DUMPSARENA

Troubleshooting Cisco Data Center Infrastructure (300-615 DCIT)

Cisco 300-615

Version Demo

Total Demo Questions: 38

Total Premium Questions: 384

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

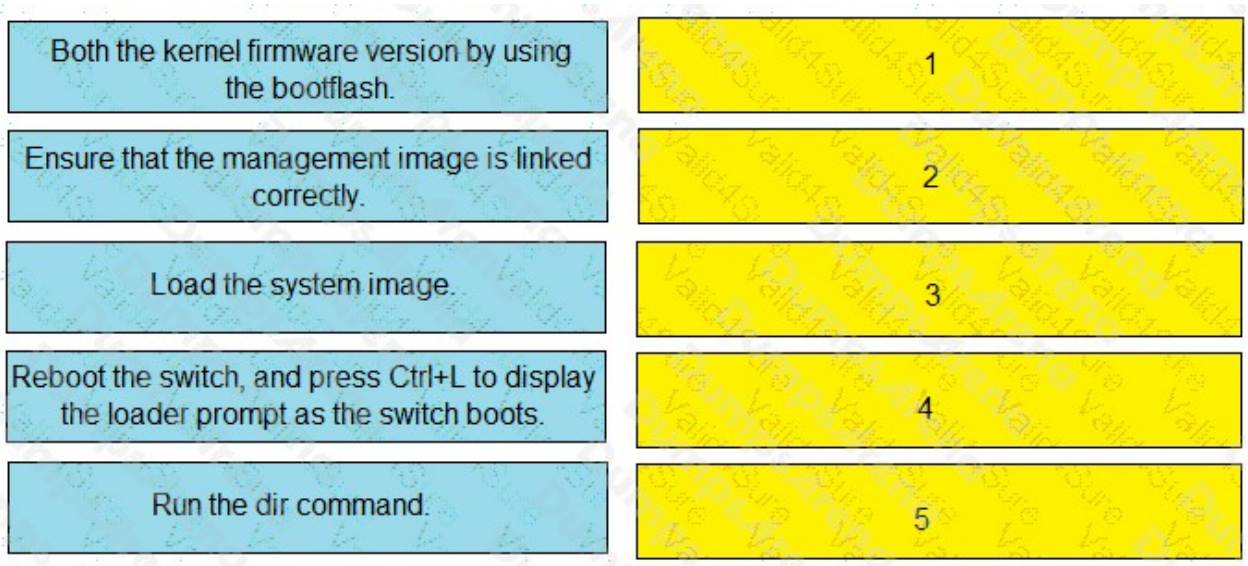
sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Troubleshoot Data Center Network	162
Topic 2, Troubleshoot Data Center Compute Platform	88
Topic 3, Troubleshoot Data Center Storage Network	79
Topic 4, Troubleshoot Data Center Automation and Orchestration	55
Total	384

QUESTION NO: 1 - (DRAG DROP)

A firmware upgrade on a fabric interconnect fails. A bootflash contains a valid image. Drag and drop the recovery steps from the left onto the correct order on the right.

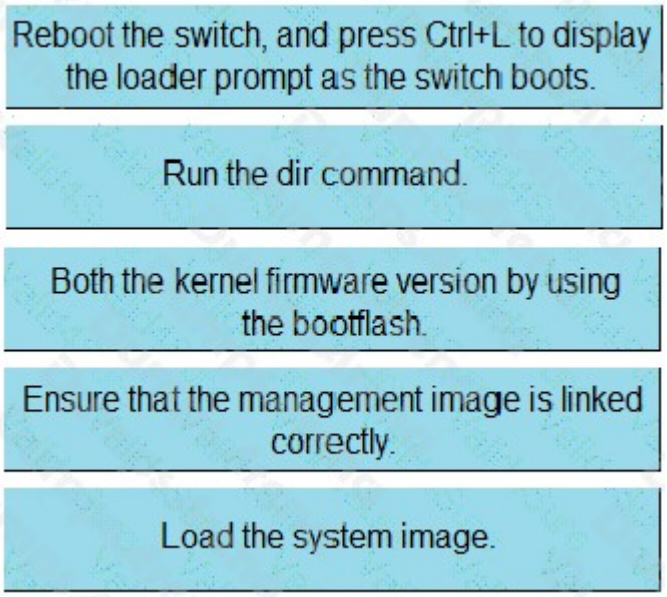


ANSWER:

Answer:

Explanation:

The correct action is the drag-and-drop mapping shown in the answer key. Cisco UCS troubleshooting depends on the relationship between policy, inventory discovery, firmware state, and server identity. The question involves firmware, fabric interconnect, so the selected option fixes the UCS component or policy that directly controls the failure. In UCS Manager, service profiles, vNIC/vHBA templates, boot policies, firmware packages, server pools, adapters, fabric interconnects, and I/O modules are applied through a managed model; a server can fail even when basic IP reachability exists if the relevant policy is incomplete or the hardware inventory has not been acknowledged correctly. The distractors typically change an unrelated object, force a disruptive recovery without fixing the policy, or apply to a different UCS mode. Cisco guidance is to verify faults, FSM stage, inventory status, association status, and the specific boot or connectivity policy before taking action. After the remediation, validate that the fault clears, the FSM completes, and the server or fabric component returns to the



expected operable state.

Explanation:

Because a valid firmware image already exists in bootflash, recovery should use the fabric interconnect's loader process to start from that known-good local image. First, reboot the switch and press Ctrl+L during startup to stop the normal boot sequence and access the loader prompt. The loader is the appropriate environment for recovering from an unsuccessful firmware upgrade because it can inspect local storage and manually initiate boot operations before the full system image is available.

At the loader prompt, run the `dir` command to confirm the bootflash contents and verify the image that will be used for recovery. Next, boot the kernel firmware version by using the image in bootflash. This starts the required kernel-level environment from the local valid image and permits the rest of the recovery boot sequence to proceed. Once that environment is active, ensure that the management image is linked correctly. The management image link identifies the image that the fabric interconnect should use for the next stage of booting, so it must point to the intended valid image before the system image is loaded.

Finally, load the system image. This completes the transition from the loader and kernel recovery stages into the normal fabric interconnect software environment. The sequence follows the dependency chain: enter recovery mode, inspect the local image, boot the recovery kernel, validate the image reference, and load the complete system software. Cisco's UCS documentation and support resources for fabric interconnect software recovery are available through Cisco UCS 6200 Series Fabric Interconnect support and the Cisco UCS Manager support portal.

QUESTION NO: 2

Refer to the exhibit.

The EEM applet generates an SNMP string every time a user executes the show ip route command After the applet was deployed, other engineers reported that they can no longer use the show ip route command to view the routing information. Which action resolves the issue?

A. Create a command alias for show ip route

B. Include a show Ip route CLI action statement

C. Modify the event-default action statement to match a syslog pattern

D. Add an event-default action statement at the end of the script

ANSWER: D

Explanation:

Add an event-default action statement at the end of the script. An Embedded Event Manager applet using the CLI event detector can intercept a command entered by a user. When the policy is configured to run synchronously, the command is held while the applet processes the event. This behavior is useful when a policy must validate, modify, or prevent a command, but it also means that the original command does not automatically proceed after the applet action completes.

The `event-default` action explicitly instructs EEM to perform the default processing for the detected CLI event. In this case, it allows the intercepted `show ip route` command to continue normally after EEM generates the SNMP notification. Engineers can therefore retain the required monitoring behavior while restoring access to routing-table output. The event-default action belongs after the notification action so that the applet performs its intended work before releasing the original command for execution. Cisco documents the CLI event detector and synchronous command handling in its Embedded Event Manager Configuration Guide; EEM action syntax is also covered in the Embedded Event Manager Command Reference.

QUESTION NO: 3

Refer to the exhibit.

```
switch1# show vpc brief
Legend:
(*) - local vPC is down, forwarding via vPC peer-link

vPC domain id : 500
Peer status : peer link is down
vPC keep-alive status : Suspended (Destination IP not reachable)
Configuration consistency status: success
vPC role : secondary, operational primary
Number of vPCs configured : 4
Peer Gateway : Disabled
Dual-active excluded VLANs : -

vPC Peer-link status
---------------------------------------------------------------------
id Port Status Active vlans
-- ---- ------ --------------------------------------------------
1 Po500 down -


switch2# show vpc brief
Legend:
(*) - local vPC is down, forwarding via vPC peer-link

vPC domain id : 20
Peer status : peer link is down
vPC keep-alive status : Suspended (Destination IP not reachable)
Configuration consistency status: success
vPC role : primary
Number of vPCs configured : 4
Peer Gateway : Disabled
Dual-active excluded VLANs : -

vPC Peer-link status
---------------------------------------------------------------------
id Port Status Active vlans
-- ---- ------ --------------------------------------------------
1 Po500 down -
```

The vPC between switch1 and switch2 does not work. Which two actions resolve the problem? (Choose two.)

A. Configure the IP address on the interface.

B. Activate VLANs on the vPC.

C. Configure one of the switches as primary for the vPC.

D. Correct the configuration of the vPC peer link and the vPC peer keepalive.

E. Match the vPC domain ID between the two devices.

ANSWER: D E

Explanation:

Correct the configuration of the vPC peer link and the vPC peer keepalive because these are the two required control connections between vPC peers. The peer link carries synchronization traffic, including MAC-address, ARP, and IGMP state, and it also supports traffic forwarding during certain failure scenarios. The peer-keepalive link is a separate, routed management or Layer 3 path used to determine whether the remote peer remains alive when peer-link connectivity is lost. Restoring both connections establishes the required vPC peer relationship and allows the vPC to become operational.

Match the vPC domain ID between the two devices because both vPC peers must use the identical domain ID. The domain ID identifies the two switches as members of the same vPC system and is part of the consistency validation performed while establishing the peer adjacency. After the peer link, peer keepalive, and matching domain ID are correctly configured, the

switches can form the vPC peer relationship and bring eligible vPC member ports up. Cisco's vPC troubleshooting guidance and vPC technology overview describe these peer-link, keepalive, and domain requirements.

