking

tracking threshold with action permit

unidirectional PBR

resilient hashing

C. PBR node tracking

tracking threshold with action permit

symmetric PBR

resilient hashing

D. PBR node tracking

tracking threshold with action deny

symmetric PBR

unidirectional PBR

Answer: A

Explanation:

To divert traffic between VM-1 and VM-2 using a Multi-Node service graph while preventing an insufficient number of available Layer 4 to Layer 7 devices in the first cluster, the following configuration set should be used:

PBR node tracking: This feature monitors the health of the nodes (Layer 4 to Layer 7 devices) in the service graph. If a node becomes unavailable, the system can take action based on the tracking threshold1.

Tracking threshold with action bypass: When the number of healthy nodes falls below the tracking threshold, the action 'bypass' ensures that traffic is not sent to the unhealthy nodes, thus preventing service disruption1.

Symmetric PBR: This ensures that return traffic follows the same path as the original traffic, maintaining session consistency and avoiding asymmetric routing issues1.

Resilient hashing: This hashing mechanism provides a consistent and optimal distribution of traffic across the available Layer 4 to Layer 7 devices, even when the number of nodes changes due to a failure or maintenance1.

Reference:

Cisco APIC Layer 4 to Layer 7 Services Deployment Guide, Release 5.2(x)

Question: 147

An engineer must add a group of 70 bare-metal ESXi servers to the Cisco ACI fabric, which is integrated with vCenter. These configuration steps are complete:

The configured pool of ESXi hosts is configured with an Attachable Access Entity Profile (AAEP) called AEP_VMM.

The new group uses the AAEP called AEP_BAREMETAL.

Which action extends functional VMM integration to the new nodes?

A. Update AAEP to AEP_VMM on all policy groups that are used toward bare-metal servers.

B. Create a new AAEP container object for policy groups for AEP_VMM.

C. Implement a separate VMM domain for the bare-metal servers by using AEP_VMM.

D. Add the VMM domain under the AEP_BAREMETAL AAEP object.

Answer: C

Explanation:

To extend functional VMM integration to the new group of 70 bare-metal ESXi servers, a separate VMM domain should be implemented using the AEP_VMM. This allows for the management and automation of network policies for the bare-metal servers in a manner similar to the virtualized environment managed by vCenter1.

Question: 148

Which two protocols are used for fabric discovery in ACI? (Choose two.)

A. LLDP

B. OSPF

C. CDP

D. DHCP

E. ISIS

Answer: A, E

Explanation:

The two protocols used for fabric discovery in Cisco ACI are LLDP (Link Layer Discovery Protocol) and ISIS (Intermediate System to Intermediate System). LLDP is used for neighbor discovery on the data link layer, while ISIS is a routing protocol used for reachability between the TEP IPs (VTEPs) within the ACI fabric23.

Reference: https://www.dclessons.com/aci-fabric-discovery

Question: 149

What is the purpose of the Overlay Multicast TEP in a Cisco ACI Multi-Site deployment?

A. to source and receive unicast VXLAN data plane traffic

B. to establish MP-BGP EVPN adjacencies with the spine nodes in remote sites

C. to encapsulate multicast traffic in a common multicast group

D. to perform head-end replication for BUM traffic

Answer: D

Explanation:

The purpose of the Overlay Multicast TEP (O-MTEP) in a Cisco ACI Multi-Site deployment is to perform head-end replication for BUM (Broadcast, Unknown unicast, Multicast) traffic. This common anycast address is shared by all the spine nodes in the same site and is used to replicate BUM traffic across the sites4.

Question: 150

Refer to the exhibit.

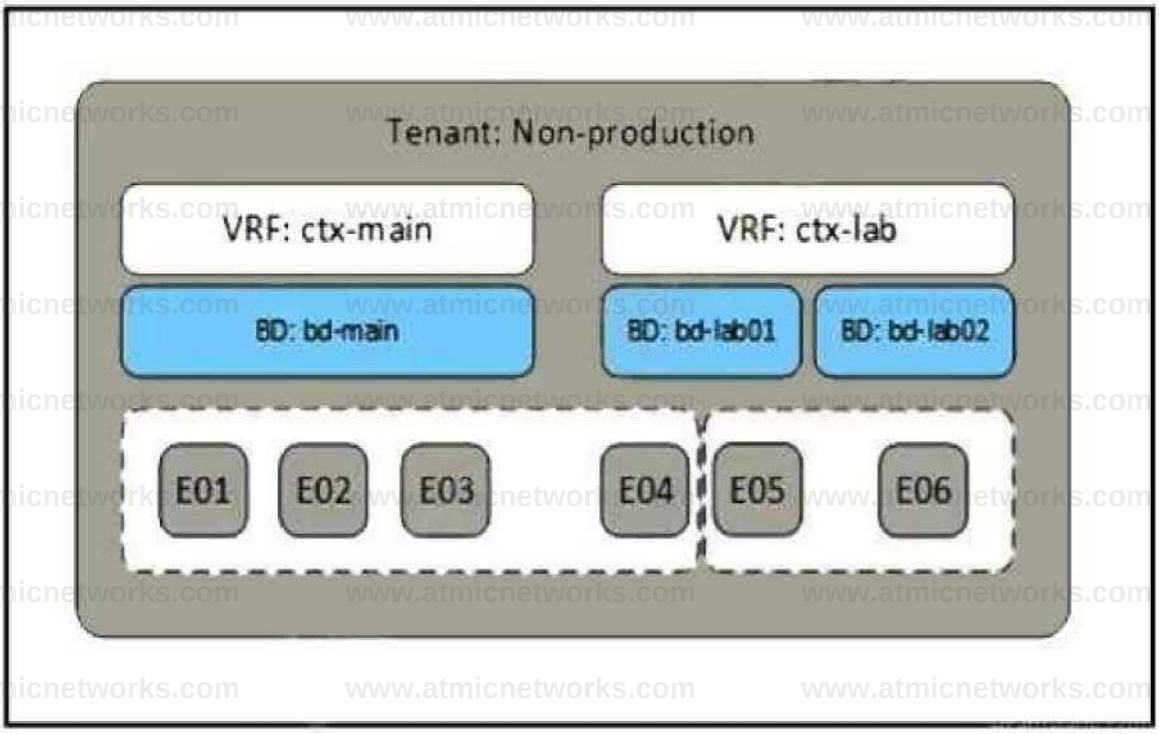

Refer to the exhibit. A network engineer must complete the Cisco ACI implementation based on the logical system design created by the systems architect. Which Cisco ACI object is required where the dotted line indicates to complete the task?

A. contract

B. application profile

C. context

D. attachable Access Entity Profile

Answer: A

Explanation:

In Cisco ACI, a contract is used to define the communication policy between EPGs (Endpoint Groups). It specifies which types of traffic are allowed to pass between EPGs and can include filters for protocols, ports, and other attributes. In the context of the logical system design, the contract would be the object that completes the communication requirements as indicated by the dotted line in the exhibit12.

Reference:

Cisco Application Centric Infrastructure (ACI) Design Guide1

Question: 151

When a pre-provision immediacy is used, when is the policy downloaded to the Cisco ACI leaf switch?

A. The policy is downloaded and programmed in the hardware policy CAM when the change is implemented on the Cisco APIC.

B. The policy is programmed in the hardware policy CAM when the policy is downloaded in the leaf software.

C. The policy is programmed in the hardware policy CAM when the first packet is received through the data path.

D. The policy is downloaded to the associated leaf switch software when the ESXi host is attached to a DVS.

Answer: A

Explanation:

When a pre-provision immediacy is used, the policy is downloaded to the Cisco ACI leaf switch and programmed into the hardware policy Content Addressable Memory (CAM) as soon as the change is implemented on the Cisco Application Policy Infrastructure Controller (APIC). This means that the policy is ready on the leaf switch before any endpoints are actually connected.

Question: 152

As part of a migration, legacy non-ACI switches must be connected to the Cisco ACI fabric. All nonACI switches run per-VLAN RSTP. After the non-ACI switches are connected to Cisco ACI, the STP convergence caused a microloop and significant CPU spike on all switches. Which configuration on the interfaces of the external switches that face the Cisco ACI fabric resolves the problem?

A. BPDU guard

B. aggressive STP timers

C. BPDU filtering

D. STP type link shared

Answer: C

Explanation:

To resolve the issue of STP convergence causing a microloop and significant CPU spike on all switches after connecting legacy non-ACI switches to the Cisco ACI fabric, BPDU filtering should be configured on the interfaces of the external switches that face the Cisco ACI fabric. BPDU filtering prevents Bridge Protocol Data Units (BPDUs) from being sent or received on those interfaces, which stops the STP process and eliminates the possibility of microloops.

Question: 153

Which two IP address types are available for transport over the ISN when they are configured from Cisco ACI Multi-Site Orchestrator? (Choose two.)

A. Management IP of APICs

B. Management IP of the MSO Node

C. Anycast Overlay Multicast TEP

D. MP-BGP EVPN Router-ID

E. Common Pervasive Gateway

Answer: C, D

Explanation:

The two IP address types that are available for transport over the Inter-Site Network (ISN) when they are configured from Cisco ACI Multi-Site Orchestrator (MSO) are the Anycast Overlay Multicast Tunnel Endpoint (TEP) and the MP-BGP EVPN Router-ID. These IP addresses are used for inter-site communication and

routing in a Multi-Site deployment.

Question: 154

A network engineer must integrate VMware vCenter cluster with Cisco ACI. The requirement is for the management traffic of the hypervisors and VM controllers to use the virtual switch associated with the Cisco Application Policy. The EPG called "Vmware-MGMT" with VLAN 300 has been created for this purpose. Which set of steps must be taken to complete the configuration?

A. Add VLAN 300 with static allocation to the VLAN POOL that is used for VMM integration.

Attach the VMM domain to the target EPG with resolution preprovision, mode static, untagged access VLAN, and Port-Encap 300.

B. Associate the target EPG with the VMM domain with default settings.

Enable Infrastructure VLAN on AAEP used toward VMware hypervisors.

C. Enable Infrastructure VLAN on AAEP used toward VMware hypervisors.

Associate the target EPG with the VMM domain with default settings.

D. Enable Infrastructure VLAN on AAEP used toward VMware hypervisors.

Create a static binding in the target EPG toward VMware hypervisors with VLAN 300, untagged access VLAN, and Untagged 802.1P mode.

Answer: D

Explanation:

To integrate the VMware vCenter cluster with Cisco ACI and ensure that the management traffic of the hypervisors and VM controllers uses the virtual switch associated with the Cisco Application Policy, the

following steps must be taken:

Enable Infrastructure VLAN on AAEP used toward VMware hypervisors: This step involves enabling the infrastructure VLAN on the Attachable Access Entity Profile (AAEP) that is used for the VMware hypervisors. This VLAN is used for carrying infrastructure traffic such as management, vMotion, and fault tolerance.

Create a static binding in the target EPG toward VMware hypervisors with VLAN 300: This step involves creating a static binding in the "Vmware-MGMT" EPG for the VMware hypervisors with VLAN 300, setting it as an untagged access VLAN, and using Untagged 802.1P mode. This ensures that the management traffic is correctly tagged and associated with the appropriate EPG.

Reference:

Cisco ACI Fabric Hardware Installation Guide

Cisco ACI Troubleshooting Guide

Cisco ACI Multi-Site Architecture White Paper

Cisco ACI Virtual Machine Manager (VMM) Integration White Paper

Question: 155

Refer to the exhibit.

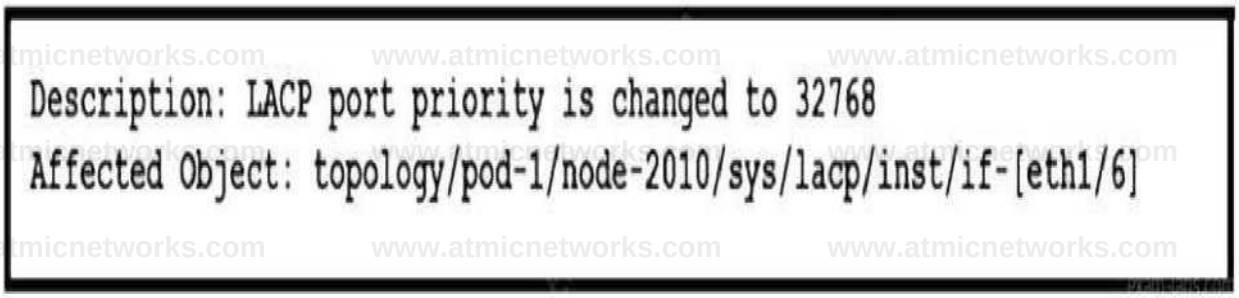

Refer to the exhibit. A Cisco ACI fabric displays this fault. Which set of actions modifies the event to be displayed as a warning in the future?

A. Navigate to the ACI Events tab.

Create a new record.

B. Navigate to the ACI Fault tab.

Create a new record.

C. Navigate to the ACI Events tab.

Change the severity level.

D. Navigate to the ACI Fault tab.

Change the severity level.

Answer: D

Explanation:

To modify the event to be displayed as a warning in the future, you need to navigate to the ACI Fault tab and change the severity level of the fault. This involves selecting the monitoring object that corresponds to the class of the Managed Object (MO) that generated the fault, then choosing the fault code you want to modify and setting an initial severity value1.

Reference:

Cisco APIC Faults, Events, and System Messages Management Guide2

How ACI Faults Are Generated and How to Selectively Prevent ... - Cisco1

Question: 156

A bridge domain for a new endpoint group in the Cisco ACI fabric must meet these requirements:

The bridge domain must function as the default gateway for the subnet so that routing remains within the Cisco ACI fabric.

ARP requests must be managed via Layer 3 unicast packets or be dropped to reduce excessive broadcast traffic.

The impact of misconfigured virtual machines must be kept to a minimum by preventing IP addresses outside of the configured subnet from being routed.

Which set of actions must be taken?

A. Disable ARP Flooding.

Enable Limit IP Learning to Subnet.

Enable Unicast Routing on the bridge domain and configure a subnet.

B. Enable Limit IP Learning to Subnet.

Enable Unicast Routing on the bridge domain and configure a subnet.

Set Multi-Destination Flooding to Flood in BD.

C. Set Endpoint Retention Policy to default.

Enable ARP Flooding.

Enable Unicast Routing on the bridge domain and configure a subnet.

D. Enable Unicast Routing on the bridge domain and configure a subnet.

Set L2 Unknown Unicast to Flood.

Disable Endpoint Retention Policy.

Answer: A

Explanation:

To meet the requirements for the new endpoint group's bridge domain in the Cisco ACI fabric, the following actions must be taken:

Disable ARP Flooding: This action prevents ARP requests from being broadcast across the entire bridge domain, which reduces excessive broadcast traffic.

Enable Limit IP Learning to Subnet: This setting restricts IP address learning to only those IPs within the configured subnet, minimizing the impact of misconfigured virtual machines.

Enable Unicast Routing on the bridge domain and configure a subnet: Enabling unicast routing allows the bridge domain to function as the default gateway for the subnet, ensuring that routing remains within the Cisco ACI fabric.

Question: 157

An engineer configures a one-armed policy-based redirect service Insertion for an unmanaged firewall. The engineer configures these Cisco ACI objects:

a contract named All_Traffic_Allowed

a Layer 4 to Layer 7 device named FW-Device

a policy-based redirect policy named FW-1Arm-Policy-Based RedirectPolicy

Which configuration set redirects the traffic to the firewall?

A.

Configure a policy-based redirect subject.

Associate the policy-based redirect subject with All_Traffic_Allowed.

B.

Configure a firewall bridge domain.

Associate the bridge domain with FW-Device.

C.

Configure a device interface policy.

Associate the device interface policy with FW-Device.

D.

Configure a service graph.

Associate the service graph with All_Traffic_Allowed.

Answer: D

Explanation:

To redirect traffic to the firewall using a one-armed policy-based redirect service insertion, the following configuration set should be used:

Configure a service graph: This involves defining the service graph template that represents the logical flow of traffic through network services.

Associate the service graph with All_Traffic_Allowed: By associating the service graph with the contract named All_Traffic_Allowed, traffic that matches the contract will be redirected through the service graph to the

firewall device.

Reference:

Cisco ACI Fabric Administration Guide, Release 4.2(1)

Cisco ACI Virtualization Guide, Release 4.2(1)

Question: 158

Refer to exhibit.

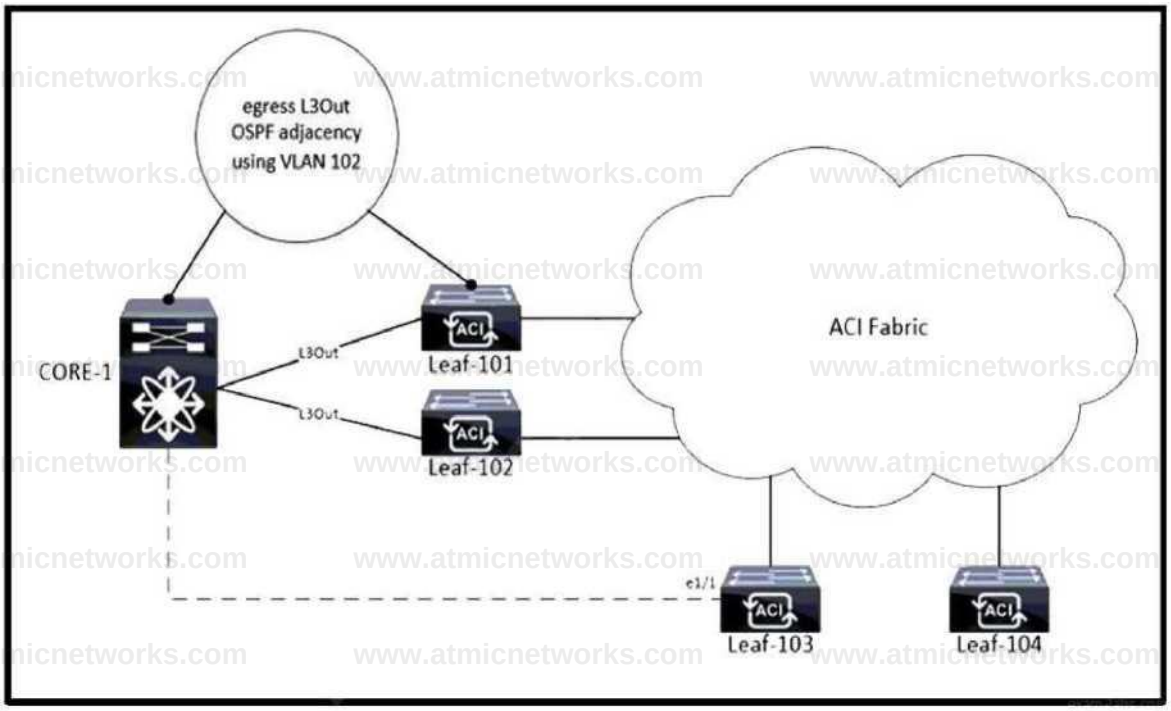

Refer to the exhibit. The Cisco ACI fabric has an egress L3Out from Leaf-101 and Leaf-102 to CORE-1. VLAN 102 is used to form the OSPF adjacency. The workloads must be migrated into EPG-101, and the static port binding is configured to Leaf-103 e1/1 with encap VLAN 101. An engineer completes the port binding and receives an MCP fault. Which action clears the fault?

A. Use VLAN 101 for OSPF adjacency on the egress L3Out.

B. Use VLAN 102 as the encap VLAN on the EPG-101 static port binding.

C. Add VLAN 102 to the VLAN pool that is used by the static port binding.

D. Prune VLAN 101 from the VLAN pool that is used by the egress L3Out.

Answer: B

Explanation:

To clear the MCP (Mis-Cabling Protocol) fault in a Cisco ACI fabric when an egress L3Out is configured from Leaf-101 and Leaf-102 to CORE-1, and VLAN 102 is used for OSPF adjacency, the encapsulation VLAN on the EPG-101 static port binding should match the VLAN used for OSPF adjacency. This means that VLAN 102 should be used as the encapsulation VLAN for the static port binding on Leaf- 103 e1/1. Using the same VLAN for both OSPF adjacency and the static port binding ensures consistency and prevents misconfiguration that could lead to MCP faults1.

Reference:

ACI Fabric L3Out White Paper - Cisco1

Question: 159

Refer to the exhibit.

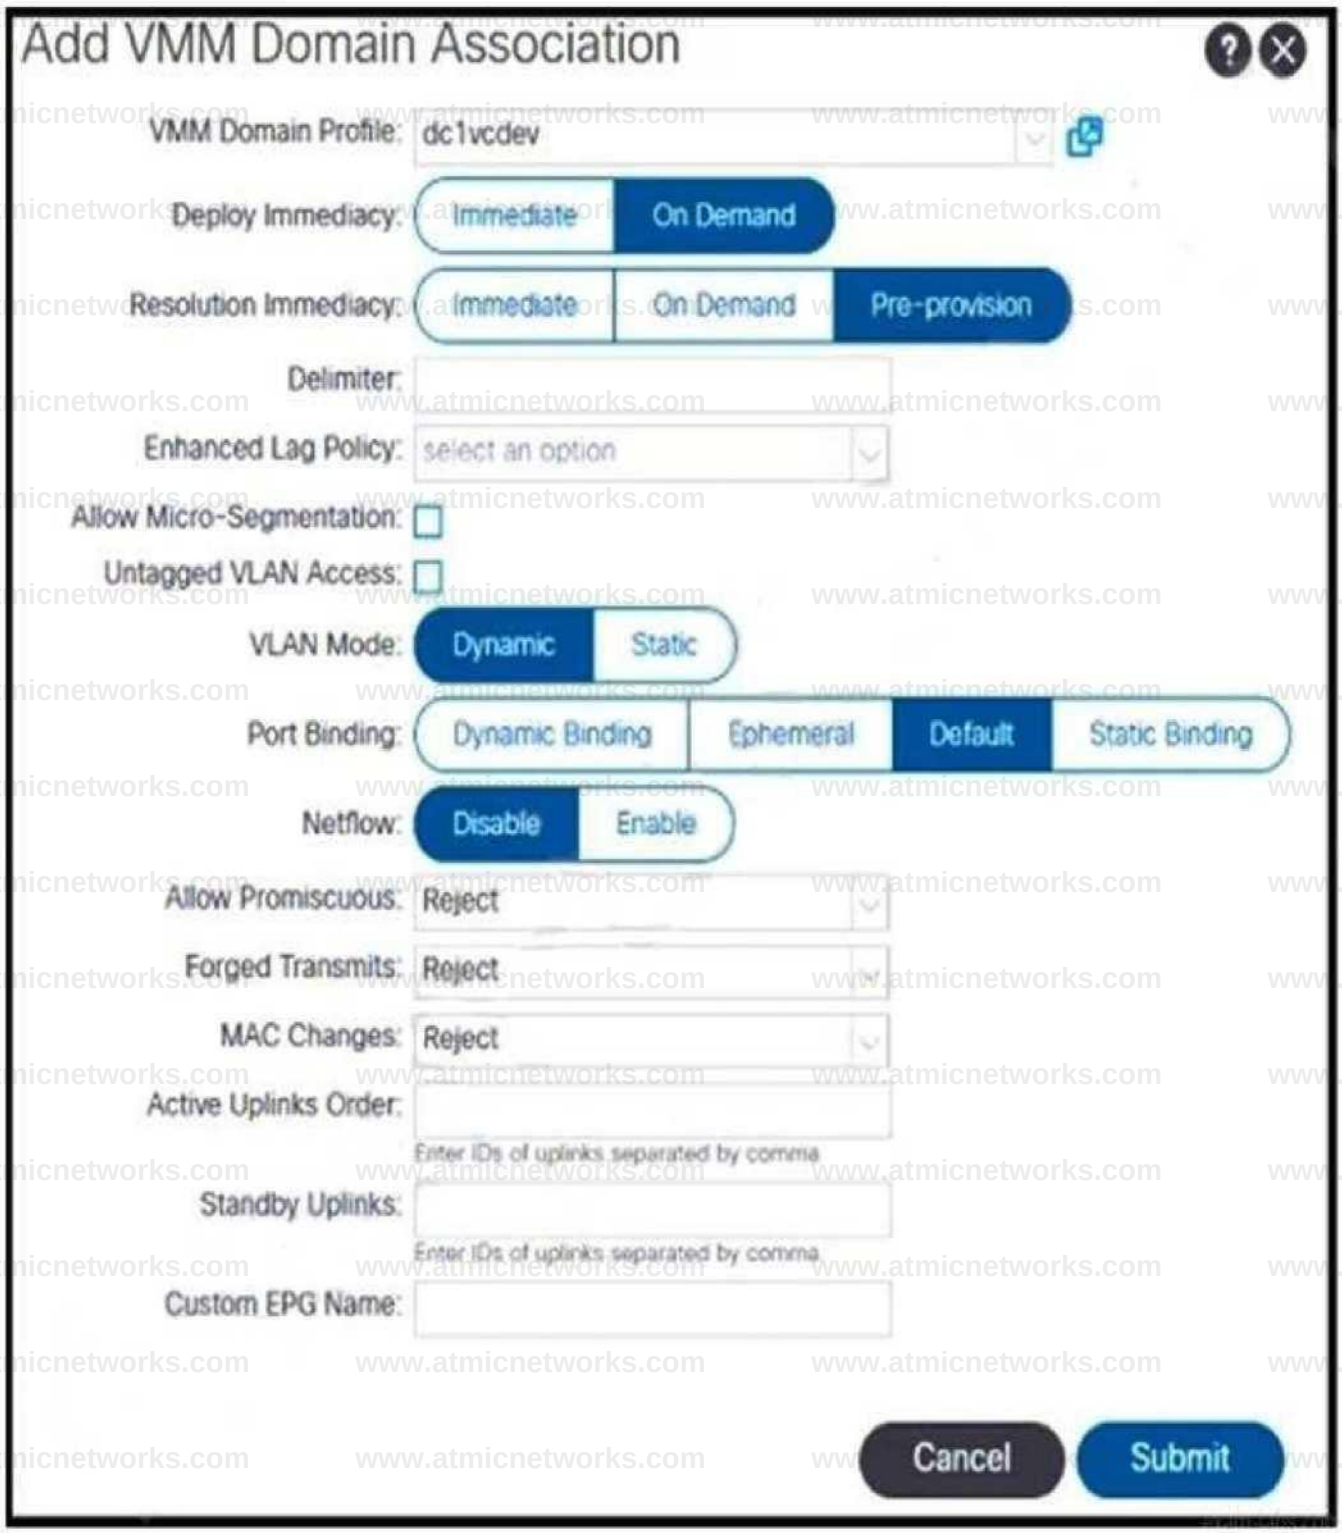

Refer to the exhibit. The EPG-100 must be extended to the vCenter as a port group with a tagged VLAN ID of 100. Which set of actions accomplishes this goal?

