complete your programming course about resources, doubts and more!

Cisco

(350-601)

Implementing and Operating Cisco Data Center Core Technologies (DCCOR)

Total: **522 Questions**

Link:

Question: 1

DRAG DROP -
An engineer is implementing NetFlow on a Cisco Nexus 7000 Series Switch.
Drag and drop the NetFlow commands from the left into the correct order they must be entered on the right.
Select and Place:

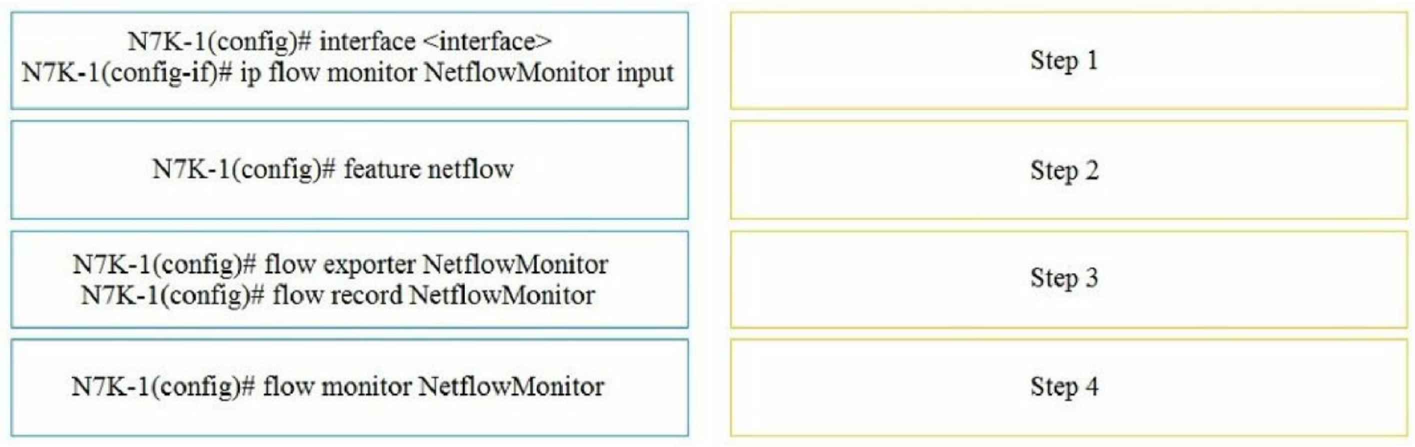

N7K-1(config)# interface <interface> N7K-1(config-if)# ip flow monitor NetflowMonitor input	Step 1
N7K-1(config)# feature netflow	Step 2
N7K-1(config)# flow exporter NetflowMonitor N7K-1(config)# flow record NetflowMonitor	Step 3
N7K-1(config)# flow monitor NetflowMonitor	Step 4

Answer:

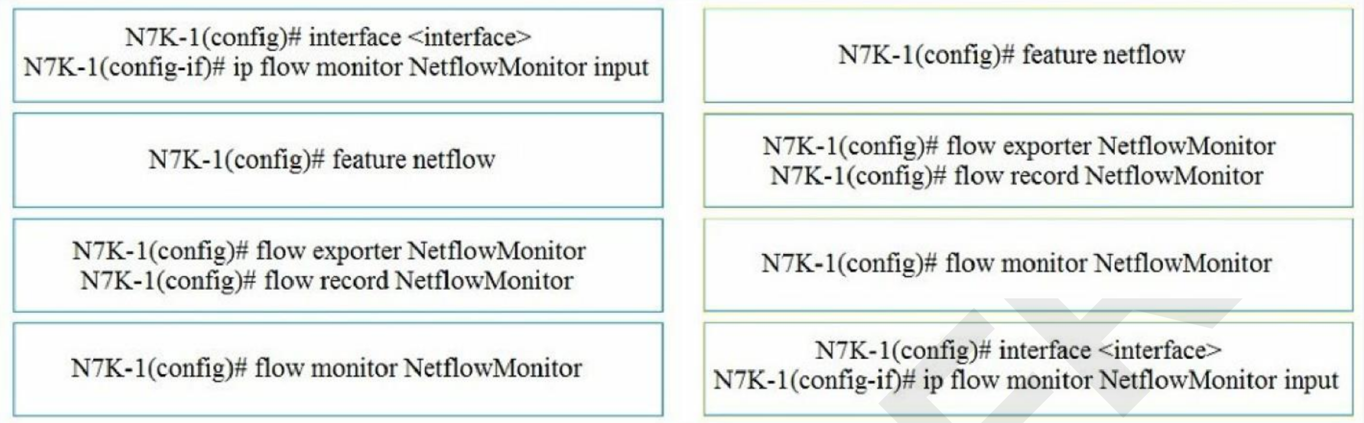

N7K-1(config)# interface <interface> N7K-1(config-if)# ip flow monitor NetflowMonitor input	N7K-1(config)# feature netflow
N7K-1(config)# feature netflow	N7K-1(config)# flow exporter NetflowMonitor N7K-1(config)# flow record NetflowMonitor
N7K-1(config)# flow exporter NetflowMonitor N7K-1(config)# flow record NetflowMonitor	N7K-1(config)# flow monitor NetflowMonitor
N7K-1(config)# flow monitor NetflowMonitor	N7K-1(config)# interface <interface> N7K-1(config-if)# ip flow monitor NetflowMonitor input

Explanation:

Reference:
https://www.cisco.com/c/en/us/support/docs/switches/nexus-7000-series-switches/112213-netflow-nexus70 00-nsox-configex.html

Question: 2

Which virtual MAC address is the default for HSRP version 2 group 10?

A. 0000.5E00.0110
B. 0000.0C9F.F00A
C. 3784.0898.1C0A
D. 0000.0C9F.F010

Answer: B

Explanation:

The correct answer is **B. 0000.0C9F.F00A**. This is because HSRP (Hot Standby Router Protocol) version 2 uses a different virtual MAC address structure than version 1. HSRPv2's MAC address is derived from the base multicast MAC address of 0000.0C07.AC00 and adds the HSRP group number to it. The specific formula is:

0000.0C07.AC00 + (group_number * 256) for the last two octets.

In this scenario, the group number is 10. Therefore, we calculate the last two octets: 10 * 256 = 2560. The hexadecimal representation of 2560 is 0x0A00. Then, add this to AC00 = AC00 + 0A00 = B600. This doesn't directly correspond. It appears to be taking the AC00 as a base, then converting 10 to hex which is 0x0A, and it's placing this in the last two octets (last byte). Then it prepends F0. This follows the HSRPv2 algorithm for constructing the MAC address, it's more like: 0000.0C9F.F0 + the HSRP group number in hex =
0000.0C9F.F00A. The 0x0A represents the decimal 10.

Option A (0000.5E00.0110) is a multicast MAC address used for IPv6 solicited-node multicast, not HSRP.

Option C (3784.0898.1C0A) doesn't follow the standard HSRP MAC address structure. Option D (0000.0C9F.F010) is the virtual MAC address for HSRP group 16 (since 16 in hex is 0x10), not group 10. HSRP ensures redundancy by having an active router forward traffic while a standby router takes over if the active one fails, using these virtual addresses to maintain seamless connectivity. The virtual MAC address provides a stable endpoint for network devices, regardless of which physical router is currently active.

Authoritative Links:

Cisco Documentation on HSRPv2 Virtual MAC Address:https://www.cisco.com/c/en/us/support/docs/ip/hot-standby-router-protocol-hsrp/10547-hsrp-vmac.html
Network Lessons HSRP:https://networklessons.com/cisco/ccna-routing-switching-icnd2-200-105/hsrp-hot-standby-router-protocol

Question: 3

Refer to the exhibit.

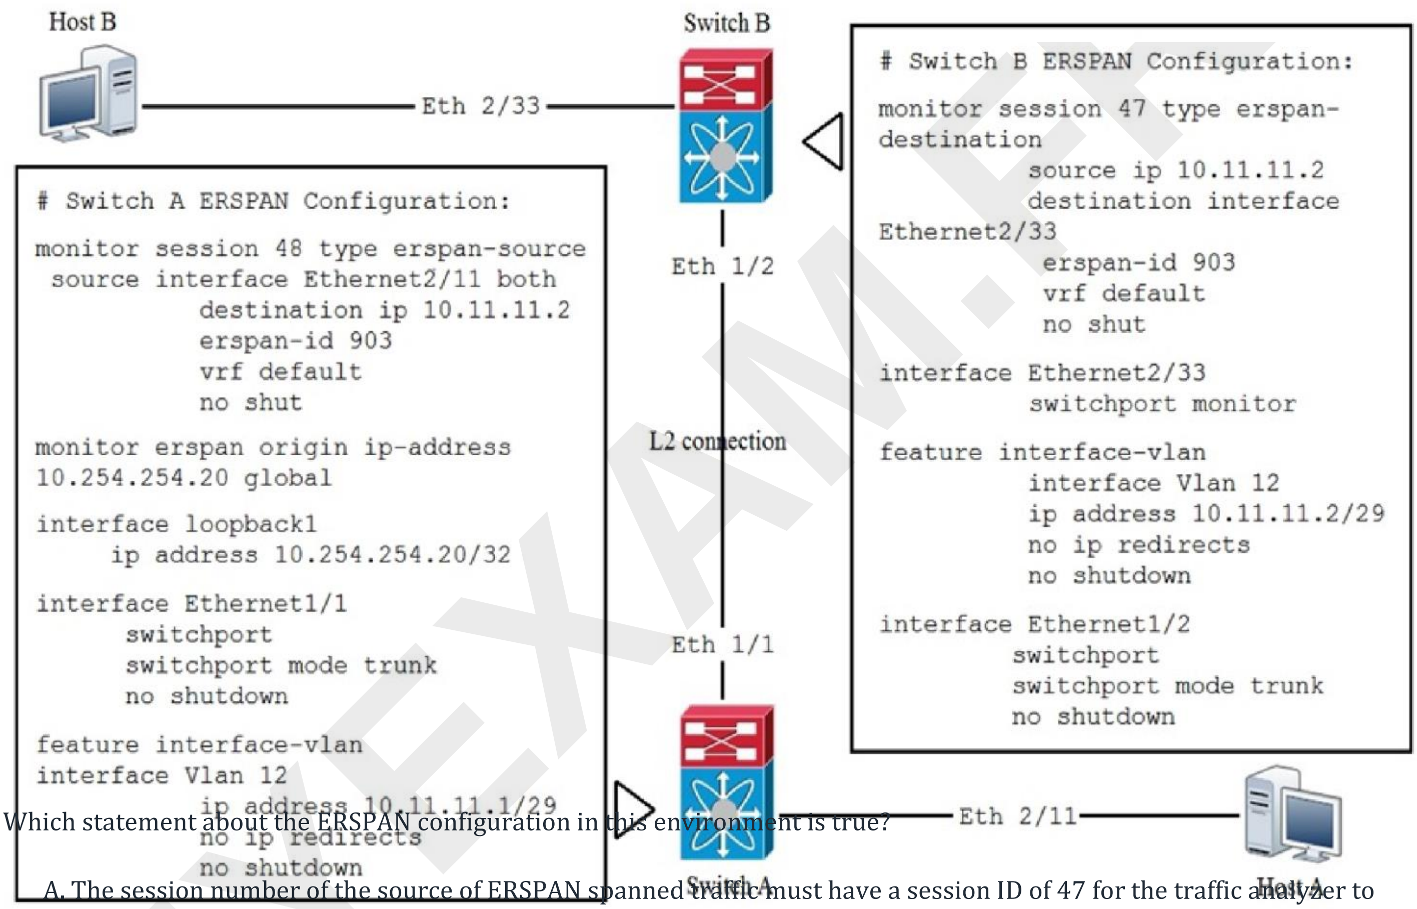

```
# Switch A ERSPAN Configuration:

monitor session 48 type erspan-source
 source interface Ethernet2/11 both
         destination ip 10.11.11.2
         erspan-id 903
         vrf default
         no shut

monitor erspan origin ip-address
10.254.254.20 global

interface loopback1
     ip address 10.254.254.20/32

interface Ethernet1/1
      switchport
      switchport mode trunk
      no shutdown

feature interface-vlan
interface Vlan 12
         ip address 10.11.11.1/29
         no ip redirects
         no shutdown
```

```
# Switch B ERSPAN Configuration:

monitor session 47 type erspan-
destination
         source ip 10.11.11.2
         destination interface
Ethernet2/33
          erspan-id 903
          vrf default
          no shut

interface Ethernet2/33
         switchport monitor

feature interface-vlan
         interface Vlan 12
         ip address 10.11.11.2/29
         no ip redirects
         no shutdown

interface Ethernet1/2
        switchport
        switchport mode trunk
        no shutdown
```

Which statement about the ERSPAN configuration in this environment is true?

A. The session number of the source of ERSPAN spanned traffic must have a session ID of 47 for the traffic analyzer to receive the traffic.

B. Host B is the source of ERSPAN spanned traffic and host A is the traffic analyzer.

C. The session number of the source of ERSPAN spanned traffic must have a session ID of 48 for the traffic analyzer to receive the traffic.

D. Host A is the source of ERSPAN spanned traffic and host B is the traffic analyzer.

Answer: D

Explanation:
The Erspan-id needs to match between the switches. The monitor session ID is only locally significant and does not need to match on each device.

Reference:
https://www.letsconfig.com/how-to-configure-erspan-on-cisco-nexus-switches/

Question: 4

Refer to the exhibit.

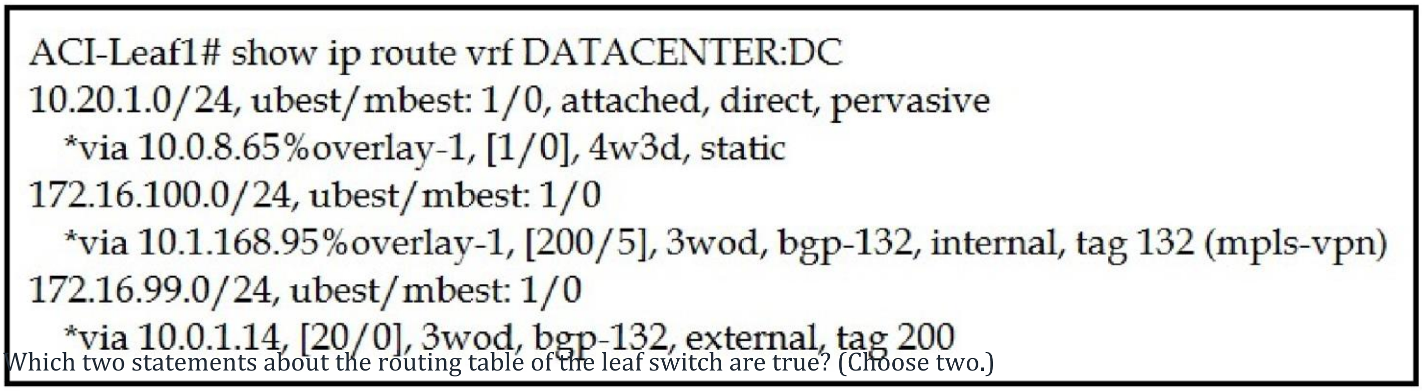

```
ACI-Leaf1# show ip route vrf DATACENTER:DC
10.20.1.0/24, ubest/mbest: 1/0, attached, direct, pervasive
  *via 10.0.8.65%overlay-1, [1/0], 4w3d, static
172.16.100.0/24, ubest/mbest: 1/0
  *via 10.1.168.95%overlay-1, [200/5], 3wod, bgp-132, internal, tag 132 (mpls-vpn)
172.16.99.0/24, ubest/mbest: 1/0
  *via 10.0.1.14, [20/0], 3wod, bgp-132, external, tag 200
```

Which two statements about the routing table of the leaf switch are true? (Choose two.)

A. 10.20.1.0/24 is a BD subnet in ACI.

B. The next hop 10.0.1.14 for route 172.16.99.0/24 is the TEP address of a border leaf in ACI.

C. 172.16.100.0/24 is a BD subnet in ACI.

D. The next hop 10.1.168.95 for route 172.16.100.0/24 is the TEP address of a border leaf in ACI.

E. The next hop 10.0.8.65 for route 10.20.1.0/24 is the TEP address of a border leaf in ACI.

Answer: AD

Explanation:
10.20.1.0/24 is a BD in ACI because of the part which says "direct, pervasive".

For the other answer choice, we see form the administrative distance and the "internal" naming, we know that this is an external subnet (172.16.100.0/24) and that as all external routes are exchanged withing the fabric through MP-BG. MP-BGP is actually used to populate the leaves in the fabric with routing on how to reach those external routes. In MP-BGP, the spines are only the route reflectors. So from this we know that the next hop of the external subnet will be a VTEP of a leaf.

Question: 5

Which mroute state is created when Bidirectional PIM is deployed at a site?

A. *,G

B. MVPN Type-6

C. MVPN Type-7

D. S,G

Answer: A

Explanation:

The correct answer is A. (,G). In Bidirectional PIM (Protocol Independent Multicast), the primary mroute state created is (,G). This state represents a shared tree, meaning that traffic destined for a specific multicast group (G) is forwarded along a single, shared path towards receivers. The asterisk () indicates that the source of the multicast traffic is not known or relevant at this point in the forwarding process. Bidirectional PIM operates on the principle of a shared, bidirectional tree, where data flows up the tree towards the Rendezvous Point (RP) and down the tree towards the receivers. Unlike Source-Specific Multicast (SSM), Bidirectional PIM does not maintain separate trees for each source. The (,G) state is fundamental to this shared-tree paradigm, allowing receivers to join the multicast group irrespective of the specific source. The traffic flows up the shared tree towards the RP and is then distributed down the shared tree to all interested receivers. The other options, MVPN Type-6 and Type-7, are related to Multicast VPN configurations, which are more complex and not directly relevant to the basic mroute state of Bidirectional PIM. The (S,G) state is associated with Source-Specific Multicast (SSM), where a tree is built for each source (S) sending to group (G). Bidirectional PIM is designed to minimize the complexity of the state by using a single shared tree, hence the (*,G) state. This makes it highly scalable for large multicast deployments.

Authoritative link:Cisco - Understanding Bidirectional PIM

Question: 6

Refer to the exhibit.

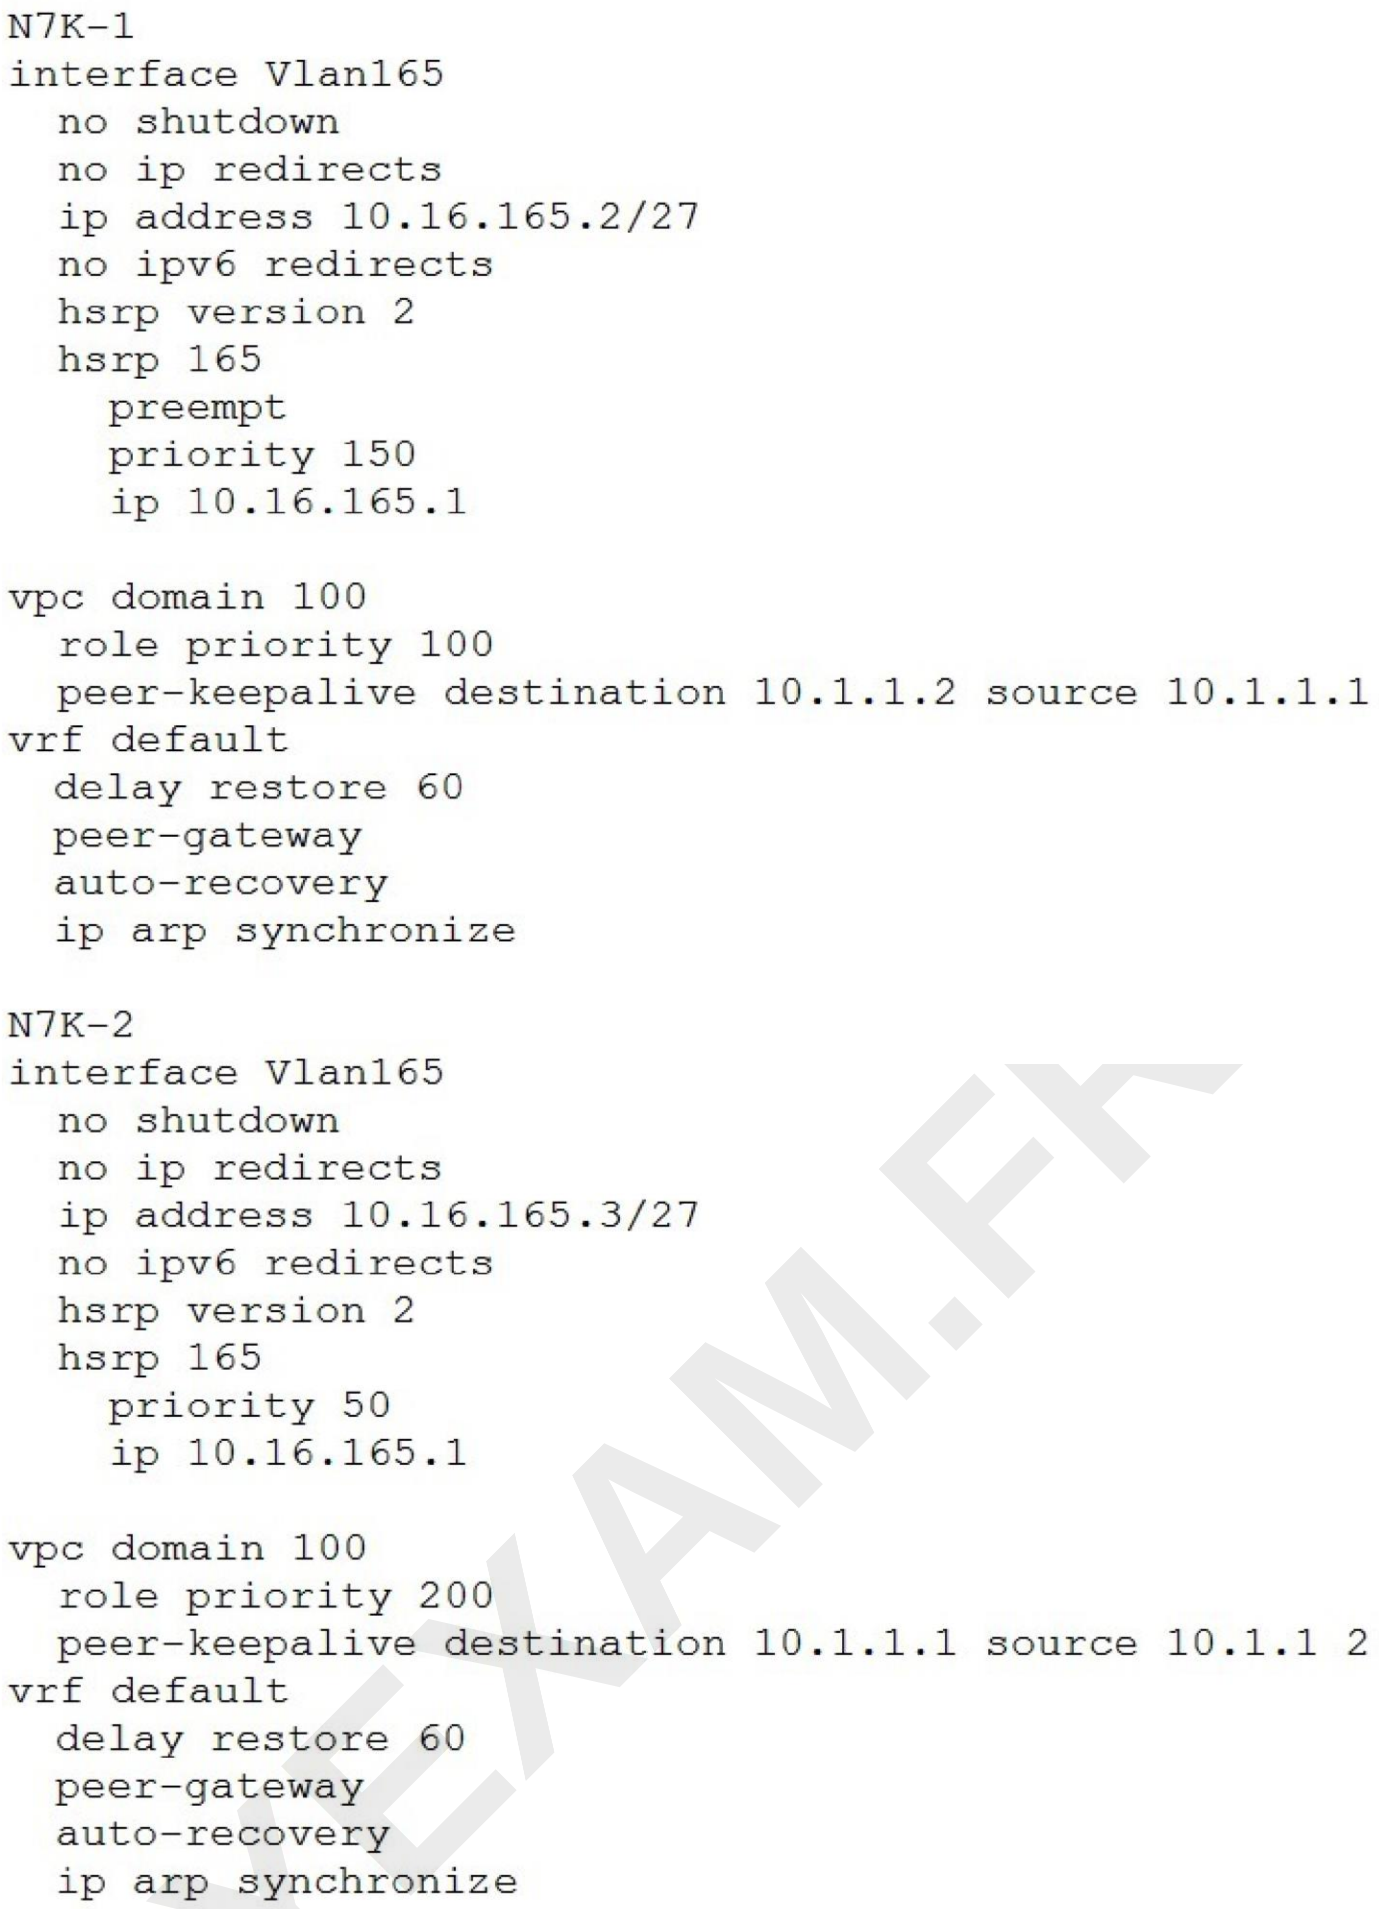

```
N7K-1
interface Vlan165
  no shutdown
  no ip redirects
  ip address 10.16.165.2/27
  no ipv6 redirects
  hsrp version 2
  hsrp 165
    preempt
    priority 150
    ip 10.16.165.1

vpc domain 100
  role priority 100
  peer-keepalive destination 10.1.1.2 source 10.1.1.1
vrf default
  delay restore 60
  peer-gateway
  auto-recovery
  ip arp synchronize

N7K-2
interface Vlan165
  no shutdown
  no ip redirects
  ip address 10.16.165.3/27
  no ipv6 redirects
  hsrp version 2
  hsrp 165
    priority 50
    ip 10.16.165.1

vpc domain 100
  role priority 200
  peer-keepalive destination 10.1.1.1 source 10.1.1 2
vrf default
  delay restore 60
  peer-gateway
  auto-recovery
  ip arp synchronize
```

Which statement about the default gateway configuration of the vPC is true?

A. Either switch can act as the active default gateway.

B. N7K-1 acts as the default gateway for all traffic.

C. N7K-2 forwards traffic that is destined for the default gateway by using the peer link.

D. N7K-2 acts as the default gateway for all traffic.

Answer: A

Explanation:
In normal Hot Standby Router Protocol operation, the active HSRP interface answers ARP requests, but with a vPC, both HSRP interfaces (active and standby) can forward traffic.

The most significant difference between the HSRP implementation of a non-vPC configuration and a vPC configuration is that the HSRP MAC addresses of a vPC configuration are programmed with the G (gateway) flag on both systems, compared with a non-vPC configuration, in which only the active HSRP interface can program the MAC address with the G flag. Given this fact, routable traffic can be forwarded by both the vPC primary device (with HSRP) and the vPC secondary device (with HSRP), with no need to send this traffic to the HSRP primary device. Without this flag, traffic sent to the MAC address would not be routed.

Reference:
https://books.google.com/books?id=5VXUDwAAQBAJ&pg=PT161&lpg=PT161&dq=%22normal+Hot+Standby+Router+Protocol+operation,+the+active+HSRP%22&source=bl&ots=dbSA3yl8BW&sig=ACfU3U1uGuplof1lJ4GF-0VVkqpjjFz0yA&hl=en&sa=X&ved=2ahUKEwia0MLTpvf4AhWmFzQIHQyCCtcQ6AF6BAgCEAM#v=onepage&q=%22normal%20Hot%20Standby%20Router%20Protocol%20operation%2C%20the%20active%20HSRP%22&f=false

Question: 7

Refer to the exhibit.

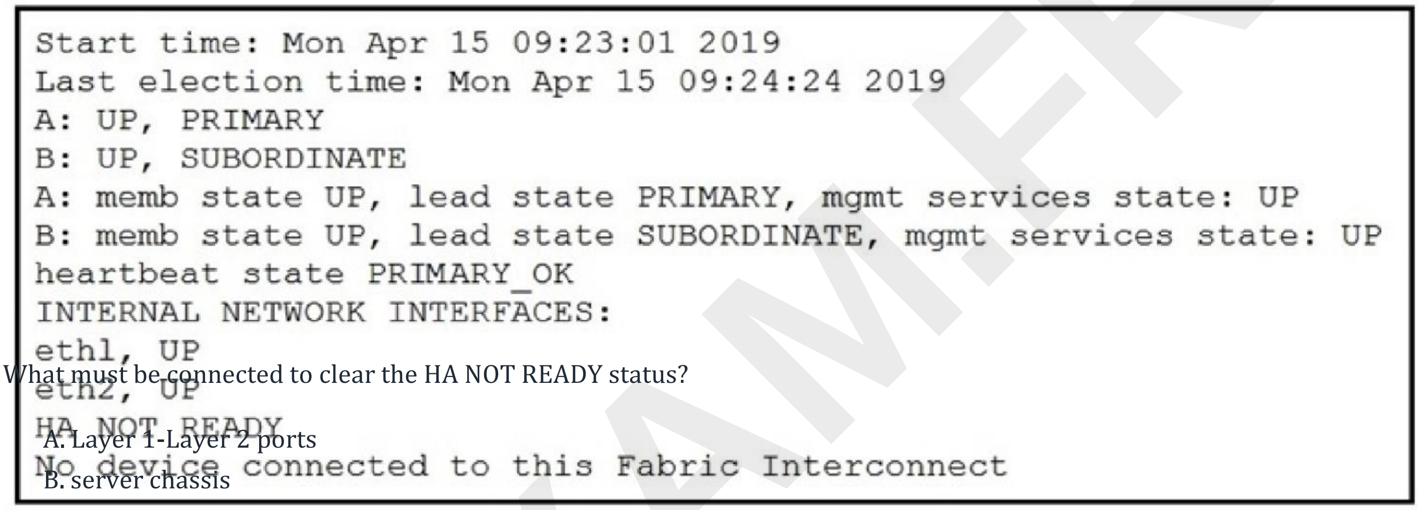

```
Start time: Mon Apr 15 09:23:01 2019
Last election time: Mon Apr 15 09:24:24 2019
A: UP, PRIMARY
B: UP, SUBORDINATE
A: memb state UP, lead state PRIMARY, mgmt services state: UP
B: memb state UP, lead state SUBORDINATE, mgmt services state: UP
heartbeat state PRIMARY_OK
INTERNAL NETWORK INTERFACES:
eth1, UP
eth2, UP
HA NOT READY
No device connected to this Fabric Interconnect
```

What must be connected to clear the HA NOT READY status?

A. Layer 1-Layer 2 ports

B. server chassis

C. management ports

D. network uplinks

Answer: B

Explanation:

Reference:
https://ucsguru.com/2012/11/07/ha-with-ucsm-integrated-rack-mounts/

Question: 8

A small remote office is set to connect to the regional hub site via NSSA ASBR. Which type of LSA is sent to the remote office OSPF area?

A. type 7 LSA
B. type 1 LSA
C. type 5 LSA
D. type 3 LSA

Answer: A

Explanation:

Okay, let's break down why the correct answer is **A. type 7 LSA** when an NSSA ASBR is involved in connecting a remote office to a regional hub site using OSPF.

An NSSA (Not-So-Stubby Area) is a special type of OSPF area designed to allow the injection of external routes while still maintaining some level of stub-area characteristics. Traditional stub areas do not allow external routes (type 5 LSAs) to be flooded into them. However, they can receive summaries or default routes via type 3 LSAs.

In this scenario, the remote office is connected via an ASBR (Autonomous System Boundary Router) that is also operating within an NSSA. When an ASBR in an NSSA receives external routes (e.g., from BGP), it doesn't generate Type 5 LSAs as it would in normal OSPF areas. Instead, the ASBR generates **Type 7 LSAs.**

Type 7 LSAs are unique to NSSAs and encapsulate external routing information. They are only flooded within the NSSA area itself. Once a Type 7 LSA reaches an ABR (Area Border Router), the ABR translates it into a standard Type 5 LSA before flooding it into the backbone or other areas of the OSPF network.

Therefore, within the remote office OSPF area, which is configured as an NSSA, the routers will receive and use type 7 LSAs to learn about external networks being advertised by the NSSA ASBR. The other options do not apply in this specific situation. Type 1 LSAs are router LSAs describing the local link and costs and are used in all areas. Type 3 LSAs are summary LSAs used for inter-area routing, and type 5 LSAs are for external routes which are not directly used within an NSSA before ABR translation.

Here's a summary:

NSSA: Allows external routes with restrictions.
ASBR: Injects external routes.
Type 7 LSA: Specific to NSSA, carries external route information. **ABR:** Translates type 7 LSAs to type 5 LSAs.

Authoritative Links for further research:

Cisco Documentation on OSPF NSSA:https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/70310-nssa-ospf.html
RFC 3101: The OSPF Not-So-Stubby Area (NSSA) Option: https://datatracker.ietf.org/doc/html/rfc3101

Question: 9

Which adjacency server configuration makes two OTV edge devices located in the same site bring up the dual-site adjacency?
A.

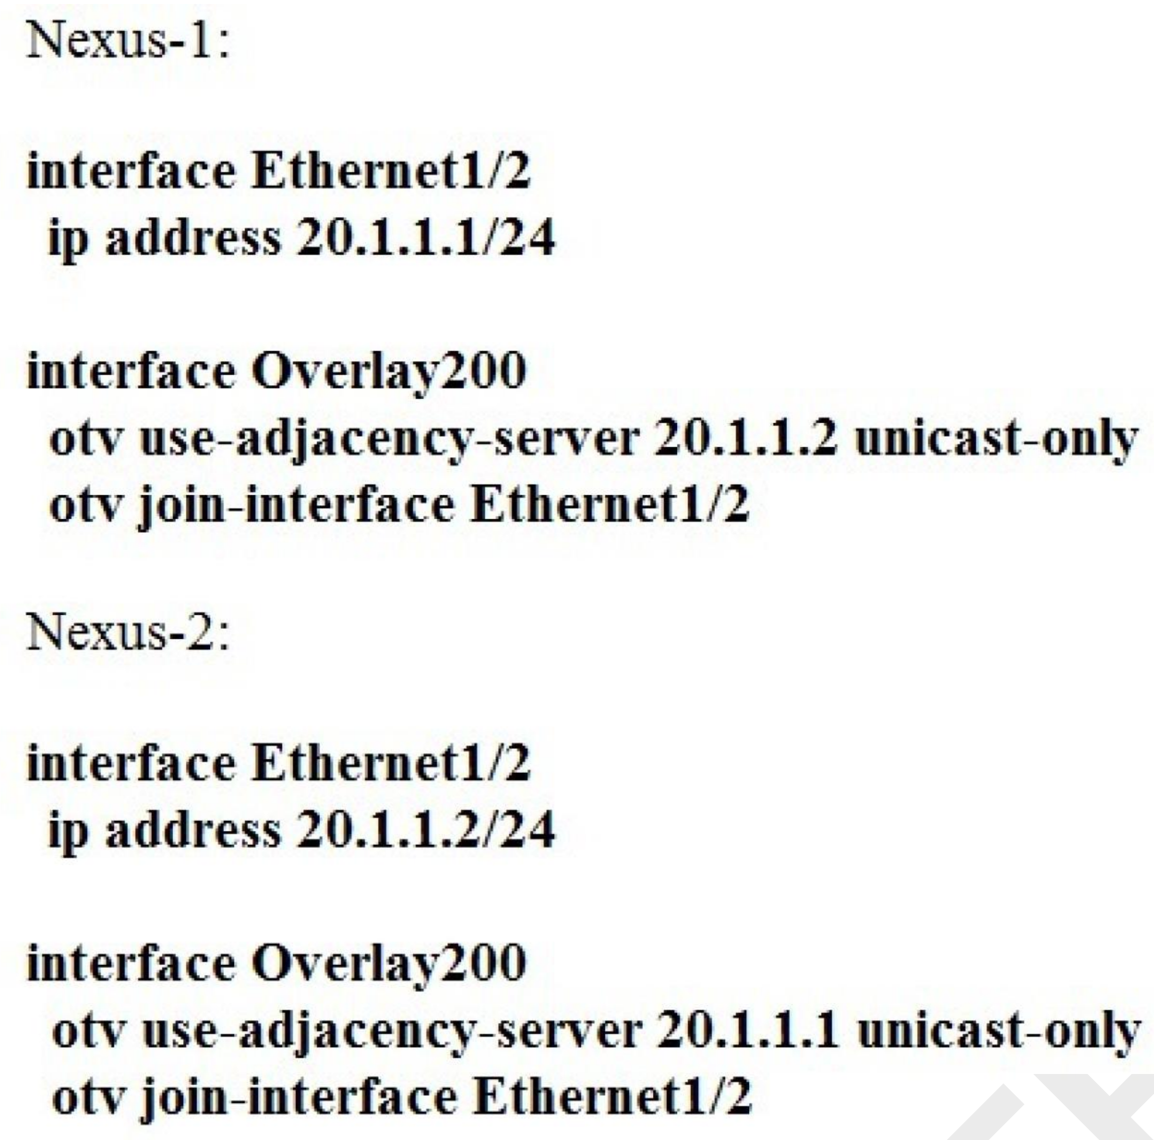

```
Nexus-1:

interface Ethernet1/2
 ip address 20.1.1.1/24

interface Overlay200
 otv use-adjacency-server 20.1.1.2 unicast-only
 otv join-interface Ethernet1/2

Nexus-2:

interface Ethernet1/2
 ip address 20.1.1.2/24

interface Overlay200
 otv use-adjacency-server 20.1.1.1 unicast-only
 otv join-interface Ethernet1/2
```

B.

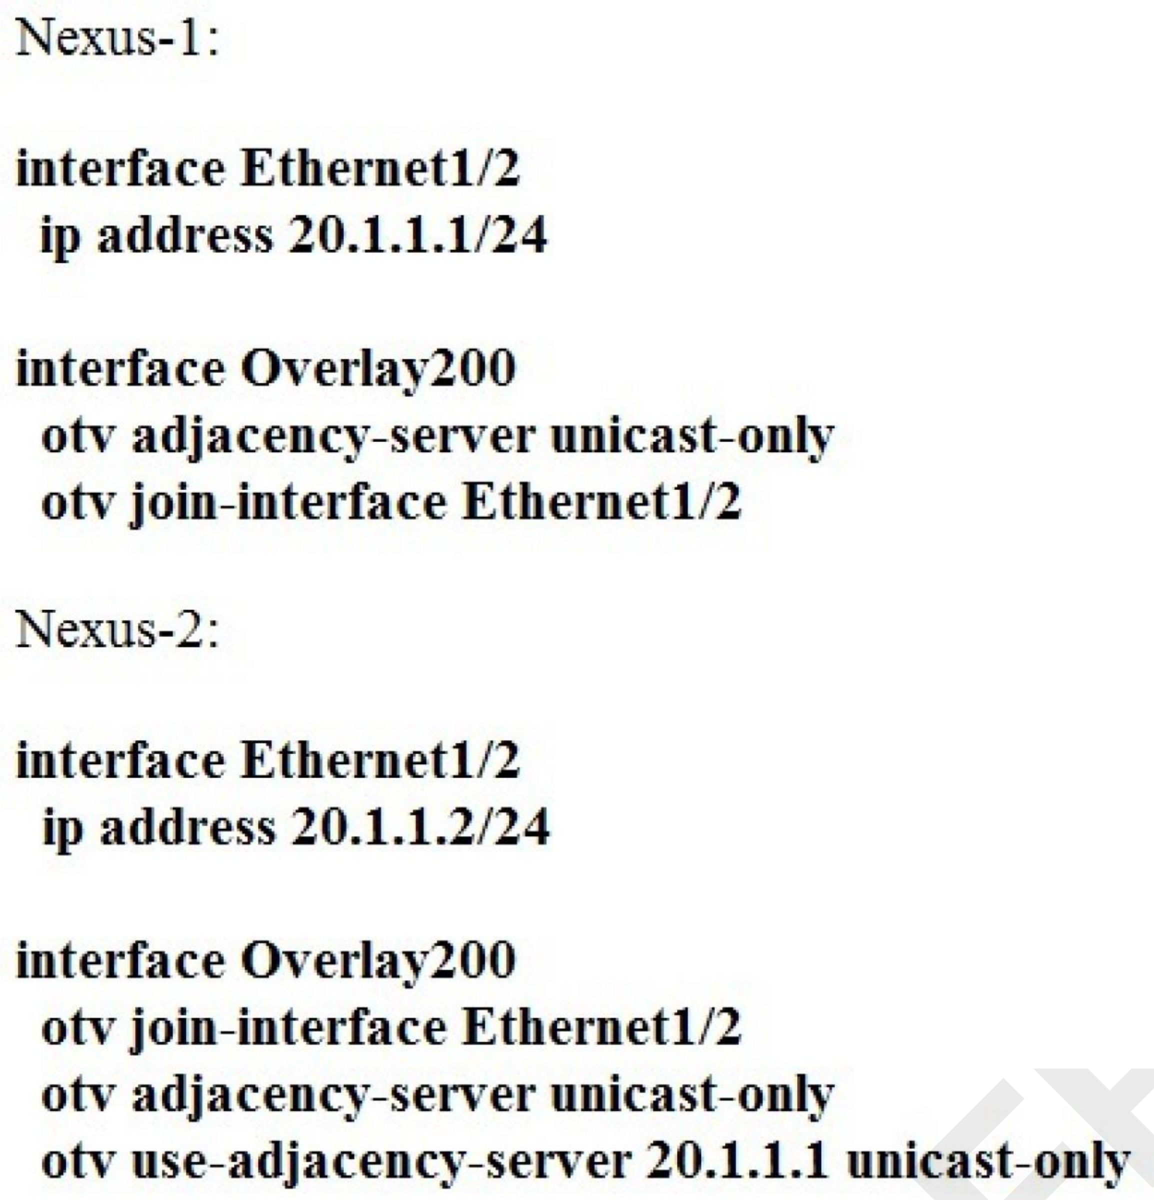

```
Nexus-1:

interface Ethernet1/2
  ip address 20.1.1.1/24

interface Overlay200
  otv adjacency-server unicast-only
  otv join-interface Ethernet1/2

Nexus-2:

interface Ethernet1/2
  ip address 20.1.1.2/24

interface Overlay200
  otv join-interface Ethernet1/2
  otv adjacency-server unicast-only
  otv use-adjacency-server 20.1.1.1 unicast-only
```

C.