QUESTION NO: 4

The script fails to convert a JSON variable into a Python dictionary. Which action resolves the issue?

```
import yaml

json_output = '''
{
  "totalCount": "26",
  "imdata": [
  {
    "fvTenant": {
      "attributes" {
        "descr": "",
        "dn": "uni/tn-LuTest",
        "lcOwn": "local",
        "modTs": "2021-01-12T17:30:56.922+00:00",
        "monPolDn": "uni/tn-common/monepg-default",
       "name": "LuTest"
       "uid": "40546834"
      }
     }
   }
  ]
}

json.loads(json_output)

--output omitted--
json.decoder.JSONDecodeError: --output omitted--
```

A. Use a json.dumps(json_output) function

B. Define a json_output as a YAML object, yaml json_output

C. Fix the JSON syntax by modifying "attribures" to "attributes":

D. Invoke a function yaml.loads(json._output ,decoder_"json")

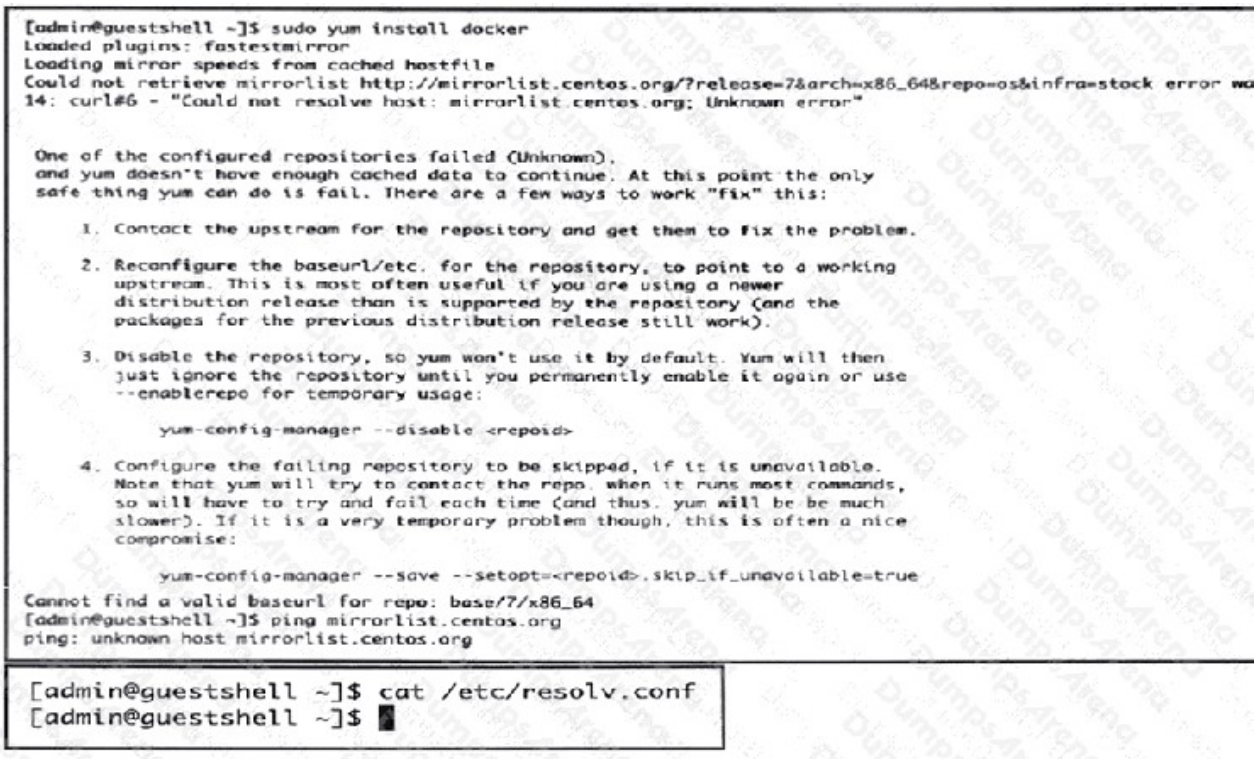
ANSWER: C

Explanation:

Fix the JSON syntax by modifying “attribures” to “attributes” because the JSON payload must use the exact key name expected by the script or API data model. After the key is corrected, Python can parse the valid JSON text into a native dictionary with `json.loads()`. JSON object member names are strings, and although a misspelled member name can still be syntactically valid JSON in isolation, it prevents the resulting dictionary from containing the expected `attributes` key. Any subsequent code that accesses that key, validates the payload, or maps it to a Cisco API object will therefore fail. Correcting the key ensures that the parsed dictionary has the structure required by the script. The Python standard library documents that `json.loads()` deserializes JSON text into Python objects, including dictionaries for JSON objects. Cisco automation workflows similarly depend on accurately structured JSON payloads and exact attribute names when exchanging data with APIs. See the Python JSON library documentation and Cisco DevNet learning resources.

QUESTION NO: 5

Refer to the exhibit.



An engineer fails to install Perl in the guest shell on a Cisco Nexus 9000 Series Switch. Which two actions must be taken in the guest shell to resolve the issue? (Choose two.)

A. Add the DNS server and suffix.

B. Add a default gateway and nameserver.

C. Install the unbound resolver package using curl.

D. Configure domain-name and nameserver under the management VRF.

E. Export the http_proxy and https_proxy environment variables

ANSWER: A E

Explanation:

Add the DNS server and suffix. is required because Guest Shell is a separate Linux container with its own network namespace and resolver configuration. Package installation tools must be able to resolve the package repository hostname, so the Guest Shell resolver needs a reachable DNS server and the appropriate search suffix when short hostnames or an

internal DNS domain are used. This configuration is made within the Guest Shell environment rather than relying solely on NX-OS name-resolution settings.

Export the http_proxy and https_proxy environment variables is also required when the Guest Shell reaches external package repositories through an enterprise web proxy. Yum and related package-download tools use these environment variables to send HTTP and HTTPS repository requests through that proxy. With working DNS resolution and the proxy variables exported, the Guest Shell can locate and reach the repository hosting the Perl package, allowing the installation to complete. Cisco documents that Guest Shell is an isolated Linux environment and that its connectivity prerequisites, including DNS and proxy access where required, must be configured for package operations. See Cisco's Guest Shell documentation and the Cisco Nexus 9000 Guest Shell configuration guide.

QUESTION NO: 6

An engineer attempts to register Cisco UCS Manager to Cisco UCS Central, but the registration fails. The engineer can ping Cisco UCS Central from UCS Manager. Which two actions must the engineer attempt to resolve the problem? (Choose two.)

A. Synchronize the date and time to NTP for Cisco UCS Central and the Cisco UCS domains.

B. Apply the UCS Central license to UCS Central.

C. Place Cisco UCS Manager on the same subnet as Cisco UCS Central.

D. Allow port 443 between Cisco UCS Manager and Cisco UCS Central.

E. Allow port 9443 between Cisco UCS Manager and Cisco UCS Central.

ANSWER: A D

Explanation:

Synchronizing the date and time to NTP for Cisco UCS Central and the Cisco UCS domains is necessary because registration establishes a secure, certificate-based HTTPS relationship. The clocks on UCS Manager and UCS Central must be sufficiently consistent for certificate validity periods and secure session establishment to be evaluated correctly. Configuring both systems to use reliable NTP sources prevents clock drift from interrupting this process.

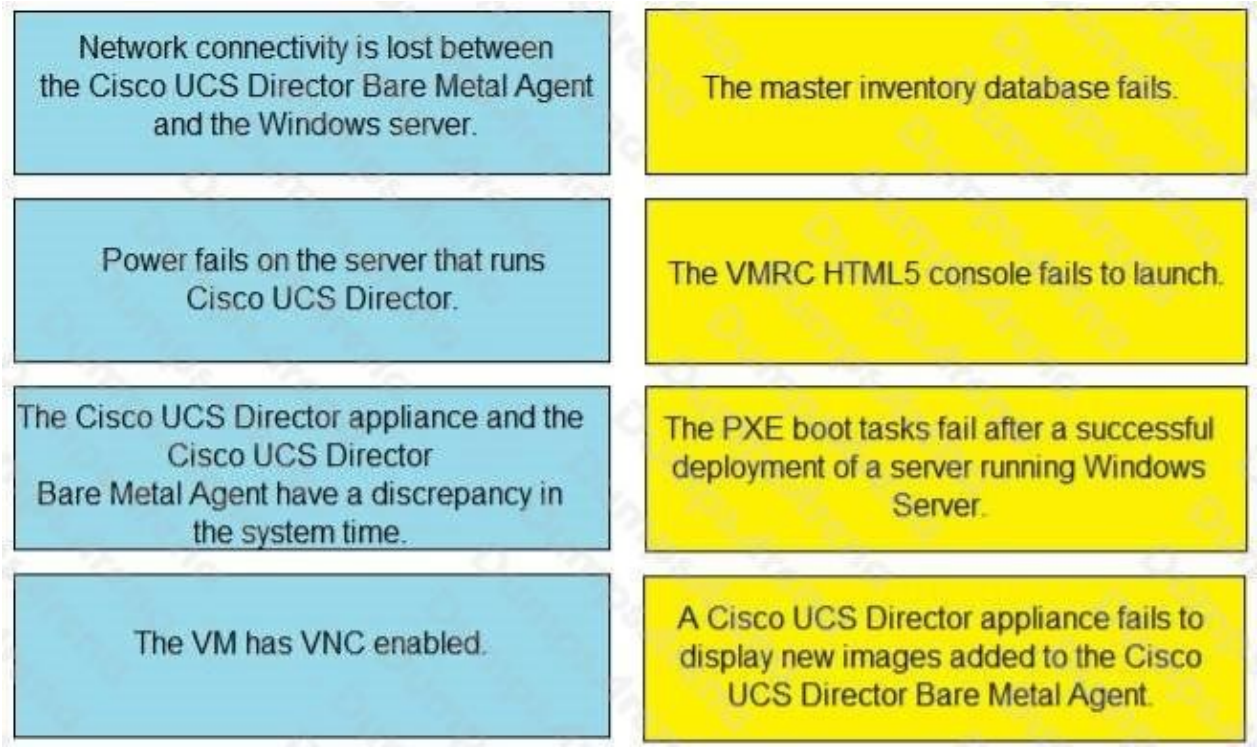
Allowing port 443 between Cisco UCS Manager and Cisco UCS Central is also required. A successful ping establishes only IP-layer reachability using ICMP; it does not prove that the TCP service used for UCS Central registration is reachable. UCS Manager communicates with UCS Central through HTTPS, which uses TCP port 443. Firewalls, ACLs, or intervening security devices must therefore permit bidirectional session establishment and return traffic for this connection. Cisco's UCS Central documentation identifies HTTPS connectivity as part of the communication path between registered UCS domains and UCS Central. See the Cisco UCS Central installation and configuration guides and the Cisco UCS Central product documentation.

QUESTION NO: 7 - (DRAG DROP)

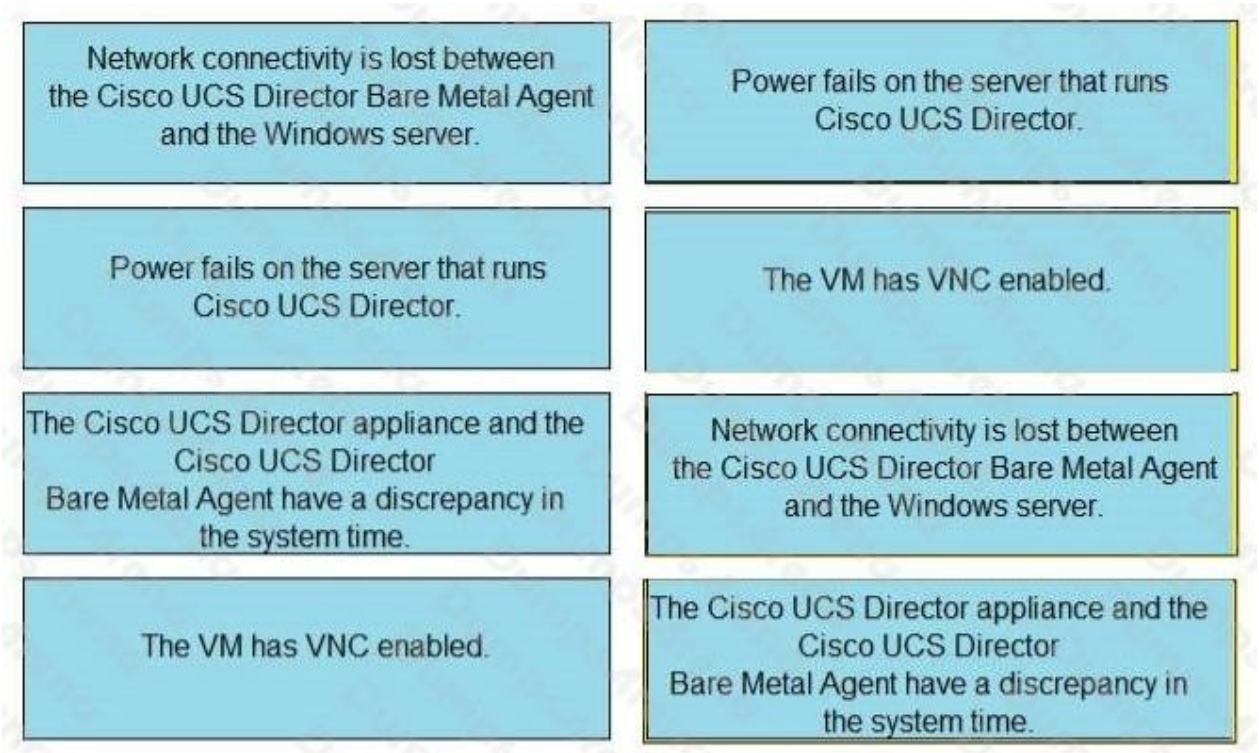
DRAG DROP

An engineer troubleshoots issues in Cisco UCS Director. Drag and drop the possible causes of the issues from the left onto the corresponding issues on the right.

