DUMPSARENA

Implementing Cisco Data Center Core Technologies (350-601 DCCOR)

Cisco 350-601

Version Demo

Total Demo Questions: 68

Total Premium Questions: 682

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Network	221
Topic 2, Software	31
Topic 3, Compute	145
Topic 4, Storage Network	178
Topic 5, Automation	107
Total	682

QUESTION NO: 1

What is an advantage of using Network File System (NFS) rather than Storage Area Network (SAN)?

A. NFS has lower latency than SAN.

B. NFS supports more types of file systems.

C. NFS supports booting the operating system.

D. NFS allows direct file access.

ANSWER: D

Explanation:

NFS allows direct file access. NFS is a file-level storage protocol: a server exports a directory (file system) and clients mount that export, then access files and directories using normal file-oriented operations. The NFS server owns and presents the file system, including its directory structure, metadata, permissions, and file-sharing behavior. This makes NFS particularly useful when multiple hosts need shared access to the same files without each host having to create and manage a separate file system on an attached storage volume.

By contrast, a SAN presents block storage, such as LUNs, to hosts. A host sees a SAN LUN much like a locally attached disk and normally must partition, format, and mount it before files can be stored or retrieved. Consequently, NFS provides a simpler direct path to shared file data for workloads requiring file-level access. The IETF NFS specification describes NFS as a distributed file-system protocol that enables a client to access files on a remote server, while Cisco's storage overview distinguishes file storage from block-oriented SAN storage. See RFC 7530: Network File System Version 4 Protocol and Cisco Storage Networking.

QUESTION NO: 2

Which two actions are needed to configure a single Cisco APIC controller for a Cisco ACI fabric for the first time? (Choose two.)

A. Register the APIC that is connected to the switch.

B. Configure the leaf switch where the Cisco APIC is connected, using CLI to allow APIC connectivity.

C. Register the switches that are discovered through LLDP.

D. Configure the first Cisco APIC controller.

E. Register all leaf and spine switches.

ANSWER: C D

Explanation:

Configuring the first Cisco APIC controller is required to initialize the fabric. During the initial APIC setup wizard, the administrator establishes fundamental fabric parameters, including the fabric name, APIC and controller-cluster settings, and management-network addressing. These settings allow the APIC to become the policy and management point for the new ACI fabric and to begin communicating with the connected fabric infrastructure.

Registering the switches that are discovered through LLDP is also required as part of fabric bring-up. After the APIC is connected to an ACI leaf switch, the fabric uses LLDP-based discovery to identify adjacent ACI switches and their physical topology. The administrator claims/registers the discovered switches in APIC so that each device is assigned its intended identity, role, node ID, and other fabric parameters. This registration step incorporates the discovered leaf and spine infrastructure into the APIC-managed fabric topology. Cisco documents initial APIC setup and switch discovery/registration as key stages of bringing up an ACI fabric. See the Cisco APIC Installation Guide and the Cisco ACI documentation and support resources.

QUESTION NO: 3

What are two types of FC/FCoE oversubscription ratios? (Choose two.)

A. switch processing power to end-node processing power

B. port bandwidth to uplink bandwidth

C. server storage to end-node count

D. edge ISL bandwidth to core ISL bandwidth

E. host bandwidth to storage bandwidth

ANSWER: D E

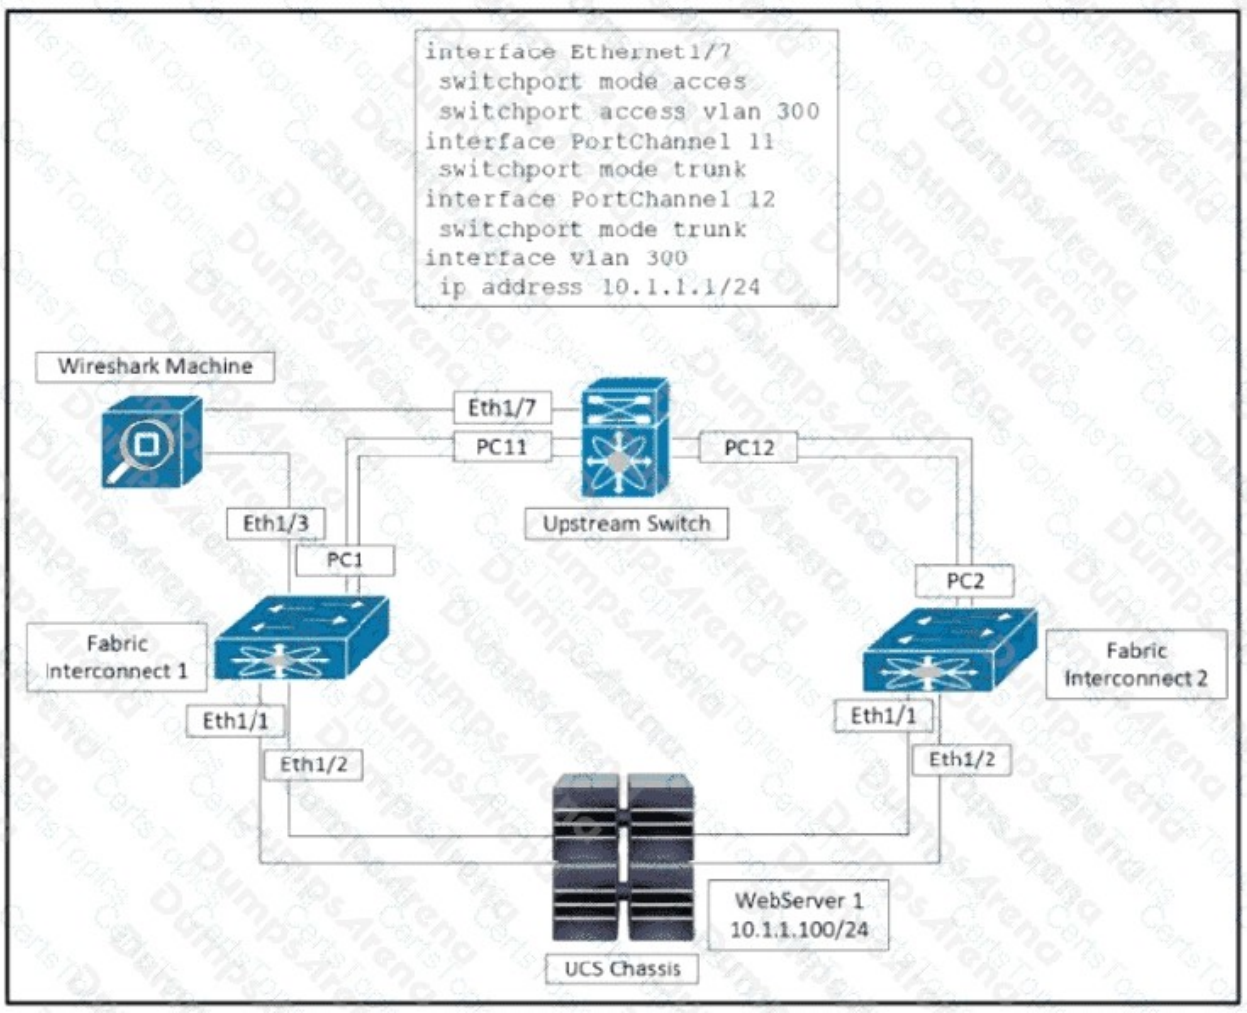
Explanation:

FC/FCoE SAN design evaluates oversubscription at both the host-to-storage access layer and across the fabric infrastructure. **host bandwidth to storage bandwidth** is the ratio between the aggregate bandwidth that attached servers can potentially generate and the aggregate bandwidth available at the storage-facing ports. This ratio helps size storage connectivity appropriately for the expected workload and prevents contention when multiple hosts access shared targets.

edge ISL bandwidth to core ISL bandwidth describes oversubscription in a core-edge SAN topology. It compares the aggregate bandwidth presented by edge-side interswitch links with the bandwidth available through the core-side interswitch links. Evaluating this ratio ensures that traffic crossing the fabric has adequate core capacity and that the core does not become a bottleneck as edge switches and connected devices scale.

These two measurements represent the principal access and fabric-path capacity considerations used when planning FC or FCoE SAN bandwidth. Cisco's SAN guidance describes core-edge topology and the need to size ISLs according to aggregate traffic requirements; Cisco's FCoE design guidance likewise addresses bandwidth planning between hosts and storage resources. See the Cisco MDS 9000 Fabric Configuration Guide and the Cisco FCoE design guide.

QUESTION NO: 4



Refer to the exhibit. An engineer is troubleshooting a connection issue from the virtual machine that is located on the Webserver 1 blade. The application team requested a packet capture for all traffic that is transferred to and from Webserver 1. The engineer confirmed that the host server vNIC1 has VLAN 300 assigned to It. Which action must be taken to capture the traffic?

A. Configure a SPAN session that is sourced from Eth1/1 with a destination that is set to port channel 1.

B. Configure a SPAN session that is sourced from Eth1/3 with a destination that is set to Eth1/2.

C. Configure a SPAN session that is sourced from port channel 1 with a destination that is set to port channel 2.

D. Configure a SPAN session that is sourced from service profile vNIC1 with a destination that is set to Eth1/3.

ANSWER: D

Explanation:

Configure a SPAN session sourced from service profile vNIC1 with a destination set to Eth1/3. In Cisco UCS, a vNIC is the logical Ethernet interface presented by the service profile to the server operating system or hypervisor. Because the requested capture must include all traffic entering and leaving the workload hosted on Webserver 1, selecting that specific vNIC as the SPAN source captures the relevant bidirectional traffic at the server-facing logical interface. The VLAN assignment confirms that vNIC1 is the interface carrying the workload's VLAN 300 traffic, so it is the appropriate capture point rather than an unrelated physical uplink or port channel. Eth1/3 is used as the physical SPAN destination interface shown for the packet-analysis device. The analyzer connected to that interface can then record the mirrored frames without becoming part of the production traffic path. Cisco UCS Manager supports SPAN configuration with vNICs as sources and physical interfaces as destinations, allowing focused troubleshooting for a particular service-profile interface. See Cisco's UCS Manager installation and configuration guides and the Cisco UCS Manager Configuration Management Guide.

QUESTION NO: 5

Which two components should be checked when a Cisco Nexus 9000 Series Switch fails to boot using POAP? (Choose two.)

A. POAP feature license

B. DHCP server to bootstrap IP information

C. image noted in the script file against switch bootflash

D. script signed with security key

E. TFTP server that contains the configuration script

ANSWER: B E

Explanation:

POAP depends on DHCP to provide the initially booting Nexus 9000 switch with its bootstrap network parameters, such as an IP address, default gateway, and—depending on the DHCP options used—the location of the POAP script. Therefore, verifying that the DHCP server is reachable on the correct VLAN, has an available lease, and returns the required POAP information is essential when the process does not start or cannot proceed.

The TFTP server that contains the configuration script must also be checked. After receiving the necessary bootstrap information, the switch downloads and executes the POAP script from the specified server. The TFTP service must be running, reachable from the management or in-band interface selected during POAP, and able to serve the script using the exact filename and path supplied through DHCP. Network reachability, VLAN configuration, routing where applicable, and any ACL or firewall policy affecting TFTP traffic should be validated. Cisco documents POAP as a DHCP-driven provisioning process in which the switch obtains a script from a network server and uses that script to complete installation and configuration. See the Cisco Nexus 9000 Series NX-OS POAP configuration guide and Cisco POAP troubleshooting documentation.

QUESTION NO: 6

An engineer is deploying Fibre Channel over Ethernet and must use Data Center Bridging Exchange Protocol to communicate Data Center Bridging configuration between directly connected devices.

Which two types of information can DCBX exchange? (Choose two.)

A. Priority Flow Control parameters

B. Enhanced Transmission Selection parameters

C. OSPF neighbor state

D. VXLAN VTEP addresses

E. Fibre Channel domain IDs

ANSWER: A B

Explanation:

DCBX is used by directly connected Data Center Bridging devices to exchange capability and configuration information. It can exchange Priority Flow Control parameters, including the traffic priorities for which lossless pause behavior is enabled.

DCBX also exchanges Enhanced Transmission Selection information, which defines bandwidth-allocation and traffic-class parameters. These exchanges help connected devices operate with compatible DCB settings for converged Ethernet traffic. Cisco documents DCB and DCBX operation in the Cisco Nexus 9000 Series QoS configuration guide.

QUESTION NO: 7

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: A B

Explanation:

vPC and HSRP provide complementary forms of redundancy that support a rolling upgrade strategy and minimize the effect of taking one Nexus switch out of service. A vPC domain connects downstream devices to two peer switches while presenting the paired links as a single port channel. Traffic can therefore continue through the remaining vPC peer when one peer is upgraded or reloaded. For a major upgrade, the peer switches can be handled sequentially, with traffic maintained by the operational peer, subject to observing the applicable Cisco upgrade procedure and compatibility requirements.

HSRP maintains first-hop default-gateway availability for connected hosts. The active HSRP device forwards traffic for the shared virtual IP and MAC address, while the standby device is prepared to assume that forwarding role. Before reloading a switch, the gateway role can be transitioned to its redundant peer; hosts continue using the same configured default gateway rather than requiring a gateway-address change. In a vPC design, HSRP gateway redundancy and dual-active forwarding paths work together to preserve both endpoint connectivity and default-gateway reachability during the maintenance window. Cisco documents vPC operation at Cisco Nexus 9000 vPC configuration guide and HSRP behavior at Cisco Nexus 9000 HSRP configuration guide.

QUESTION NO: 8

An engineer requires a solution to automate the configuration and deployment of remote network devices for a customer. The engineer must keep these considerations in mind

The customer ' s environment is based on industry-accepted standards and requires a solution that meets these standards.

The security requirements mandate the use of a secure transport mechanism between the automation software and target devices such as SSH or TLS.

The solution must be implemented using a human-readable language and provide the functionality to format data in XML or JSON.

Which solution must be used to meet these requirements?

A. Ansible

B. SNMP

C. REST API

D. NETCONF

ANSWER: D

Explanation:

NETCONF is the standards-based network-configuration protocol that directly matches these automation requirements. It defines a structured client-server mechanism for retrieving, modifying, and managing device configuration and operational

state. Its operations and data are encoded in XML, a human-readable structured format, and the data models used with NETCONF are commonly defined with YANG. This gives automation software a consistent, programmatic way to deploy configuration to remote devices while using an industry-standard protocol rather than device-specific command-line parsing.

Secure transport is an explicit part of the NETCONF architecture. The base standard requires support for NETCONF over SSH, providing authenticated and encrypted management sessions, and NETCONF can also be carried over TLS when that transport is required by the environment. Cisco platforms expose NETCONF as a programmable management interface, allowing tools to send structured configuration requests and receive structured replies. Therefore, NETCONF satisfies the requirements for standards alignment, secure remote transport, and human-readable structured configuration data. The NETCONF protocol specification is available in IETF RFC 6241, and the TLS transport mapping is defined in IETF RFC 7589.

QUESTION NO: 9

A new employee must be granted access to add VLANs into an existing Cisco UCS Manager and configure NTP synchronization with date and time zone settings. Which two privileges must be granted to the employee to complete the task? (Choose two.)