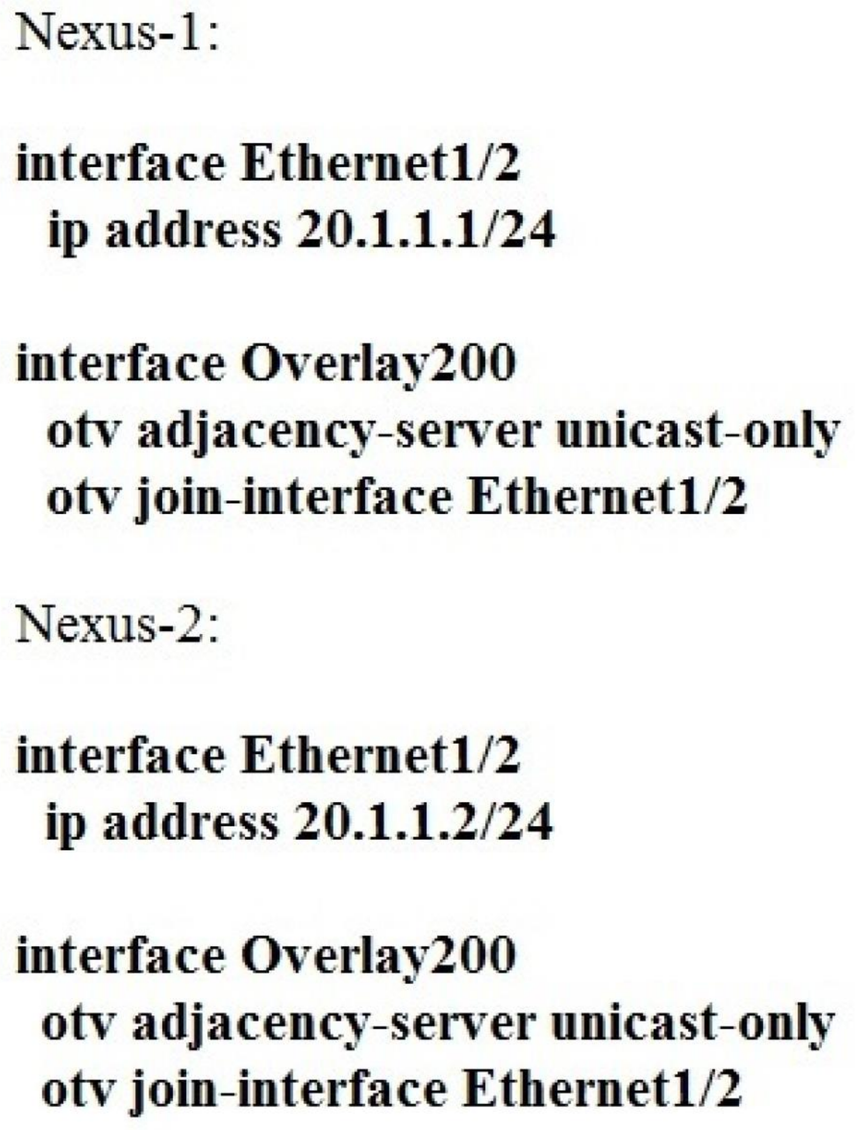

Nexus-1:

```
interface Ethernet1/2
  ip address 20.1.1.1/24

interface Overlay200
  otv adjacency-server unicast-only
  otv join-interface Ethernet1/2
```

Nexus-2:

```
interface Ethernet1/2
  ip address 20.1.1.2/24

interface Overlay200
  otv adjacency-server unicast-only
  otv join-interface Ethernet1/2
```

D.

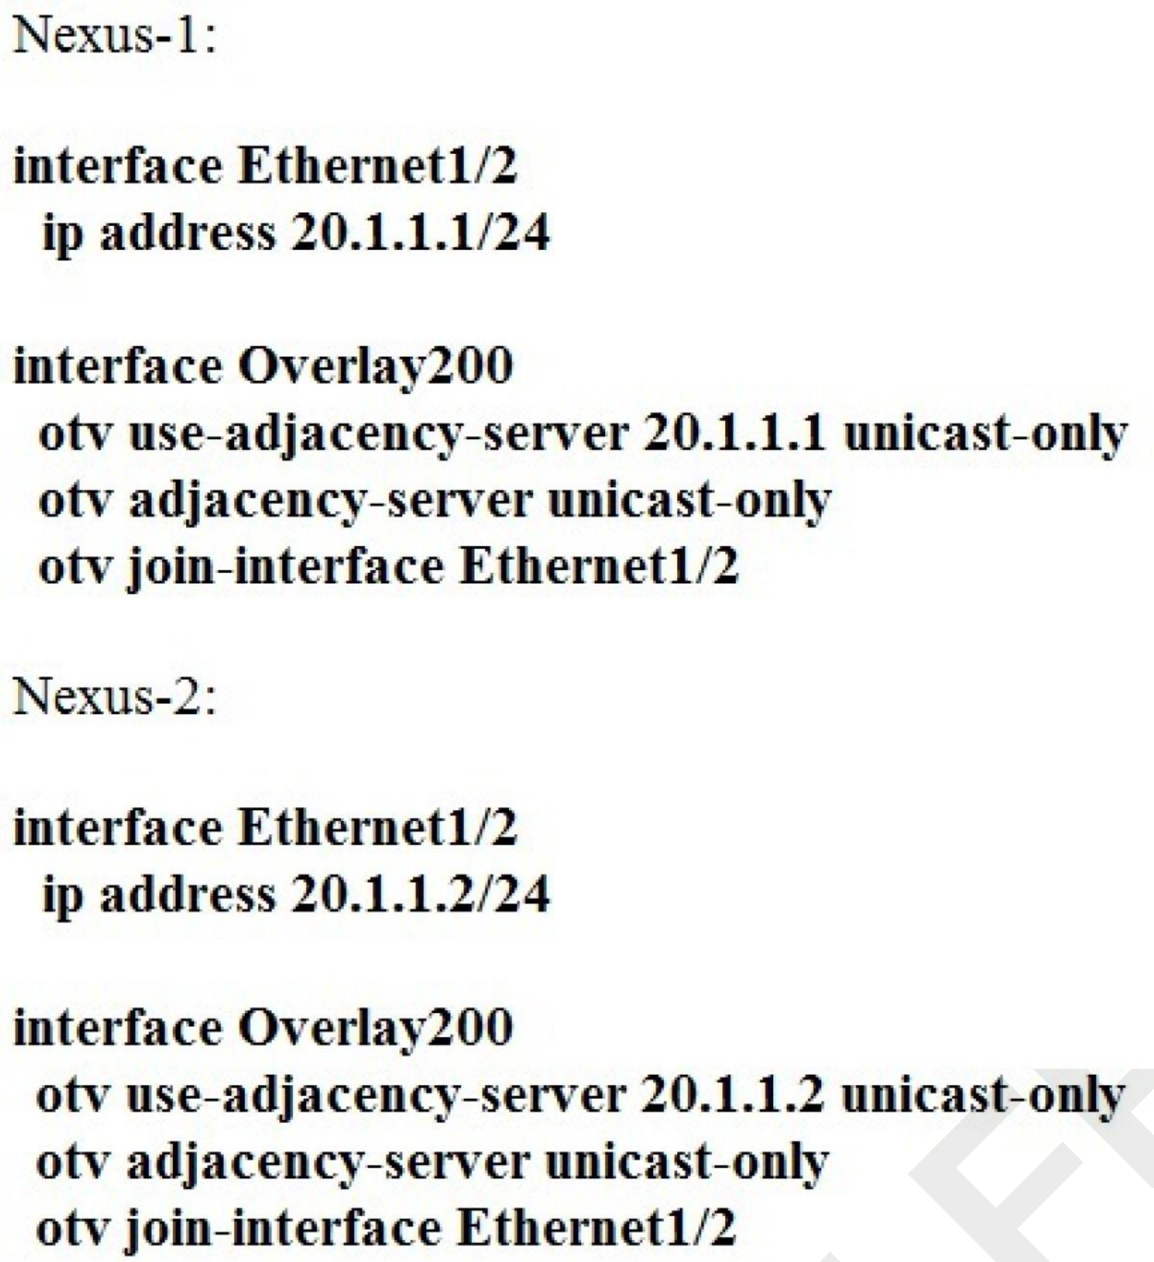

```
Nexus-1:

interface Ethernet1/2
  ip address 20.1.1.1/24

interface Overlay200
  otv use-adjacency-server 20.1.1.1 unicast-only
  otv adjacency-server unicast-only
  otv join-interface Ethernet1/2

Nexus-2:

interface Ethernet1/2
  ip address 20.1.1.2/24

interface Overlay200
  otv use-adjacency-server 20.1.1.2 unicast-only
  otv adjacency-server unicast-only
  otv join-interface Ethernet1/2
```

Answer: B

Explanation:
We need to setup a primary adjacency and a secondary adjacency. On primary, we need to setup 'otv-adjacency-server unicast-only' and on the secondary we need to configure 'otv use-adjacency-server <IP address from primary join interface> unicast only.

Question: 10

Refer to the exhibit.

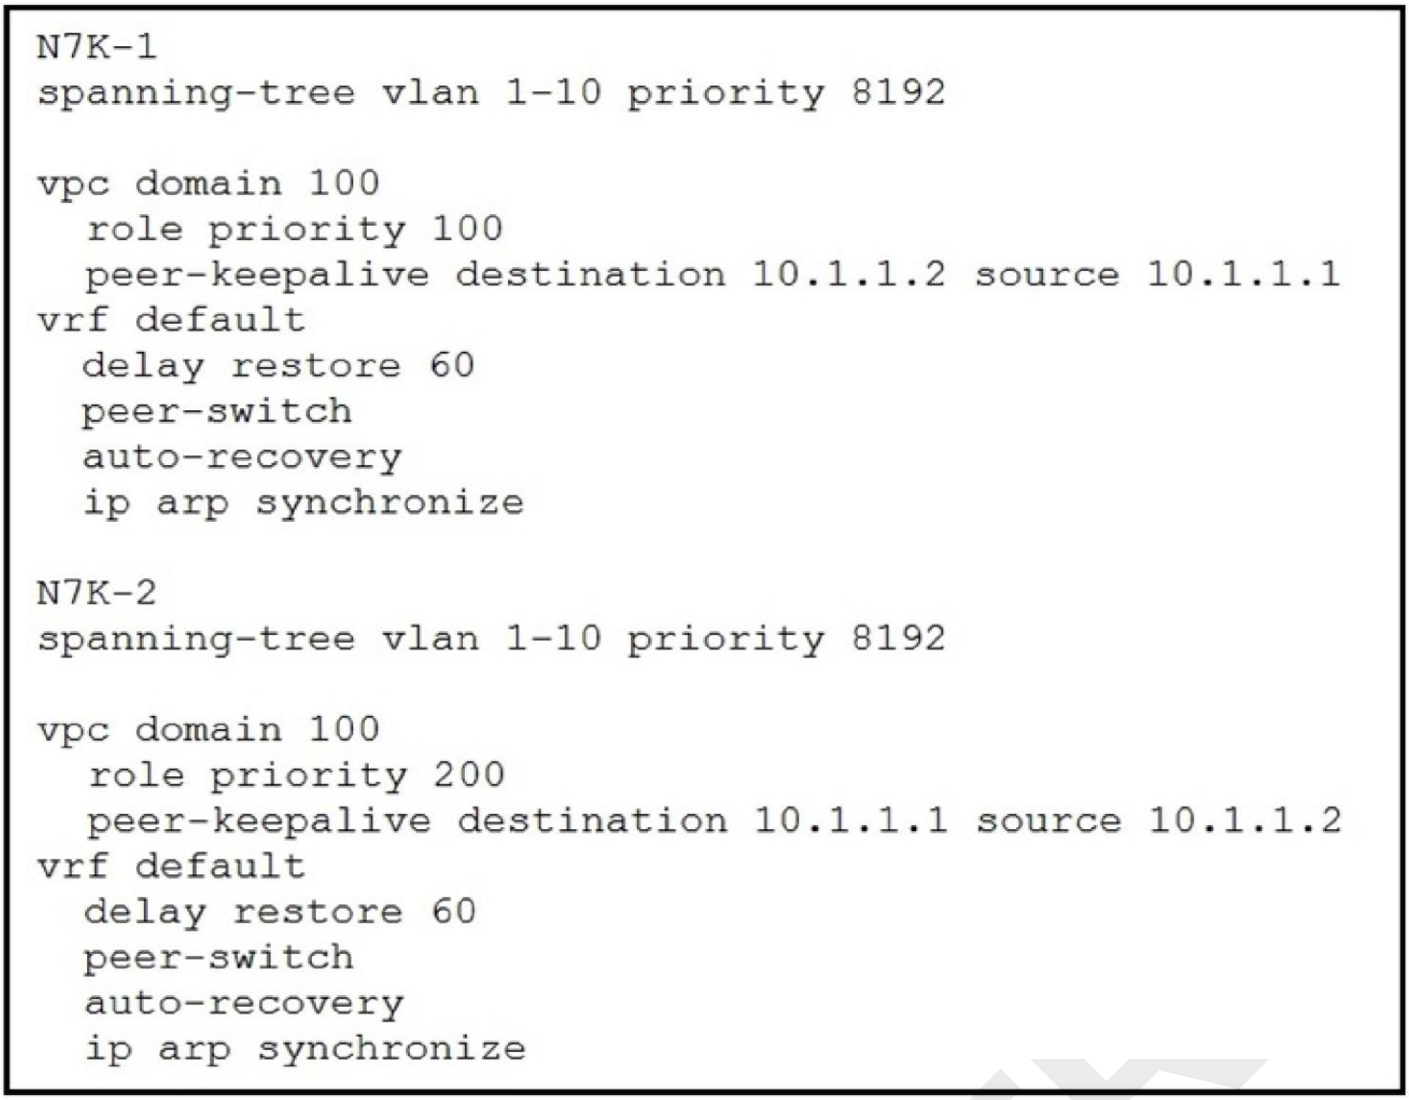

```
N7K-1
spanning-tree vlan 1-10 priority 8192

vpc domain 100
  role priority 100
  peer-keepalive destination 10.1.1.2 source 10.1.1.1
vrf default
  delay restore 60
  peer-switch
  auto-recovery
  ip arp synchronize

N7K-2
spanning-tree vlan 1-10 priority 8192

vpc domain 100
  role priority 200
  peer-keepalive destination 10.1.1.1 source 10.1.1.2
vrf default
  delay restore 60
  peer-switch
  auto-recovery
  ip arp synchronize
```

The STP priority of N7K-1 and N7K-2 are the lowest in the network.
Which statement describes STP on the vPC?

A. N7K-1 and N7K-2 appear as a single STP root.
B. N7K-1 appears as the STP root.
C. N7K-1 preempts N7K-2 as the STP root.
D. N7K-2 appears as the STP root.

Answer: A

Explanation:
Peer-switch - Enables the vPC switch pair to appear as a single STP root in the Layer 2 topology. Use the no form of the command to disable the peer switch vPC topology. If both switches do not have peer-switch enable, the question is not informing the bridge-ID (mac-address + priority) to see which switch would be the root anyway.

Question: 11

The engineer must configure SPAN on a Cisco Nexus 5000 Series Switch to get a capture of the traffic from these applications for an in-depth packet analysis.
Which two SPAN characteristics must be considered? (Choose two.)

A. SPAN source ports can be monitored in multiple SPAN sessions.

B. The Ethernet, FC, vFC, port channel, SAN port channel can be used as SPAN source ports. C. A
SPAN source port cannot be a destination SPAN port.

D. Only Ethernet, FC, vFC, port channel port types can be used as SPAN source ports.

E. The rx/tx option is available for VLAN or VSAN SPAN sessions.

Answer: BC

Explanation:

Here's a detailed justification for why options B and C are the correct answers, and why options A, D, and E are incorrect, when configuring SPAN on a Cisco Nexus 5000 Series switch:

B. The Ethernet, FC, vFC, port channel, SAN port channel can be used as SPAN source ports. This statement is correct. Cisco Nexus switches, including the 5000 series, support a wide range of interfaces as SPAN sources. These include standard Ethernet ports, Fibre Channel (FC) interfaces, virtual Fibre Channel (vFC) interfaces, and both regular and SAN port channels (aggregated links). This allows for comprehensive traffic monitoring across different types of network connectivity within the data center environment. This flexibility is essential for capturing traffic from various application flows and storage networks.

C. A SPAN source port cannot be a destination SPAN port. This is also a crucial restriction within Cisco SPAN implementations. A single port cannot simultaneously act as both the source of mirrored traffic and the destination receiving the mirrored traffic. This is to avoid loops and potential network instability. This limitation ensures the integrity of the monitoring process and prevents the mirrored traffic from affecting the original network flow.

A. SPAN source ports can be monitored in multiple SPAN sessions. This statement is incorrect. A SPAN source port can only be a part of one SPAN session. If a port needs to be monitored across multiple sessions, you must use a tool that can capture traffic from a single port and redistribute it to multiple destinations (such as a network tap or specialized network analysis tool), not through multiple SPAN sessions directly.

D. Only Ethernet, FC, vFC, port channel port types can be used as SPAN source ports. This statement is incorrect. While these are common source port types, Cisco Nexus switches may also support other interface types as SPAN sources depending on the specific model and software version. Therefore, 'only' is inaccurate and too restrictive of what the Nexus can support.

E. The rx/tx option is available for VLAN or VSAN SPAN sessions. This statement is incorrect. The rx (receive) and tx (transmit) options are used for specifying which traffic directions to mirror when defining SPAN on physical interfaces, not VLANs or VSANs. VLAN/VSAN SPAN sessions capture all traffic within a VLAN or VSAN. To capture only RX or TX on the underlying interfaces for a VLAN you have to monitor the physical interface instead of the VLAN.

Key Concepts and Why They Matter:

SPAN (Switched Port Analyzer), also known as port mirroring, is a crucial network monitoring feature in cloud environments. It enables network administrators to copy traffic from specified ports (the source) to a designated port (the destination) for analysis. This captured data is used for tasks such as performance monitoring, troubleshooting, and security analysis. The ability to select various source ports (physical interfaces and virtual channels) enhances the administrator's capability to capture the required network data for an effective analysis.

Authoritative Links for Further Research:

Cisco Nexus 5000 Series Switch Configuration Guide: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/configuration/guide/cli/CLIConfigurat
Cisco SPAN Configuration Best Practices:https://www.cisco.com/c/en/us/support/docs/switches/catalyst-6500-series-switches/10588-48.html (while for Catalyst 6500, the concepts of SPAN still apply)
Understanding Nexus SPAN:https://community.cisco.com/t5/networking-documents/understanding-nexus-span/ta-p/3135373

Question: 12

Host1 is in VLAN100 located in DataCenter1 and Host2 is in VLAN200 located in DataCenter2. Which OTV VLAN mapping configuration allows Layer 2 connectivity between these two hosts? A.

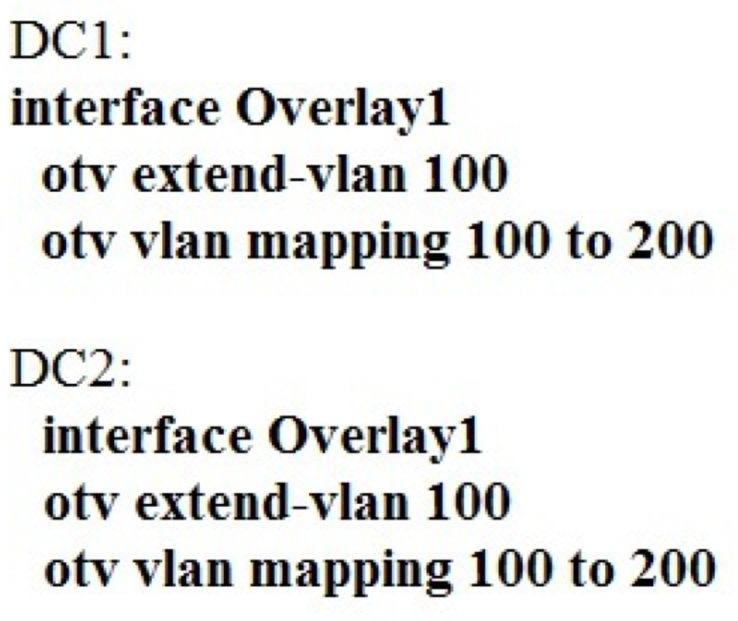

```
DC1:
interface Overlay1
  otv extend-vlan 100
  otv vlan mapping 100 to 200

DC2:
  interface Overlay1
  otv extend-vlan 100
  otv vlan mapping 100 to 200
```

B.

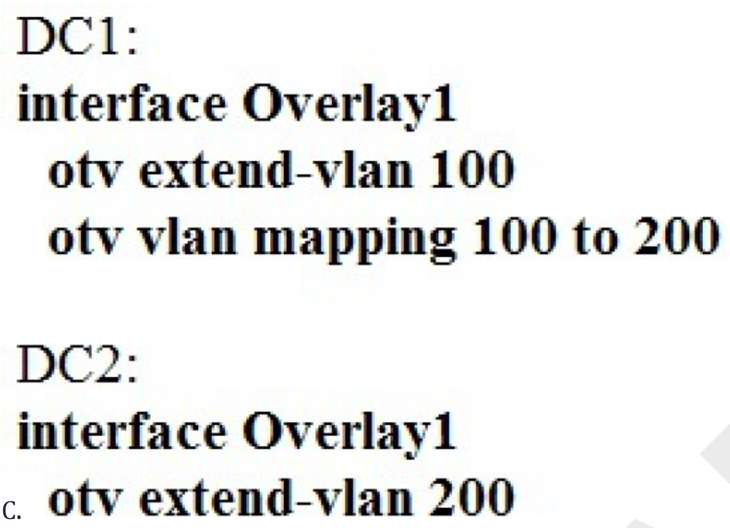

```
DC1:
interface Overlay1
  otv extend-vlan 100
  otv vlan mapping 100 to 200

DC2:
interface Overlay1
  otv extend-vlan 200
```

C.

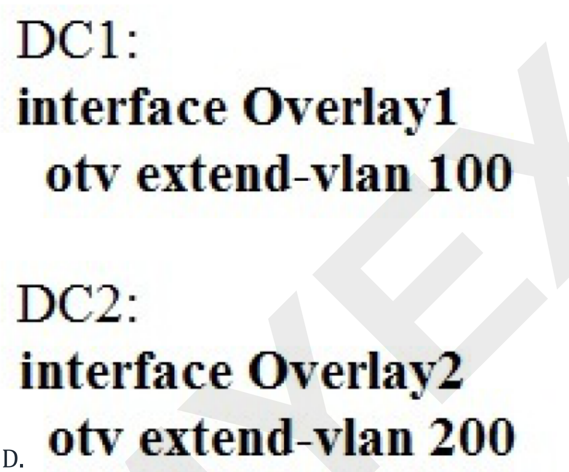

```
DC1:
interface Overlay1
  otv extend-vlan 100

DC2:
interface Overlay2
  otv extend-vlan 200
```

D.

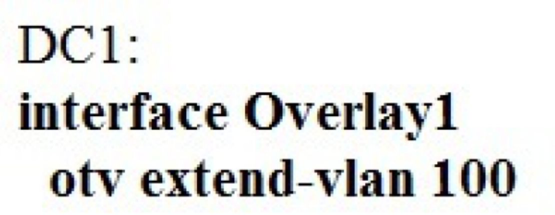

```
DC1:
interface Overlay1
  otv extend-vlan 100
```

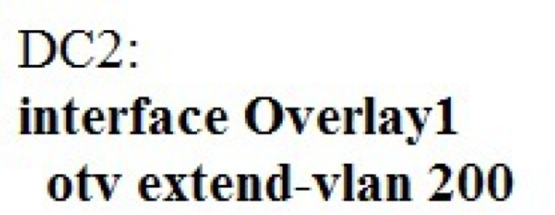

```
DC2:
interface Overlay1
  otv extend-vlan 200
```

Answer: B

Explanation:
You can map a VLAN on the local site to a VLAN with a different VLAN ID on the remote site. When you map two VLANs with different VLAN IDs across sites, they get mapped to a common VLAN called the transport VLAN. For example, when you map VLAN 1 on Site A to VLAN 2 on Site B, both VLANs are mapped to a transport VLAN. All traffic originating from VLAN 1 on Site A is translated as going from the transport VLAN. All traffic arriving at Site B from the transport VLAN is translated to VLAN 2.

Reference:
https://www.cisco.com/c/en/us/support/docs/switches/nexus-7000-series-switches/200998-Nexus-7000-OT V-VLAN-Mapping-on-Overlay.html? dtid=osscdc000283

Question: 13

Refer to the exhibit.

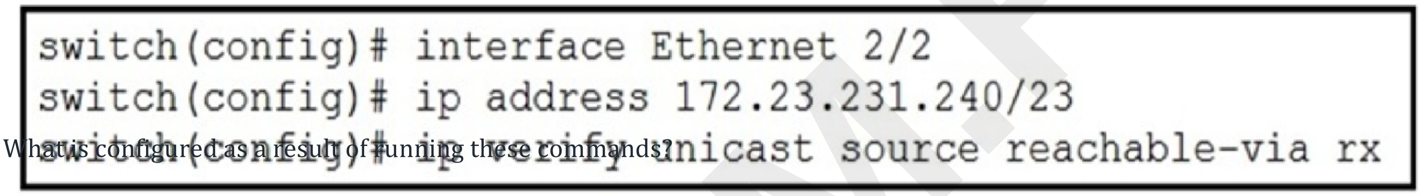

```
switch(config)# interface Ethernet 2/2
switch(config)# ip address 172.23.231.240/23
switch(config)# ip verify unicast source reachable-via rx
```

What is configured as a result of running these commands?

A. loose unicast RPF

B. strict unicast RPF

C. IP Source Guard

D. reverse lookup for outbound packets

Answer: B

Explanation:
The ip verify unicast source reachable-via rx command enables Unicast RPF in strict mode. To enable loose mode, administrators can use the any option to enforce the requirement that the source IP address for a packet must appear in the routing table.

Reference:
https://www.certyiq.com/discussions/cisco/view/62014-exam-350-601-topic-1-question-13-discussion/

Question: 14

Which configuration implements static ingress replication?

A.

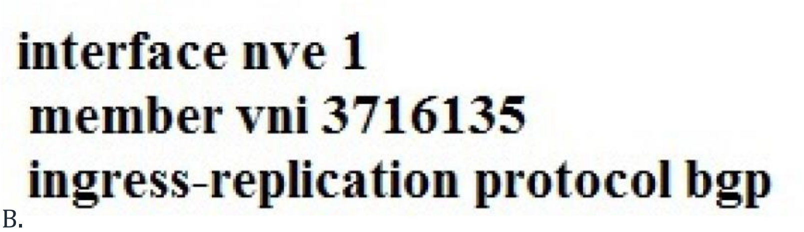

```
interface nve 1
 member vni 3716135
 ingress-replication protocol bgp
```

B.

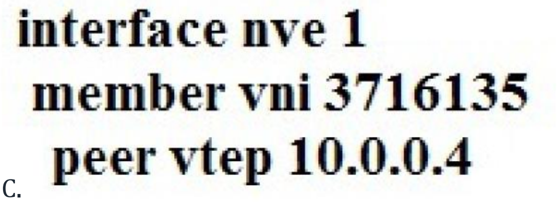

```
interface nve 1
 member vni 3716135
  peer vtep 10.0.0.4
```

C.

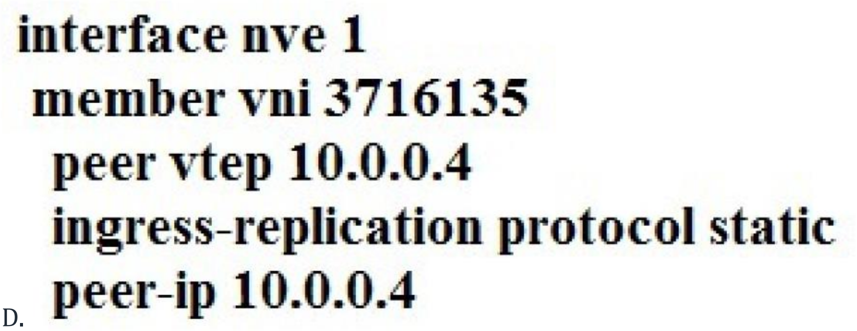

```
interface nve 1
 member vni 3716135
  peer vtep 10.0.0.4
  ingress-replication protocol static
  peer-ip 10.0.0.4
```

D.

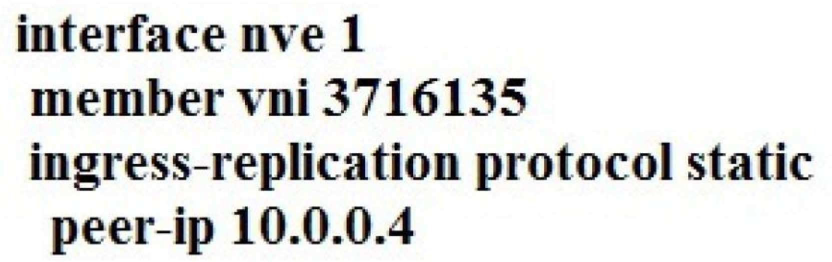

```
interface nve 1
 member vni 3716135
 ingress-replication protocol static
  peer-ip 10.0.0.4
```

Answer: D

Explanation:

The following enables static ingress replication for peers.

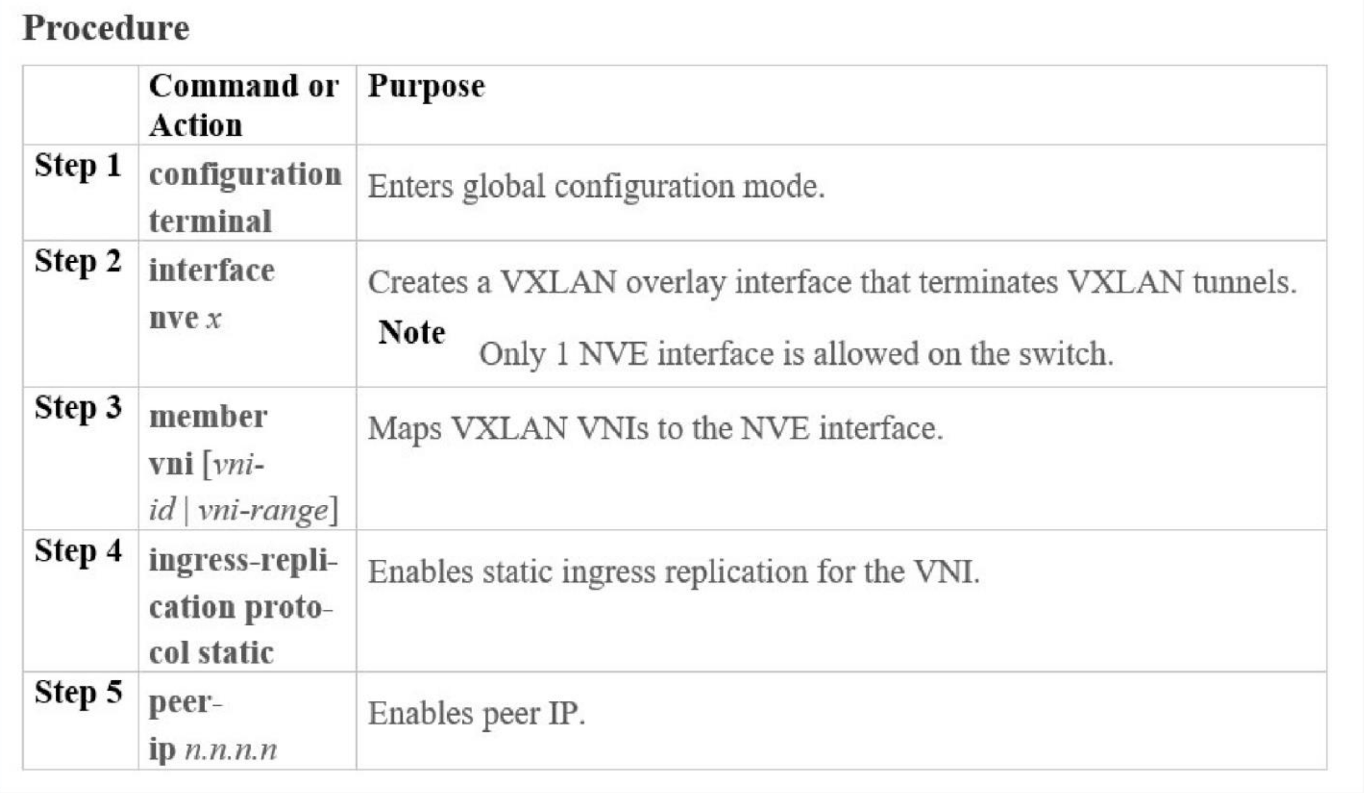

Procedure

	Command or Action	Purpose
Step 1	**configuration terminal**	Enters global configuration mode.
Step 2	**interface nve** *x*	Creates a VXLAN overlay interface that terminates VXLAN tunnels. **Note** Only 1 NVE interface is allowed on the switch.
Step 3	**member vni** [*vni-id* | *vni-range*]	Maps VXLAN VNIs to the NVE interface.
Step 4	**ingress-replication protocol static**	Enables static ingress replication for the VNI.
Step 5	**peer-ip** *n.n.n.n*	Enables peer IP.

Reference:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/vxlan/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_VXLAN_Configuration_Guide_7x/b_Cisco_Nexus_9000_Series_NX-OS_VXLAN_Configuration_Guide_7x_chapter_011.html

Question: 15

Refer to the exhibit.

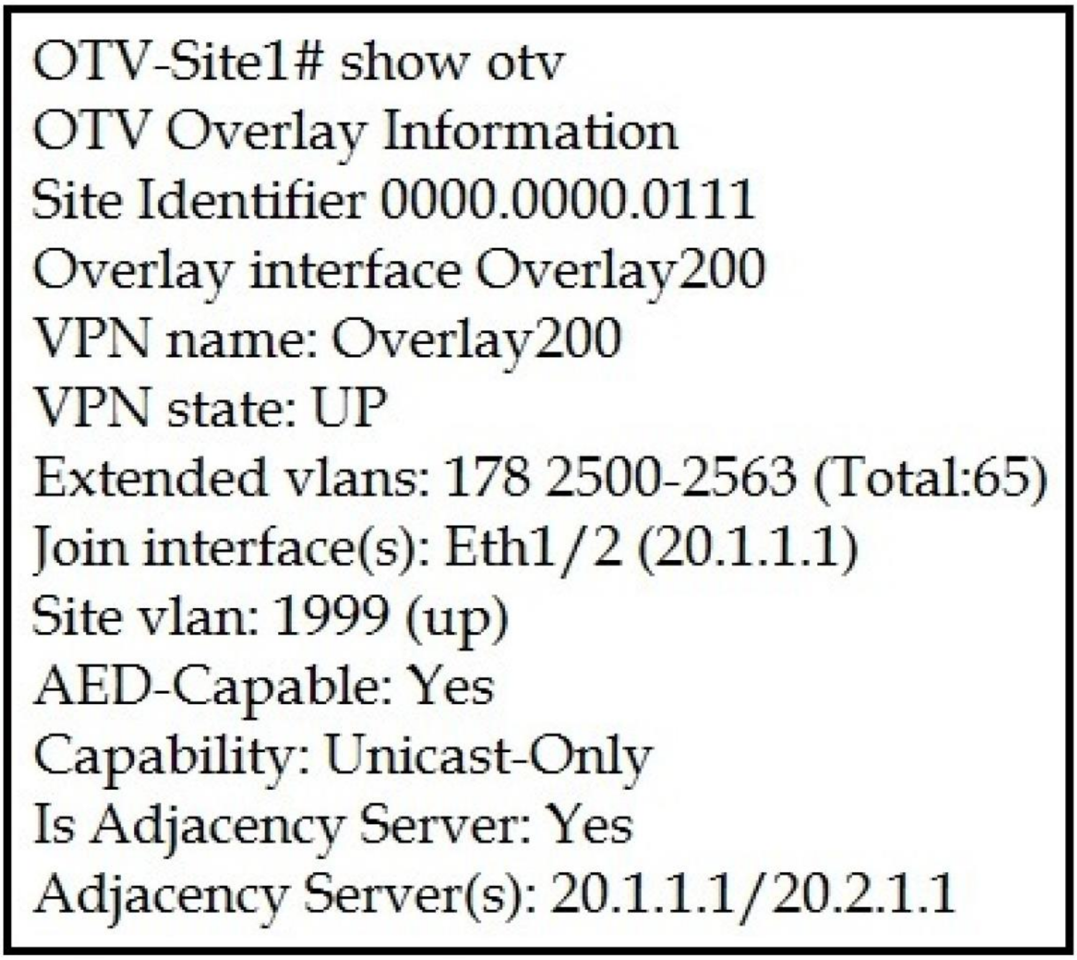

```
OTV-Site1# show otv
OTV Overlay Information
Site Identifier 0000.0000.0111
Overlay interface Overlay200
VPN name: Overlay200
VPN state: UP
Extended vlans: 178 2500-2563 (Total:65)
Join interface(s): Eth1/2 (20.1.1.1)
Site vlan: 1999 (up)
AED-Capable: Yes
Capability: Unicast-Only
Is Adjacency Server: Yes
Adjacency Server(s): 20.1.1.1/20.2.1.1
```

A network engineer is setting up a multihomed OTV network. The first site has been set up with a primary and secondary adjacency server.
Which configuration must be added on the remote OTV AEDs site? A.

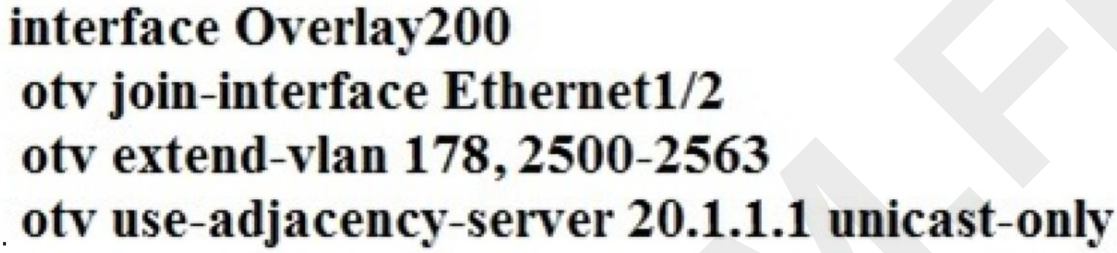

```
interface Overlay200
 otv join-interface Ethernet1/2
 otv extend-vlan 178, 2500-2563
 otv use-adjacency-server 20.1.1.1 unicast-only
```

B.

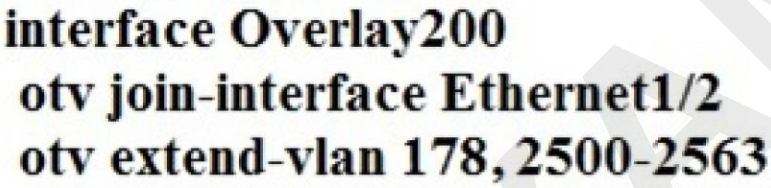

```
interface Overlay200
 otv join-interface Ethernet1/2
 otv extend-vlan 178, 2500-2563
```

C.

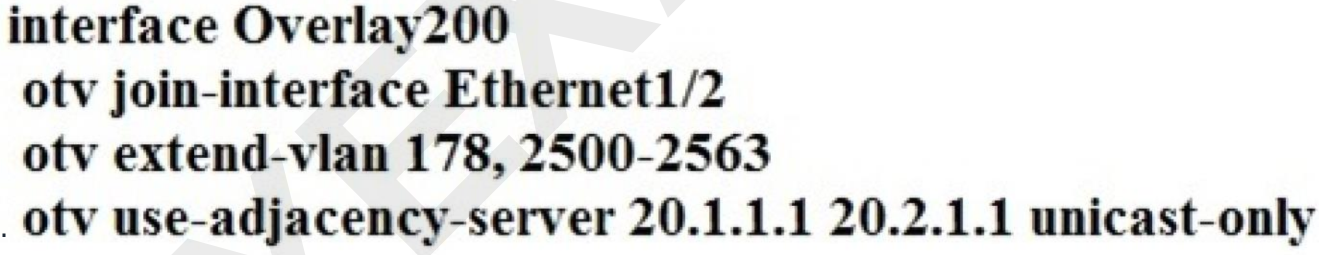

```
interface Overlay200
 otv join-interface Ethernet1/2
 otv extend-vlan 178, 2500-2563
 otv use-adjacency-server 20.1.1.1 20.2.1.1 unicast-only
```

D.

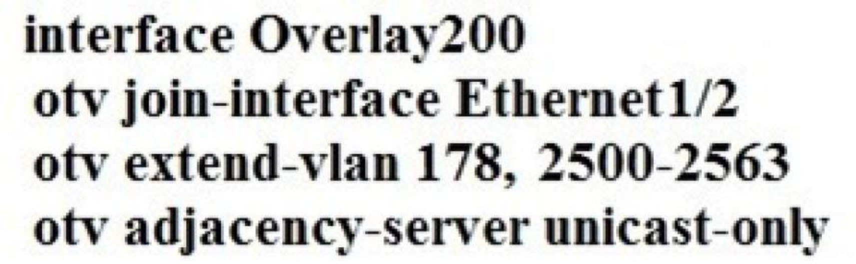

```
interface Overlay200
 otv join-interface Ethernet1/2
 otv extend-vlan 178, 2500-2563
 otv adjacency-server unicast-only
```

Answer: C

Explanation:
We need to add both IP's for the primary and secondary adjacency server when using a unicast only design.

Question: 16

A customer has a requirement to deploy a cloud service and needs to have full control over the underlying OS, data and applications.
Which cloud model meets this requirement?

A. MaaS
B. PaaS
C. SaaS
D. IaaS

Answer: D

Explanation:

The correct answer is Infrastructure as a Service (IaaS). IaaS provides the fundamental building blocks of computing infrastructure – servers, storage, and networking – via the internet. This model grants the highest level of control to the customer because it essentially delivers the raw hardware on which they install their chosen operating system, applications, and data. Unlike other cloud models, IaaS doesn't pre-configure or manage higher-level components like the operating system, allowing for complete customization. This contrasts with Platform as a Service (PaaS), which provides a platform for application development and deployment, abstracting away the underlying infrastructure and OS. Software as a Service (SaaS) delivers ready-to-use applications over the internet, granting minimal control over the underlying technology.

Managed as a Service (MaaS), though sometimes used for specific purposes like managed devices, is not a standard cloud computing model and doesn't fit this particular control-based requirement. Therefore, IaaS is the sole option enabling the customer to retain full control over the OS, data, and applications.

Further research can be conducted using resources like:

NIST Definition of Cloud Computing:https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf (Page 2 defines the cloud models)
Microsoft Azure Cloud Models:https://azure.microsoft.com/en-us/overview/what-is-cloud-computing/#cloud-service-models
Amazon Web Services (AWS) Cloud Models:https://aws.amazon.com/what-is-cloud-computing/ (Explore the sections on different cloud models)

Question: 17

Refer to the exhibit.