Select and Place:



ANSWER:



Explanation:

The correct troubleshooting relationships follow the dependencies used by Cisco UCS Director, its database services, and the Bare Metal Agent workflow. The master inventory database is hosted by the Cisco UCS Director installation, so a power failure on the server running Cisco UCS Director interrupts the database service and causes the master inventory database failure. Cisco documentation describes UCS Director as an appliance-based management platform whose operational services and inventory functions depend on the appliance being available and healthy. See the Cisco UCS Director installation guides for the appliance architecture and service deployment context.

VMRC access relies on the expected console connection method for the virtual machine. With VNC enabled, the VM console configuration can prevent the VMRC HTML5 console from launching as expected. The Cisco UCS Director Virtual Machine

Manager Console integration is intended to provide browser-based access through the configured VM console workflow, so the VM's console settings must be compatible with that workflow.

PXE deployment requires network communication between the Cisco UCS Director Bare Metal Agent and the deployed Windows server. If that connectivity is lost, the post-deployment PXE boot task cannot complete successfully even though the Windows Server deployment itself completed. The Bare Metal Agent participates in provisioning and image-related operations, making its reachability essential to this stage of the process.

Finally, image discovery between the Cisco UCS Director appliance and the Bare Metal Agent depends on synchronized system time. A time discrepancy can invalidate time-sensitive communication, synchronization, and service requests, leaving newly added Bare Metal Agent images absent from the appliance display. Maintaining consistent time across integrated infrastructure components is therefore necessary for reliable inventory and image synchronization. Cisco's UCS Director support documentation provides the product documentation and operational guidance for these integrated services.

QUESTION NO: 8

Refer to the exhibit.



An engineer fails to implement a Cisco UCS Manager integration manager. The credentials and IP connectivity between Cisco APIC and UCS Manager are configured as expected. Which action resolves the issue?

A. Disable HTTP to HTTPS redirection in Cisco UCS Manager

B. Enable JSON API on the Cisco UCS Manager

C. Change the Integration Manager name to FQDN of the Cisco UCS Manager

D. Install ExternalSwitch app in the APIC controller

ANSWER: D

Explanation:

Installing the **ExternalSwitch** app in the APIC controller resolves the integration issue because this APIC application provides the capability APIC uses to integrate with and manage the external-switch-facing aspects of a Cisco UCS domain. APIC applications are not merely optional GUI components: they supply services, object models, and workflows required for particular controller integrations. Therefore, valid UCS Manager credentials and working IP reachability alone are insufficient when the required APIC application is absent or not operational. After the application is installed and enabled, APIC can complete the UCS Manager integration workflow and use the UCS domain for the intended ACI policy and virtual-machine networking integration. The engineer should also verify that the application is in a healthy running state on the APIC cluster before retrying the integration. Cisco documents APIC applications and controller operations in the Cisco APIC System Configuration Guide, and describes UCS-related virtual machine networking integration in the Cisco APIC Virtual Machine Networking Configuration Guide.

QUESTION NO: 9

A Cisco UCS C-Series Rack Server is not discovered after it is connected directly to a Cisco UCS Fabric Interconnect. The physical link is up, but no server inventory is shown in Cisco UCS Manager. Which two actions resolve the issue? (Choose two.)

A. Configure the fabric-interconnect port as a server port.

B. Configure the server for Cisco UCS Manager management mode.

C. Configure the fabric-interconnect port as an Ethernet uplink port.

D. Configure the server in standalone CIMC mode.

E. Enable NPIV on the fabric-interconnect Ethernet port.

ANSWER: A B

Explanation:

The fabric-interconnect port must be configured as a server port. Cisco UCS uses server ports to identify directly attached rack servers and to establish the management and data connectivity required for discovery. An uplink port or appliance port does not provide the expected rack-server discovery behavior.

The server must be configured for Cisco UCS Manager management mode. A C-Series server that remains in standalone CIMC mode is managed independently and is not available for UCS Manager discovery through the fabric interconnect. After the server port and management mode are corrected, Cisco UCS Manager can discover the server and display its inventory. Cisco documents integrated C-Series server connectivity in the Cisco UCS C-Series Integrated Management documentation.

QUESTION NO: 10

During a boot process of a Cisco UCS C-Series Rack Server, an engineer receives a " No Signal " message from the vKVM and physical connection. Which set of steps resolves the issue?

A. 1- Disconnect the power cord2- Confirm that all cards are properly seated3- Connect the power cord and power on the server

B. 1- Power off the server and disconnect the power cord2- Confirm that all cards are properly seated3- Connect the power cord and power on the server

C. 1- Power off the server and disconnect the power cord2- Confirm that all cards are available3- Connect the power cord and power on the server

D. 1- Disconnect the power cord2- Confirm that all cards are properly seated3- Connect the power cord

ANSWER: B

Explanation:

“1- Power off the server and disconnect the power cord 2- Confirm that all cards are properly seated 3- Connect the power cord and power on the server” is the appropriate recovery procedure. A simultaneous No Signal condition through both vKVM and a locally attached physical display indicates that video output is not being initialized during POST, rather than an isolated browser, vKVM-session, or display-cable issue. Fully powering off the rack server and removing AC power places the system in a safe state for hardware inspection. The engineer can then check that expansion cards and other internally installed components are firmly and correctly seated, because an improperly seated card can interfere with successful hardware initialization. Restoring AC power and booting the server allows the system to reinitialize its hardware and video subsystem during POST. This sequence also ensures that the physical inspection is performed only after server power has been removed, consistent with Cisco’s hardware maintenance guidance. Cisco documents C-Series server hardware installation and component handling procedures in its UCS C220 M5 Installation Guide; additional platform diagnostic and fault-isolation guidance is available in the Cisco UCS C-Series installation and configuration guides.

QUESTION NO: 11

After upgrading a client’s fabric interconnect to the latest version, an engineer discovers that the appliance has lost its entire configuration. Which two locations should an engineer check to confirm that the appliance had backed up the configuration prior to the upgrade in order to retrieve the lost configuration? (Choose two.)

A. Equipment FSM page, InternalBackup stage

B. Service profiles page, InternalBackup stage

C. Admin page, PollInternalBackup stage

D. Equipment FSM page, PollInternalBackup stage

E. Admin page, InternalBackup stage.

ANSWER: A C

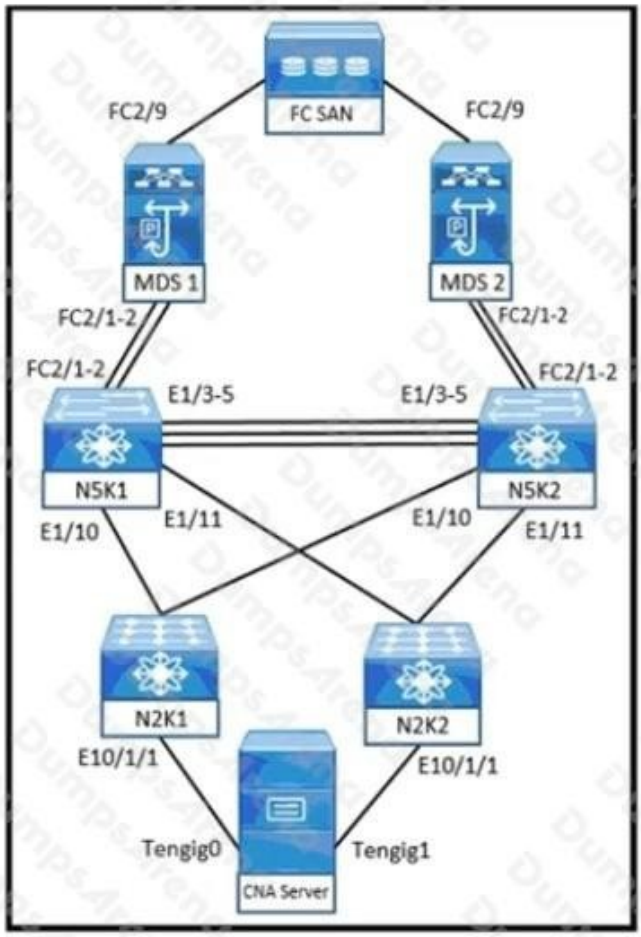
Explanation:

Cisco UCS Manager performs an internal configuration backup as part of an infrastructure firmware upgrade workflow. The **Equipment FSM page, InternalBackup stage** is the relevant workflow state under the equipment and firmware-management side of UCS Manager. A successful InternalBackup stage indicates that the upgrade process created its internal backup before proceeding with the fabric-interconnect upgrade operations.

The same activity can also be validated from the administrative backup workflow. The **Admin page, PollInternalBackup stage** polls and reports the completion state of the internal-backup operation. Reviewing both FSM locations provides confirmation that the automatic backup was requested and that its completion was observed. If these stages completed successfully, the engineer can use the internally created backup as the recovery source for restoring the UCS Manager configuration after the upgrade.

Cisco documents that configuration backups should be created before upgrades and that UCS Manager backup and restore operations are managed through its administration and firmware workflows. See the Cisco UCS Manager Configuration Guides and the Cisco UCS Manager Upgrade Guide.

QUESTION NO: 12




```
N5K1# show platform software fcoe_mgr event-history errors
1) Event:E_DEBUG, length:93, at 203962 usecs after Wed Jun 12
20:33:04 2019 [102] fcoe_mgr_vfc_ac_eval(4946):
DEBUG:shut:Sending event to delete protos of vfc1 due to 62
2) Event:E_DEBUG, length:119, at 197505 usecs after Wed Jun 12
20:33:04 2019  [102] fcoe_mgr_fc2_msg_handler(5706): proto
if_index 1e000000 p_proto (nil)and oxid 8805fc2 usrhandle 0[0]
iuhdr type:1
3) Event:E_DEBUG, length:91, at 197053 usecs after Wed Jun 12
20:33:04 2019 [102] fcoe_mgr_proto_ac_eval(1847): >Bringing down
PROTO 1e000000 due to
truly missing fka
```

Refer to the exhibit. An engineer discovered a VFC interface in a down state on a Cisco Nexus Switch. Which action resolves the problem?

A. Change the cable between the switch and the server

B. Reset the Cisco Nexus port toward the file server

C. Update the CNA firmware and driver versions

D. Modify the fip keepalive message time to 30 seconds

ANSWER: C

Explanation:

Update the CNA firmware and driver versions is the appropriate resolution because a virtual Fibre Channel interface depends on successful FCoE operation between the Nexus switch and the server converged network adapter. The adapter firmware and host driver must support the deployed FCoE features and interoperate with the switch software release. A mismatch can prevent the CNA from correctly establishing or maintaining FCoE control-plane operation, leaving the VFC interface down even when the physical Ethernet path is present. Updating the CNA firmware and driver to Cisco-supported, compatible versions addresses this endpoint interoperability condition and allows the adapter to complete the required FCoE initialization and login process. Cisco's Nexus FCoE troubleshooting guidance specifically identifies CNA driver and firmware compatibility as a prerequisite when investigating FCoE connectivity and VFC operational state. The supported hardware and software compatibility information should be checked before selecting the adapter image and driver package, then the update should be performed according to the server and CNA vendor's documented maintenance procedure. See Cisco's Nexus 5000 FCoE troubleshooting guide and the Cisco device support and compatibility resources.