A. Service Profile Compute (ls-compute)

B. Ext LAN Policy (ext-lan-policy)

C. Service Profile Network Policy (ls-network-policy)

D. Service Profile Config (ls-config)

E. Ext LAN Config (ext-lan-config)

ANSWER: D E

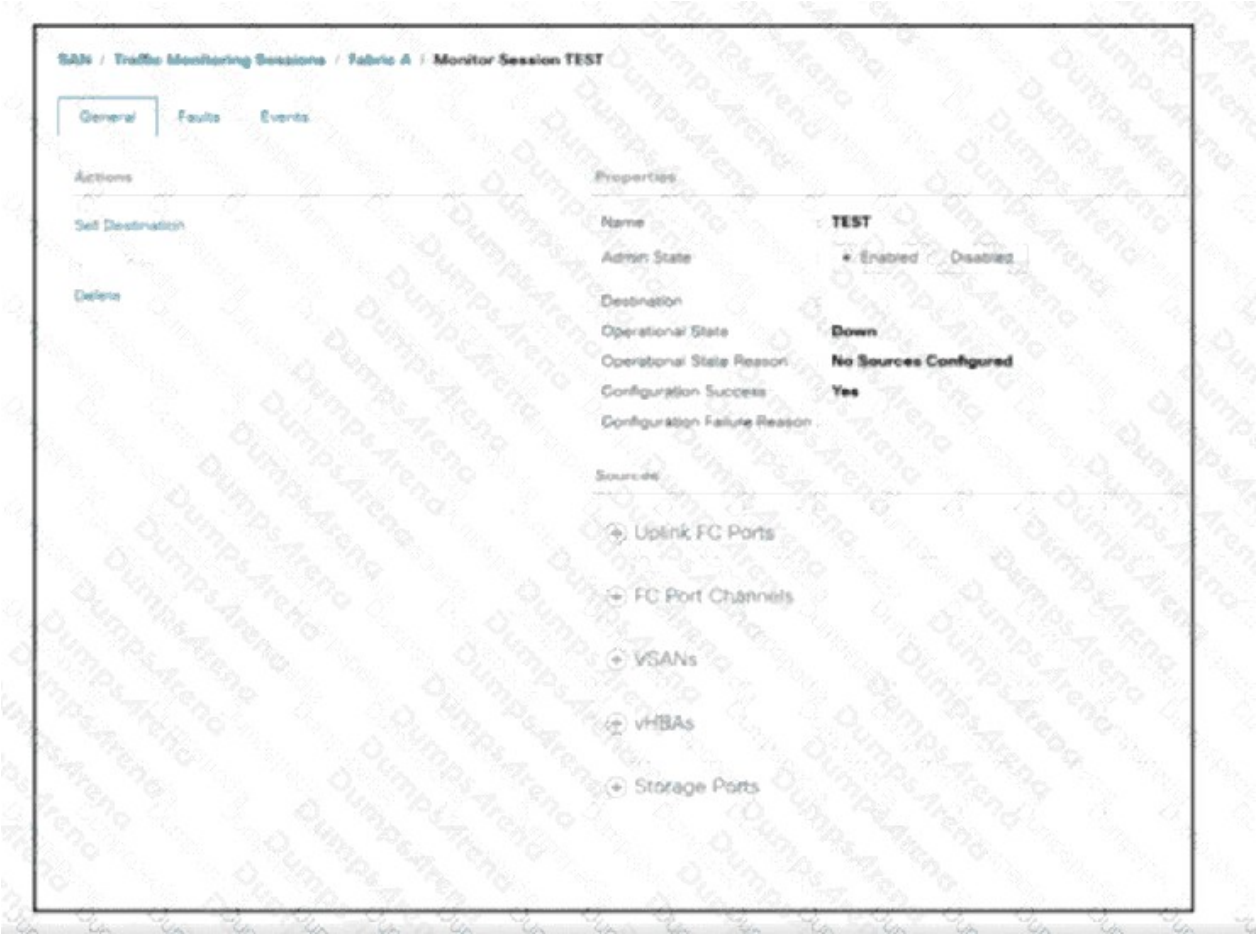
Explanation:

Service Profile Config (ls-config) is required for the service-profile-related configuration authority needed when administering UCS Manager configuration elements associated with the requested operational setup. This privilege is part of the Cisco UCS Manager role-based access-control model and allows the delegated administrator to make the relevant configuration changes without receiving unrestricted administrative access.

Ext LAN Config (ext-lan-config) is required for configuration of the external LAN environment, including the fabric-level LAN configuration required to add VLANs. VLANs are defined in UCS Manager's LAN configuration and can subsequently be used by service-profile networking constructs. Granting this privilege therefore delegates the necessary fabric networking authority, while Service Profile Config (ls-config) supplies the associated configuration permission required by the task.

Cisco UCS Manager evaluates privileges individually for managed objects and operations, so a custom role should include both privileges when the employee must perform both portions of this change. Cisco documents the UCS Manager privilege definitions and their RBAC usage in the Cisco UCS Manager Privileges Reference. Additional background on UCS Manager authorization is available in the Cisco UCS Manager Administration Management Guide.

QUESTION NO: 10



Refer to the exhibit. Users report SAN traffic packet loss on a blade server that is managed by using Cisco UCS Manager. An engineer must implement a SPAN feature to an external Fibre Channel analyzer. Which configuration source in the TEST monitoring session must the engineer select?

A. all the storage ports

B. all the virtual host bus adapters from the server

C. all the uplink Fibre Channel ports

D. all the VSANs from the server

ANSWER: C

Explanation:

all the uplink Fibre Channel ports is the appropriate source selection for this Fibre Channel SPAN session. SAN traffic from the blade server's virtual HBAs traverses the fabric interconnect and is forwarded through its uplink Fibre Channel interfaces toward the upstream SAN fabric. Monitoring all uplink FC ports provides the external Fibre Channel analyzer with visibility of the FC traffic crossing the fabric-interconnect-to-SAN boundary, including the traffic relevant to investigating reported packet loss. This lets the engineer inspect FC frames and associated transport behavior at the point where server SAN connectivity leaves the UCS domain, while using the analyzer-connected interface as the monitoring-session destination. Cisco UCS Manager supports Fibre Channel SPAN monitoring for troubleshooting FC connectivity and traffic flows, with uplink Fibre Channel interfaces available as FC SPAN sources. The monitoring configuration should be applied on the applicable fabric(s) so that the SAN paths used by the affected blade are represented in the capture. Cisco's UCS Manager configuration documentation describes monitoring sessions and their supported source and destination types; see the Cisco UCS Manager Configuration Guides. For FC monitoring concepts and operational context, see the Cisco UCS Manager GUI Configuration Guide.

QUESTION NO: 11

A network engineer must restore the configuration of a Cisco Nexus 9300 Series Switch The backup file is stored on an external server The current configuration must be replaced without having to reload the switch Which command must be used to accomplish this goal?



A. Option A

B. `configure replace bootflash:backup.cfg`

C. Option C

D. Option D

ANSWER: B

Explanation:

`configure replace bootflash:backup.cfg` is the NX-OS command used to replace the active running configuration with the contents of a saved configuration file without reloading the Nexus 9300 switch. Before issuing the replacement command, the engineer copies the backup from the external server to local bootflash storage, using an appropriate supported transfer method such as SCP, SFTP, FTP, or TFTP. The local filename used in the command must match that copied configuration file. Configuration replacement is specifically intended for restoring a known-good configuration as a complete replacement rather than merging individual lines into the current running configuration. NX-OS evaluates the replacement configuration and applies the necessary changes to the running configuration, allowing the device to remain operational during the restore process. The engineer should ensure that the backup contains reachable management connectivity or use console access, because replacing the configuration can alter the management interface, routing, AAA, or other access-related settings. Cisco documents configuration replacement under the NX-OS configuration-management features and provides the command syntax in the Nexus 9000 command reference. See the Cisco Nexus 9000 NX-OS Fundamentals Configuration Guide and the Cisco Nexus 9000 NX-OS Command Reference.

QUESTION NO: 12

A network automation engineer must configure a Cisco NX-OS REST API on a Cisco 9000 Series Switch with these requirements:

NX-API REST and the DME database must be enabled.

The NX-API REST must listen for HTTP traffic on port 83.

The HTTPS certificate must be uploaded.

Which configuration set must be used to accomplish these goals?

A)

```
feature nxapi
nxapi {http } port 83
nxapi certificate {httpscrt | httpskey}
nxapi certificate enable
```

B)

```
feature nxapi
nxapi (http) port 83
nxapi certificate {httpscrt}
```

C)

```
feature nxapi
nxapi (tcp) port 83
nxapi certificate {httpscrt}
nxapi certificate enable
```

D)

```
feature nxapi
nxapi (tcp) port 83
nxapi certificate {httpscrt}
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: B

Explanation:

Option B is correct because it combines the required NX-OS NX-API REST service activation, Data Management Engine (DME) availability, custom HTTP listener configuration, and certificate installation. Enabling the NX-API feature activates the NX-API infrastructure and its REST interface, which uses the DME object model to expose switch configuration and operational state through REST requests. The HTTP port must then be explicitly set to `83`; otherwise, NX-API uses its default HTTP listener behavior rather than satisfying the stated port requirement. The configuration also includes the HTTPS certificate upload or certificate-file association required for the switch to present the configured certificate when clients establish secure NX-API connections. This is important for trusted HTTPS automation sessions and for avoiding reliance on the switch's default self-signed certificate. Together, these elements meet every stated requirement: REST/DME enablement, HTTP service on port 83, and HTTPS certificate deployment. Cisco documents NX-API REST configuration, including enabling NX-API and setting HTTP/HTTPS parameters, in the Cisco Nexus 9000 NX-OS NX-API configuration guide. Cisco's NX-API REST SDK documentation also describes the DME-based REST API model.

QUESTION NO: 13

An engineer is implementing VSAN 10 on multiple Cisco Nexus 5600 Series Switches and must ensure that the full zone set and active zone set are identical across the fabric. Which configuration must be implemented to meet this requirement?

A. switch(config)# zoneset distribute vsan 10

B. switch(config)#zoneset activate name ZONE10 vsan 10

C. switch(config)# zoneset distribute full vsan 10

D. switch(config)# zone-attribute-group name ATTR1 vsan 10

ANSWER: C

Explanation:

`switch(config)# zoneset distribute full vsan 10` is required because the `full` keyword distributes the complete zoning database for VSAN 10 throughout the fabric. This includes every configured zone and zoneset, rather than only the currently active zoneset. As a result, all participating Nexus switches maintain the same full zoneset information as well as the same active zoning state, which is the stated operational requirement. Cisco zoning distribution uses the fabric control-plane to synchronize this information between switches, helping ensure that a subsequent zoneset activation is based on a consistent database everywhere in the VSAN. The command should be configured in global configuration mode and applies specifically to VSAN 10, so it does not unintentionally alter zoning-distribution behavior for other VSANs. Cisco documents the distinction between normal zoneset distribution and full zoneset distribution in its SAN switching guidance; full distribution is the appropriate choice when both configured and active zoning content must be synchronized across all switches. See the Cisco Nexus 5600 Series SAN Switching Configuration Guide and Cisco's zoning configuration documentation.

QUESTION NO: 14

What is a characteristic of EPLD updates on Cisco MDS 9000 Series Switches?

A. EPLD bundles are released separately from a Cisco MDS NX-OS release

B. EPLD packages update hardware functionality on a device

C. EPLD updates are nondisruptive to traffic flow

D. EPLD updates are installed only via the Cisco DCNM GUI

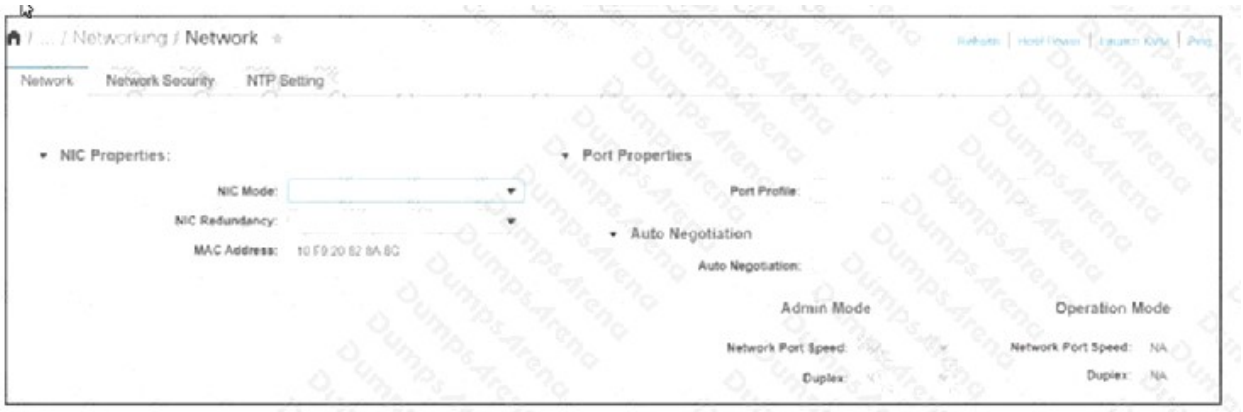
ANSWER: B

Explanation:

EPLD packages update hardware functionality on a device. EPLD stands for erasable programmable logic device and refers to programmable hardware logic used by Cisco MDS 9000 Series modules and switches. An EPLD image contains firmware for hardware components such as interfaces, control logic, and other programmable devices on the platform. Updating this image ensures that the switch hardware operates with the intended logic revisions, including hardware fixes, feature support, and compatibility required by the applicable Cisco MDS NX-OS software release.

Cisco treats EPLD maintenance as a hardware-firmware operation rather than a normal NX-OS software configuration change. Before performing an update, administrators should use the EPLD verification facilities to identify components whose installed versions differ from the image versions, then plan the maintenance according to the module and platform requirements. Cisco's MDS software installation documentation describes EPLD image management and the procedures used to upgrade programmable-device firmware. The `install epld` workflow applies the relevant firmware images to supported hardware components. See the Cisco MDS 9000 Series NX-OS Software Installation and Upgrade Guide: Upgrading EPLD Images and Cisco's MDS 9000 Series support documentation for platform-specific guidance.

QUESTION NO: 15



Refer to the exhibit. An engineer configures a Cisco C-Series UCS server. The company policy states that all C-Series servers must be configured so that Cisco IMC is reachable only from the management port of the server. Which NIC configuration must the engineer apply to meet this objective?

A. NIC Mode: Shared OCPNIC Redundancy: Active-Active

B. NIC Mode: DedicatedNIC Redundancy: None

C. NIC Mode: Shared LOMNIC Redundancy: Active-Standby

D. NIC Mode: Cisco CardNIC Reundancy: Active-Standby

ANSWER: B

Explanation:

NIC Mode: DedicatedNIC Redundancy: None meets the requirement because Dedicated NIC mode assigns Cisco Integrated Management Controller (Cisco IMC) network access exclusively to the server's physical, dedicated management Ethernet port. This separates out-of-band management traffic from the host's production-facing LAN-on-motherboard, Cisco card, or OCP network interfaces. Consequently, administrators can reach Cisco IMC only through the management network connected to that dedicated port, as required by the stated policy.

The *None* redundancy setting is appropriate because a dedicated management interface is a single physical interface and does not use a secondary data NIC as a failover path. This configuration provides a clear management-plane boundary: Cisco IMC functions such as remote KVM, power control, hardware inventory, event logs, and firmware administration remain available through the isolated management connection without exposing the controller through shared host network interfaces. Cisco documents the dedicated management port as the interface used for out-of-band Cisco IMC connectivity and configuration. See the Cisco UCS C-Series GUI Configuration Guide and the Cisco UCS C-Series Rack Servers product documentation.

QUESTION NO: 16

An engineer performs a Cisco UCS C480 ML M5 installation from a Cisco UCS Manager As a part of the design the server must perform an iSCSI bool from the storage array Which configuration must be performed in the SAN settings in the root organization to satisfy that condition?

A. Implement a WWPN pool.

B. Create a virtual iSCSI vNIC.

C. Enable a host SAN pool.

D. Create an IQN suffix pool.

ANSWER: D

Explanation:

Create an IQN suffix pool. For iSCSI boot, the server requires a unique iSCSI Qualified Name (IQN) so that the storage system can identify the initiator and associate it with the appropriate boot LUN. In Cisco UCS Manager, an IQN suffix pool is configured under the root organization's SAN settings and provides the unique suffix portion of the IQN assigned to each service-profile-based server. UCS Manager combines the configured IQN prefix with an allocated suffix from this pool to create the initiator identity used by the virtual iSCSI adapter. This pool-based allocation is especially important for service profile mobility and consistent, automated server provisioning, because the iSCSI initiator identity follows the service profile according to its configuration. The iSCSI boot policy can then reference the virtual iSCSI interface and define the storage target and boot-LUN information. Cisco's UCS Manager guidance for iSCSI boot specifically includes creating an IQN suffix pool as part of the SAN configuration before configuring the iSCSI boot policy and service profile. See the Cisco UCS Manager GUI Configuration Guide and the Cisco UCS Manager installation and configuration guides.

QUESTION NO: 17

An engineer must configure OSPF in the data center. The external routes have already been redistributed into OSPF. The network most meets 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:

`default-information originate` is required because it causes the OSPF autonomous system boundary router to advertise an OSPF default route only when a default route, `0.0.0.0/0`, is present in its local routing table. A default learned dynamically from the upstream router therefore acts as the condition that enables the Internet exit advertisement. When that upstream-learned default route is removed, the router no longer meets the prerequisite for originating the OSPF default, so the default advertisement is withdrawn. This behavior ensures that data center devices receive a usable route toward cloud services only while the upstream path is available. The already redistributed external routes remain available through the existing redistribution configuration; default origination specifically supplies the catch-all route needed for destinations that do not match more-specific internal or external prefixes. Cisco documents that the OSPF `default-information originate` command originates a default external route when a default route exists in the routing table. See the Cisco IOS XE OSPF Configuration Guide and the OSPF default-route behavior defined in RFC 2328.

QUESTION NO: 18

An architect must select a scripting language for automating Cisco Nexus Series Switches and Cisco UCS Manager. The language must comply with these requirements:

â€¢Can run in interactive mode for testing

â€¢Can be run from different operating systems without changing code â€¢Extensible with a diverse set of libraries

Which scripting language meets these requirements?

A. Bash

B. Python

C. TCL

D. JavaScript

ANSWER: B

Explanation:

Python meets all the stated automation requirements. Its interactive interpreter (the Python REPL) lets engineers execute commands, inspect objects, test API calls, and validate logic incrementally before placing a script into production. Python is designed to be portable: the same source code can generally run on Windows, Linux, and macOS when a compatible Python runtime and required dependencies are available. This makes it well suited to automation environments where administrators use different workstation or server operating systems.

Python also has a mature and extensive ecosystem of standard-library modules and third-party packages. For Cisco data-center automation, this includes libraries and SDKs that support structured API interactions, data formats such as JSON and XML, HTTP communication, authentication, logging, testing, and configuration management. Cisco provides Python-focused interfaces and examples for NX-OS programmability, and the Cisco UCS Python SDK enables programs to manage UCS domains through the UCS Manager API. These capabilities allow reusable scripts to automate inventory collection, configuration deployment, policy management, and operational validation across Nexus switches and UCS Manager. See Cisco's NX-OS programmability documentation and the Cisco UCS Developer Center.

QUESTION NO: 19