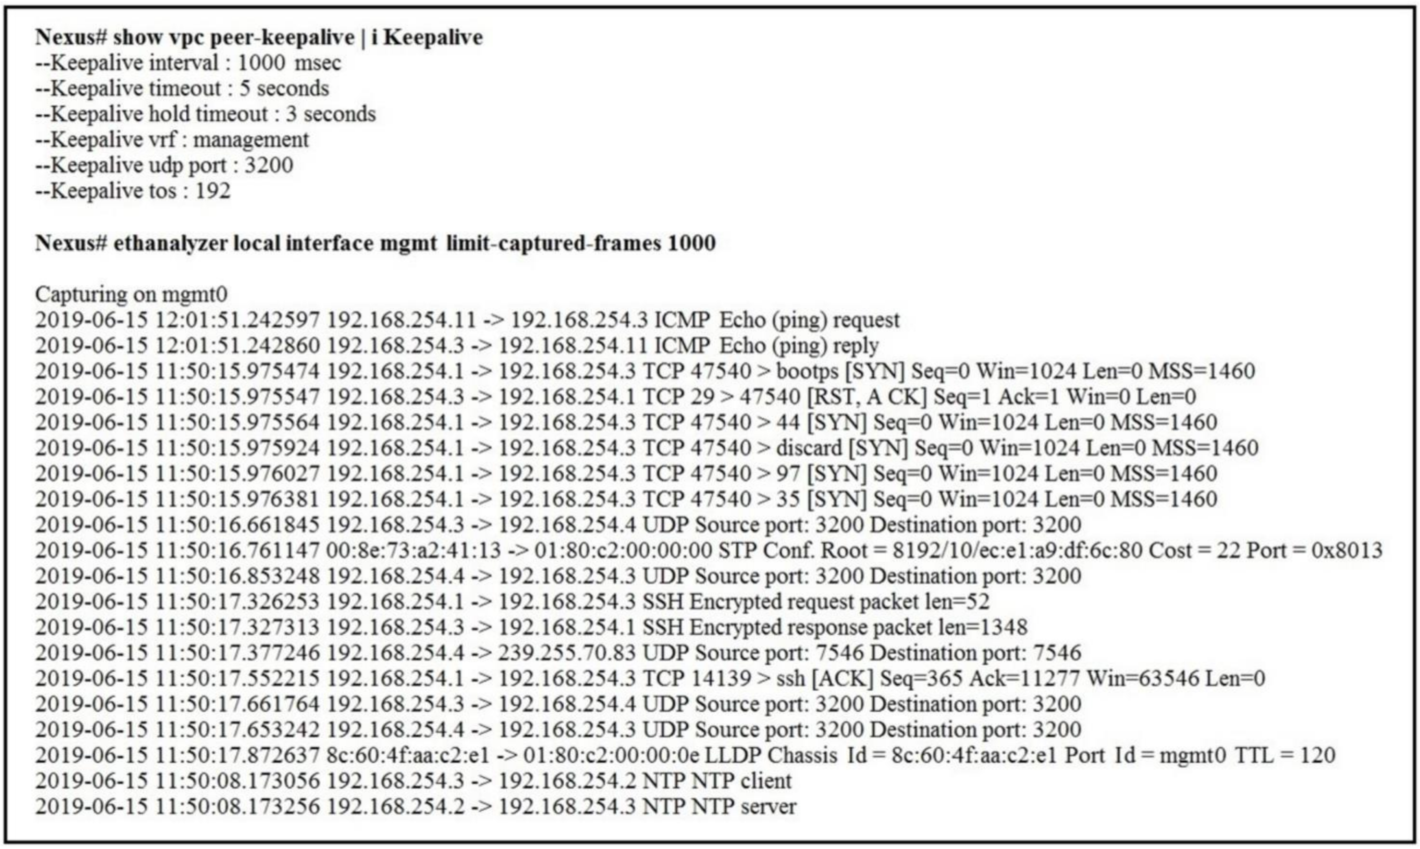

```
Nexus# show vpc peer-keepalive | i Keepalive
--Keepalive interval : 1000 msec
--Keepalive timeout : 5 seconds
--Keepalive hold timeout : 3 seconds
--Keepalive vrf : management
--Keepalive udp port : 3200
--Keepalive tos : 192

Nexus# ethanalyzer local interface mgmt limit-captured-frames 1000

Capturing on mgmt0
2019-06-15 12:01:51.242597 192.168.254.11 -> 192.168.254.3 ICMP Echo (ping) request
2019-06-15 12:01:51.242860 192.168.254.3 -> 192.168.254.11 ICMP Echo (ping) reply
2019-06-15 11:50:15.975474 192.168.254.1 -> 192.168.254.3 TCP 47540 > bootps [SYN] Seq=0 Win=1024 Len=0 MSS=1460
2019-06-15 11:50:15.975547 192.168.254.3 -> 192.168.254.1 TCP 29 > 47540 [RST, A CK] Seq=1 Ack=1 Win=0 Len=0
2019-06-15 11:50:15.975564 192.168.254.1 -> 192.168.254.3 TCP 47540 > 44 [SYN] Seq=0 Win=1024 Len=0 MSS=1460
2019-06-15 11:50:15.975924 192.168.254.1 -> 192.168.254.3 TCP 47540 > discard [SYN] Seq=0 Win=1024 Len=0 MSS=1460
2019-06-15 11:50:15.976027 192.168.254.1 -> 192.168.254.3 TCP 47540 > 97 [SYN] Seq=0 Win=1024 Len=0 MSS=1460
2019-06-15 11:50:15.976381 192.168.254.1 -> 192.168.254.3 TCP 47540 > 35 [SYN] Seq=0 Win=1024 Len=0 MSS=1460
2019-06-15 11:50:16.661845 192.168.254.3 -> 192.168.254.4 UDP Source port: 3200 Destination port: 3200
2019-06-15 11:50:16.761147 00:8e:73:a2:41:13 -> 01:80:c2:00:00:00 STP Conf. Root = 8192/10/ec:e1:a9:df:6c:80 Cost = 22 Port = 0x8013
2019-06-15 11:50:16.853248 192.168.254.4 -> 192.168.254.3 UDP Source port: 3200 Destination port: 3200
2019-06-15 11:50:17.326253 192.168.254.1 -> 192.168.254.3 SSH Encrypted request packet len=52
2019-06-15 11:50:17.327313 192.168.254.3 -> 192.168.254.1 SSH Encrypted response packet len=1348
2019-06-15 11:50:17.377246 192.168.254.4 -> 239.255.70.83 UDP Source port: 7546 Destination port: 7546
2019-06-15 11:50:17.552215 192.168.254.1 -> 192.168.254.3 TCP 14139 > ssh [ACK] Seq=365 Ack=11277 Win=63546 Len=0
2019-06-15 11:50:17.661764 192.168.254.3 -> 192.168.254.4 UDP Source port: 3200 Destination port: 3200
2019-06-15 11:50:17.653242 192.168.254.4 -> 192.168.254.3 UDP Source port: 3200 Destination port: 3200
2019-06-15 11:50:17.872637 8c:60:4f:aa:c2:e1 -> 01:80:c2:00:00:0e LLDP Chassis Id = 8c:60:4f:aa:c2:e1 Port Id = mgmt0 TTL = 120
2019-06-15 11:50:08.173056 192.168.254.3 -> 192.168.254.2 NTP NTP client
2019-06-15 11:50:08.173256 192.168.254.2 -> 192.168.254.3 NTP NTP server
```

A flapping link issue has been reported on the vPC keepalive link. A packet capture has been activated on the Cisco Nexus switch.
What is the destination IP address of the vPC keepalive packets that are sent by the switch?

A. 192.168.254.4
B. 192.168.254.1
C. 192.168.254.2
D. 239.255.70.83

Answer: A

Explanation:
3200 is the default UDP port for keepalive packets. Just look for the first line with port 3200 and note the destination IP, which is 192.168.254.4.

Reference:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus3000/sw/interfaces/92x/b-cisco-nexus-3 000-nx-os-interfaces-configuration-guide-92x/b-cisco-nexus-3000-nx-os-interfaces-configuration-guide-92x_chapter_0110.html

Question: 18

Due to a major version change, an engineer must perform a software upgrade on a Cisco Nexus Series switch. Which two technologies should be implemented to reduce disruptions to the network during the upgrade? (Choose two.)

A. vPC
B. HSRP
C. VDC
D. MLAG
E. PAgP

Answer: AB

Explanation:

The correct answer is A and B, vPC and HSRP, because they are technologies designed to enhance network availability during software upgrades on Cisco Nexus switches. vPC (virtual PortChannel) allows a switch to appear as a single logical device to connected downstream devices, and when a vPC peer switch undergoes an upgrade, traffic can seamlessly transition to the operational peer. This is achieved by the vPC peer link maintaining a synchronized forwarding state between the peers. HSRP (Hot Standby Router Protocol) provides first-hop redundancy, where multiple routers share a virtual IP address. If the active router (in this case, a Nexus switch) fails or is being upgraded, the standby router takes over, ensuring continued traffic flow. While VDC (Virtual Device Contexts) provide logical partitioning of the switch, this does not prevent disruption during software upgrades on the physical switch. MLAG (Multi-Chassis Link Aggregation) is essentially a non-Cisco term for vPC and therefore functionally the same as vPC, so it is redundant to list both.

Lastly, PAgP (Port Aggregation Protocol) is used for link aggregation but does not provide any high availability features that would reduce disruption to the network during software upgrades on a switch itself.

Therefore, vPC (for physical link redundancy) and HSRP (for routing redundancy) are the correct choices for minimizing network downtime during the upgrade process.

Cisco vPC

Cisco HSRP

Question: 19

Refer to the exhibit.

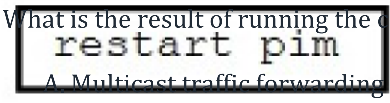

```
restart pim
```

What is the result of running the command presented?

A. Multicast traffic forwarding is suspended.

B. MRIB is flushed.

C. The PIM database is deleted.

D. PIM join messages are suspended.

Answer: C

Explanation:

When you restart PIM, the following tasks are performed:

✑ The PIM database is deleted.

✑ The MRIB and MFIB are unaffected and forwarding of traffic continues.

✑ The multicast route ownership is verified through the MRIB.

✑ Periodic PIM join and prune messages from neighbors are used to repopulate the database.

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/6-x/multicast/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_Multicast_Routing_Configuration_Guide/b_Cisco_Nexus_9000_Series_NX-OS_Multicast_Routing_Configuration_Guide_chapter_011.html

Question: 20

Refer to the exhibit.

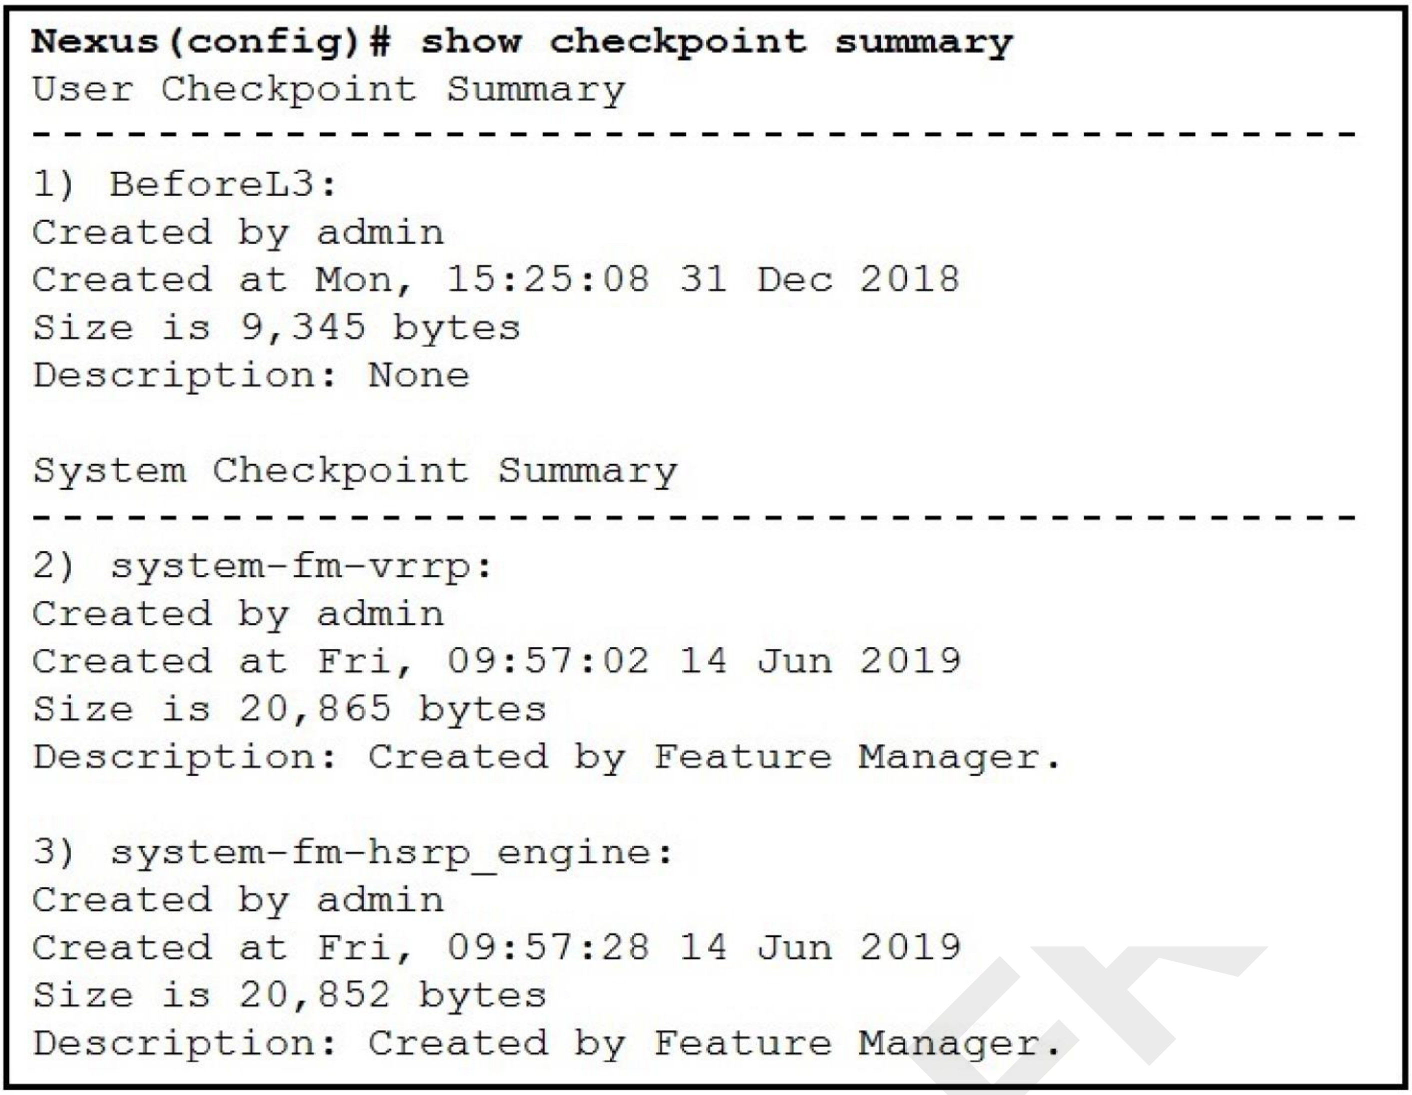

```
Nexus(config)# show checkpoint summary
User Checkpoint Summary
--------------------------------------------------------------
1) BeforeL3:
Created by admin
Created at Mon, 15:25:08 31 Dec 2018
Size is 9,345 bytes
Description: None

System Checkpoint Summary
--------------------------------------------------------------
2) system-fm-vrrp:
Created by admin
Created at Fri, 09:57:02 14 Jun 2019
Size is 20,865 bytes
Description: Created by Feature Manager.

3) system-fm-hsrp_engine:
Created by admin
Created at Fri, 09:57:28 14 Jun 2019
Size is 20,852 bytes
Description: Created by Feature Manager.
```

What is the reason the system-fm-vrrp checkpoint was created?

A. The VRRP process crashed and the checkpoint was automatically created.
B. The VRRP service restarted and the checkpoint was automatically created.
C. The network administrator manually created it.
D. The VRRP-enabled feature has been disabled.

Answer: D

Explanation:
Reasons that could trigger automated system checkpoints are highlighted below:
✑ License expiration of a feature
✑ Disabling a feature with the no feature command
✑ Removing an instance of a Layer 3 protocol
The system-generated checkpoint name convention has the format system-fm-feature. To help illustrate this automated feature we attempted to disable the VRRP feature on our Nexus 5000 therefore triggering the system to create a checkpoint. First we confirm the VRRP feature is enabled by issuing the show feature | include vrrp command then disable it and then verify it has been disabled:

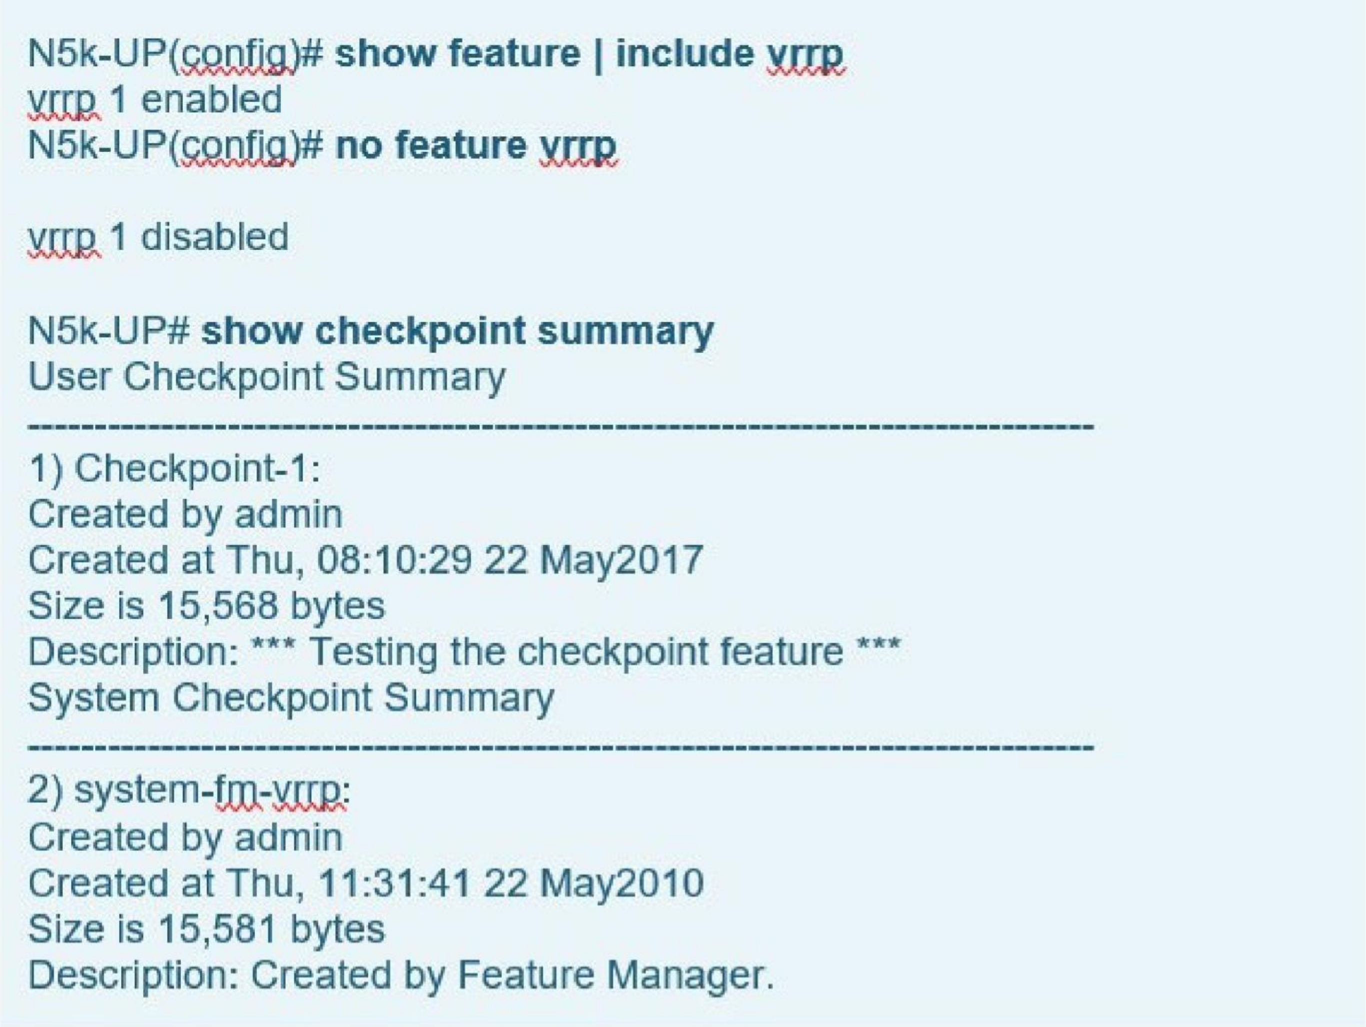

```
N5k-UP(config)# show feature | include vrrp
vrrp 1 enabled
N5k-UP(config)# no feature vrrp

vrrp 1 disabled

N5k-UP# show checkpoint summary
User Checkpoint Summary
--------------------------------------------------------------------------------
1) Checkpoint-1:
Created by admin
Created at Thu, 08:10:29 22 May2017
Size is 15,568 bytes
Description: *** Testing the checkpoint feature ***
System Checkpoint Summary
--------------------------------------------------------------------------------
2) system-fm-vrrp:
Created by admin
Created at Thu, 11:31:41 22 May2010
Size is 15,581 bytes
Description: Created by Feature Manager.
```

Notice that the system now shows a second checkpoint system-fm-vrrp which did not previously exist. This second checkpoint was created automatically by the
Nexus as soon as we disabled the vrrp feature.

Reference:
https://www.firewall.cx/cisco-technical-knowledgebase/cisco-data-center/1202-cisco-nexus-checkpoint-roll back-feature.html

Question: 21

What are two capabilities of the Cisco Network Assurance Engine? (Choose two.)

A. It validates that devices comply with network security policies.
B. It predicts the impact of changes to the network.
C. It predicts the network load on a data center.
D. It ensures that network performance meets an SLA.
E. It verifies the speed of network packet flows by using telemetry.

Answer: AB

Explanation:

The Cisco Network Assurance Engine (NAE) is a sophisticated tool designed to enhance the reliability and security of data center networks. Option A is correct because NAE continuously analyzes network configurations against defined security policies. It flags deviations from these policies, helping maintain a secure infrastructure and mitigating potential vulnerabilities. This is a crucial aspect of network security, ensuring that devices adhere to organizational security standards. Option B is also correct; NAE utilizes its understanding of the network topology and configurations to simulate the impact of proposed changes. This predictive capability enables network administrators to assess the consequences of modifications before implementation, reducing the risk of outages and other adverse effects. Options C, D, and E, however, do not accurately describe the primary functions of NAE. While NAE indirectly contributes to improved network performance, it doesn't directly predict network load or ensure SLA compliance in the way these options suggest. Furthermore, NAE does not focus on direct real-time packet flow speed measurement using telemetry data; rather, its primary function is to analyze intent and detect policy violations or potential disruptions. The combination of validation against security policies and change impact prediction makes options A and B the correct choices.

Authoritative Links for further research:

Cisco Network Assurance Engine:https://www.cisco.com/c/en/us/products/cloud-systems-management/network-assurance-engine/index.html
Cisco's solution overview:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/network-assurance-engine/index.html

Question: 22

What is an advantage of streaming telemetry over SNMP?

A. on-change traps sent to a receiver
B. periodic push-based subscription messages
C. periodic polling of the device status
D. MD5-based authentication on polling

Answer: B

Explanation:

Streaming telemetry offers a significant advantage over SNMP by utilizing a push-based model, where devices proactively send data at specified intervals or upon certain events. This is reflected in option B, "periodic push-based subscription messages." In contrast, SNMP relies on a pull-based mechanism, requiring a management station to periodically poll devices for information. This polling, described in option C ("periodic polling of the device status"), introduces latency and scalability limitations. Traps in SNMP, as mentioned in option A ("on-change traps sent to a receiver"), are only sent when a specific event occurs, lacking the continuous visibility offered by streaming telemetry. While SNMP can use authentication, option D ("MD5-based authentication on polling") does not distinguish it from streaming telemetry's security capabilities. The push model of streaming telemetry reduces network load and allows for real-time data analysis, crucial for modern cloud environments. Streaming telemetry's continuous data flow allows for better anomaly detection and proactive network management compared to the sporadic nature of SNMP polling.

Here are a few authoritative links for further research on this topic:

1. **Cisco: Streaming Telemetry:**https://www.cisco.com/c/en/us/solutions/data-center-virtualization/streaming-telemetry.html
2. **Juniper Networks: What is Streaming Telemetry?**https://www.juniper.net/us/en/research-library/glossary/what-is-streaming-telemetry.html
3. **IETF: RFC 7950, The Network Configuration Protocol (NETCONF):** https://datatracker.ietf.org/doc/html/rfc7950 (while not directly about streaming telemetry, it's related to the configuration and data retrieval mechanisms often used with it).

These links provide more information on the concepts and benefits of streaming telemetry in modern networking.

Question: 23

APIC EPG Resolution Immediacy is set to "Immediate".
Which statement is true about the Deployment Immediacy for VMM domains associated to EPGs?

A. If "On demand" is selected, the policy is programmed in the hardware only when the first packet is received through the data path.

B. If "Immediate" is selected, the policy is programmed in the hardware as soon as the leaf is booted.

C. The "Immediate" and "On demand" options require a port group to be created on the VDS.

D. If "On demand" is selected, the policy is programmed in the hardware only when the APIC detects a VM created in the EPG.

Answer: A

Explanation:

Okay, let's break down why option A is the correct answer concerning APIC EPG resolution immediacy and VMM domain deployment immediacy.

The question focuses on the interaction between Application Policy Infrastructure Controller (APIC) EPG resolution immediacy ("Immediate") and the deployment immediacy of associated VMM domains. When an EPG's resolution immediacy is set to "Immediate," it means that the endpoint to EPG mapping is programmed in the hardware immediately upon learning the MAC/IP address, not on a per-packet basis.

However, this "Immediate" setting at the EPG level doesn’t directly dictate how the VMM domain policies are deployed. VMM domain deployment immediacy operates independently, offering two options: "Immediate" and "On Demand".

Option A states: "If 'On demand' is selected, the policy is programmed in the hardware only when the first packet is received through the data path." This is accurate. When the VMM domain's deployment immediacy is set to "On Demand," the policy relevant to that domain (like VLAN configurations and port group association) is not pre-programmed in the hardware. Instead, the policy is activated and programmed in the hardware the first time a packet from that specific virtual machine (VM) in that EPG attempts to transmit. This “on-demand”approach optimizes resource usage by only programming policies for active endpoints.

Options B, C, and D are incorrect. Option B is wrong because the EPG's immediate resolution does not imply the VMM domain policies are programmed upon boot. They are programmed based on their own deployment immediacy settings. Option C is misleading; while port groups are generally created, it’s not explicitly a requirement for both "Immediate" and "On demand," it's generally a configuration step done when associating the VMM domain to the APIC. Option D is also inaccurate because on-demand deployment is triggered by the first packet, not a VM creation event. The APIC detects the endpoint activity from the data path.

In summary, while the EPG resolution immediacy determines when endpoint mappings are programmed, the VMM domain deployment immediacy controls when the policy related to that VMM domain is programmed.
"On Demand" specifically defers this programming until the first packet is seen from that associated endpoint.

Authoritative Links for Further Research:

Cisco ACI Fundamentals Guide:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/fundamentals/b-aci-fundamentals.html (Refer to sections on EPGs and VMM domains)
Cisco ACI VMM Integration Guide:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b-apic-vmm-integration-kb.html (Focus on deployment modes)

Question: 24

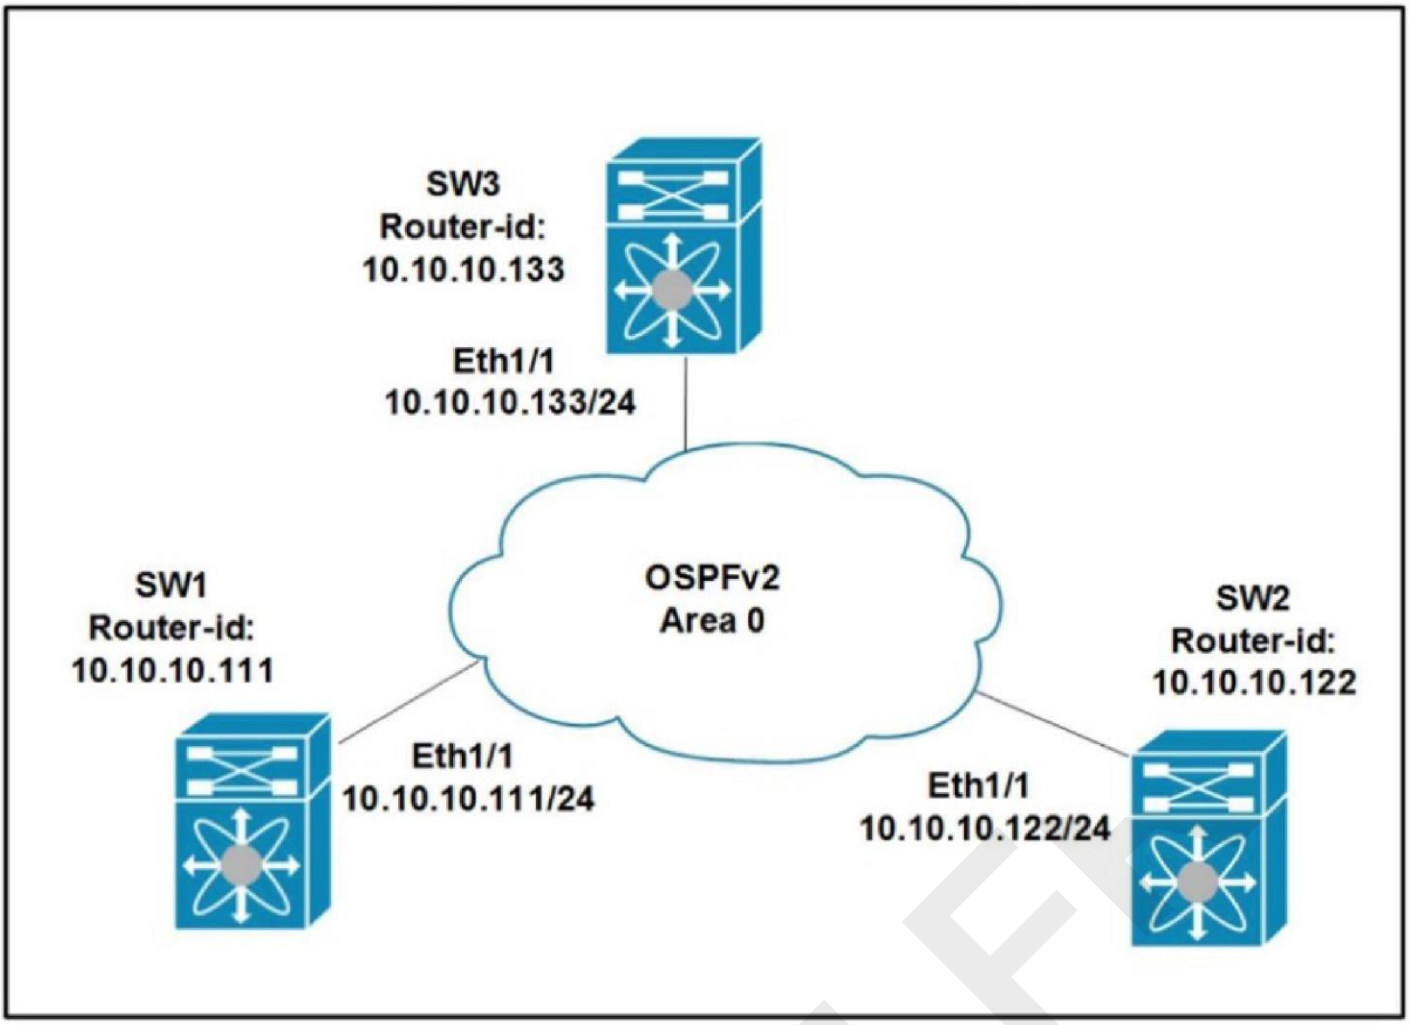

Refer to the exhibit. All switches are configured with the default OSPF priority. Which configuration should be applied to ensure that the SW2 Cisco Nexus switch controls the LSA floods and advertises the network to the remaining nodes in the OSPFv2 area?

A. SW2# configure terminal SW2 (config)# interface ethernet 1/1 SW2 (config-if)# ip ospf priority 255 B. SW2# configure terminal SW2 (config)# interface ethernet 1/1 SW2 (config-if)# ip ospf priority 1 C. SW2# configure terminal SW2 (config)# router ospf 1 SW2 (config-router)# router-id 10.10.10.22 D. SW2# configure terminal SW2 (config)# interface ethernet 1/1 SW2 (config-if)# ip ospf priority 0

Answer: A

Explanation:

The correct answer is A.

A = 255 is the highest router priority

B = It is the same (as we are doing nothing)

C = Is not the highest router ID

D = Ineligible to become the DR/BDR

Reference:

https://www.cisco.com/c/m/en_us/techdoc/dc/reference/cli/nxos/commands/ospf/ip-ospf-priority.html

Question: 25

DRAG DROP -

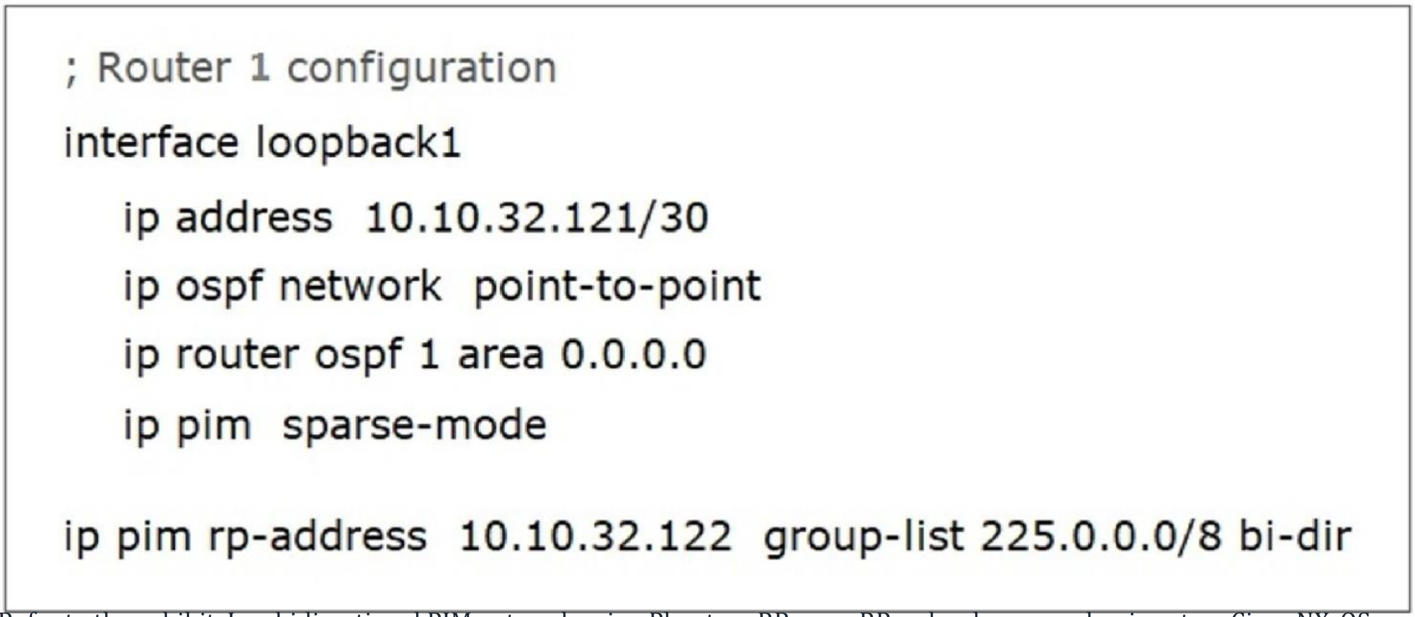

```
; Router 1 configuration
interface loopback1
   ip address  10.10.32.121/30
   ip ospf network  point-to-point
   ip router ospf 1 area 0.0.0.0
   ip pim  sparse-mode

ip pim rp-address  10.10.32.122  group-list 225.0.0.0/8 bi-dir
```

Refer to the exhibit. In a bidirectional PIM network using Phantom RP as an RP redundancy mechanism, two Cisco NX-OS routers have these requirements:

➾ R1 must be the active RP.

➾ R2 must be the backup RP that is used only if R1 is not reachable.

Drag and drop the configuration steps to complete the configuration for Router 2. Not all configuration steps are used.

Select and Place:

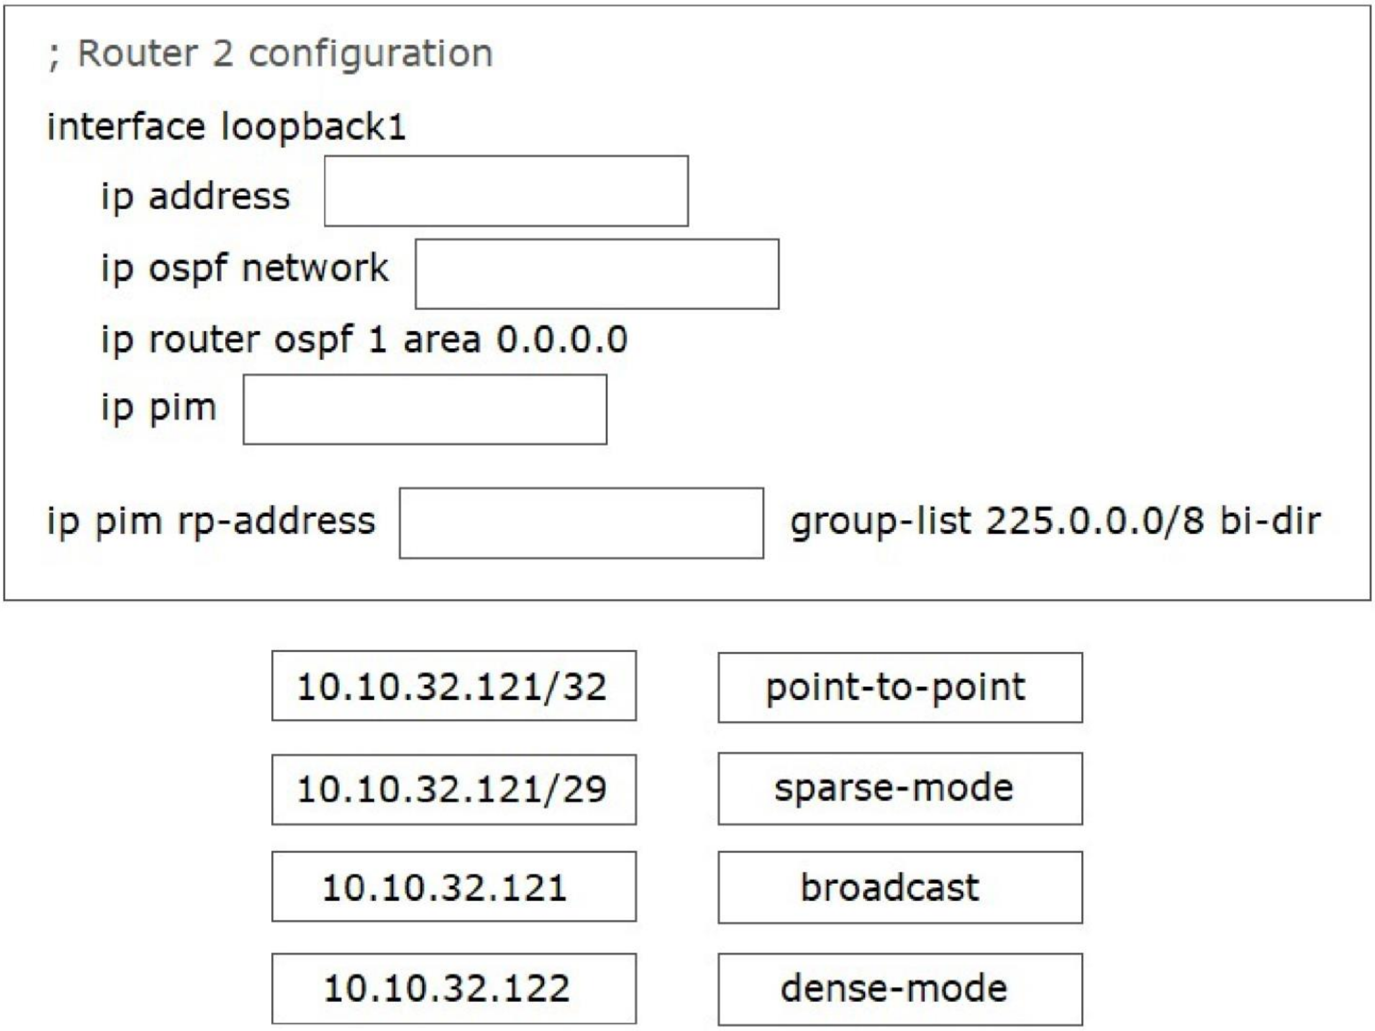

```
; Router 2 configuration
interface loopback1
    ip address
    ip ospf network
    ip router ospf 1 area 0.0.0.0
    ip pim

ip pim rp-address                group-list 225.0.0.0/8 bi-dir
```

10.10.32.121/32

point-to-point

10.10.32.121/29

sparse-mode

10.10.32.121

broadcast

10.10.32.122

dense-mode

Answer:

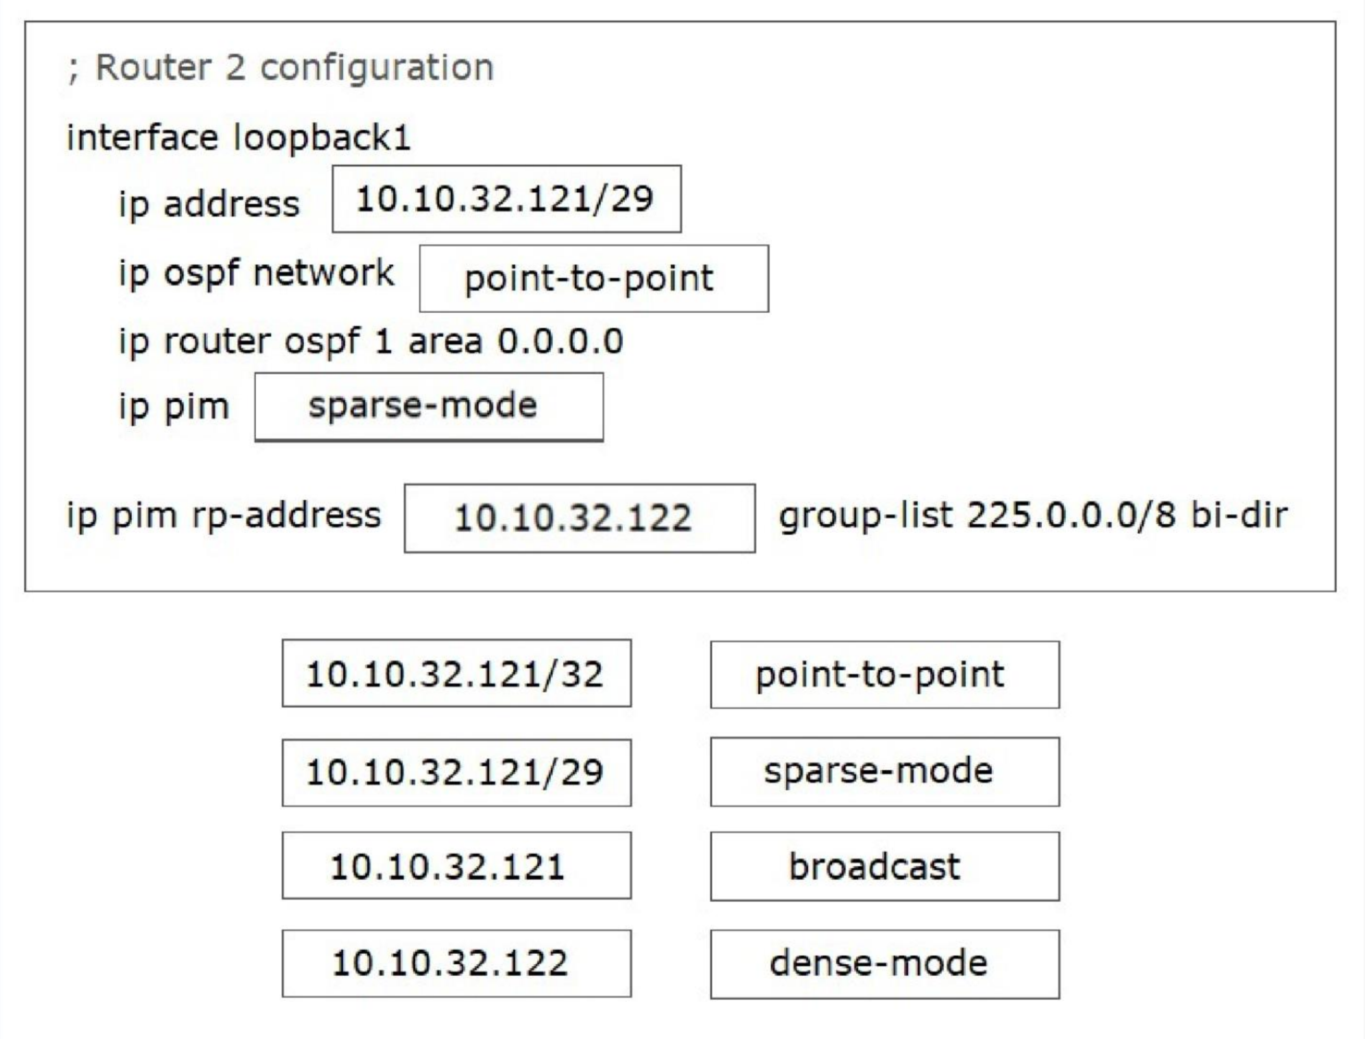

```
; Router 2 configuration
interface loopback1
    ip address  [10.10.32.121/29]
    ip ospf network  [point-to-point]
    ip router ospf 1 area 0.0.0.0
    ip pim  [sparse-mode]

ip pim rp-address  [10.10.32.122]  group-list 225.0.0.0/8 bi-dir
```

10.10.32.121/32	point-to-point
10.10.32.121/29	sparse-mode
10.10.32.121	broadcast
10.10.32.122	dense-mode

Explanation:

Reference:
https://community.cisco.com/t5/networking-documents/rp-redundancy-with-pim-bidir-phantom-rp/ta-p/3117191

Question: 26

An engineer must configure a Nexus 7000 series switch for HSRP on VLAN 100. When fully functional, the router must be the active master. Which set of commands must be used to implement the scenario?