QUESTION NO: 13

A customer configures HSRP between two data centers that are interconnected with OTV. The configuration succeeds, but traffic between two ESXi virtual hosts on the same site is routed suboptimally through the OTV overlay. Which two actions optimize the traffic? (Choose two.)

A. Disable first-hop redundancy.

B. Filter HSRP traffic by using a Layer 3 VACL on the OTV edge devices.

C. Filter HSRP by using a Layer 2 MAC list on the ESXi vSwitch.

D. Filter HSRP traffic by using a Layer 3 VACL on the ESXi vSwitch.

E. Filter HSRP by using a Layer 2 MAC list on the OTV edge devices.

ANSWER: B E

Explanation:

Filtering HSRP traffic by using a Layer 3 VACL on the OTV edge devices and filtering HSRP by using a Layer 2 MAC list on the OTV edge devices both provide first-hop redundancy protocol isolation across the OTV overlay. The objective is to prevent HSRP control packets from extending between sites. Once the HSRP adjacency is localized, each data center can maintain a local active first-hop gateway for its attached hosts rather than allowing a remote device to become the active gateway for a stretched VLAN. This keeps traffic sourced and destined within one site from being sent to the remote data center merely to reach the active HSRP gateway, then potentially returning through the overlay.

A Layer 3 VLAN access map can identify and drop HSRP IP/UDP control traffic, including the HSRP multicast group and UDP port used by the deployed HSRP version. A Layer 2 MAC access list can instead identify the relevant HSRP multicast MAC destination and prevent those frames from crossing the OTV edge. Applying either filtering method at the OTV edge is appropriate because that is the boundary at which intersite Layer 2 extension occurs. Cisco documents OTV first-hop redundancy isolation and the use of VLAN access maps for this purpose in its Nexus OTV Configuration Guide. HSRP packet formats and multicast addressing are documented in Cisco IOS HSRP documentation.

QUESTION NO: 14

The external routes fail to propagate to leaf switches in Cisco ACI fabric Which two actions resolve the issue? (Choose two.)

A. Enable the VTEP pool in the fabric

B. Assign an MP-BGP AS number to the fabric.

C. Specify the spine nodes as route reflectors

D. Associate the correct contract to the L30ut.

E. Configure an MP-BGP area.

ANSWER: B C

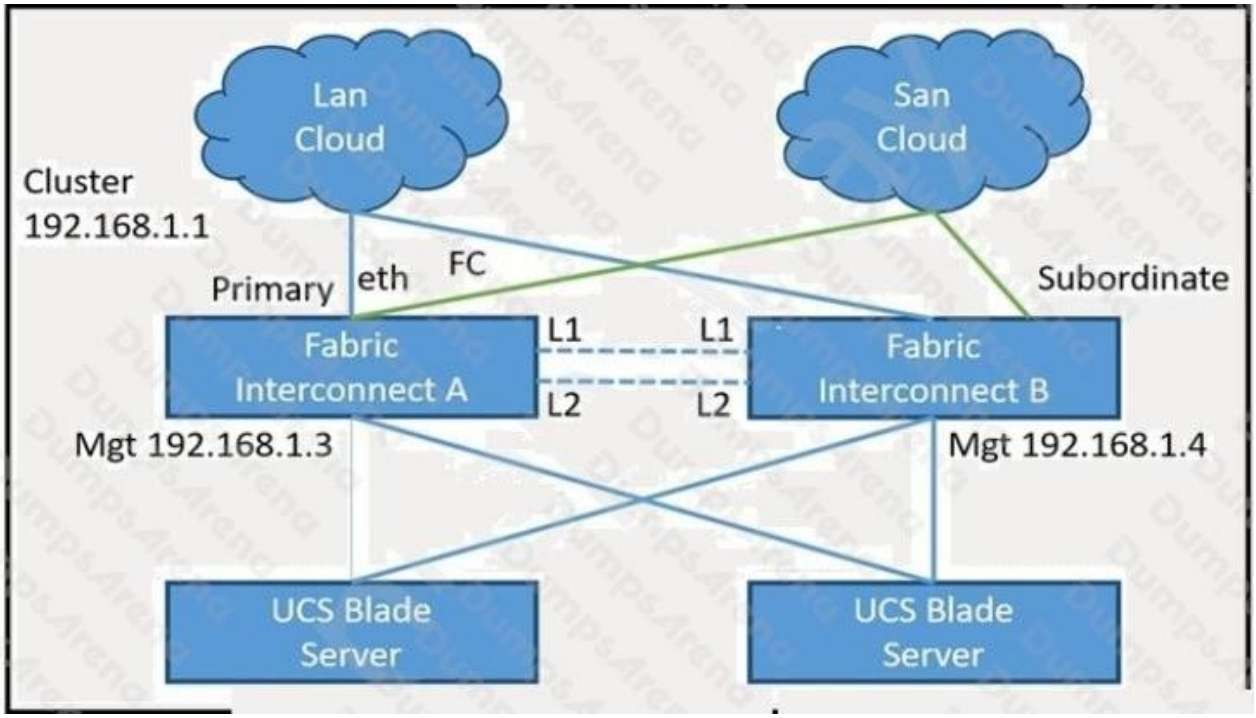
Explanation:

Assigning an MP-BGP AS number to the fabric establishes the autonomous-system identity used by the APIC-managed fabric control plane. This fabric ASN is required for Multiprotocol BGP operation that distributes routing information across the ACI fabric, including prefixes learned through an external routed connection. It also provides the BGP context necessary for route advertisements and loop-prevention behavior between the fabric and external peers.

Specifying spine nodes as route reflectors enables scalable propagation of externally learned routes to the leaf switches that require them. In ACI, route-reflector spine switches reflect MP-BGP routes between leaf switches, avoiding a full mesh of BGP sessions while ensuring relevant external prefixes can be distributed throughout the fabric. The route-reflector policy is applied at the fabric level and must identify the intended spine nodes, which then provide the internal BGP route-reflection function for participating leaf nodes. Together, the fabric MP-BGP ASN and spine route reflectors provide the required internal BGP framework for external-route dissemination. Cisco documents these fabric BGP and route-reflector requirements in the Cisco APIC Layer 3 Configuration Guide and the Cisco APIC Basic Configuration Guide.

QUESTION NO: 15

Refer to the exhibit.



A network engineer finds one of the fabric interconnects offline when connecting L1 and L2 ports on both fabric interconnects. Which action resolves the issue?

A. 1. Connect to Fabric Interconnect B.
2. Verify the cluster status and HA election.
3. Validate Fabric Interconnect B hardware issues.

B. 1. Connect to Fabric Interconnect A.
2. Verify the cluster status and HA election.
3. Validate Fabric Interconnect A hardware issues.

C. 1. Connect Fabric Interconnect B and change the role to primary.
2. Reboot Fabric Interconnect B.
3. Add the Fabric Interconnect A as subordinate fabric to the cluster.

D. 1. Connect Fabric Interconnect B and execute “Erase configuration”
2. Reboot Fabric Interconnect B.
3. Add the subordinate Fabric Interconnect to the cluster.

ANSWER: D

Explanation:

Connect Fabric Interconnect B and execute “Erase configuration,” reboot it, and add the subordinate Fabric Interconnect to the cluster. A Cisco UCS fabric-interconnect pair operates as a clustered system with one primary and one subordinate fabric interconnect. For the subordinate unit to join cleanly, it must not retain a conflicting standalone or previous cluster configuration. Erasing the configuration resets the fabric interconnect so that it can be discovered and incorporated by the established cluster after the L1 and L2 cluster links are connected. Rebooting applies the reset state and allows the fabric interconnect to complete its startup and cluster-join process. Once it joins, Cisco UCS Manager can synchronize the required cluster configuration and restore high availability between the two fabric interconnects. This is the appropriate recovery workflow when the intended subordinate fabric interconnect remains offline because of an existing or inconsistent configuration. Cisco documents UCS installation and configuration procedures, including fabric-interconnect clustering, in its UCS Manager installation and configuration guides. Additional operational guidance is available in the Cisco UCS Troubleshooting Guide.

QUESTION NO: 16

Endpoints in an application EPG must access a network through an existing Cisco ACI Layer 3 Out connection. The Layer 3 Out routing protocol is operational, and the external prefix is learned by the fabric, but traffic from the application EPG to the external network is dropped. Which two actions resolve the issue? (Choose two.)

A. Associate a contract with the application EPG as consumer and the external EPG as provider.

B. Configure the external destination prefix as an external subnet under the Layer 3 Out external EPG.

C. Configure the application EPG as a route reflector client.

D. Configure the bridge domain as an unenforced VRF.

E. Assign the external prefix as a static MAC address in the application EPG.

ANSWER: A B

Explanation:

The internal application EPG and the external EPG must be associated with a contract in the correct provider-consumer direction. For traffic initiated from the application EPG toward the external network, the application EPG consumes the contract and the external EPG provides it. The contract must contain a filter that permits the required protocol and port.

The external destination prefix must also be configured as an external subnet under the Layer 3 Out external EPG. This associates the routed external address space with the external EPG so that policy can be applied to traffic destined for that prefix. A learned route alone does not create the required policy classification for the external destination.

Cisco describes Layer 3 Out external EPGs, external subnets, and contracts in the Cisco APIC Layer 3 Configuration Guide.

QUESTION NO: 17

An upgrade of protected RPMs from the Bash shell does not take effect. Which action resolves this issue?

A. Restart the Bash shell.

B. Reload the switch.

C. Upgrade the RPMs from the Guest shell.

D. Disable and re enable the Bash feature.

ANSWER: D

Explanation:

Disable and re enable the Bash feature. is the required action because NX-OS protects RPM packages that are part of, or used by, the Bash environment. When Bash is enabled, its shell environment and associated protected package state can remain active, so an RPM update initiated from that environment may be staged without becoming active. Disabling the Bash feature stops and removes the active Bash environment. Enabling it again reinitializes Bash and loads the currently installed versions of the protected RPMs, allowing the upgrade to take effect.

This behavior is particularly important on Cisco NX-OS platforms because Bash is an optional, feature-controlled shell rather than the standard NX-OS CLI. Feature state governs the lifecycle of the Bash service and its runtime components. Administrators should therefore perform the feature disable/enable operation during an approved maintenance window and confirm that Bash-dependent automation or operational workflows are not running before restarting the environment. Cisco documents Bash as a feature that must be explicitly enabled and managed within NX-OS configuration. See the Cisco Nexus 9000 NX-OS Programmability Guide: Bash and the Cisco NX-OS Software support documentation.

QUESTION NO: 18

The external routes fail to propagate to leaf switches in Cisco ACI fabric. Which two actions resolve the issue? (Choose two.)