```
NX-Core# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
NX-Core(config)# interface fc2/1
NX-Core(config-if)# switchport speed 4000
NX-Core(config-if)# no shutdown
NX-Core(config)# interface fc2/2
NX-Core(config-if)# no shutdown
NX-Core(config)# monitor session 1
NX-Core(config-monitor)# source interface fc2/1
NX-Core(config-monitor)# destination interface fc2/2
NX-Core(config-monitor)# no shutdown
```

An engineer implements SPAN configuration on an Nexus 5000 series switch. Which two commands are necessary to complete the SPAN session configuration? (Choose two)

A. Configure the switchport mode SD under interface fc2/1.

B. Configure switchport speed 4000 under interface fc2/2

C. Configure the switchport monitor under interface fc2/2.

D. Configure the switchport mode SD under interface fc2/2.

E. Configure session type fc under monitor session 1.

ANSWER: B D

Explanation:

Fibre Channel SPAN on a Nexus 5000 uses a dedicated SPAN destination port. The destination interface must be placed into SD (SPAN destination) mode with `switchport mode SD`. This mode reserves the Fibre Channel interface for receiving copied traffic from the monitor session rather than participating as a normal fabric-facing switch port. Therefore, *Configure the switchport mode SD under interface fc2/2.* is required for the designated destination interface.

The SPAN destination port must also operate at a compatible Fibre Channel speed for the monitored traffic and the attached analyzer. Configuring `switchport speed 4000` on the destination interface establishes the required 4-Gbps operational speed in this configuration. Consequently, *Configure switchport speed 4000 under interface fc2/2* completes the necessary destination-port preparation for the FC SPAN session. Cisco's SAN switching guidance describes configuring FC SPAN destination interfaces in SD mode and configuring the relevant interface attributes before use: Cisco Nexus 5000 Series NX-OS SAN Switching Configuration Guide. Additional NX-OS SPAN configuration information is available in the Cisco Nexus 5000 System Management Configuration Guide.

QUESTION NO: 20

Which LDAP rule must be used in Cisco UCS Manager to search for a user continuously down the tree hierarchy until the user is found?

A. disable secure socket layer

B. enable group authorization

C. set group recursion to recursive

D. Check box of use primary group

ANSWER: C

Explanation:

set group recursion to recursive is correct because Cisco UCS Manager uses the Group Recursion setting in an LDAP group rule to evaluate nested LDAP group membership recursively. With recursive processing enabled, UCS Manager continues traversing the directory's group hierarchy rather than limiting its lookup to only the immediately associated group. This is important when enterprise directory designs use nested groups to assign roles indirectly: a user can belong to a child group, which belongs to another group that is mapped to a UCS Manager role. Recursive group evaluation allows UCS Manager to follow that membership relationship through the hierarchy and determine whether the user is authorized through the configured LDAP group rule. The setting is configured as part of the LDAP group-rule configuration after an LDAP provider has been defined. Cisco documents Group Recursion as an LDAP group-rule attribute used to support recursive group membership lookup. See Cisco's UCS Manager Security Configuration Guide and the Cisco UCS documentation portal for the applicable UCS Manager release.

QUESTION NO: 21

An engineer creates a service profile in Cisco UCS Manager and must assign a policy that reboots blades when changes are applied. The changes must be applied only after user acknowledgment. Which two policies must be configured to meet these requirements? (Choose two.)

A. Reboot Policy

B. Boot Policy

C. Power Control Policy

D. Global Policy

E. Maintenance Policy

ANSWER: A E

Explanation:

Reboot Policy and **Maintenance Policy** must be configured together. A Reboot Policy defines the reboot behavior for a server associated with a service profile, ensuring that a required restart can occur when a configuration change must take effect. This makes the reboot action an explicit part of the service-profile configuration rather than relying on an administrator to handle it manually.

A Maintenance Policy controls when disruptive service-profile changes are deployed to the blade. Setting its deployment behavior to user acknowledgment holds pending changes until an authorized user confirms their application. Once acknowledged, UCS Manager can apply the pending configuration and use the assigned reboot behavior where the change requires a restart. This combination supports controlled change management: the blade is restarted as necessary, but only after the user has approved deployment of the disruptive update.

Cisco documents maintenance policies as the mechanism for selecting user acknowledgment and other deployment timing methods for changes that can disrupt server operation. See the Cisco UCS Manager Configuration Guide and the Cisco UCS Manager GUI Configuration Guide library.

QUESTION NO: 22

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: A C

Explanation:

Change the vPC system-priority of the replacement chassis to a higher value than the peer and configure auto-recovery to the disable state on both peers. The replacement peer must not supersede the currently operating peer's vPC control role when it rejoins the domain. Assigning the replacement chassis a higher system-priority preserves the incumbent peer's selection and prevents an unnecessary change in the vPC/LACP system identity seen by connected downstream devices. Keeping that identity stable allows existing vPC-enabled LACP bundles to retain their established negotiation state rather than being torn down and initialized again.

Auto-recovery must also be disabled on both peers before the isolated replacement device is reloaded. This prevents a peer that starts without peer-link communication from independently recovering and assuming an operational role that could alter vPC state during the replacement sequence. Together, these settings ensure that the surviving peer remains authoritative while the replacement chassis is introduced and synchronization is completed. Cisco's vPC configuration and recovery guidance is available in the Cisco Nexus 7000 Series vPC Configuration Guide and Cisco's vPC configuration documentation.

QUESTION NO: 23

An engineer must start a software upgrade on a Cisco Nexus 5000 Series Switch during a zone merge. What is the result of this action?

A. The upgrade stops

B. The zone merge pauses until the upgrade completes

C. The zone merge stops

D. The zone merge executes and then the upgrade completes

ANSWER: A

Explanation:

The upgrade stops. Cisco Nexus 5000 Series switch software-upgrade procedures include a safeguard for Fibre Channel zoning activity because a zone merge is a fabric-wide process that reconciles zoning databases when switches or fabrics join. Starting an upgrade while that process is underway is not permitted to continue normally: the upgrade operation is stopped so that the active zoning operation can complete without an overlapping disruptive software-maintenance workflow. The engineer should allow the zone merge to finish and confirm that the fabric has reached a stable state before initiating the upgrade again. This behavior protects SAN control-plane consistency and helps avoid introducing additional state changes while zoning information is being synchronized across the fabric. Cisco documents this as a software-upgrade prerequisite and operational consideration for Nexus 5000 platforms with Fibre Channel features. The applicable upgrade and downgrade guide is available from Cisco Nexus 5000 Series Upgrade and Downgrade Guide. Cisco's SAN zoning overview also describes the purpose of zoning and the importance of maintaining consistent zone-set information across a fabric: Cisco SAN zoning documentation.

QUESTION NO: 24

Port security is statically configured on a Cisco Nexus 7700 Series switch and F3 line card. The switch is configured with an Advanced Services license. Which two actions delete secured MAC addresses from the interface? (Choose two.)

A. The address must be removed from the configuration.

B. Shutdown and then no shutdown must be run on the interface.

C. The device must be restarted manually.

D. The address must reach the age limit that is configured for the interface.

E. The interface must be converted to a routed port.

ANSWER: A E

Explanation:

With static port security on a Nexus 7700 F3 interface, a secure MAC address is a configured Layer 2 port-security entry rather than a learned entry that is handled through normal aging behavior. Removing the address from the configuration explicitly removes the static secure-MAC binding from that interface. This is the direct administrative mechanism for deleting an individually configured secure address while retaining the interface as a switchport.

Converting the interface to a routed port also deletes its secured MAC addresses. Port security is a Layer 2 switchport feature, and making the interface a routed interface removes the Layer 2 switchport context in which the static port-security entries exist. Cisco documents static secure MAC address handling and the effects of changing an interface's Layer 2 mode in the Nexus 7000 Series security configuration documentation. The Advanced Services license and the F3 line card provide the platform support for this port-security capability; they do not change the deletion behavior of statically configured addresses. See Cisco's Nexus 7000 Series NX-OS Security Configuration Guide and the Nexus 7000 Series NX-OS Command Reference.

QUESTION NO: 25

Which two methods are available to manage an ACI REST API session authentication when a user is unauthenticated? (Choose two.)

A. POST to aaaUserLogin

B. GET to aaaListDomains

C. GET aaaRefresh

D. DELETE to aaaLogout

E. POST to aaaLogin

ANSWER: B E

Explanation:

An unauthenticated Cisco Application Policy Infrastructure Controller REST client can use `GET to aaaListDomains` to retrieve the authentication domains configured on the controller. This supports a login workflow because the client can discover the domain context before submitting user credentials. The request is made to the `/api/aaaListDomains.json` endpoint and does not require an existing authentication token or session cookie.

`POST to aaaLogin` is the session-establishment operation. The client sends a POST request to `/api/aaaLogin.json` with an `aaaUser` payload containing the username and password, optionally including the selected domain. On successful authentication, the controller returns an `aaaLogin` response and provides the token/cookie used to authenticate subsequent REST requests. These two operations therefore provide the discovery and authentication capabilities available before a REST session exists. Cisco documents the APIC REST authentication process and login request format in the APIC REST API Configuration Guide. Cisco's Using the REST API documentation also describes REST URL structure and API request usage.

QUESTION NO: 26

A service profile in Cisco UCS Manager must be configured with these requirements:

Cisco UCS Manager must identify the interfaces on the VMware ESXi host that connect to the Cisco ACI fabric switches.

The vNIC must support traffic management when congestion occurs.

Which two policies meet these requirements? (Choose two.)

A. vNIC/vHBA Placement

B. Network Control

C. Serial Over LAN

D. LAN Connectivity

E. QoS

ANSWER: B E

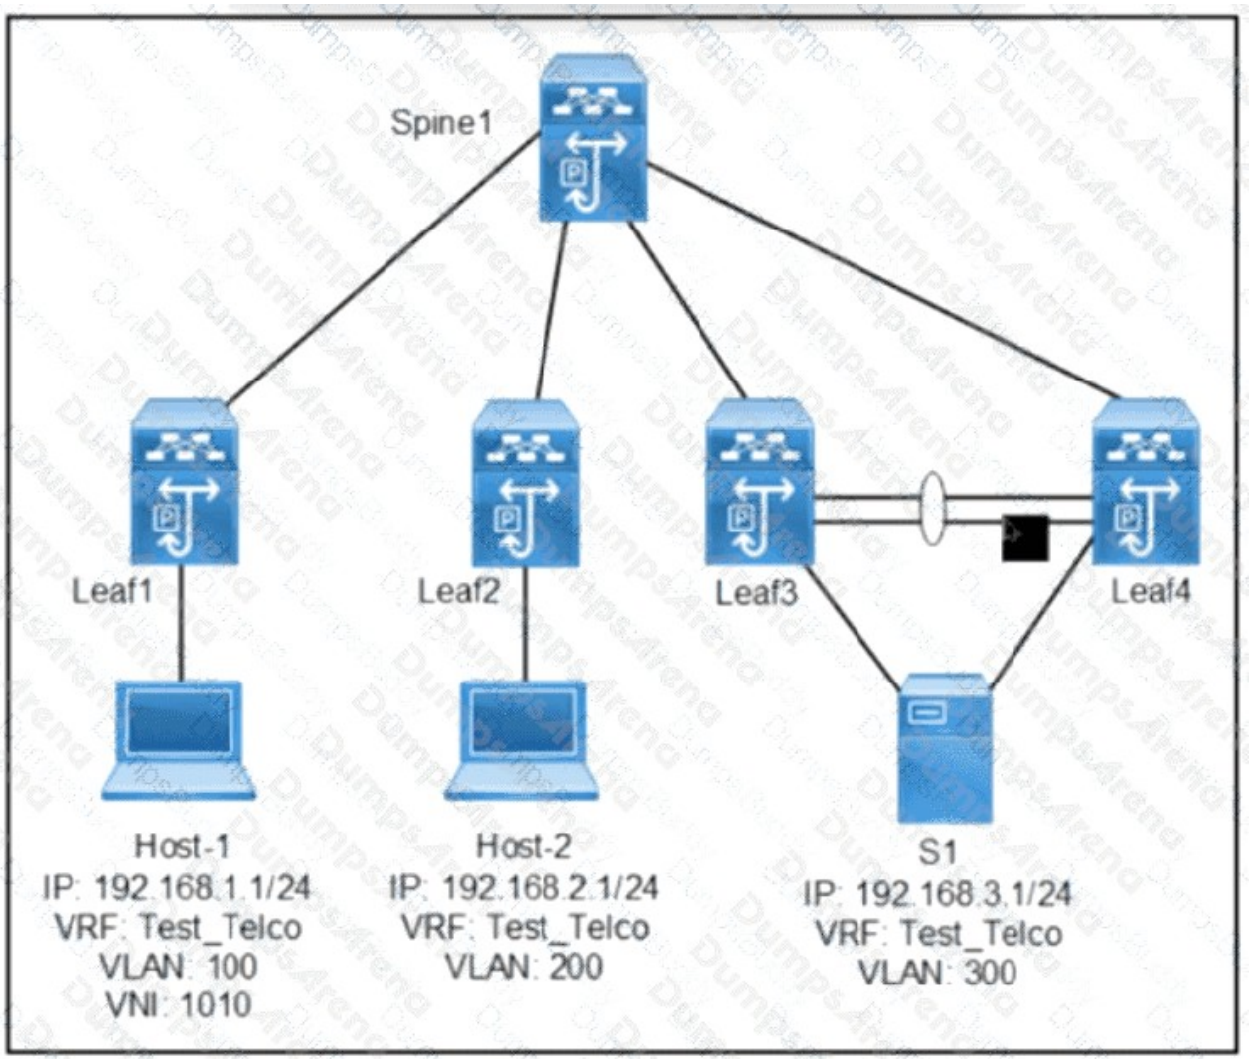
Explanation:

Network Control is required because a UCS network control policy defines Layer 2 control settings for vNICs, including CDP and LLDP. Enabling LLDP allows the connected infrastructure to advertise and discover neighbor and port information. This provides the interface-identification capability needed when ESXi host uplinks connect into the Cisco ACI fabric. Cisco UCS Manager can apply the selected network control policy to the vNIC through the service profile's LAN connectivity configuration.

QoS is required for congestion-related traffic handling. A UCS QoS policy assigns the vNIC traffic class and defines QoS characteristics such as priority, bandwidth allocation or limits, and packet-drop behavior. These settings determine how traffic is treated when adapter or fabric resources are contended, allowing important traffic to receive the intended service level during congestion. The QoS policy is associated with the vNIC in the service profile so that the policy follows the vNIC configuration consistently.

Cisco documents network control policies and their CDP/LLDP settings in the Cisco UCS Manager GUI Configuration Guide. Cisco also describes vNIC QoS configuration and traffic-class behavior in the Cisco UCS Manager configuration guides.

QUESTION NO: 27



Refer to the exhibit A VXLAN data center fabric has three hosts mapped to three different VLANs. Ingress and egress VTEPs perform Layer 2 and Layer 3 lookups VLAN 100 is mapped to VNI 1010 on Lead. Which set of actions allow communication between Host-1 and Host-2?

A. MAP VLAN 200 to Layer 2 VNI 2020 on Leaf2. Attach Test_Telco_VRF to Layer 2 VNI 1010.

B. Suppress ARP for Layer 2 VNI under interface NVE. Configure anycast gateway feature on interface vlan of L3 VNI

C. Enable prefix-based routing under interface vlan of L3 VNI. MAP VLAN 200 to layer 2 VNI 2020 on Leaf2

D. Attach Test_Telco_VRF to Layer 3 VNI 1010. Suppress ARP for Layer 3 VNI under interface NVE

ANSWER: C

Explanation:

“Enable prefix-based routing under interface vlan of L3 VNI. MAP VLAN 200 to layer 2 VNI 2020 on Leaf2” provides the required Layer 2 and Layer 3 VXLAN functions for traffic between endpoints in different VLANs. Mapping VLAN 200 to Layer 2 VNI 2020 on Leaf2 extends the Host-2 VLAN across the VXLAN overlay and gives the egress VTEP a local VNI-to-VLAN association for delivery to Host-2. Because Host-1 and Host-2 reside in separate VLANs, their traffic also requires routing through the tenant VRF’s Layer 3 VNI. Enabling prefix-based routing on the Layer 3 VNI SVI permits the VTEP to make the routed lookup for the destination subnet and forward the packet toward the destination Layer 2 VNI. This matches the VXLAN routing model in which the ingress VTEP performs the required lookup and encapsulates traffic to the remote VTEP, while the remote VTEP decapsulates it and performs the local VLAN forwarding needed to reach the attached host. Cisco’s VXLAN configuration guidance describes the use of Layer 2 VNIs for VLAN extension and Layer 3 VNIs for inter-VNI routing within a VRF. See the Cisco Nexus 9000 VXLAN Configuration Guide and Cisco VXLAN BGP EVPN configuration documentation.

QUESTION NO: 28

Which component is disrupted when the Cisco Integrated Management Controller is upgraded on a Cisco UCS Series Server?

A. SAN traffic

B. Cisco UCS Manager

C. KVM sessions

D. data traffic

ANSWER: C

Explanation:

KVM sessions is correct. The Cisco Integrated Management Controller (CIMC) supplies the server's out-of-band management functions, including the HTML-based and Java-based KVM console services. When CIMC firmware is upgraded, the management controller itself restarts so that the newly installed firmware can run. That restart temporarily makes the CIMC management interface unavailable and terminates any active remote KVM connections. An administrator using the console must reconnect after the CIMC has completed its reboot and becomes reachable again.

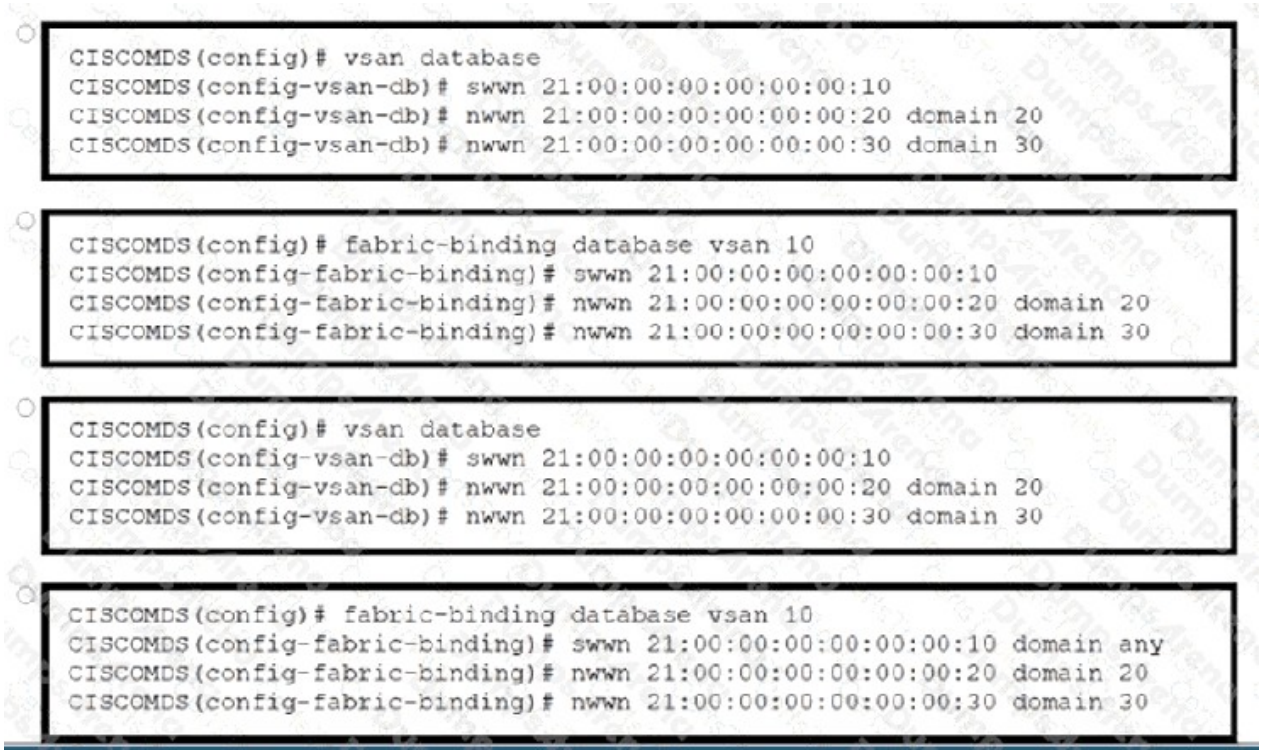
This behavior is important when scheduling maintenance because a KVM console may be providing the only remote view of a server's boot process, operating-system console, or BIOS configuration. Cisco's C-Series firmware-upgrade documentation identifies the CIMC reboot as part of the controller upgrade process, and Cisco's CIMC documentation describes KVM as a service hosted through the management controller. Plan for a brief loss of remote console access and reconnect the session after the upgrade status is complete. See the Cisco UCS C-Series Firmware Upgrade Guide and the Cisco UCS C-Series Rack Servers support documentation.

QUESTION NO: 29