A. feature hsrp interface vlan100 ip address 10.1.1.2 255.255.255.0 priority 255 preempt hsrp version 2 hsrp 1000 ip 10.1.1.1

B. feature-set hsrp interface vlan100 ip address 10.1.1.2 255.255.255.0 priority 20 preempt hsrp version 2 hsrp 1000 ip 10.1.1.1

C. feature-set hsrp interface vlan100 ip address 10.1.1.2 255.255.255.0 priority 80 preempt hsrp version 2 hsrp 1000 ip 10.1.1.1

D. feature hsrp interface vlan100 ip address 10.1.1.2 255.255.255.0 priority 240 preempt hsrp version 2 hsrp 1000 ip 10.1.1.1

Answer: A

Explanation:

Okay, let's break down why option A is the correct configuration for HSRP on a Nexus 7000 and why the other options are incorrect.

Justification for Option A:

feature hsrp: This command enables the Hot Standby Router Protocol (HSRP) feature on the switch, which is a prerequisite for using HSRP. Without this command, HSRP functionality is not available.

interface vlan100: This command specifies that we are configuring HSRP on VLAN 100. The router must operate as the active master within this specific VLAN.

ip address 10.1.1.2 255.255.255.0: This assigns the physical IP address to the VLAN 100 interface on the router. This is the router's actual IP address within the subnet.

priority 255: This is the critical part of achieving master status. HSRP uses a priority value to determine which router becomes active. The router with the highest priority wins. In HSRP, the maximum priority value is 255, this would be the preferred priority for the active master role.

preempt: This keyword means that if a router with a higher priority comes online after another becomes active, it will assume the active role. This ensures the highest priority router is always the active master, even if it was not first online. This ensures mastership is consistent and that the configuration achieves its purpose.

hsrp version 2: Specifies using HSRP version 2, which allows for better tracking and communication. This version is widely preferred over version 1.

hsrp 1000 ip 10.1.1.1: This creates the HSRP group "1000" and gives it a virtual IP address of 10.1.1.1. This is the IP address that other devices on VLAN 100 will use as their default gateway, it allows for redundancy, ensuring traffic is always forwarded.

Why Other Options are Incorrect:

Option B and C use feature-set hsrp: The feature-set command was used on older platforms. It is important to use feature hsrp on a Nexus 7000 as it is the command to enable the feature properly. The priority values of 20 and 80 would not ensure the router would be the master, a higher value is required.

Option D has a priority of 240: While higher than the priority value in options B and C it is not the maximum and may lead to other routers with higher priorities taking over the active role.

All Options B, C, and D do not ensure the router will be the master: The priority needs to be the highest possible (255) to ensure the router takes over the master role.

Lack of Proper Feature Activation: Options B and C utilize an outdated method of enabling HSRP, which would not function as intended.

Authoritative Links for Further Research:

Cisco HSRP Configuration Guide:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus7000/sw/layer2/configuration/guide/b_Cisco_N
Cisco HSRP Command Reference:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp/command/ipapp-cr-book/ipapp-h1.html

In summary, Option A is the only configuration that will enable HSRP, set the correct priority to make the device a master, and allows for the master role to always be maintained when a higher priority device comes online.

Question: 27

Which MAC address is an HSRP version 2?

A. 3842.4250.0000

B. 0000.0C07.AC1H

C. 0000.0C9F.F0C8

D. 0100.5E7F.FFFF

Answer: C

Explanation:

The correct MAC address for an HSRP version 2 virtual router is **0000.0C9F.F0C8 (Option C)**. Here's the breakdown:

HSRP (Hot Standby Router Protocol) uses a virtual MAC address that is shared between active and standby routers within a group. This address allows hosts to maintain communication even if the active router fails. The MAC addresses used for HSRP vary depending on the version of the protocol.

For HSRP version 1, the virtual MAC address always begins with 0000.0C07.AC, followed by a hexadecimal representation of the HSRP group number (e.g., 0000.0C07.ACxx). Option A (3842.4250.0000) is not a valid HSRP MAC address prefix, and option B (0000.0C07.AC1H) contains an invalid hex character.

HSRP version 2 uses a different, extended multicast MAC address range. HSRP v2 multicast MAC addresses always start with 0000.0C9F.F0, and end with the HSRP group number. So, the third octet is F0 + the HSRP group number. Option C (0000.0C9F.F0C8) follows this format, making it a valid HSRP version 2 MAC address with a hexadecimal representation.

Option D (0100.5E7F.FFFF) is a multicast MAC address, but it is not used by HSRP. This address is typically associated with IPv6 multicast.

In summary, HSRP version 2 uses the multicast MAC address prefix 0000.0C9F.F0, and Option C is the only address that follows this format, making it the correct answer. The other provided addresses are invalid or belong to other protocols.

Authoritative Links:

Cisco Documentation on HSRP:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp/configuration/15-mt/ipapp-15-mt-book/ipapp-hsrp.html
HSRP MAC Address Explanation:https://networklessons.com/cisco/ccna-routing-switching/hsrp-hot-standby-router-protocol (This resource provides a clear breakdown of HSRP MAC addresses)

Question: 28

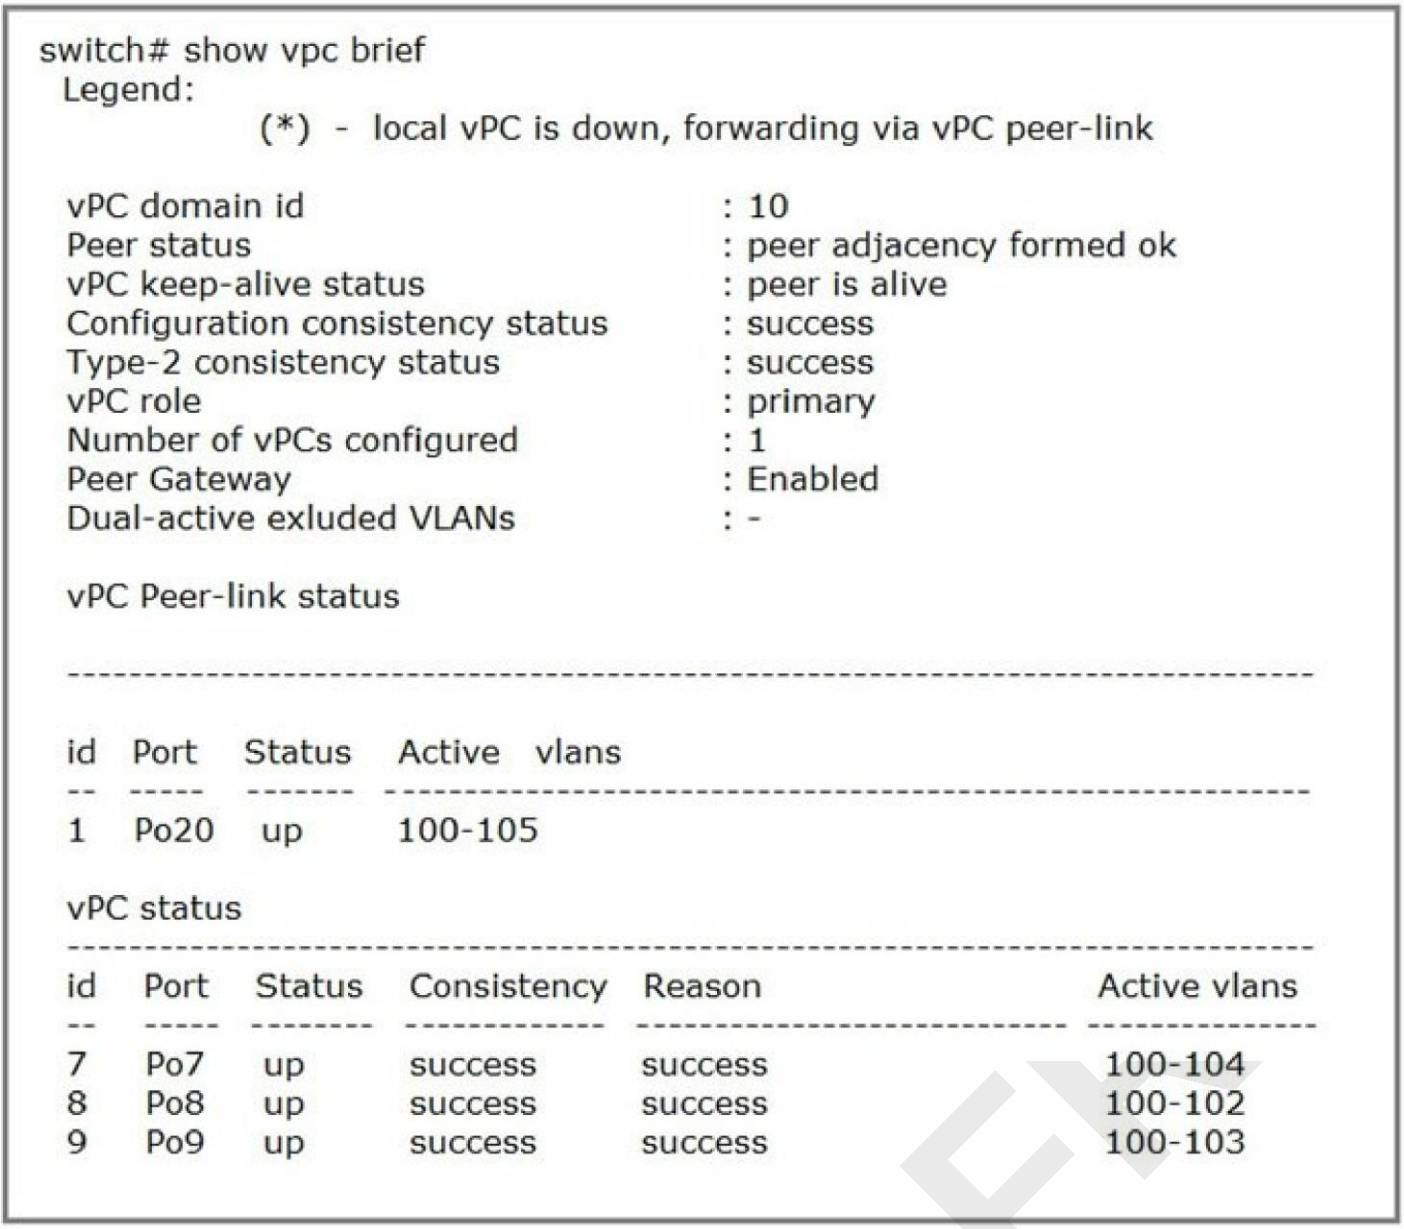

```
switch# show vpc brief
  Legend:
              (*)  -  local vPC is down, forwarding via vPC peer-link

  vPC domain id                          : 10
  Peer status                            : peer adjacency formed ok
  vPC keep-alive status                  : peer is alive
  Configuration consistency status       : success
  Type-2 consistency status              : success
  vPC role                               : primary
  Number of vPCs configured              : 1
  Peer Gateway                           : Enabled
  Dual-active exluded VLANs              : -

  vPC Peer-link status

  ---------------------------------------------------------------------

  id   Port    Status    Active   vlans
  --   -----   -------   --------------------------------------------------
  1    Po20     up       100-105

  vPC status
  ---------------------------------------------------------------------
  id   Port    Status    Consistency   Reason                     Active vlans
  --   -----   --------  ------------  -------------------------  -------------
  7    Po7      up       success       success                     100-104
  8    Po8      up       success       success                     100-102
  9    Po9      up       success       success                     100-103
```

Refer to the exhibit. Which VLANs are capable to be assigned on vPC interfaces? A. 100-
102

B. 100-103

C. 100-104

D. 100-105

Answer: D

Explanation:

D is correct.

The PortChannel connecting the vPC peers should carry all the VLANs used by the vPC member ports. So the peer-link should carry all the vlans used via the other ports.

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/layer2/503_n2_1/503_n2_1nw/
Cisco_n5k_layer2_config_gd_rel_503_N2_1_chapter8.html

Question: 29

An engineer need to implement a solution that prevents loops from occurring accidentally by connecting a switch to interface Ethernet1/1. The port is designated to be used for host connectivity. Which configuration should be implemented?

A. switch# configure terminal switch(config)# interface Ethernet1/1 switch(config-if)# spanning-tree bpduguard enable

B. switch# configure terminal switch(config)# interface Ethernet1/1 switch(config-if)# spanning-tree guard loop

C. switch# configure terminal switch(config)# interface Ethernet1/1 switch(config-if)# spanning-tree loopguard default

D. switch# configure terminal switch(config)# interface Ethernet1/1 switch(config-if)# spanning-tree bpdufilter enable

Answer: A

Explanation:

The correct answer is **A. spanning-tree bpduguard enable**. This command, applied on a port intended for host connectivity (access port), provides a security mechanism against accidental loop formation. Specifically, it protects against the situation where a user might connect a switch to this port, potentially causing a spanning tree loop.

Here's the breakdown:

Spanning Tree Protocol (STP): STP is a Layer 2 protocol that prevents loops in Ethernet networks by blocking redundant paths. However, it assumes that all switches participate in STP.

BPDU (Bridge Protocol Data Unit): Switches exchange BPDUs to establish the loop-free topology.

spanning-tree bpduguard enable: This command enables BPDU Guard on the specified interface. When the interface receives a BPDU, it's immediately placed in an err-disabled state. This action effectively shuts down the port, preventing any potential loop that would be caused by the unauthorized switch.

Purpose: The primary purpose of BPDU Guard is to safeguard access ports from unauthorized switches that could disrupt the spanning tree topology. Access ports should only connect to end devices (hosts).

Alternative Options:

B. spanning-tree guard loop: This command enables Loop Guard which is used to protect against unidirectional link failures and isn't appropriate for preventing a switch from being connected to a host port.

C. spanning-tree loopguard default: This command enables Loop Guard globally on the switch, not on an individual interface. It's also not appropriate for preventing loops from unauthorized switches on access ports. **D. spanning-tree bpdufilter enable**: This command enables BPDU Filtering, which can potentially disrupt STP operation and might not prevent loops. It should be used with extreme caution and is rarely the correct choice for access ports.

In summary, spanning-tree bpduguard enable is the best method to prevent loops on a host-connected port, safeguarding the network from accidental misconfigurations.

Authoritative Links:

Cisco Documentation on BPDU Guard:
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960x/software/15_2_2_e/configuration/guide/scg29
Spanning Tree Protocol (STP): https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/10590-79.html

Question: 30

A network engineer repeatedly saves a configuration on Catalyst switches to NVRAM using the write memory command. Which action should be taken to implement the same action on Nexus switches?

A. Use the alias command to use the write memory command.
B. Use the wri command to use the copy running-config startup-config command.
C. Use the exit command to leave the configuration mode and save the configuration automatically.
D. Use the write memory command to save the configuration.

Answer: A

Explanation:

The correct answer is **A. Use the alias command to use the write memory command.** Here's the justification:

Cisco Catalyst switches traditionally use the write memory command (often abbreviated as wr) to save the running configuration to NVRAM, making it the startup configuration used on reboot. Cisco Nexus switches, however, utilize a different command structure for this function. They use the command copy running-config startup-config to achieve the same result, explicitly copying the currently active configuration to the stored startup configuration. The wri shortcut is also not native to Nexus platforms. Direct use of the write memory command will not work on a Nexus switch.

Option B, suggesting the use of wri to invoke the copy running-config startup-config command, is incorrect. While copy running-config startup-config is the correct command, the wri shortcut is not built in. Option C, stating that the exit command automatically saves the configuration, is also incorrect; the exit command only exits the configuration mode. It does not, by itself, save any changes.

Option D stating direct use of write memory on a nexus switch is incorrect since this command will result in an error message.

To bridge the familiarity of Catalyst users with the Nexus command structure, Cisco allows the creation of aliases. Aliases enable you to define a custom command that represents another, more complex command. By creating an alias that maps write memory to copy running-config startup-config, Nexus switches can mimic the traditional Catalyst behavior. This is achieved using the alias command in global configuration mode. For example, the configuration line alias exec write copy running-config startup-config would enable a user to type write memory and have the Nexus switch execute the copy running-config startup-config command.

Therefore, using an alias to map the familiar write memory command to the Nexus equivalent is the best approach to provide a user-friendly experience to those transitioning from Catalyst to Nexus platforms. This doesn't change the underlying command but rather provides a wrapper to use a known command on the new platform.

Authoritative Links for Further Research:

Cisco Nexus Command Reference Guide:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/command/b-cisco-nexus-9000-series-command-reference.html (Search for "copy running-config startup-config" and "alias" for specific command details)
Cisco Nexus Configuration Fundamentals:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/configuration/fundamentals/b-cisco-nexus-9000-series-fundamentals-configuration-guide.html

Question: 31

What occurs when running the command install deactivate <filename> while a software maintenance upgrade is performed on a Cisco Nexus 9000 Series switch?

A. The current set of packages is committed.
B. The package features for the line card are disabled.
C. The package is removed from the switch.
D. The current upgrade stops.

Answer: B

Explanation:

The install deactivate <filename> command, when executed during a software maintenance upgrade on a Cisco Nexus 9000 Series switch, specifically deactivates the designated package's features for the line card. This action does not remove the package from the switch; instead, it disables its functionality. Option A is incorrect as package commitment happens through different commands. Option C is also incorrect as the package remains on the switch, though its functionality is disabled. Option D is incorrect; the upgrade process would not typically be stopped by a deactivation command during an upgrade, though the features of the
deactivated package would become unavailable. The deactivate option in Cisco NX-OS context implies disabling functionality of the installed package; the software image itself is not removed by this command. This action allows for granular control over specific package features and enables administrators to manage their operational status independent of the underlying software installation. It also aids in troubleshooting by temporarily deactivating a faulty package without affecting the entire system. The concept of module disabling is fundamental in maintaining high availability in network devices by controlling software features and functionality. The process maintains a flexible approach to updating and controlling specific features without interrupting overall switch operations.

Further research on Cisco NX-OS software management, including image management and package management, is available at:

Cisco NX-OS Software Upgrade and Downgrade Guide
Cisco NX-OS Install and Uninstall Operations

Question: 32

An engineer installed a new Nexus switch with the mgm0 interface in vrf management. Connectivity to the rest of the network needs to be tested from the guest shell of the NX-OS. Which command tests connectivity from the guest shell of the NX-OS?

A. [[email protected] ~]$ dohost ''ping vrf management 173.37.145.84''
B. [[email protected] ~]$ chvrf management ping 173.37.145.84
C. [[email protected] ~$ ping 173.37.145.84 vrf management
D. [[email protected] ~]$ iping vrf management ip 173.37.145.84

Answer: B

Explanation:

The correct answer is **B. [[email protected] ~]$ chvrf management ping 173.37.145.84**.

Here's why:

The guest shell in Cisco NX-OS operates within a Linux environment, separate from the main NX-OS control plane. When an interface is assigned to a VRF (Virtual Routing and Forwarding instance), like mgm0 in vrf management, you need to specify that VRF when running commands from the guest shell that require routing or network access. Standard Linux networking commands like ping will use the default routing table, which is not aware of the VRF configuration within NX-OS.

Option B uses chvrf, which is the specific command within the guest shell that allows you to execute commands within a particular VRF context. chvrf management sets the context to the management VRF, ensuring that the ping command is executed with the appropriate routing information associated with the mgm0 interface. This will allow the guest shell to successfully reach the target IP address, 173.37.145.84, via the configured VRF.

Option A, dohost 'ping vrf management 173.37.145.84', is incorrect. The dohost command is used to execute NX-OS commands from within the guest shell, and not the other way around. Moreover, ping vrf management is not a valid NX-OS command syntax.Option C, ping 173.37.145.84 vrf management, is incorrect because the VRF specification is not supported in the standard Linux ping command syntax and will simply treat it as another argument.Option D, iping vrf management ip 173.37.145.84, uses an invalid command iping.

Therefore, chvrf management ping <ip> is the correct way to initiate a ping test from the guest shell using the specified VRF, ensuring proper network routing and connectivity testing.

Further research:

Cisco NX-OS Guest Shell Guide:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/93x/programmability/guest_shell/guid OS_Guest_Shell_Guide_93x/m_accessing_guestshell.html (Look for information regarding VRF context within the Guest Shell)
Cisco NX-OS VRF Configuration:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/system_management/configuration/guide/b_Cisco_Nexus_9000_Series_System_Management_Configuration_G

Question: 33

A DNS server with IP address 192.168.1.1 is deployed in a data center. A network engineer must configure a Cisco UCS Fabric Interconnect to use this DNS.
Which configuration should be applied?

A. ficl-mgmt-A# scope fabric-interconnect a ficl-mgmt-A /fabric-interconnect # set name 192.168.1.1 ficl-mgmt-A /fabric-interconnect # scope system ficl-mgmt-A /system # commit-buffer

B. ficl-mgmt-A# scope system ficl-mgmt-A /system # create dns 192.168.1.1 ficl-mgmt-A /system* # commit-buffer

C. ficl-mgmt-A# scope fabric-interconnect a ficl-mgmt-A /fabric-interconnect # set name 192.168.1.1 ficl-mgmt-A /fabric-interconnect* # commit-buffer

D. ficl-mgmt-A# scope system ficl-mgmt-A /system # scope services ficl-mgmt-A /system/services # create dns 192.168.1.1 ficl-mgmt-A /system/services* # commit-buffer

Answer: D

Explanation:

The correct answer is **D** because it demonstrates the correct CLI commands to configure a DNS server within the Cisco UCS Fabric Interconnect's system settings. Let's break down why.

Options A and C attempt to modify the Fabric Interconnect's name with the IP address of the DNS server, which is incorrect. The 'set name' command is used for setting the name of the Fabric Interconnect itself, not a DNS server. These commands are operating at the wrong level within the UCS CLI hierarchy. Option B incorrectly tries to create a DNS entry directly under the system scope.

Option D accurately targets the correct location for DNS configuration within the UCS CLI. The command sequence scope system, moves the context to the system configuration settings. Next, scope services navigates to the service settings area where you configure DNS, NTP, SNMP, etc. Finally, the create dns 192.168.1.1 command creates a DNS server entry using the specified IP. commit-buffer commits the changes to the running configuration. This action configures the Fabric Interconnect to utilize the provided DNS server for name resolution. The provided command hierarchy is consistent with the Cisco UCS CLI structure and is necessary for effective DNS configuration.

For further research into this configuration, refer to the official Cisco UCS documentation. Specifically, look at guides on Fabric Interconnect administration and networking, focusing on system configuration and services.

The Cisco UCS Manager Configuration Guide is a great starting point.https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-manager/GUI-User-Guides/Network-Mgmt/4-2/b_UCSM_Network_Management_Guide_4_2/b_UCSM_Network_Management_Guide_4_2_chapter_01001.html (This link provides general guidance on network management using Cisco UCS Manager.)https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-manager/CLI-User-Guides/Networking/4-1/b_UCSM_CLI_Network_Management_Guide_4_1/b_UCSM_CLI_Network_Management_Guide_4_1_chapter_01001 (This link is a specific guide for CLI usage.)

Question: 34

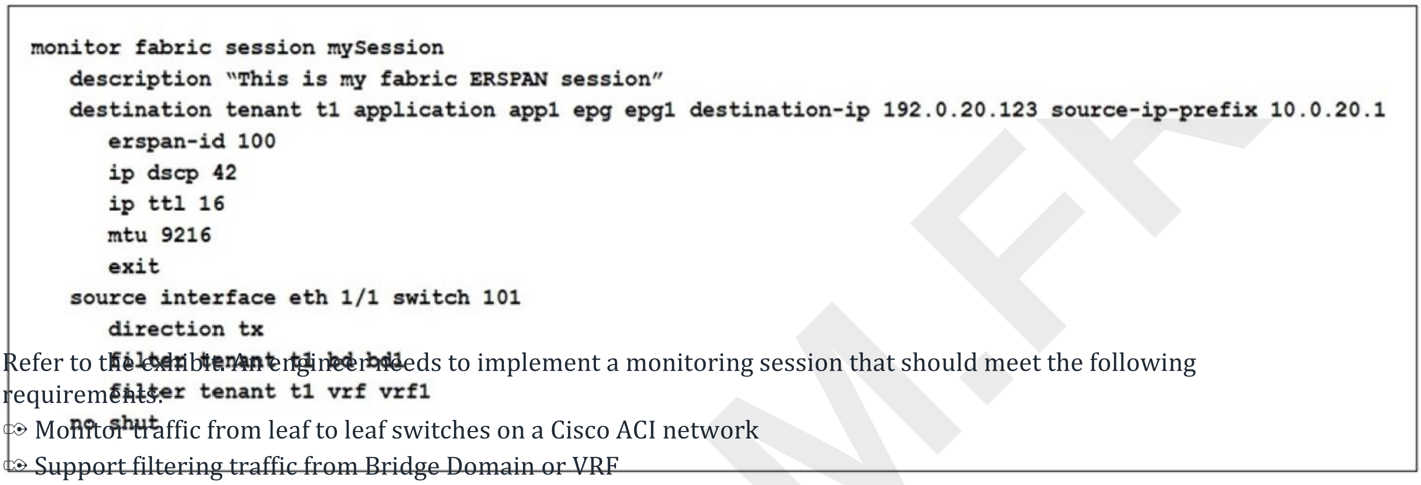

```
monitor fabric session mySession
   description "This is my fabric ERSPAN session"
   destination tenant t1 application app1 epg epg1 destination-ip 192.0.20.123 source-ip-prefix 10.0.20.1
      erspan-id 100
      ip dscp 42
      ip ttl 16
      mtu 9216
      exit
   source interface eth 1/1 switch 101
      direction tx
      filter tenant t1 bd bd1
      filter tenant t1 vrf vrf1
      no shut
```

Refer to the exhibit. An engineer needs to implement a monitoring session that should meet the following requirements:

✑ Monitor traffic from leaf to leaf switches on a Cisco ACI network

✑ Support filtering traffic from Bridge Domain or VRF

Which configuration must be added to meet these requirements?

A. interface eth 1/2 leaf 101

B. application epg epg1 app1

C. interface eth 1/2 switch 101

D. application app1 epg epg1

Answer: D

Explanation:

Reference:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/troubleshooting/b_APIC_Troubleshooting/ b_APIC_Troubleshooting_chapter_0110.html

Question: 35

An engineer is implementing OTV on a transport that supports multicast. The solution needs to meet the following requirements:

✑ Establish adjacency to the remote peer by using multicast.

Enable OTV advertisements for VLAN 100 to the other site.

▪

Which two commands should be configured to meet these requirements? (Choose two.)

A. otv extend-vlan 100
B. otv site-vlan 100
C. otv use-adjacency-server 172.27.255.94
D. otv data-group 232.2.2.0/28
E. otv control-group 232.1.1.1

Answer: AE

Explanation:

A & E correct!

VLAN 100 to be extendedThe correct ans is A and E

Question: 36

An engineer needs to create a new user in the local user database on Cisco UCS Fabric Interconnect. The user needs permissions to change the following configuration inside UCS Manager version 3.1:

✑ vNIC and vHBA profiles

✑ Fan speed and power redundancy profile of UCS Manager

Which two roles must be assigned to a user to achieve this goal? (Choose two.)

A. facility-manager
B. server-equipment
C. operations
D. server-profile
E. server-compute

Answer: AD

Explanation:

The correct answer is **A. facility-manager and D. server-profile**. Let's break down why:

Cisco UCS Manager uses a role-based access control (RBAC) system to grant specific permissions to users. This is a common practice in cloud environments to maintain security and ensure users have only the necessary access. facility-manager and server-profile roles provide the needed permissions for the described tasks.

The facility-manager role provides access to hardware-related configurations at the system level. These include functions such as managing fans, power redundancy, and overall system settings. Changing fan speed and power redundancy directly relates to these hardware management features controlled by the facility-manager role. This is a critical aspect of data center operations, as ensuring proper cooling and power management is essential for the reliability of the infrastructure.

The server-profile role, on the other hand, allows users to work with server-related configuration, specifically the creation and modification of service profiles and their components. This includes Virtual Network Interface Cards (vNICs) and Virtual Host Bus Adapters (vHBAs) profiles. These are essential for connecting servers to the network and storage. Manipulating these configurations is under the scope of the server-profile role, which is essential for configuring the networking and storage aspects of the server's logical identity.

The other roles are not sufficient. The server-equipment role typically provides access to view the system's physical hardware but does not include modification permissions, so it's not suitable for changing fan settings. The operations role usually covers day-to-day tasks like checking server status and managing user sessions but not for configuring hardware profiles. The server-compute role allows for the management of the compute aspect of servers but does not grant access to configure networking (vNIC) or storage (vHBA) profiles or manage the facility-related functions.

Therefore, to configure both the vNIC/vHBA profiles and the fan/power redundancy settings, the user must possess both the facility-manager and server-profile roles. These roles provide the necessary permissions in a least privilege manner, in line with best practices for security and access control.

Authoritative Links:

Cisco UCS Manager Role Based Access Control:
https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-manager/GUI-User-Guides/Admin-Mgmt/3-1/b_UCSM_Admin_Mgmt_Guide_3_1/b_UCSM_Admin_Mgmt_Guide_3_1_chapter_0100.html
Cisco UCS Roles and Privileges:https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-manager/GUI-User-Guides/Admin-Mgmt/4-1/b_UCSM_Admin_Mgmt_Guide_4_1/b_UCSM_Admin_Mgmt_Guide_4_1_chapter_0100.html#task_E517405B03904

Question: 37

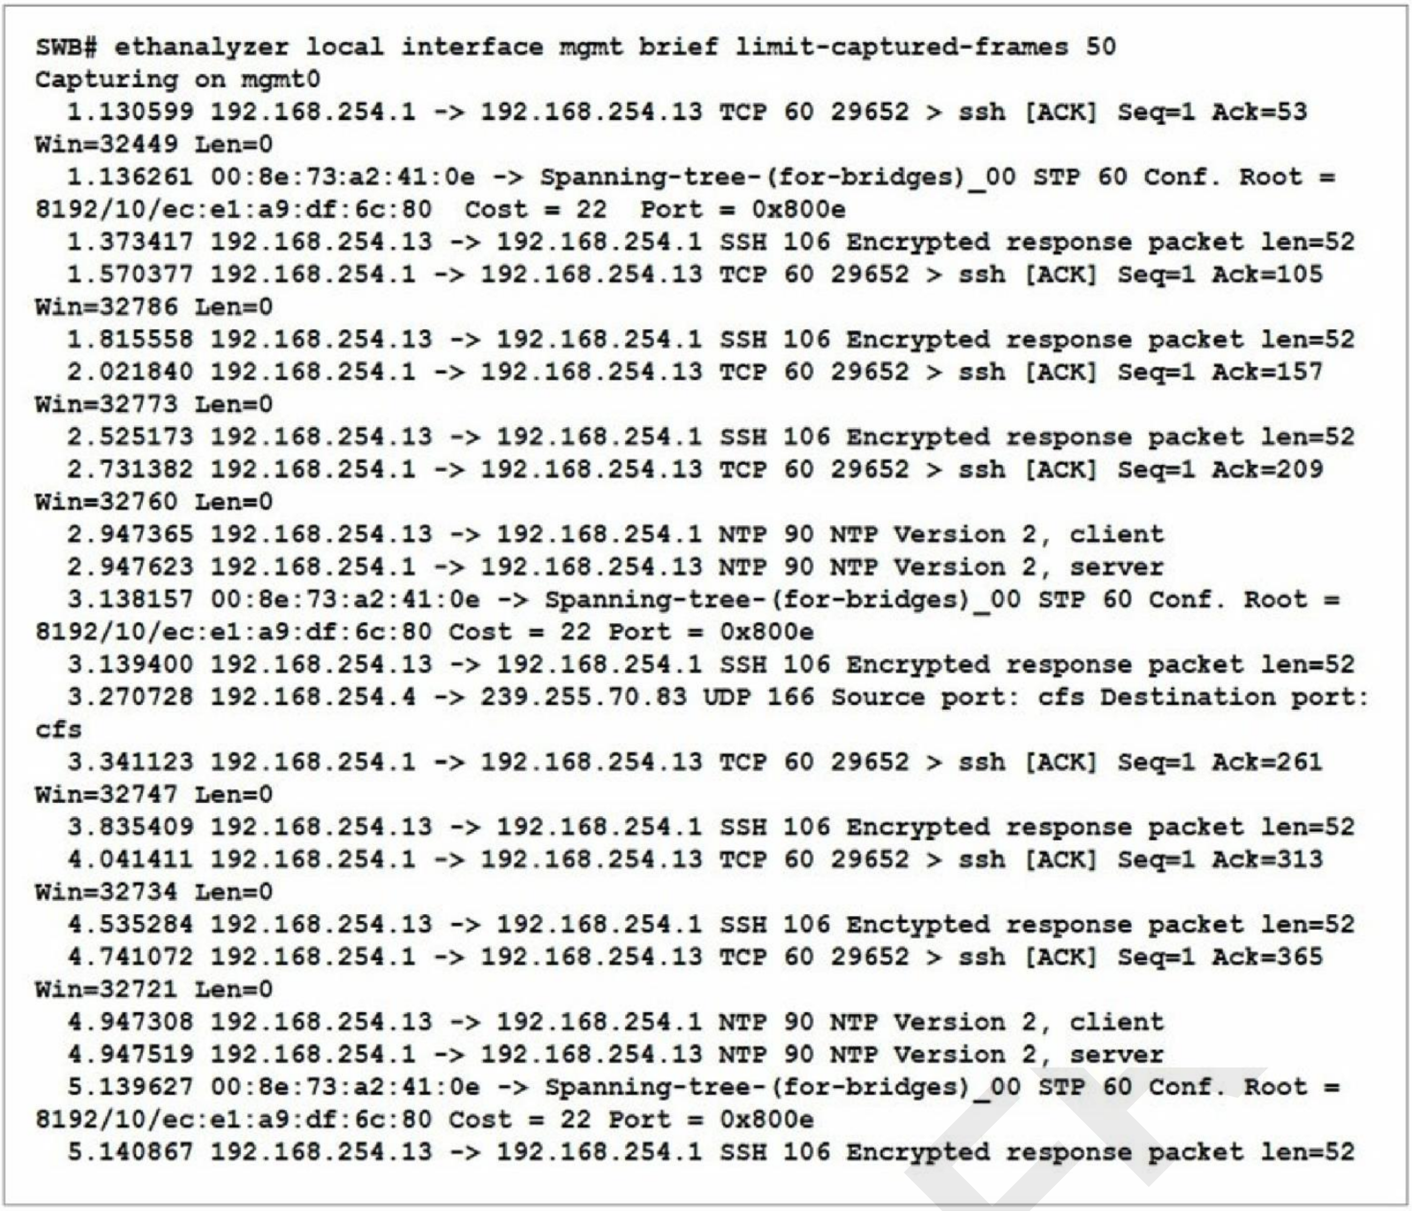

```
SWB# ethanalyzer local interface mgmt brief limit-captured-frames 50
Capturing on mgmt0
  1.130599 192.168.254.1 -> 192.168.254.13 TCP 60 29652 > ssh [ACK] Seq=1 Ack=53
Win=32449 Len=0
  1.136261 00:8e:73:a2:41:0e -> Spanning-tree-(for-bridges)_00 STP 60 Conf. Root =
8192/10/ec:e1:a9:df:6c:80  Cost = 22  Port = 0x800e
  1.373417 192.168.254.13 -> 192.168.254.1 SSH 106 Encrypted response packet len=52
  1.570377 192.168.254.1 -> 192.168.254.13 TCP 60 29652 > ssh [ACK] Seq=1 Ack=105
Win=32786 Len=0
  1.815558 192.168.254.13 -> 192.168.254.1 SSH 106 Encrypted response packet len=52
  2.021840 192.168.254.1 -> 192.168.254.13 TCP 60 29652 > ssh [ACK] Seq=1 Ack=157
Win=32773 Len=0
  2.525173 192.168.254.13 -> 192.168.254.1 SSH 106 Encrypted response packet len=52
  2.731382 192.168.254.1 -> 192.168.254.13 TCP 60 29652 > ssh [ACK] Seq=1 Ack=209
Win=32760 Len=0
  2.947365 192.168.254.13 -> 192.168.254.1 NTP 90 NTP Version 2, client
  2.947623 192.168.254.1 -> 192.168.254.13 NTP 90 NTP Version 2, server
  3.138157 00:8e:73:a2:41:0e -> Spanning-tree-(for-bridges)_00 STP 60 Conf. Root =
8192/10/ec:e1:a9:df:6c:80 Cost = 22 Port = 0x800e
  3.139400 192.168.254.13 -> 192.168.254.1 SSH 106 Encrypted response packet len=52
  3.270728 192.168.254.4 -> 239.255.70.83 UDP 166 Source port: cfs Destination port:
cfs
  3.341123 192.168.254.1 -> 192.168.254.13 TCP 60 29652 > ssh [ACK] Seq=1 Ack=261
Win=32747 Len=0
  3.835409 192.168.254.13 -> 192.168.254.1 SSH 106 Encrypted response packet len=52
  4.041411 192.168.254.1 -> 192.168.254.13 TCP 60 29652 > ssh [ACK] Seq=1 Ack=313
Win=32734 Len=0
  4.535284 192.168.254.13 -> 192.168.254.1 SSH 106 Enctypted response packet len=52
  4.741072 192.168.254.1 -> 192.168.254.13 TCP 60 29652 > ssh [ACK] Seq=1 Ack=365
Win=32721 Len=0
  4.947308 192.168.254.13 -> 192.168.254.1 NTP 90 NTP Version 2, client
  4.947519 192.168.254.1 -> 192.168.254.13 NTP 90 NTP Version 2, server
  5.139627 00:8e:73:a2:41:0e -> Spanning-tree-(for-bridges)_00 STP 60 Conf. Root =
8192/10/ec:e1:a9:df:6c:80 Cost = 22 Port = 0x800e
  5.140867 192.168.254.13 -> 192.168.254.1 SSH 106 Encrypted response packet len=52
```

Refer to the exhibit. Cisco Fabric Services is enabled in the network. Which type of IP address is used by the Cisco Fabric Services protocol?

A. IPv4 multicast address
B. IPv4 unicast address
C. IPv4 gateway address
D. IPv4 anycast address

Answer: A

Explanation:

CFS uses the defaults IP multicast address of 239.255.70.83. Although this is configurable. So answer is A.

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/configuration/guide/fm/FabricManage

Question: 38

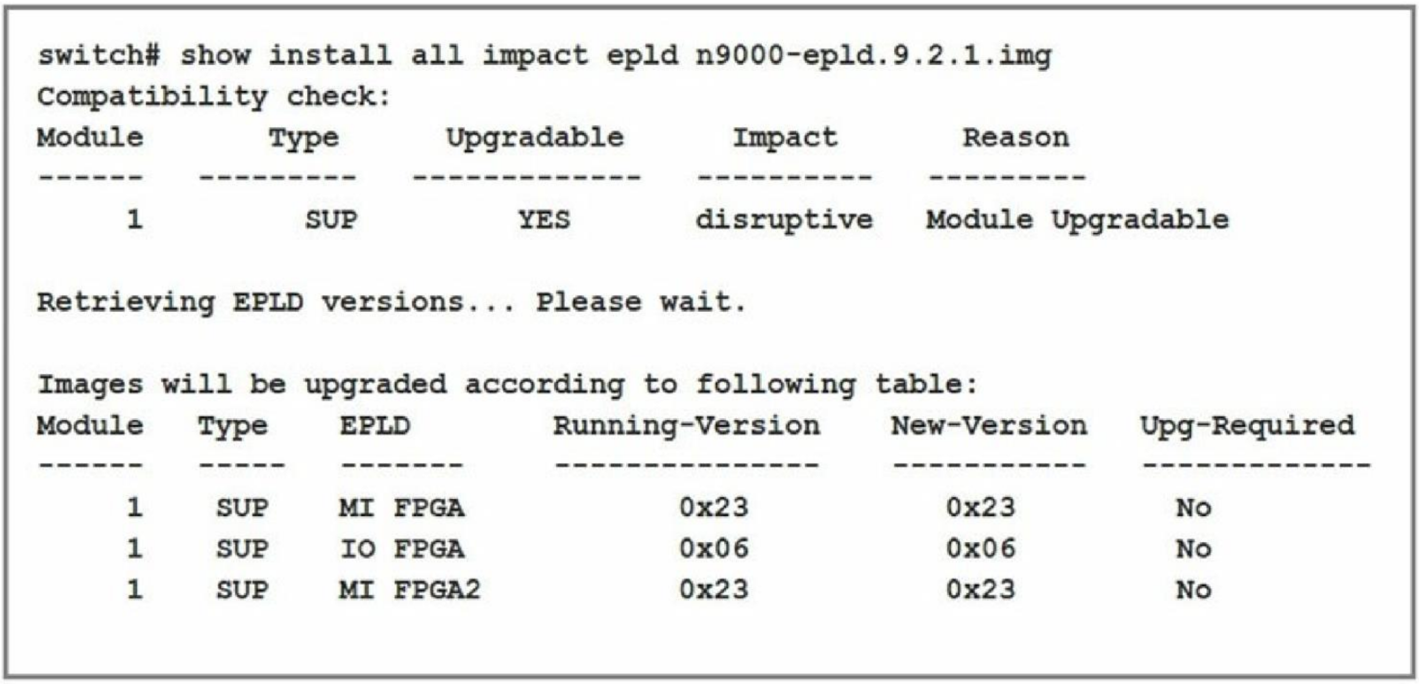

```
switch# show install all impact epld n9000-epld.9.2.1.img
Compatibility check:
Module       Type      Upgradable      Impact       Reason
------   ---------   -------------   ----------   ---------
     1         SUP          YES       disruptive   Module Upgradable

Retrieving EPLD versions... Please wait.

Images will be upgraded according to following table:
Module   Type    EPLD        Running-Version    New-Version   Upg-Required
------   -----   -------     ---------------    -----------   -------------
     1    SUP    MI FPGA          0x23              0x23          No
     1    SUP    IO FPGA          0x06              0x06          No
     1    SUP    MI FPGA2         0x23              0x23          No
```

Refer to the exhibit. A network engineer plans to upgrade the current software images of a Cisco Nexus switch. What is the impact of starting the EPLD upgrade?