A. Enable the VTEP pool in the fabric.

B. Assign an MP-BGP AS number to the fabric.

C. Specify the spine nodes as route reflectors.

D. Associate the correct contract to the L3Out.

E. Configure an MP-BGP area.

ANSWER: B C

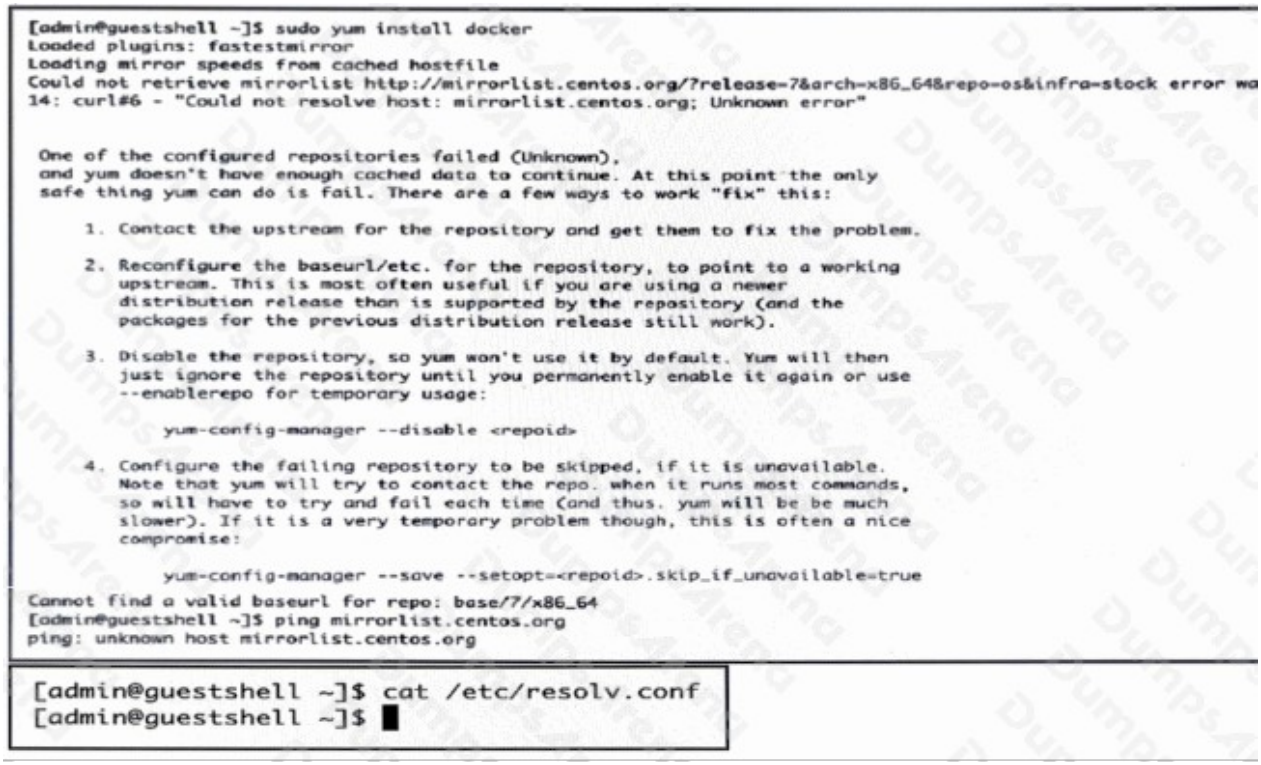
Explanation:

Cisco ACI distributes endpoint and external-routing reachability information across the fabric by using the fabric's MP-BGP EVPN control plane. Therefore, assigning an MP-BGP AS number to the fabric is required to establish the fabric-wide BGP identity used for this control-plane exchange. The fabric BGP autonomous system number is configured as a fabric-wide setting and must be present before the relevant BGP control-plane functions can operate correctly.

Spine switches must also be specified as route reflectors. In an ACI fabric, leaf switches act as BGP route-reflector clients, while spine switches provide route-reflector services. The route reflectors receive reachability advertisements and reflect them to the appropriate leaf switches, allowing routes learned through a Layer 3 Out connection to become available throughout the fabric. A resilient deployment normally designates two spine nodes as route reflectors so that route distribution continues if one reflector becomes unavailable. Cisco documents the required fabric BGP AS configuration and route-reflector node policy under fabric infrastructure configuration. See the Cisco APIC Basic Configuration Guide and the Cisco APIC Layer 3 Configuration Guide.

QUESTION NO: 19

Refer to the exhibit.



An engineer fails to install Perl in the guest shell on a Cisco Nexus 9000 Series Switch. Which two actions must be taken in the guest shell to resolve the issue? (Choose two.)

A. Add the DNS server and suffix.

B. Add a default gateway and nameserver.

C. Install the unbound resolver package using curl.

D. Configure domain-name and nameserver under the management VRF.

E. Export the http_proxy and https_proxy environment variables

ANSWER: A E

Explanation:

Adding the DNS server and suffix is required because the Guest Shell is a separate Linux environment and must be able to resolve the repository and proxy hostnames used by its package manager. Supplying a reachable resolver in the Guest Shell's resolver configuration, along with the appropriate search suffix when short hostnames are used, enables name resolution for the external resources needed to download the Perl package. The Guest Shell does not simply rely on a successful command prompt on NX-OS; its own Linux networking and resolver settings must support the requested connection.

Exporting the `http_proxy` and `https_proxy` environment variables is also required when the management network reaches external package repositories through an HTTP/HTTPS proxy. The package-management tools and underlying download utilities read these variables to send repository traffic through the corporate proxy. With both resolver configuration and proxy variables present, `yum` can resolve repository/proxy names and establish the outbound connection needed to retrieve Perl. Cisco documents Guest Shell as an isolated Linux container with its own operational configuration in the Cisco Nexus 9000 NX-OS Programmability Guide; proxy-variable behavior is defined in the curl documentation.

QUESTION NO: 20

Refer to the exhibit.

A client configures an upgrade of its Cisco Nexus switches that are connected to the Cisco ACI controller, such that the switches are upgraded one at a time. After the upgrade is run, it is discovered that both Cisco Nexus switches were upgraded simultaneously. Which two actions ensure that the switches upgrade one at a time? (Choose two.)

A. Configure the *Do not pause on failure and do not wait on cluster health " Run Mode

B. Choose the ' Pause upon upgrade failure " Run Mode

C. Select the " Graceful Maintenance " checkbox

D. Place each leaf switch in a different upgrade group

E. Check the Ignore Compatibility " checkbox

ANSWER: C D

Explanation:

Selecting the " Graceful Maintenance " checkbox enables the APIC to use a controlled, rolling maintenance workflow for the nodes covered by the maintenance policy. This is especially important for leaf switches that provide redundant connectivity, because the controller can take a node through maintenance while preserving forwarding through the remaining available topology. Graceful maintenance is the setting that instructs the upgrade process to account for operational continuity rather than treating all selected switches as an unrestricted parallel update set.

Placing each leaf switch in a different upgrade group establishes separate maintenance boundaries. Cisco ACI uses maintenance groups to define the set of switches addressed by a firmware upgrade operation. Separating the switches prevents both switches from being treated as members of the same group for the upgrade workflow, allowing the administrator to apply the upgrade in a sequential, switch-by-switch manner. Combining separate upgrade groups with graceful maintenance provides both explicit operational separation and a controlled maintenance process. Before execution, the administrator should verify group membership and the selected firmware maintenance policy, then monitor the APIC upgrade status before proceeding to the next group. Cisco documents maintenance-group-based switch firmware upgrades and graceful upgrade behavior in the Cisco APIC Upgrade Guide and the Cisco APIC Configuration Guide.

QUESTION NO: 21

A vPC Type-1 inconsistency between two vPC peers in a VXLAN EVPN setup is discovered. Which two actions need to be attempted to resolve the issue? (Choose two.)

A. Configure the NVE interfaces to be Up on both switches.

B. Set a different distributed gateway virtual MAC address.

C. Set a different secondary IP addresses on NVE source-interface.

D. Configure the same VNI to multicast group maping.

E. Set a different primary IP addresses on NVE source-interface.

ANSWER: A D

Explanation:

Configure the NVE interfaces to be Up on both switches and configure the same VNI to multicast group maping because both settings are subject to vPC Type-1 consistency validation for VXLAN. A Type-1 mismatch is a critical vPC consistency failure: the affected vPC can be suspended to prevent forwarding behavior from differing across the two peers. Therefore, restoring an operational NVE interface on each peer is necessary before the vPC pair can provide consistent VXLAN encapsulation and decapsulation behavior.

For multicast-based BUM replication, each VNI must have an identical multicast-group association on both members of the vPC pair. Matching the VNI-to-multicast-group mapping ensures that both peers replicate and receive broadcast, unknown-unicast, and multicast VXLAN traffic consistently. After correcting the relevant settings, verify the peer consistency status and vPC state with commands such as `show vpc consistency-parameters vpc` and `show vpc`. Cisco documents

VXLAN EVPN and vPC design and configuration requirements in the Cisco Nexus 9000 Series VXLAN Configuration Guide and the Cisco Nexus 9000 Series vPC Configuration Guide.

QUESTION NO: 22

```
Nexus# show running-config | section role

role feature-group name Group1
  feature radius
role name User1
  rule 4 deny read-write feature-group Group1
  rule 3 permit read
  rule 2 deny read-write
  rule 1 deny command clear users

Nexus(config)# no role feature-group name Group1
ERROR: Feature group is in use. Can't remove the feature group
```

Refer to the exhibit.

An engineer must delete the feature-group role on the Cisco Nexus 9000 Series Switch but the CLI rejects the command Which command set resolves the issue?

A. rote name User1rule 4 permit read-write feature-group Group1

B. role name User1 no rule 4

C. role name Userno rule 1

D. role name User1rule1 permit read-write feature-group Group1

ANSWER: B

Explanation:

`role name User1 no rule 4` is the required command set because the feature group is still referenced by rule 4 in the `User1` RBAC role. On Cisco Nexus NX-OS, feature groups can be used within role rules to grant a privilege level, such as read-write access, to a collection of features. NX-OS preserves this dependency and does not allow removal of an object that remains in use by an active role rule. The administrator must first enter the relevant role context and remove the rule that contains the feature-group reference. Removing rule 4 eliminates the dependency between `User1` and `Group1`, after which the feature group can be deleted successfully. This sequence also maintains a consistent RBAC policy state: permissions are explicitly removed from the user role before the referenced permission group is removed. Afterward, verify the role configuration and confirm that rule 4 is no longer listed before issuing the feature-group deletion command. Cisco documents NX-OS role-based access control, including role rules and feature-based authorization, in the Cisco Nexus 9000 Series NX-OS Security Configuration Guide and the NX-OS Security Command Reference.

QUESTION NO: 23

```
N5k-A-199# sh install all impact system
N5000-uk9.7.1.4.N1.1.bin kickstart
N5000-uk9-kickstart.7.3.4.N1.1.bin
Verifying image bootflash:/n5000-uk9-
kickstart.7.3.4.N1.1.bin for boot
variable "kickstart".
[####################] 100% -- SUCCESS
Verifying image bootflash:/n5000-
uk9.7.1.4.N1.1.bin for boot variable
"system".
[# ] 0% -- FAIL. Return code 0x4045001F
(image MD5  checksum error).
Install has failed. Image verification
failed (0x40930011)
```

Refer to an exhibit. An engineer is troubleshooting an upgrade failure on a switch. Which action resolves the issue?