```
switch# show fabric-binding database active vsan 10

-----------------------------------------------
Vsan   Logging-in Switch WWN      Domain-id
-----------------------------------------------
10     21:00:00:00:00:00:00:10     Any
10     21:00:00:00:00:00:00:20     20
10     21:00:00:00:00:00:00:30     30
[Total 3 entries]
```

Refer to the exhibit An engineer must configure fabric binding on a Cisco MDS 9000 Series Switch to match the output Which fabric database configuration accomplishes this goal?



A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: D

Explanation:

Option D is the correct selection for the fabric-binding database shown in the exhibit. Fabric binding is a Cisco MDS switch security feature that restricts E ports to authorized switches by maintaining an explicit list of permitted peer switch identities in the fabric-binding database. To reproduce the required output, the configuration must define the applicable switch WWNs in the correct fabric-binding scope and associate each authorized peer with the required domain information. The running configuration is then activated or committed so that the active fabric-binding database reflects those authorized members.

This controls which switches may form or retain interswitch connections, helping prevent an unauthorized switch from joining the SAN fabric. Fabric binding is configured under the relevant VSAN context, and the configured database must correspond exactly to the peer identities and domain assignments displayed by the requested output. Cisco documents fabric binding as part of SAN switch security and describes the fabric-binding database and activation workflow in the Cisco MDS 9000 Series Security Configuration Guide. Additional MDS SAN security guidance is available from Cisco MDS 9000 Series documentation.

QUESTION NO: 30

An engineer must permit communication from a web EPG to an application EPG in a Cisco ACI tenant. The application EPG must allow only the TCP ports defined in a contract filter.

Which two actions must be taken to meet the requirement? (Choose two.)

A. Configure the web EPG as a contract consumer.

B. Configure the application EPG as a contract provider.

C. Associate both EPGs with separate VRFs.

D. Configure the application EPG as a contract consumer.

E. Disable unicast routing on the bridge domain.

ANSWER: A B

Explanation:

The web EPG must consume the contract, and the application EPG must provide the contract. In Cisco ACI, a contract defines the policy rules that permit communication between EPGs. The consuming EPG initiates communication that is subject to the filters in the contract, while the providing EPG offers the permitted service.

After the contract is applied in both directions, the filter entries in the contract determine the allowed protocols and ports. Without a provider and consumer relationship, the default inter-EPG policy denies the traffic. Cisco documents contract relationships in the Cisco APIC Tenant Configuration Guide.

QUESTION NO: 31

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 configuration `spanning-tree bpduguard enable` on Ethernet1/1 is appropriate for a port intended exclusively for host connectivity. BPDU Guard protects an edge-facing port from an accidental Layer 2 connection to another switch. A normal endpoint such as a server, workstation, or appliance should not transmit Spanning Tree Protocol bridge protocol data units (BPDUs). If the interface receives a BPDU, that indicates that a switch or another bridged Layer 2 device may have been connected. BPDU Guard immediately places the port into the err-disabled state, stopping forwarding on that interface and preventing the unexpected device from participating in the Layer 2 topology. This provides a direct safeguard against accidental switching loops and unauthorized topology changes at the access edge. On Cisco Nexus platforms, BPDU Guard can be configured per interface, as shown, and is commonly paired operationally with edge-port behavior for host-facing links. Cisco documents that enabling BPDU Guard causes the port to shut down when it receives BPDUs, protecting the network from potential loops. See Cisco's NX-OS Spanning Tree Configuration Guide and the spanning-tree bpduguard command reference.

QUESTION NO: 32

An engineer must implement a set of changes in a Cisco Nexus 9000 Series Switch so that:

Which set of actions meets these requirements?

A)

B)

C)

D)

A. Option A

B. Option B

C. Option C

D. Option D

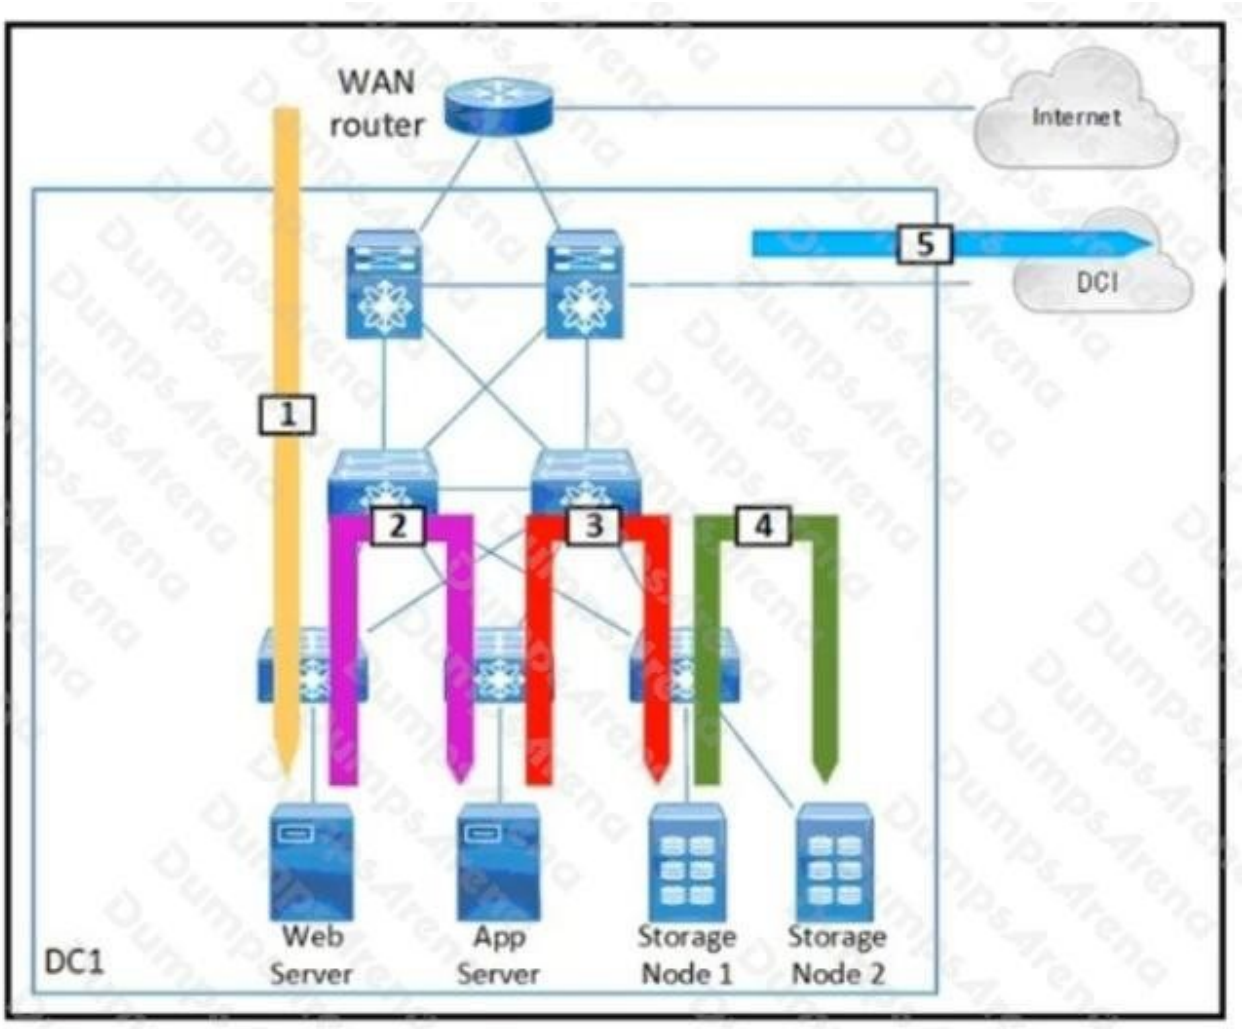
ANSWER: B

Explanation:

The supplied question does not include the required outcomes after “so that,” and the answer choices contain only placeholders rather than configuration actions. Consequently, there is no technical information from which to validate a Nexus 9000 implementation, such as the relevant NX-OS feature, required commands, platform mode, or operational verification criteria. The provided answer key identifies “Option B” as the intended single correct response, so it remains the correct selection to preserve the source item’s designated answer. A technically valid review would require the omitted requirements and the full text of each proposed action set; without them, assigning correctness based on Cisco NX-OS behavior would be speculative. Cisco Nexus 9000 configuration behavior and supported features must be evaluated against the applicable NX-OS release documentation and configuration guide, available through the Cisco Nexus 9000 Series installation and configuration guides. Cisco’s Nexus 9000 Series support page is also the appropriate starting point for release-specific documentation.

QUESTION NO: 33 - (DRAG DROP)

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:



ANSWER:



Explanation:

Number 1 is correctly mapped to North-South because this flow describes traffic that enters or exits the data center from an external-facing network, such as users, clients, WAN services, or Internet-connected resources reaching application workloads. In Cisco data-center designs, north-south is the conventional direction used for traffic moving between the data-center environment and outside networks. Number 2 is correctly mapped to East-West because it represents communication between workloads within the same data center, including application tiers, virtual machines, containers, or servers

exchanging data across the fabric. Cisco explains the importance of scalable east-west traffic handling in modern leaf-spine architectures in its data-center networking overview.

Number 3 correctly identifies Storage Traffic. This is the host-to-storage data path used for normal read and write operations, such as storage access over Fibre Channel, iSCSI, or network-attached storage. It is a distinct traffic class because it commonly requires predictable latency, bandwidth, and loss characteristics. Number 4 correctly identifies Storage Replication, which is the data movement used to synchronize copies between storage systems for resilience, backup, or disaster recovery. Cisco storage networking guidance discusses the role of dedicated and reliable storage connectivity in its storage networking portfolio.

Number 5 correctly identifies Inter-Data Center traffic because it crosses between separate data-center locations. This flow supports services such as workload mobility, replication, business continuity, and site-to-site application communication. The unused North-West and South-West labels are not recognized Cisco data-center traffic-flow categories; the five placed labels accurately cover the depicted paths.

QUESTION NO: 34

A company provides applications and database hosting services to multiple customers using isolated infrastructure-as-a-service services within the same data center environment. The environment is based on Cisco MDS 9000 Series Switches. The requirement is to manage the environment by using Fibre Channel Security Protocol and to enable user authentication when the centralized AAA server is unreachable. All communication between the MDS switches and the remote servers must be encrypted. Which command set must be used to meet these requirements?

A. aaa group server tacacs+ TacacsServerlaaa authentication login console TacacsServerl

B. aaa group server tacacs+ TacacsServerl _aaa authentication dhchap default group Tacattserveii

C. aaa group server radius RadiusServeMaaa authentication login default RadiusServeM

D. aaa group server radius RadiusServer1
aaa authentication dhchap default group RadiusServer1

ANSWER: D

Explanation:

`aaa group server radius RadiusServer1`
`aaa authentication dhchap default group RadiusServer1` is the appropriate command set because FC-SP uses DH-CHAP for Fibre Channel entity authentication, and Cisco MDS NX-OS integrates external DH-CHAP authentication through a configured RADIUS server group. DH-CHAP is a challenge-handshake mechanism: it verifies the shared secret without sending that secret in clear text. This is the FC-SP-specific AAA authentication method, rather than the normal switch administrative-login authentication method. When the configured external RADIUS service cannot be reached, Cisco MDS uses its local DH-CHAP credentials as the fallback authentication source, preserving authentication availability for the fabric. The RADIUS server group identifies the centralized AAA infrastructure, while the `aaa authentication dhchap` statement applies that group to the default FC-SP DH-CHAP authentication policy. Cisco documents DH-CHAP configuration under FC-SP and its use of RADIUS-based AAA in the Cisco MDS security configuration guidance. See the Cisco MDS 9000 Series NX-OS Security Configuration Guide and the DH-CHAP specification.

QUESTION NO: 35

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. server-compute

B. facility-manager

C. operations

D. server-equipment

E. server-profile

ANSWER: B E

Explanation:

The **server-profile** role is required because vNIC and vHBA configuration is part of service-profile administration in Cisco UCS Manager. Service profiles define the logical server identity and connectivity configuration applied to a server, including its virtual network interfaces and virtual host-bus adapters. Assigning this role lets the user create and modify those profile-level definitions without requiring unrestricted administrative access.

The **facility-manager** role is required for the fan-speed and power-redundancy settings. These controls are part of the UCS domain's facility and environmental management functions: they govern the power and cooling behavior of the physical infrastructure rather than an individual server's logical identity. Combining server-profile with facility-manager therefore grants the least-privilege role set that covers both requested configuration areas.

Cisco documents the available UCS Manager RBAC roles and their permissions in the Cisco UCS Manager User Management Guide, Release 3.1. Related service-profile and connectivity configuration concepts are documented in the Cisco UCS Manager GUI Configuration Guide, Release 3.1.

QUESTION NO: 36

In a Cisco SAN fabric, UCS fabric interconnects and blades are all monitored by Cisco DCNM. The requirement is for the Cisco UCS manager to be monitored by Cisco DCNM. Which two

configuration steps accomplish this goal? (Choose two.)

A. Configure NTP in the equipment policies.

B. Create an SNMP user using SHA and AES-128.

C. Configure high availability for the uplinks.

D. Add a locally authenticated read-only user.

E. Configure the DCNM IP address as a seed IP

ANSWER: A B

Explanation:

To enable Cisco DCNM to monitor a Cisco UCS Manager domain, UCS Manager must provide the secure management telemetry and consistent timestamps that DCNM expects. Creating an SNMP user with SHA authentication and AES-128 privacy enables secure SNMPv3 communication between DCNM and UCS Manager. DCNM uses the configured SNMP credentials to authenticate to the UCS Manager management endpoint and collect the inventory, fault, and operational information needed for monitoring. SHA and AES-128 provide message authentication and encryption, respectively, which is the recommended secure SNMP configuration for management-system integration.

Configuring NTP in the UCS Manager equipment policies synchronizes the UCS domain clock with an authoritative time source. Accurate, consistent timestamps are essential when DCNM correlates UCS Manager events, faults, and monitored status with data collected from the SAN fabric, fabric interconnects, and blades. Time synchronization also makes operational troubleshooting and event chronology reliable across the monitored environment. Cisco documents SNMP configuration and monitoring behavior for UCS domains in the Cisco UCS Manager Monitoring and Management Guide; DCNM deployment and device-management guidance is available in the Cisco DCNM configuration guides.

QUESTION NO: 37

Refer to the exhibit. SVR-A in EPG-A/BD-A fails to reach SVR-B in EPG-B/BD-B. Which two conditions should be verified to analyze the traffic between the two EPGs? (Choose two.)

A. A contract exists between EPG-A and EPG-B and allows the traffic.

B. Limit Endpoint Learning is disabled on the bridge domain for EPG-B.

C. SVR-B is learned as an endpoint in EPG-B from Node-112.

D. EPG-A and EPG-B are excluded from the preferred group.

E. The bridge domain for EPG-B floods ARP packets.

ANSWER: A C

Explanation:

Inter-EPG communication in Cisco ACI is controlled by policy. Therefore, a contract must be deployed between EPG-A and EPG-B, with the applicable provider/consumer relationship and a contract subject whose filters permit the required protocol, ports, and direction of the SVR-A-to-SVR-B flow. Confirming that the contract exists and allows the traffic validates the fundamental policy requirement for communication across EPGs in different bridge domains.

It is also necessary to verify that SVR-B is learned in EPG-B from Node-112. ACI uses endpoint learning to associate an endpoint's MAC and IP addresses with its EPG, bridge domain, interface, and leaf location. This endpoint record determines where the fabric forwards traffic for SVR-B. Confirming the expected learning location ensures that the destination address is classified into EPG-B and that the fabric has valid endpoint-location information for forwarding. Cisco documents the EPG contract model and its use for permitted communication in the Cisco APIC Layer 3 Networking Configuration Guide. Endpoint learning and endpoint-table verification are described in the Cisco APIC Basic Configuration Guide.

QUESTION NO: 38

An engineer is asked to use SaaS to manage Cisco computing.

Which two items are needed to accomplish the task? (Choose two.)

A. Intersight

B. UCS Central

C. Device/Claim ID

D. Node name/Serial Number

E. UCS Manager

ANSWER: A C

Explanation:

Cisco Intersight is Cisco's SaaS-based infrastructure lifecycle-management platform for Cisco compute environments. It provides cloud-hosted visibility, inventory, configuration, monitoring, and operations for supported Cisco UCS and HyperFlex resources. Therefore, an Intersight organization/account is the cloud service required when the stated management method is SaaS.

A device must also be securely claimed into that Intersight account before it can be managed. The claiming workflow uses the identifying information generated or displayed by the device connector, commonly described as the Device ID together with a claim code. This credential pair associates the physical Cisco compute target with the correct Intersight tenant and establishes the trusted management relationship. In practical deployment, the administrator obtains the device identifier and claim information from the target's Device Connector interface and enters it in Intersight during the claim process. Once claimed, the target appears in the Intersight inventory and can be managed through the SaaS portal.

Cisco describes Intersight as a cloud operations platform for Cisco infrastructure on its Cisco Intersight product page. The device onboarding and claim workflow is covered in the Cisco Intersight Getting Started Guide.

QUESTION NO: 39

An engineer needs to implement a solution that will provide telemetry of MDS 9000 switches in a SAN fabric. The solution should use Cisco Data Center Network Manager (DCNIM) with SAN insights feature. Which three steps are required to deploy the solution? (Choose Three.)

A. Select a target Fabric to be monitored.

B. Activate ENTERPRISE_PKG license on target switches

C. Select the target ports to be monitored for telemetry data

D. Configure name resolution between the devices

E. Select a target VSAN to be monitored

F. Activate SAN_ANALYTICS_PKG license on target switches

ANSWER: A C F

Explanation:

DCNM SAN Insights is deployed by defining the SAN scope that DCNM will collect and analyze, choosing the interfaces that generate the required flow telemetry, and ensuring that SAN Analytics is licensed on the participating MDS switches. “Select a target Fabric to be monitored.” establishes the fabric in which SAN Insights is enabled and provides the management context for telemetry collection and visualization. “Select the target ports to be monitored for telemetry data” is required because SAN Insights collects analytics from explicitly selected supported switch interfaces; this permits focused collection for the host and storage-facing paths of interest. “Activate SAN_ANALYTICS_PKG license on target switches” enables the MDS SAN Analytics capability that produces the telemetry data consumed by SAN Insights. With these elements in place, DCNM can collect and display metrics such as I/O operations, throughput, latency, and SCSI command-level behavior for the selected flows. Cisco’s SAN Insights configuration documentation describes selecting fabrics and ports as part of enabling analytics, while the MDS licensing documentation identifies the SAN Analytics package required for the feature. See the Cisco DCNM SAN Insights Configuration Guide and the Cisco MDS NX-OS Licensing Configuration Guide.

QUESTION NO: 40

Service degradation is reported on a VM that is deployed on a Cisco UCS blade server. The traffic from the vNIC is required to SPAN in both directions to a packet analyzer that is connected to UCS-A slot 2 port 12. Which two commands are needed to complete the configuration? (Choose two.)

A. UCS-A /eth-traffic-mon/fabric/eth-mon-session* # create dest-interface 2 12

B. UCS-A/eth-traffic-mon/fabric/eth-mon-session ' # create eth-mon-sessioaTdest-interface 2 12

C. UCS-A /org/service-profile/vnic/mon-src ' 8 set direction receive transmit

D. UCS-A /org/service-profile/vnic/mon-src* # set direction both

E. UCS-A/eth-traffic-rnon/fabnc/eth-mon-session n activate

ANSWER: A D

Explanation:

To send the vNIC traffic to the packet analyzer, the monitoring session needs a physical destination interface on the appropriate fabric and the vNIC monitoring source must be configured for bidirectional capture. `create dest-interface 2 12`, when entered in the Ethernet monitoring-session context for Fabric A, defines the analyzer’s attachment point as slot

2, port 12. This provides the egress destination for the copied packets. The source-monitor configuration must then use `set direction both` under the vNIC's monitoring-source context. The `both` direction captures traffic received by the vNIC and traffic transmitted by the vNIC, satisfying the requirement to analyze the complete conversation rather than just one traffic direction. These settings are part of Cisco UCS Manager Ethernet traffic monitoring, in which a monitoring session associates source interfaces such as vNICs with a destination interface connected to an external analyzer. Cisco documents Ethernet traffic monitoring configuration and CLI command contexts in the Cisco UCS Manager CLI Configuration Guide; the applicable UCS Manager documentation is also available from Cisco's installation and configuration guides page.

QUESTION NO: 41

An engineer must assign a device alias of Marketing to a Fibre Channel port that has a worldwide name of 21:00:00:00:11:11:22:22 on a Cisco MDS 9000 Series Switch. Which two commands accomplish this goal? (Choose two.)

A. vsan database

B. device-alias database

C. device-alias name Marketing pwwn 21:00:00:00:11:11:22:22

D. device-alias name Marketing wwn 21:00:00:00:11:11:22:22

E. device-alias rename wwn Marketing

ANSWER: B C

Explanation:

On Cisco MDS 9000 Series switches, a device alias is an administrative, user-friendly name associated with a storage device's port WWN (pWWN). Configuration begins by entering the device-alias database submode with `device-alias database`. This mode provides the context in which device-alias entries are created and maintained.

Within that database submode, `device-alias name Marketing pwwn 21:00:00:00:11:11:22:22` creates the required mapping. The `name` field defines the alias, Marketing, and the `pwwn` field identifies the specific Fibre Channel endpoint by its port worldwide name. A pWWN is the appropriate identifier because the requested alias applies to an individual Fibre Channel port rather than to a node as a whole. In a production fabric using Cisco Fabric Services distribution, the configured database is typically committed afterward with `device-alias commit` so the change is activated and distributed; however, the two commands that establish the requested alias definition are entering the device-alias database and adding the named pWWN entry. Cisco documents device-alias configuration and the `device-alias name ... pwwn ...` syntax in the Cisco MDS Device Alias configuration guide and the Cisco MDS NX-OS Command Reference.

QUESTION NO: 42

What is the impact of an EPLD upgrade on a Cisco MDS 9000 Series Switch?

A. The active supervisor traffic is disrupted.

B. The upgrade disrupts the management connectivity to the switch.

C. The standby supervisor module reloads multiple times.

D. The upgrade process disrupts only the module that is being upgraded.

ANSWER: D

Explanation:

The upgrade process disrupts only the module that is being upgraded. EPLD images contain firmware for programmable hardware devices, such as FPGAs and CPLDs, on an MDS module. Applying an EPLD image requires the affected hardware to be reprogrammed and the target module to be reset or reloaded so that it can begin operating with the new

programmable-logic firmware. Consequently, interfaces and functions supplied by that specific module are briefly unavailable during its upgrade window. This makes EPLD maintenance a module-scoped disruptive operation and requires normal change planning for any traffic or services carried by the target module. In a redundant MDS design, administrators can plan the sequence around available redundancy and module roles to limit the operational effect of that module reset. Cisco's MDS EPLD documentation describes EPLD upgrades as requiring reload/reset activity for the hardware being programmed and provides the supported procedures and platform-specific considerations. See the Cisco MDS 9000 Series EPLD Upgrade Guide and Cisco's MDS 9000 installation and configuration guides for the applicable release and hardware.

QUESTION NO: 43

An engineer configures multiple blade servers for OS installation. The requirements are to use the internal HDDs inside the Cisco UCS blade server and for the OS to remain operational even if one of the HDDs is removed for maintenance. Each blade has two internal HDDs. Which two actions must be taken to meet these requirements? (Choose two.)

A. Set Enforce vNIC/vHBA/iSCSI Name to disabled

B. Set the Local Disk Policy Mode to RAID 1.

C. Set the First Boot Order to be Local LUN.

D. Set Boot Policy Mode to UEFI

E. Set the Protect Configuration to enabled

ANSWER: B C

Explanation:

Setting the Local Disk Policy Mode to RAID 1 creates a mirrored virtual drive from the blade's two internal HDDs. RAID 1 writes equivalent data to both physical drives, so the operating system volume remains available when either one disk is removed or fails. The surviving disk continues servicing I/O, providing the required single-disk fault tolerance while the failed or removed drive can be replaced and the mirror rebuilt. In Cisco UCS, this local-disk behavior is defined in the service profile's local disk configuration policy and is applied to the server's supported storage controller.

Setting the First Boot Order to be Local LUN makes the internal RAID virtual drive the preferred OS boot target. A Local LUN is the logical disk presented by the controller, rather than either raw physical HDD individually. Therefore, after the RAID 1 logical drive is created, selecting Local LUN in the boot policy ensures that the blade boots its installed operating system from the mirrored internal storage. Together, the local disk policy supplies resilient storage and the boot-policy setting directs the server to use that resilient local boot volume. Cisco documents local disk and boot-policy configuration in the Cisco UCS Manager Configuration Management Guide and provides RAID concepts in its Cisco UCS B-Series Blade Servers documentation.

QUESTION NO: 44

Which two statements describe the process of performing an EPLD upgrade on a Cisco MDS 9000 Series Switch? (Choose two.)

A. The upgrade is performed from the standby supervisor module.

B. The upgrade process disrupts only the module that is being upgraded.

C. The upgrade is performed on the active supervisor.

D. If an upgrade is interrupted, the upgrade continues after a connection is restored.

E. Modules must be online to be upgraded.

ANSWER: A B

Explanation:

An EPLD upgrade updates the programmable-logic firmware used by hardware components in an MDS switch, rather than the NX-OS system image itself. In a switch with redundant supervisors, the EPLD upgrade procedure begins with the standby supervisor module. This sequencing preserves the active supervisor's control-plane role while the standby supervisor's programmable devices are updated, supporting the redundant-supervisor upgrade workflow described for MDS platforms.

The scope of disruption is limited to the hardware module undergoing the EPLD update. Applying new EPLD firmware requires the affected module to be reset or power-cycled so that its programmable logic loads the new image. Consequently, interfaces and services hosted by that module are temporarily unavailable during its upgrade, while modules not being upgraded remain operational. Administrators should therefore plan the activity around the traffic and services carried by each affected module and confirm EPLD status before and after the operation. Cisco documents EPLD upgrade behavior and procedures in the Cisco MDS 9000 Series NX-OS Software Upgrade and Downgrade Guide and related platform guidance in the Cisco MDS 9000 NX-OS documentation.

QUESTION NO: 45

Which two features are used to allow a nondisruptive upgrade for the ISSU process? (Choose two.)

A. connectivity management processor

B. nonstop forwarding

C. graceful restart

D. stateful switchover

E. advanced services license

ANSWER: B D

Explanation:

In-Service Software Upgrade (ISSU) relies on **nonstop forwarding** and **stateful switchover** to preserve traffic forwarding while control-plane software is upgraded. Stateful switchover maintains synchronized control-plane and protocol state between the active and standby supervisor or route processor. If the active supervisor changes during the upgrade, the standby supervisor can assume the active role with the operational state already available, minimizing the effect of the switchover on network operation. Nonstop forwarding keeps the data plane forwarding packets using the existing forwarding information while the control plane is restarting or transitioning. This allows line cards and forwarding hardware to continue sending traffic rather than waiting for routing protocols and the control plane to fully reconverge. Together, stateful switchover provides resilient supervisor continuity and nonstop forwarding preserves packet forwarding during that continuity event, which are the high-availability foundations that enable an ISSU workflow to be nondisruptive when the platform, software release path, and configuration meet ISSU requirements. Cisco documents these ISSU high-availability dependencies in its Nexus high-availability documentation and describes the related capabilities in the Cisco IOS high-availability command reference.

QUESTION NO: 46

A network engineer is adding a Cisco HyperFlex data platform to the Cisco Intersight management portal.

Which two components are required for Intersight to claim the Cisco HyperFlex data platform? (Choose two.)

A. device FQDN

B. device claim code

C. device public IP address

D. device ID

E. device serial number

ANSWER: B D

Explanation:

Cisco Intersight uses a unique device identity together with a time-limited authorization value to establish ownership of a Cisco HyperFlex cluster. The **device ID** identifies the specific HyperFlex system that is being registered. This identifier is displayed by the HyperFlex management interface when Intersight services are enabled or when the cluster is prepared for claiming.

The **device claim code** is the corresponding one-time code used to authorize the claim operation. Entering the device ID and claim code in Cisco Intersight associates the HyperFlex platform with the appropriate Intersight account and enables cloud-based inventory, monitoring, lifecycle management, and advisory capabilities. The claim code is designed to provide proof that the administrator performing the registration has authorized access to the target HyperFlex system, rather than relying solely on an identifier. Cisco's HyperFlex documentation describes obtaining the device ID and claim code from the HyperFlex Connect interface for Intersight registration. See the Cisco HyperFlex Data Platform Installation Guide and the Cisco Intersight target-claim documentation.

QUESTION NO: 47

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

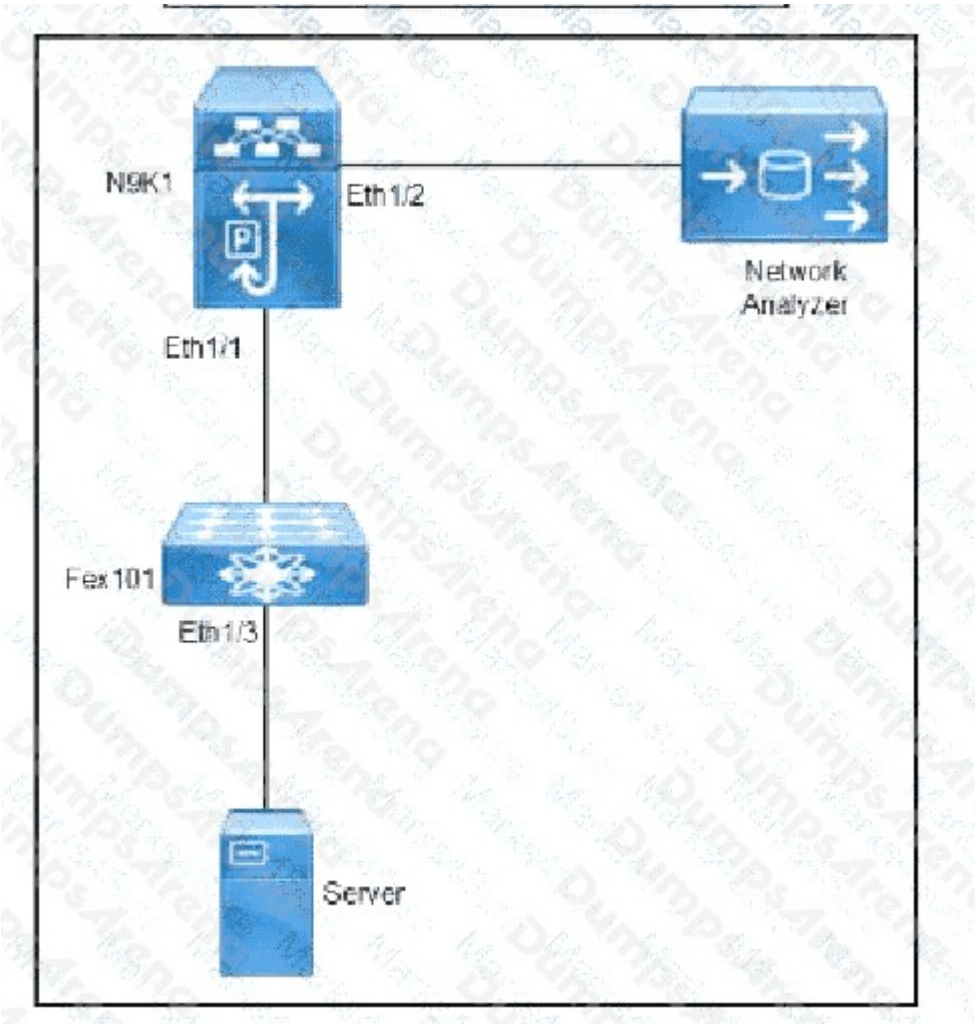
D. 0000.0C9F.F010

ANSWER: B

Explanation:

`0000.0C9F.F00A` is the default HSRP version 2 virtual MAC address for group 10. For IPv4 HSRP version 2, Cisco derives the virtual MAC address using the fixed prefix `0000.0C9F.F`, followed by the HSRP group number represented as a three-digit hexadecimal value. HSRP version 2 supports group numbers from 0 through 4095, which requires three hexadecimal digits and distinguishes its virtual MAC addressing format from HSRP version 1. Group 10 expressed in hexadecimal is `00A`. Appending that value to the HSRPv2 prefix produces `0000.0C9F.F00A`. This MAC address is shared by the active HSRP router for the group and is used as the destination MAC identity associated with the configured virtual IPv4 gateway. Hosts use the virtual IP address as their default gateway, while HSRP provides first-hop redundancy by allowing the active router to answer for that virtual gateway identity. Cisco documents the HSRPv2 virtual MAC format as `0000.0C9F.Fxxx`, where `xxx` is the group number in hexadecimal. See Cisco's HSRP configuration guide and the HSRP overview.

QUESTION NO: 48



Refer to the exhibit. An engineer must configure SPAN to monitor the traffic sent to the server. The copied traffic must be sent to the network analyzer. Which configuration on N9K1 accomplishes these goals?

A. source interface eth1/3 tx

B. source interface eth101/1/3 tx

C. source interface eth1/1 tx

D. source interface eth1/1 rx

ANSWER: C

Explanation:

`source interface eth1/1 tx` is correct because the server-bound traffic exits N9K1 through Ethernet1/1 on its way toward the downstream switch and attached server. In an NX-OS SPAN session, the configured source interface identifies the point at which packets are copied, and the direction qualifier selects whether the switch mirrors ingress, egress, or both directions of traffic at that interface. The `tx` qualifier mirrors frames transmitted by N9K1 out of Ethernet1/1. Therefore, it captures the traffic that N9K1 sends toward the server, satisfying the stated monitoring requirement. The SPAN session destination interface can then forward those copied frames to the directly attached network analyzer for inspection; SPAN does not reroute the production traffic itself. Cisco NX-OS documents that SPAN source interfaces support directional monitoring and defines transmit traffic as traffic sent from the switch out of the source interface. This makes the egress direction on the N9K1 interface facing the server path the appropriate observation point. See Cisco's Nexus 9000 Series NX-OS SPAN configuration guide and the NX-OS SPAN command reference.

QUESTION NO: 49

A POAP-enabled Cisco Nexus switch will not enter POAP mode.

Which two conditions should be verified? (Choose two.)

A. No startup configuration is available.

B. The switch is in bootup process.

C. The license file is missing on the switch.

D. Bootflash must contain a special directory named POAP with poap.py file.

E. No Cisco NX-OS image is present on the bootflash.

ANSWER: A B

Explanation:

POAP is evaluated as part of the Cisco Nexus switch startup sequence. Therefore, verify that the switch is in bootup process, because POAP is an automatic, boot-time provisioning mechanism rather than a feature that is normally initiated after the switch has completed startup. The switch must also have no startup configuration available. During boot, NX-OS checks for an existing startup configuration; when none is found, a POAP-capable switch can enter POAP mode and begin its discovery/provisioning workflow. This behavior is intended to enable initial, zero-touch deployment of an unconfigured Nexus switch.

In practical troubleshooting, confirm that the device has actually rebooted through the normal NX-OS startup path and inspect whether a startup configuration exists in the active configuration location. If a configuration is present, the switch treats itself as already provisioned and does not invoke the automatic POAP workflow. Cisco documents that POAP is triggered during boot when a startup configuration is absent, and describes POAP as an automated deployment process for initially provisioning Nexus devices. See Cisco's Nexus 9000 POAP configuration guide and the Cisco NX-OS fundamentals documentation.

QUESTION NO: 50

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?