A.

Define a static VLAN range (from 100-200) under a VLAN pool that is associated with the dc1vcdev domain.

Associate the dc1vcdev domain with EPG and select these settings:

Untagged VLAN Access: unselected

VLAN Mode: Static with Encap: 100

B.

Define a static VLAN range (from 100-200) under a VLAN pool that is associated with the dc1vcdev domain.

Associate the dc2vcdev domain with EPG and select these settings:

Untagged VLAN Access: selected

VLAN Mode: Static with Encap: 100

C.

Define a dynamic VLAN range (from 100-200) under a VLAN pool that is associated with the del vdev domain.

Associate the dc1vcdev domain with EPG and select these settings:

Untagged VLAN Access: unselected

VLAN Mode: Static with Encap: 100

D.

Define a dynamic VLAN range (from 100-200) under a VLAN pool that is associated with the dc1vdev domain.

Associate the dc2vcdev domain with EPG and select these settings:

Untagged VLAN Access: selected

VLAN Mode: Static with Encap: 100

Answer: A

Explanation:

To extend the EPG-100 to the vCenter as a port group with a tagged VLAN ID of 100, the following actions should be taken:

Define a static VLAN range (from 100-200) under a VLAN pool: This step involves creating a VLAN pool that

includes the desired VLAN range. The VLAN pool is then associated with the VMM domain that corresponds to the vCenter integration1.

Associate the dc1vcdev domain with EPG: The VMM domain named 'dc1vcdev' should be associated with the EPG-100. This association allows for the automatic creation of port groups in vCenter when the EPG is associated with a VMM domain1.

Select the appropriate settings for the domain association:

Untagged VLAN Access: This should be unselected because the VLAN ID needs to be tagged.

VLAN Mode: Set to 'Static' to specify the encapsulation VLAN ID.

Encap: Set to '100' to define the specific VLAN ID that should be tagged for the port group1.

Reference:

Cisco APIC Layer 2 Networking Configuration Guide1

Cisco APIC Layer 2 Networking Configuration Guide1

Question: 160

A customer must deploy three Cisco ACI based data centers. Each site must be separated from the others. Which characteristic of Cisco ACI Multi-Pod makes it unsuitable for this deployment?

A. creates a virtual pod in the remote location

B. requires all pods to share the same Cisco APIC cluster

C. has distance and scale limitations

D. places leaf switches in the remote site that belong to the same fabric as at the headquarters site

Answer: B

Explanation:

The characteristic of Cisco ACI Multi-Pod that makes it unsuitable for deploying three separate data centers is that it requires all pods to share the same Cisco APIC cluster. In a Multi-Pod deployment, all the pods are part of the same fabric and are managed by a single APIC cluster, which means that they are not completely

isolated from each other.

Question: 161

A network engineer must configure a new SNMP configuration and syslog servers. The requirement is for all faults and events related to endpoint groups, bridge domains, and VRFs to be sent to it. Which action must be taken to meet the requirements?

A. Enable access monitoring policies on the required endpoint groups, bridge domains, and VRFs.

B. Utilize common tenant monitoring policies in the Cisco APIC.

C. Configure fabric monitoring policies and attach to the spine switch in the fabric.

D. Implement fabric-wide monitoring policies on all nodes.

Answer: B

Explanation:

To meet the requirement of sending all faults and events related to endpoint groups, bridge domains, and VRFs to the new SNMP configuration and syslog servers, the network engineer should utilize common tenant monitoring policies in the Cisco APIC. The common tenant is used for global configurations that apply across the entire fabric, including monitoring policies.

Question: 162

An engineer must configure a Layer 3 connection to the WAN router. The hosts in production VRF must access WAN subnets. The engineer associates EPGs in the production VRF with the external routed domain.

Which action completes the task?

A. Configure the Export Route Control Subnet scope for the external EPG.

B. Configure the External Subnets for the External EPG scope for the external EPG.

C. Configure the Import Route Control Subnet scope for the external EPG.

D. Configure the Shared Route Control Subnet scope for the external EPG.

Answer: A

Explanation:

To complete the task of configuring a Layer 3 connection to the WAN router and allowing hosts in the production VRF to access WAN subnets, the engineer should configure the Export Route Control Subnet scope for the external EPG. This setting allows the subnets associated with the external EPG to be advertised to external routers, enabling connectivity to the WAN.

Question: 163

Refer to exhibit.

Fault Code: F3222

Severity: '.vain ng

Last Transition: 2021-02-08T22 08 45 469*00 00

Lifecycle: Raised

Affected Object: **id/subj-DefaultS/ns-cnat-west-vmm-vlp**

Description: Fault delegate Resouice Pool has been used till the threshold

Type: Operational

Cause: resouice-pool-consumed

Change Set: usageStatus Old green New real

Created: 2021-02-08T21 59 45 876'00 00

Code: F3222

Number of Occurrences: -

Original Seventy: warning

Previous Seventy: cleared

Highest Seventy: warn ng

Refer to the exhibit. A Cisco APIC raises an error when the EPG must accept endpoints from a VMM domain created. Which action clears the fault?

A. Expand the VLAN pool for the VMM domain.

B. Create a bridge domain for the VMM domain.

C. Associate the EPG with the VMM domain.

D. Associate the VLAN pool with the VMM domain.

Answer: C

Explanation:

To clear the fault raised by the Cisco APIC when the EPG must accept endpoints from a VMM domain, the EPG should be associated with the VMM domain. This association is necessary for the APIC to recognize the endpoints within the VMM domain as part of the EPG. Without this association, the APIC cannot apply the policies defined in the EPG to the endpoints, which can result in errors1.

Reference:

Question: 164

A Cisco ACI endpoint group must have its gateway address migrated out of the ACI fabric. An engineer configures EPG-TEST with a static port binding and configures the encap VLAN with the required VLAN. Which configuration set must be used on the bridge domain to meet these requirements?

A. L2 Unknown Unicast: Hardware Proxy

Unicast Routing: Disabled

ARP Flooding: Enabled

B. L2 Unknown Unicast: Hardware Proxy

Unicast Routing: Disabled

ARP Flooding: Disabled

C. L2 Unknown Unicast: Flood

Unicast Routing: Disabled

ARP Flooding: Enabled

D. L2 Unknown Unicast: Flood

Unicast Routing: Enabled

ARP Flooding: Enabled

Answer: A

Explanation:

To migrate the gateway address out of the ACI fabric for an endpoint group (EPG), the bridge domain configuration must ensure that routing is contained within the ACI fabric and that ARP requests are managed efficiently. The correct configuration set is:

L2 Unknown Unicast: Set to Hardware Proxy. This setting allows the ACI fabric to use the spine proxy function to handle unknown unicast traffic, which helps to reduce flooding within the fabric.

Unicast Routing: Set to Disabled. Since the gateway is being migrated out of the ACI fabric, unicast routing should be disabled to prevent the ACI fabric from attempting to route traffic for the EPG.

ARP Flooding: Set to Enabled. This allows ARP requests to be flooded within the bridge domain, which is

necessary when the unicast routing is disabled, to ensure that ARP requests can still be resolved.

Reference:

Cisco ACI Bridge Domain Configuration Guide

Cisco ACI Best Practices Guide

Question: 165

What is a requirement for Cisco ACI IPN to manage multidestination traffic?

A. pervasive gateway

B. unicast routing

C. anycast gateway

D. multicast routing

Answer: D

Explanation:

For Cisco ACI IPN (Inter-Pod Network) to manage multidestination traffic, multicast routing is a requirement. This is because multidestination traffic includes broadcast, unknown unicast, and multicast (BUM) traffic, which needs to be efficiently transported across different pods in a Cisco ACI Multi-Pod environment. Multicast routing protocols like Bidirectional Protocol-Independent Multicast (Bidir PIM) are used to handle this type of traffic1.

Reference:

Cisco Application Centric Infrastructure Multi-Pod Configuration White Paper1

Question: 166

An organization deploys active-active data centers and active-standby firewalls in each data center. Which action should be taken in a Cisco ACI Multi-Pod to maintain traffic symmetry through the firewalls?

A. Disable Resilient Hashing.

B. Disable service node Health Tracking.

C. Enable Pod ID Aware Redirection.

D. Enable Endpoint Dataplane Learning.

Answer: C

Explanation:

In a Cisco ACI Multi-Pod environment with active-active data centers and active-standby firewalls in each data center, enabling Pod ID Aware Redirection is the action that should be taken to maintain traffic symmetry through the firewalls. This feature ensures that traffic is consistently directed through the active firewall based on the pod identifier, thus maintaining the desired traffic flow and symmetry2.

Reference:

Cisco ACI Multi-Pod and Service Node Integration White Paper

Question: 167

Refer to the exhibit.

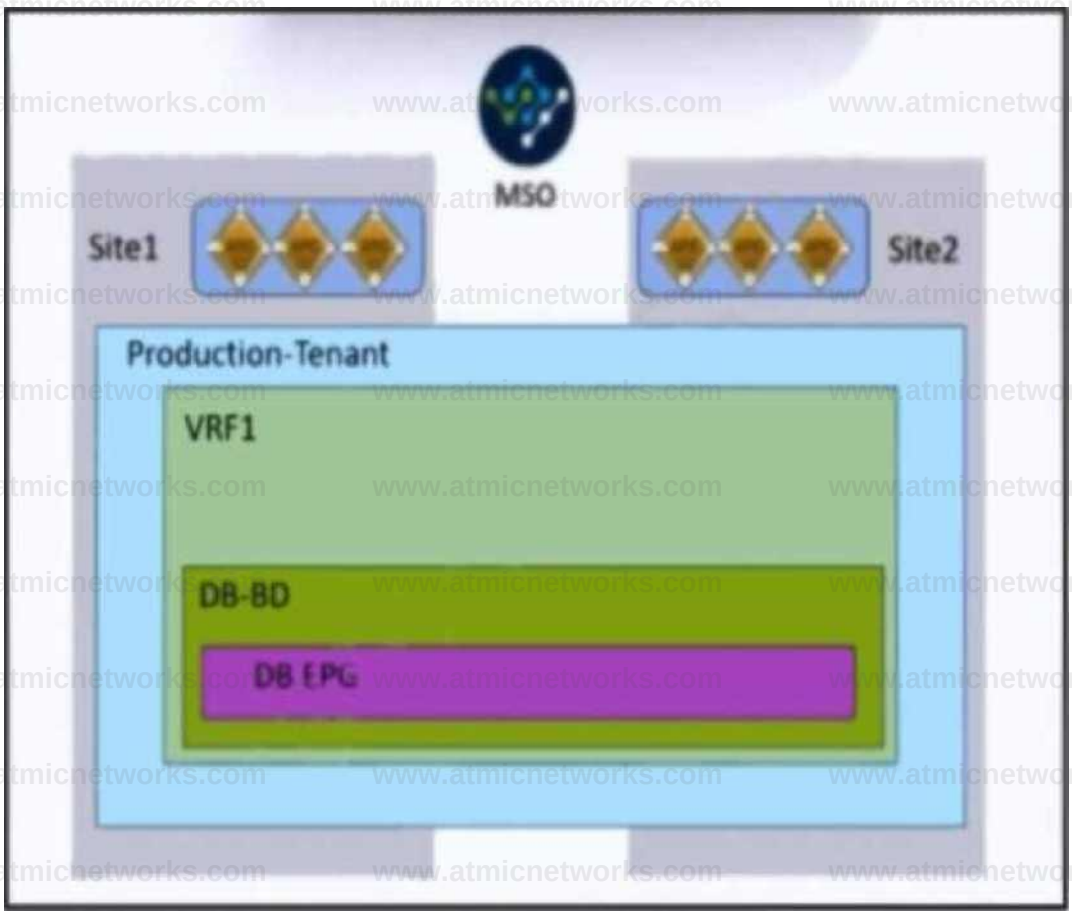

Refer to the exhibit. An engineer must allow IP mobility between Site1 and Site2 in a Cisco ACI MultiSite orchestrator. The design must meet these requirements:

A disaster recovery (DR) solution must exist between the sites that do not require vMotion support.

The application must be started at a DR site without having to re-IP the application servers.

The solution must avoid any broadcast storms between the sites.

Which two actions meet these criteria? (Choose two.)

A. Define a unique bridge domain subnet per site.

B. Configure STP between Cisco ACI fabrics.

C. Deploy a local EPG for Site1 and Site2.

D. Disable Inter-site BUM Traffic.

E. Apply the L2 Stretch feature.

Answer: D, E

Explanation:

To allow IP mobility between Site1 and Site2 in a Cisco ACI Multi-Site orchestrator while meeting the disaster recovery (DR) requirements and avoiding broadcast storms, the following actions should be taken:

Disable Inter-site BUM Traffic: Disabling Broadcast, Unknown unicast, and Multicast (BUM) traffic between sites helps to avoid broadcast storms that can occur when extending Layer 2 domains across multiple sites1.

Apply the L2 Stretch feature: The Layer 2 Stretch feature allows subnets to be extended across multiple sites without the need to re-IP the application servers. This supports IP mobility and enables applications to be started at a DR site using the same IP address space1.

These actions ensure that the design supports a DR solution without requiring vMotion support, allows applications to be started at a DR site without re-IPing servers, and avoids broadcast storms between the sites.

Reference:

Cisco ACI Multi-Site Architecture White Paper1

Cisco Nexus Dashboard Orchestrator Overview2

Question: 168

Refer to the exhibit.

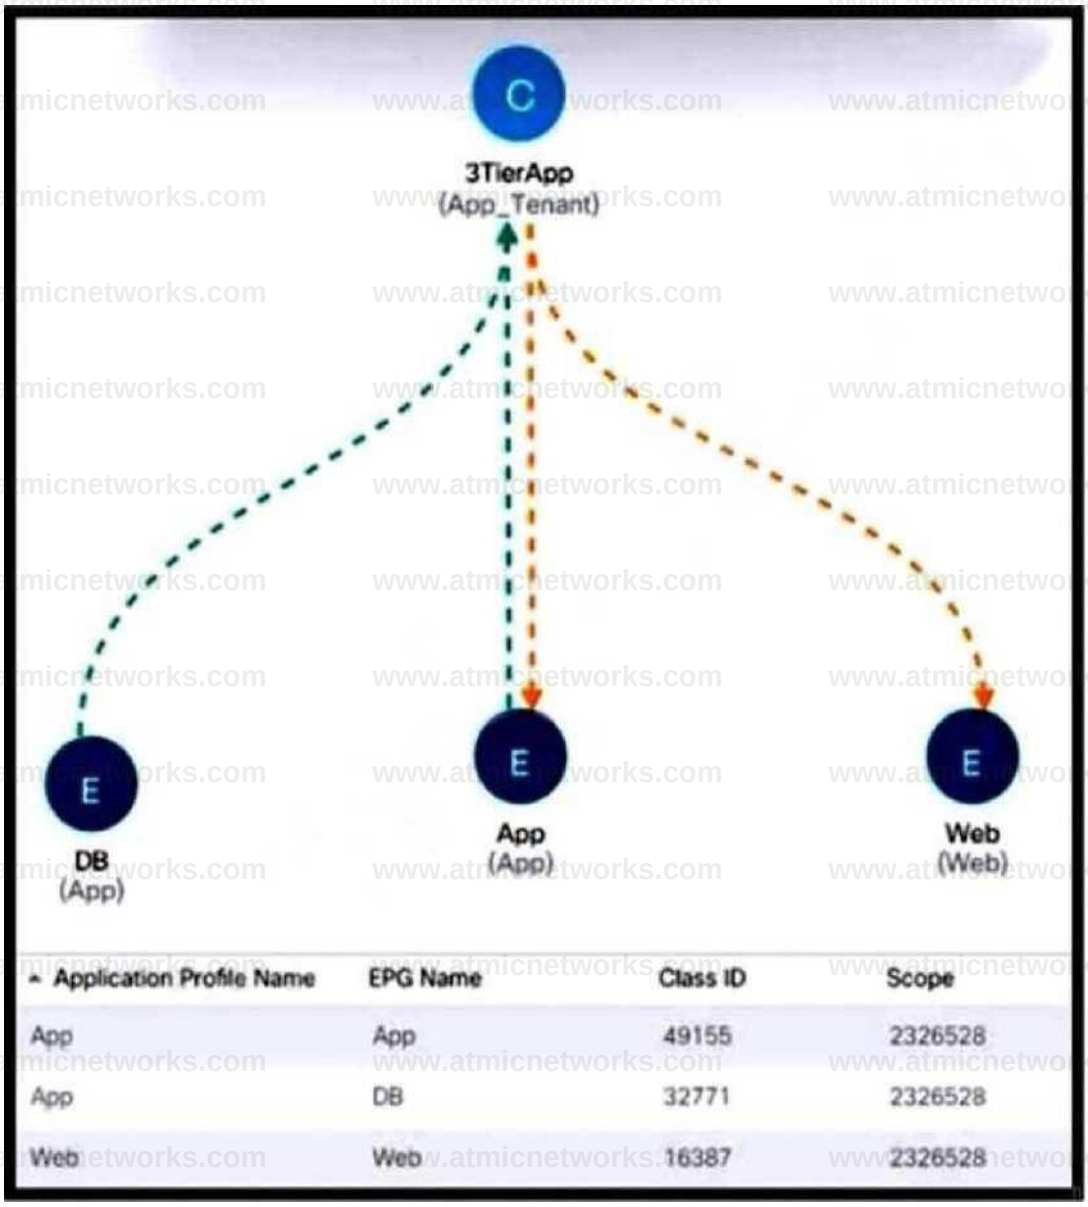

Application Profile Name	EPG Name	Class ID	Scope
App	App	49155	2326528
App	DB	32771	2326528
Web	Web	16387	2326528

Refer to the exhibit. New e-commerce software is deployed on Cisco ACI fabric. The environment must meet these requirements:

The overall number of contracts must be reduced by reusing the existing contracts within a VRF when possible.

The e-commerce software must communicate only with software EPGs that are part of the same ANP.

The e-commerce software must be prevented from communicating with applications in different ANPs.

Which scope must be selected to meet these requirements?

A. Application Profile
B. Endpoint Group

C. Tenant

D. Global

Answer: A

Explanation:

To meet the requirements for the new e-commerce software deployed on the Cisco ACI fabric, the scope of the contracts should be set to “Application Profile”. This scope ensures that the contracts are applied within the same Application Network Profile (ANP), allowing the e-commerce software to communicate only with software Endpoint Groups (EPGs) that are part of the same ANP. It also prevents communication with applications in different ANPs, thus maintaining the necessary isolation and reducing the overall number of contracts by reusing existing ones within a VRF.

Reference:

Cisco ACI Contracts and Scopes Documentation

Cisco ACI Best Practices Guide

Question: 169

Refer to the exhibit.

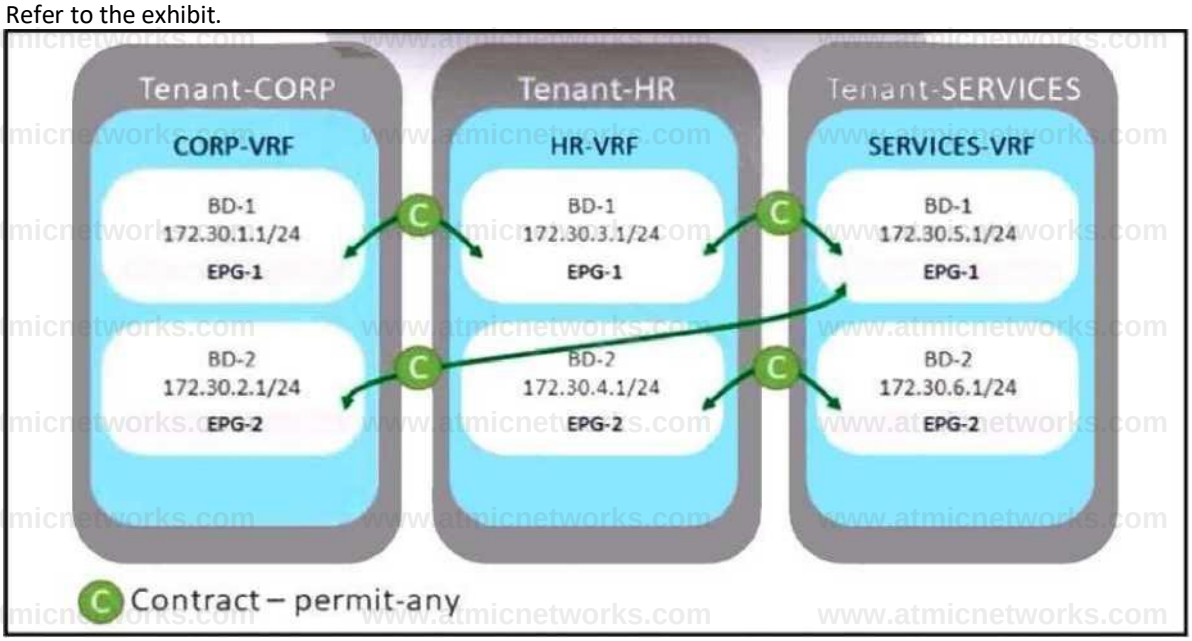

Refer to the exhibit. A company merges three of its departments: CORP, HR, and SERVICES, Currently, the connectivity between departments is achieved by using VRF route leaking. The requirement is to redesign the Cisco ACI networking architecture to communicate between EPGs and BDs from any tenant without configuring contracts or VRF route leaking. Which configuration meets these criteria?

A. Configure an unenforced VRF in the user tenant and map all required EPGs to it.

B. Implement an enforced VRF in the common tenant and map all required BDs to it.

C. Configure an enforced VRF in the user tenant and map all required EPGs to it.

D. Implement an unenforced VRF in the common tenant and map all required BDs to it.

Answer: D

Explanation:

To enable communication between EPGs and BDs from any tenant without the need for contracts or VRF route leaking, the best approach is to implement an unenforced VRF within the common tenant. By doing so, all bridge domains that are associated with this VRF will be able to communicate with each other without additional configurations. This method simplifies the network architecture and reduces the complexity associated with contract management and VRF route leaking1.

Reference:

Cisco ACI Contract Guide White Paper1

Question: 170

What two actions should be taken to deploy a new Cisco ACI Multi-Pod setup? (Choose two.)

A. Configure MP-BGP on IPN routers that face the Cisco ACI spines.

B. Connect all spines to the IPN.

C. Configure anycast RP for the underlying multicast protocol

D. Configure the TEP pool of the new pod to be routable across the IPN.

E. Increase interface MTU for all IPN routers to support VXLAN traffic.

Answer: D, E

Explanation:

When deploying a new Cisco ACI Multi-Pod setup, it's crucial to ensure that the Tunnel Endpoint (TEP) pool of the new pod is routable across the Inter-Pod Network (IPN). This allows the TEP addresses to be reachable across the pods, facilitating communication between them. Additionally, increasing the Maximum Transmission Unit (MTU) for all IPN routers is necessary to support the larger frame size required by VXLAN encapsulation, which is typically above the standard Ethernet MTU of 1500 bytes.

Reference:

Cisco ACI Multi-Pod Configuration Guide

Cisco ACI Best Practices Guide

Question: 171

Refer to the exhibit.