A. The switch reboots one time for the new EPLD versions to take effect.

B. The switch reboots multiple times for the new EPLD versions to take effect.

C. The switch skips the EPLD upgrade for each EPLD device of the switch.

D. The switch skips the EPLD upgrade for MIFPGA and proceeds with the other EPLD devices.

Answer: C

Explanation:
Determining Whether to Upgrade EPLD Images
EPLD image number, you can skip the upgrade.

To determine the EPLD upgrades needed for a Cisco Nexus 3000 Series switch, use the show install impact epld bootflash: command on that switch and indicate the n9000-epld.9.2.4 image. In the following example, the MIFPGA, and IOFPGA EPLD images do not need to be upgraded.

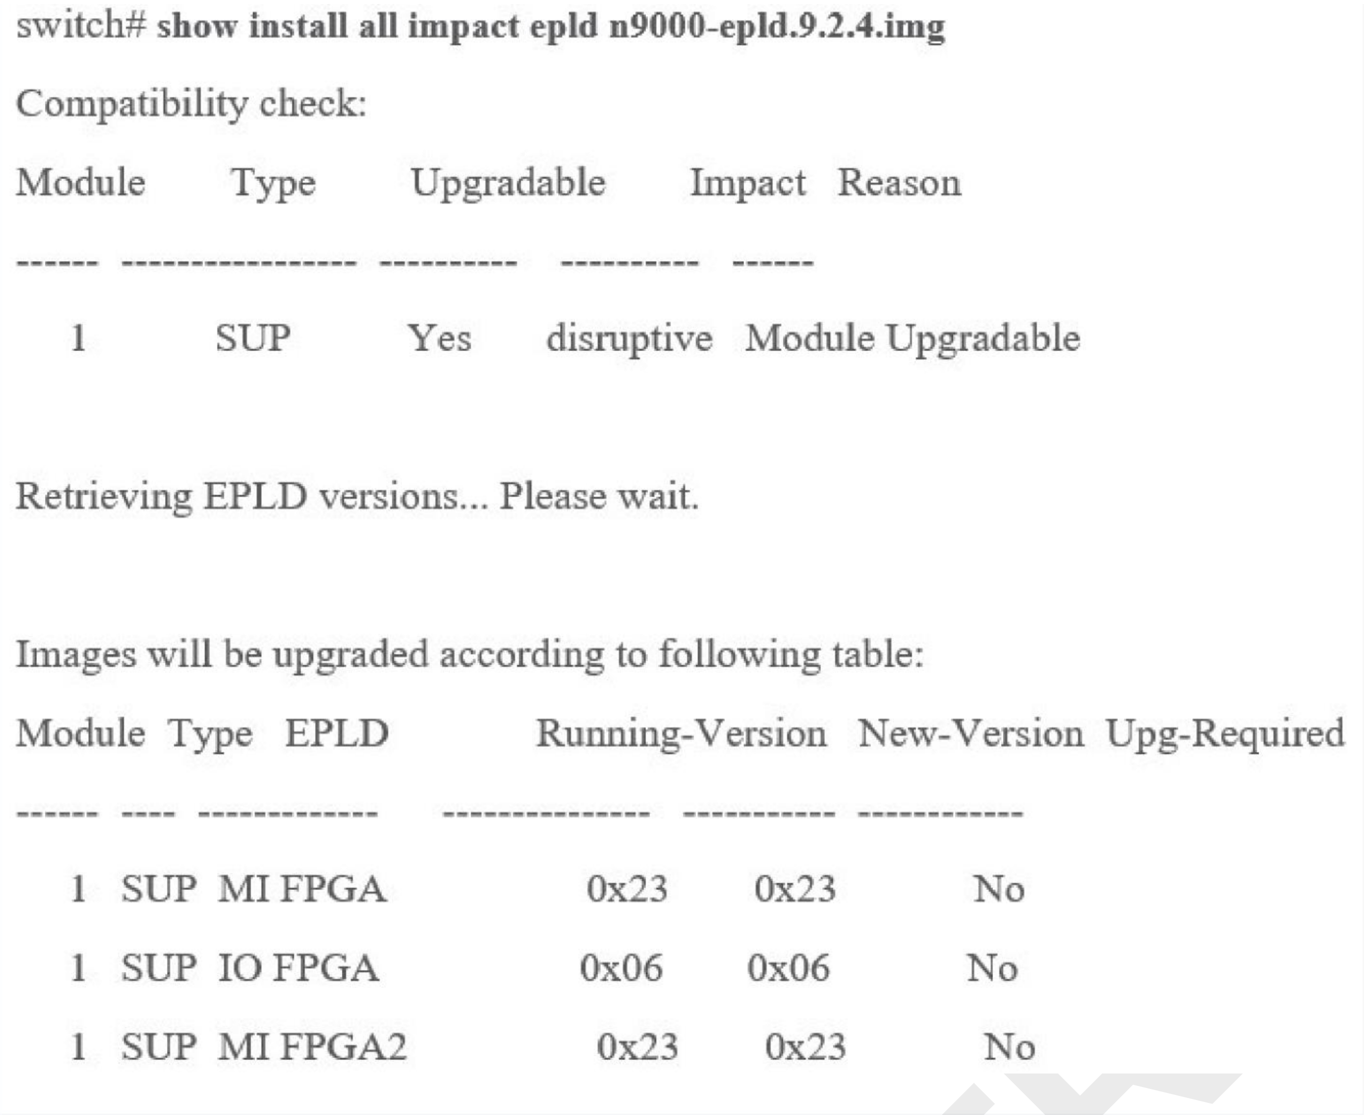

```
switch# show install all impact epld n9000-epld.9.2.4.img

Compatibility check:

Module      Type       Upgradable       Impact   Reason

------  ----------------  ----------   ----------   ------

    1         SUP         Yes      disruptive   Module Upgradable


Retrieving EPLD versions... Please wait.


Images will be upgraded according to following table:

Module  Type   EPLD            Running-Version  New-Version  Upg-Required

------  ----  ------------     ---------------   -----------  ------------

    1   SUP  MI FPGA              0x23        0x23           No

    1   SUP  IO FPGA              0x06        0x06           No

    1   SUP  MI FPGA2              0x23        0x23           No
```

Reference:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/92x/epld-rn/nxos_n9K_epldRN_924.html

Question: 39

Which behavior defines streaming telemetry as a push model in Cisco devices?

A. Events and network changes generate telemetry data
B. Monitoring clients are pulling data from the network to see real-time statistics
C. JSON encoded telemetry data is transported using the gRPC protocol
D. The network devices send data in JSON or GPB format to configure endpoints

Answer: D

Explanation:

The correct answer is **D. The network devices send data in JSON or GPB format to configured endpoints.** This behavior accurately defines a "push" model within the context of streaming telemetry.

Here's why:

Streaming telemetry, in general, can be implemented using either a "push" or "pull" model. In a push model, devices actively initiate sending data to pre-configured collectors or destinations. The devices "push" data outwards. In the given scenario, these devices are Cisco network elements that stream telemetry data. These data streams are encoded in formats like JSON (JavaScript Object Notation) or GPB (Google Protocol Buffers), which are both common for structured data representation. The fact that the data is sent from the device to a specific endpoint is the key characteristic of the push mechanism.

Option A, while related, doesn't define the push behavior itself. Events triggering data generation are a pre-requisite for any form of telemetry, but the mechanism of transmission (push or pull) is what we are evaluating here. Option B describes a "pull" model, where monitoring clients request information from network devices, which is the opposite of what push models achieve. Option C highlights the technical aspects of data encoding (JSON) and transport (gRPC). While important, they are implementation details of the communication, not the core concept of pushing data.

In contrast, Option D clearly states that the devices are sending data to specified destinations, making it a clear description of a push based model. This actively sends, or “pushes” data to configured locations, without waiting for a request from a client. This aligns with the idea of streaming telemetry as it is a continuous and unsolicited flow of data initiated by the source.

Authoritative Links:

Cisco's Streaming Telemetry Documentation:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/prog/configuration/1612/b_16_12_programmability_config_guide/b_16_12_programmability_config_guide_c (This link describes push model within Cisco environments)
General Overview of Telemetry:https://en.wikipedia.org/wiki/Telemetry (Provides general information on telemetry concepts)
gRPC Documentation:https://grpc.io/docs/ (Provides info on gRPC protocol)
JSON Documentation:https://www.json.org/ (Provides information on JSON)
Google Protocol Buffer Documentation:https://developers.google.com/protocol-buffers (Provides information on GPB)

Question: 40

After a Cisco Nexus 7000 Series Switch chassis replacement, the administrator discovers that all vPC-enabled LACP port channels are reinitialized. The administrator wants to prevent this issue the next time the chassis is replaced. Which two actions must be taken to meet this requirement before the isolated device is reloaded?
(Choose two.)

A. Change the vPC system-priority of the replacement chassis to a higher value than the peer
B. Set the vPC MAC address to a higher value than the peer
C. Configure auto-recovery to the disable state on both peers
D. Set the vPC MAC address to a lower value than the peer
E. Change the vPC system-priority of the replacement chassis to a lower value than the peer

Answer: BC

Explanation:

The core issue here revolves around the behavior of vPC (virtual Port Channel) when a primary peer device is replaced in a Cisco Nexus environment. During a chassis replacement, the newly installed Nexus switch comes up with a default vPC configuration, including system-priority and MAC address. This default state can lead to a disruptive re-election process, causing LACP (Link Aggregation Control Protocol) port channels to re-initialize as the new switch attempts to assert itself as the primary peer. To avoid this disruption, it is essential to proactively manage vPC parameters before bringing the replacement chassis online.

Option B is correct because setting a lower vPC MAC address on the new device relative to the existing peer will make sure that when the new switch comes online it will not attempt to assume the role of a primary. vPC roles depend on the lowest MAC address. This avoids the primary role transfer and prevents LACP port channel re-initialization, thus minimizing disruptions.

Option C is also correct because enabling the auto-recovery feature can also lead to a primary role takeover upon the reloaded switch coming online. Auto-recovery allows a secondary vPC device to automatically become primary if its peer fails for too long and then reloads, which causes the secondary device to become primary and force a re-initiation of LACP port channels. By disabling auto-recovery before reloading, you prevent this automatic takeover scenario, keeping the original primary vPC device and preserving the existing LACP configuration.

Options A and E are incorrect because the vPC system priority is designed to determine the vPC primary role during a live device failure, not for replacement scenarios. Changing the vPC system-priority of the replacement chassis to a higher (option A) or lower (option E) value would trigger a re-election process. Option D is incorrect because setting the vPC MAC address to a lower value than the peer would make it a primary, and thus a disruptive re-election could occur.

In summary, configuring a lower vPC MAC address and disabling auto-recovery on the new replacement device prevents unwanted re-election, ensuring a smooth transition and minimizing disruption to LACP port channels during a chassis replacement.

Further Reading:

1. Cisco Nexus 7000 Series NX-OS Virtual Port Channel Configuration Guide: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus7000/sw/configuration/n7k_clis/virtual OS_Virtual_Port_Channel_Configuration_Guide_7x/b_Cisco_Nexus_7000_Series_NX-OS_Virtual_Port_Channel_Configuration_Guide_7x_chapter_011.html
2. Understanding vPC and its operation: https://www.cisco.com/c/en/us/support/docs/switches/nexus-7000-series-switches/119350-technote-vpc-00.html

Question: 41

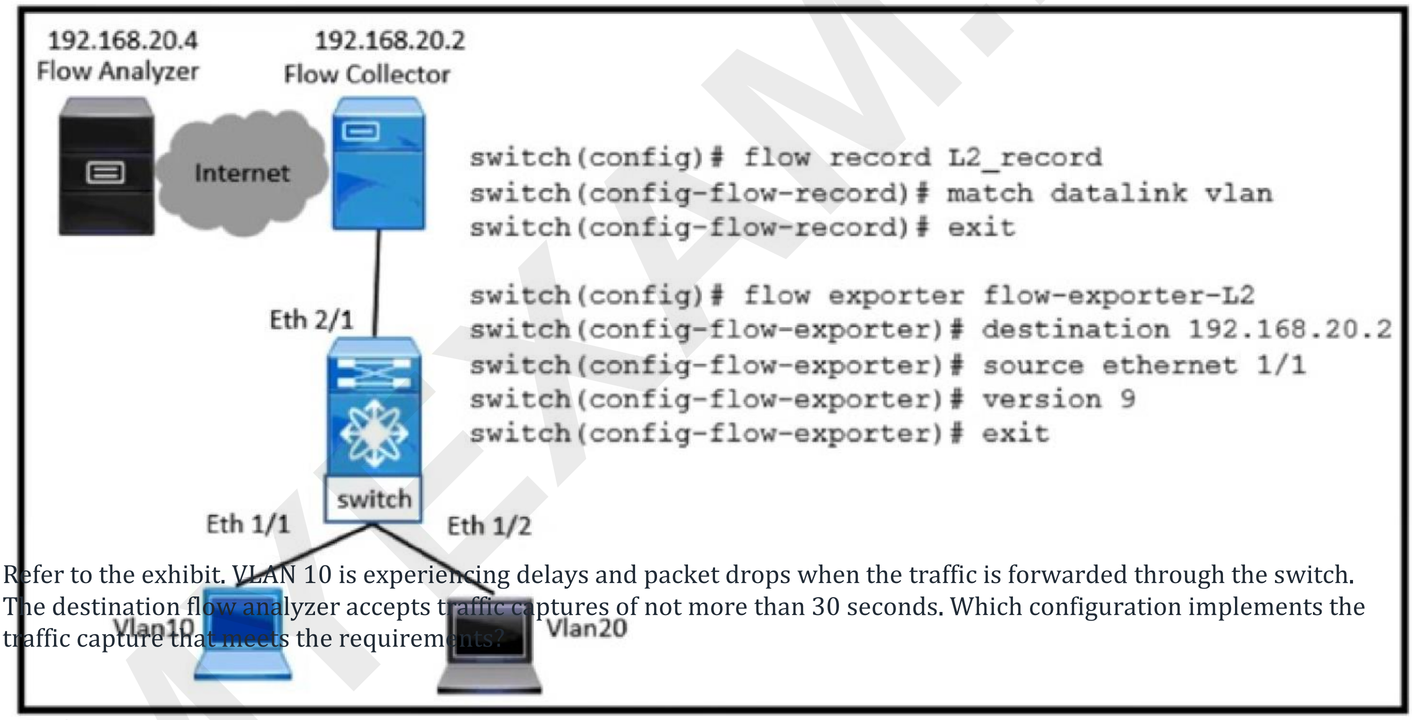

```
switch(config)# flow record L2_record
switch(config-flow-record)# match datalink vlan
switch(config-flow-record)# exit

switch(config)# flow exporter flow-exporter-L2
switch(config-flow-exporter)# destination 192.168.20.2
switch(config-flow-exporter)# source ethernet 1/1
switch(config-flow-exporter)# version 9
switch(config-flow-exporter)# exit
```

Refer to the exhibit. VLAN 10 is experiencing delays and packet drops when the traffic is forwarded through the switch. The destination flow analyzer accepts traffic captures of not more than 30 seconds. Which configuration implements the traffic capture that meets the requirements?

A.

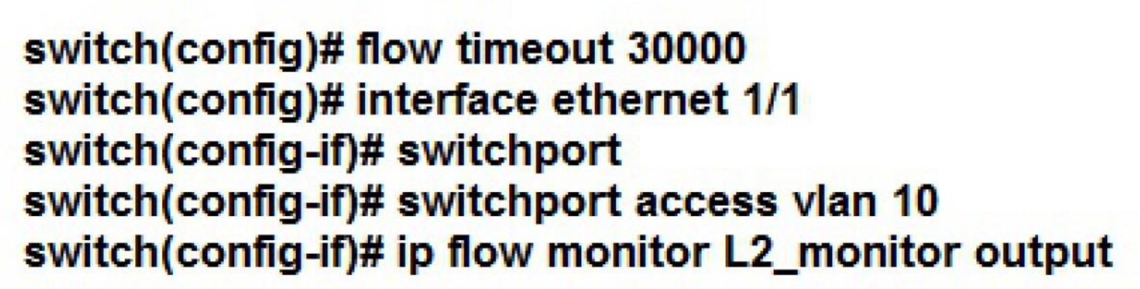

```
switch(config)# flow timeout 30000
switch(config)# interface ethernet 1/1
switch(config-if)# switchport
switch(config-if)# switchport access vlan 10
switch(config-if)# ip flow monitor L2_monitor output
```

B.

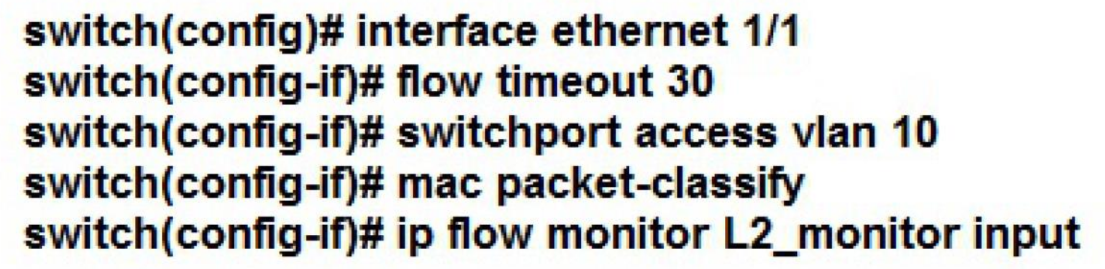

```
switch(config)# interface ethernet 1/1
switch(config-if)# flow timeout 30
switch(config-if)# switchport access vlan 10
switch(config-if)# mac packet-classify
switch(config-if)# ip flow monitor L2_monitor input
```

C.

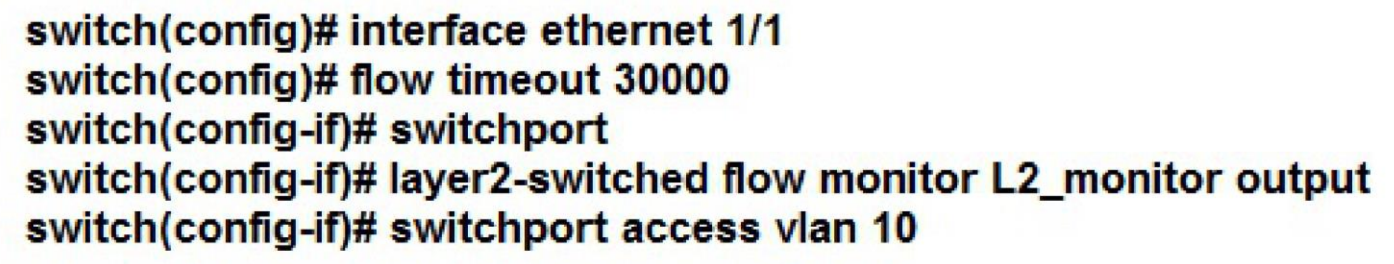

```
switch(config)# interface ethernet 1/1
switch(config)# flow timeout 30000
switch(config-if)# switchport
switch(config-if)# layer2-switched flow monitor L2_monitor output
switch(config-if)# switchport access vlan 10
```

D.

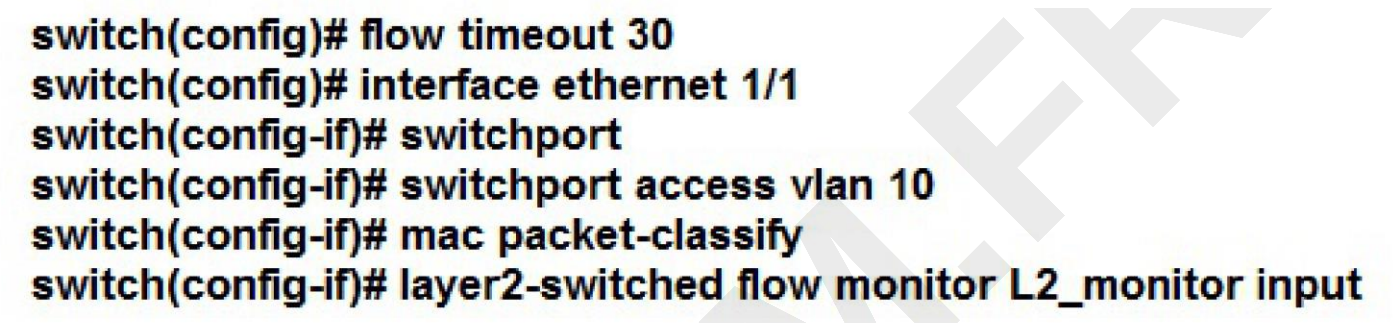

```
switch(config)# flow timeout 30
switch(config)# interface ethernet 1/1
switch(config-if)# switchport
switch(config-if)# switchport access vlan 10
switch(config-if)# mac packet-classify
switch(config-if)# layer2-switched flow monitor L2_monitor input
```

Answer: D

Explanation:

Reference:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/system_management/configuration/guide/ b_Cisco_Nexus_9000_Series_NX-OS_System_Management_Configuration_Guide_7x/b_Cisco_Nexus_9000_Series_NX-OS_System_Management_Configuration_Guide_7x_chapter_011100.html

Question: 42

An engineer performs a set of configuration changes for the vPC domain using Session Manager. Which two commands are used to verify the configuration and apply the device changes when no errors are returned? (Choose two.)

A. commit

B. verify
C. apply
D. checkpoint
E. write

Answer: AB

Explanation:

Let's break down why options A and B are the correct commands within the context of Cisco's Session Manager for vPC (virtual PortChannel) configuration changes. Session Manager acts as a transactional system, staging configurations before they're actually implemented on the device.

First, an engineer makes changes within the Session Manager environment. These changes are not yet live on the networking devices themselves. After these modifications, it's essential to **verify** the intended configuration. This process, initiated using the **verify** command (option B), checks for syntax errors, inconsistencies, or other issues in the proposed configurations, preventing problems before deployment. It does not implement the changes, it only checks if they can be. If no errors are found during the verification step, the configurations are considered safe for application.

Next, the **commit** command (option A) is the key to implementing the verified changes on the physical vPC devices within the domain. The commit command is the action that takes the verified staged configuration and pushes it live onto the device. It's the final step that makes the changes operational.

Commands like **apply** (option C) and **checkpoint** (option D) have other uses within configuration management. 'Apply' typically refers to a broader action in other contexts, but it is not directly linked to applying changes staged with Cisco Session Manager. 'Checkpoint' creates backups, but doesn't apply configurations to live devices. The **write** command (option E), is also associated with saving the current configuration and not for a session manager commit process.

In summary, the sequential actions of verification with **verify** followed by implementation with **commit** are critical steps in the Cisco Session Manager workflow. This systematic approach ensures that changes are error-free and are applied in a controlled and predictable manner, reducing the risk of network disruptions.

Authoritative links for further research:

Cisco Nexus 9000 Series NX-OS Programmability Guide, Release 10.1(x):
https://www.cisco.com/c/en/us/td/docs/dcn/nxos/nexus9000/programmability/10x/b-cisco-nexus-9000-series-nx-os-programmability-guide-10x/m-session-manager-10x.html
Cisco Nexus 9000 Series NX-OS Configuration Command Reference, Release 10.1(x):
https://www.cisco.com/c/en/us/td/docs/dcn/nxos/nexus9000/command/10x/b-cisco-nexus-9000-series-nx-os-command-reference-10x/m-session-mgr-cmds-10x.html

Question: 43

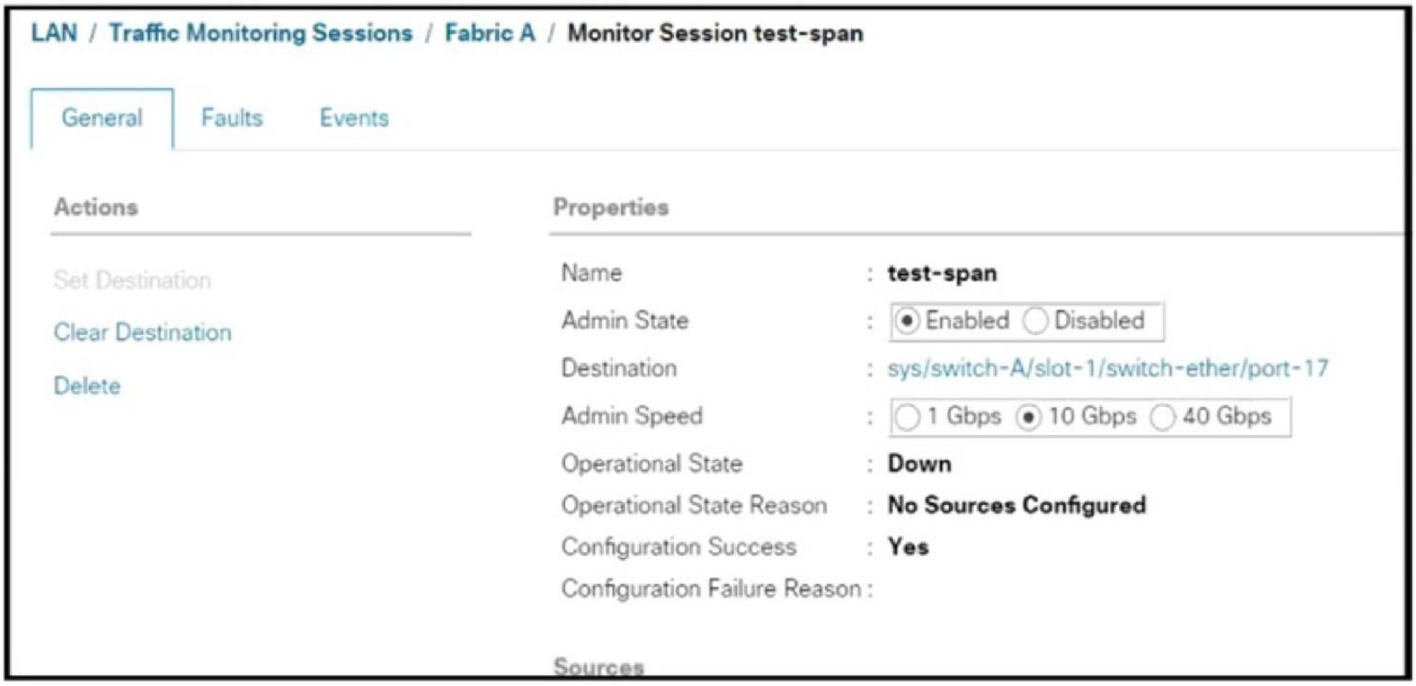

LAN / Traffic Monitoring Sessions / Fabric A / Monitor Session test-span

Name : **test-span**
Admin State : ● Enabled ○ Disabled
Destination : sys/switch-A/slot-1/switch-ether/port-17
Admin Speed : ○ 1 Gbps ● 10 Gbps ○ 40 Gbps
Operational State : **Down**
Operational State Reason : **No Sources Configured**
Configuration Success : **Yes**
Configuration Failure Reason :

Refer to the exhibit. An engineer must monitor all LAN traffic on Fabric A from a blade server. Which source should be configured in the test-span monitor session to complete this task?

A. all uplink FCoE ports
B. all vNICs from the service profile that correspond to this server C. all uplink Ethernet ports
D. all vHBAs from the service profile that correspond to this server

Answer: B

Explanation:

Reference:
https://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/141/UCSM_GUI_Configuratio n

Question: 44

An engineer configured an environment that contains the vPC and non-vPC switches. However, it was noticed that the downstream non-vPC switches do not receive the same STP bridge ID from the upstream vPC switch peers. Which vPC feature must be implemented to ensure that vPC and non-vPC switches receive the same STP bridge ID from the upstream vPC switch peers?

A. vpc local role-priority 4000
B. peer-switch
C. system-mac 0123.4567.89ab
D. peer-gateway

Answer: B

Explanation:

The correct answer is **B. peer-switch**.

The peer-switch command is crucial for ensuring consistent STP behavior in a vPC (virtual PortChannel) environment when interacting with non-vPC switches. Without this command, each vPC peer switch will independently participate in STP, using its own unique bridge ID. This can lead to suboptimal STP convergence and potential loop issues. When peer-switch is enabled on both vPC peer devices, they share a single, virtual bridge ID for STP purposes. This effectively makes the vPC pair appear as a single logical STP entity to downstream non-vPC switches. Consequently, all downstream devices will receive and operate on the same STP bridge ID from the vPC domain, simplifying STP topology and eliminating the risk of inconsistencies. This feature ensures deterministic path selection, prevents loops, and facilitates faster convergence. Options A, C, and D are related to vPC but don't address the STP bridge ID consistency issue between vPC and non-vPC switches. The vpc local role-priority command (A) is related to vPC role selection, system-mac (C) configures a specific system MAC address, and peer-gateway (D) deals with gateway functionality within the vPC domain. They do not influence the STP behavior across the vPC border.

For further research, consult the following Cisco documentation:

Cisco Nexus 9000 Series NX-OS vPC Configuration Guide: Search for "peer-switch" in this guide for detailed explanations. (Specific links will vary depending on the NX-OS version)
Cisco Live presentations on vPC: These often have practical insights on vPC configuration and troubleshooting.

These resources will provide in-depth knowledge on vPC implementation, including the function and importance of peer-switch in relation to STP.

Question: 45

A network architect wants to propose a scalable network monitoring solution in which data is repeatedly acquired from network devices. The solution must use a push model and provide close to real-time access to operational data. Which technology must be used to meet these requirements?

A. CLI-based scripting
B. logging
C. SNMPv3
D. streaming telemetry

Answer: D

Explanation:

The correct answer is D, streaming telemetry. Here's why:

Streaming telemetry is designed for high-frequency, continuous data collection from network devices, fitting the requirement for repeated data acquisition. It uses a push model, where devices actively send data to a collector, rather than waiting for requests. This push approach provides near real-time access to operational data, aligning with the architect's need for close to real-time visibility.

CLI-based scripting (A) relies on polling, which is not a push model and is inefficient for real-time data retrieval. Logging (B) captures events but doesn't continuously stream operational data. SNMPv3 (C), while secure, is a pull-based protocol, where the management station requests data from devices. This does not satisfy the push model and can introduce latency for real-time data acquisition.

Streaming telemetry employs technologies like gRPC and Protocol Buffers for efficient data transmission and encoding, contributing to scalability and speed. In a modern data center, continuous monitoring of key metrics is crucial for identifying issues quickly and maintaining service quality. The push nature of streaming telemetry enables proactive monitoring, allowing the architect to identify issues before they significantly impact network performance.

For further research, refer to:

Cisco's documentation on streaming telemetry: https://www.cisco.com/c/en/us/solutions/data-center/streaming-telemetry.html
IETF's RFC for gRPC: https://grpc.io/
Protocol Buffers documentation: https://developers.google.com/protocol-buffers

Question: 46

A Cisco MDS 9000 Series Switch is configured for SAN Analytics and SAN Telemetry Streaming. An engineer must enable analytics for NVMe on interfaces in the range of fc1/1-12. Due to a large amount of traffic generated in the SAN environment, the data must be collected at regular intervals of 60 seconds from ports fc1/1-12 and then for ports fc1/13-24 for the next 60 seconds. Ports in the range fc1/13-24 were already enabled for analytics. Which set of commands must be used to meet these requirements?

A.

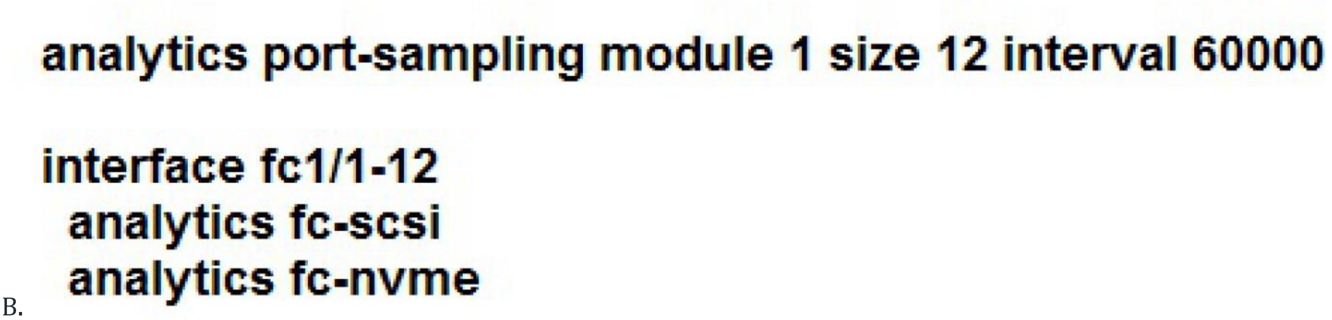

```
analytics port-sampling module 1 size 12 interval 60000

interface fc1/1-12
  analytics fc-scsi
  analytics fc-nvme
```

B.

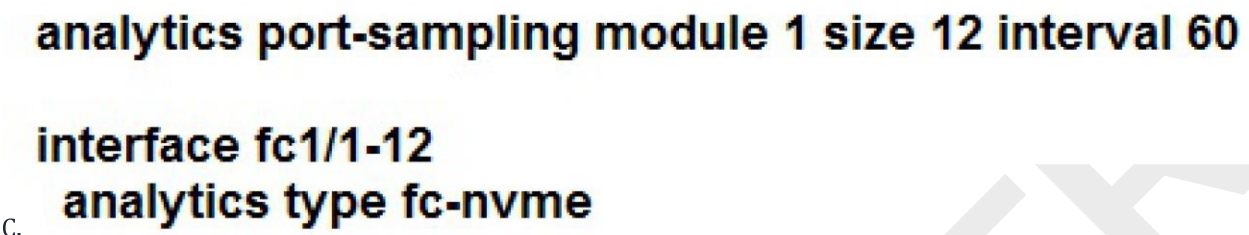

```
analytics port-sampling module 1 size 12 interval 60

interface fc1/1-12
  analytics type fc-nvme
```

C.

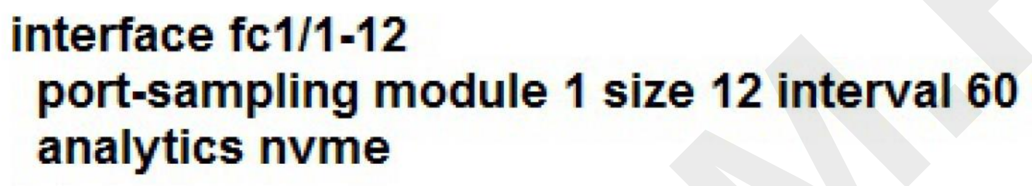

```
interface fc1/1-12
  port-sampling module 1 size 12 interval 60
  analytics nvme
```

D.

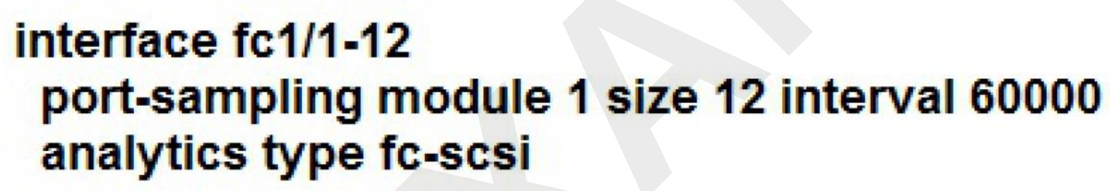

```
interface fc1/1-12
  port-sampling module 1 size 12 interval 60000
  analytics type fc-scsi
```

Answer: B

Explanation:

Reference:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/8_x/config/san_analytics/cisco-mds9000-san-analytics-telemetry-streaming- config-guide-8x/configuring_san_telemetry-streaming.html

Question: 47

A customer needs a tool to take advantage of the CI/CD model to streamline its operations and optimize cost. The customer wants to integrate the solution with the Cisco products it currently uses, including Cisco ACI networking and Cisco UCS Series servers. The solution should also provide on-premises Kubernetes and AppDynamics performance monitoring. Because of the security requirements, the solution should not install a local client on products under management. Which orchestration solution meets these requirements?

A. Cisco UCS Director

B. Cisco CloudCenter

C. Cisco APIC

D. Cisco DCNM

Answer: B

Explanation:

The correct answer is **B. Cisco CloudCenter**. Here's why: Cisco CloudCenter is specifically designed to be a multi-cloud management and orchestration platform that fits the customer's requirements. It allows users to model, deploy, and manage applications across various environments, including on-premises infrastructure like Cisco ACI and UCS, as well as public clouds. This aligns with the customer's desire for a CI/CD model to streamline operations and optimize costs. Furthermore, CloudCenter offers integration with Cisco ACI for network provisioning and UCS for server management, fulfilling the customer's need to leverage existing Cisco products. It also provides on-premises Kubernetes support, enabling containerized application deployment, and integrates with AppDynamics for performance monitoring. Crucially, CloudCenter operates agentlessly, meaning it doesn't require installing a local client on managed devices, addressing the customer's security requirements. This agentless approach centralizes management and minimizes the security footprint. Cisco UCS Director, while a powerful tool for infrastructure automation, is primarily focused on UCS and doesn't offer the broad multi-cloud capabilities, Kubernetes support, or agentless operation provided by CloudCenter. APIC is the central control point for ACI, not an orchestration platform. DCNM focuses on data center network management, not application orchestration. CloudCenter's focus on application orchestration, multi-cloud management, agentless operation, and integration with the mentioned Cisco products makes it the optimal solution.

Further Research:

Cisco CloudCenter:https://www.cisco.com/c/en/us/products/cloud-systems-management/cloudcenter/index.html
Cisco ACI:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/index.html
Cisco UCS:https://www.cisco.com/c/en/us/products/servers-unified-computing/index.html

Question: 48

A customer wants to offload some of its order processing to a public cloud environment. The customer environment is based on Cisco ACI and uses Puppet with containerized applications. The operations team requires a solution to orchestrate and optimize the cost of the new solution. Which product must be used to meet these requirements?

A. Cisco Workload Optimization Manager

B. Cisco CloudCenter

C. Cisco Intersight

D. Cisco Data Center Network Manager

Answer: B

Explanation:

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

Cisco CloudCenter (now part of Intersight) is specifically designed for multi-cloud application management and orchestration. The scenario describes a hybrid cloud environment – the customer has an on-premises Cisco ACI infrastructure and wants to extend order processing to a public cloud. This necessitates a tool that can manage deployments across both environments. CloudCenter excels at this by providing a unified platform for deploying, managing, and optimizing applications across various clouds, including public, private, and hybrid setups. It supports application modeling, deployment automation, and lifecycle management. Its orchestration capabilities enable the customer's operations team to control the application deployment process, thereby optimizing costs, resource usage, and performance across their chosen environments. The described customer uses Puppet for configuration management within their containers, this can be integrated with the features offered by CloudCenter for seamless automation. While other options have overlaps, only CloudCenter offers the specific orchestration and hybrid cloud application management required to fulfill the request. Cisco Workload Optimization Manager is focused on resource optimization within a single datacenter, not multi-cloud application orchestration. Cisco Intersight is a management platform for Cisco infrastructure but isn't focused on application lifecycle management and orchestration across cloud platforms. Similarly, Cisco Data Center Network Manager is primarily focused on managing the network fabric and not application deployments. CloudCenter, due to its application-centric and hybrid focus, correctly satisfies the specific customer requirement to orchestrate and optimize the cost of their application across their on-premises ACI environment and the public cloud.