```
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

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

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

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
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: B

Explanation:

The configuration represented by “Option B” is correct because it implements a switched-traffic capture for VLAN 10 and constrains the capture operation to the analyzer’s 30-second acceptance limit. Capturing the VLAN as traffic is forwarded through the switch provides visibility into the frames associated with the reported delay and loss, allowing the analyzer to inspect the relevant packet flow rather than an unrelated interface or control-plane-only view. The bounded capture period is important operationally: it ensures that the receiving analyzer can ingest and process the capture without rejecting it for exceeding its supported duration. Cisco NX-OS SPAN functionality is designed to copy selected traffic from a source, including VLAN-based sources, to a destination for analysis. Cisco documentation also describes packet-capture controls that limit capture scope and duration, which helps avoid unnecessary collection while preserving evidence of a transient forwarding problem. This approach provides a focused, analyzer-compatible capture of the affected VLAN traffic. See Cisco’s NX-OS SPAN configuration documentation and the NX-OS Ethanalyzer documentation.

QUESTION NO: 51

Which two statements describe modifying Cisco UCS user accounts? (Choose two.)

A. Disabling a user account maintains all of the data in the Cisco UCS Fabric Interconnect.

B. The admin account is used only to log on by using SSH.

C. The password of the user account must contain a minimum of 10 characters.

D. Local user accounts override the same account on a remote authentication server, such as TACACS, RADIUS, or LDAP.

E. The password of the user account expires in 30 days.

ANSWER: A D

Explanation:

Disabling a user account maintains all of the data in the Cisco UCS Fabric Interconnect. Disabling prevents that identity from logging in, while retaining the account definition and associated configuration information. This approach is appropriate when access must be temporarily suspended without losing the account's details, role assignments, or other stored user information. The account can subsequently be enabled again if access is required.

Local user accounts override the same account on a remote authentication server, such as TACACS, RADIUS, or LDAP. Cisco UCS Manager checks its locally configured user database first when a username is presented. Therefore, if the same username exists both locally and in a configured remote authentication service, the local account takes precedence. This behavior lets administrators preserve local administrative access, including during remote AAA service outages or when remote-directory settings are being changed. It also means that local usernames should be managed carefully to avoid unintentionally superseding a remote user identity. Cisco documents local and remote authentication behavior in its Cisco UCS Manager Administration Management Guide and the UCS documentation collection is available from Cisco UCS Manager installation and configuration guides.

QUESTION NO: 52 - (DRAG DROP)

Drag and drop the hyperflex characteristics from the left onto the appropriate. HyperFlex Node types on the right. Not all options are used.



ANSWER:

Explanation:

Cisco HyperFlex hybrid nodes use a tiered disk design intended to balance usable capacity and cost. Solid-state drives provide the caching tier, where frequently accessed data and write activity can be handled with low latency. Hard-disk drives provide the capacity tier, allowing the cluster to deliver substantially more storage capacity at a lower cost per gigabyte. This architecture makes a hybrid node appropriate for deployments in which capacity efficiency and budget are important and where the applications do not require consistently high all-flash performance.

An all-flash HyperFlex node is designed for workloads that need more predictable, repeatable performance across a larger active working set. Its capacity tier uses SSDs rather than HDDs, avoiding the mechanical-disk latency that can limit sustained storage operations. The caching function is provided by SSD or NVMe storage, with NVMe offering a high-performance media and interface option on supported platforms. Combining flash-based caching with a flash-based capacity layer enables the node to sustain lower latency and higher I/O performance for demanding enterprise applications. Cisco describes HyperFlex as a hyperconverged platform that combines compute and distributed storage while supporting configurations selected for their capacity and performance requirements; see the Cisco HyperFlex Data Platform architecture overview.

Therefore, the correct placement groups the SSD-cache, HDD-capacity, capacity-oriented recommendation together for Hybrid Node, while grouping SSD-or-NVMe caching, SSD capacity, and the sustained high-performance recommendation together for All Flash Node.

QUESTION NO: 53

```
vlan 10
 fcoe vsan 10

interface ethernet 1/3
 spanning-tree port type edge trunk
 switchport mode trunk
 switchport trunk allowed vlan 1, 10
```

Refer to the exhibit. An engineer must configure FCoE on a Cisco Nexus switch. Which two command sets complete this configuration? (Choose two.)

A. vsan database
vsan 10 interface vfc 10

B. interface vfc 10 bind interface ethernet 1/3
no shut

C. vsan database
vsan 10 interface ethernet 1/3

D. interface vfc 10 bind interface fcoe-vfc 1/3
no shut

E. interface fcoe 10 bind interface ethernet 1/3 no shut

ANSWER: A B

Explanation:

The command set `interface vfc 10 bind interface ethernet 1/3 no shut` creates virtual Fibre Channel interface 10, associates it with the Ethernet interface that carries the lossless FCoE traffic, and administratively enables the vFC interface. A vFC interface is the Fibre Channel-facing logical interface used by an FCoE-capable Nexus switch; it must be bound to its underlying Ethernet interface before it can forward FCoE traffic.

The command set `vsan database vsan 10 interface vfc 10` then places that vFC interface into VSAN 10. VSAN membership is required because Fibre Channel services, including fabric segmentation, zoning scope, and Fibre Channel

forwarding, operate within a VSAN. Together, these command sets establish the required relationship from Ethernet transport to the vFC interface and from the vFC interface to the intended SAN fabric.

Cisco documents that a vFC interface is configured by binding it to an Ethernet interface, and that vFC interfaces are assigned to VSANs through the VSAN database. See the Cisco Nexus FCoE configuration guide and Cisco NX-OS SAN Configuration Guide.

QUESTION NO: 54

Refer to the exhibit An engineer must restrict users assigned to the sangroup role on the Cisco MDS 9000 Series Switch from issuing commands

on VSANs 15 to 20 Which command must the engineer run to achieve

this objective?