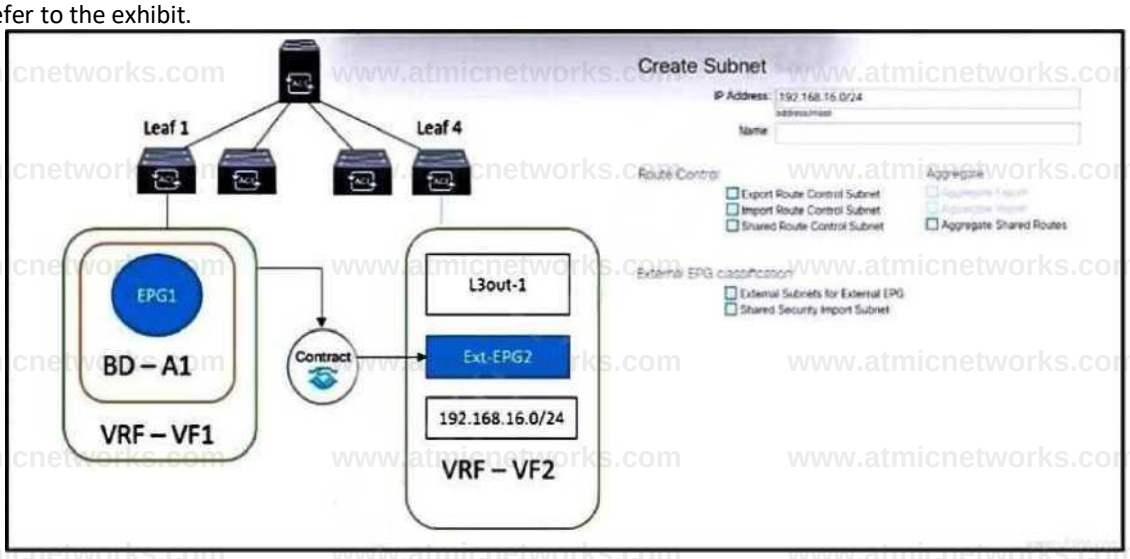

Refer to the exhibit. The external subnet and internal EPG1 must communicate with each other, and the L3Out traffic must leak into the VRF named "VF1". Which configuration set accomplishes these goals?

A.

Export Route Control Subnet

Import Route Control Subnet

Aggregate Shared Routes

B.

External Subnets for External EPG

Shared Route Control Subnet

Shared Security Import Subnet

C.

External Subnets for External EPG

Import Route Control Subnet

Shared Route Control Subnet

D.

Export Route Control Subnet

Shared Security Import Subnet

Aggregate Shared Routes

Answer: B

Explanation:

To enable communication between the external subnet and internal EPG1, and to allow L3Out traffic to leak into the VRF named “VF1”, the following configuration set should be used:

External Subnets for External EPG: This configuration defines the subnets that are external to the ACI fabric but need to be reachable from within the fabric. It is necessary to specify which subnets are to be considered part of the L3Out1.

Shared Route Control Subnet: This setting allows the subnets to be shared across different VRFs, enabling communication between EPGs that are in different tenants but within the same VRF1.

Shared Security Import Subnet: This configuration ensures that the security policies (contracts) associated with the shared subnets are also imported, allowing for the necessary traffic to flow between the internal and external endpoints1.

By applying this configuration set, the external subnet and internal EPG1 will be able to communicate, and the L3Out traffic will be properly leaked into the designated VRF, meeting the specified goals.

Reference:

ACI Inter VRF/Tenant Route Leaking Configuration Example1

Question: 172

Refer to the exhibit. VM1 and VM2 are in Cisco ACI POD1 and communication takes place. Which event is triggered when VM2 is live migrated from POD1 to POD2?

A. Leaf 102 installs a bounce entry for VM2 pointing to the PTEP address of leaf 201.

B. Leaf 201 creates a tunnel with leaf 102 because of the bounced traffic that is destined to VM2.

C. Spines from POD2 send an MP-BGP EVPN update to the leaves in POD1 about the new location of VM2.

D. An MP-BGP EVPN update is received by spines in POD1 announcing the reachability of VM2 via the proxy VTEP address of the spines in POD2.

Answer: C

Explanation:

When VM2 is live migrated from POD1 to POD2, the spines in POD2 will send an MP-BGP EVPN update to the leaves in POD1. This update informs the leaves in POD1 about the new location of VM2, allowing them to update their forwarding tables and ensure that traffic destined for VM2 is correctly routed to its new location in POD2.

Question: 173

An engineer configures a Cisco ACI Multi-Pod for disaster recovery. Which action should be taken for the new nodes to be discoverable by the existing Cisco APICs?

A. Configure IGMPv3 on the interfaces of IPN routers that face the Cisco ACI spine.

B. Enable subinterfaces with dot1q tagging on all links between the IPN routers.

C. Enable DHCP relay on all links that are connected to Cisco ACI spines on IPN devices.

D. Configure BGP as the underlay protocol in IPN.

Answer: C

Explanation:

For the new nodes in a Cisco ACI Multi-Pod deployment to be discoverable by the existing Cisco APICs, DHCP relay should be enabled on all links that are connected to Cisco ACI spines on Inter-Pod Network (IPN) devices. This allows the APICs to dynamically assign IP addresses to the new nodes, facilitating their discovery and integration into the fabric.

Reference:

Cisco ACI Multi-Pod Configuration Guide

Cisco ACI Best Practices Guide

Question: 174

Refer to the exhibit.

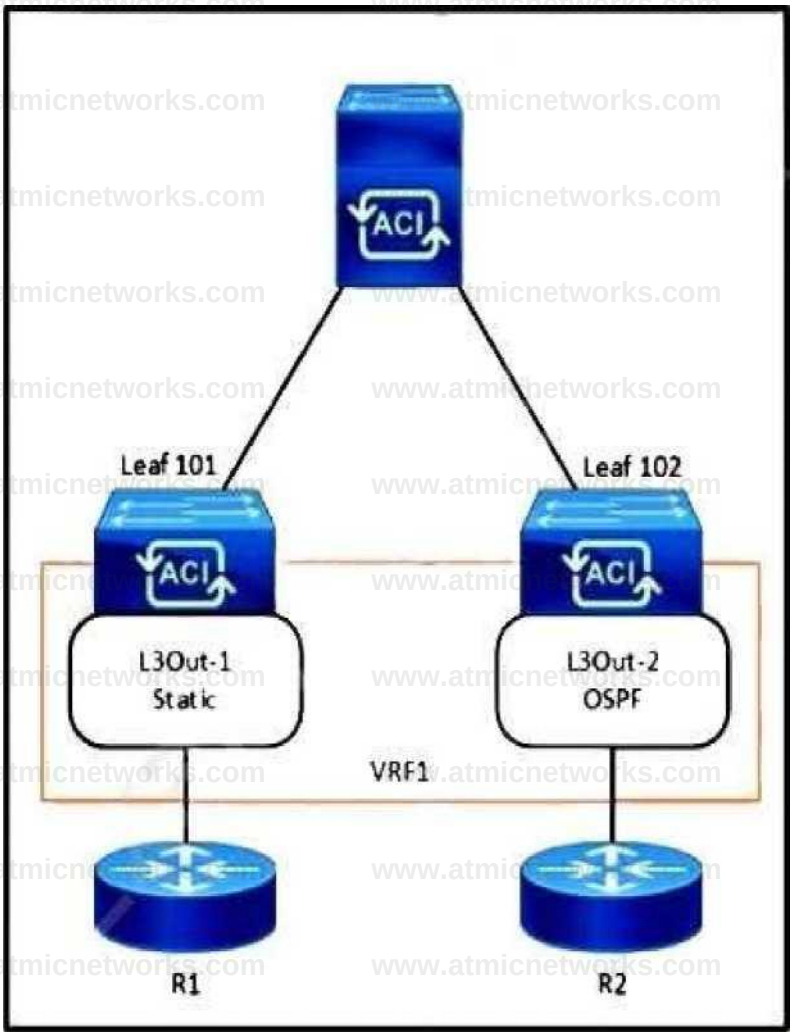

Refer to the exhibit. The 0.0.0.0/0 is configured as a default static route on L3Out-1. Which action should be taken for the 0.0.0.0/0 prefix to advertise out on L3Out-2 OSPF?

A. Enable Export Route Control Subnet.

B. Enable Shared Security Import Subnet.

C. Enable Shared Route Control Subnet.

D. Enable Aggregate Export Subnet.

Answer: A

Explanation:

To advertise the 0.0.0.0/0 prefix out on L3Out-2 OSPF, when it is configured as a default static route

on L3Out-1, the action that should be taken is to enable the Export Route Control Subnet. This setting allows the default route to be advertised to other L3Outs, enabling the ACI fabric to act as a transit network1.

Reference:

ACI Fabric L3Out White Paper - Cisco1

Question: 175

A packet is routed between two endpoints on different Cisco ACI leaf switches. Which VXLAN VNID is applied to the packet?

A. FD

B. EPG

C. VRF

D. BD

Answer: D

Explanation:

In Cisco ACI, when a packet is routed between two endpoints on different leaf switches, the VXLAN VNID that is applied to the packet corresponds to the Bridge Domain (BD). The BD VNID is used to encapsulate the packet for Layer 2 transport across the fabric. This VNID ensures that the packet is delivered to the correct bridge domain, maintaining the segmentation and policy enforcement required by the ACI fabric.

Reference:

Cisco ACI Fabric Forwarding White Paper

Cisco ACI Basic Configuration Guide

Question: 176

Refer to the exhibit.

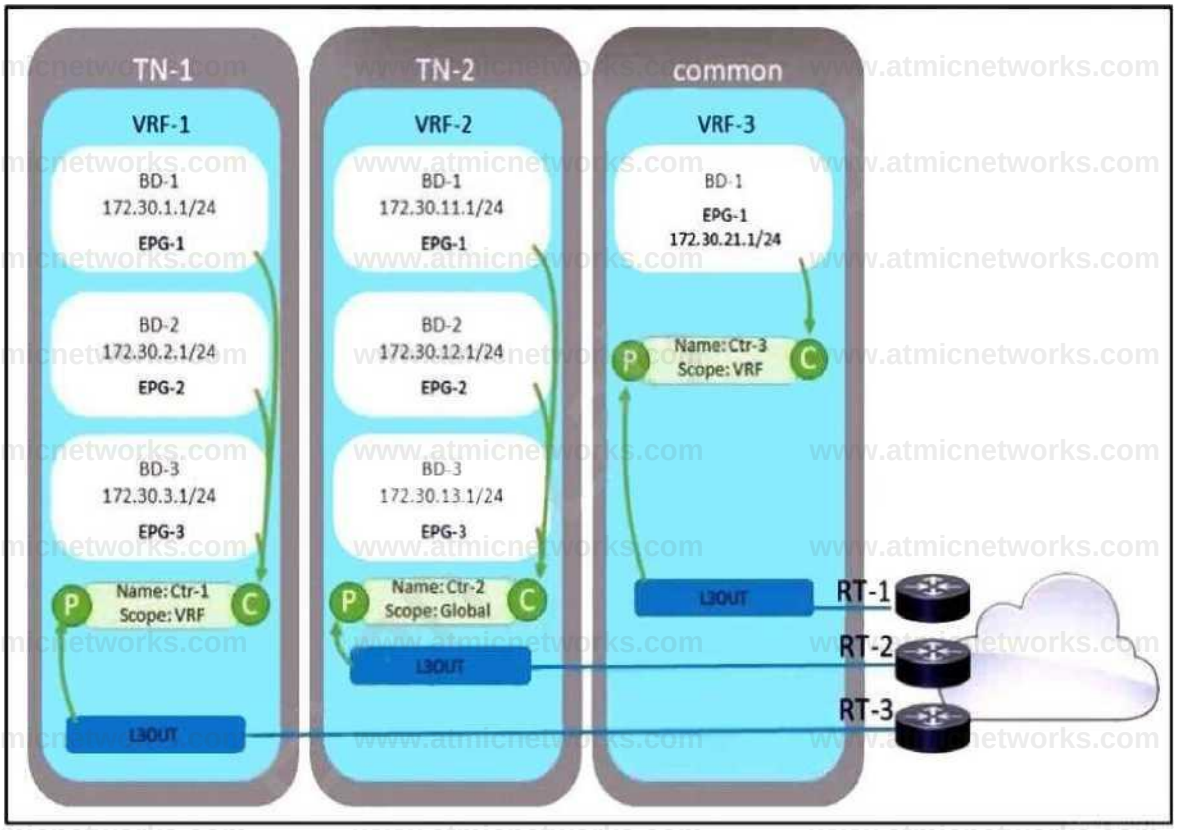

Refer to the exhibit. A company decided to decrease its routing footprint and remove RT-2 and RT-3 devices from its data center. Because of that, the exit point must be created from all the tenants by using the common tenant. Which two configuration tasks must be completed to meet these requirements? (Choose two.)

A. Move subnets from all the bridge domains to the EPG level and mark them with flag Shared between VRFs.

B. Update the L3Out ExtEPG subnet in the common tenant with flag Shared Route Control Subnet and Aggregate Shared Routes.

C. Mark all subnets with flag Shared between VRFs and attach contract Ctr-3 as a provider to all the EPGs.

D. Change contract Ctr-3 scope to Global, consume it by all EPGs, and flag all subnets with flag Shared between VRFs.

E. Export contract Ctr-2 into the tenant TN-1 and attach it as a consumer to all the EPGs in the tenant TN-1.

Answer: B, D

Explanation:

To create an exit point from all tenants using the common tenant and decrease the routing footprint, the

following configuration tasks must be completed:

Update the L3Out ExtEPG subnet in the common tenant with flag Shared Route Control Subnet and Aggregate Shared Routes: This configuration allows the subnets to be shared across different VRFs within the common tenant, enabling communication between EPGs that are in different tenants1.

Change contract Ctr-3 scope to Global, consume it by all EPGs, and flag all subnets with flag Shared between VRFs: By changing the scope of contract Ctr-3 to Global, it can be consumed by all EPGs across the fabric. Additionally, flagging all subnets with Shared between VRFs ensures that the subnets can be used by multiple tenants1.

Reference:

Cisco Community Discussion on Common Tenant1

Cisco ACI Basic Configuration Guide2

Cisco ACI Configuring Shared L3Outs Documentation

Question: 177

A company must connect three Cisco ACI data centers by using Cisco ACI Multi-Site. An engineer must configure the Inter-Site Network (ISN) between the existing sites. Which two configuration steps must be taken to implement the ISN? (Choose two.)

A. Configure OSPF on subinterfaces on routers that are directly connected with spine nodes.

B. Configure ISN site extension on Cisco routers in the network.

C. Configure OSPF on all ISN routers.

D. Configure BIDIR-PIM on all ISN routers.

E. Configure encapsulation VLAN-4 between the routers and spine nodes.

Answer: C, D

Explanation:

To implement the Inter-Site Network (ISN) for a Cisco ACI Multi-Site deployment, the following configuration steps are essential:

Configure OSPF on all ISN routers: OSPF (Open Shortest Path First) is a routing protocol that is used to ensure that all routers in the ISN have the necessary routing information to forward packets between sites.

Configure BIDIR-PIM on all ISN routers: BIDIR-PIM (Bidirectional Protocol Independent Multicast) is used for efficient multicast traffic forwarding across the ISN. This is particularly important for Cisco ACI Multi-Site deployments as it supports the replication of broadcast, unknown unicast, and multicast (BUM) traffic across sites.

Reference:

Cisco ACI Multi-Site Configuration Guide

Cisco ACI Multi-Site Architecture White Paper

Question: 178

Refer to the exhibit.

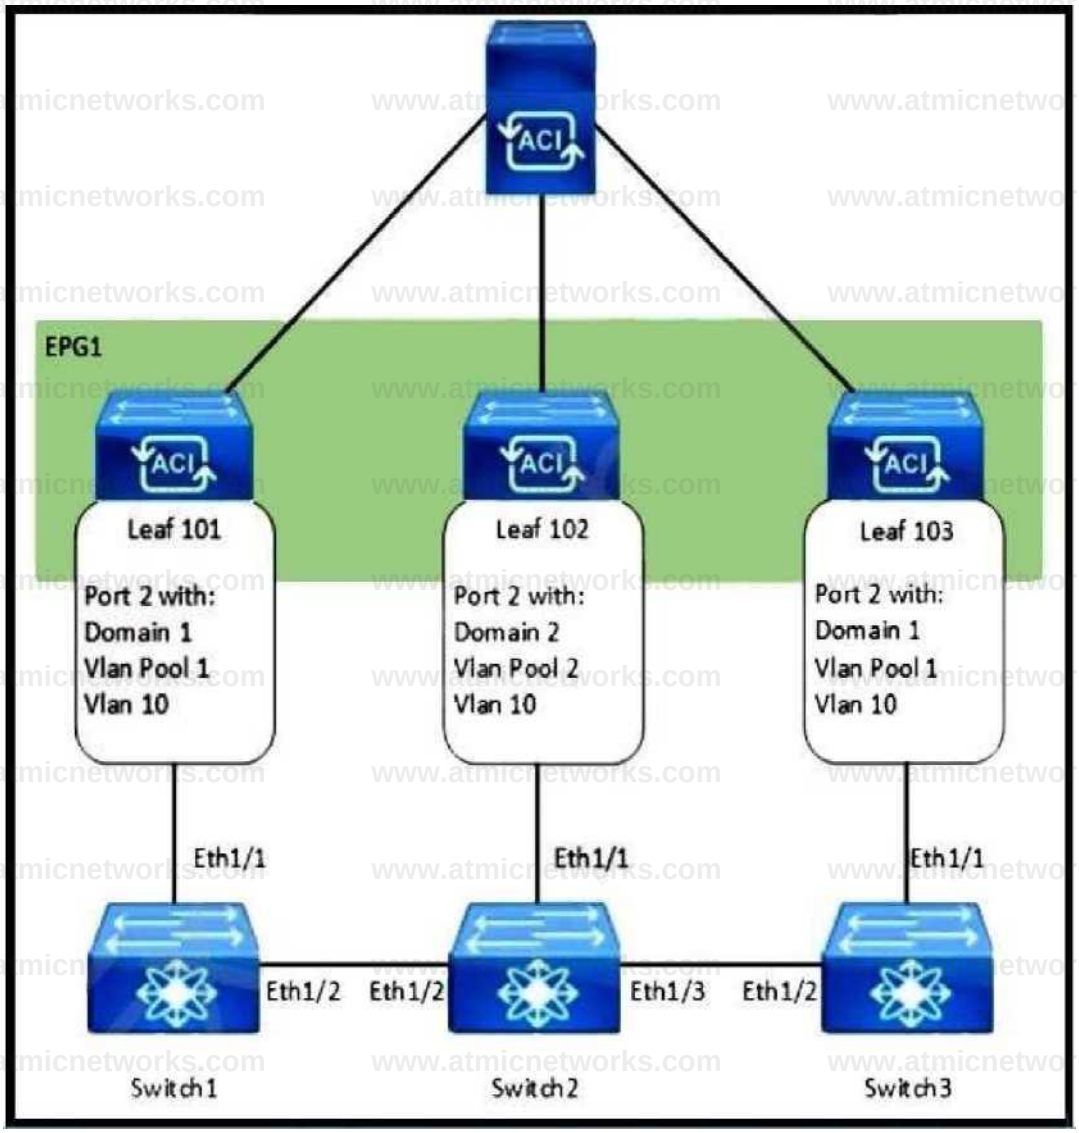

Refer to the exhibit. How are the STP BPDUs forwarded over Cisco ACI fabric?

A. Cisco ACI acts as the STP root for all three external switches.

B. STP BPDUs that are generated by Switch2 are received by Switch1 and Switch3

C. STP BPDUs that are generated by Switch1 are received only by Switch3.

D. Cisco ACI fabric drops all STP BPDUs that are generated by the external switches.

Answer: B

Explanation:

In a Cisco ACI fabric, Spanning Tree Protocol (STP) Bridge Protocol Data Units (BPDUs) that are generated by an external switch, such as Switch2 in this scenario, are indeed forwarded across the fabric. This allows externally

connected switches to maintain a loop-free topology. While Cisco ACI does not participate in STP as it uses its own protocols and mechanisms for loop prevention within the fabric, it will forward STP BPDUs across End Point Groups (EPGs) on which they are received1. This ensures that the connected external switches can still use STP for their loop prevention mechanisms.

Reference:

Cisco Community Discussion on STP BPDUs in ACI1

Question: 179

Refer to the exhibit.

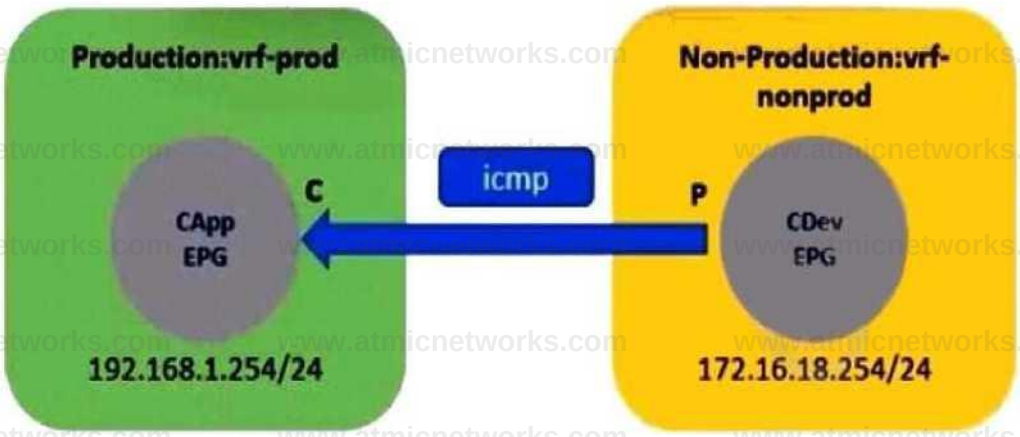


```
leaf-01* show ip route vrf Production:vrf-prod
IP Route Table for VRF "Production:vrf-prod" '•' denotes best ucast next-hop '••' denotes best
mcast next-hop *[x/y]' denotes [preference/metric]
 %<string>' In via output denotes VRF <string>

172.16.18.    0/24, ubest/mbest: 1/0, attached, direct, pervasive
      •via 10.0.64.66%overlayl, [1/0], 00:02:10, static, tag 4294967294
192.168.1.    0/24, ubest/abest: 1/0, attached, direct, pervasive
      •via 10.0.64.66%overlay-l, [1/0], ld00h, static
192.168.1.254/32, ubest/mbest: 1/0, attached, pervasive
      •via 192.168.1.254, vlanl4, [0/0], ld00h, local, local

leaf-01* show ip route vrf Non-Production:vrf-nonprod
IP Route Table for VRF "Non-Production:vrf-nonprod" '•' denotes best ucast next-hop '••■
denotes best mcast next-hop '[x/y]' denotes [preference/oetric] "%<stringy" in via output
denotes VRF <$tring>

172.16.18.    0/24, ubest/mbest: 1/0, attached, direct, pervasive
      •via 10.0.64.66%overlay-l, [1/0], 00:02:15, static, tag 4294967294
172.16.18.254/32, ubest/mbest: 1/0, attached, pervasive
      •via 172.16.18.254, vlanl6, [0/0], 00:21:33, local, local
```

Refer to the exhibit. An administrator configures inter-VRF route leaking between Production:vrf- prod and Non-Production:vrf-nonprod. However, the route in the Non-Production:vrf-nonprod VRF to the production tenant is missing. Which action resolves the VRF route leaking issue?

A. Change the contract scope to Global.

B. Enable the Shared between VRFs option for the BD subnet in the production VRF.

C. Enable the Shared between VRFs option for the EPG subnet in the non-production VRF.

D. Export the contract from provider to consumer tenant.

Answer: B

Explanation:

To resolve the VRF route leaking issue and ensure that the route in the Non-Production:vrf-nonprod VRF to the production tenant is present, the action that should be taken is to enable the “Shared between VRFs” option for the bridge domain (BD) subnet in the production VRF. This setting allows the subnet to be shared across different VRFs within the same tenant or across tenants, enabling communication between EPGs that are in different VRFs1.

Reference:

ACI Inter VRF/Tenant Route Leaking Configuration Example

Question: 180

How is broadcast forwarded in Cisco ACI Multi-Pod after ARP flooding is enabled?

A. Ingress replication is used on the spines to forward broadcast frames in the IPN infrastructure.

B. Within a pod, the ingress leaf switch floods the broadcast frame on all fabric ports.

C. Broadcast frames are forwarded inside the pod and across the IPN using the multicast address that is associated to the bridge domain.

D. For the specific bridge domain, all spines forward the broadcast frames to IPN routers.

Answer: C

Explanation:

After ARP flooding is enabled in Cisco ACI Multi-Pod, broadcast frames are forwarded within the pod and across the Inter-Pod Network (IPN) using the multicast address associated with the bridge domain. If the setting is ‘Flood’, the leaf switch floods to the Group IP (GIPo) multicast group allocated for the bridge domain, ensuring that both local and remote pods receive a flooded copy

Question: 181

What are two PBR characteristics of the Cisco ACI Active-Active Across Pods deployment mode in Cisco ACI Multi-Pod design? (Choose two.)