Authoritative links:

Cisco CloudCenter (now part of Intersight):https://www.cisco.com/c/en/us/products/cloud-systems-management/cloudcenter/index.html
Cisco Intersight:https://www.cisco.com/c/en/us/products/servers-unified-computing/intersight/index.html

Question: 49

DRAG DROP -

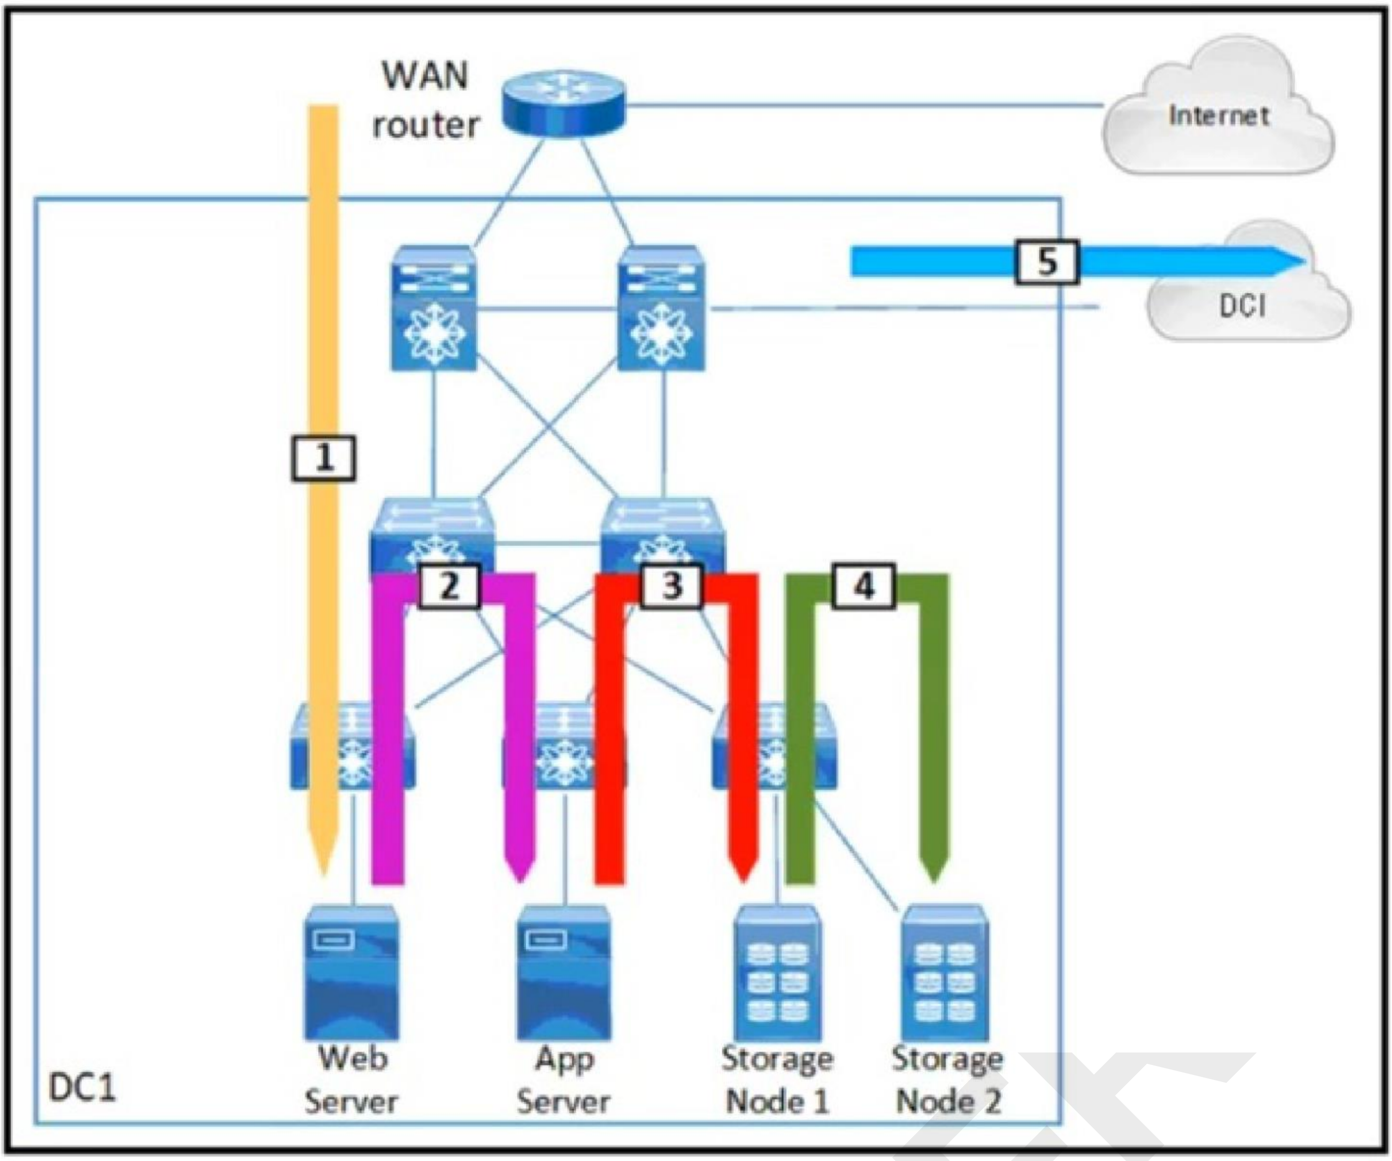

Refer to the exhibit. Drag and drop each traffic flow type from the left onto the corresponding number on the right. Not all traffic flow types are used.
Select and Place:

Inter-Data Center

East-West

North-West

North-South

South-West

Storage Traffic

Storage Replication

1

2

3

4

5

Answer:

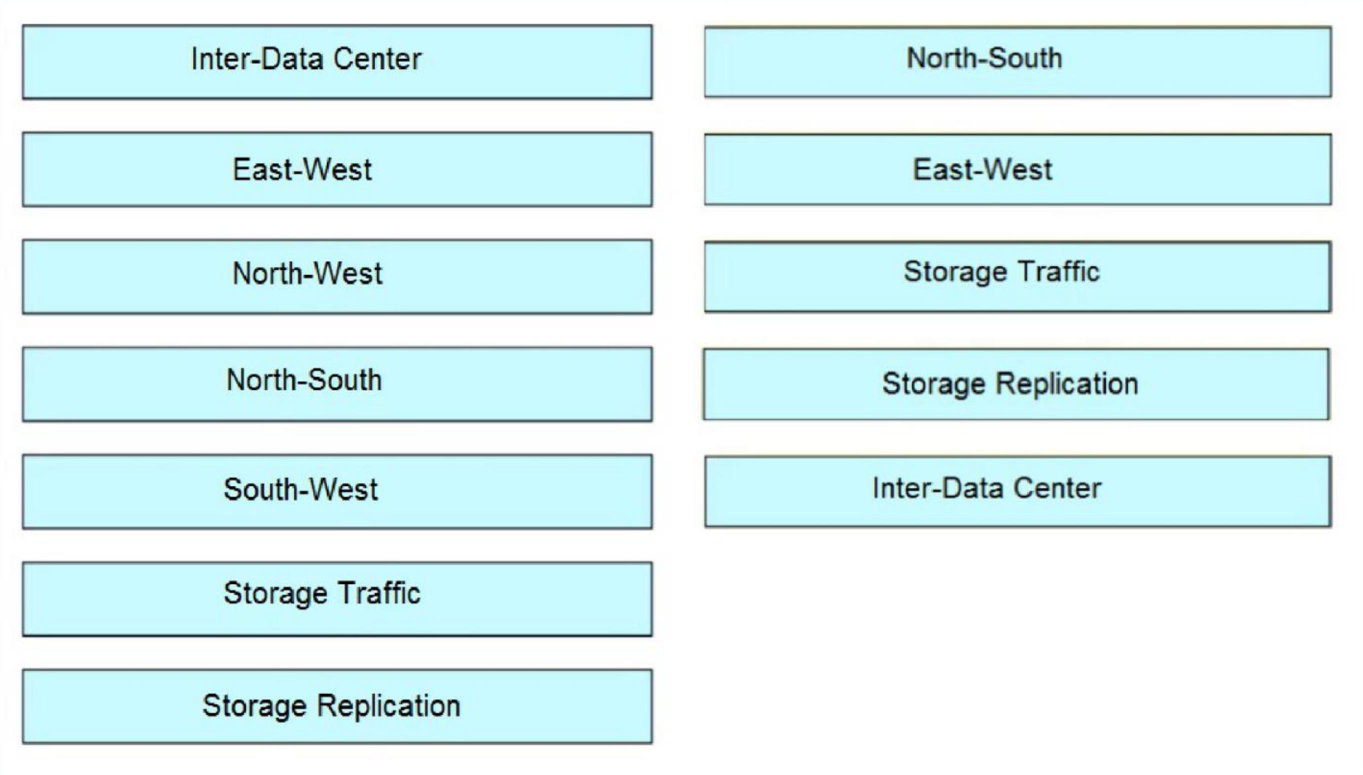

Inter-Data Center	North-South
East-West	East-West
North-West	Storage Traffic
North-South	Storage Replication
South-West	Inter-Data Center
Storage Traffic	
Storage Replication	

Explanation:

North-South = 1

East-West = 2

Storage Traffic = 3

Storage Replication = 4

Inter- Data Center (DCI) = 5

Question: 50

DRAG DROP -

A network engineer is asked to describe the cloud infrastructure models from the perspective of their operation and access to resources. Drag and drop the descriptions from the left onto the appropriate characteristics on the right.

Select and Place:

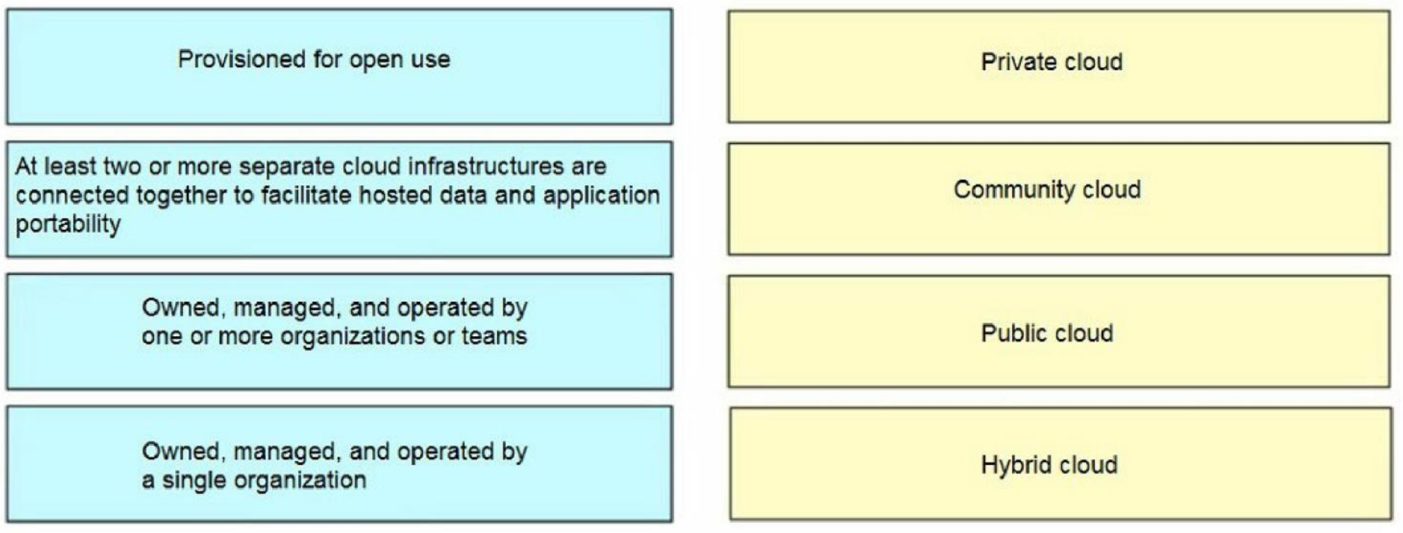

Provisioned for open use	Private cloud
At least two or more separate cloud infrastructures are connected together to facilitate hosted data and application portability	Community cloud
Owned, managed, and operated by one or more organizations or teams	Public cloud
Owned, managed, and operated by a single organization	Hybrid cloud

Answer:

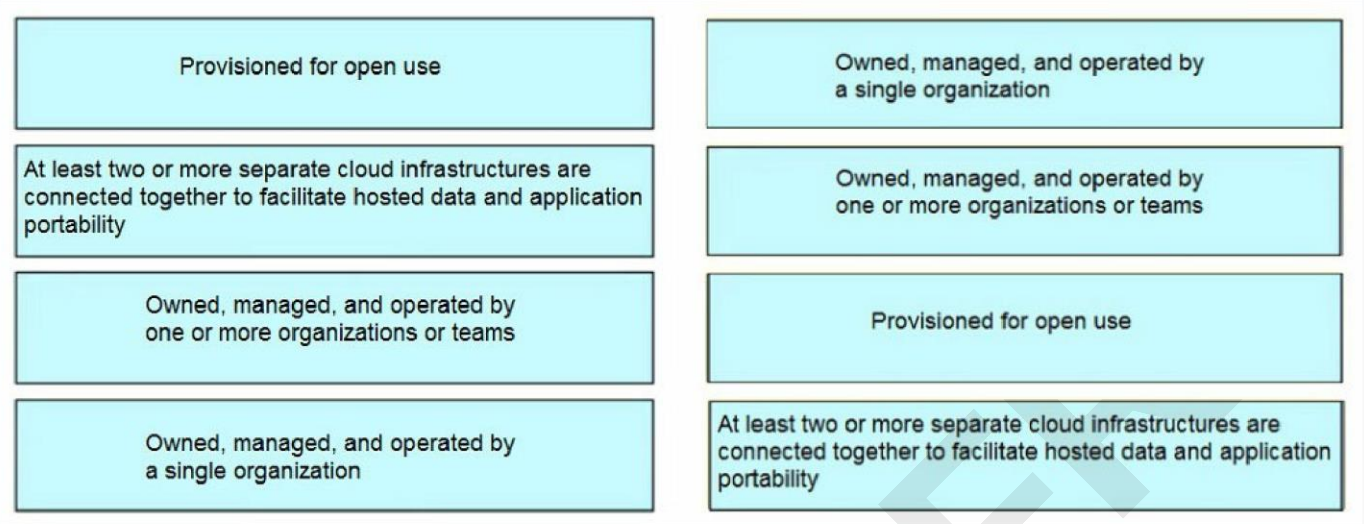

Provisioned for open use	Owned, managed, and operated by a single organization
At least two or more separate cloud infrastructures are connected together to facilitate hosted data and application portability	Owned, managed, and operated by one or more organizations or teams
Owned, managed, and operated by one or more organizations or teams	Provisioned for open use
Owned, managed, and operated by a single organization	At least two or more separate cloud infrastructures are connected together to facilitate hosted data and application portability

Question: 51

An engineer must configure OSPF routing on Cisco Nexus 9000 Series Switches. The IP subnet of the Eth1/2 interfaces for both switches must be advertised via
OSPF. However, these interfaces must not establish OSPF adjacency or send routing updates. The current OSPF adjacency over the interfaces Eth1/1 on SW1 and Eth1/1 on SW2 must remain unaffected. Which configuration must be applied to both Nexus switches to meet these requirements?

A.

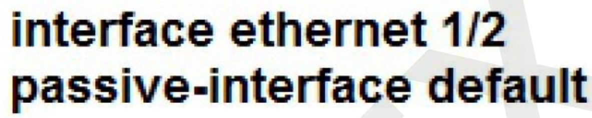

```
interface ethernet 1/2
passive-interface default
```

B.

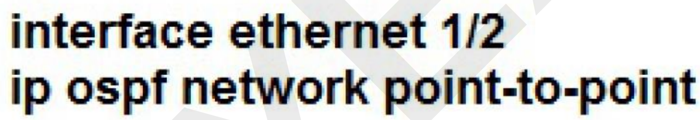

```
interface ethernet 1/2
ip ospf network point-to-point
```

C.

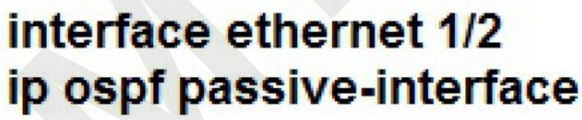

```
interface ethernet 1/2
ip ospf passive-interface
```

D.

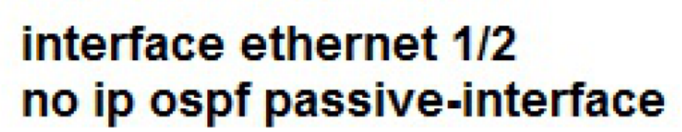

```
interface ethernet 1/2
no ip ospf passive-interface
```

Answer: C

Explanation:

C is correct

(config-if)# ip ospf passive-interface > disables neighborship but still advertise the network only on interface ethernet 1/2 as the question asked.

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/6-x/unicast/configuration/guide/l3_cli_nxos/l3_ospf.html

Question: 52

A network engineer must perform a backup and restore of the Cisco Nexus 5000 Series Switch configuration. The backup must be made to an external backup server. The only protocol permitted between the Cisco Nexus Series switch and the backup server is UDP. The backup must be used when the current working configuration of the switch gets corrupted. Which set of steps must be taken to meet these requirements?

A. 1. Perform a running-config backup to an SFTP server. 2. Copy backup-config from the SFTP server to the running-config file.

B. 1. Perform a startup-config backup to a TFTP server. 2. Copy backup-config from the backup server to the running-config file.

C. 1. Perform a running-config backup to an SCP server. 2. Copy running-config in the boot flash to the running-config file.

D. 1. Perform a startup-config backup to an FTP server. 2. Copy startup-config in the boot flash to the running-config file.

Answer: B

Explanation:

The correct answer is **B** because it aligns with the requirements of backing up the switch configuration using UDP and restoring it in case of corruption. Let's break down why:

Backup Protocol: The question explicitly states that only UDP is permitted. TFTP (Trivial File Transfer Protocol) is the only protocol among the options that utilizes UDP for data transfer, making it the sole suitable choice for the backup server connection. SFTP, SCP, and FTP rely on TCP, which is prohibited here.

Backup Target: We need to back up the configuration, specifically the "startup-config", as this configuration is loaded during the switch's boot-up process. Backing up the "running-config" would only save the current operational settings, not the configuration that will be loaded on a reboot. The startup-config is the correct candidate for restoring the device to a known good state during system corruption.

Restoration Process: To restore the switch's configuration after a corruption, the backed-up startup-config from the TFTP server needs to be copied to the switch's running configuration. This will overwrite the corrupted running configuration and ensure proper device operation upon next reboot. Options that copy from bootflash do not fulfill the restore requirement and are irrelevant to the backup made to the TFTP server. **Option A** proposes SFTP, which is TCP-based, violating the UDP-only requirement. Additionally, it incorrectly suggests backing up the running-config.

Option C proposes SCP, also TCP-based, and does not involve the use of the backed-up configuration. Furthermore, it uses the wrong backup config file.

Option D proposes FTP, another TCP-based protocol, and it uses the wrong location for restoring the configuration.In summary, **option B** correctly identifies TFTP as the only UDP-based protocol, recommends backing up the startup configuration, and offers a proper procedure for restoring a corrupt configuration.

Authoritative Links:

Cisco Nexus 5000 Series Configuration Guide: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/configuration/guide/config/n5k_config (This provides overall documentation for the Nexus 5000, including configuration backup and restore)
Understanding TFTP:https://www.ibm.com/docs/en/zvm/7.2?topic=protocols-tftp-trivial-file-transfer-protocol (Provides an overview of the TFTP protocol and its characteristics)

Question: 53

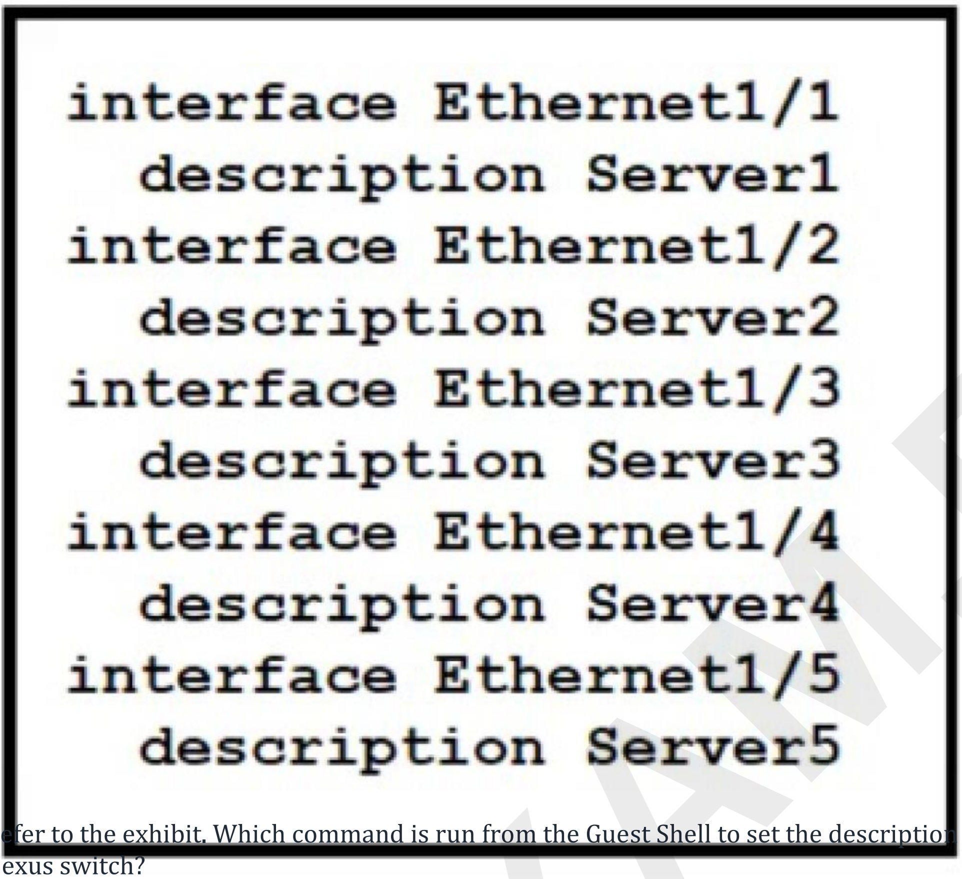

```
interface Ethernet1/1
  description Server1
interface Ethernet1/2
  description Server2
interface Ethernet1/3
  description Server3
interface Ethernet1/4
  description Server4
interface Ethernet1/5
  description Server5
```

Refer to the exhibit. Which command is run from the Guest Shell to set the description on the first five interfaces of the Cisco Nexus switch?

A. [[email protected] ~]$ for x in 1..5 ; dohost conf t ; interface eth1/$x ; description Server$x;
B. [[email protected] ~]$ for x in 1..5 ; do dohost conf t ; interface eth1/$x ; description Server$x; done
C. [[email protected] ~]$ for x in (1..5); do dohost conf t; interface eth1/$x ; description Server$x; done
D. [[email protected] ~]$ for x in (1..5); dohost conf t ; interface eth1/$x ; description Server$x;

Answer: B

Explanation:

Correct B

```
NXs1# guestshell
[[email protected] ~]$ for x in 1..5 ; do dohost "conf t ; interface eth1/$x ; description Server$x" ; done Enter
configuration commands, one per line. End with CNTL/Z.

Enter configuration commands, one per line. End with CNTL/Z.

Enter configuration commands, one per line. End with CNTL/Z.

Enter configuration commands, one per line. End with CNTL/Z.

Enter configuration commands, one per line. End with CNTL/Z.

[[email protected] ~]$
```

Question: 54

An engineer must configure multiple EPGs on a single access port in a large Cisco ACI fabric without using VMM integration. The relevant access policies and tenant policies have been created. A single AAEP is used to configure the access ports in the fabric. Which two additional steps must be taken to complete the configuration? (Choose two.)

A. The EPGs must link directly to the corresponding AAEP
B. A contract must be defined between the EPGs
C. The EPGs must be configured as static ports
D. The corresponding bridge domains must be configured in legacy mode
E. The EPGs must be linked to the correct physical domain

Answer: CE

Explanation:

Here's a detailed justification for the correct answer (C and E) to the Cisco DCCOR question, explaining why other options are incorrect:

The question focuses on configuring multiple Endpoint Groups (EPGs) on a single access port within a Cisco Application Centric Infrastructure (ACI) fabric without Virtual Machine Manager (VMM) integration. To achieve this, two key steps are required:

C. The EPGs must be configured as static ports: When not using VMM integration, EPGs can't dynamically learn endpoints. Instead, endpoints connected to specific ports must be statically associated with EPGs. This tells ACI which endpoint belongs to which EPG on that port, allowing for proper policy enforcement. In other words, you have to configure which VLAN to which EPG is assigned on a particular port.

E. The EPGs must be linked to the correct physical domain: A physical domain maps to a particular VLAN pool or range. By linking the EPGs to the correct physical domain, you're telling ACI to use specific VLAN tags for the traffic associated with each EPG. This ensures that each EPG's traffic is properly tagged and segregated on the physical network. In this static port configuration without VMM, the physical domain provides the mechanism for identifying and segregating different traffic types using VLAN tags.

Let's discuss why the other options are incorrect:

A. The EPGs must link directly to the corresponding AAEP: The Attachment Access Entity Profile (AAEP) defines the access policy associated with a port. EPGs do not directly link to the AAEP. Instead, the AAEP is associated with interface policy groups, which are then associated with the physical interfaces. EPGs are instead linked to a domain, which is then associated with the AAEP.

B. A contract must be defined between the EPGs: While contracts are required for inter-EPG communication, they're not directly related to the configuration of multiple EPGs on a single port. Contracts define how EPGs can communicate with each other but aren't necessary for the initial setup of EPGs on a single port. Contracts are not necessary for allowing traffic within a single EPG.

D. The corresponding bridge domains must be configured in legacy mode: Bridge domains can operate in legacy mode but it's not required for the EPG setup described. Legacy mode is typically for compatibility with older Layer 2 protocols and isn't essential to configure multiple EPGs on a single access port. A bridge domain must be configured for each EPG, but legacy mode is not required.

In summary, correctly configuring multiple EPGs on a single access port requires static port configuration and associating those EPGs to a physical domain that allows VLAN tagging. This is the core of the non-VMM configuration within ACI.

Authoritative Links for Further Research:

Cisco ACI Fundamentals:https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/5x/fundamentals/cisco-aci-fundamentals-52x.html
Cisco ACI Policy Model:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/white-paper-c11-733972.html
Cisco ACI Static Binding:https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/6x/l2-configuration/cisco-apic-layer-2-configuration-60x/m-config-static-binding.html

Question: 55

Several production and development database servers exist in the same EPG and IP subnet. The IT security policy is to prevent connections between production and development. Which attribute must be used to assign the servers to different microsegments?

A. data center
B. VM name
C. IP address
D. VMM domain

Answer: B

Explanation:

The correct answer is **B. VM name**. Here's a detailed justification:

Microsegmentation, a crucial security practice in modern data centers, aims to create fine-grained security policies within a network. The goal is to isolate workloads, limiting the impact of breaches and controlling lateral movement. In this scenario, the objective is to prevent communication between production and development database servers within the same EPG (Endpoint Group) and IP subnet. Since these servers reside in the same network constructs, traditional network-based segmentation via VLANs or subnets isn't sufficient.

Option A, "data center," is too broad and not applicable for microsegmenting servers within the same data center. Option C, "IP address," is inadequate for this purpose because servers in the same subnet often share IP address ranges, making isolation difficult using just IP filtering. Option D, "VMM domain," while relevant to virtualization, is typically used for managing the virtualization infrastructure itself, not for application-level microsegmentation within an EPG.

VM name, on the other hand, provides a specific, unique, and manageable attribute for each virtual machine.

Cisco ACI (Application Centric Infrastructure), or similar SDN solutions, leverage this attribute to create policy-based segmentation. By creating contracts (rules) that specify that VMs with names containing "production" can only connect to each other and VMs with names containing "development" can only connect to each other, the desired isolation is achieved. This approach allows for dynamic changes and updates, as policies are tied to the names rather than static IPs or network segments. This enables logical grouping of VMs regardless of physical location or network configuration.

Microsegmentation using VM names or tags is a common and practical method for achieving this type of isolation within an EPG. It leverages the inherent manageability of virtualized environments.

Authoritative Links for Further Research:

Cisco ACI Microsegmentation:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/white-paper-c11-734375.html
VMware NSX Microsegmentation:https://www.vmware.com/solutions/cloud-native-apps/microsegmentation.html
What is Microsegmentation?https://www.paloaltonetworks.com/cyberpedia/what-is-microsegmentation

Question: 56

An environment consists of a Cisco MDS 9000 Series Switch that uses port channels. An engineer must ensure that frames between the source and the destination follow the same links for a specific flow. Subsequent flows are allowed to use a different link. Which load balancing method should be used to accomplish this goal?

A. src-id/dst-id
B. src-dst-oui
C. src-dst-port
D. src-id/dst-id/oxid

Answer: D

Explanation:

The correct load balancing method for ensuring that frames of a specific flow consistently follow the same link within a port channel on a Cisco MDS 9000 Series Switch, while allowing subsequent flows to utilize different links, is **D. src-id/dst-id/oxid**.

Here's why:

Port channel load balancing distributes traffic across multiple physical links within the channel. The goal is efficient utilization and redundancy. Different load balancing algorithms are available, each using different parameters to determine which link should carry the traffic. Options A (src-id/dst-id), B (src-dst-oui) and C (src-dst-port) operate at Layer 2 and 3 level and are not optimal for Fibre Channel where source and destination identifiers may not consistently identify the flow.

src-id/dst-id/oxid: This method uses the source Fibre Channel ID (FCID), the destination FCID, and the Exchange Originator ID (OXID). FCID are the unique identifies for fibre channel devices, and OXID represents the exchange within the flow. Because OXID is the per-flow identifier, using it as part of the load balancing parameters makes sure packets from the same flow follow the same path, while packets from different flows (even if originating from same source, and targetting same destination), will take a different path.Options A, B, and C provide coarser load balancing. They might map different flows to the same link, failing to ensure path consistency for a particular flow. Since the question explicitly asks to ensure that a particular flow will always follow the same link while allowing subsequent flows to different links, the usage of unique identifier that is tied with the flow itself becomes important.

Therefore, **src-id/dst-id/oxid** is the correct approach because it incorporates unique identifiers tied to the specific flow, enabling consistent path selection while permitting subsequent flows to take different routes. This is important for maintaining flow integrity in Fibre Channel environments.

Authoritative Links:

Cisco MDS 9000 Series NX-OS Fabric Configuration Guide: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/nx-os/fabric/configuration/guide/cisco_mds9000_fabric_config/port_channels.html (Search for "Load Balancing" in the document for more information)
Cisco MDS 9000 Family Command Reference: (Search for port channel command) https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/nx-os/command/reference/index.html

Question: 57

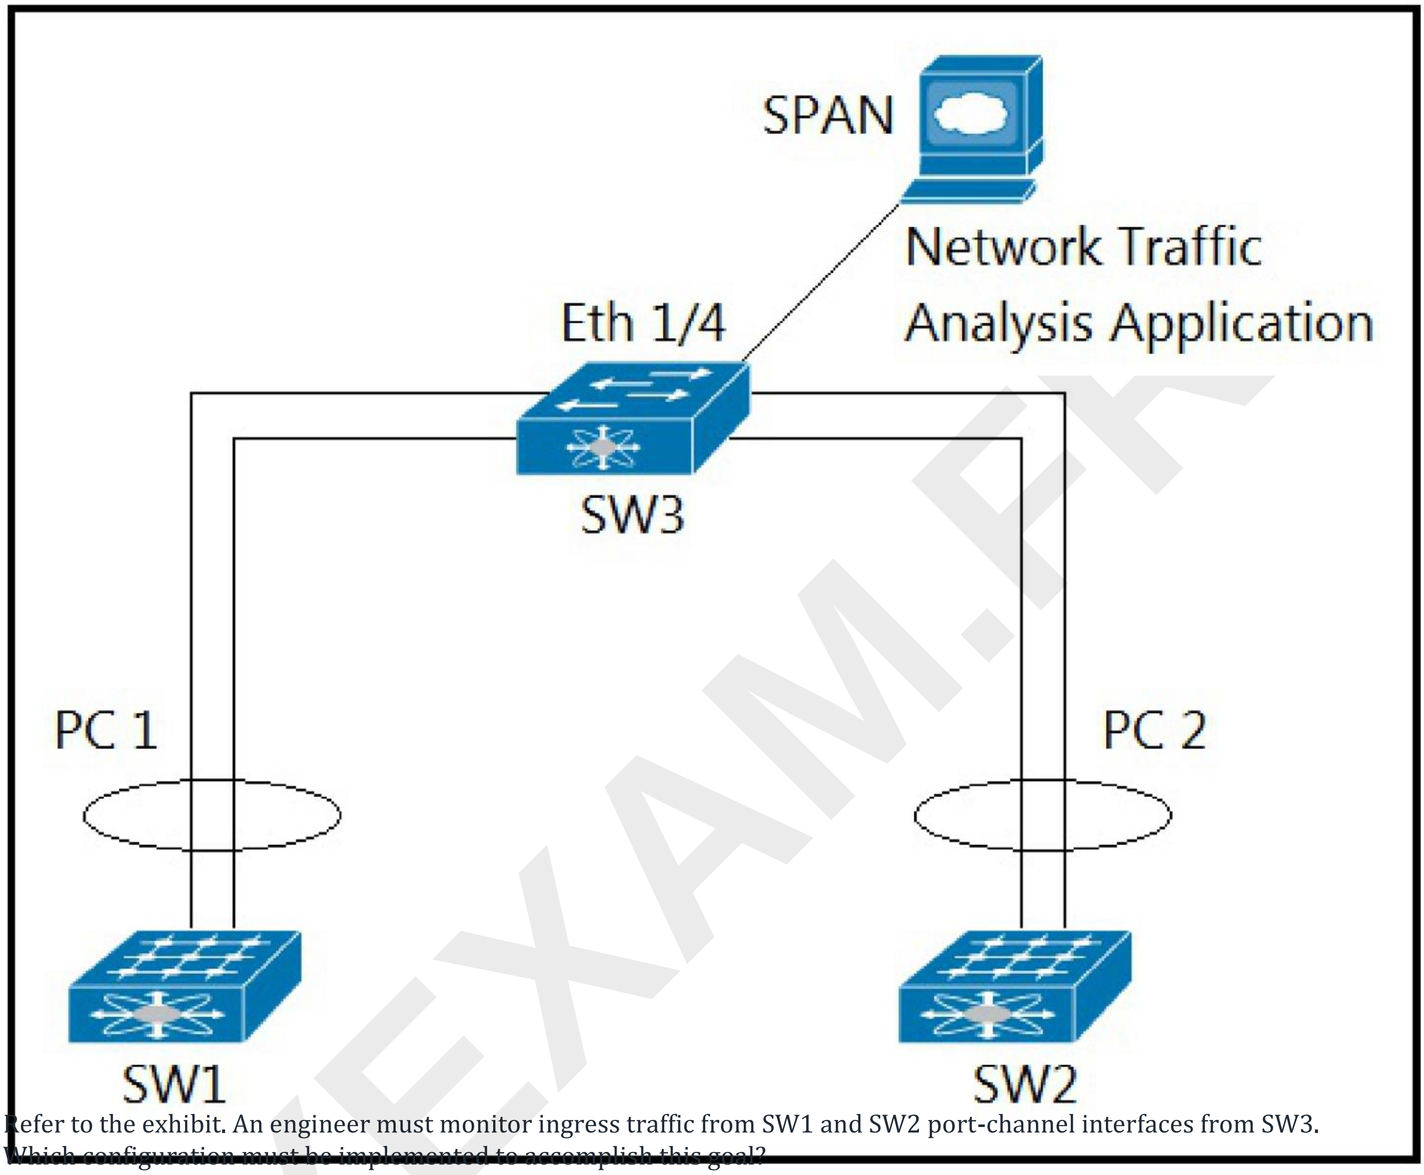

Refer to the exhibit. An engineer must monitor ingress traffic from SW1 and SW2 port-channel interfaces from SW3. Which configuration must be implemented to accomplish this goal?

A. SW3# configure terminal SW3(config)# source interface port-channel 1 rx SW3(config)# source interface port-channel 2 tx SW3(config)# interface ethernet 1/4 SW3(config-if)# destination interface ethernet 1/4 SW3(config-if)# exit SW3(config)# interface Port-channel 1-2 SW3(config-if)# switchport monitor SW3(config-if)# exit

B. SW3# configure terminal SW3(config)# monitor session 2 SW3(config-monitor)# source interface port-channel 1 rx SW3(config-monitor)# source interface port-channel 2 rx SW3(config)# interface ethernet 1/4 SW3(config-if)# switchport monitor SW3(config-if)# exit SW3(config)# monitor session 2 SW3(config-monitor)# destination interface ethernet 1/4

C. SW3# configure terminal SW3(config)# monitor session 2 SW3(config-monitor)# source interface port-channel 1 tx SW3(config-monitor)# source interface port-channel 2 tx SW3(config)# interface Port-channel 1 SW3(config-if)# switchport monitor SW3(config-if)# exit SW3(config)# interface Port-channel 2 SW3(config-if)# switchport monitor SW3(config-if)# exit

D. SW3# configure terminal SW3(config)# monitor session 2 SW3(config-monitor)# source interface port-channel 1 tx SW3(config-monitor)# source interface port-channel 2 rx SW3(config)# interface port-channel 1 SW3(config-if)# switchport monitor SW3(config-if)# exit SW3(config)# interface port-channel 2 SW3(config-if)# switchport monitor SW3(config-if)# exit

Answer: B

Explanation:
B is correct since the ingress traffic (traffic received from SW1 and SW2) must be the source interface on SW3:

```
SW3# configure terminal -

SW3(config)# monitor session 2 -
SW3(config-monitor)# source interface port-channel 1 rx
SW3(config-monitor)# source interface port-channel 2 rx
SW3(config)# interface ethernet 1/4
SW3(config-if)# switchport monitor

SW3(config-if)# exit -

SW3(config)# monitor session 2 -
SW3(config-monitor)# destination interface ethernet 1/4
```

Question: 58

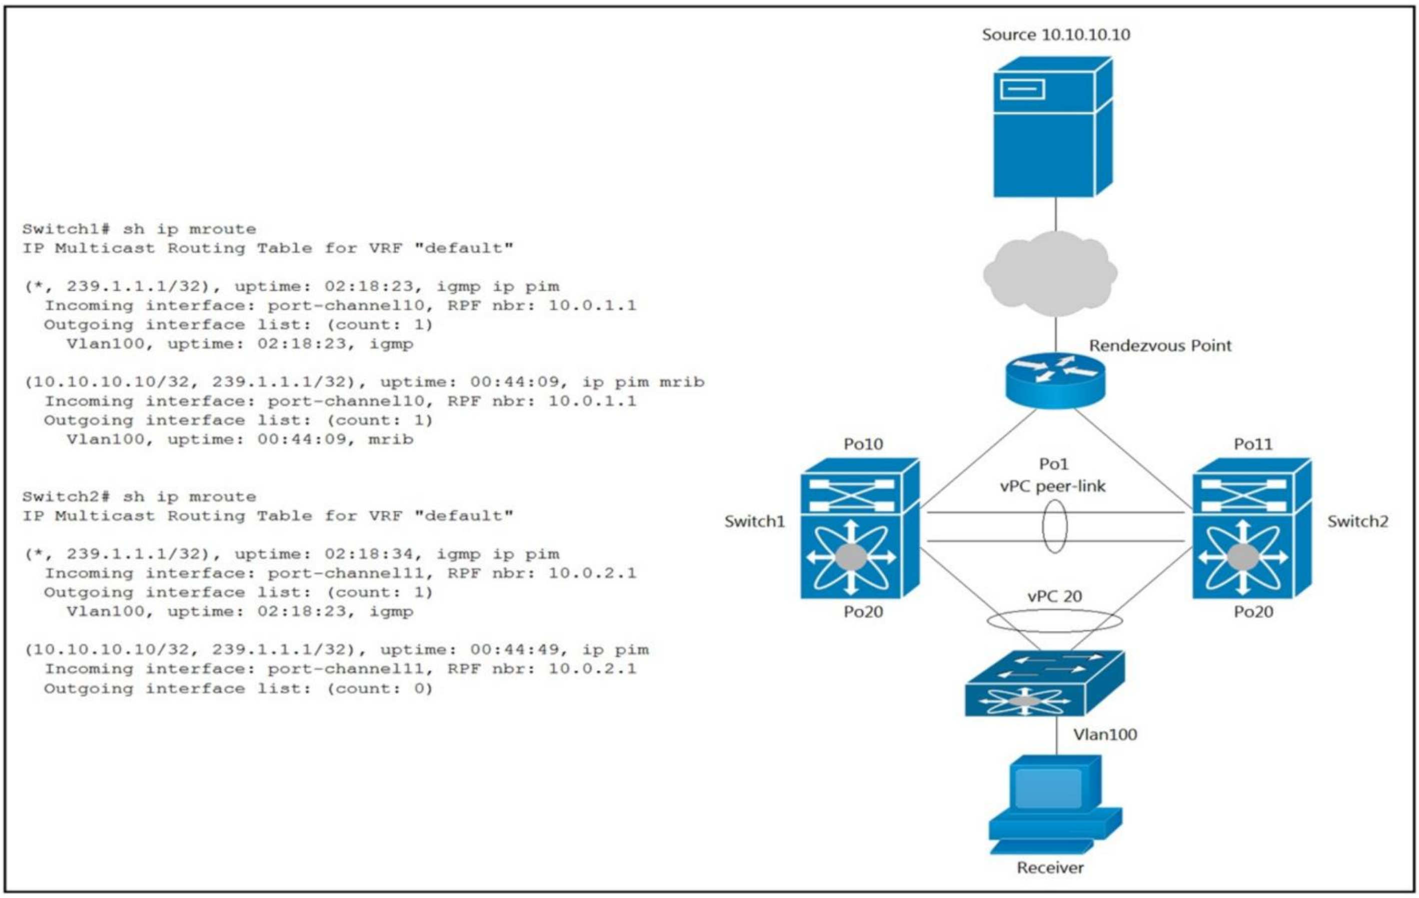

```
Switch1# sh ip mroute
IP Multicast Routing Table for VRF "default"

(*, 239.1.1.1/32), uptime: 02:18:23, igmp ip pim
  Incoming interface: port-channel10, RPF nbr: 10.0.1.1
  Outgoing interface list: (count: 1)
    Vlan100, uptime: 02:18:23, igmp

(10.10.10.10/32, 239.1.1.1/32), uptime: 00:44:09, ip pim mrib
  Incoming interface: port-channel10, RPF nbr: 10.0.1.1
  Outgoing interface list: (count: 1)
    Vlan100, uptime: 00:44:09, mrib


Switch2# sh ip mroute
IP Multicast Routing Table for VRF "default"

(*, 239.1.1.1/32), uptime: 02:18:34, igmp ip pim
  Incoming interface: port-channel11, RPF nbr: 10.0.2.1
  Outgoing interface list: (count: 1)
    Vlan100, uptime: 02:18:23, igmp

(10.10.10.10/32, 239.1.1.1/32), uptime: 00:44:49, ip pim
  Incoming interface: port-channel11, RPF nbr: 10.0.2.1
  Outgoing interface list: (count: 0)
```

Refer to the exhibit. A host with a source address 10.10.10.10 sends traffic to multicast group 239.1.1.1. How do the vPC switches forward the multicast traffic?

A. If multicast traffic is received on Switch1 over the vPC peer-link, the traffic is dropped B. If multicast traffic is received on Po10 Switch1, the traffic is forwarded out on Po1 and Po20 C. If multicast traffic is received on Po11 Switch2, the traffic is forwarded out only on Po20 D. If multicast traffic is received on Po11 Switch2, the traffic is dropped.

Answer: D

Explanation:
Switch2 will never get the stream due to a missing OIF on Switch 2

Reference:
https://www.cisco.com/c/en/us/support/docs/ip/multicast/214140-multicast-forwarding-in-vpc-based-on-loc.html#anc8

Question: 59

A network architect must redesign a data center network based on OSPFv2. The network must perform fast reconvergence between directly connected switches.
Which two actions must be taken to meet the requirements? (Choose two.)

A. Set low OSPF hello and DEAD timers.

B. Configure all links on AREA 0.

C. Enable BFD for failure detection.

D. Use OSPF point-to-point links only.

E. Implement a virtual link between the switches.

Answer: CD

Explanation:

The question requires optimizing OSPFv2 for rapid reconvergence between directly connected data center switches.

Option C: Enable BFD for failure detection. BFD (Bidirectional Forwarding Detection) is a low-overhead protocol designed for fast failure detection. By enabling BFD between OSPF neighbors, link failures can be detected much faster than relying solely on OSPF hello timers. This quick detection allows OSPF to react rapidly, triggering faster reconvergence. https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/configuration/15-sy/iro-15-sy-book/ospf-bfd.html

Option D: Use OSPF point-to-point links only. Point-to-point links simplify the OSPF neighbor discovery process. Unlike broadcast networks, OSPF does not need to elect a Designated Router (DR) and Backup Designated Router (BDR) on a point-to-point link. The absence of DR/BDR election, as well as the lack of neighbor discovery in the same way, speeds up neighbor adjacencies and thus leads to quicker convergence. The elimination of DR/BDR election also reduces overhead on the switch.
https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html

Why other options are not suitable:

Option A: Set low OSPF hello and DEAD timers. While reducing these timers can improve convergence speed, doing so aggressively can cause network instability due to false link failure detections. BFD provides more efficient and faster failure detection, making it superior.

Option B: Configure all links on AREA 0. While a good practice, placing all links in area 0 doesn't directly address fast reconvergence between directly connected switches. Areas help with summarization but don't impact the speed at which a directly connected switch recognizes a failed link.

Option E: Implement a virtual link between the switches. Virtual links are used to connect non-backbone OSPF areas. They are not applicable to directly connected switches and they do not help with fast reconvergence between switches.

In summary, enabling BFD (C) and using point-to-point links (D) are key actions for achieving rapid reconvergence in a data center OSPFv2 network as they offer the fastest and most effective way to detect link failures and react accordingly.