A. vsan policy deny vsan 15-20

B. no permit vsan 15-20

C. permit vsan 15-20

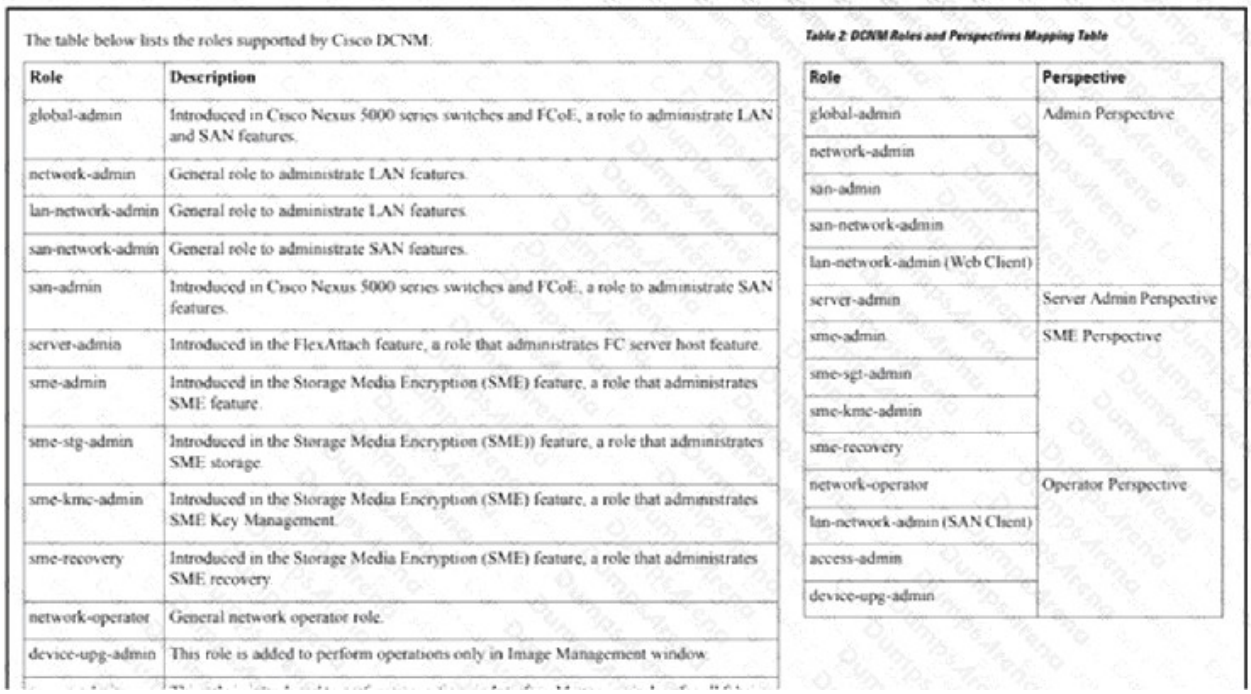
D. no vsan policy deny

ANSWER: A

Explanation:

The `vsan policy deny` configuration followed by `vsan 15-20` is correct because Cisco MDS role-based access control supports VSAN-based policies in addition to command-based rules. A deny VSAN policy explicitly identifies the VSAN range to which members of the role must not have access. Applying the policy to VSANs 15 through 20 prevents users assigned to the `sangroup` role from executing role-authorized operations against those VSANs, while leaving the role's permissions for VSANs outside that range unchanged. The configuration is entered in the role configuration context: first create or enter the deny VSAN-policy submode with `vsan policy deny`, then specify the protected range using `vsan 15-20`. This approach is appropriate when access must be withheld only for a defined set of VSANs rather than broadly removing command privileges from the role. Cisco documents VSAN policies as part of MDS RBAC and describes using permit or deny policies to control a role's VSAN scope. See the Cisco MDS NX-OS RBAC configuration guide and the Cisco MDS NX-OS command reference.

QUESTION NO: 55



The table below lists the roles supported by Cisco DCNM.

Role	Description
global-admin	Introduced in Cisco Nexus 5000 series switches and FCoE, a role to administrate LAN and SAN features.
network-admin	General role to administrate LAN features.
lan-network-admin	General role to administrate LAN features.
san-network-admin	General role to administrate SAN features.
san-admin	Introduced in Cisco Nexus 5000 series switches and FCoE, a role to administrate SAN features.
server-admin	Introduced in the FlexAttach feature, a role that administrates FC server host feature.
sme-admin	Introduced in the Storage Media Encryption (SME) feature, a role that administrates SME feature.
sme-stg-admin	Introduced in the Storage Media Encryption (SME)) feature, a role that administrates SME storage.
sme-kmc-admin	Introduced in the Storage Media Encryption (SME) feature, a role that administrates SME Key Management.
sme-recovery	Introduced in the Storage Media Encryption (SME) feature, a role that administrates SME recovery.
network-operator	General network operator role.
device-upg-admin	This role is added to perform operations only in Image Management window.

Table 2: DCNM Roles and Perspectives Mapping Table

Role	Perspective
global-admin	Admin Perspective
network-admin	
san-admin	
san-network-admin	
lan-network-admin (Web Client)	
server-admin	Server Admin Perspective
sme-admin	SME Perspective
sme-sgt-admin	
sme-kmc-admin	
sme-recovery	
network-operator	Operator Perspective
lan-network-admin (SAN Client)	
access-admin	
device-upg-admin	

Refer to the exhibit. An engineer is configuring Cisco Data Center Network Manager to automate the provisioning of Cisco Nexus 9000 Series Switches. The engineer must configure user access for network engineers to permit device operations in Interface Manager. The solution must hide Admin and Config Menu items in Interface Manager Which two roles must be assigned to the network engineers to meet these requirements? (Choose two.)

A. San-admin

B. network-operator

C. network-admin

D. access-admin

E. global-admin

ANSWER: B D

Explanation:

The required role combination is **network-operator** and **access-admin**. In Cisco DCNM, the network-operator role grants operational access needed to work with managed network devices, including performing the device-oriented activities exposed through Interface Manager. This provides the network engineers with the appropriate operational scope without granting full administrative control over DCNM or its configuration features.

The access-admin role is used with the network-operator role to apply the Interface Manager access model required by the question. This combination presents the operational functions that engineers need while suppressing the Admin and Config menu items in Interface Manager. Consequently, the users can carry out permitted device operations but do not receive the broader configuration and administration menus that would expose management-level capabilities.

Cisco documents DCNM/Nexus Dashboard Fabric Controller user roles and role-based access control as the mechanism for assigning permissions to users and restricting the application functions they can access. For additional role and access-control context, see Cisco's Managing User Accounts documentation and the Nexus Dashboard Fabric Controller documentation portal.

QUESTION NO: 56

An engineer must create a configuration checkpoint on a Cisco Nexus 9000 Series Switch and restore that checkpoint if a subsequent change causes an issue.

Which two commands are used to perform these tasks? (Choose two.)

A. checkpoint prechange

B. rollback running-config checkpoint prechange

C. copy running-config checkpoint prechange

D. restore startup-config prechange

E. commit checkpoint prechange

ANSWER: A B

Explanation:

The `checkpoint` command saves a named checkpoint of the current running configuration. For example, `checkpoint prechange` creates a checkpoint named `prechange`.

The `rollback running-config checkpoint` command restores the saved checkpoint into the running configuration. For example, `rollback running-config checkpoint prechange` rolls the running configuration back to that saved state. Cisco documents checkpoints and rollback operations in the Cisco Nexus 9000 Series NX-OS rollback configuration guide.

QUESTION NO: 57

An engineer is duplicating an existing Cisco UCS setup at a new site.

What are two characteristics of a logical configuration backup of a Cisco UCS Manager database? (Choose two.)

A. contains the configured organizations and locales

B. contains the VLAN and VSAN configurations

C. contains the AAA and RBAC configurations

D. contains all of the configurations

E. contains a file with an extension.tgz that stores all of the configurations

ANSWER: A B

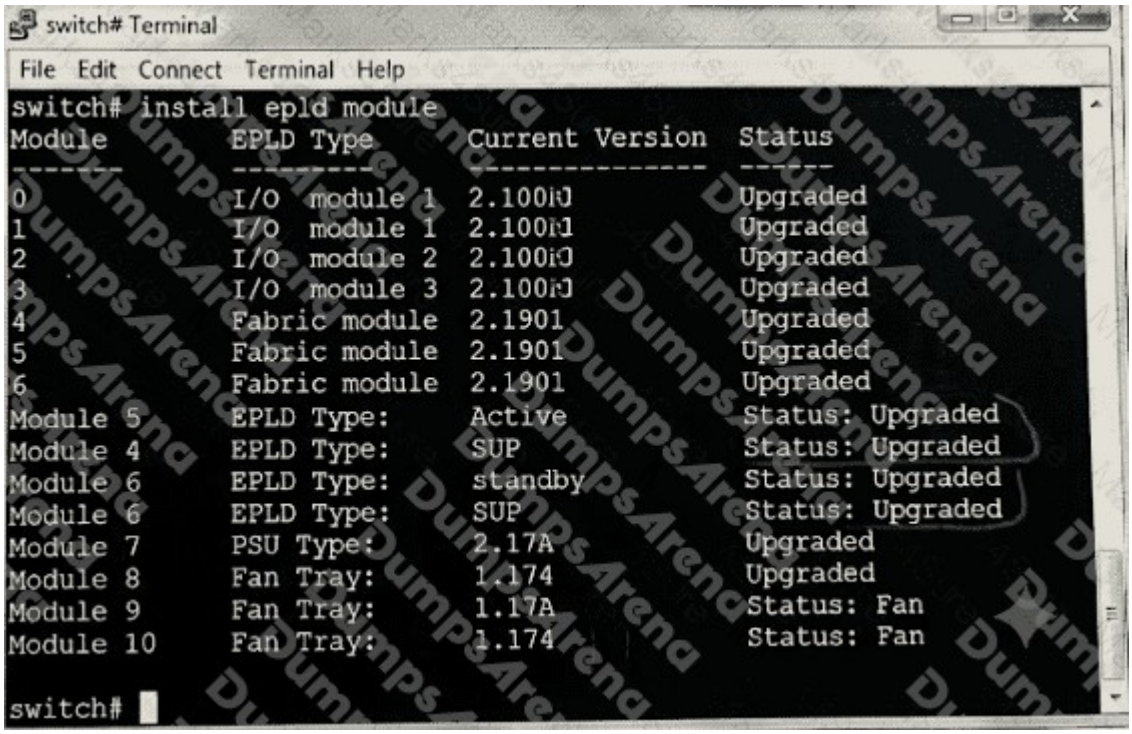
Explanation:

A logical configuration backup is intended to preserve the portable, logical portions of a Cisco UCS Manager deployment so that the configuration can be recreated or applied in another UCS domain. It includes the configured organizational hierarchy and locales. Organizations provide the hierarchical structure in which service profiles, policies, pools, and templates are organized; locales define the organizational scopes available for role-based access assignments. Preserving these objects maintains the logical structure needed when reproducing the environment.

The backup also includes VLAN and VSAN configurations. These define the Ethernet and Fibre Channel network constructs that policies and service profiles reference, so retaining them is necessary for the duplicated configuration to preserve its intended connectivity definitions. Cisco distinguishes this XML-based logical configuration backup from a full-state backup, which captures the complete UCS Manager state in a compressed archive and is generally used for restoring the same UCS domain. Cisco documents the available backup types and their uses in the Cisco UCS Manager Administration Guide. Additional UCS Manager configuration and backup documentation is available through Cisco UCS Manager Maintenance Guides.

QUESTION NO: 58

Refer to the exhibit.



Which Cisco MDS 9700 Series Switch module EPLDs are upgraded?

A. PSU

B. fabric

C. standby supervisor

D. fan

ANSWER: C

Explanation:

The upgraded module is the **standby supervisor**. The exhibit identifies the affected hardware as a SUP module and shows that its EPLD is upgradable and that the EPLD upgrade completes successfully. In a Cisco MDS 9700 director with redundant supervisors, the supervisor that is not currently active is the standby supervisor. EPLD maintenance is performed against this standby SUP so that the active supervisor continues to provide control-plane operation while the standby module's programmable logic is updated. The module status and successful completion shown in the output therefore identify a supervisor EPLD upgrade, specifically on the standby supervisor rather than on a power supply, fan, or fabric module. Cisco documents EPLD upgrades as a hardware-programmable-logic maintenance procedure and provides platform-specific procedures and impact information in the NX-OS upgrade documentation. See the Cisco MDS 9000 NX-OS Software Upgrade and Downgrade Guide and the Cisco MDS 9700 Series documentation overview for MDS 9700 hardware and software maintenance context.

QUESTION NO: 59

Which action must be performed before renumbering a Cisco UCS chassis?

A. Run the shut and no shut command on the connected ports.

B. Move the chassis to new ports on fabric interconnect.

C. Decommission the chassis.

D. Re-acknowledge the chassis.

ANSWER: C

Explanation:

Decommission the chassis. In Cisco UCS Manager, a chassis identifier is part of the chassis's managed inventory and its association with server slots, I/O modules, and fabric connectivity. Cisco requires the chassis to be decommissioned before its chassis ID can be changed, ensuring that UCS Manager cleanly removes the current managed configuration before the chassis is discovered and managed under the new identifier. After the renumbering operation, the chassis can be recommissioned so that UCS Manager rediscovers it and restores it to service. This workflow protects the consistency of the UCS domain inventory and avoids retaining stale chassis-specific state while the identifier changes. Cisco's UCS Manager chassis-management documentation describes decommissioning as the required preliminary action when changing a chassis ID. See the Cisco UCS Manager GUI Configuration Guide and the Cisco UCS Manager CLI Configuration Guide for chassis administration procedures.

QUESTION NO: 60

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 vlan 100
ip address 10.1.1.2 255.255.255.0
hsrp version 2
hsrp 1000
priority 255
preempt
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:

The command set beginning with `feature hsrp`, configuring `interface vlan 100`, and assigning HSRP group 1000 a priority of 255 with preemption enabled is correct. On Cisco Nexus 7000 switches, HSRP must first be enabled with the `feature hsrp` command. The SVI is then configured with its unique Layer 3 address, while the shared default-gateway address is configured in the HSRP group. HSRP version 2 supports group numbers beyond 255, so it is required for group 1000. A priority of 255 is the highest permitted HSRP priority and ensures that this switch has the preferred election value. The `preempt` command is also necessary so that, after the switch recovers or becomes fully operational, it takes over the active HSRP role rather than allowing a lower-priority peer that is currently active to remain active. Cisco documents HSRP configuration, including priority and preemption behavior, in its Nexus 7000 HSRP configuration guide and describes the related commands in the Nexus 7000 NX-OS command reference.

QUESTION NO: 61

A network engineer plans to upgrade the firmware of a Cisco UCS B-Series chassis by using the Auto Install feature. Which component is upgraded during the infrastructure firmware upgrade stage?

A. Adapter

B. Cisco IMC

C. I/O Modules

D. BIOS

ANSWER: C

Explanation:

I/O Modules are upgraded during the infrastructure firmware upgrade stage of Cisco UCS B-Series Auto Install. In a B-Series chassis, the I/O Modules provide the chassis-level connectivity between installed blade servers and the fabric interconnects. Because they are shared infrastructure components rather than server-specific components, their firmware is handled in the infrastructure portion of the automated upgrade workflow. Cisco UCS Manager uses the firmware package and upgrade sequence to bring these chassis connectivity modules to the target release before or as part of preparing the chassis infrastructure for subsequent server-related firmware actions. This staging helps maintain compatible firmware across the chassis data path and UCS domain. The Auto Install process distinguishes shared chassis infrastructure firmware from firmware associated with individual blade-server hardware, allowing the administrator to apply the appropriate package and activation behavior for each upgrade stage. Cisco's UCS firmware-management documentation describes the infrastructure firmware workflow and identifies I/O Module firmware as part of the chassis infrastructure scope. See the Cisco UCS Firmware Management Guide and the Cisco UCS B-Series overview.

QUESTION NO: 62

An administrator is implementing DCNM so that events are triggered when monitored traffic exceeds the configured percent utilization threshold. The requirement is to configure a maximum limit of

39913690 bytes that applies directly to the statistics collected as a ratio of the total link capacity. Which DCNM performance monitoring configuration parameter must be implemented to achieve this result?