A. Save the system image in NVRAM.

B. Use the same system image as the kickstart image.

C. Load a new system image.

D. Reload the same system image.

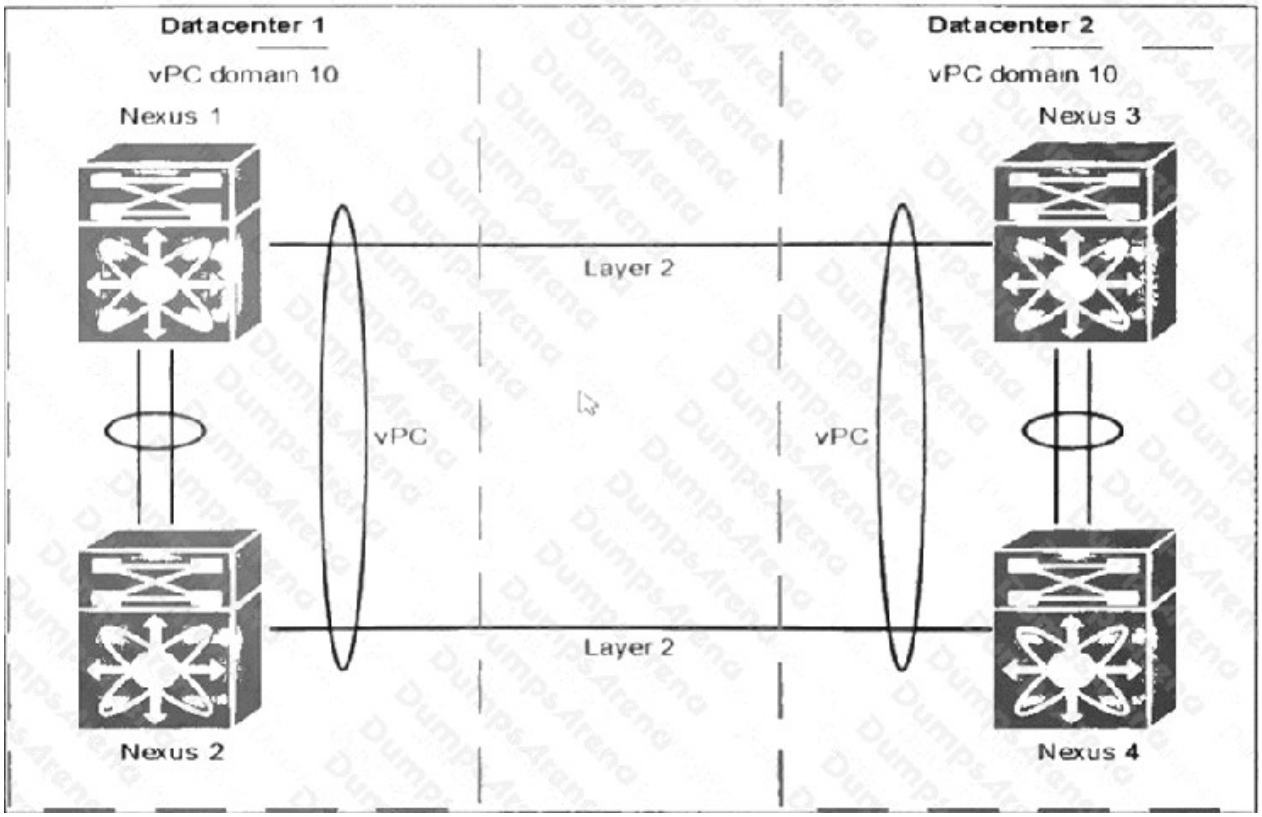
ANSWER: C

Explanation:

Loading a new system image resolves the upgrade failure because the system image currently selected for the installation cannot be used successfully by the switch. In NX-OS upgrade workflows, the system image must be a valid, complete image file that is supported by the platform and compatible with the intended software release. A failed image validation or installation indicates that the appropriate remediation is to obtain the correct image again, copy it to local storage, and perform the upgrade using that replacement file.

Before initiating the installation, the engineer should verify the image integrity with the published checksum and confirm that the target release is supported for the switch model and its current upgrade path. The image should be stored in a location accessible to the switch, commonly `bootflash:`, and its availability and integrity should be confirmed before rerunning the upgrade. Cisco documents the required image-validation and installation workflow in its Nexus 9000 Series NX-OS Software Upgrade and Downgrade Guide. Cisco also provides software image downloads and associated release information through the Cisco Software Download portal.

QUESTION NO: 24



Refer to the exhibit. The administrator set up two pairs of Cisco Nexus switches. The administrator set different vPC priorities on all four Nexus switches. As soon as the administrator activates the vPC between the two pairs the network faces

different issues. The problems range from both Nexus pairs declare themselves as root bridge, as well as spanning-tree inconsistencies and traffic forwarding issues. Which action resolves these issues?

A. Change the role priority on one pair.

B. Configure peer-config-check-bypass on both pairs.

C. Configure a vPC peer link between both peers.

D. Change domain ID on one pair to a different ID.

ANSWER: D

Explanation:

Change domain ID on one pair to a different ID. A vPC domain ID must be unique for each independent vPC pair in the same Layer 2 environment. Cisco Nexus uses the vPC domain ID as an input to the vPC system MAC address. The two switches within one vPC domain deliberately present this shared system MAC to connected devices so that they behave as a single logical spanning-tree bridge. If two separate vPC pairs use the same domain ID, they can derive the same vPC system MAC and therefore advertise an identical spanning-tree bridge ID. This duplicate identity prevents spanning tree from making a reliable root-bridge and port-role decision, which can produce the reported root declarations, consistency conditions, and forwarding disruption. Assigning a distinct domain ID to one vPC pair causes that pair to use a different vPC system MAC and restores a unique bridge identity for each logical vPC switch. After the change, verify the vPC domain and system MAC with `show vpc`, then confirm the expected root bridge and forwarding roles with `show spanning-tree`. Cisco documents the relationship between the vPC domain ID and system MAC in the Nexus vPC configuration guide; spanning-tree operation is covered in Cisco's Layer 2 switching guide.

QUESTION NO: 25

Refer to the exhibit.

```
switch1# show vpc brief
Legend:
(*) - local vPC is down, forwarding via vPC peer-link

vPC domain id : 500
Peer status : peer link is down
vPC keep-alive status : Suspened (Destination IP not reachable)
Configuration consistency status : success
vPC role : secondary, operational primary
Number of vPCs configured : 4
Peer Gateway : Disabled
Dual-active exluded VLANs : -

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status  Active  vlans
---  -----  ------- ------------------------------------------------
1    Po500  down    -

switch2# show vpc brief
Legend:
(*) - local vPC is down, forwarding via vPC peer-link

vPC domain id : 20
Peer status : peer link is down
vPC keep-alive status : Suspened (Destination IP not reachable)
Configuration consistency status : success
vPC role : secondary, operational primary
Number of vPCs configured : 4
Peer Gateway : Disabled
Dual-active exluded VLANs : -

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status  Active  vlans
---  -----  ------- ------------------------------------------------
1    Po500  down    -
```

vPC between switch1 and switch2 is not working. Which two actions are needed to fix the problem? (Choose two.)

A. Match vPC domain ID between the two devices.

B. Configure IP address on the interface.

C. Activate VLANs on the vPC.

D. Configure vPC peer link and vPC peer keepalive correctly.

E. Configure one of the switches as primary for the vPC.

ANSWER: A C

Explanation:

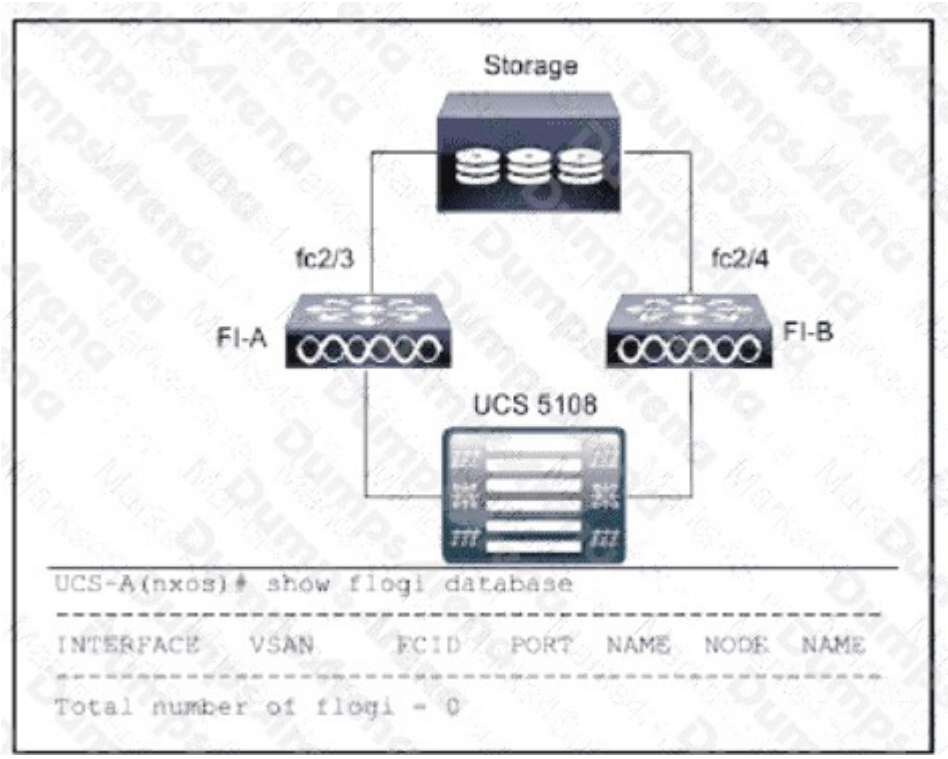
Match vPC domain ID between the two devices because both Nexus switches must be configured as peers in the same vPC domain. The domain configuration establishes the vPC relationship and contains the peer-keepalive parameters and system-level settings used to coordinate the two switches. A domain-ID inconsistency prevents the switches from forming a valid

vPC domain relationship, so the participating vPC cannot become operational. Cisco documents the vPC domain as a required common configuration element on both peers: Cisco Nexus vPC configuration guide.

Activate VLANs on the vPC because a VLAN used for traffic over a vPC must be present and active on both vPC peers and permitted on the relevant vPC member port channel. VLAN availability must also be consistent through the vPC topology so that either peer can forward traffic without creating a split forwarding condition. After making these changes, validate that the vPC peer adjacency and peer link are operational, verify VLAN state and trunk VLAN allowance on both switches, and use `show vpc` and `show vpc consistency-parameters vlan` to confirm a consistent operational state. Cisco's vPC troubleshooting guidance describes these consistency and VLAN validation checks: Cisco vPC troubleshooting technote.

QUESTION NO: 26

Refer to the exhibit.



The Fabric Interconnect fails to communicate with the storage device. Which type of port must interface fc2/3 be configured as to resolve the issue?

A. FCoE uplink

B. FC storage

C. uplink

D. appliance

ANSWER: B

Explanation:

FC storage is the required port type when interface fc2/3 on a Cisco UCS Fabric Interconnect is physically connected directly to an FC storage array. An FC storage port is a SAN-facing target connection: it allows the Fabric Interconnect to establish Fibre Channel logins with storage target ports and present the array's LUNs to UCS service-profile initiators through the applicable VSAN and zoning configuration. After configuring the interface as FC storage, the administrator should ensure that the port is assigned to the intended VSAN, that the storage target is reachable and logged in, and that the SAN zoning configuration permits the required initiator-to-target communication. Cisco UCS documentation distinguishes storage connections from SAN uplinks because each role participates differently in the Fibre Channel topology and control plane. This direct-storage role is specifically intended for attaching FC target devices such as storage arrays to the Fabric

Interconnect. Cisco's UCS configuration documentation and platform guides are available through the Cisco UCS Manager installation and configuration guides. Additional UCS Fibre Channel and storage connectivity guidance is available in the Cisco UCS Manager support documentation.

QUESTION NO: 27

Refer to the exhibit.

An engineer must strictly enforce the port security configuration database in the Fibre Channel fabric. A new device has been connected to the Fibre Channel storage interface that was configured for port security. The port security was enabled for the VSAN that the port belongs to. Despite the configuration, the rogue devices can log into the fabric. Which action resolves the issue?

A. Commit the active database with the port-security database copy vsan 1 command

B. Disable auto learn in VSAN 1 with the no port-security auto-learn vsan 1 command

C. Disable CFS distribution with the no port-security distribute vsan 1 command

D. Commit the changes with the port-security commit copy vsan 1 command

ANSWER: B

Explanation:

Disable auto learn in VSAN 1 with the `no port-security auto-learn vsan 1` command. Cisco MDS port security restricts fabric logins by validating a device's pWWN against the port-security database for the relevant VSAN. However, auto-learn changes that enforcement behavior: when an unrecognized device logs in, its pWWN can be learned dynamically and added to the database. Consequently, merely enabling port security does not provide a fixed allow-list while auto-learn remains enabled. Disabling auto-learn prevents new, unapproved devices from being automatically admitted through database learning. The administrator can then explicitly control authorized pWWNs in the port-security database and distribute the approved policy as appropriate for the fabric. This produces the required strict enforcement, because a newly connected device whose pWWN is not already authorized cannot successfully log in merely by appearing on the storage-facing interface. Cisco documents auto-learning as a port-security capability that learns device WWNs and notes that it must be disabled when the database is intended to be manually controlled. See the Cisco MDS Security Configuration Guide: Configuring Port Security and the Cisco MDS 9000 Series Command Reference.