Question: 60

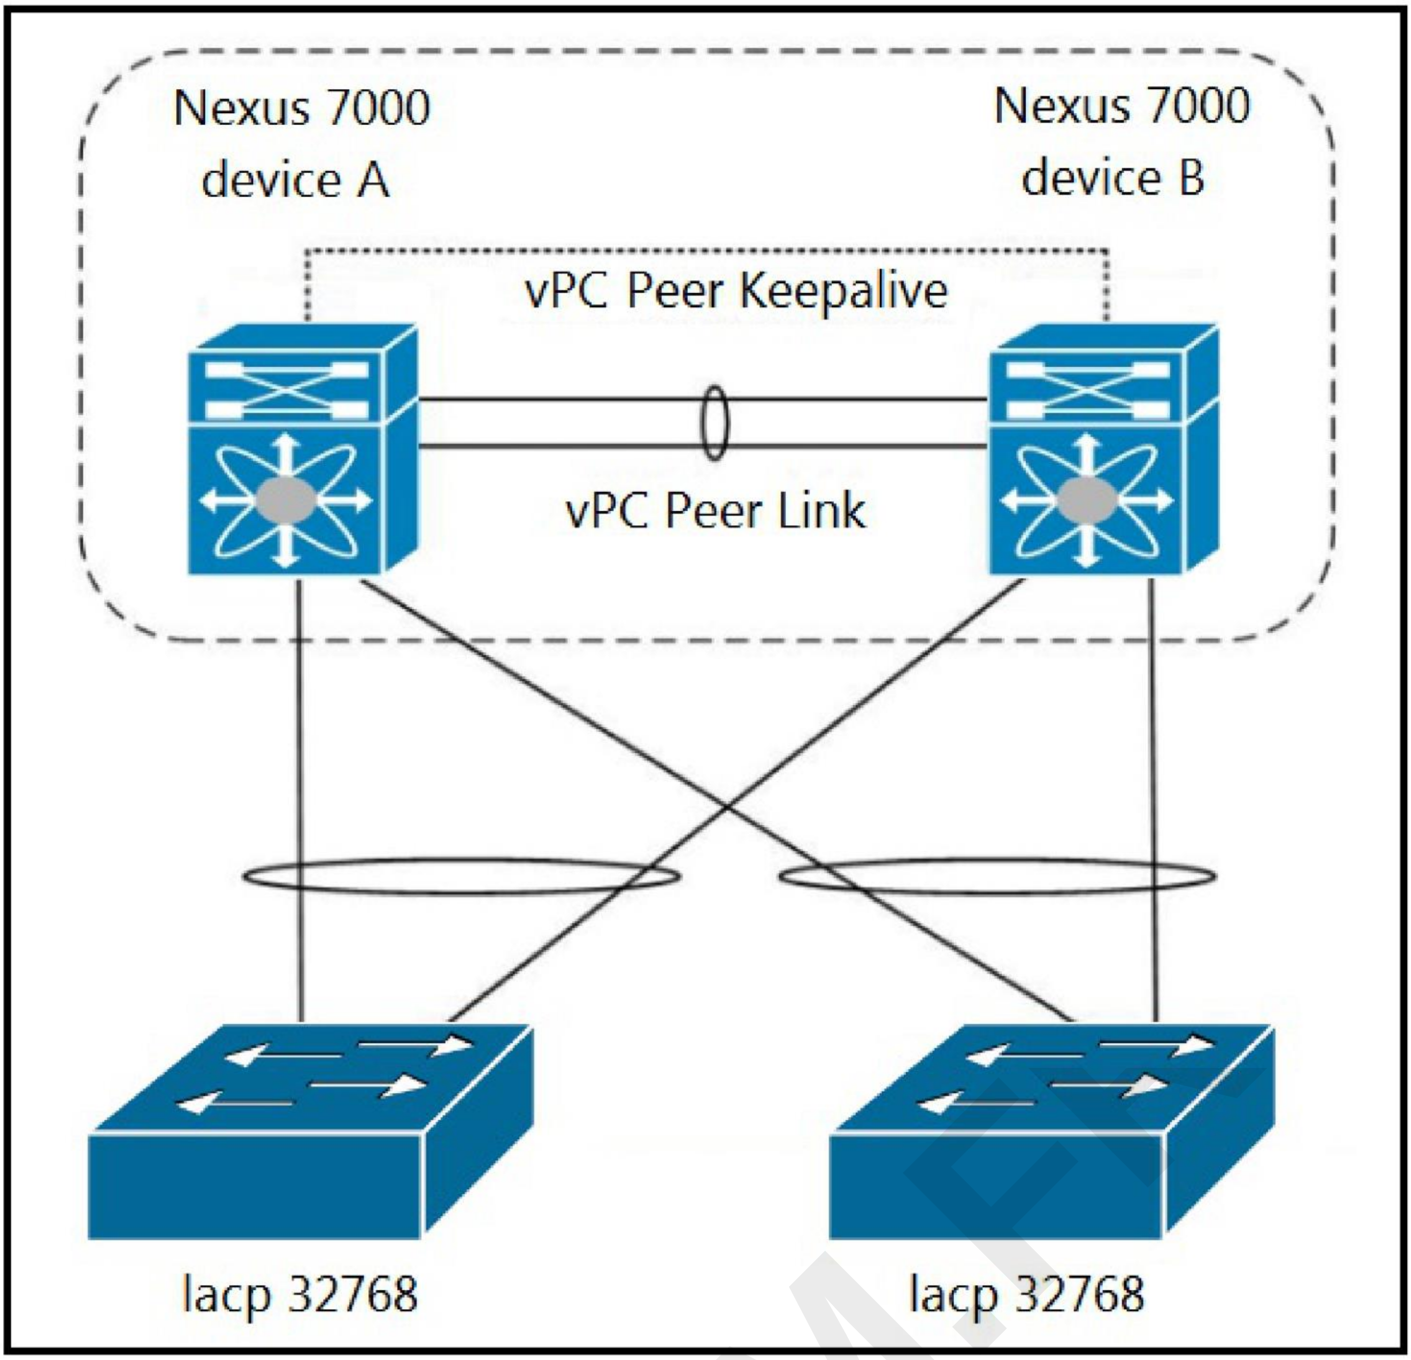

Refer to the exhibit. Which configuration ensures that the Cisco Nexus 7000 Series Switches are the primary devices for LACP?

A. N7K_A(config-vpc-domain)# role priority 1 N7K_B(config-vpc-domain)# role priority 2
B. N7K_A(config-vpc-domain)# system-priority 32768 N7K_B(config-vpc-domain)# system-priority 32768
C. N7K_A(config-vpc-domain)# system-priority 100 N7K_B(config-vpc-domain)# system-priority 200
D. N7K_A(config-vpc-domain)# system-priority 4000 N7K_B(config-vpc-domain)# system-priority 4000

Answer: D

Explanation:
You should manually configure the vPC system priority when you are running Link Aggregation Control Protocol (LACP) to help ensure that the vPC peer devices are the primary devices on LACP.

When you manually configure the system priority, ensure that you configure the same priority value on both vPC peer devices. If these values do not match, vPC will not come up. The defaut priority is 32667, and a lower value will become the primary.

Reference:
https://community.cisco.com/t5/server-networking/vpc-system-priority/td-p/1650914

Question: 61

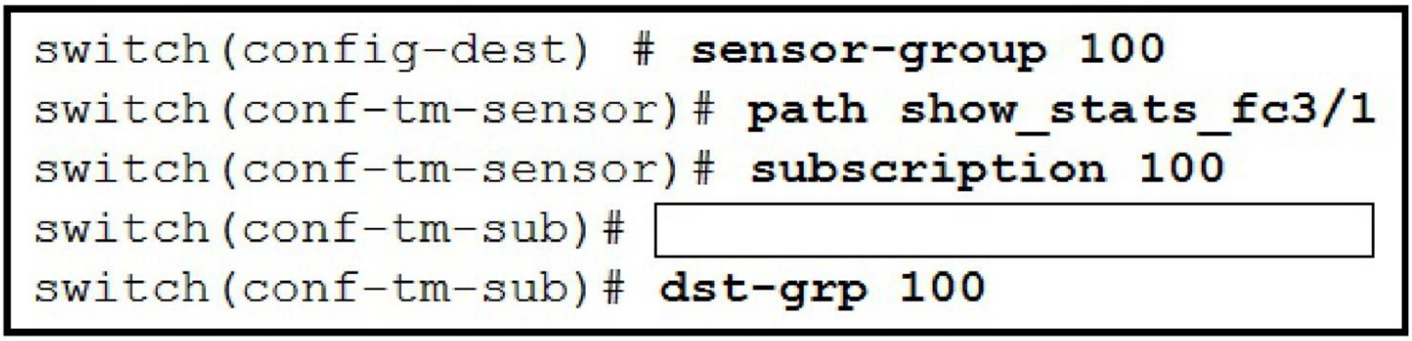

```
switch(config-dest) # sensor-group 100
switch(conf-tm-sensor)# path show_stats_fc3/1
switch(conf-tm-sensor)# subscription 100
switch(conf-tm-sub)#
switch(conf-tm-sub)# dst-grp 100
```

Refer to the exhibit. An engineer needs to implement streaming telemetry on a Cisco MDS 9000 Series Switch. The requirement is for the show command data to be collected every 30 seconds and sent to receivers. Which command must be added to the configuration to meet this requirement?

A. snsr-grp 200 sample-interval 30000
B. snsr-grp 200 sample-interval 30
C. sensor-grp 200 sample-period 30
D. sensor-grp 200 sample-period 30000

Answer: A

Explanation:
Link the sensor group with an ID to the subscription node and set the data streaming sample interval in milliseconds: switch(conf-tm-sub)# snsr-grp id sample-interval interval
Note: The minimum streaming sample interval that is recommended is 30000.

Reference:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/8_x/config/san_analytics/cisco-mds9000-san-analytics-telemetry-streaming- config-guide-8x/configuring-san-telemetry-streaming.html

Question: 62

An engineer must configure OSPF in the data center. The external routes have already been redistributed into OSPF. The network must meet these criteria:
* The data center servers must reach services in the cloud and the services behind the redistributed routes. * The exit point toward the Internet should be propagated only when there is a dynamically learned default route from the upstream router.
Which feature is required?

A. stubby area
B. totally stubby area
C. default-information originate
D. default-information originate always

Answer: C

Explanation:

The correct answer is **C. default-information originate**. Here's why:

The scenario requires OSPF to propagate a default route (0.0.0.0/0) specifically when it's learned dynamically from an upstream router. This dynamic learning usually happens through protocols like BGP, which connects the data center to the internet. OSPF itself doesn't inherently learn default routes from outside; they must be injected. default-information originate is the command in OSPF that allows for this behavior, meaning it conditionally injects a default route into the OSPF domain when the router itself has a valid default route in its routing table, learned through another source.

Option A, **stubby area**, while limiting external route advertisements, doesn't dynamically inject a default route based on upstream conditions. A stub area blocks type 5 LSAs. Option B, **totally stubby area**, also focuses on route filtering, blocking both type 5 and type 3 LSAs, forcing all traffic towards the ABR. However, it still doesn't inherently fulfill the condition of originating a default route only when one is dynamically learned.

Option D, **default-information originate always**, will inject a default route regardless of whether the OSPF router has learned one dynamically or not. This doesn't meet the requirement of conditionally advertising the default route based on upstream connectivity. The core need of the question involves originating a default route when a specific condition is met - and default-information originate allows this. It provides a way to conditionally advertise the default route based on the router's active routing table, making it the most suitable option for the scenario. Using default-information originate allows the data center to rely on the upstream connection and ensures the default route is advertised only when a dynamically learned internet route exists.

Authoritative Links for further research:

Cisco Documentation on OSPF: https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-ospf-default.html
OSPF Default Route Injection: https://networklessons.com/ospf/ospf-default-route

Question: 63

An engineer must configure a VXLAN routing on a Cisco Nexus 9000 Series Switch. The engineer requires a solution where all the leaf switches have the same gateway MAC and IP addresses. Which configuration set accomplishes this task?

A. NX9K(config)# fabric forwarding anycast-gateway-mac AA:BB:AA:BB:AA:BB NX9K(config)# interface VLAN-interface-name NX9K(config-if)# vrf member vrf-name NX9K(config-if)# fabric forwarding mode anycast-gateway

B. NX9K(config)# install feature-set fabric NX9K(config)# feature-set fabric NX9K(config)# fabric forwarding anycast-gateway-mac AA:BB:AA:BB:AA:BB NX9K(config)# interface VLAN-interface-name NX9K(config-if)# fabric forwarding mode anycast-gateway

C. NX9K(config)# install feature-set fabric NX9K(config)# feature-set fabric NX9K(config)# fabric forwarding anycast-gateway-mac AA:BB:AA:BB:AA:BB NX9K(config)# interface VLAN-interface-name NX9K(config-if)# vrf member vrf-name NX9K(config-if)# fabric forwarding mode anycast-gateway

D. NX9K(config)# fabric forwarding anycast-gateway-mac AA:BB:AA:BB:AA:BB NX9K(config)# interface VLAN-interface-name NX9K(config-if)# fabric forwarding mode anycast-gateway

Answer: A

Explanation:

The correct answer is A. This configuration implements an anycast gateway, a crucial concept in VXLAN environments for achieving active-active forwarding and simplified gateway management. The fabric forwarding anycast-gateway-mac AA:BB:AA:BB:AA:BB command globally sets the common MAC address for the anycast gateway across all leaf switches in the fabric. Each leaf switch then uses this common MAC address when acting as the gateway for a particular VLAN. The interface VLAN-interface-name command enters interface configuration mode for the specific VLAN. Next, vrf member vrf-name associates this VLAN interface with the desired VRF. Finally, fabric forwarding mode anycast-gateway enables the anycast gateway functionality on the interface, allowing the interface to act as an anycast gateway using the predefined MAC.

Option B lacks the VRF association, which is necessary in a VXLAN environment that uses VRFs for tenant segmentation. Option C includes extra install feature-set fabric and feature-set fabric commands, which are not required for the task, and the VRF association is not in right place. Option D misses the vrf member vrf-name configuration, which will not configure properly. In summary, option A provides the minimal, correct steps to enable the anycast gateway on a Cisco Nexus 9000, achieving the requirement of a common MAC and IP for gateway functionality across all leaf switches, while enabling per VRF isolation.

For further research, consider these resources:

Cisco Nexus 9000 Series VXLAN Configuration Guide: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/vxlan/configuration/guide/b_Cisco_Ne (Focus on anycast gateway section).

Cisco Data Center VXLAN Overview:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/index.html

Question: 64

An engineer implements an ACI fabric and must implement microsegmentation of endpoints within the same IP subnet using a network-based attribute. The attribute mapping must allow IP subnet independence. Which attribute must be selected?

A. IP
B. tag
C. custom
D. MAC address

Answer: D

Explanation:

The correct answer is **D. MAC address**. Here's why:

Microsegmentation aims to isolate workloads or endpoints at a granular level, often within the same IP subnet. While IP addresses can be used for this, they're typically bound to a subnet, making microsegmentation less flexible if you need IP subnet independence. Using IP addresses alone would tie your segmentation policies to the IP addressing scheme, reducing agility and requiring policy updates upon IP address or subnet changes.

MAC addresses, being hardware addresses, are unique to each endpoint and remain constant regardless of IP addressing. This makes them an ideal attribute for network-based microsegmentation when IP subnet independence is needed. You can define policies based on MAC addresses, ensuring that segmentation rules apply even if endpoints move between subnets or receive different IP addresses. This decouples your segmentation strategy from the IP addressing scheme, making it more adaptable and easier to manage.

ACI utilizes Endpoint Groups (EPGs) to group endpoints based on attributes. You can define policies between EPGs, and these policies are enforced based on the criteria defined for membership. If you were to use IP addresses, a change of IP would result in the device no longer falling under the policy. However, because a MAC address will not change as an address change, policies based on this will remain in effect and operate within multiple IP subnets. Using MAC addresses as an attribute provides a consistent way to identify and segment workloads, regardless of their IP configuration, which is crucial for microsegmentation that requires IP subnet independence. Tags, while useful for other types of segmentation, do not inherently provide the necessary granularity for isolating individual endpoints in this context. Custom attributes are not a native and universal characteristic.

Therefore, MAC addresses are the preferred attribute for microsegmentation in ACI when aiming for IP subnet independence due to their consistent, endpoint-specific nature.

Supporting Links:

Cisco Application Centric Infrastructure (ACI) Fundamentals:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/index.html (Refer to documentation on EPGs and contract definitions).

Cisco ACI Endpoint Groups:https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/getting-started/b_Getting_Started_Guide_2_x/b_Getting_Started_Guide_2_x_chapter_0100.html (Explore details about endpoint group membership.)

Question: 65

What is a characteristic of the install all command on the Cisco Nexus Series Switch?

A. automatically checks the image integrity
B. continues the upgrade process if any step in the sequence fails
C. upgrades only certain modules
D. impacts data plane traffic

Answer: A

Explanation:

The correct answer is **A. automatically checks the image integrity**. Here's a detailed justification:

The install all command on Cisco Nexus Series switches is designed to streamline the process of upgrading the entire system, including the operating system (NX-OS) and any other software components. A critical function of this command is to ensure the new software image is valid and not corrupted before initiating the upgrade.

This is achieved by a built-in mechanism that verifies the image's integrity using a checksum or other cryptographic method. This step is vital to prevent installing a flawed image that could render the switch unstable or inoperable. The other options are incorrect. Option B is wrong because the install all command, by default, will halt the upgrade process if any step encounters an issue. This prevents the switch from moving forward with a potentially incomplete or problematic upgrade. Option C is wrong as install all upgrades all modules and software components within the scope of a regular upgrade operation. It does not allow for selective module upgrades. Lastly, Option D is incorrect; the install all command does cause a brief impact to the data plane during the reboot process, since the system must briefly go offline to perform the upgrade and boot into the new image. Therefore, image integrity validation is a defining characteristic of the install all command.

For further research, refer to the official Cisco documentation on NX-OS Software Upgrade and Downgrade:

Cisco NX-OS Software Upgrade and Downgrade Guide
Cisco Nexus 9000 Series NX-OS Configuration Command Reference (Search for install all for specific command details)

Question: 66

DRAG DROP -
An engineer must configure the HSRP protocol to implement redundancy using two Cisco Nexus Series Switches, in addition, me HSRP must meet these requirements:
* switch1 must retain the primary role if switch2 goes offline.
* switch1 must retain the primary role until normal conditions are restored.
* switch1 and switch2 must ensure that the routing tables are converged before taking the active role.
* switch2 must retain the primary role if the default gateway is not reachable.

Drag and drop the configuration commands from the right to the left to meet the requirements. The commands are used more than once. Not all commands are used.

Select and Place:

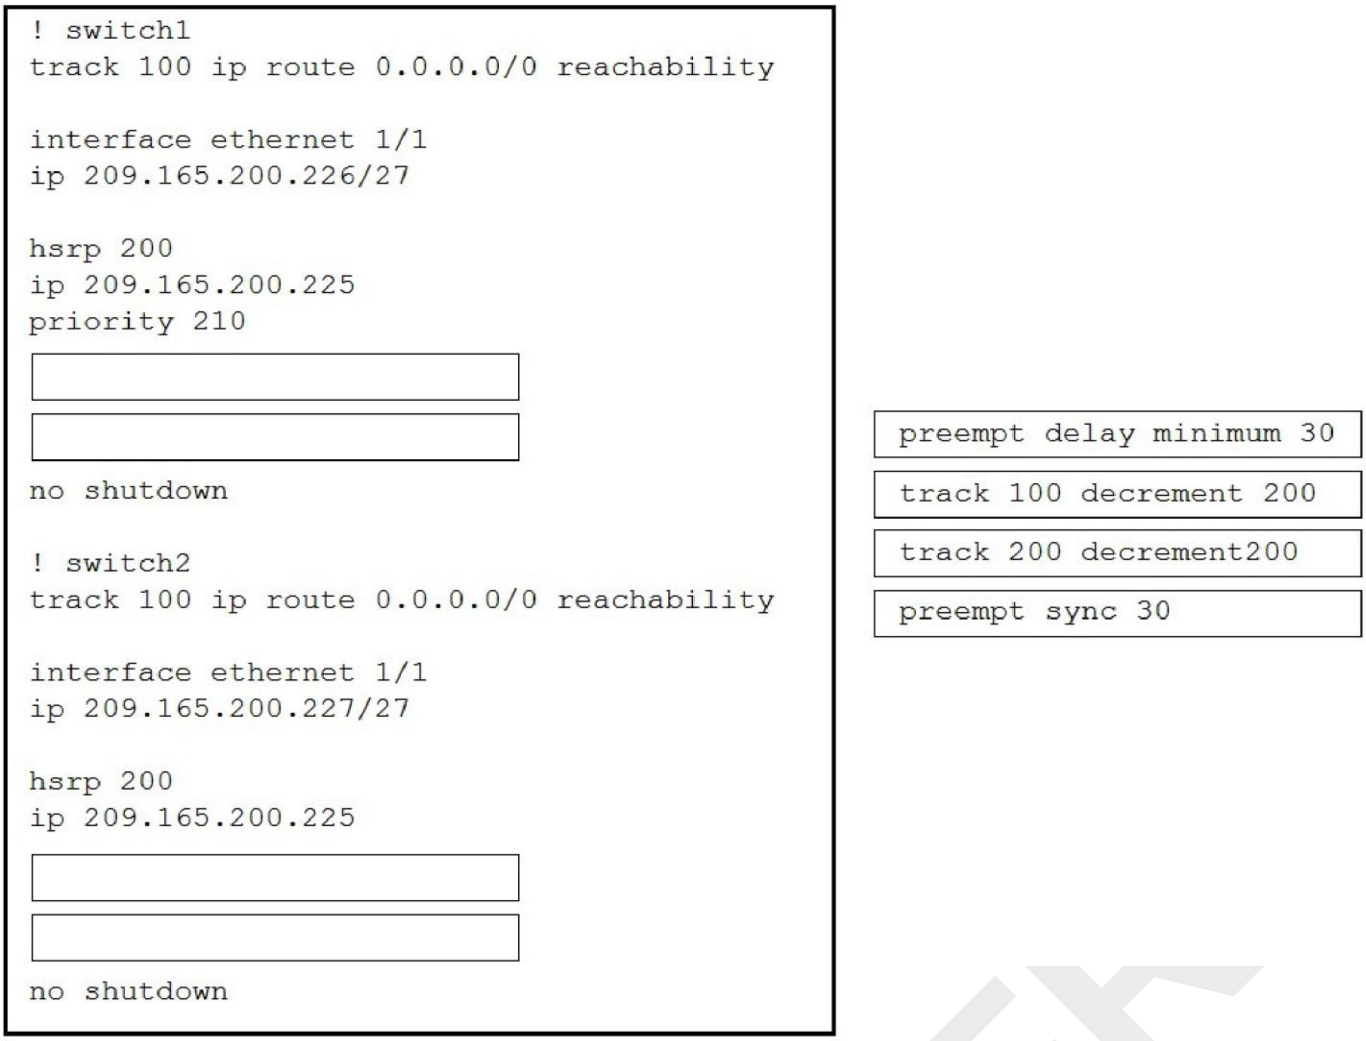

```
! switch1
track 100 ip route 0.0.0.0/0 reachability

interface ethernet 1/1
ip 209.165.200.226/27

hsrp 200
ip 209.165.200.225
priority 210
□
□
no shutdown

! switch2
track 100 ip route 0.0.0.0/0 reachability

interface ethernet 1/1
ip 209.165.200.227/27

hsrp 200
ip 209.165.200.225
□
□
no shutdown
```

```
preempt delay minimum 30
track 100 decrement 200
track 200 decrement200
preempt sync 30
```

Answer:

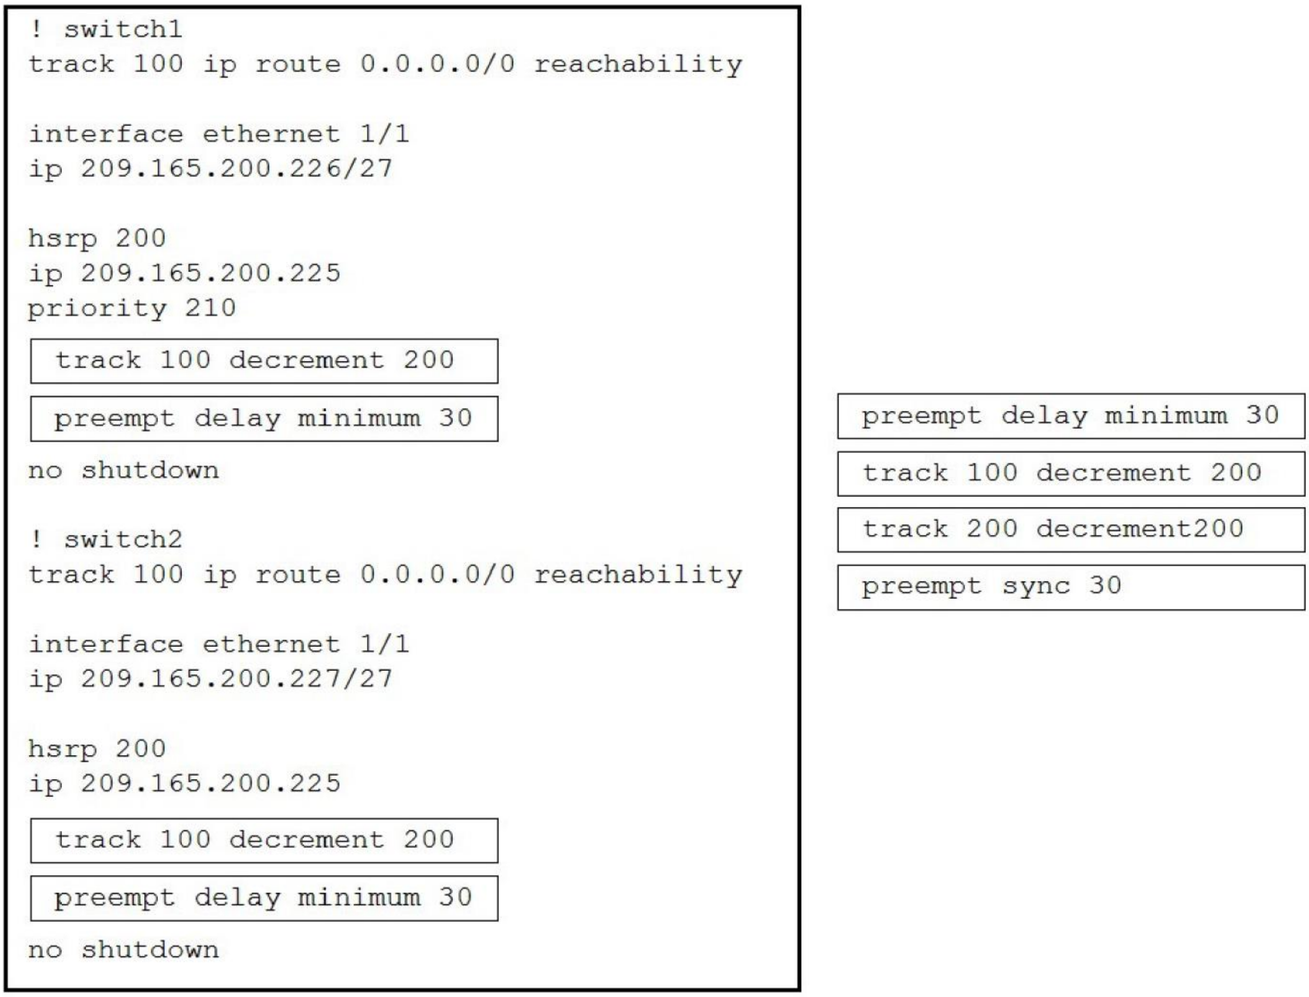

```
! switch1
track 100 ip route 0.0.0.0/0 reachability

interface ethernet 1/1
ip 209.165.200.226/27

hsrp 200
ip 209.165.200.225
priority 210
 track 100 decrement 200
 preempt delay minimum 30
no shutdown

! switch2
track 100 ip route 0.0.0.0/0 reachability

interface ethernet 1/1
ip 209.165.200.227/27

hsrp 200
ip 209.165.200.225
 track 100 decrement 200
 preempt delay minimum 30
no shutdown
```

```
preempt delay minimum 30
track 100 decrement 200
track 200 decrement200
preempt sync 30
```

Question: 67

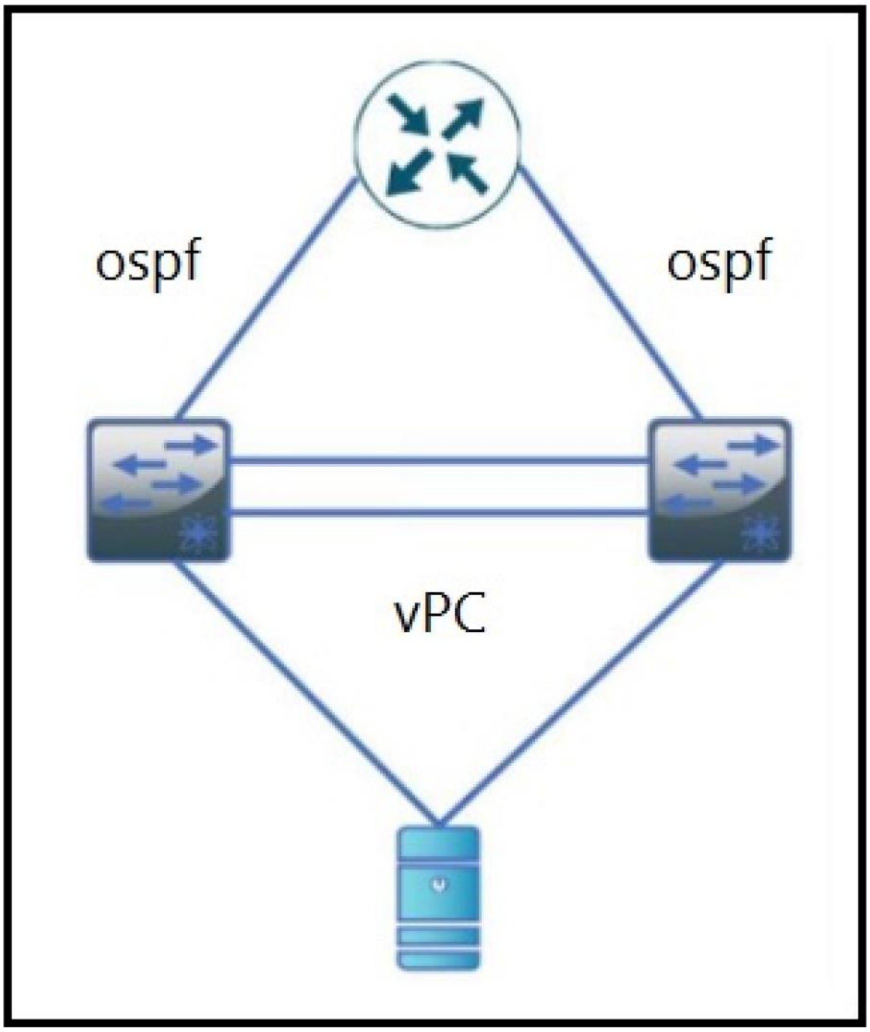

Refer to the exhibit. During a vPC peer switch reload, there is packet loss between the server and the router. Which action must be taken to prevent this behavior during future reloads?

A. Set the routed uplink ports of the Cisco Nexus peers as orphans.
B. Increase the vPC delay restore timer.
C. Disable vPC ARP synchronize on the vPC peers.
D. Decrease the OSPF hello and dead interval timers.

Answer: B

Explanation:
After a vPC peer device reloads and comes back up, the routing protocol needs time to reconverge. The recovering vPCs leg may black-hole routed traffic from access to aggregation/core until uplink Layer 3 connectivity is re-established. vPC Delay Restore feature delays vPCs leg bring-up on the recovering vPC peer device. vPC Delay Restore allows for Layer 3 routing protocols to converge before allowing any traffic on vPC leg. This results in a more graceful restoration and zero packet loss during the recovery phase (traffic still gets diverted on the alive vPC peer device). This feature is enabled by default with a vPC restoration default timer of 30 seconds. The timer can be tuned according to a specific Layer 3 convergence baseline from 1 to 3600 seconds.

Reference:
https://www.cisco.com/c/en/us/support/docs/ios-nx-os-software/nx-os-software/212589-understanding-vpc-election-process.html

Question: 68

A Cisco ACI engineer must configure an access port on a Cisco ACI leaf switch. The access policy should be configured so that it supports L3Out connectivity concurrently with several EPGs that are mapped to this interface with the static port method. How should the access policy be configured?

A. by linking the interface policy group to multiple Attachable Access Entity Profiles
B. with two interface policy groups that are linked to the same interface profile port selector
C. by mapping a single physical domain, mapped to the L3out and EPG interfaces
D. with a single Attachable Access Entity Profile that is linked to the corresponding routed domain and physical domain

Answer: D

Explanation:

The correct approach for configuring a Cisco ACI access port to support both L3Out connectivity and multiple EPGs using static port mapping involves utilizing a single Attachable Access Entity Profile (AAEP) linked to both a routed domain and a physical domain. An AAEP acts as a container for policies that govern how endpoints connect to the fabric. The physical domain is essential for EPG connectivity and signifies the VLAN or encapsulation used for those EPGs on that interface. The routed domain, conversely, is crucial for enabling L3Out connectivity, associating the interface with a Layer 3 network and routing capabilities.

Option D correctly captures this need. Linking a single AAEP to both types of domains allows the same physical interface to participate in both routed and bridged (EPG) contexts. An interface policy group, while used, doesn't define which connectivity is supported - it refers to the policy set related to the interface (speed, duplex). Option A suggests multiple AAEPs, and this is incorrect, as only one AAEP can be linked to an interface. Options B and C suggest the creation of more policies, while a simpler solution with one AAEP can be used. Using different interface groups does not solve the core problem of the routed domain, since the same interface port can belong to an interface profile.

Therefore, option D is the most efficient and straightforward method in Cisco ACI to fulfill the given requirement. The routed domain handles external routing policies, while the physical domain manages EPG assignments to this specific interface. This unified approach simplifies management and ensures consistent connectivity.

Authoritative Links for Further Research:

Cisco ACI Fundamentals:https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/white-paper-c11-737927.html (Specifically, section on access policies)
Cisco ACI Configuration Guide:https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/4x/configuration/l3-config/Cisco-APIC-Layer-3-Configuration-Guide-401/aci-config-l3out-401.html (Focus on L3Out configuration with access policies.)
Cisco ACI Policy Model:https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/white_papers/Cisco-ACI-Policy-Model.html (For a more in-depth understanding of policy structures and interaction)

Question: 69

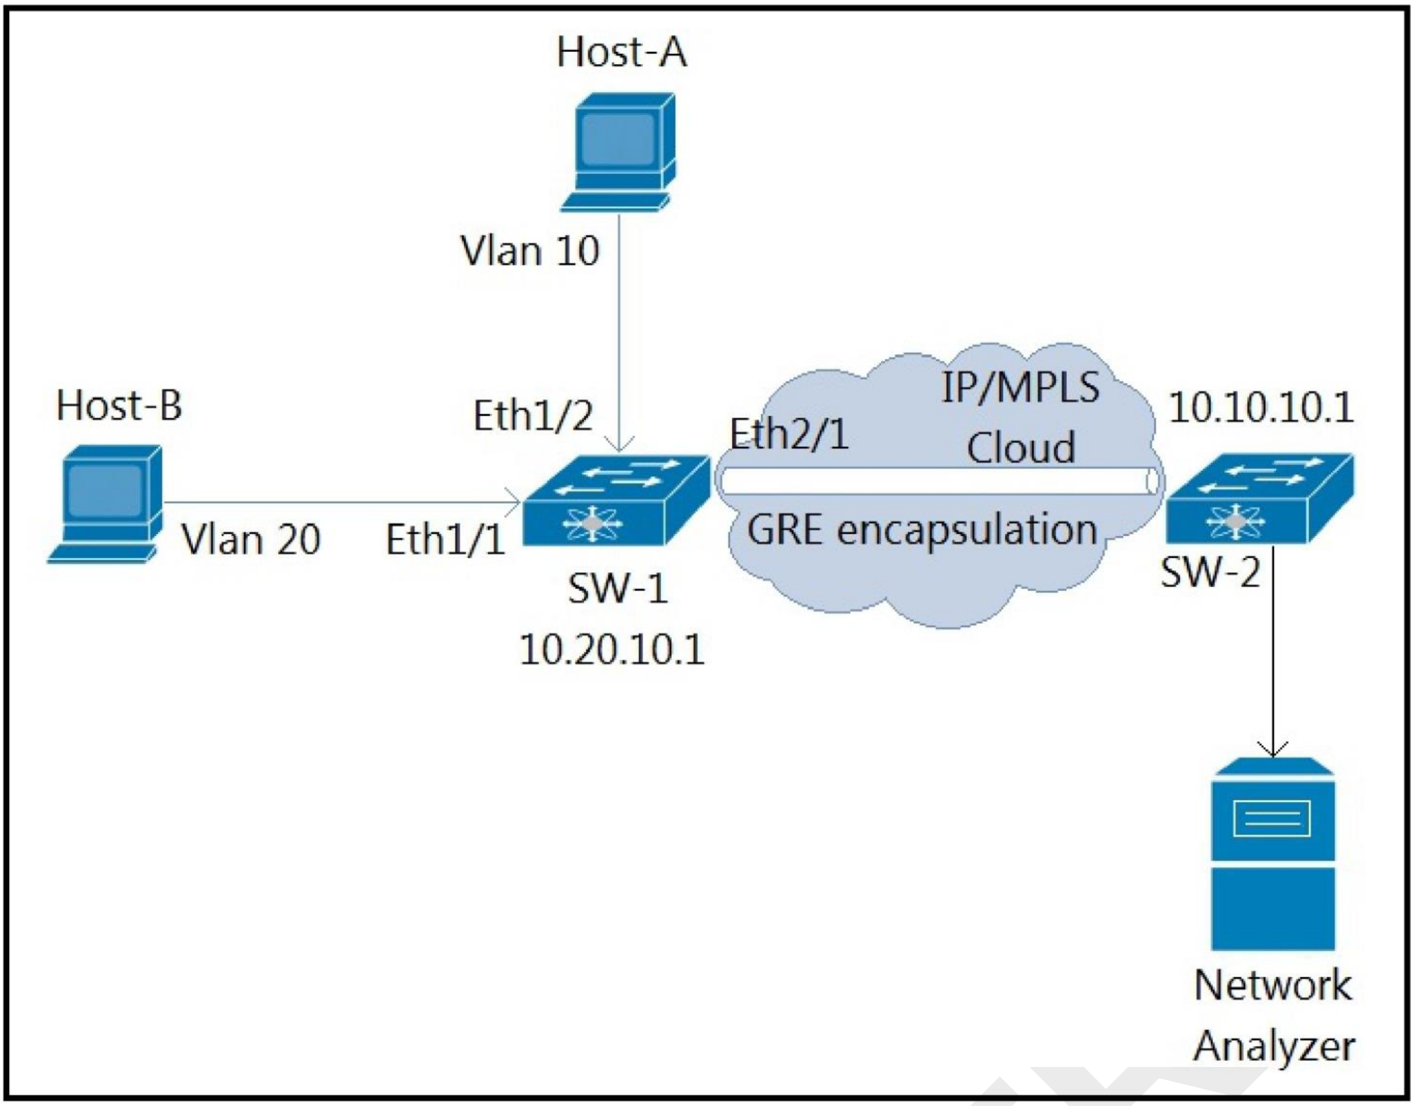

Refer to the exhibit. An engineer monitors ingress traffic from Host-A and all traffic for VLAN 20 from Host-B. Which configuration must be implemented to monitor and export traffic to Network Analyzer?

A. SW-1(config)# monitor session 5 SW-1(config-monitor)# source interface ethernet 1/2 rx SW-1(config-monitor)# source interface ethernet 1/1 tx SW-1(config-monitor)# source vlan 20 both SW-1(config-monitor)# destination ip 10.10.10.1 SW-1(config-monitor)# no shut

B. SW-1(config)# monitor session 5 type erspan-source SW-1(config-erspan-src)# source interface ethernet 1/2 rx SW-1(config-erspan-src)# source vlan 20 both SW-1(config-erspan-src)# destination ip 10.10.10.1 SW-1(config-erspan-src)# erspan-id 111 SW-1(config-erspan-src)# no shut

C. SW-1(config)# monitor session 5 SW-1(config-monitor)# source interface ethernet 1/2 rx SW-1(config-monitor)# source interface ethernet 1/1 both SW-1(config-monitor)# source vlan 20 both SW-1(config-monitor)# destination interface ethernet 2/1 SW-1(config-monitor)# no shut

D. SW-1(config)# monitor session 5 type erspan SW-1(config-erspan-src)# source interface ethernet 1/2 rx SW-1(config-erspan-src)# source interface ethernet 1/1 both SW-1(config-erspan-src)# destination ip 10.10.10.1 SW-1(config-erspan-src)# mtu 1000 SW-1(config-erspan-src)# no shut

Answer: B

Explanation:

B is correct

SW-1(config)# monitor session 5 type erspan-source

SW-1(config-erspan-src)# source interface ethernet 1/2 rx

SW-1(config-erspan-src)# source vlan 20 both

SW-1(config-erspan-src)# destination ip 10.10.10.1

SW-1(config-erspan-src)# erspan-id 111

SW-1(config-erspan-src)# no shut

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5600/sw/system_management/7x/b_5600_System_Mgmt_Config_7x/ b_6k_System_Mgmt_Config_7x_chapter_01110.html

Question: 70

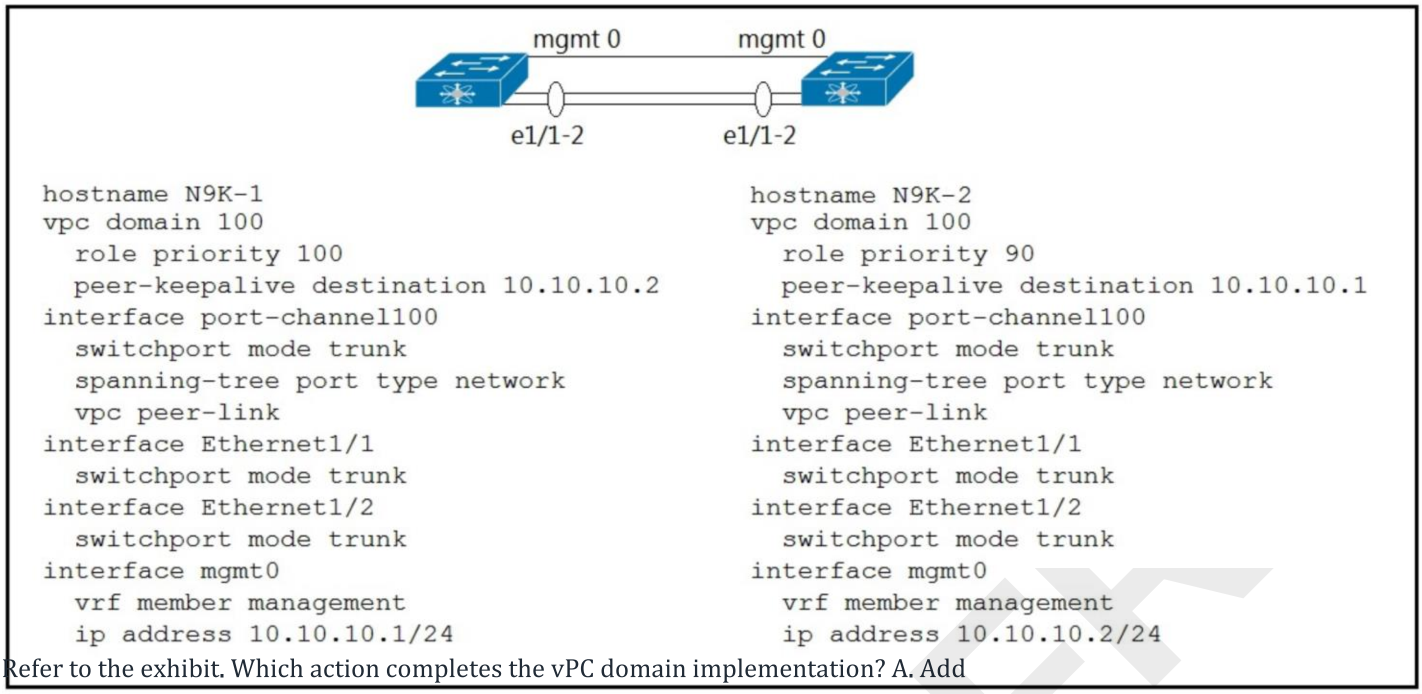

```
hostname N9K-1                                    hostname N9K-2
vpc domain 100                                    vpc domain 100
  role priority 100                                 role priority 90
  peer-keepalive destination 10.10.10.2             peer-keepalive destination 10.10.10.1
interface port-channel100                         interface port-channel100
  switchport mode trunk                             switchport mode trunk
  spanning-tree port type network                   spanning-tree port type network
  vpc peer-link                                     vpc peer-link
interface Ethernet1/1                             interface Ethernet1/1
  switchport mode trunk                             switchport mode trunk
interface Ethernet1/2                             interface Ethernet1/2
  switchport mode trunk                             switchport mode trunk
interface mgmt0                                   interface mgmt0
  vrf member management                             vrf member management
  ip address 10.10.10.1/24                          ip address 10.10.10.2/24
```

Refer to the exhibit. Which action completes the vPC domain implementation? A. Add the vPC member ports to the vPC channel group.