A. Traffic is dynamically redirected to the firewall that owns the connection.

B. Deployment occurs in transparent mode.

C. The connection state is unsynchronized.

D. Deployment occurs in go-to mode only.

E. This mode causes the traffic to flow asymmetrically.

Answer: A, D

Explanation:

In the Cisco ACI Active-Active Across Pods deployment mode, one of the key characteristics of Policy-Based Redirect (PBR) is that traffic is dynamically redirected to the firewall that owns the connection. This allows for load distribution across multiple Layer 4 to Layer 7 devices, also known as symmetric PBR. Additionally, deployment in this mode occurs in go-to mode only, which means that traffic is explicitly directed to a service node or firewall for processing

Question: 182

Refer to the exhibit.

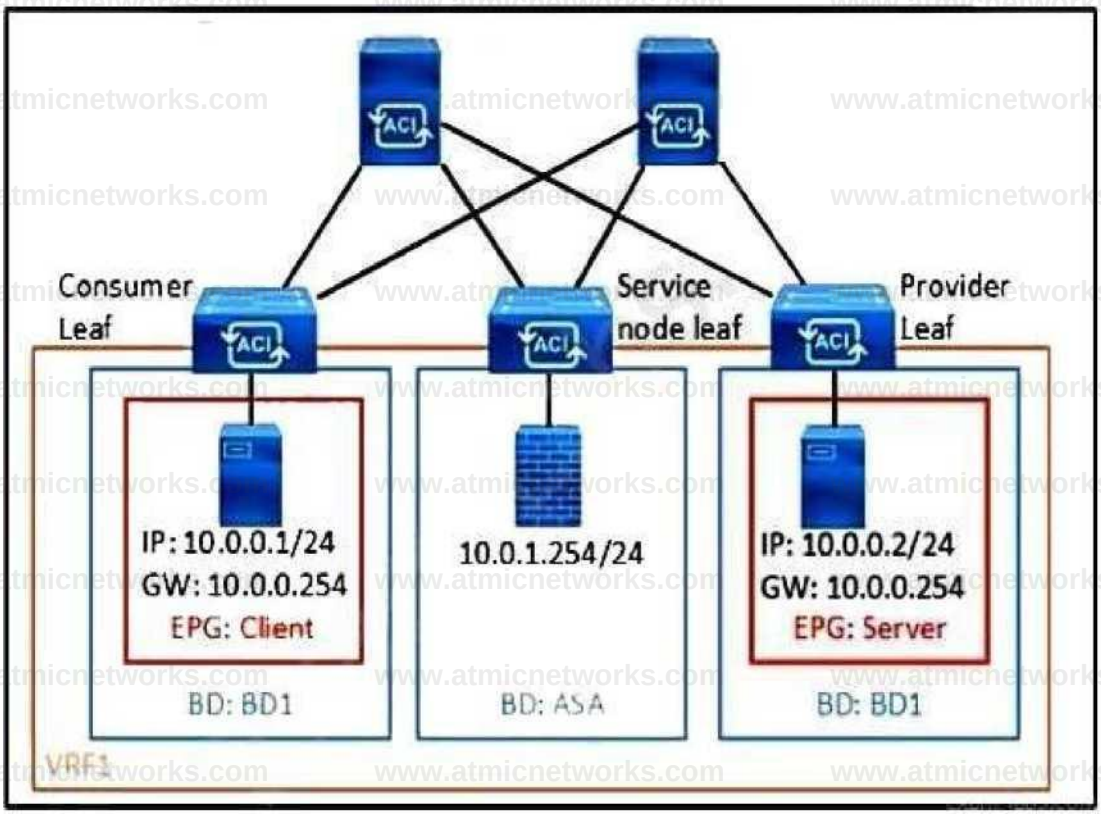

Refer to the exhibit. What must be configured in the service graph to redirect HTTP traffic between the EPG client and EPG server to go through the Cisco ASA firewall?

A. precise filter to allow only HTTP traffic

B. permit-all contract filter

C. contract with no filter

D. contract filter to allow ARP and HTTP.

Answer: A

Explanation:

To redirect HTTP traffic between the EPG client and EPG server through the Cisco ASA firewall using a service graph in Cisco ACI, a precise filter must be configured to allow only HTTP traffic. This filter will specify the protocol (HTTP) and the Layer 4 port (typically port 80 for HTTP) to ensure that only HTTP traffic is redirected to the firewall for inspection or processing. The service graph can then be associated with the relevant EPGs to enforce the redirection1.

Reference:

Service Graph Design with Cisco ACI (Updated to Cisco APIC Release 5.2) White Paper

Question: 183

What is the advantage of implementing an active-active firewall cluster that is stretched across separate pods when anycast services are configured?

A. A cluster is capable to be deployed in transparent mode across pods.

B. A different MAC/IP configuration combination is configurable for the firewall in each pod.

C. Local traffic in a pod is load-balanced between the clustered firewalls.

D. The local pod anycast node is preferred by the local spines.

Answer: C

Explanation:

The advantage of implementing an active-active firewall cluster that is stretched across separate pods when anycast services are configured is that local traffic within a pod is load-balanced between the clustered firewalls. This means that traffic destined for the anycast IP address of the service node (firewall) will preferentially be directed to the local node within the same pod, thus optimizing traffic flows and reducing latency by avoiding unnecessary traffic tromboning to a remote pod1.

Question: 184

Refer to the exhibit.

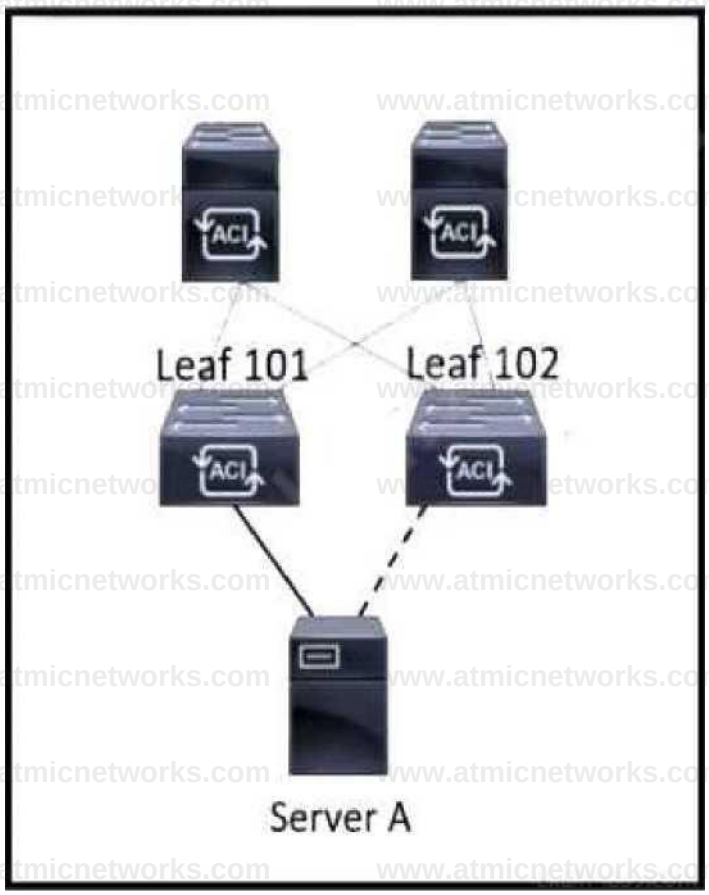

Refer to the exhibit. Server A is connected to the Cisco ACI fabric using two teamed interfaces. One interface in a team is configured as active and the other remains in standby mode. When a failover occurs and the standby interface becomes active, it uses its built-in MAC address to send traffic. Which bridge domain configuration must be applied to resolve the issue?

A. Configure Hardware proxy.

B. Set L2 Unknown Unicast to Flood.

C. Enable ARP flooding.

D. Activate Limit IP Learning to Subnet.

Answer: C

Explanation:

In the scenario where Server A is connected to the Cisco ACI fabric using teamed interfaces with one active and one standby, enabling ARP flooding in the bridge domain configuration is necessary to resolve the issue that occurs when the standby interface becomes active and uses its built-in MAC

address to send traffic. Enabling ARP flooding allows the fabric to learn the new MAC address without waiting for the ARP timeout, which helps to ensure that traffic continues to flow to the correct destination after a failover event1.

Reference:

Discuss Cisco 300-620 Exam Topic 9 Question 71 | Pass4Success1

Question: 185

In a Cisco ACI Multi-Site fabric, the Inter-Site BUM Traffic Allow option is enabled in a specific stretched bridge domain. What is used to forward BUM traffic to all endpoints in the same broadcast domain?

A. ingress replication on the spines in the source site

B. egress replication on the destination leaf switches

C. egress replication on the source leaf switches

D. ingress replication on the spines in the destination site

Answer: A

Explanation:

In a Cisco ACI Multi-Site fabric, when the Inter-Site BUM Traffic Allow option is enabled for a stretched bridge domain, ingress replication on the spines in the source site is used to forward Broadcast, Unknown unicast, and Multicast (BUM) traffic to all endpoints in the same broadcast domain. This method ensures that BUM traffic is replicated and sent out through the spines to the destination sites.

Question: 186

The engineer notices frequent MAC and IP address moves between different leaf switch ports. Which action prevents this problem from occurring?

A. Disable enforce subnet check.

B. Enable endpoint loop protection.

C. Enable rogue endpoint control.

D. Disable IP bridge domain enforcement.

Answer: C

Explanation:

To prevent frequent MAC and IP address moves between different leaf switch ports, enabling rogue endpoint control is the recommended action. This feature helps in identifying and mitigating the movement of endpoints that could potentially cause loops or other issues within the network.

Question: 187

A customer is deploying a new application across two ACI pods that is sensitive to latency and jitter. The application sets the DSCP values of packets to AF31 and CS6, respectively. Which configuration changes must be made on the APIC to support the new application and prevent packets from being delayed or dropped between pods?

A. disable DSCP mapping on the IPN devices

B. disable DSCP translation policy

C. align the ACI QoS levels and IPN QoS policies

D. align the custom QoS policy on the EPG site in the customer tenant

Answer: C

Explanation:

To support a new application that is sensitive to latency and jitter and sets the DSCP values of packets to AF31 and CS6, it is necessary to align the ACI Quality of Service (QoS) levels and Inter-Pod

Network (IPN) QoS policies. This alignment ensures that the DSCP values are preserved and respected across the ACI fabric and the IPN, preventing packets from being delayed or dropped between pods.

Reference:

Cisco ACI Multi-Site Architecture White Paper

Cisco ACI Best Practices Guide

Cisco ACI Quality of Service Configuration Guide

Question: 188

What controls communication between EPGs?

A. Inter-EPG communication is controlled by BGP.

B. Inter-EPG communication is controlled by contracts.

C. Inter-EPG communication is controlled by IS-IS.

D. Inter-EPG communication is controlled by VXLAN.

Answer: B

Explanation:

In Cisco ACI, communication between Endpoint Groups (EPGs) is controlled by contracts. Contracts define the types of traffic and the rules that govern how EPGs can communicate with each other. They act as a policy framework that enforces security and segmentation within the ACI fabric.

Question: 189

Which feature should be disabled on a bridge domain when a default gateway for endpoints is on an external device instead of a Cisco ACI bridge domain SVI?

A. unknown unicast flooding

B. ARP flooding

C. unicast routing

D. proxy ARP

Answer: C

Explanation:

When the default gateway for endpoints is on an external device instead of a Cisco ACI bridge domain Switched Virtual Interface (SVI), the unicast routing feature should be disabled on the bridge domain. This prevents the ACI fabric from attempting to route traffic for the endpoints, which should instead be directed to the external default gateway.

Question: 190

Where are STP BPDUSs flooded in Cisco ACI fabric?

A. in the access encapsulation VLAN part of different VLAN pools

B. in the bridge domain VLAN

C. in the native VLAN ID

D. in the VNID that is assigned to the FD VLAN

Answer: B

Explanation:

In the Cisco ACI fabric, Spanning Tree Protocol (STP) Bridge Protocol Data Units (BPDUs) are flooded within the bridge domain VLAN. This allows the BPDUs to be propagated throughout the bridge domain, enabling STP to function correctly for the associated VLAN.

Question: 191

Refer to the exhibit.

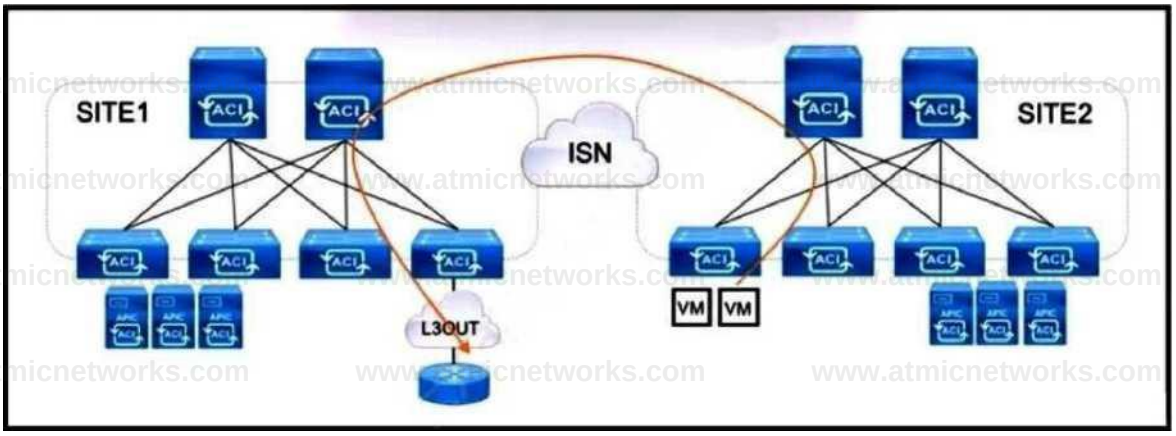

Refer to the exhibit. An engineer is configuring a production Multi-Site solution to provide connectivity from EPGs from a specific site to networks reachable through a remote site L3OUT. All required schema and template objects are already defined. Which additional configuration must be implemented in the Multi-Site Orchestrator to support the cross-site connectivity?

A. Configure a routable TEP pool for SITE1.

B. Enable CloudSec for intersite traffic encryption.

C. Add a new stretched external EPG to the existing L3OUT.

D. Implement a policy-based redirect using a service graph.

Answer: C

Explanation:

To support cross-site connectivity in a production Multi-Site solution, where connectivity from EPGs from a specific site to networks reachable through a remote site L3OUT is required, the additional configuration that must be implemented in the Multi-Site Orchestrator is to add a new stretched

external EPG to the existing L3OUT1. This allows the EPGs from one site to communicate with the networks defined in the L3OUT at the remote site, facilitating the required connectivity across sites.

Reference:

Cisco ACI Multi-Site Configuration Guide1

Question: 192

Refer to the exhibit.

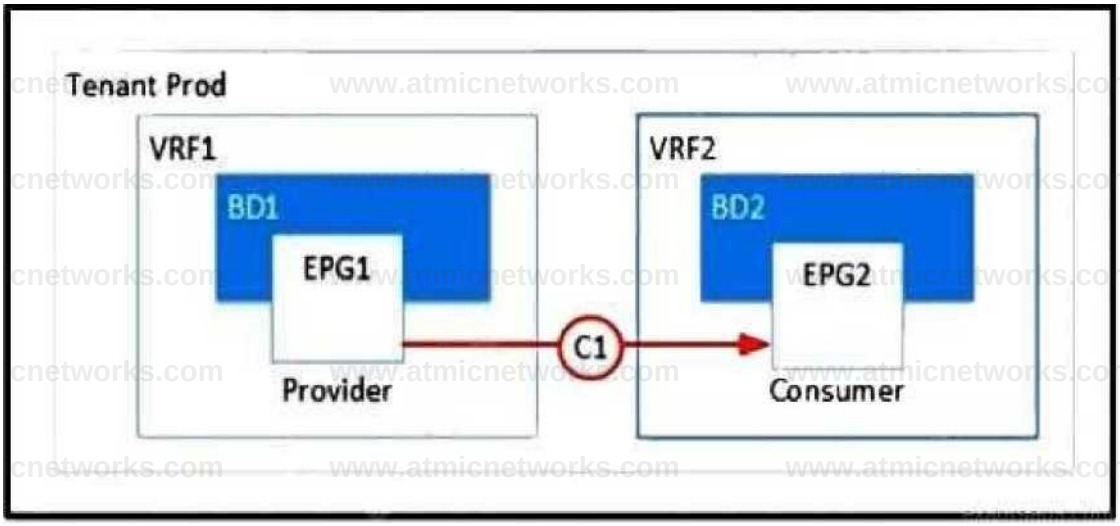

Refer to the exhibit. Which two configurations enable inter-VRF communication? (Choose two.)

A. Set the subnet scope to Shared Between VRFs.

B. Enable Advertise Externally under the subnet scope.

C. Export the contract and import as a contract interface.

D. Change the contract scope to Tenant.

E. Change the subject scope to VRF.

Answer: A, C

Explanation:

To enable inter-VRF communication, the following configurations are necessary:

Set the subnet scope to Shared Between VRFs: This allows the subnets to be shared across different VRFs within the same tenant or across tenants, enabling communication between EPGs that are in different VRFs1.

Export the contract and import as a contract interface: By exporting a contract from the provider tenant and importing it as a contract interface in the consumer tenant, you establish a relationship that allows for controlled communication between EPGs in separate VRFs or tenants1.

Reference:

ACI Inter VRF/Tenant Route Leaking Configuration Example1

Question: 193

A network engineer is implementing a Layer 3 Out in the Cisco ACI fabric. The data center core switches must connect to a pair of leaf switches and exchange routes via a routing protocol. In addition, the implementation must meet these criteria;

- The external switch interface must use 802.1Q tagging.
- Access to the internet for the ACI fabric must be the L30ut.
- The L30ut must use a routing protocol that has rapid convergence time and low CPU usage.

Which configuration set meets these requirements?

A. Configure the OSPF Protocol policy with an area of 0.

Set up the Routed External Network object and Node Profile and select OSPF. Create the Switch profile and select VPC with the appropriate interfaces. Create the default network and associate it with the Routed Outside object.

B. Configure the BGP Protocol policy with the appropriate Autonomous System number. Configure an Interface policy and an External Bridged Domain. Create an External Bridged Network and use the configured VLAN pool. Build the Leaf profile and select the Routed sub-interface with the appropriate VLAN.

C. Implement the IS-IS Protocol policy with the selected Autonomous System number. Create the Routed Outside object and Node Profile and select IS-IS. Configure the Interface profile and select the Routed Interface with the appropriate interfaces. Create the External Network object.

D. Implement the EIGRP Protocol policy with the selected Autonomous System number. Create Routed Outside object and Node Profile and select EIGRP as the routing protocol. Build the Interface profile and select SVI and the appropriate VPC. Configure the External Network object with a network of 0.0.0.070.

Answer: D

Explanation:

Question: 194

Refer to the exhibit.

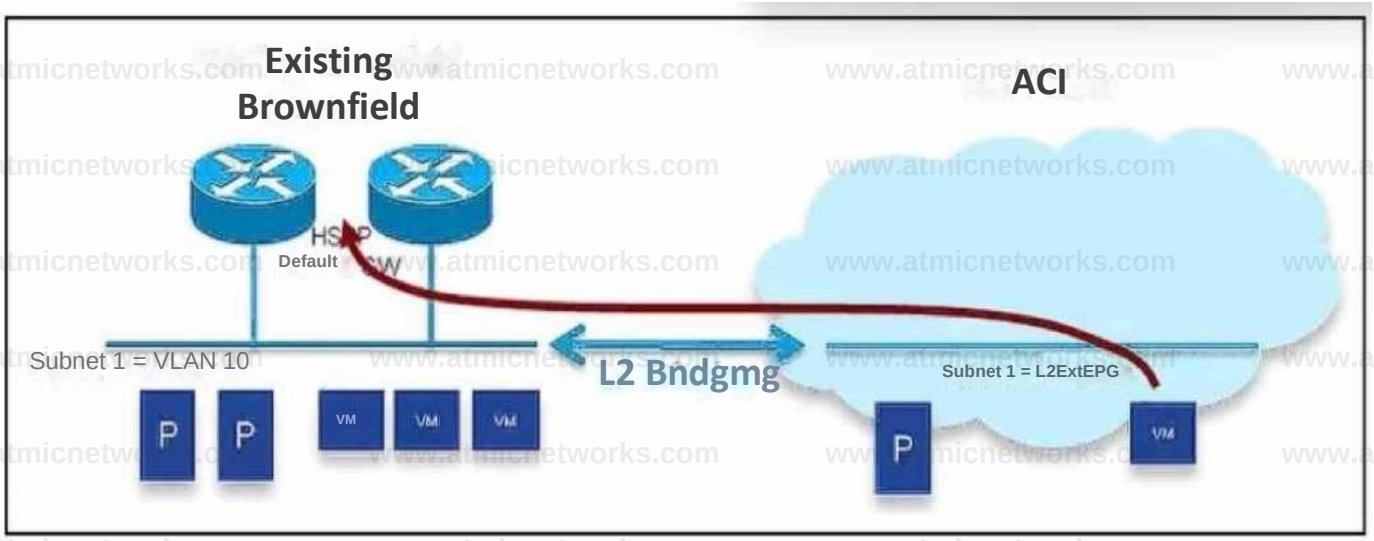

An engineer must migrate workloads from the brownfield network to the Cisco ACI fabric. The VLAN 10 default gateway remains in the router located in the brownfield Network. The bridge domain has already been associated with L20ut. Which two actions must be taken to migrate the workloads? (Choose two.)

A. Select Limit IP Learning to Subnet.

B. Configure Multi-Destination Flooding Flood in Encapsulation.

C. Set L2 Unknown Unicast Flood.

D. Map the MAC address of the default gateway to the bridge domain

E. Enable ARP Flooding

Answer: C, E

Explanation:

Question: 195

How does Cisco ACI detect the IP address of a silent host that moved from one location to another without

notifying a Cisco ACI leaf?

A. ARP requests are flooded in the bridge domain.

B. Bounce entries are installed on the leaf switch.

C. Endpoint announce messages are sent to COOP.

D. Silent hosts are detected by the ACI fabric.

Answer: A

Explanation:

In a Cisco ACI fabric, when a silent host moves from one location to another without notifying the ACI leaf switch (e.g., via Gratuitous ARP), the detection mechanism depends on how the bridge domain is configured. Specifically, ARP flooding can help detect such silent hosts.

Question: 196

Cisco ACI fabric is integrated with a VMware environment. The engineer must back up the current configuration of the fabric and restore the vCenter password when the configuration is ... Which action accomplishes this goal?

A. Select SCP protocol for the remote location.

B. Create a Configuration Import Policy.

C. Enable the Global AES Encryption setting.

D. Set the Authentication type to Use Password.

Answer: C

Explanation:

By enabling Global AES Encryption and ensuring secure data inclusion, you can safely back up and restore the Cisco ACI fabric configuration, including the vCenter password.

Question: 197

An engineer associates EPG-A with a VMM domain and sets the Deployment and Resolution preferences to Immediate. The host that will generate endpoints for EPG-A is attached to Leaf-and Leaf-102 using etht1/1. However, no configuration for EPG-A appears to have been pushed to the leaf switches. Which action must be taken for the configuration to be pushed to f-101 and Leaf-102?

A. Enable CDP or LLDP on the host.

B. Configure both ports for trunking.

C. Enable LACP on the leaf switch ports.

D. Disable and enable eth1/1 on both leaf switches

Answer: A

Explanation:

The scenario involves associating EPG-A with a Virtual Machine Manager (VMM) domain, setting the

Deployment and Resolution preferences to Immediate, and attaching the host (generating endpoints for EPG-A) to Leaf-101 and Leaf-102 via eth1/1. However, no configuration for EPG-A is pushed to the leaf switches. This suggests an issue with the discovery or communication between the ACI fabric and the host. Let's analyze the options based on Cisco ACI documentation.

Requirement Analysis

The VMM domain integration (e.g., VMware vCenter) relies on protocols like Cisco Discovery Protocol (CDP) or Link Layer Discovery Protocol (LLDP) to detect and map virtualized endpoints to the correct EPG.

The "Immediate" preference ensures that policies are applied as soon as endpoints are detected, but this requires proper protocol support for host discovery.

No configuration on the leaf switches indicates a failure in endpoint detection or policy application.

Option Evaluation

A . Enable CDP or LLDP on the host:

ACI uses CDP or LLDP to discover hosts and their interfaces when integrated with a VMM domain. If the host

does not have CDP or LLDP enabled, the ACI fabric cannot detect the host's attachment to Leaf-101 and Leaf-102, preventing EPG-A configuration from being pushed. Enabling these protocols on the host resolves the issue.

Reference: Cisco APIC VMware Integration Guide, "VMM Domain Integration with CDP/LLDP."

B . Configure both ports for trunking:

Configuring ports as trunk ports allows multiple VLANs, which is necessary for EPG encapsulation. However, this is typically handled by the VMM domain integration and does not address the lack of configuration push if discovery fails due to missing CDP/LLDP.