QUESTION NO: 28

You have an ACI environment with three APICs, two spine switches, and four-leaf switches. You wipe and reboot all APICs first then leaf and spine switches one-by-one and successfully go through the Initial Setuo dialogue on the APIC 1 CIMC KVM console. When you log in to the APIC1 WebGUI, you notice that you do not see any directly connected leaf switches being discovered under Fabric > Inventory > Fabric Membership. What is the cause of the issue?

A. The leaf nodes were not erased properly, which caused a fabric parameters mismatch with the APIC1.

B. Rebooting the APICs and the leaf and spine switches after wiping them is not required.

C. You forgot to enter the TEP Pool value during the Initial Setup dialogue on APIC1.

D. The same Fabric Name value should be used before after wiping all devices.

ANSWER: A

Explanation:

The leaf nodes were not erased properly, which caused a fabric parameters mismatch with the APIC1. ACI switches retain local fabric-related information until their configuration is correctly cleaned. This information can include the prior fabric identity and bootstrap parameters used to join the earlier APIC-controlled fabric. After APIC1 is rebuilt and initialized as a

new fabric, a leaf that still holds inconsistent prior fabric state cannot successfully complete the discovery and registration process. Consequently, it does not appear in Fabric Membership even when it is physically connected to the fabric.

The appropriate recovery action is to use the documented ACI switch-cleaning procedure on the affected leaf nodes, reload them, and allow them to bootstrap against the newly initialized APIC fabric. During successful bootstrap, the APIC assigns and provisions the fabric-specific information required for the node to join and become visible in the inventory. Cisco's guidance for investigating undiscovered nodes emphasizes validating switch cleanup and the parameters used during fabric discovery. See Cisco: Troubleshoot ACI Switches That Are Not Discovered by the APIC and the Cisco APIC Basic Configuration Guide.

QUESTION NO: 29

Refer to the exhibit.

An engineer troubleshoots the HSRP configuration and notes that the remote end uses HSRP version 1. The engineer sets the local HSRP to version 1, but the problem continues. Which action resolves the issue?

A. Change the local group number to 255 or less and request that the remote data center matches the group number in its configuration.

B. Set the local group number as 300 in the remote data center configuration

C. Run version 1 and version 2 HSRP on VLAN300 in the local data center.

D. Reduce the priority of the local HSRP to below 100 to force the local HSRP to standby and then change the version to version 1.

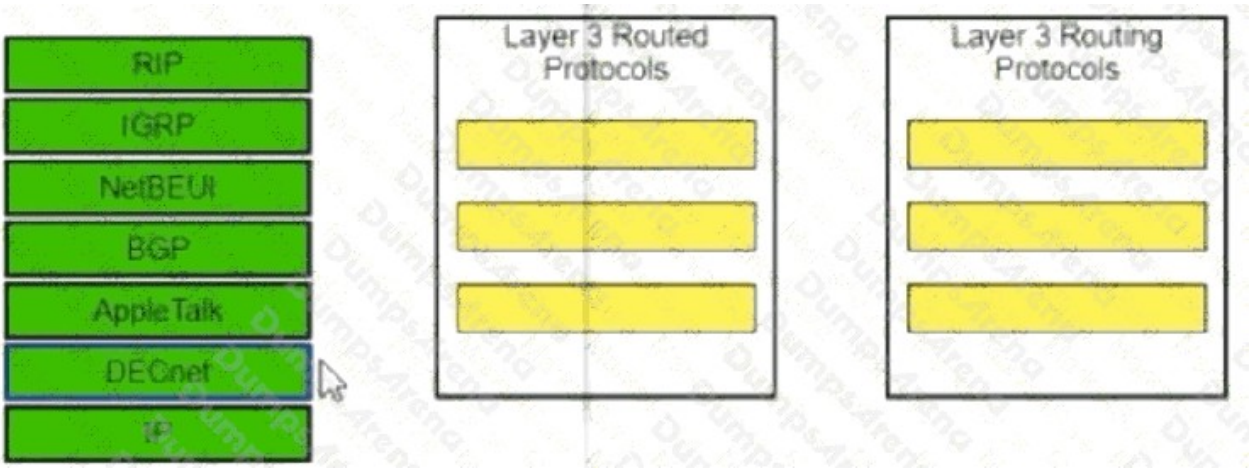
ANSWER: A

Explanation:

Change the local group number to 255 or less and ensure that the remote data center uses the same group number in its configuration. HSRP routers form a redundancy relationship only when the relevant HSRP parameters align on the shared Layer 2 segment, including the HSRP version and group number. Although HSRP version 2 supports a substantially expanded group-number range, HSRP version 1 supports group numbers only from 0 through 255. Therefore, changing only the local router to version 1 does not create a valid HSRP relationship if the configured group remains above 255, such as group 300. The local group must be moved into the version 1-supported range, and the peer must be configured for that same valid group so both devices participate in the same HSRP group and elect active and standby roles correctly. This compatibility requirement follows the HSRP version 1 protocol definition, which encodes the group field as an 8-bit value, and Cisco's HSRP guidance on version-specific behavior and configuration. After making the change, verify the group state and peer discovery with the appropriate `show hsrp` command. See RFC 2281: Cisco Hot Standby Router Protocol and Cisco HSRP Frequently Asked Questions.

QUESTION NO: 30 - (SIMULATION)

Drag each item listed on the left to its proper category on the right Not all items will fit the categories.



ANSWER: See the explanation for the answer

Explanation:

Place the items as follows:

`NetBEUI` does not fit either category. It is a nonroutable protocol and is not a Layer 3 routed protocol or a routing protocol.

Routed protocols define addressing and carry end-user data between networks, whereas routing protocols exchange reachability information to build routing tables.

QUESTION NO: 31

Refer to the exhibit.

```
MDS-1#
zoneset name ESX vsan 10
 zone name ESX01_B vsan 10
  * fcid 0x010a0e [pwwn 20:00:00:25:b5:aa:41:ba]
    fcid 0x840001 [pwwn 52:4a:93:7e:47:1d:ce:02]
    fcid 0x840021 [pwwn 52:4a:93:7e:47:1d:ce:06]
VSAN 10:
FCID    TYPE PWWN                (VENDOR)     FC4-TYPE:FEATURE
0x840041  N    52:4a:93:7e:47:1d:ce:04         scsi-fcp:target
0x840061  N    52:4a:93:7e:47:1d:ce:08         scsi-fcp:target
0x010a0e  N    20:00:00:25:b5:aa:41:ba         scsi-fcp:init
VSAN 20:
FCID    TYPE PWWN                (VENDOR)     FC4-TYPE:FEATURE
0x840001  N    52:4a:93:7e:47:1d:ce:02         scsi-fcp:target
0x840021  N    52:4a:93:7e:47:1d:ce:06         scsi-fcp:target
FI-A /org/service-profile #
vHBA:
    Name: vHBA1
    Fabric ID: B
    Dynamic WWPN: 20:00:00:25:B5:BB:41:BA
    Equipment: sys/chassis-1/blade-1/adaptor-1/host-fc-1
    Template Name: vHBA_SAN_B
    Oper Nw Templ Name: org-root/san-conn-templ-vHBA_SAN_B
    Adapter Policy: Windows
    Oper Adapter Policy: org-root/fc-profile-Windows
vHBA Template:
    Name                 Type              Fabric ID
    -------------------- ----------------- ---------
    vHBA_SAN_B           Updating Template B
FI-B(nxos)#
-------------------------------------------------------------------------------
INTERFACE VSAN FCID             PORT NAME               NODE NAME       EXTERNAL
-------------------------------------------------------------------------------
vfc2981   10   0x010a0e 20:00:00:25:b5:bb:41:ba 20:00:00:25:b5:a4:00:c9 San-po11
```

A blade that is connected to FI-B fails to boot from a SAN LUN. The boot policy is configured to boot by using vHBA1 as the SAN Primary. Which set of actions resolves the issue?

A. â€¢ Set ESX01_B zone to VSAN 20.â€¢ Create a new vHBA and assign VSAN 20.

B. â€¢ vHBA1 to VSAN 20 in Service Profile.â€¢ VMware Adapter Policy on vHBA1.

C. â€¢Apply VMware Adapter Policy on vHBA1â€¢Change vHBA1 to VSAN 20 in Service Profile

D. â€¢ Set ESX01_B zone to VSAN 20.â€¢ Change vHBA1 to VSAN 20 in Service Profile

ANSWER: D

Explanation:

Set ESX01_B zone to VSAN 20 and change vHBA1 to VSAN 20 in the Service Profile. For SAN boot to work, the vHBA used by the boot policy must be associated with the same VSAN as the Fibre Channel fabric, target storage port, and active zoning configuration. In this case, the blade uses Fabric Interconnect B and the configured primary boot path is vHBA1. Assigning vHBA1 to VSAN 20 places its initiator login in the intended SAN fabric. Moving the ESX01_B zone to VSAN 20 ensures that the initiator represented by that host is zoned with its required storage target ports in that same VSAN. Once both the vHBA and zone use VSAN 20, the fabric can perform name-server registration and zoning-based communication between the server initiator and storage target, allowing the boot LUN to be discovered through the primary SAN path. Cisco UCS service profiles define vHBA VSAN membership, while Fibre Channel zoning is scoped to an individual VSAN; these two elements must align for a boot-from-SAN path to become operational. See the Cisco UCS Manager Configuration Guide and the Cisco MDS NX-OS Fabric Configuration Guide.

QUESTION NO: 32

```
Nexus5k#show system internal vpcm info global | i Sticky|Reload
Sticky Master: TRUE
Reload timer started: FALSE
Reload restore configured: TRUE, timer :240
Nexus5k#
```

Refer to the exhibit. During a maintenance window of the Nexus switch pair, one of the Cisco Nexus chassis was brought offline for maintenance. The primary Cisco Nexus switch was stable during the maintenance window. After finishing the

tasks on the chassis, the administrator wanted to bring the second switch online. The administrator checked the status of the disconnected switch that was removed from the vPC during the maintenance. The engineer must connect the peer chassis seamlessly with the primary vPC peer without experiencing traffic interruption. Which action accomplishes this goal?

A. Change the role priority, then connect the peer switch.

B. Change the auto recovery to disabled, and enable auto recovery.

C. Change the auto recovery timer to active if not already set, then re-establish the connectivity.

D. Change the auto timer to a higher value, then connect the peer switch.

ANSWER: D

Explanation:

Changing the auto timer to a higher value before connecting the peer switch provides the surviving vPC peer sufficient time to maintain its established operational role and forwarding state while the returning chassis initializes and re-establishes the vPC control-plane relationships. The vPC auto-recovery mechanism is designed to allow a peer to recover service after its peer has been unavailable for the configured recovery interval. In a planned maintenance scenario, the interval must account for the expected boot, configuration, and peer-link restoration time of the returning chassis. Using a higher timer prevents premature recovery behavior while the peer is still coming back online, helping preserve the active vPC state and avoiding an unnecessary forwarding-state transition for downstream vPC-attached devices. Once peer connectivity and consistency

are restored, the returning switch can rejoin the vPC domain in an orderly manner, with the stable peer continuing to forward traffic throughout the process. Cisco documents auto-recovery as a vPC resiliency feature and describes configuring its reload-delay value according to the recovery requirements of the deployment. See the Cisco Nexus LAN Configuration Guide: vPC configuration and the Cisco Nexus vPC auto-recovery documentation.