A. Absolute Values

B. Baseline

C. Util%

D. Per Port Monitoring

ANSWER: C

Explanation:

Util% is the appropriate DCNM performance-monitoring parameter because it evaluates monitored interface traffic relative to the interface's total available link capacity. DCNM derives utilization from the traffic statistics it collects, such as byte counters over the collection interval, and expresses the result as a percentage of the port bandwidth. An event can therefore be generated when that calculated percentage exceeds the configured utilization threshold. The stated byte value is part of the underlying traffic-statistics measurement used in the calculation; the relevant thresholding behavior is the comparison of that measured traffic rate with the link's capacity, rather than merely recording a raw counter value. This makes utilization-based monitoring suitable where equivalent traffic levels should be interpreted differently on links with different speeds—for example, the same traffic rate represents a much higher load on a lower-capacity interface. Cisco documents DCNM/Nexus Dashboard Fabric Controller monitoring and management capabilities in its Nexus Dashboard Fabric Controller product documentation and publishes related configuration material through the Cisco Nexus Dashboard Fabric Controller support portal.

QUESTION NO: 63

```
NEXUS(config)# radius-server timeout 15
NEXUS(config)# ip radius source-interface Vlan100
NEXUS(config)# radius-server host 10.1.50.31 key 7 "Fewhg1234"
authentication accounting
NEXUS(config)# aaa group server radius RADIUS
NEXUS(config-radius)# server 10.1.50.31
NEXUS(config-radius)# deadtime 10
NEXUS(config-radius)# use-vrf management
NEXUS(config-radius)# source-interface mgmt 0
NEXUS(config)# aaa authentication login default group RADIUS
```

Refer to the exhibit. What is the result of implementing the configuration?

A. The RADIUS server timeout value is 15 milliseconds.

B. RADIUS traffic is sourced from the VLAN 200 interface.

C. Users specify the RADIUS server when they log in.

D. Only the RADIUS server is used for authentication.

ANSWER: D

Explanation:

Only the RADIUS server is used for authentication. In Cisco NX-OS AAA, the `aaa authentication login default group radius` method list applies RADIUS to default login authentication. Because the method list specifies only the RADIUS server group, no secondary local authentication method is configured as a fallback. Consequently, login requests are sent to the configured RADIUS server group and authentication depends on the availability of that external RADIUS service. This behavior is important operationally: if administrators require access when the RADIUS infrastructure cannot be reached, they must explicitly include an additional method, such as `local`, in the authentication method list. The configured default login method list is used for standard device login sessions unless a different named method list is applied to the relevant access service. Cisco documents AAA method lists and RADIUS server groups in its NX-OS security configuration guidance: Cisco Nexus 9000 Series NX-OS AAA configuration guide and Cisco Nexus 9000 Series NX-OS RADIUS configuration guide.

QUESTION NO: 64

A new employee must be granted access to add VLANs into an existing Cisco UCS Manager and configure NTP synchronization with date and time zone settings. Which two privileges must be granted to the employee to complete the task? (Choose two.)

A. Ext LAN Security (ext-lan-security)

B. Service Profile Network Policy (ls-network-policy)

C. Ext LAN Policy (ext-lan-policy)

D. Service Profile Compute (ls-compute)

E. Service Profile Config (ls-config)

ANSWER: A C

Explanation:

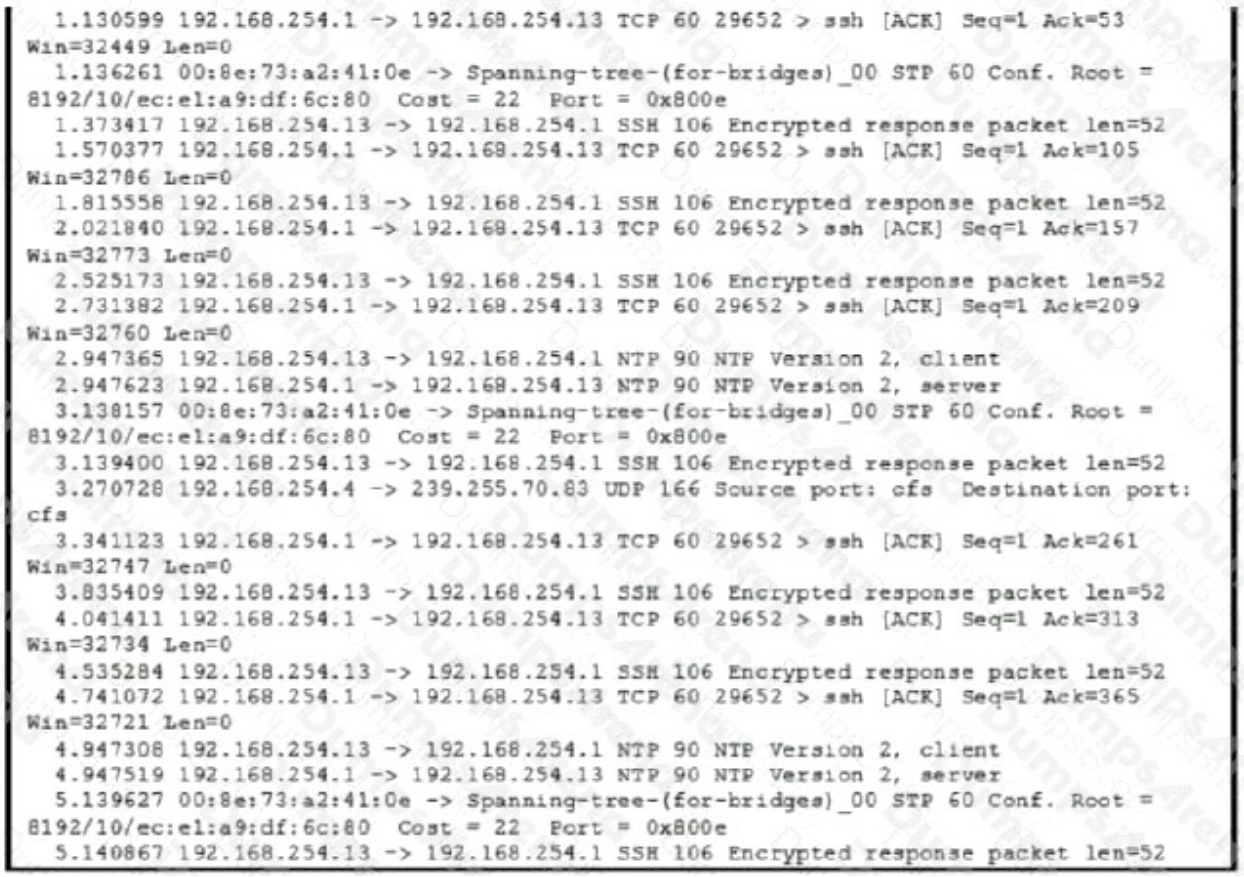
Ext LAN Security (ext-lan-security) is required because Cisco UCS Manager associates this privilege with administration of external LAN security-related global settings, including configuration of NTP providers and the system date, time, and time zone. Granting this privilege allows the employee to establish the time synchronization configuration required for consistent timestamps, authentication-related operations, logging, and scheduled UCS domain functions.

Ext LAN Policy (ext-lan-policy) is required for the VLAN portion of the task. VLANs are fabric-level LAN configuration objects in Cisco UCS Manager rather than settings owned by an individual service profile. This privilege grants access to configure external LAN policy elements, including the VLAN definitions and associated LAN policy configuration used by the fabric interconnects. Together, these privileges provide the distinct permissions needed to manage global VLAN configuration and the UCS domain's NTP/date/time-zone configuration without granting unrelated service-profile administration rights.

For Cisco UCS Manager configuration and RBAC guidance, consult Cisco's UCS Manager installation and configuration guides and the UCS Manager user guides.

QUESTION NO: 65

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 unicast address

B. IPv4 multicast address

C. IPv4 gateway address

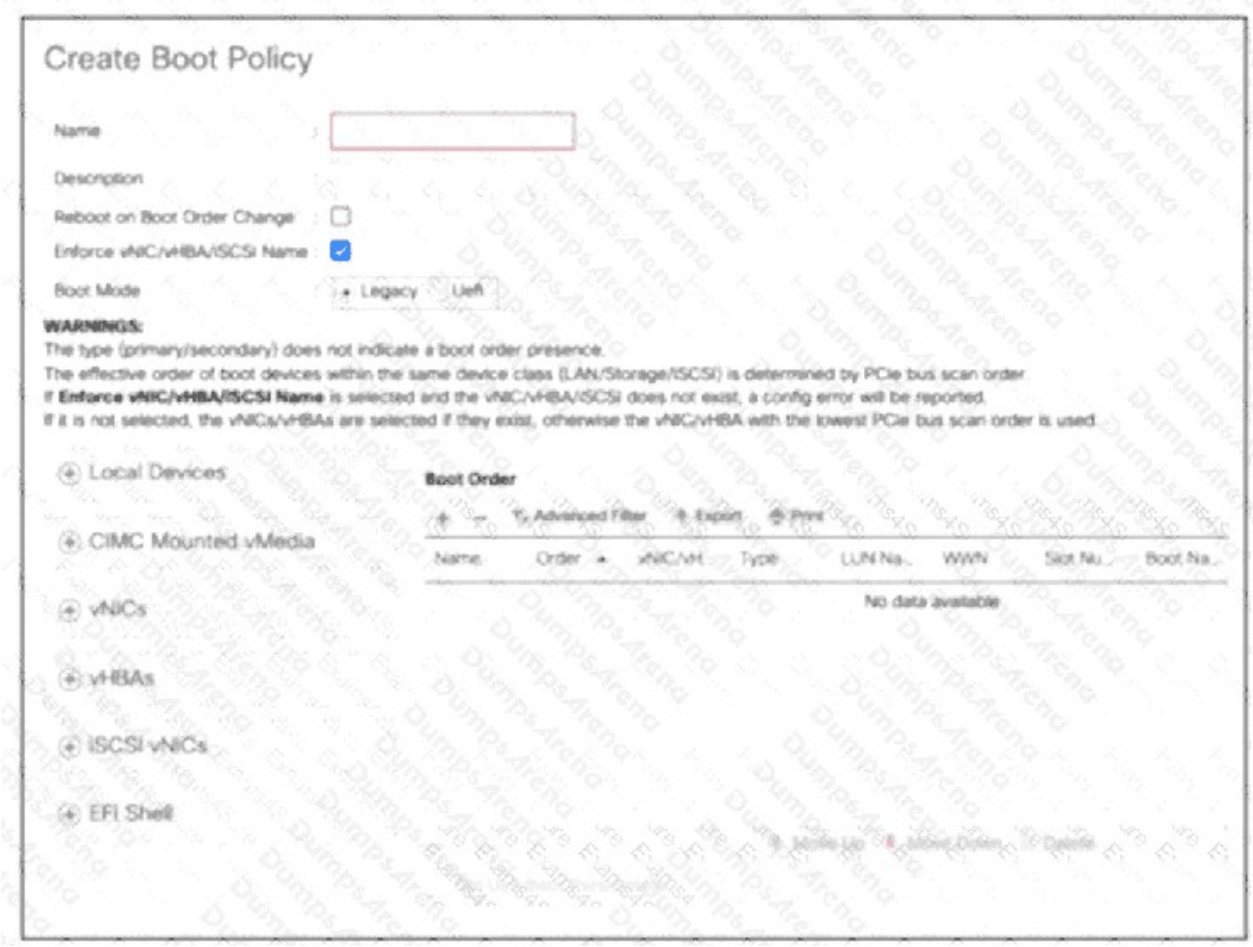
D. IPv4 anycast address

ANSWER: B

Explanation:

IPv4 multicast address is correct because Cisco Fabric Services (CFS) distributes configuration information and state among Cisco Nexus switches by sending messages to a multicast group rather than targeting one individual peer. This distribution model lets all participating switches in the CFS scope receive the same update efficiently, which is essential for synchronizing features such as device-alias databases, NTP configuration, VLAN information, and other CFS-capable application data. CFS operates as an internal fabric distribution service and uses multicast transport for its control and synchronization traffic; it does not require a gateway-oriented addressing model to identify participating fabric members. Cisco documents CFS as a mechanism for distributing configuration changes across switches in a fabric and describes its multicast-based communication behavior in the NX-OS system-management documentation. See the Cisco Nexus 9000 Series NX-OS Cisco Fabric Services configuration guide and the Cisco NX-OS Cisco Fabric Services documentation.

QUESTION NO: 66



Refer to the exhibit. An engineer installs a new Cisco UCS B-Series server. The requirement is to install an operating system on a LUN that is provided by the Fibre Channel storage array. Which two steps must be configured in the boot policy to accomplish this goal? (Choose two)

Diss We Enforce vNIC/VHBA/iSCSI Name.

A. Add Local CD/DVD and SAN Boot.

B. Configure a SAN Boot Target with a WWPN.

C. Set Boot Order to External USB and LAN Boot Create .

D. SCSI VNICs and add Local LUN.

ANSWER: A B

Explanation:

“Add Local CD/DVD and SAN Boot.” is required because the boot policy needs a SAN boot device so that the service profile can direct the blade to boot from Fibre Channel-attached storage. A local CD/DVD boot entry provides installation media precedence, allowing the server to start the operating-system installer while retaining the SAN LUN as the intended installation and subsequent boot destination. The boot order can then be adjusted as appropriate after installation.

“Configure a SAN Boot Target with a WWPN.” is also required because a SAN boot entry must identify the storage target available through the server’s vHBA. The target WWPN identifies the Fibre Channel storage port presented by the array, and the SAN boot configuration associates that target with the LUN used as the boot device. Cisco UCS applies this policy through the service profile, allowing the blade’s vHBA to discover and boot from the specified Fibre Channel LUN during server startup.

Cisco documents SAN boot configuration as creating a SAN boot device and specifying its target WWPN and LUN details in the boot policy. See Cisco’s Configure Boot from SAN on Cisco UCS guide and the Cisco UCS Manager GUI Configuration Guide.

QUESTION NO: 67

Which two statements describe EVPN route types used in a VXLAN EVPN fabric? (Choose two.)

A. Type 2 routes advertise MAC and optional IP address bindings.

B. Type 3 routes advertise IP prefixes for routed traffic.

C. Type 5 routes advertise IP prefixes for routed traffic.

D. Type 1 routes advertise default gateway MAC addresses.

E. Type 4 routes advertise ARP suppression entries.

ANSWER: A C

Explanation:

An EVPN type 2 route is a MAC/IP Advertisement route. It distributes endpoint MAC reachability and can include the endpoint IP address, which enables ARP and neighbor-discovery suppression functions in supported VXLAN EVPN deployments.

An EVPN type 5 route is an IP Prefix route. It advertises IP-prefix reachability for routed traffic between VTEPs. These route types provide distinct endpoint and prefix advertisement functions in EVPN. Their definitions are specified in RFC 7432 and RFC 9136.

QUESTION NO: 68

An engineer must remove the running configuration of a Cisco Nexus 9000 Series Switch and replace it with a backup file without restarting the switch. The backup file is stored on an SFTP server that is reachable via the mgmt0 interface Which set of commands accomplishes this task?

A. copy sftp://admin@192.168.10.100/backup.cfg startup-config vrf management copy startup-config running-config

B. configure replace sftp://admin@192.168.10.100/backup.cfg vrf management copy running-config bootflash:startup-config

C. copy sftp://admin@192.168.10.100/backup.cfg bootflash:backup.cfg vrf management configure replace bootflash:backup.cfg

D. copy sftp://admin@192.168.10.100/backup.cfg startup-config vrf management reload in 10

ANSWER: C

Explanation:

`copy sftp://admin@192.168.10.100/backup.cfg bootflash:backup.cfg vrf management` followed by `configure replace bootflash:backup.cfg` is correct because it first transfers the saved configuration from the remote server into a local Nexus file system using the management virtual routing and forwarding instance. The management interface is associated with the `management` virtual routing and forwarding instance on Nexus switches, so specifying that instance ensures that the remote server is reached through the intended out-of-band path. After the file is local, `configure replace` applies it as a replacement for the active configuration. This operation is designed for configuration replacement and rollback workflows: it compares the target file with the active configuration and makes the required changes without rebooting the switch. Staging the file on `bootflash:` is also preferable operationally because the replacement source remains locally available while the replacement is processed. If the backup configuration contains management connectivity settings, engineers should verify console access or use an appropriate rollback plan before applying it, because configuration replacement can alter the management path being used. Cisco documents configuration-file copy and replacement operations in the Cisco Nexus 9000 Series NX-OS Fundamentals Configuration Guide and management-interface behavior in the NX-OS Interfaces Configuration Guide.