Reference: Cisco ACI Interface Configuration Guide, "Trunk Port Configuration."

C. Enable LACP on the leaf switch ports:

Link Aggregation Control Protocol (LACP) is used for port channels (e.g., vPC), but it is not required for basic EPG deployment with VMM domains. The issue is related to host discovery, not link aggregation.

Reference: Cisco ACI vPC Configuration Guide.

D . Disable and enable eth1/1 on both leaf switches:

This is a troubleshooting step to reset the port state, but it does not address the root cause (lack of CDP/LLDP for host detection). It is a reactive measure, not a solution.

Reference: Cisco APIC Troubleshooting Guide.

Final Answer Justification

A is correct because enabling CDP or LLDP on the host allows the ACI fabric to detect the host and push the EPG-A configuration to Leaf-101 and Leaf-102. This is a standard requirement for VMM domain integration in ACI.

Primary Cisco Reference:

Cisco APIC Layer 4 to Layer 7 Services Deployment Guide, "VMM Domain Configuration."

Cisco ACI Best Practices, "Host Discovery with CDP/LLDP."

Question: 198

An engineer is implementing an out-of-band (OOB) management access for the Cisco ACI fabric. The secure access must meet these requirements:

- Only GUI and secure shell must be allowed to access the management interfaces of the ACIs.

- The only IP ranges that must be permitted to connect the fabric will be 10.10.10.0724 and 192.168.15.0/24.

Which configuration set meets these requirements?

A. Implement HTTPS and SSH protocol filters in the OOB contract. Add the required subnets to the external network instance profile.

B. Create an out-of-band EPG in the external management entity. Associate the management profile with the OOB contract.

C. Set up static IPs on the management interfaces from the required IP range. Add the required subnets to the external network instance profile.

D. Create an out-of-band EPG in the common tenant. Associate the external network instance profile with the OOB contract.

Answer: A

Explanation:

The engineer is implementing out-of-band (OOB) management access for the Cisco ACI fabric with the following requirements:

Only GUI (HTTPS) and Secure Shell (SSH) must be allowed to access the management interfaces.

Only IP ranges 10.10.10.0/24 and 192.168.15.0/24 must be permitted to connect.

This requires configuring access control and restricting IP ranges for OOB management.

Requirement Analysis

OOB management in ACI is typically handled via the Management Tenant (mgmt) and an OOB contract to define allowed protocols and sources.

The external network instance profile defines the permitted IP ranges for external access.

Option Evaluation

A . Implement HTTPS and SSH protocol filters in the OOB contract. Add the required subnets to the external network instance profile:

An OOB contract can specify allowed protocols (HTTPS on port 443 and SSH on port 22) to restrict access to GUI and SSH only. Adding the subnets 10.10.10.0/24 and 192.168.15.0/24 to the external network instance profile limits the source IP ranges, meeting both requirements.

Reference: Cisco APIC Management Guide, "Out-of-Band Management Configuration" and "Contract Configuration."

B . Create an out-of-band EPG in the external management entity. Associate the management profile with the OOB contract:

This approach creates an EPG for OOB management, but it does not specify protocol filters (HTTPS/SSH) or IP range restrictions. The management profile alone does not enforce these requirements.

Reference: Cisco ACI External Management Configuration Guide.

C . Set up static IPs on the management interfaces from the required IP range. Add the required subnets to the external network instance profile:

Assigning static IPs to management interfaces is a configuration step, but it does not enforce protocol restrictions (HTTPS/SSH) or limit source IP ranges via a contract. This is incomplete.

Reference: Cisco APIC Interface Configuration Guide.

D . Create an out-of-band EPG in the common tenant. Associate the external network instance profile with the OOB contract:

The common tenant can host an OOB EPG, but this option lacks explicit protocol filtering (HTTPS/SSH) and relies on the external network instance profile, which may not fully address the GUI/SSH restriction.

Reference: Cisco ACI Tenant Configuration Guide.

Final Answer Justification

A is correct because it directly addresses both requirements: using an OOB contract to filter HTTPS and SSH protocols and adding the specified subnets to the external network instance profile to restrict IP ranges.

Primary Cisco Reference:

Cisco APIC Management Tenant Configuration Guide, "OOB Management Access."

Cisco ACI Security Guide, "Contract-Based Access Control."

Question: 199

An engineer configures port-12 on Leaf-101 and Leaf-102 to connect to a new server, SVR-12. The new server will belong to EPG-12 and use encap VLAN-1212. The engineer configured SVR-12 as a VPC member port and statically bound the VPC member port to EPG-12. Which additional step must the engineer take to configure connectivity?

A. Create a VPC Explicit Protection Group for EPG-12 and VLAN-1212.

B. Associate a domain with EPG-12 that is associated with VLAN-1212.

C. Select VLAN-1212 on the EPG-12 Interface Policy Group.

D. Configure an LACP Interface Policy and apply it to EPG-12.

Answer: B

Explanation:

The engineer configures port 1/2 on Leaf-101 and Leaf-102 to connect to a new server (SVR-12),

which belongs to EPG-12 using VLAN-1212 encapsulation. The server is configured as a vPC member port and statically bound to EPG-12. The task is to ensure connectivity, indicating a missing configuration step.

Requirement Analysis

Static binding of a vPC member port to an EPG requires proper VLAN association and domain mapping.

vPC ensures redundancy, and the VLAN (1212) must be allocated and associated with the EPG via a domain (e.g., VMM or physical domain).

Option Evaluation

A . Create a VPC Explicit Protection Group for EPG-12 and VLAN-1212:

vPC Explicit Protection Groups are used for specific vPC configurations, but this is not a standard requirement for EPG binding and VLAN association.

Reference: Cisco ACI vPC Configuration Guide.

B . Associate a domain with EPG-12 that is associated with VLAN-1212:

Associating a domain (e.g., VMM or physical domain) with EPG-12 and mapping VLAN-1212 ensures the VLAN is allocated and recognized by the fabric for the vPC ports. This is a critical step for static binding to work.

Reference: Cisco APIC EPG Configuration Guide, "Domain and VLAN Association."

C . Select VLAN-1212 on the EPG-12 Interface Policy Group:

The Interface Policy Group defines port settings, but VLAN selection is managed via the domain association,

not the policy group alone. This is insufficient.

Reference: Cisco ACI Interface Policy Configuration Guide.

D . Configure an LACP Interface Policy and apply it to EPG-12:

LACP is optional for vPC and not required for static EPG binding. The issue is VLAN/domain association, not link aggregation.

Reference: Cisco ACI LACP Configuration Guide.

Final Answer Justification

B is correct because associating a domain with EPG-12 and VLAN-1212 ensures the fabric recognizes the VLAN encapsulation and applies the EPG configuration to the vPC ports.

Primary Cisco Reference:

Cisco APIC vPC Deployment Guide.

Cisco ACI EPG and Domain Configuration Guide.

Question: 200

Refer to the exhibit.

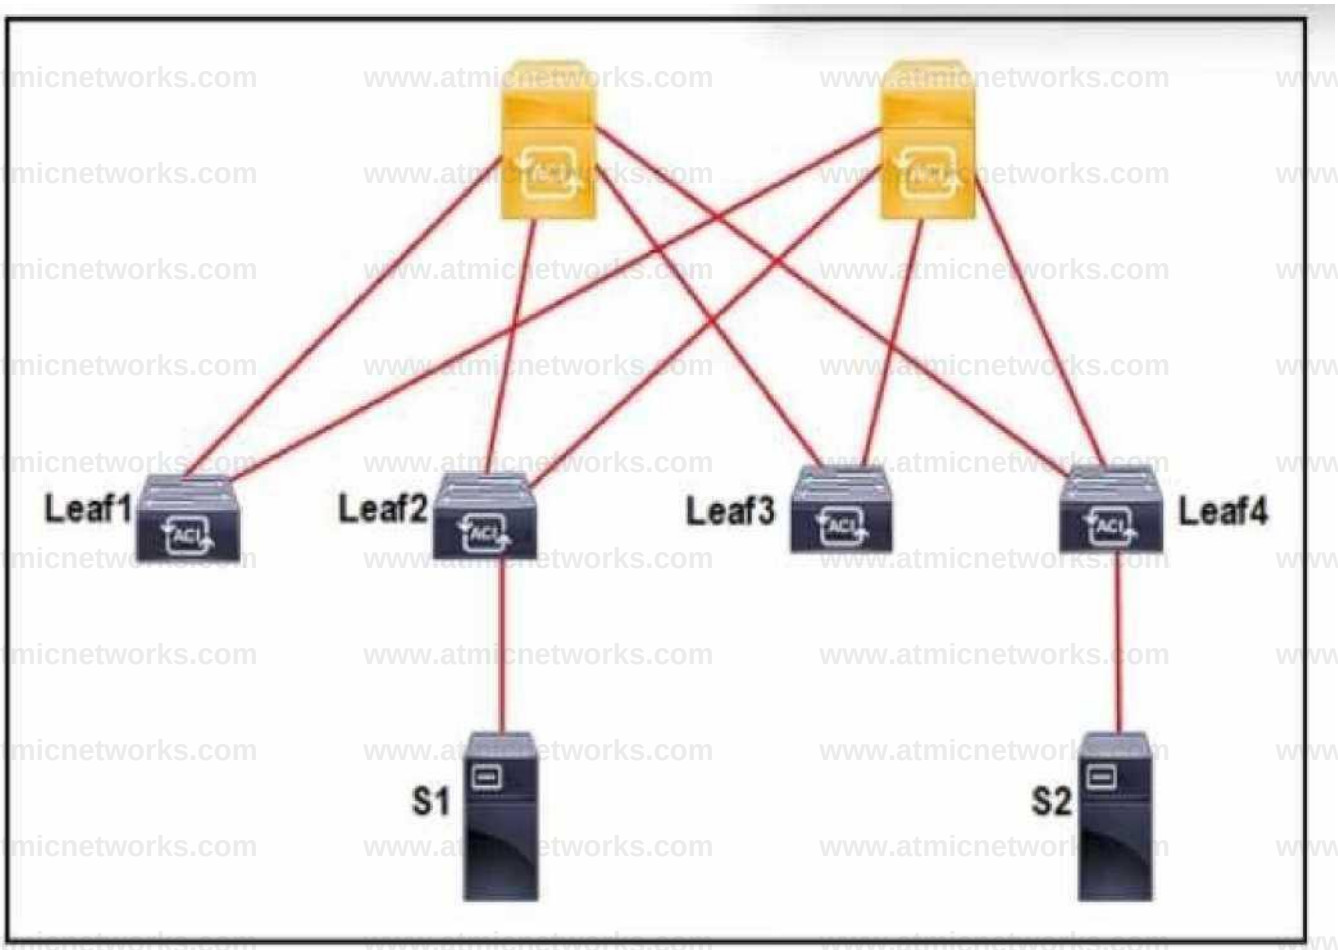

An application called App_1 is hosted on the server called S1. A silent host application. App_2. is hosted on S2. Both applications use the same VLAN encapsulation, which action forces Cisco ACI fabric to learn App_2 on ACI leaf 2?

A. Set Multi-Destination Flooding to Drop.

B. Set Unicast Routing to Hardware Proxy.

C. Set L2 Unknown Unicast to Flood.

D. Set L3 Unknown Multicast to Optimized flood.

Answer: C

Explanation:

The scenario involves an application (App_1) on server S1 and a silent host application (App_2) on S2, both using the same VLAN encapsulation. The task is to force the ACI fabric to learn App_2 on Leaf-2. A silent host does not generate traffic, so special handling is needed.

Requirement Analysis

A silent host requires flooding (e.g., ARP or unknown unicast) to be learned by the fabric when it moves or is detected.

The goal is to ensure Leaf-2 learns App_2's endpoint.

Option Evaluation

A . Set Multi-Destination Flooding to Drop:

Dropping multi-destination traffic prevents learning, which is counterproductive for a silent host.

Reference: Cisco ACI Bridge Domain Configuration Guide.

B . Set Unicast Routing to Hardware Proxy:

Hardware proxy optimizes unicast routing but does not force learning of a silent host via flooding.

Reference: Cisco ACI Routing Configuration Guide.

C . Set L2 Unknown Unicast to Flood:

Enabling L2 unknown unicast flooding causes the fabric to flood traffic (e.g., ARP) across the bridge domain, allowing Leaf-2 to learn App_2's MAC address even if it is silent.

Reference: Cisco APIC Bridge Domain Settings, "L2 Unknown Unicast Flooding."

D . Set L3 Unknown Multicast to Optimized Flood:

This applies to multicast traffic and is irrelevant for learning a silent host's MAC address.

Reference: Cisco ACI Multicast Configuration Guide.

Final Answer Justification

C is correct because flooding L2 unknown unicast traffic ensures the ACI fabric learns App_2 on Leaf-2 by propagating ARP or other discovery traffic.

Primary Cisco Reference:

Cisco ACI Endpoint Learning Guide, "Flooding for Silent Hosts."

Cisco APIC Bridge Domain Configuration Guide.

Question: 201

An engineer configures an L30ut in VRF-1 that was configured for Import Route Control Enforcement. The L30ut uses OSPF to peer with a core switch. The L30ut has one external EPG, it has been configured with a subnet 10.1.0.0/24. Which scope must be set to force 10.1.0.0/24 to populate in the routing table for VRF-1?

A. External Subnet for External EPG

B. Export Route Control Subnet

C. Shared Route for External EPG

D. Import Route Control Subnet

Answer: D

Explanation:

The "Import Route Control Subnet" scope is used to control which external routes are imported into the ACI fabric's routing table

Question: 202

What is a characteristic of a Cisco ACI Multi-Pod?

A. It eliminates the need to deploy multicast in the Layer 3 network that interconnects the pods.

B. Spines use BGP peering with IPN to send out the TEP pool prefix for the local pod.

C. It manages the configuration of different Cisco ACI pods using a single common Cisco APIC cluster.

D. A VPNv4 address family is used to exchange endpoint information between spines.

Answer: C

Explanation:

In a Cisco ACI Multi-Pod architecture, multiple pods (each with its own leaf-and-spine topology) are interconnected via an Inter-Pod Network (IPN). The key characteristic of the Multi-Pod setup is that it is managed as a single fabric by a single APIC cluster, simplifying operations and maintaining consistency across all pods.

Question: 203

An engineer discovered an outage on the mgmt0 port of Leaf113 and Leaf114. Both leaf switches were recently registered in the fabric and have health scores of 100. The engineer overs there is no IP address assigned to the mgmt0 interface of the switches. Which action resolves the outage?

A. Statically bind the mgmt0 interface of Leaf113 and Leaf114 to the oob-default EPG.

B. Enable Leaf 113 and Leaf 114 mgmt0 under the leaf switch.

C. Associate the oobbrc-default contract to Leaf113 and Leaf114.

D. Add Leaf113 and Leaf114 to the node management address policy.

Answer: D

Explanation:

In Cisco ACI, the mgmt0 interface is used for out-of-band (OOB) management, and its IP address must be explicitly assigned through the Node Management Address Policy under the mgmt tenant. When new leaf switches are registered in the fabric, they do not automatically receive an IP address for their mgmt0 interfaces unless they are added to this policy.

Question: 204

When Layer 3 routed traffic is destined to a Cisco ACI fabric, which mechanism does ACI use to detect silent hosts?

A. gratuitous ARP

B. inverse ARP

C. ARP gleaning

D. proxy ARP

Answer: C

Explanation:

The question asks about the mechanism Cisco ACI uses to detect silent hosts when Layer 3 routed traffic is destined to the ACI fabric. A "silent host" refers to an endpoint that does not initiate traffic or send ARP requests, making detection challenging without specific mechanisms. Let's evaluate the options based on Cisco ACI documentation and best practices.

Requirement Analysis

Layer 3 routed traffic implies that the ACI fabric is handling IP routing, and the detection mechanism must work within the context of the bridge domain and endpoint learning.

Silent hosts require passive detection methods since they do not generate active traffic (e.g., ARP requests).

The solution must align with ACI's endpoint learning and proxy mechanisms.

Option Evaluation

A . Gratuitous ARP:

Gratuitous ARP (GARP) is a mechanism where a host announces its IP-to-MAC mapping to update network devices. However, this requires the host to send the GARP, which a silent host does not do. ACI can use GARP when hosts are active, but it is not the primary method for silent hosts.

Reference: Cisco ACI Endpoint Learning Guide, "Gratuitous ARP Handling."

B . Inverse ARP:

Inverse ARP is used in Frame Relay or ATM networks to map DLCI/VPI-VCI to IP addresses. It is not applicable to ACI's IP-based Layer 3 routing or silent host detection.

Reference: Not relevant to ACI documentation.

C . ARP Gleaning:

ARP gleaning is a passive mechanism in ACI where the fabric learns the IP-to-MAC binding of a silent host by inspecting ARP traffic destined to or from other devices (e.g., routers or proxies) within the same bridge domain. When Layer 3 routed traffic passes through the ACI fabric, the leaf switches can glean the silent host's information from ARP responses or related traffic, even if the host itself is silent.

Reference: Cisco APIC Bridge Domain Configuration Guide, "ARP Gleaning for Silent Hosts" and Cisco ACI Endpoint Learning White Paper.

D . Proxy ARP:

Proxy ARP is used by ACI to respond to ARP requests on behalf of silent hosts when the hardware proxy feature is enabled. While this helps with connectivity, it is a response mechanism, not a detection method. Detection still relies on gleaning or other passive learning.

Reference: Cisco ACI Routing Configuration Guide, "Proxy ARP."

Final Answer Justification

C is correct because ARP gleaning is the specific mechanism ACI uses to detect silent hosts when Layer 3 routed traffic is involved. The fabric passively learns the host's IP and MAC addresses from ARP traffic generated by other devices (e.g., routers or active hosts) within the bridge domain, which is particularly effective in routed environments.

Primary Cisco Reference:

Cisco APIC Layer 3 Networking Configuration Guide, "Endpoint Learning with ARP Gleaning."

Cisco ACI Best Practices, "Silent Host Detection in Routed Networks."

Question: 205

Which switch type is discovered first in the Cisco ACI fabric discovery process?

A. leaf

B. access

C. distribution

D. spine

Answer: A

Explanation:

The question asks which switch type is discovered first in the Cisco ACI fabric discovery process. The ACI fabric consists of spine and leaf switches, and the discovery process is a critical initial step to establish the fabric topology. Let's analyze this based on Cisco ACI documentation.

Requirement Analysis

The ACI fabric discovery process involves the Application Policy Infrastructure Controller (APIC) identifying and registering switches to form the network topology.

The process relies on protocols like Link Layer Discovery Protocol (LLDP) and Cisco Discovery Protocol (CDP) to detect connected devices.

The sequence of discovery is determined by the physical connectivity and the role of switches in the fabric.

Option Evaluation

A . leaf:

In the Cisco ACI fabric, the discovery process begins with the leaf switches. When the APIC is powered on and connected to the fabric, it first discovers the leaf switches directly attached to it or other initial points of connectivity. Leaf switches are the access layer devices that connect endpoints and are the starting point for fabric topology discovery. Once leaf switches are identified, they assist in discovering the spine switches.

Reference: Cisco APIC Installation and Setup Guide, "Fabric Discovery Process" and Cisco ACI Fabric Fundamentals, "Switch Discovery."

B . access:

"Access" is not a specific switch type in the ACI fabric context. It may refer to external access switches, but these are not part of the initial fabric discovery process, which focuses on spine and leaf switches.

Reference: Not applicable to ACI fabric discovery.

C . distribution:

"Distribution" is a traditional network layer, but it does not apply to the ACI fabric, which uses a spine-leaf architecture. This option is irrelevant.

Reference: Not applicable to ACI.

D . spine:

Spine switches are discovered after the leaf switches because the leaf switches provide the connectivity path

to the spines. The APIC uses the leaf switches to map the spine layer in the fabric topology.

Reference: Cisco ACI Architecture White Paper, "Spine-Leaf Discovery Sequence."

Final Answer Justification

A is correct because the ACI fabric discovery process starts with the leaf switches. The APIC initially detects leaf switches connected to it or other initial points, and then uses this information to discover the spine switches, establishing the full fabric topology.

Primary Cisco Reference:

Cisco APIC Getting Started Guide, "Fabric Bring-Up and Discovery."

Cisco ACI Design Guide, "Fabric Discovery Process."

Question: 206

An Cisco ACI leaf switch learns the source IP address of a packet that enters the front panel port of the switch. Which bridge domain setting is used?

A. Unicast Routing

B. L3 Unknown Multicast Flooding - Flood

C. ARP Flooding

D. Unknown Unicast - Hardware proxy

Answer: A

Explanation:

The question asks which bridge domain setting is used when a Cisco ACI leaf switch learns the source IP address of a packet entering the front panel port. This involves understanding how ACI handles endpoint learning and IP address association within a bridge domain.

Requirement Analysis

When a packet enters a leaf switch's front panel port, ACI learns the source IP and MAC address of the endpoint to populate its endpoint table.

The bridge domain settings control how IP addresses are learned and routed, especially for Layer 3 traffic.

The correct setting must enable the leaf switch to associate the source IP with the endpoint.

Option Evaluation

A . Unicast Routing:

The "Unicast Routing" setting in a bridge domain enables Layer 3 routing and allows the leaf switch to learn the source IP address of a packet by associating it with the MAC address and VLAN. When enabled, the switch performs IP-to-MAC mapping and updates the endpoint database, which is the standard mechanism for learning source IP addresses from incoming traffic.

Reference: Cisco APIC Bridge Domain Configuration Guide, "Unicast Routing Enablement" and Cisco ACI Endpoint Learning White Paper.

B . L3 Unknown Multicast Flooding - Flood:

This setting controls how unknown multicast traffic is handled at Layer 3 (e.g., flooding or dropping). It is unrelated to learning the source IP address of a unicast packet entering the front panel port.

Reference: Cisco ACI Multicast Configuration Guide.

C . ARP Flooding:

ARP Flooding allows ARP requests to be flooded within the bridge domain, which helps resolve IP-to- MAC mappings for silent hosts or external devices. However, it is not the primary setting for learning the source IP of an active packet; it is a supplementary mechanism.

Reference: Cisco APIC Bridge Domain Settings, "ARP Flooding Configuration."

E. Unknown Unicast - Hardware Proxy:

The "Hardware Proxy" setting optimizes unicast traffic by using a proxy to respond to ARP requests, reducing flooding. While it aids in endpoint management, it is not the setting that directly enables learning the source IP address from incoming packets.

Reference: Cisco ACI Routing Configuration Guide, "Hardware Proxy."

Final Answer Justification

A is correct because enabling "Unicast Routing" in the bridge domain allows the leaf switch to learn the source IP address of a packet entering the front panel port by associating it with the source MAC address and VLAN. This is the foundational setting for Layer 3 endpoint learning in ACI.

Primary Cisco Reference:

Cisco APIC Layer 3 Networking Configuration Guide, "Bridge Domain Unicast Routing."

Cisco ACI Endpoint Group Configuration Guide, "IP Address Learning."

Question: 207

Refer to the exhibit.

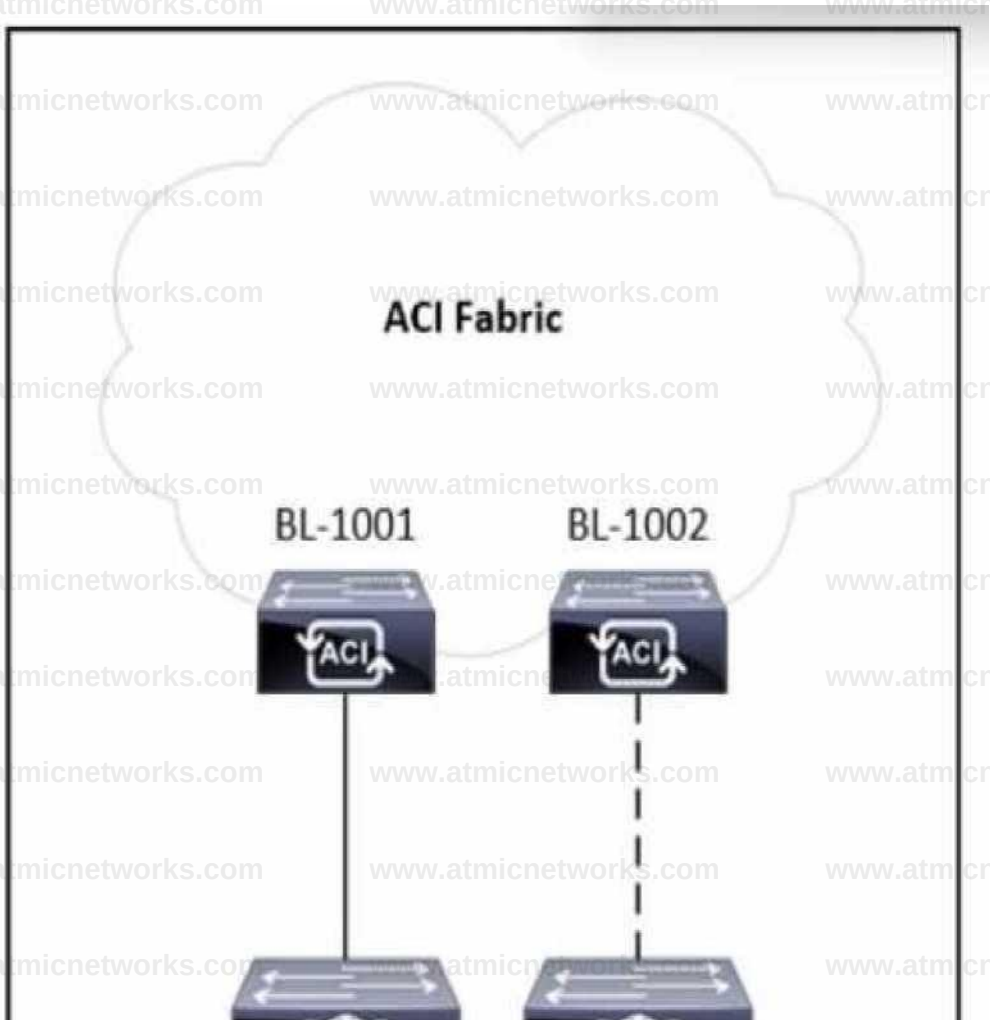

Refer to the exhibit. A tenant is configured with a single L30ut and a single-homed link to the core router called Core-1. An engineer must add a second link to the L30ut that connects to Core-2 router. Which action allows the traffic from Core-2 to BL-1002 to have the same connectivity as the traffic from Core-1 to BL-1001?