B. Configure the system MAC on the vPC domain.

C. Include the VRF management on the vPC domain.

D. Allow VLANs on the vPC peer link member interfaces.

Answer: A

Explanation:

Technically both A & C are right. But would opt for A as it's clearly missing memebrs.

A is correct

Question: 71

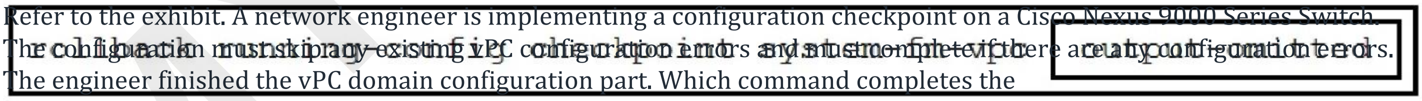

```
rollback running-config checkpoint system-fm-vpc [illegible]
```

Refer to the exhibit. A network engineer is implementing a configuration checkpoint on a Cisco Nexus 9000 Series Switch. The configuration must skip any existing vPC configuration errors and must complete if there are any configuration errors. The engineer finished the vPC domain configuration part. Which command completes the checkpoint implementation?

A. best-effort
B. stop-at-first-failure
C. atomic
D. verbose

Answer: A

Explanation:
The available options are:
1. atomic(default) Rollback only if there are no errors
2. Best-Effort: Rollback and skip any errors
3. Stop at First Failure : Rollback and stop if an error occurs

Reference:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/system_management/configuration/guide/ b_Cisco_Nexus_9000_Series_NX-OS_System_Management_Configuration_Guide_7x/b_Cisco_Nexus_9000_Series_NX-OS_System_Management_Configuration_Guide_7x_chapter_01000.pdf

Question: 72

An engineer must implement a VXLAN-based data center interconnect. The long-distance transport provided by a service provider is IP-based, supports a maximum MTU of 1554 bytes, and does not support outbound traffic replication. Which action must be taken to build the data center interconnect?

A. Announce host reachability over BGP.
B. Create an IP access list and associate it with VNI to replicate traffic to remote VTEPs.
C. Configure a route map to associate the IPs of the remote VTEPs.
D. Implement BGP EVPN ingress replication.

Answer: D

Explanation:

The correct answer is **D. Implement BGP EVPN ingress replication.** Here's why:

VXLAN (Virtual Extensible LAN) is used for overlay networking to extend Layer 2 networks over Layer 3 infrastructure, essential for data center interconnects (DCIs). In this scenario, the service provider's transport network has a limited MTU (1554 bytes), which means we must account for the VXLAN overhead. Also, it doesn't support outbound replication. Traditional multicast replication is not an option.

BGP EVPN (Border Gateway Protocol Ethernet VPN) provides a control plane for VXLAN, facilitating MAC address learning and distribution between VTEPs (VXLAN Tunnel Endpoints). Ingress replication addresses the limitation of multicast replication by having the ingress VTEP duplicate the VXLAN-encapsulated traffic for each remote VTEP.

Why other options are incorrect:

A. Announce host reachability over BGP: While BGP is crucial for control plane signaling in EVPN, just announcing host reachability doesn't handle the replication requirement.

B. Create an IP access list and associate it with VNI to replicate traffic to remote VTEPs: Access lists don't perform replication. They filter traffic.

C. Configure a route map to associate the IPs of the remote VTEPs: Route maps manipulate routing attributes. They aren't used to replicate traffic.

BGP EVPN ingress replication is the only method that fits the requirements by handling both the lack of multicast support from the service provider and allowing for the distribution of Layer 2 traffic across the Layer 3 transport. It efficiently sends a copy of the frame to each remote VTEP directly from the source VTEP, overcoming limitations of the transport network. This makes it the suitable approach.

Authoritative Links:

Cisco BGP EVPN Ingress Replication:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l2vpn/configuration/16-12/mp-l2vpn-16-12-book/mp-l2vpn-evpn-ovly.html#GUID-E1A7C088-C86E-47A5-B772-6E16F118874A
Understanding VXLAN and BGP EVPN:https://networkdirection.net/articles/networking/evpn-and-vxlan-the-ultimate-guide/

Question: 73

An engineer must perform an initial configuration of VXLAN Tunnel End-Point functionality on the Cisco Nexus 9000 Series platform. All the necessary features are already enabled on the switch. Which configuration set must be used to accomplish this goal?

A. NX9K(config)# interface loopback0 NX9K(config-if)# ip address 10.10.10.11/32 NX9K(config)# interface VLAN10 NX9K(config-if)# ip unnumbered loopback0 NX9K(config)# interface nve1 NX9K(contig-if-nve)# source-interface loopback0

B. NX9K(config)# interface loopback0 NX9K(config-if)# ip address 10.10.10.11/32 NX9K(config)# interface overlay0 NX9K(config-if-overlay)# source-interface loopback0 NX9K(config)# interface VLAN10 NX9K(contig-if-nve)# source-interface loopback0

C. NX9K(config)# interface loopback0 NX9K(config-if)# ip address 10.10.10.11/32 NX9K(config)# interface tunnel1 NX9K(config-if)# tunnel source loopback0 NX9K(config)# interface ethernet1/1 NX9K(contig-if)# ip unnumbered loopback0

D. NX9K(config)# interface loopback0 NX9K(config-if)# ip address 10.10.10.11/32 NX9K(config)# interface nve1 NX9K(config-if-nve)# source-interface loopback0 NX9K(config)# interface ethernet1/1 NX9K(contig-if)# ip unnumbered loopback0

Answer: D

Explanation:

The correct answer is D because it accurately reflects the essential steps for configuring a VXLAN Tunnel Endpoint (VTEP) on a Cisco Nexus 9000 switch. First, a loopback interface (loopback0) is created and assigned an IP address (10.10.10.11/32). This loopback interface serves as the stable source address for the VTEP. Next, the Network Virtualization Endpoint (NVE) interface (nve1) is created and configured to use the loopback0 interface as its source. This binds the VTEP to the assigned IP address, enabling it to participate in VXLAN tunnels. The configuration also shows a basic Ethernet interface (ethernet1/1) and assigns the loopback IP to it. Although this interface doesn't need to be configured with a loopback IP in this case, the IP Unnumbered command is commonly used to reduce the overhead of configuring IPs on each interface. Option A fails because it uses a VLAN interface with an unnumbered configuration. Option B has an incorrect interface type with the overlay0. Similarly, Option C is incorrect as it uses a tunnel interface and incorrectly refers to the tunnel source. The use of loopback interfaces for VTEP source addresses is a standard practice for establishing stable VTEP endpoints, as they are logical interfaces unaffected by physical interface failures. For further reading, refer to Cisco's documentation on VXLAN configuration on Nexus 9000 platforms: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/vxlan/configuration/guide/b-cisco-nexus-9000-vxlan-config-guide/b-cisco-nexus-9000-vxlan-config-guide_chapter_010.html and https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/93x/configuration/guides/interface/b-cisco-nexus-9000-interface-config-93x/configuring-nve-interfaces.html.

Question: 74

An engineer must perform a software upgrade on a production Cisco Nexus 7000 Series Switch. Before the upgrade activity, the requirement is for all ports to be shut down and routing protocols to terminate gracefully. Which feature must be used to meet these requirements?

A. Switch Profile
B. Service Profile Template
C. Maintenance Mode Profile
D. Configuration Profile

Answer: C

Explanation:

The correct answer is **C. Maintenance Mode Profile**. Maintenance Mode Profile is specifically designed for planned downtime activities, like software upgrades, on Cisco Nexus devices. It provides a controlled way to gracefully shut down ports and terminate routing protocols before maintenance. When enabled, a maintenance mode profile will shut down all interfaces on the device and gracefully shut down all routing protocol processes. This ensures that no traffic is dropped abruptly during the maintenance period. Other options don't fit the requirements; Switch Profiles (A) are used for configuration consistency across switches, Service Profile Templates (B) are used for server profiles and Configuration Profiles (D) are more focused on baseline configuration management. While they might allow configuration changes, they do not provide a dedicated and controlled mechanism for shutting down ports and routing protocols in a maintenance scenario.

Maintenance mode uses timers to allow protocols to gracefully terminate adjacency. For instance, the maintenance mode can be configured to use the BGP graceful shutdown timer which allows an adequate time for a BGP neighbor to learn and adjust to the shutdown event. It avoids hard-drop of packets during upgrades and ensures that after upgrade, the protocols resume their operations gracefully. By using Maintenance Mode profile, the engineer can ensure minimal disruption to the network. In short, Maintenance Mode Profile is the intended mechanism for preparatory steps prior to maintenance activities, which aligns perfectly with the upgrade requirement, thus the correct answer.

Further Reading:

Cisco Nexus 7000 Series NX-OS System Management Configuration Guide:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus7000/sw/system_management/config/guide/b_OS_System_Management_Configuration_Guide/b_Cisco_Nexus_7000_Series_NX-OS_System_Management_Configuration_Guide_chapter_011.html - Look for information on 'Maintenance Mode'.

Question: 75

DRAG DROP -

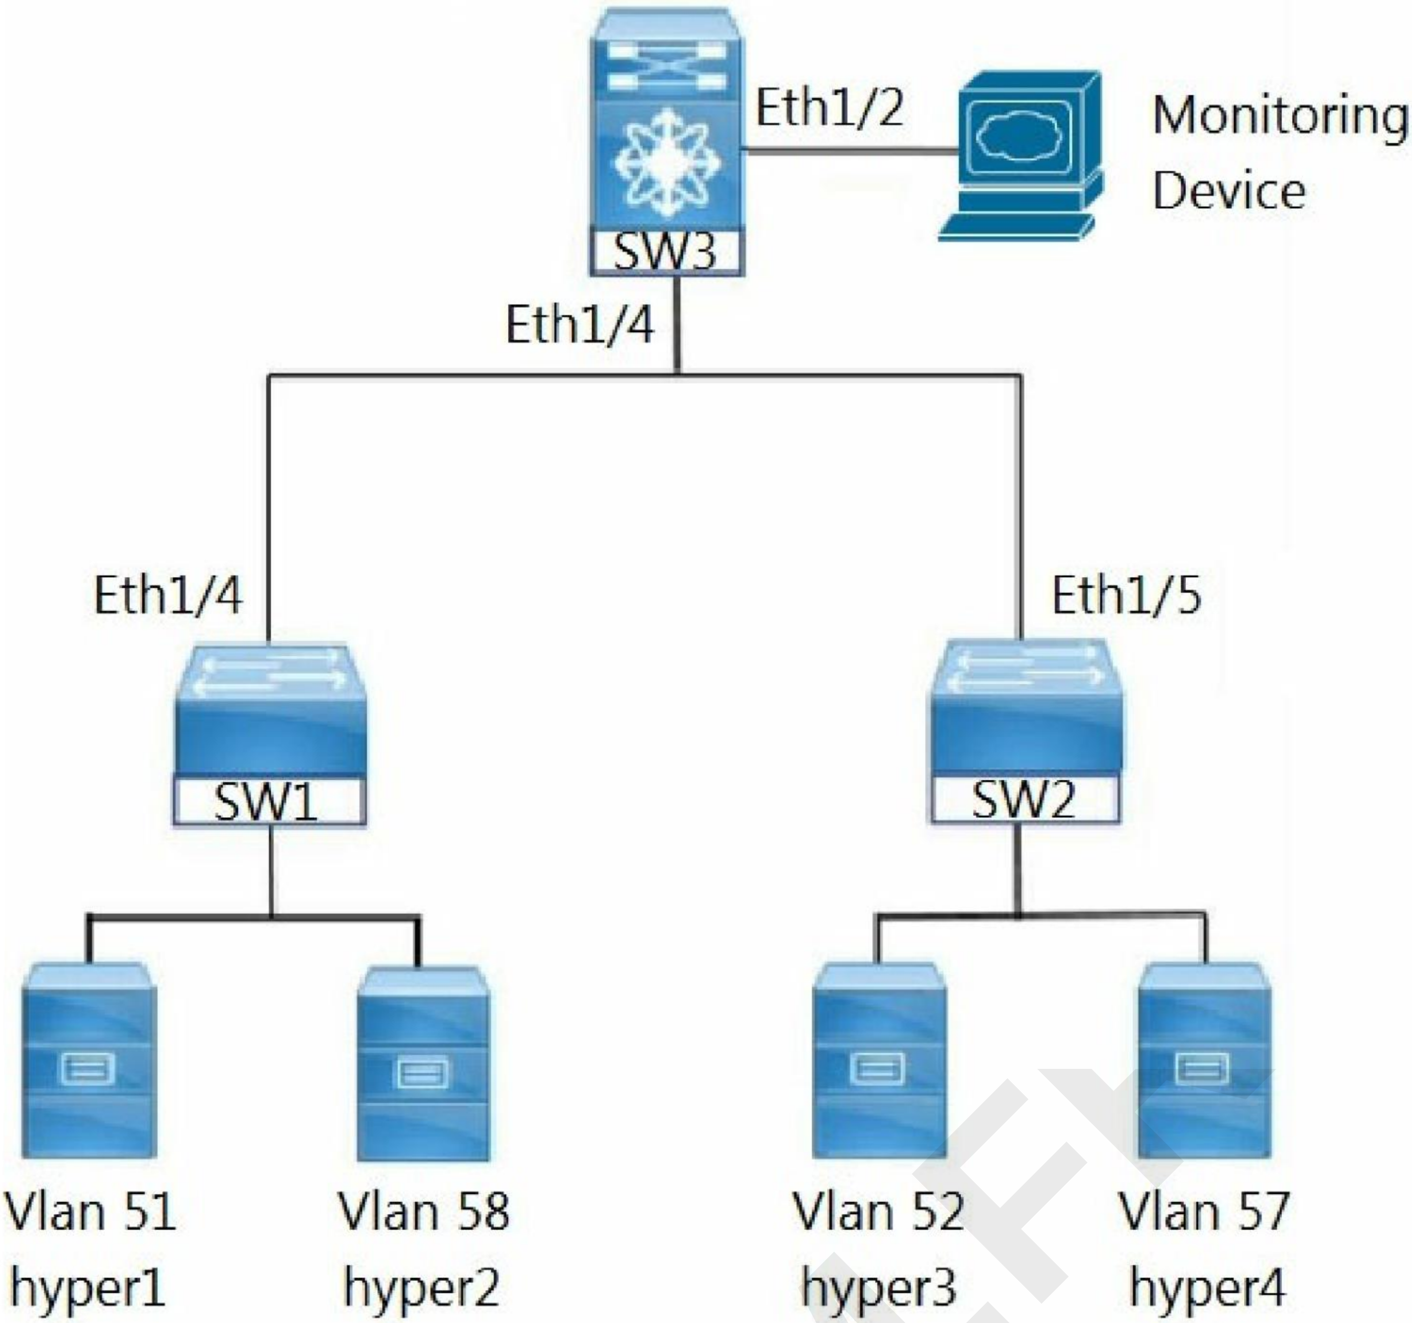

Refer to the exhibit. An engineer must monitor the Ethernet port and the corresponding VLAN traffic for the hyper4 device. The SW3 device is a Cisco Nexus 7000
Series Switch. Drag and drop the code snippets from the right into the boxes in the configuration to meet these requirements.
Select and Place:

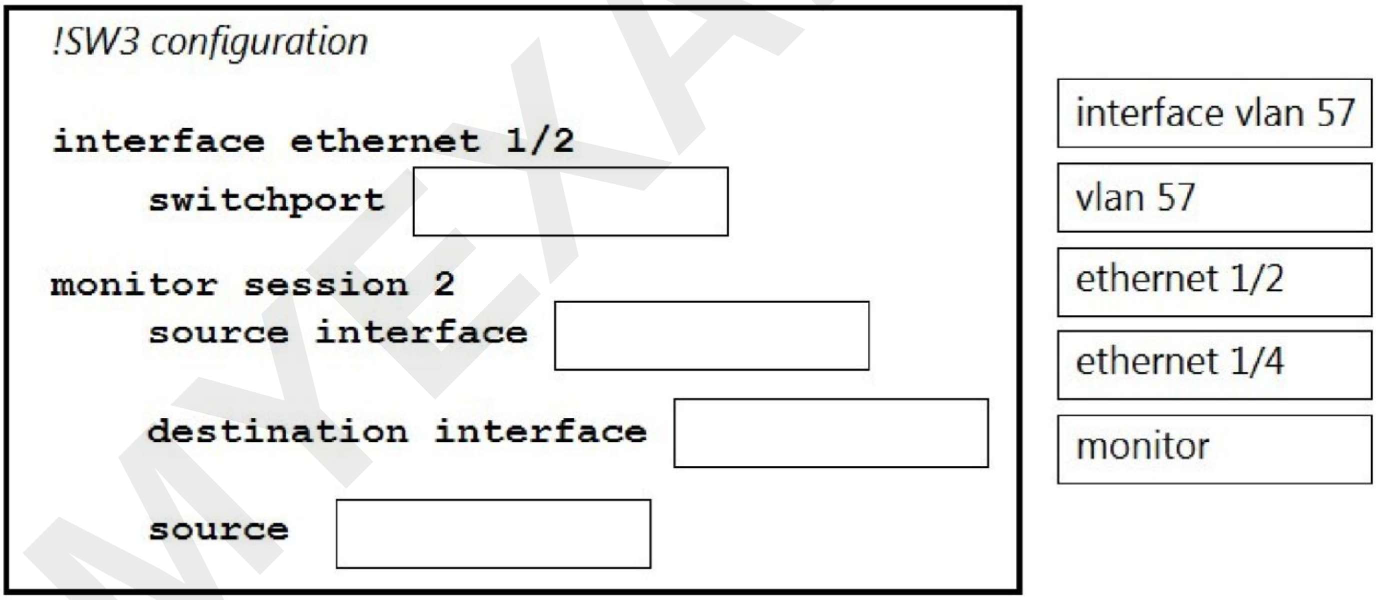

!SW3 configuration

```
interface ethernet 1/2
    switchport □
monitor session 2
    source interface □

    destination interface □

    source □
```

- interface vlan 57
- vlan 57
- ethernet 1/2
- ethernet 1/4
- monitor

Answer:

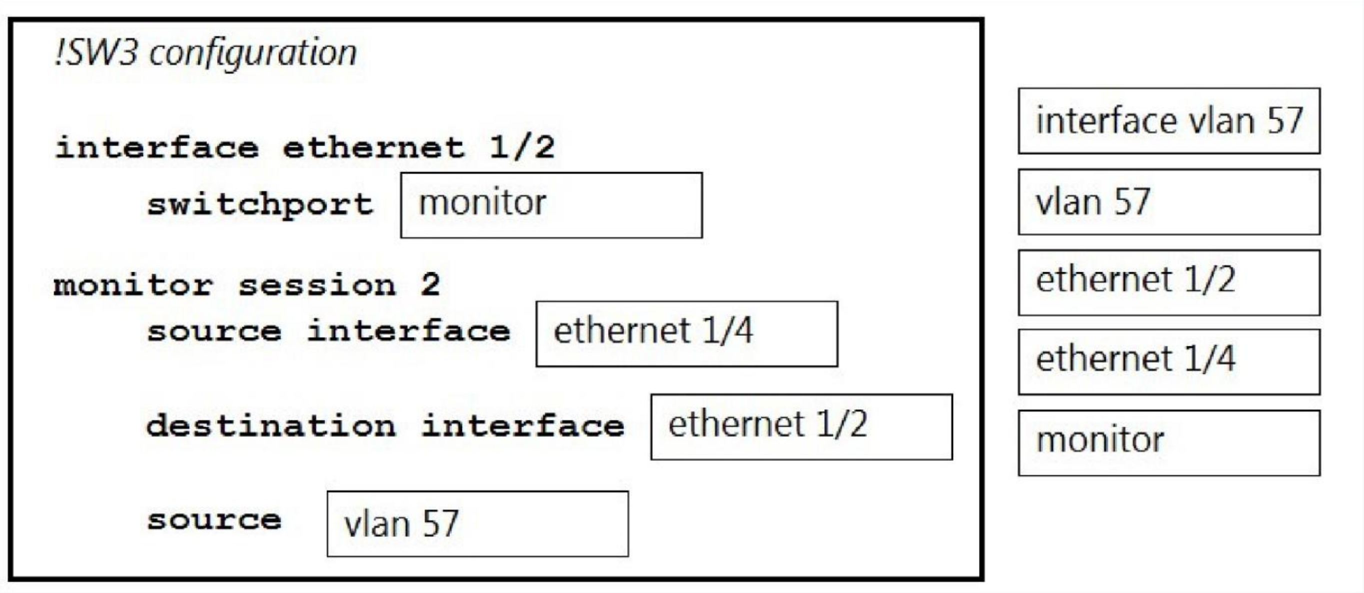

```
!SW3 configuration

interface ethernet 1/2
    switchport monitor

monitor session 2
    source interface ethernet 1/4

    destination interface ethernet 1/2

    source vlan 57
```

interface vlan 57

vlan 57

ethernet 1/2

ethernet 1/4

monitor

Question: 76

DRAG DROP -
Drag and drop the network assurance concepts from the left onto the corresponding benefits on the right.
Select and Place:

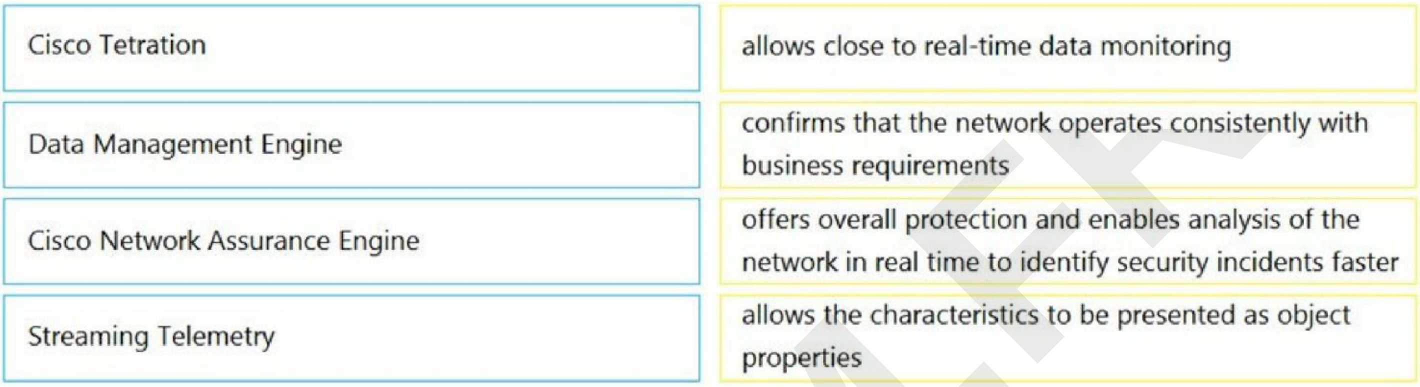

Cisco Tetration	allows close to real-time data monitoring
Data Management Engine	confirms that the network operates consistently with business requirements
Cisco Network Assurance Engine	offers overall protection and enables analysis of the network in real time to identify security incidents faster
Streaming Telemetry	allows the characteristics to be presented as object properties

Answer:

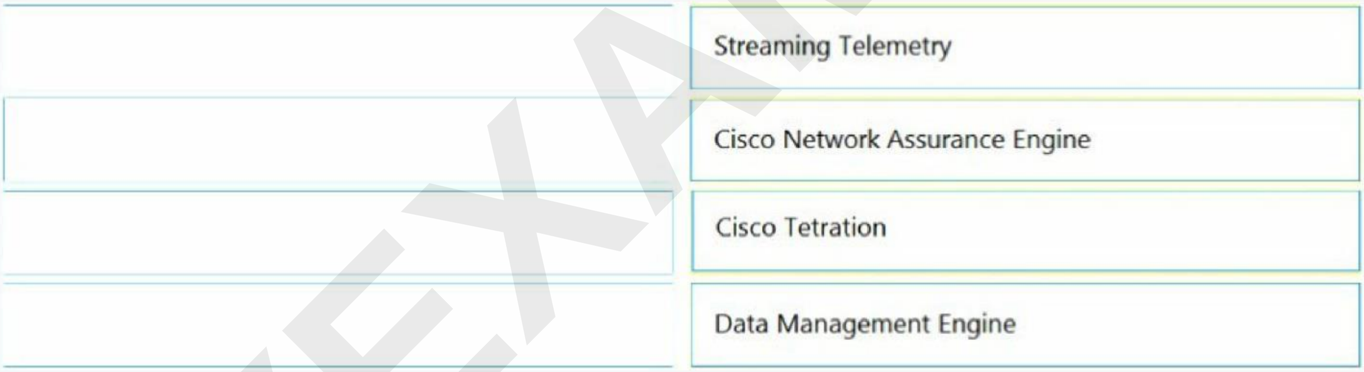

	Streaming Telemetry
	Cisco Network Assurance Engine
	Cisco Tetration
	Data Management Engine

Question: 77

DRAG DROP -
An engineer must configure HSPR protocol on two Cisco Nexus 9000 Series Switches running a virtual port channel.
In addition, the HSRP implementation must meet these requirements:
➾ It must allow more than 500 groups.
➾ switch1 must act as the primary switch.

✑ Both switches must use a user-defined hardware address.
Drag and drop the commands from the right to complete a configuration of the HSRP on the left. The commands are used more than once. Not all commands are used.
Select and Place:

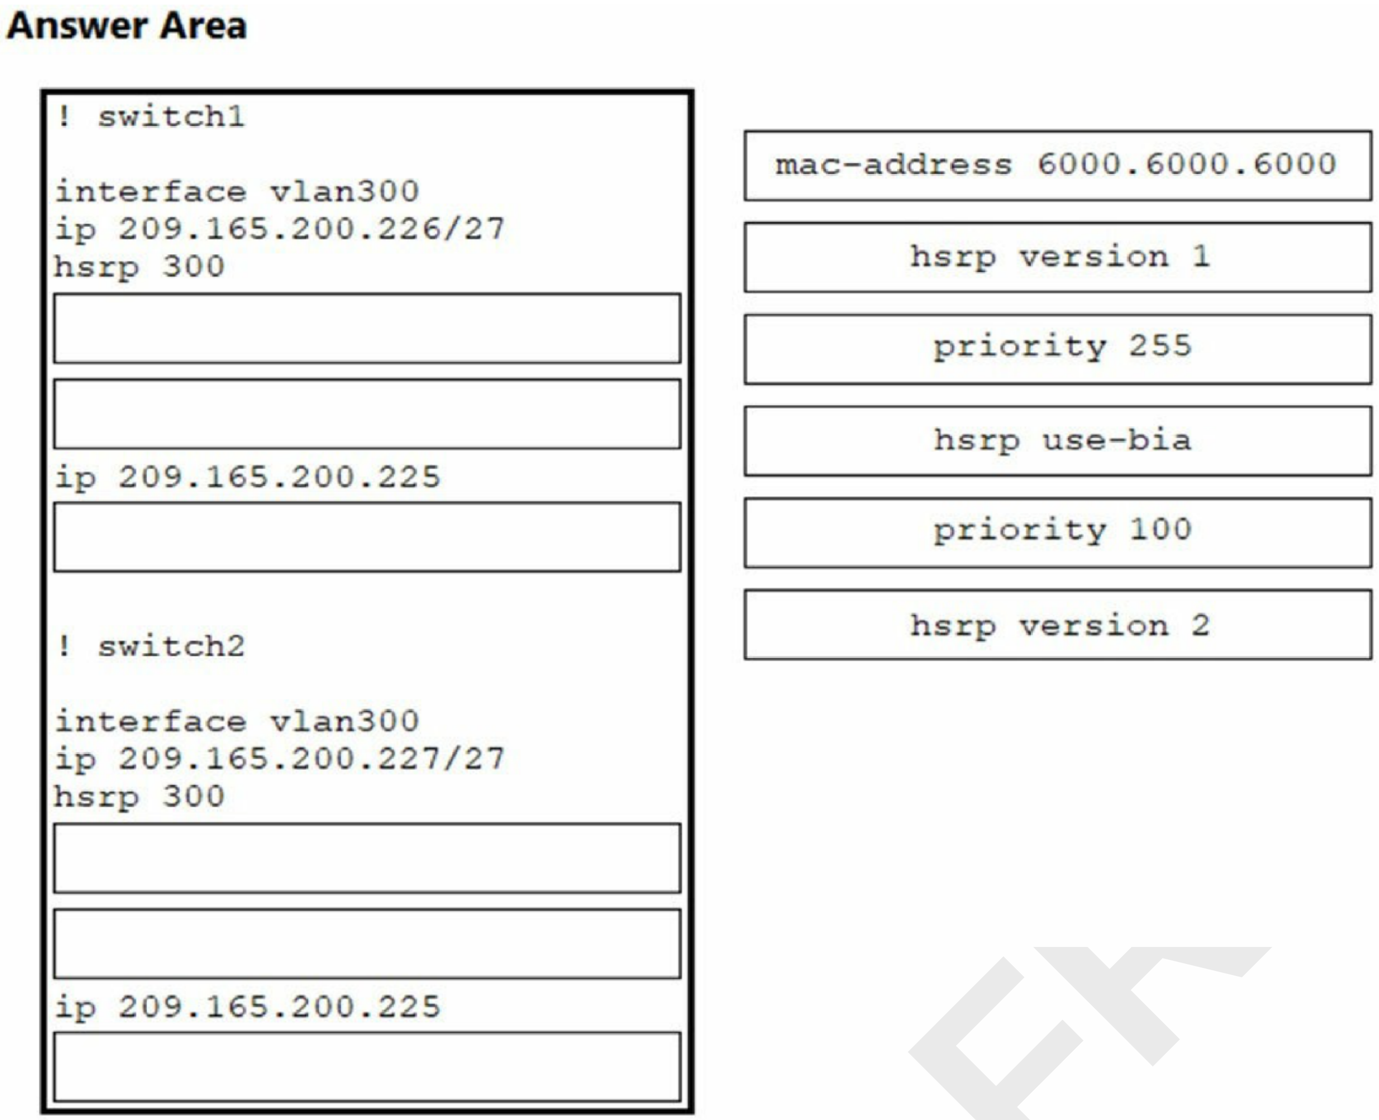

Answer Area

```
! switch1

interface vlan300
ip 209.165.200.226/27
hsrp 300


ip 209.165.200.225


! switch2

interface vlan300
ip 209.165.200.227/27
hsrp 300


ip 209.165.200.225
```

```
mac-address 6000.6000.6000
hsrp version 1
priority 255
hsrp use-bia
priority 100
hsrp version 2
```

Answer:

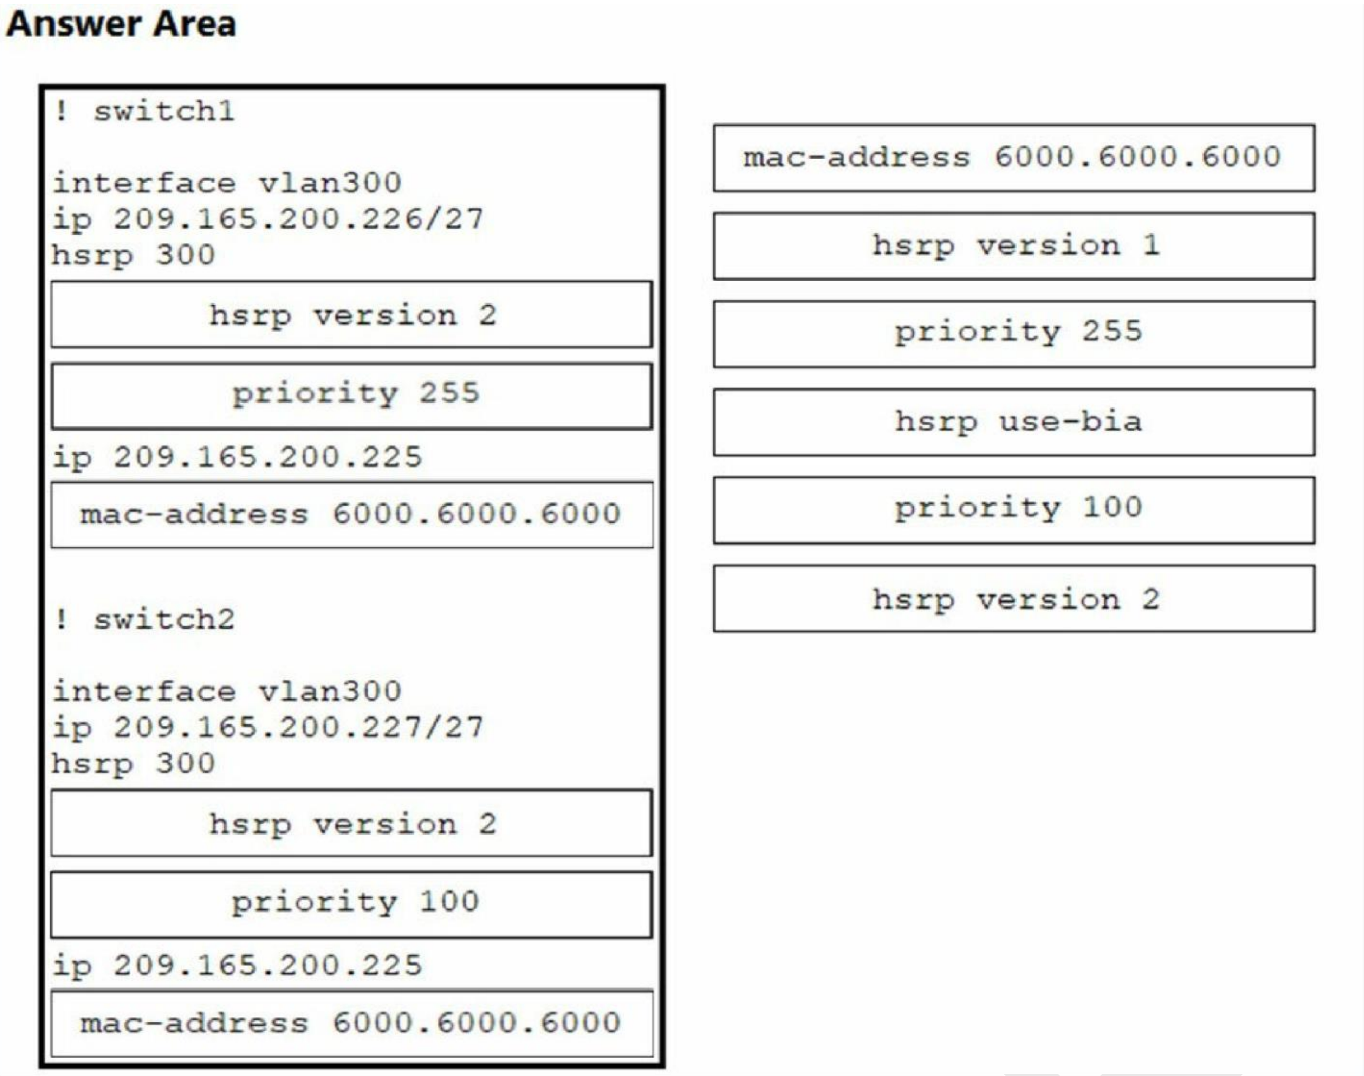

Answer Area

```
! switch1

interface vlan300
ip 209.165.200.226/27
hsrp 300
        hsrp version 2
        priority 255
ip 209.165.200.225
  mac-address 6000.6000.6000

! switch2

interface vlan300
ip 209.165.200.227/27
hsrp 300
        hsrp version 2
        priority 100
ip 209.165.200.225
  mac-address 6000.6000.6000
```

```
mac-address 6000.6000.6000
hsrp version 1
priority 255
hsrp use-bia
priority 100
hsrp version 2
```

Explanation:

Reference:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/6-x/unicast/configuration/guide/l3_cli_nxos/l3_hsrp.html

Question: 78

An engineer must configure OSPFv2 connectivity between a pair of Cisco Nexus switches. The connection between the switches must meet these requirements:
✑ Use unicast for updates.
✑ Use decentralized communication of updates.
✑ Full adjacency between switches.
Which configuration is needed to meet these requirements?

A. filter list between the neighbors
B. stub area between the neighbors
C. virtual link between the neighbors
D. interface with a type of p2p network

Answer: C

Explanation:

The correct answer is **C. virtual link between the neighbors**. Here's why:

OSPFv2's primary purpose is to distribute routing information within an autonomous system. It uses different network types and structures to achieve this.

Why C is Correct:

Unicast Updates: While OSPF typically uses multicast for updates on broadcast and NBMA networks, it defaults to unicast on point-to-point links. Virtual links are specifically designed for connectivity over non-backbone areas and operate by establishing a logical point-to-point connection between two area border routers using unicast for communication.

Decentralized Communication: Virtual links facilitate decentralized update propagation by bypassing area boundaries. Rather than relying on a designated router for route aggregation, each area border router participating in the virtual link forwards routing updates directly to the adjacent router in another area. This avoids a single point of failure or route manipulation by area border routers.

Full Adjacency: Virtual links allow for direct, full adjacency establishment between area border routers that might not otherwise have direct physical connectivity within the backbone. This ensures complete routing information exchange and avoids complexities related to multiple paths through different areas.

Why other options are incorrect:

A. Filter list between the neighbors: Filter lists primarily control which routes are advertised or received. While this can limit updates, it does not directly impact the unicast, decentralized, or adjacency requirements specified in the question. Filtering can even block full adjacency.

B. Stub area between the neighbors: A stub area restricts routing information exchange, it only allows default route propagation from Area Border Router, This requirement is the exact opposite of what the engineer needed (full adjacency). Also it does not use unicast communication only.

D. Interface with a type of p2p network: Directly configuring the interface as p2p does configure unicast traffic but does not accomplish decentralized communication of updates or the full adjacency goal that the virtual link is able to accomplish, The Virtual Link is the best solution as its purpose is the exact use case the question described.

In summary: A virtual link provides unicast communication between area border routers, decentralizes the routing update process by directly propagating updates across area boundaries, and facilitates full adjacency over non-contiguous areas. This directly addresses all requirements laid out by the problem.

Authoritative Links for Further Research:

Cisco's OSPF documentation:https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html - Explains various OSPF topics, including virtual links and their use cases.

OSPF Virtual Links:https://www.firewall.cx/networking-topics/routing/ospf-routing-protocol/1127-ospf-virtual-links.html - A detailed explanation of OSPF Virtual Links.

Question: 79

An infrastructure architect is analyzing the deployment type for an application. Several companies must be able to operate the application environment and integrate it with a third-party protocol. The development team also must have maximum control over their development environment, including the operating system used, Python interpreter, and corresponding libraries. Which cloud deployment solution meets these requirements?

A. private cloud Platform as a Service

B. public cloud Storage as a Service

C. hybrid cloud Software as a Service

D. community cloud Infrastructure as a Service

Answer: C

Explanation:

The correct answer is **C. hybrid cloud Infrastructure as a Service**. Here's why:

The scenario requires a solution that balances control with external access and integration. A **community cloud** (Option D) is designed for shared use among organizations with similar interests, which doesn't align with "several companies" operating the environment and needing third-party protocol integration. While **private cloud PaaS** (Option A) offers more control than public PaaS, it doesn't inherently facilitate integration with diverse external entities like the requirement does, it's typically focused for internal use. **Public cloud SaaS** (Option B) primarily offers pre-built applications, not the required granular control over development environment.

Hybrid cloud combines the benefits of public and private cloud models. Infrastructure as a Service (IaaS) provides the necessary granular control, allowing the development team to manage operating systems, Python interpreters, and libraries as needed. The integration requirements and external company access are managed by the public cloud components, providing connectivity and third party API's. This hybrid IaaS model enables companies to each operate independently while still integrating with third party systems.

In summary, hybrid cloud IaaS best fits the scenario as it grants development teams maximum control over their environment while providing the necessary access for multiple companies and third-party integration.

Authoritative Links:

NIST Definition of Cloud Computing:https://csrc.nist.gov/publications/detail/sp/800-145/final
Cloud Deployment Models:https://azure.microsoft.com/en-us/resources/cloud-computing-dictionary/what-are-the-different-cloud-deployment-models/#:~:text=There%20are%20four%20main%20types,cloud%2C%20and%20hybrid%20cloud. **Cloud Service Models:**https://azure.microsoft.com/en-us/resources/cloud-computing-dictionary/what-are-the-different-cloud-service-models/

Question: 80

An engineer is performing an ISSU upgrade on the Cisco MDS 9000 Series Switch. What is the impact on the control plane of the switch?

A. It is down for less than 80 seconds.
B. It remains down for more than 80 seconds but less than 140 seconds.
C. It remains down until the next reboot.
D. It remains up throughout the update.

Answer: A

Explanation:

The correct answer is A. An In-Service Software Upgrade (ISSU) on a Cisco MDS 9000 Series switch is designed to minimize downtime and disruption to the network. During an ISSU, the control plane of the switch experiences a brief interruption, typically less than 80 seconds. This controlled interruption is necessary as the switch migrates from the old software image to the new one. The process involves loading the new software onto a redundant supervisor module, switching over to that module, and then upgrading the original module. This minimizes impact because data plane forwarding continues using hardware resources on the module not being upgraded. The control plane is what handles protocols, routing, and configuration management, and thus is briefly offline to allow the software transition. While the control plane is briefly down, the data plane continues to forward traffic with minimal interruption. After the control plane is back, routing protocols and other control functions resume operation. Option B is incorrect because the control plane interruption during a successful ISSU should be shorter than that. Option C is incorrect because ISSU specifically aims to avoid a full reboot. Option D is incorrect because the control plane is briefly impacted.

Authoritative links for further research:

Cisco MDS 9000 Series NX-OS Software Upgrade and Downgrade Guide:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/nx-os/upgrade/guide/b_MDS_NX-OS_Upgrade_Guide/m_issu.html
Cisco MDS 9000 Family In-Service Software Upgrade - Cisco:
https://www.cisco.com/c/en/us/products/collateral/storage-networking/mds-9000-family/white_paper_c11-690736.html