A. Add a second path to the logical interface profile of the existing L30ut

B. Add a second subnet to the external EPG to the existing L30ut.

C. Add a second OSPF interface profile to the logical interface profile.

D. Add a second interface to the external domain to the existing L30ut.

Answer: A

Explanation:

The scenario involves a tenant configured with a single L3Out and a single-homed link from the ACI fabric (via border leaf BL-1001) to a core router (Core-1). The engineer must add a second link to the L3Out, connecting to Core-2 via BL-1002, ensuring that traffic from Core-2 to BL-1002 has the same connectivity as traffic from Core-1 to BL-1001. The exhibit shows a single-homed setup with BL-1001 and BL-1002 as potential border leaves, with a dashed line indicating a planned second link.

Requirement Analysis

The existing L3Out is single-homed to Core-1 via BL-1001, and the goal is to extend it to a multihomed configuration with Core-2 via BL-1002.

"Same connectivity" implies that the second link must be integrated into the existing L3Out configuration, sharing the same routing policies, external EPG, and subnet reachability.

The solution must leverage ACI's L3Out framework to add the new path without creating a separate L3Out or altering the current routing setup unnecessarily.

Option Evaluation

A . Add a second path to the logical interface profile of the existing L3Out:

The logical interface profile in an L3Out defines the interfaces (e.g., routed ports or sub-interfaces) and their associations with nodes (border leaves). Adding a second path to this profile allows the inclusion of the new link from BL-1002 to Core-2, ensuring it operates under the same L3Out configuration (e.g., routing protocol, external EPG). This maintains consistent connectivity and leverages ACI's multi-homing support.

Reference: Cisco APIC Layer 3 Configuration Guide, "Configuring Logical Interface Profiles for L3Out" and "Multi-Homing L3Out."

B . Add a second subnet to the external EPG to the existing L3Out:

Adding a subnet to the external EPG defines additional networks advertised or learned via the L3Out, but it does not address the addition of a new physical link or path. This is irrelevant to the connectivity requirement for the second link.

Reference: Cisco ACI External EPG Configuration Guide.

C . Add a second OSPF interface profile to the logical interface profile:

While the L3Out uses OSPF (implied by the context), adding a second OSPF interface profile would create a separate routing instance, which is unnecessary and could disrupt consistency. The existing

OSPF configuration can be extended via the logical interface profile.

Reference: Cisco APIC OSPF Configuration Guide.

D . Add a second interface to the external domain to the existing L3Out:

The external domain associates VLAN pools, but it does not define paths or interfaces for L3Out connectivity. This option is incorrect as it confuses domain configuration with interface path management.

Reference: Cisco ACI Domain Configuration Guide.

Final Answer Justification

A is correct because adding a second path to the logical interface profile of the existing L3Out integrates the new link from BL-1002 to Core-2 into the same routing and policy framework, ensuring equivalent connectivity to the Core-1 to BL-1001 link. This aligns with ACI's support for multi-homed L3Outs.

Primary Cisco Reference:

Cisco APIC Layer 3 Networking Configuration Guide, "L3Out Multi-Homing."

Cisco ACI Best Practices, "Expanding L3Out Connectivity."

Question: 208

An engineer wants to configure Cisco ACI switches to use authenticated ZMQ when communicating with the proxy spine. Which configuration allows MD5 ZMQ messages only?

A. IS-IS password using MD5

B. COOP Group policy in strict mode

C. COOP Group policy in compatible mode

D. BGP password using MD5

Answer: B

Explanation:

https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/security-configuration/cisco-apic-security-configuration-guide-release-52x/protocol-authentication-52x.html

Question: 209

Which feature is used to program policy CAM on a leaf switch without sending traffic from VM to the leaf?

A. immediate resolution immediacy

B. immediate deployment immediacy

C. on-demand deployment immediacy

D. on-demand resolution immediacy

Answer: B

Explanation:

Question: 210

Refer to the exhibit.

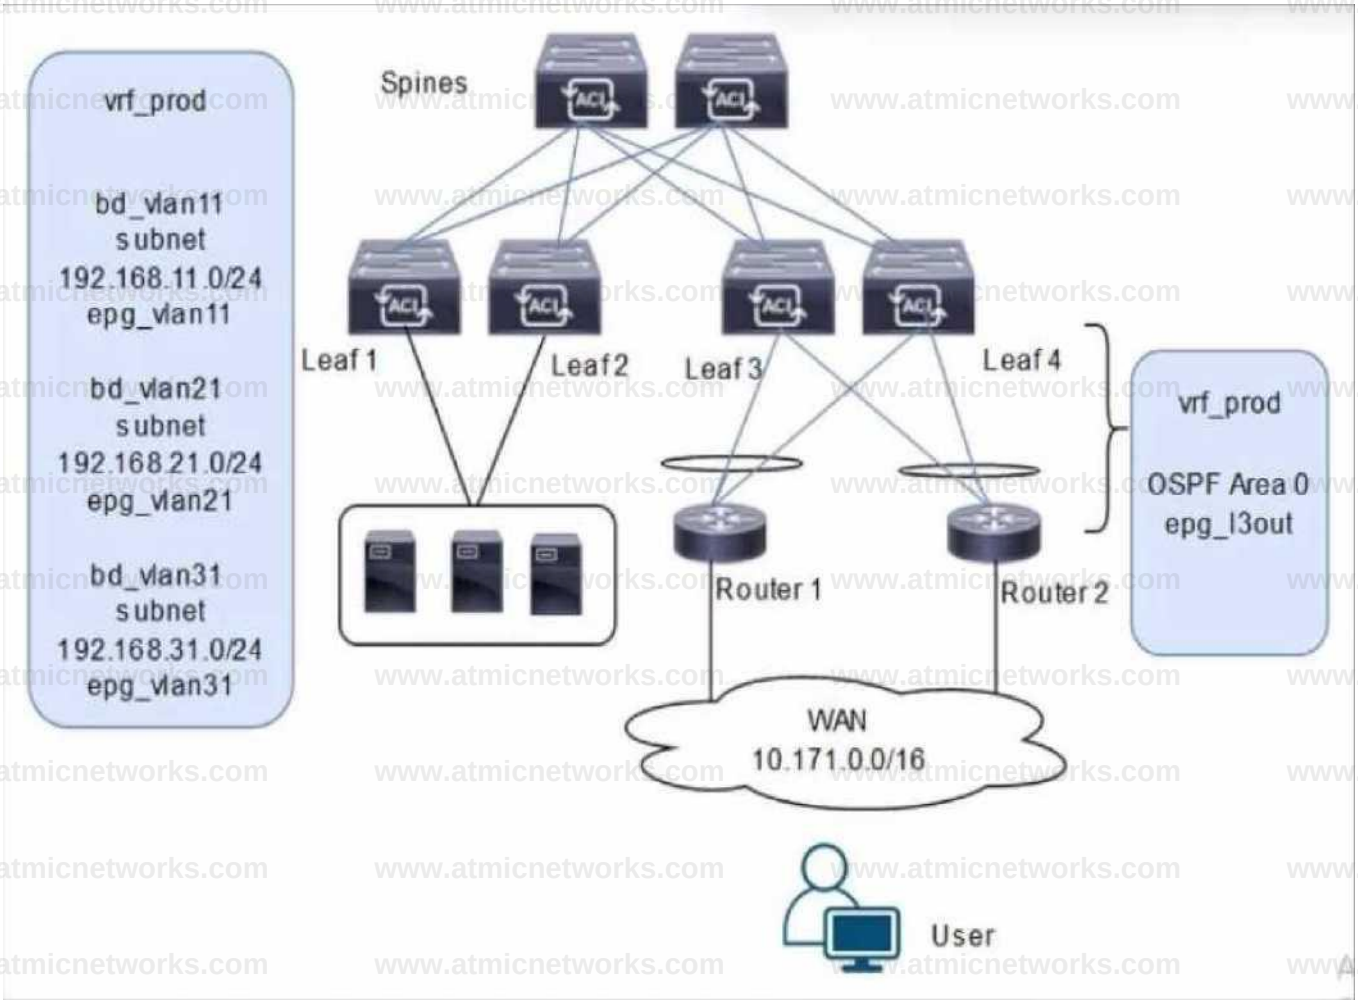

A customer is deploying a WAN with these requirements: •Routers 1 and 2 must receive only routes 192.168.11.0/24 and 192.168.21.0724 from the Cisco ACI fabric •Reachability to the WAN users must be permitted only for the servers that are located in vrf_prod.

Which settings must be configured to meet these objectives?

A. Configure the subnets 192.168.11.0/24 and 192.168.21.0/24 as Private to VRF Configure the subnet 192.168.31.0/24 as Advertised Externally. Configure an EPG subnet 0.0.0.0/0 as External Subnets for External EPG.

B. Configure the subnets 192.168.11.0/24 and 192.168.21.0/24 as Private to VRF. Configure the subnet 192.168.31.0/24 as Advertised Externally. Configure an EPG subnet 0.0.0.0/0 as Shared Route Control Subnet.

C. Configure the subnets 192.168.11.0/24 and 192.168.21.0/24 as Advertised Externally.

Configure the subnet 192.168.31.0/24 as Private to VRF.

Configure an EPG subnet 0.0.0.0/0 as Shared Route Control Subnet.

D. Configure the subnets 192.168.11.0/24 and 192.168.21.0/24 as Advertised Externally.

Configure the subnet 192.168.31.0/24 as Private to VRF.

Configure an EPG subnet 0.0.0.0/0 as External Subnets for External EPG.

Answer: D

Explanation:

The scenario involves deploying a WAN with Cisco ACI, where Routers 1 and 2 (connected via an L3Out with OSPF Area 0) must receive specific routes (192.168.11.0/24 and 192.168.21.0/24) from the ACI fabric, and reachability to WAN users must be permitted only for servers in vrf_prod. The diagram shows three bridge domains (bd_vlan11, bd_vlan21, bd_vlan31) with their respective subnets and EPGs, all under vrf_prod, along with an L3Out (epg_l3out) for WAN connectivity.

Requirement Analysis

Routers 1 and 2 must receive only routes 192.168.11.0/24 and 192.168.21.0/24:

These subnets belong to bd_vlan11 and bd_vlan21, respectively. To advertise these routes to Routers 1 and 2 via the L3Out, they must be marked with the appropriate scope in the bridge domain configuration.

In ACI, the "Advertised Externally" scope on a subnet ensures that it is advertised to external routers via the L3Out routing protocol (OSPF in this case).

Reachability to WAN users must be permitted only for servers in vrf_prod:

This implies that only the subnets in vrf_prod (192.168.11.0/24, 192.168.21.0/24, and 192.168.31.0/24) should be accessible, but WAN users should only reach specific subnets based on policy.

The external EPG (epg_l3out) represents the WAN users (10.171.0.0/16), and its subnet scope must control inbound reachability.

The subnet 192.168.31.0/24 (bd_vlan31) should not be advertised to the WAN, as it is not listed in the routes Routers 1 and 2 should receive.

Option Evaluation

A . Configure the subnets 192.168.11.0/24 and 192.168.21.0/24 as Private to VRF. Configure the subnet 192.168.31.0/24 as Advertised Externally. Configure an EPG subnet 0.0.0.0/0 as External Subnets for External EPG:

Setting 192.168.11.0/24 and 192.168.21.0/24 as "Private to VRF" means they are not advertised externally, which fails the requirement for Routers 1 and 2 to receive these routes.

Setting 192.168.31.0/24 as "Advertised Externally" incorrectly advertises this subnet to the WAN, which is

not desired.

The "External Subnets for External EPG" scope on 0.0.0.0/0 allows WAN users to reach all subnets in vrf_prod, which is correct for reachability.

Conclusion: Fails the first requirement (route advertisement).

Reference: Cisco APIC Layer 3 Configuration Guide, "Subnet Scope Configuration."

B . Configure the subnets 192.168.11.0/24 and 192.168.21.0/24 as Private to VRF. Configure the subnet 192.168.31.0/24 as Advertised Externally. Configure an EPG subnet 0.0.0.0/0 as Shared Route Control Subnet:

Similar to Option A, setting 192.168.11.0/24 and 192.168.21.0/24 as "Private to VRF" prevents their advertisement to the WAN, failing the first requirement.

Setting 192.168.31.0/24 as "Advertised Externally" incorrectly advertises this subnet.

The "Shared Route Control Subnet" scope allows route leaking between VRFs, which is irrelevant here since there is only one VRF (vrf_prod) and no route leaking is required.

Conclusion: Fails both requirements (route advertisement and reachability control).

Reference: Cisco ACI VRF Configuration Guide, "Route Leaking with Shared Subnets."

C . Configure the subnets 192.168.11.0/24 and 192.168.21.0/24 as Advertised Externally. Configure the subnet 192.168.31.0/24 as Private to VRF. Configure an EPG subnet 0.0.0.0/0 as Shared Route Control Subnet:

Setting 192.168.11.0/24 and 192.168.21.0/24 as "Advertised Externally" ensures these subnets are advertised to Routers 1 and 2 via OSPF, meeting the first requirement.

Setting 192.168.31.0/24 as "Private to VRF" prevents its advertisement to the WAN, which aligns with the requirement since only 192.168.11.0/24 and 192.168.21.0/24 should be advertised.

The "Shared Route Control Subnet" scope on 0.0.0.0/0 in the external EPG is incorrect for controlling reachability. This scope is used for route leaking, not for defining which subnets are accessible from the external EPG.

Conclusion: Meets the first requirement but fails the second (reachability control).

Reference: Cisco ACI External EPG Configuration Guide, "Subnet Scopes."

D . Configure the subnets 192.168.11.0/24 and 192.168.21.0/24 as Advertised Externally. Configure the subnet 192.168.31.0/24 as Private to VRF. Configure an EPG subnet 0.0.0.0/0 as External Subnets for External EPG:

Setting 192.168.11.0/24 and 192.168.21.0/24 as "Advertised Externally" ensures these subnets are advertised to Routers 1 and 2 via OSPF, meeting the first requirement.

Setting 192.168.31.0/24 as "Private to VRF" prevents its advertisement to the WAN, which is correct since

only the specified subnets should be advertised.

The "External Subnets for External EPG" scope on 0.0.0.0/0 in the external EPG (epg_l3out) allows WAN users (10.171.0.0/16) to reach all subnets in vrf_prod, which includes 192.168.11.0/24, 192.168.21.0/24, and 192.168.31.0/24. This satisfies the second requirement, as servers in vrf_prod are accessible, and contracts can further restrict access if needed (though not specified in the question).

Conclusion: Meets both requirements (route advertisement and reachability).

Reference: Cisco APIC Layer 3 Networking Configuration Guide, "L3Out Subnet Scopes" and "External EPG Configuration."

Final Answer Justification

D is correct because:

It ensures that only 192.168.11.0/24 and 192.168.21.0/24 are advertised to Routers 1 and 2 by setting their scope to "Advertised Externally."

It keeps 192.168.31.0/24 private to vrf_prod by setting its scope to "Private to VRF."

It allows WAN users to reach all vrf_prod subnets (including servers) by setting 0.0.0.0/0 as "External Subnets for External EPG," fulfilling the reachability requirement.

Primary Cisco Reference:

Cisco APIC Layer 3 Configuration Guide, "Configuring Subnets for L3Out."

Cisco ACI Routing and Forwarding Guide, "External EPG and Subnet Scopes."
Cisco ACI Best Practices, "Controlling Route Advertisement and Reachability."

Question: 211

Cisco ACI fabric has three different endpoints S1, S2. and S3. These endpoints must communicate with each other without contracts. These objects have been created in APIC:

- Two EPGs named DNS_EPG and Database_EPG
- Two application profiles. PROD_App and Data_App
- Two bridge domains DNS_BD and Database_BD
- PROD_APP and Database_BD mapped to Tenant PROD
- Data_App and DNS_BD mapped to Tenant Data

Which set of actions completes the fabric configuration?

A. Add S1, S2, S3 under Database_EPG.

MAP Database_EPG under PROD_ App.

Associate Datbase_EPG with DNS_BD.

B. Add S1, S2, S3, under DNS_EPG.

MAP DNS_EPG to Data_App.

Associate DNS_EPG with Dns_BD.

C. Add S1, S2, S3 under DNS_EPG.

MAP DNS_EPG to Data_App.

Associate DNS_EPG with Database_BD.

D. Add S1, S2, S3 under Database_EPG.

MAP Database_EPG under Data_App.

Associate Datbase_EPG with Database_BD.

Answer: B

Explanation:

Question: 212

A company is implementing a new security policy to track system access, configuration, and changes. The network engineer must enable the log collection to track user login and logout attempts. In addition, any configuration changes such as a fabric node failure must be collected in the logs. The syslog policy is configured to send logs to the company SEIM appliance.

Which two log types must be enabled to meet the security requirements? (Choose two.)

A. error

B. audit

C. event

D. health

E. fault

Answer: BE

Explanation:

Question: 213

Cisco ACI fabric contains 10 standalone leaf switches. An engineer must configure only the first two leaf switches in a VPC. Which VPC protection type must be configured to accomplish goal?

A. serial

B. explicit

C. reciprocal

D. consecutive

Answer: B

Explanation:

The scenario involves a Cisco ACI fabric with 10 standalone leaf switches, and the engineer must configure only the first two leaf switches (e.g., Leaf-1 and Leaf-2) in a Virtual Port Channel (vPC) configuration. The goal is to select the appropriate vPC protection type to achieve this specific
pairing while leaving the other eight leaf switches standalone.

Requirement Analysis

In Cisco ACI, vPC is used to provide active-active link redundancy and load balancing between two leaf switches connected to the same set of endpoints (e.g., servers or external devices).

The fabric contains 10 standalone leaf switches, meaning no pre-existing vPC pairs are configured, and the task is to form a vPC specifically between the first two leaf switches.

The vPC protection type determines how the pair is defined and protected within the ACI fabric, ensuring

the configuration is limited to the intended switches.

Option Evaluation

A . serial:

The "serial" protection type is not a valid vPC protection option in Cisco ACI. This term might be confused with serial link configurations or other networking contexts, but it does not apply to vPC in ACI.

Reference: Cisco ACI vPC Configuration Guide (no mention of "serial" protection).

B . explicit:

The "explicit" vPC protection type allows the engineer to manually specify the exact pair of leaf switches to form a vPC. In this case, the engineer can configure Leaf-1 and Leaf-2 as an explicit vPC pair, leaving the other eight leaf switches standalone. This method provides precise control over which switches are paired, aligning with the requirement to configure only the first two.

Reference: Cisco APIC vPC Deployment Guide, "Explicit vPC Protection Group Configuration" and Cisco ACI Best Practices, "vPC Pairing."

C . reciprocal:

The "reciprocal" protection type is not a standard vPC protection option in ACI. This term might imply mutual protection, but it is not documented as a specific vPC configuration mode in ACI.

Reference: No reference in Cisco ACI vPC documentation.

D . consecutive:

The "consecutive" protection type is not a recognized vPC protection option in ACI. It might suggest pairing switches based on their IDs (e.g., consecutive node IDs), but ACI does not use this as a formal protection mechanism. The fabric could interpret this incorrectly, potentially pairing unintended switches.

Reference: No reference in Cisco ACI vPC Configuration Guide.

Final Answer Justification

B is correct because the "explicit" vPC protection type allows the engineer to manually designate Leaf-1 and Leaf-2 as a vPC pair, ensuring that only these two switches are configured in a vPC while the remaining eight leaf switches remain standalone. This method provides the precision needed to meet the goal and is a standard feature in ACI vPC deployments.

Primary Cisco Reference:

Cisco APIC Layer 2 Networking Configuration Guide, "Configuring vPC with Explicit Protection."

Cisco ACI Design Guide, "vPC Deployment Options."

Cisco ACI vPC Best Practices, "Explicit vPC Pair Configuration."

Question: 214

Refer to the exhibit.

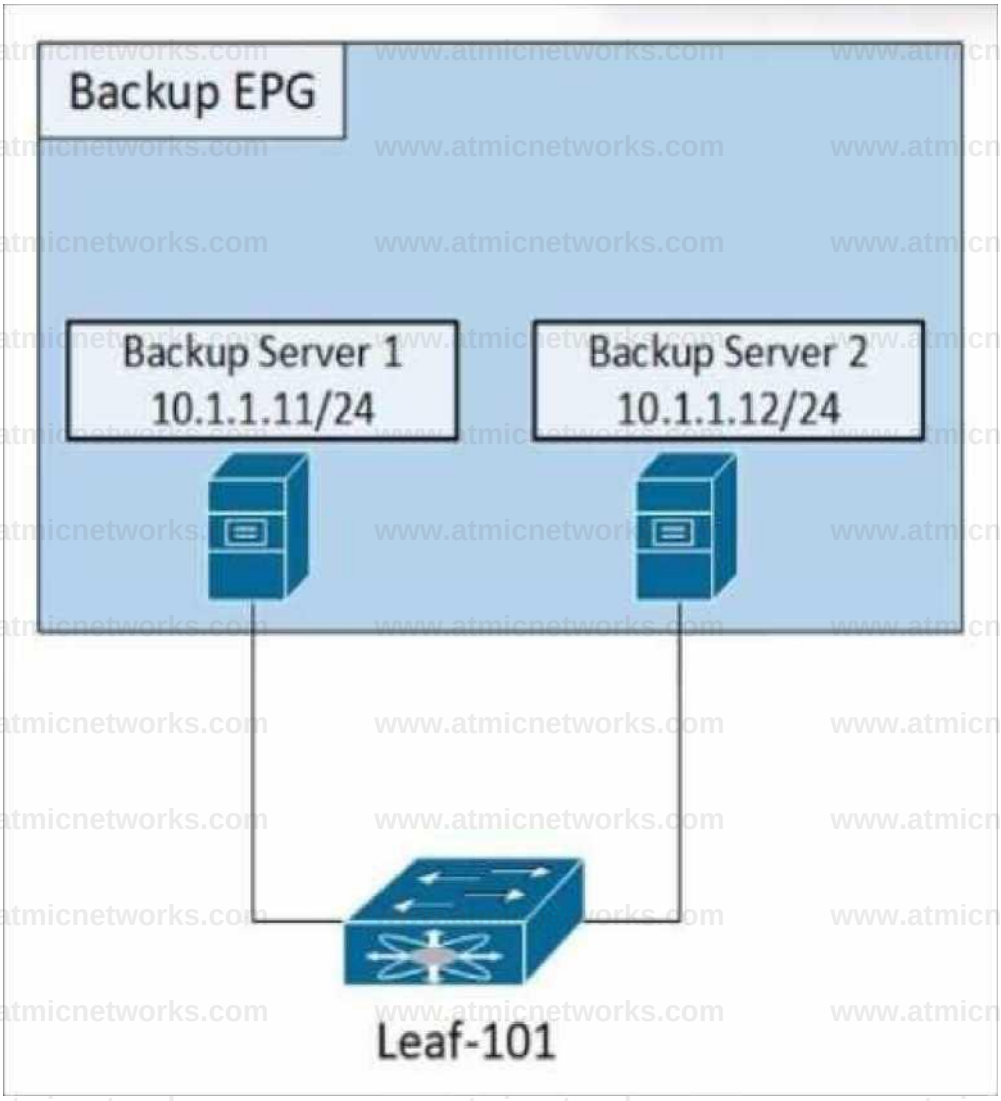

Refer to the exhibit. An engineer must disable the communication between the two backup servers in the backup EPG. Which action accomplishes this goal?

A. Set Preferred Group Member to Excluded.

B. Set the physical domain to None.

C. Set a different static binding for the encap VLAN.

D. Set Intra EPG Isolation to Enforced.

Answer: D

Explanation:

In Cisco ACI, Intra-EPG Isolation is a feature that prevents communication between endpoints within the same EPG. By default, endpoints in the same EPG can communicate freely without requiring contracts. To disable communication between two backup servers within the same EPG (e.g., in the "Backup EPG"), you need to enforce Intra-EPG Isolation.

Question: 215

Which Cisco ACI setting corresponds to the VMware MAC pinning?

A. route based on IP hash

B. route based on originating virtual port

C. route based on physical NIC load

D. route based on MAC hash

Answer: B

Explanation:

In VMware vSphere, "MAC Pinning" refers to the "Route based on originating virtual port" loadbalancing algorithm. This method assigns each virtual machine (VM) vNIC to a specific uplink (pNIC) and ensures that all traffic from that VM consistently exits through the same uplink.

Question: 216

Refer to the exhibit.

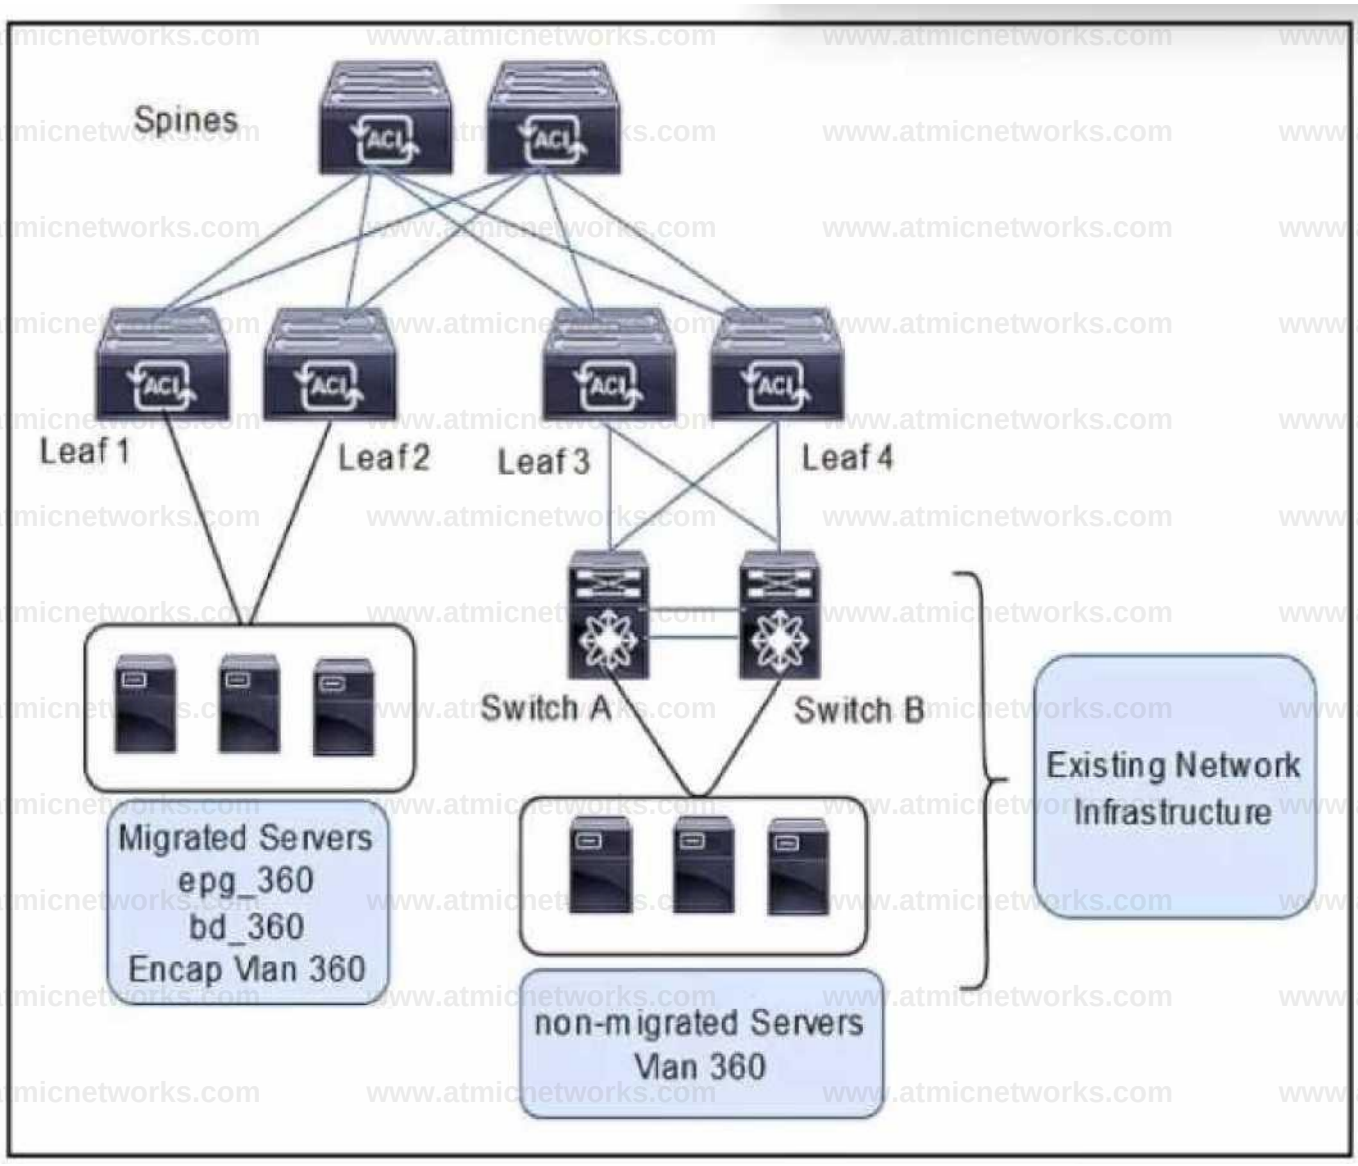

Refer to the exhibit. An engineer is migrating legacy servers into the Cisco ACI environment. The requirement is to ensure that all endpoints and MAC addresses are learned properly in legacy and Cisco ACI switches. Which configuration set must be configured under the bridge domain called bd 360 to accomplish this goal?

A. L2 Unknown Unicast: Hardware Proxy ARP Flooding: Disabled

B. L2 Unknown Unicast: Hardware Proxy ARP Flooding: Enabled

C. L2 Unknown Unicast: Flood ARP Flooding: Disabled

D. L2 Unknown Unicast: Flood

ARP Flooding: Enabled

Answer: D

Explanation:

Question: 217

Refer to the exhibit.

Refer to the exhibit. An engineer configures a Layer 4 to Layer 7 device object. The device is a virtual firewall with a single network adapter and it must be deployed in routed mode. Which .. completes the configuration of the device object?

A. Change Function Type to GoTo.

B. Add an outside interface to the cluster interfaces.

C. Change context awareness to Multiple.

D. Enable Promiscuous Mode.

Answer: A

Explanation:

https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/6x/l4-l7-configuration/cisco-apic-layer-4-to- layer-7-services-deployment-guide-60x/defining-a-logical-device-60x.html

Question: 218

Refer to the exhibit.

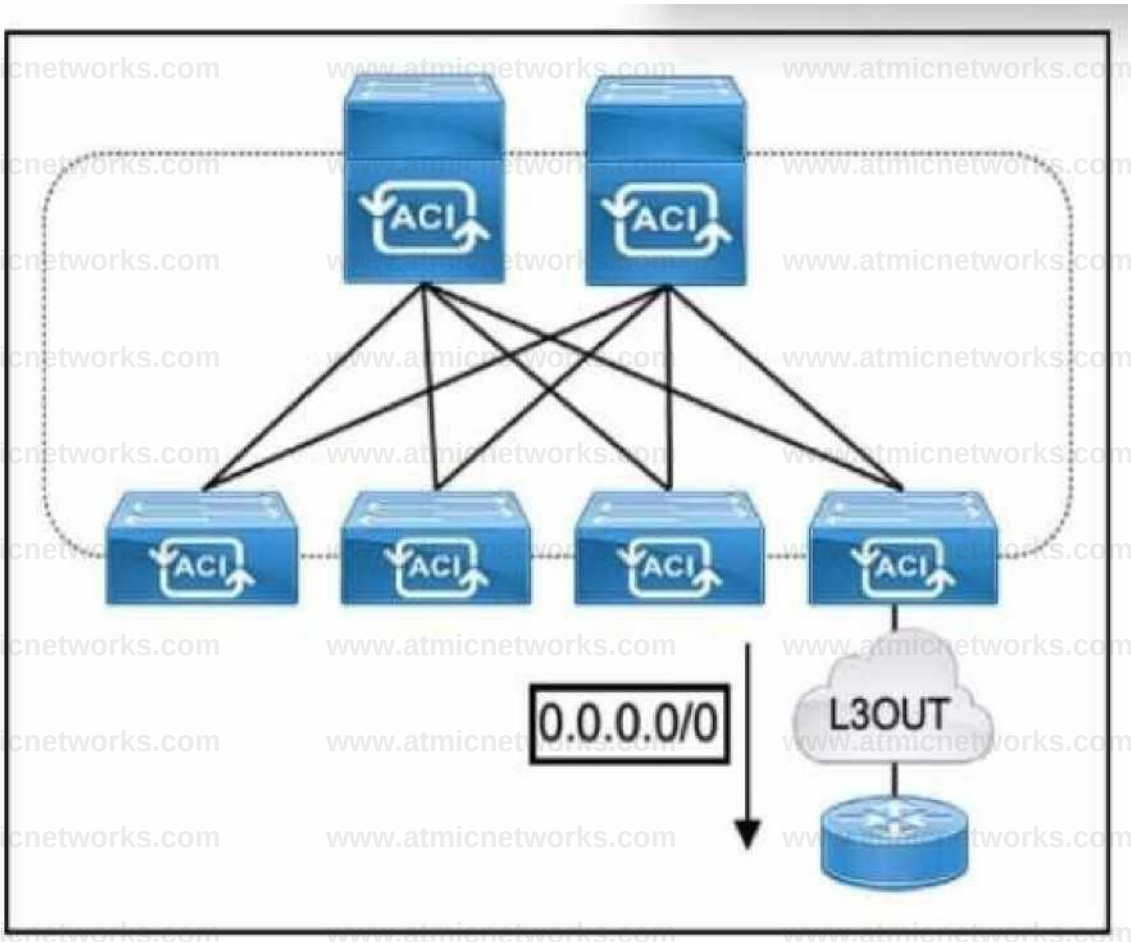

Refer to the exhibit. The default route is not present in the routing tables of the Cisco ACI leaf switches. All static and direct routes are currently being redistributed and advertised. Which jn must be taken to advertise a default route on the eBGP L30ut?

A. Configure a static default route on the ACI node profiles with next-hop null.

B. Create a Default Route Leak Policy on the L3Qut.

C. Enable a BGP peer prefix policy set to Always.

D. Implement an export route map matching 0.0.0.0/0.

Answer: B

Explanation:

For external connections to the fabric that only require a default route, there is support for originating a default route for OSPF, EIGRP, and BGP L3Out connections. If a default route is received from an external peer, this route can be redistributed out to another peer following the transit export route control as described earlier in this article. A default route can also be advertised out using a Default Route Leak policy. This policy supports advertising a default route if it is present in the routing table or it always supports advertising a default route. The Default Route Leak policy is configured in the L3Out connection.

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_010100.html

Question: 219

An engineer must attach an ESXi host to the Cisco ACI fabric. The host is connected to Leaf 1 and has its gateway IP address 10.10.10.254/24 configured inside the ACI fabric. A new wall is attached to Leaf 2 and mapped to the same EPG and BD as the ESXi host. The engineer must migrate the gateway of the ESXi host to the firewall. Which configuration set complishes this goal?

A. Disable unicast routing.

Configure IP address 10.10.10.254/24 on the ACI BD.

B. Enable unicast routing.

Configure IP address 10.10.10.254/24 on the ACI EPG.

C. Disable unicast routing.

Define IP address 10.10.10.254/24 on the firewall.

D. Enable unicast routing.

Set IP address 10.10.10.254/24 on the firewall.

Answer: C

Explanation:

Question: 220

Refer to the exhibit.

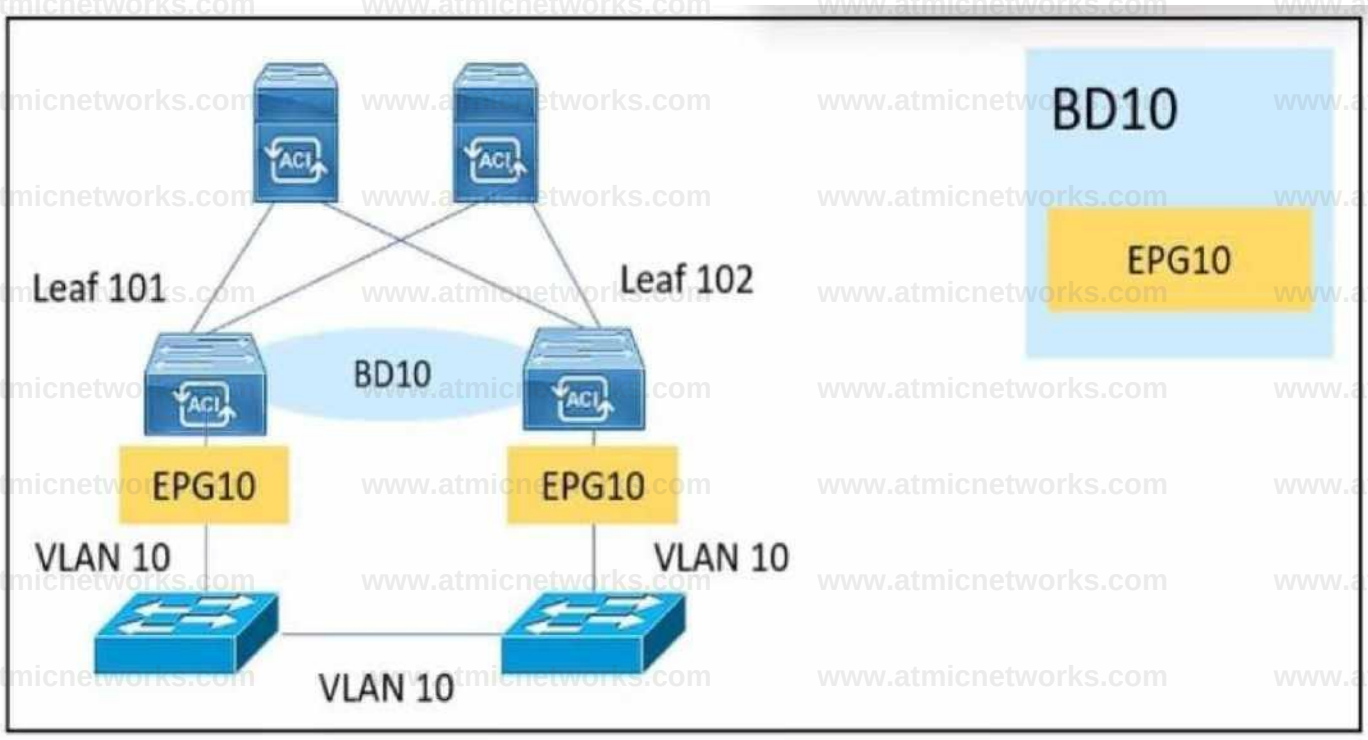

An engineer is implementing a BPDU filter on external switch interfaces That face the Cisco ACI fabric to prevent excessive TCNs from impacting the fabric. Which Configuration must be applied on Cisco ACI to avoid a Layer 2 loop?

A. Apply an MSTP instance on Cisco ACI.

B. Configure MCP globally

C. implement BPDU Guard.

D. Enable STP on downlinks.

Answer: B

Explanation:

Question: 221

What is the name of the automatically configured VLAN 3600 presented during Cisco ACI fabric discovery?

.F101# show ip int brief vrf overlay-1
utput truncated for brevity...)

iterface Status for VRF "overlay-1 "(4)
lace Address Interlace Status
/49 unassigned protocol-up/link-up/admin-up
Z49.34 unnumbered protocol-up/link-up/admin-up

/50 unassigned protocol-up/link-up/admin-up
/50.35 unnumbered protocol-up/link-up/admin-up

8 10 233.44 30/27 protocol-up/link-up/admin-up
10.233.46.32/32 protocol-up/link-up/admin-up
23 10.233.44.32/32 protocol-up/link-up/admin-up

,F1O1# show vlan extended
,N Name Encap Ports

ia default vxlan-4 6 9660132 Eth1/1. Eth1'2, Eth1/47
-3600

A. Transit VLAN

B. Infrastructure VLAN

C. Loopback VLAN

D. Fabric VLAN

Answer: B

Explanation:

The Infrastructure VLAN is automatically configured during the Cisco ACI fabric discovery process. This VLAN, often referred to as the “infra VLAN,” is used to enable internal communication between the ACI fabric components, such as the leaf and spine nodes. It is vital for the overlay network and the control plane functionality. In the provided output, VLAN 3600 represents the Infrastructure VLAN, which is configured for inter-node communication.

Question: 222

An engineer configures SNMP for an ACI fabric and created an SNMP Monitoring Destination Group called snmp_dgroup1. Snmp_dgroup1 is configured with the server hostname and Community password. An SNMP

policy called snmp_podpolicy1 is configured to enable SNMP and add an SNMP Client Group Profile called snmp_clgroup1. Snmp_podpolicy1 is associated default pod profile via a pod policy group named podl. Which configuration set must the engineer enable to complete the SNMP configuration?

A. Configure an SNMP management contract to permit all traffic. Associate snmp_podpolicy1 with an SNMP pod profile.

B. Configure the OOB management contract to permit all traffic. Associate snmp_clgroup1 with the SNMP management EPG.

C. Configure the OOB management contract to permit UDP 162. Associate snmp_dgroupl with the OOB management EPG.

D. Configure an SNMP management contract to permit UDP 162. Associate the SNMP Source to snmp_clgroup1.

Answer: C

Explanation:

https://community.cisco.com/t5/documentos-data-center/configuraci%C3%B3n-snmp-para-aci/ta-p/4680520

Question: 223

Engineer must configure SNMP inside a Cisco ACI fabric. The engineer has created an SNMP Policy, called SNMP-policy and an SNMP Monitoring Group called SNMP-group1 that Contains five trap receivers. Which configuration set completes the configuration?

A. Edit oobbrc to permit traffic using UDP port 16. Associate the client group policy to SNMP-group1.

B. Permit OOB management traffic using UDP port 161. Associate client group policy with the OOB management EPG.

C. Allow all OOB management traffic. Configure three trap receivers on SNMP-group1.

D. Create an OOB management contract. Include the SNMP server in the OOB management EPG.

Answer: B

Explanation:

Question: 224

Which two external entities are referenced by an AEP? (Choose two.)

A. VMware vCenter server

B. VMM domain

C. Layer 3 domain

D. Hypervisor

E. Fibre Channel switch

Answer: B, C

Explanation:

Question: 225

Which two hardware models are supported as fixed spine in Cisco ACI fabrics? (Choose two.)

A. Cisco Nexus 9508

B. Cisco Nexus 9236C

C. Cisco Nexus 9364C

D. Cisco Nexus 9336C-FX2

E. Cisco Nexus 9332C

Answer: CE

Explanation:

Question: 226

An engineer must configure a service graph for the policy-based redirect to redirect traffic to a transparent firewall. The policy must be vendor-agnostic to support any firewall appliance, Which two actions accomplish these goals? (Choose two.)

A. Set the Service Type to Other.

B. Set Promiscuous Mode to True.

C. Set Function Type to L2.

D. Set Managed to True.

E. Set Context Aware to Single.

Answer: AC

Explanation:

Question: 227

Cisco ACI fabric must send a packet between two pods in a Cisco AC1 Multi-Pod topology where ARP flooding is disabled within the bridge domain. How does a Cisco ACI spine switch .. ARP messages from a leaf switch in POD1 to POD2?

A. The ARP message is dropped and connectivity is lost between the endpoints.

B. ARP optimization is applied and sends ARP to remote anycast.

C. A proxy ARP message is sent to destination group 225.224.0.0.

D. An ARP Glean message is sent to multicast address 239.255.255.240.

Answer: D

Explanation:

Question: 228

Refer to the exhibit.

Refer to the exhibit. An engineer created a local user named User on Cisco ACI. The engineer must configure the fabric so that the User can access only common and PROD tenants, ch set of actions accomplishes the goal?

A. Add security domain "all" to User.

Associate security domain "all" under PROD tenant.

B. Add security domain 'Tenant" to User.

Associate security domain "Tenant" under PROD tenant.

C. Add security domain "common" to User.

Associate security domain "common" under PROD tenant.

D. Add security domain "mgmt" to User

Associate security domain "mgmt" under PROD tenant.

Answer: C

Explanation:

Question: 229

An engineer created a monitoring policy called Test in a Cisco ACI fabric and had to change the severity level of the monitored object Call home source. Which set of actions prevent the event from appearing in event reports?

A. Select Event Severity Assignment Policies. Set severity level to cleared.

B. Select Faults Severity Assignment Policies. Set severity level to cleared.

C. Select Event Severity Assignment Policies. Set severity level to squelched.

D. Select Faults Severity Assignment Policies. Set severity level to squelched.

Answer: C

Explanation:

Question: 230

Cisco ACI fabric is integrated with VMware VDS. The fabric must apply a security policy to check the integrity of traffic out of the network adapter. Which action must be taken to drop the .. when the ESXi host discovers a mismatch between the actual source MAC address transmitted by the guest operating system and the effective MAC address of the virtual machine ....?

A. Reject MAC changes.

B. Accept forged transmits.

C. Accept MAC changes.

D. Reject forged transmits.

Answer: B

Explanation:

https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.security.doc/GUID-7DC6486F-5400-44DF-8A62-6273798A2F80.html

Question: 231

Refer to the exhibit.

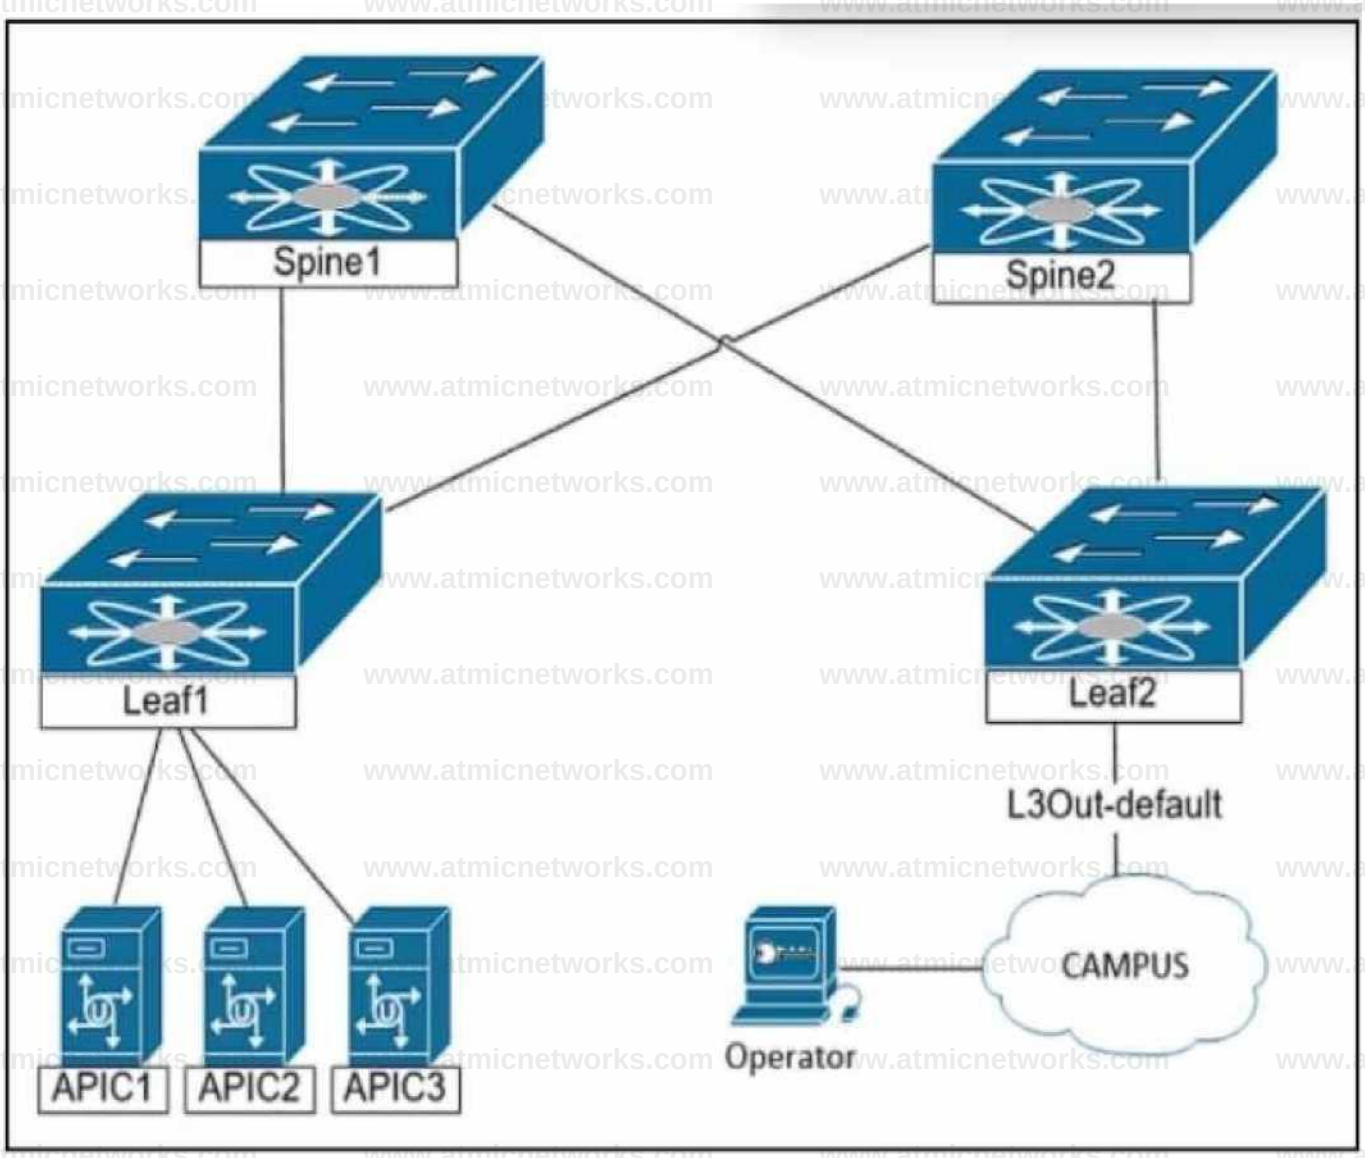

The engineer is planning to configure in-band management for the Cisco ACI fabric. The goal is to allow the network operators to reach the Cisco APIC servers and fabric switches from the in-band network. Which configuration must be applied on the bridge domain to accomplish these goals?

A. Enable Unicast Routing. Configure a virtual IP address.

B. Enable Unicast Routing. Set scope to Advertised Externally.

C. Scope: Shared between VRF. Set the IP address as primary.

D. Make this IP address primary. Configure an L30ut for Route Profile.

Answer: B

Explanation:

https://www.cisco.com/c/en/us/support/docs/cloud-systems-management/application-policy-infrastructure-controller-apic/221867-configure-in-band-management-in-aci.html Scope - Choose according to the route leakage method you use. Here choose to use L3out, and then click Advertised Externally.

Question: 232

Which protocol is used in a Multi-Pod topology to synchronize reachability information across pods?

A. IS-IS

B. MP-BGP EVPN

C. OSPF

D. COOP

Answer: B

Explanation:

https://learnduty.com/cisco-aci/aci-multi-pod-overview-and-basics/

Question: 233

Cisco ACI fabric contains a tenant called Prod. User_1 must have write access to tenant Prod and full access to the fabric access policy. Which set of actions must be taken to meet these requirements?

A. Associate User_1 to the fabric access policy.

Associate the security domain to the fabric access policy.

Create RBAC for the distinguished name of tenant Prod.

B. Associate User_1 to tenant Prod.

Associate the security domain to the distinguished name of the fabric access policy.

Create RBAC for the distinguished name of security domain.

C. Associate User_1 to the distinguished name of the fabric access policy.

Associate the security domain to RBAC.

Create RBAC for the distinguished name of User 1.

D. Associate User_1 to the security domain.

Associate the security domain to tenant Prod.

Create RBAC for the distinguished name of fabric access policy.

Answer: D

Explanation:

Question: 234

Refer to the exhibit.

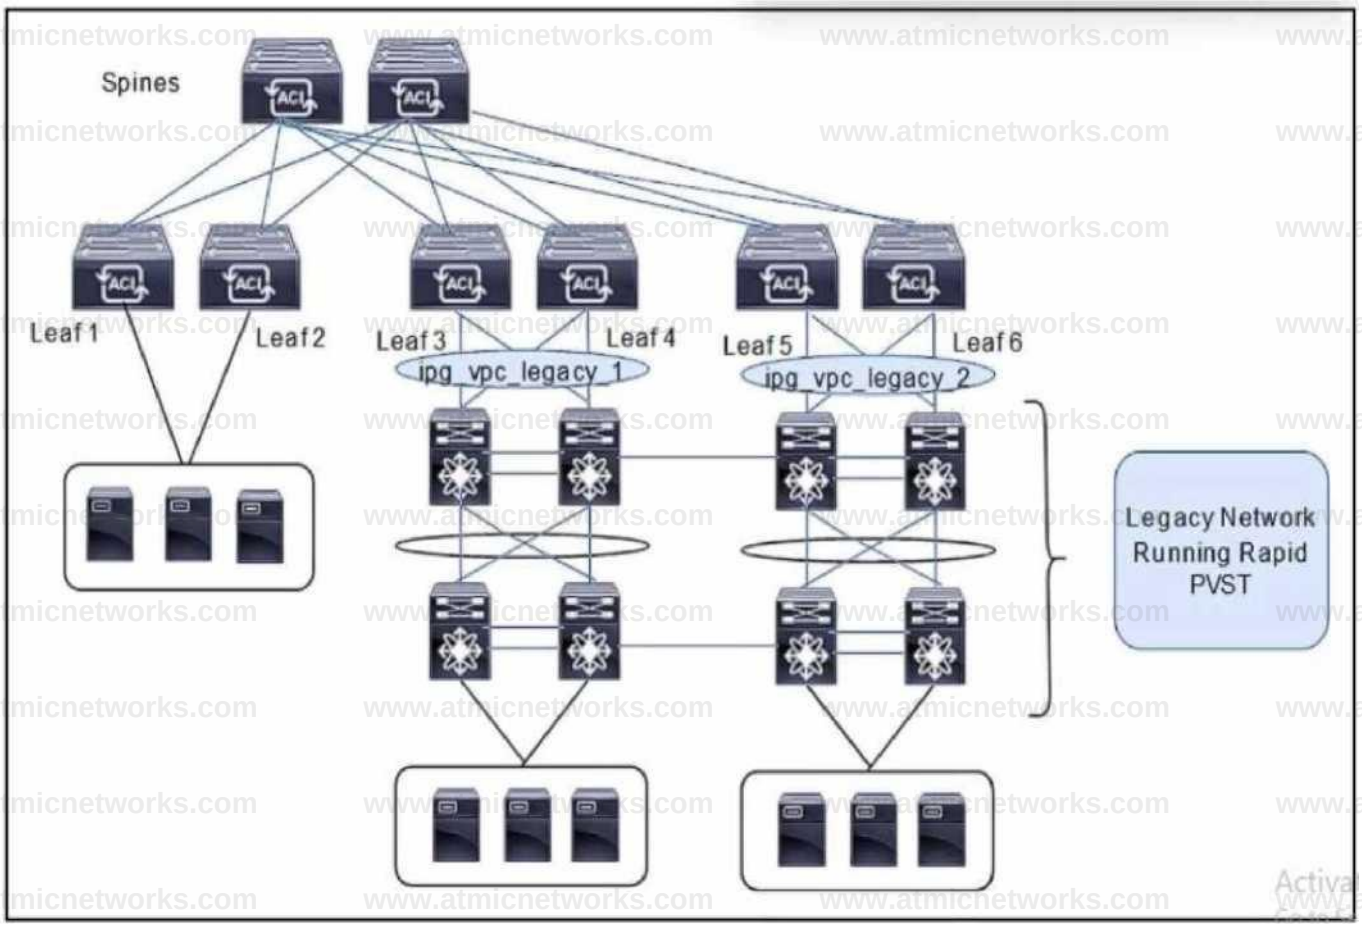

Refer to the exhibit. A client is configuring a new Cisco ACI fabric. All VLANs will be extended during the migration phase using the VPC connections on leaf switches 3. 4 and leaf switches toward the legacy network. The migration phase has these requirements;

* If The legacy switches must be able to transfer BPDUs through the ACI fabric.

* If the legacy switches fail to break a loop. Cisco ACI must break the loop.

Which group settings must be configured on VPC interface policy groups ipg_vpc-legacy_1 and ipg_vpc-legacy_2 to meet these requirements?

A. MCP: enabled

BPDU Guard: enabled

BPDU Filter: disabled

B. MCP: enabled

BPDU Guard: disabled

BPDU Filter: disabled

C. MCP: disabled

BPDU Guard: disabled

BPDU Filter: enabled

D. MCP: disabled

BPDU Guard: enabled

BPDU Filter enable

Answer: B

Explanation:

Question: 235

What is the result of selecting the On Demand attribute in the Deploy Immediacy feature during VMM domain association to an EPG?

A. The EPG policy is downloaded to the leaf when a hypervisor is connected, and a VM is placed in a port group.

B. The EPG policy is programmed in the hardware policy CAM only when the first packet is received through the data path.

C. The EPG policy is programmed in the hardware policy CAM as soon as the policy is downloaded in the

leaf software.

D. The EPG policy is downloaded to the leaf when a hypervisor is attached to a DVS. and CDP or LLDP adjacency is formed.

Answer: B

Explanation:

Question: 236

Refer to the exhibit.

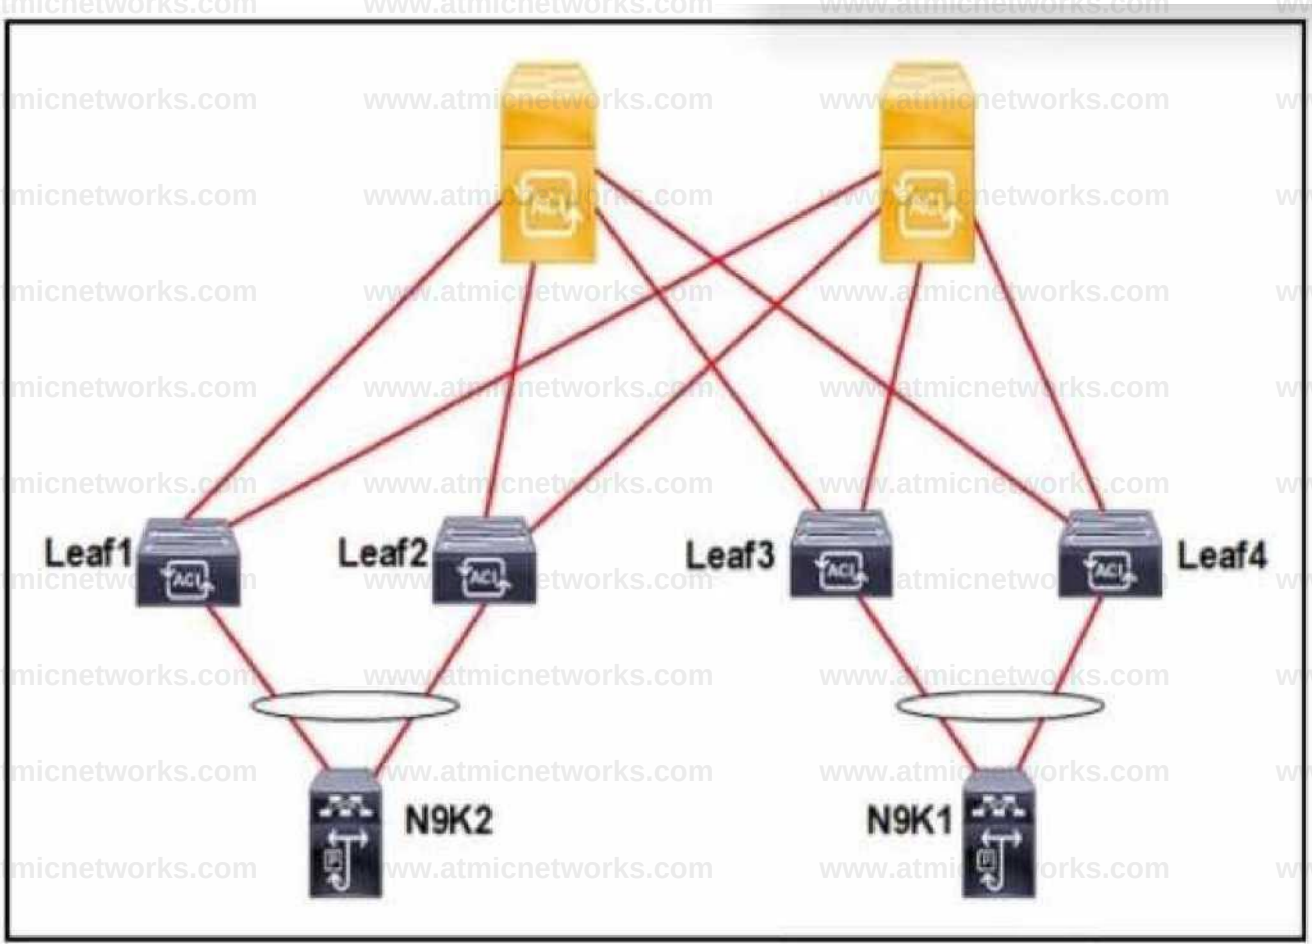

The Cisco ACI fabric is built with L20ut to the N9K1 and N9K2 switches. The switches run the RSTP protocol.

The requirement is for the Cisco ACI fabric to detect 5 from the N9K and for the fabric to be protected against loops. Which set of actions must be taken to meet the requirements?

A. Configure the N9K STP link type as point-to-point link. Enable MCP on ACI globally.

B. Configure the N9K STP link type as a point-to-point Enable MCP on the ACI leaf interfaces.

C. Configure the N9K STP link type as a shared link. Enable MCP on the ACI leaf interfaces.

D. Configure the N9K STP link type as a shared link. Enable MCP on ACI globally.

Answer: D

Explanation:

Question: 237

Engineer resolves an underlying condition of a fault but notices that the fault was not deleted from the Faults view. Which two actions must be taken to remove the fault? (Choose two.)

A. The fault is deleted after the retention interval.

B. Acknowledge the fault as an administrator.

C. The raised condition ceases.

D. The soaking timer expires.

Answer: AB

Explanation:

Question: 238

An engineer plans a Cisco ACI firmware upgrade. The ACI fabric consists of three Cisco APIC controllers, two spine switches, and four leaf switches. Two leaf switches have 1-Gb copper s for bare metal servers, and the other two leaf switches have 10-Gb SFP ports to connect storage. Which set of actions accomplishes an upgrade with minimal disruptions?

A. Upgrade the APIC controllers by selecting the desired firmware and choosing Upgrade Now.

Divide the switches into two upgrade groups: spines and leaves.

Start the firmware upgrade on the spine upgrade group and then proceed with the leaf upgrade group.

B. Upgrade the APIC controllers by initiating the upgrade process that uses the most recent uploaded firmware.

Divide the switches into three upgrade groups: spines, 1-Gb switches, and 10-Gb switches.

Start the firmware upgrade on the spine upgrade group and then proceed with the other two groups.

C. Upgrade the APIC controllers by selecting the desired firmware and choosing Upgrade Now.

Divide the switches into two upgrade groups with one spine, one 1-Gb switch, and one 10-Gb switch per group.

Start the firmware upgrade on the first upgrade group and when it finishes, start the second upgrade group.

D. Upgrade the APIC controllers as a single group by selecting the firmware and choosing Upgrade Now.

Divide the switches into four upgrade groups with one switch per group.

Start the firmware upgrade on each upgrade group in succession until all four are complete.

Answer: C

Explanation:

Question: 239

How many ARP requests are sent from leaf switches to perform host tracking for local endpoints?

A. 1

B. 2

C. 3

D. 4

Answer: A

Explanation:

Question: 240

Cisco ACI fabric must detect all silent endpoints for the Layer 3 bridge domain. Which actions accomplish

this goal?

A. Disable Unicast Routing.

Enable L2 Unknown Unicast Hardware Proxy.

B. Disable Unicast Routing.

Enable L2 Unknown Unicast Flood.

C. Enable Unicast Routing. Disable ARP Flooding.

D. Enable Unicast Routing. Enable ARP Flooding.

Answer: C

Explanation:

Question: 241

Network engineer configured a Cisco ACI fabric as follows:

- An EPG called EPG-A is created and associated with a VMM domain called North. •The EPG-A is associated with BD-A and is in an application profile called Apps-A.
- The BD-A is associated with VRF-1 in the Prod tenant.

Which port group must be selected to place VMs in EPG-A?

A. Prod|VRF-1 |Apps-A|EPG-A

B. Prod|Business_Apps|BD-A|EPG-A

C. Prod|Apps-A|North|EPG-A

D. Prod|Apps-A|EPG-A

Answer: D

Explanation:

Question: 242

An engineer implements a configuration backup on the Cisco APIC. The backup job must meet these requirements:

- The backup must transfer the encrypted data to the remote server.
- The transfer must be resumed if the connection is interrupted.

Which configuration set meets these requirements?

A. Select protocol HTTP in Create Remote Location. Choose JSON format in Configuration Export Policy.

B. Select protocol TFTP in Create Remote Location. Choose JSON format in Configuration Export Policy

C. Select protocol FTP in Create Remote Location. Choose XML format in Configuration Export Policy.

D. Select protocol SFTP in Create Remote Location. Choose XML format in Configuration Export Policy.

Answer: D

Explanation:

Question: 243

Refer to the exhibit.

A network engineer must improve the configuration backup process and the configuration restore process. The

current ACI solution is integrated with VMMs and third-y.. L4-L7 devices. The process requires that no additional information be re-entered when importing the configuration for a fully- functional state. Which configuration configures the port policy?

A. Enable the Global AES Encryption Setting.

B. Select the JSON data format to be used when exporting

C. Create target DNs for all tenants.

D. Configure a local snapshot.

Answer: A

Explanation:

Enabling AES encryption ensures that sensitive data, such as credentials for VMMs and third-party integrations, is securely encrypted in the backup file. This is essential for a fully functional restore without requiring re-entry of sensitive details.

Question: 244

An engineer needs to avoid loops in the ACI network and needs an ACI leaf switch to error-disable an interface if the interface receives an ACI-generated packet. Which action meets these requirements?

A. Enable the Loop Indication by MCP event in the Error Disabled Recovery Policy.

B. Set Rogue EP Control in the Endpoint Controls Policy.

C. Uncheck the Loop Protection Action check box in MCP Instance Policy.

D. Change the default administrative state of the global MCP Instance Policy.

Answer: D

Explanation:

MisCabling Protocol (MCP) detects loops from external sources (i.e., misbehaving servers, external

networking equipment running STP, etc.) and will err-disable the interface on which ACI receives its own packet. Enabling this feature is a best practice, and it should be enabled globally and on all interfaces, regardless of the end device. For MCP to be enabled, you need to have it enabled globally and on a per-interface basis. While MCP is enabled on all interfaces by default, it is not turned "on" until you also enable it globally. The global configuration knob for MCP can be enabled by configuring the global settings here: Fabric > Access Policies > Global Policies > MCP Instance Policy default.
https://www.cisco.com/c/dam/en/us/solutions/collateral/data-center-virtualization/application- centric-infrastructure/aci-guide-using-mcp-mis-cabling-protocol.pdf

Question: 245

A Cisco ACI fabric is integrated with a Cisco ASA firewall using a service graph under the tenant called Operations. The fabric must permit the firewall used on tenant Operations to be referenced by the tenant called Management. Which export action must be used to accomplish this goal?

A. Layer4-Layer7 device
B. router configurations
C. service graph template
D. device selection policies

Answer: A

Explanation:

In Cisco ACI, when a service graph is deployed under one tenant (e.g., Operations) and needs to be referenced by another tenant (e.g., Management), the Layer 4-Layer 7 (L4-L7) device export action is used. This allows the firewall or other L4-L7 devices defined in the service graph to be shared across tenants. By exporting the L4-L7 device, the configuration enables the Management tenant to reference and use the firewall deployed in the Operations tenant.

Question: 246

Refer to the exhibit.

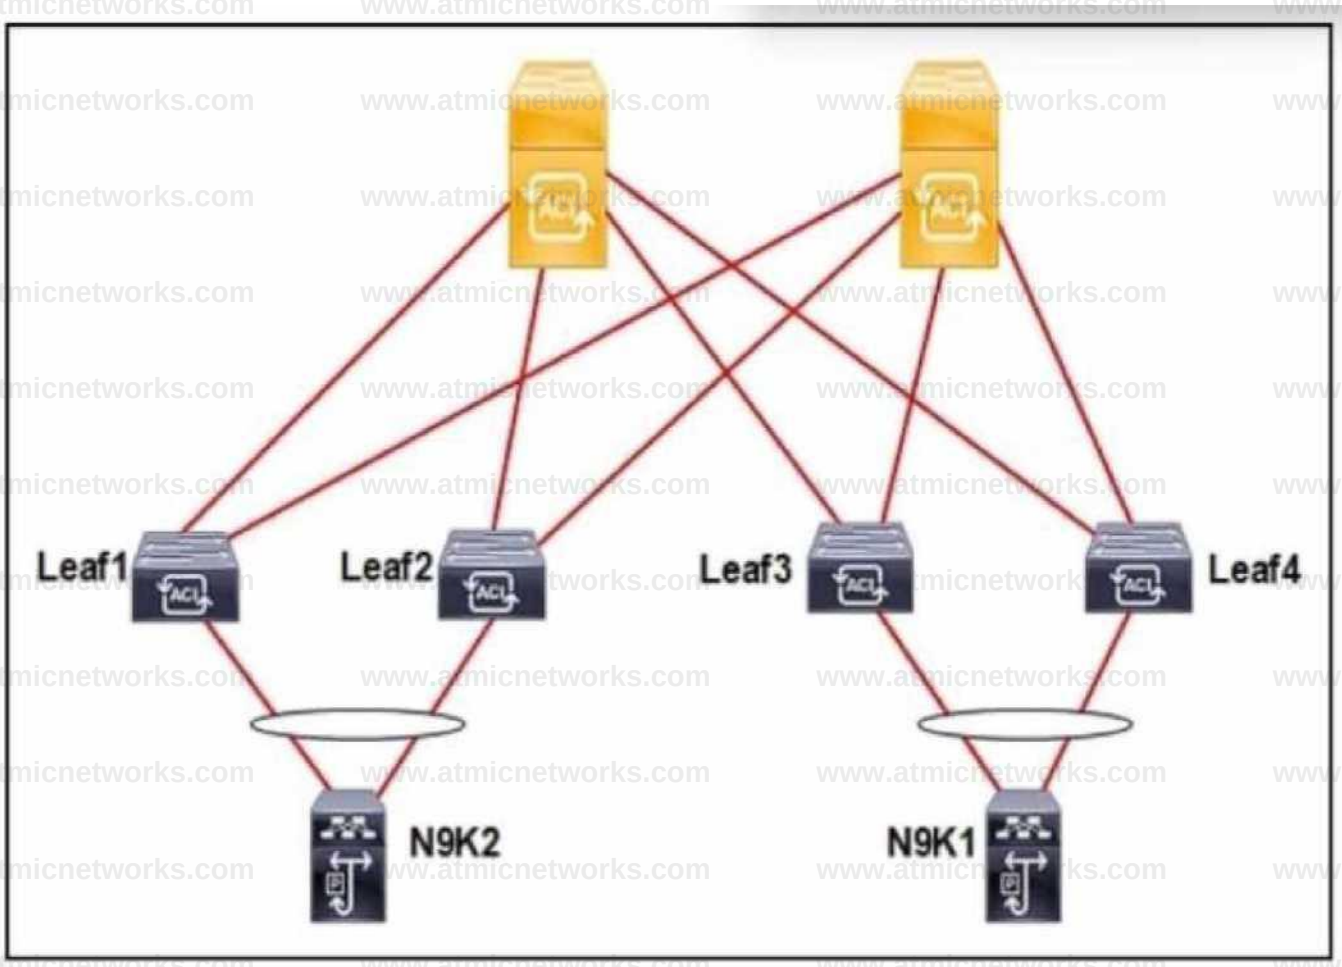

Refer to the exhibit. An engineer connects a Cisco ACI fabric to two different Cisco Nexus 9000 Series Switches. The fabric must be configured to ensure a loop-free topology and N9K1 be configured as the root bridge for VLAN 10. Which action meets these requirements?

A. Enable STP on ports between the leaf and spine.

B. Activate MCP on ports between the leaf and Nexus 9000 Series Switches.

C. Enable Cisco Discovery Protocol on ports between the leaf and spine.

D. Set BPDU Guard on ports between the leaf and Nexus 9000 Series Switches.

Answer: B

Explanation:

Question: 247

What is the maximum number of sites connected using spine back-to-back with a direct link in a Cisco ACI Multi-Site fabric?

A. 2
B. 3
C. 4
D. 5

Answer: A

Explanation:

In a Cisco ACI Multi-Site setup, back-to-back spine connectivity is limited to a direct connection between two sites. This design simplifies inter-site communication by avoiding the need for an intermediate Inter-Pod Network (IPN) or Multi-Site Orchestrator.