QUESTION NO: 33

```
vrf context management
 ip name-server 4.2.2.2
 ip route 0.0.0.0/0 192.168.30.2

interface mgmt0
 ip address dhcp
 vrf member management

N9K-Core# ping google.com vrf management
PING google.com (216.58.209.238): 56 data bytes
64 bytes from 216.58.209.238: icmp_seq=0 ttl=127 time=151.982 ms
64 bytes from 216.58.209.238: icmp_seq=1 ttl=127 time=136.198 ms
64 bytes from 216.58.209.238: icmp_seq=2 ttl=127 time=224.796 ms
64 bytes from 216.58.209.238: icmp_seq=3 ttl=127 time=148.458 ms
64 bytes from 216.58.209.238: icmp_seq=4 ttl=127 time=129.98 ms

- - - google.com ping statistics - - -
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 129.98/158.282/224.796 ms
```

Refer to the exhibit. A network engineer connects the Cisco Nexus switch management port to the Internet using DHCP to allow the Guest shell that runs on the switch to download Python packages. The engineer can ping google.com from the Cisco Nexus switch, but the Guest shell fails to download Python packages. Which action resolves the problem?

A. Update the Python packages directly on the Cisco Nexus switch.

B. Manually configure DNS in the Guest shell, even if it is claimed on the Cisco Nexus switch through DHCP.

C. Manually configure NTP in the Guest shell.

D. Connect the Guest shell to data plane interfaces to be able to connect to the networks outside of the Cisco Nexus switch.

ANSWER: B

Explanation:

Manually configure DNS in the Guest shell, even if it is claimed on the Cisco Nexus switch through DHCP. The NX-OS management interface and the Guest Shell are separate environments with independent network configuration. DHCP information obtained by NX-OS for the management interface, including DNS-server settings, enables the switch itself to resolve a name such as `google.com`, but those settings are not automatically inherited by the Guest Shell container. Package-management tools such as `pip` must resolve repository hostnames before they can download packages. Therefore, the Guest Shell needs its own valid resolver configuration, typically by adding reachable DNS servers to its `/etc/resolv.conf` file or configuring DNS through the supported Guest Shell network setup. Once the Guest Shell has DNS reachability through its management network connection, it can resolve Python package repository names and retrieve packages. Cisco documents that Guest Shell is a Linux container hosted on NX-OS and uses its own networking and configuration context; NX-OS system-management behavior, including management-interface addressing and DNS, is documented separately. See the Cisco Nexus 9000 Series NX-OS Programmability Configuration Guide and the Cisco Nexus 9000 Series NX-OS System Management Configuration Guide.

QUESTION NO: 34

An engineer investigates an FCoE connectivity failure between a Cisco Nexus switch and a converged network adapter. The Ethernet interface is up, but the virtual Fibre Channel interface remains down. Which two conditions must be verified for FCoE initialization to complete? (Choose two.)

A. Verify that the bound Ethernet interface is operational.

B. Verify that the FCoE VLAN is allowed and tagged on the Ethernet trunk.

C. Configure the FCoE VLAN as the native VLAN.

D. Disable LLDP on the bound Ethernet interface.

E. Configure the vFC interface as a routed interface.

ANSWER: A B

Explanation:

The Ethernet interface bound to the virtual Fibre Channel interface must be operational. A vFC interface depends on its bound Ethernet interface to transport FIP and encapsulated Fibre Channel traffic, so it cannot come up if that Ethernet interface is down or incorrectly bound.

The FCoE VLAN must be allowed and tagged on the Ethernet trunk. FIP uses the FCoE VLAN during discovery and virtual-link initialization. If the VLAN is not carried across the trunk, or if it is configured as the native VLAN and sent untagged, the CNA cannot complete the expected FIP exchange. Cisco documents vFC binding and FCoE VLAN requirements in the Cisco Nexus 5000 FCoE Configuration Guide.

QUESTION NO: 35

A fabric interconnect fails to start, and the console displays the loader prompt. Which two actions resolve the issue? (Choose two.)

A. Load an uncorrupt bootloader image.

B. Load an uncorrupt kickstart image.

C. Reconnect Layer 1 and Layer 2 cables between the FIs.

D. Reformat the fabric interconnect.

E. Load the correct version of the boot image.

ANSWER: B E

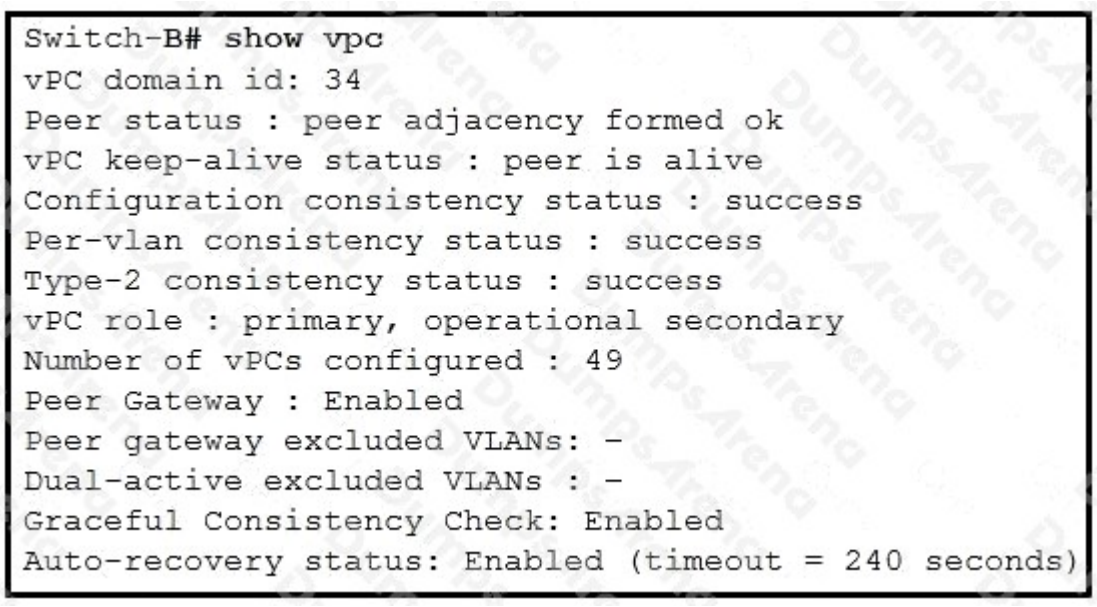
Explanation:

The loader prompt indicates that the fabric interconnect has reached its low-level boot environment but cannot continue through the normal NX-OS startup sequence. Recovery therefore requires valid software images that the loader can locate and execute. **Load an uncorrupt kickstart image.** is required because the kickstart image initializes the platform hardware and provides the initial NX-OS services needed to proceed with booting. If that image is damaged or unavailable, the fabric interconnect cannot advance beyond the loader environment.

Load the correct version of the boot image. is also required because the image used after kickstart must be a valid, compatible release for the fabric interconnect and its installed software set. The kickstart and subsequent boot/system image operate as a matched image pair; using the intended version allows the device to complete initialization and load Cisco UCS software normally. Image recovery from the loader is a software restoration procedure, typically using an accessible local or network image source, followed by verification that the configured boot variables point to the recovered images. Cisco documents the UCS software image and upgrade architecture in the Cisco UCS upgrade documentation and provides platform support resources at Cisco Unified Computing System support.

QUESTION NO: 36

Refer to the exhibit.



```
Switch-B# show vpc
vPC domain id: 34
Peer status : peer adjacency formed ok
vPC keep-alive status : peer is alive
Configuration consistency status : success
Per-vlan consistency status : success
Type-2 consistency status : success
vPC role : primary, operational secondary
Number of vPCs configured : 49
Peer Gateway : Enabled
Peer gateway excluded VLANs: -
Dual-active excluded VLANs : -
Graceful Consistency Check: Enabled
Auto-recovery status: Enabled (timeout = 240 seconds)
```

After a failover occurs, which two actions must be performed on Switch-B to manually preempt the operational primary role back to Switch-A? (Choose two.)

A. Configure the local vPC role priority to have a lower value than Switch-A.

B. Configure the local vPC role priority to have a higher value than Switch-A.

C. Disable and then re-enable the vPC peer-keepalive link.

D. Configure the local vPC role priority to have the same value as Switch-A.

E. Disable and then re-enable the vPC peer link.

ANSWER: B E

Explanation:

Configuring the local vPC role priority to have a higher value than Switch-A ensures that Switch-A has the numerically lower vPC role-priority value. In vPC, the switch with the lower configured role priority is elected as the primary after a role election. Therefore, assigning Switch-B a higher value establishes Switch-A as the preferred primary switch for the next election. The priority change alone does not necessarily cause an immediate operational role transition while the peer relationship remains established. Disabling and then re-enabling the vPC peer link forces the required peer-role reevaluation, allowing the configured preference to take effect and restoring Switch-A to the operational primary role. This is a controlled manual preemption procedure: first establish the intended role-election preference, then trigger a new election through peer-link reset. Cisco documents that vPC role priority is used to elect the primary peer and that the lower value wins the election. Review the peer status and operational role afterward with `show vpc role` and `show vpc` to confirm that Switch-A is primary and that the peer link has returned to a healthy state. See Cisco's Nexus 9000 vPC Configuration Guide and the Nexus 9000 vPC configuration documentation.

QUESTION NO: 37

Refer to the exhibit.



A Cisco UCS Fabric Interconnect fails during the upgrade process. The working images of the fabric interconnect are stored on the bootflash. Which set of commands recovers the fabric interconnect?

A. loader> dir
loader> boot ucs-6300-k9-kickstart.5.0.2.N1.3.02d56.bin
switch(boot)# init system
switch(boot)# reload
switch(boot)# load ucs-6300-k9-system.5.0.2.N1.3.02d56.bin

B. loader> dir
loader> boot ucs-6300-k9-kickstart.5.0.2.N1.3.02d56.bin
switch(boot)# copy ics-manager-k9.1.4.1k.bin nuova-sim-mgmt-nsg.0.1.001.bin
switch(boot)# boot ucs-6300-k9-system.5.0.2.N1.3.02d56.bin

C. loader> dir
loader> boot ucs-6300-k9-kickstart.5.0.2.N1.3.02d56.bin
switch(boot)# init system
switch(boot)# reload
switch(boot)# boot ucs-6300-k9-system.5.0.2.N1.3.02d56.bin

D. loader> dir
loader> boot ucs-6300-k9-kickstart.5.0.2.N1.3.02d56.bin
switch(boot)# copy ics-manager-k9.1.4.1k.bin nuova-sim-mgmt-nsg.0.1.001.bin
switch(boot)# load ucs-6300-k9-system.5.0.2.N1.3.02d56.bin

ANSWER: D

Explanation:

The recovery sequence that boots `ucs-6300-k9-kickstart.5.0.2.N1.3.02d56.bin`, copies `ics-manager-k9.1.4.1k.bin` to `nuova-sim-mgmt-nsg.0.1.001.bin`, and then loads `ucs-6300-k9-system.5.0.2.N1.3.02d56.bin` is correct. At the loader prompt, the kickstart image supplies the minimal NX-OS environment required to access bootflash and perform recovery actions. Cisco UCS Manager relies on its management image being available under the expected internal filename, so the copy operation restores that required image association before the full system image is loaded. The `load` command then starts the system image from bootflash, allowing the Fabric Interconnect to complete startup with the known-working software set rather than continuing the failed upgrade path. This procedure preserves the recovery focus on locally available images and avoids requiring network-based image retrieval

while the Fabric Interconnect is in loader mode. Cisco documents Fabric Interconnect software and recovery-related procedures in the Cisco UCS Manager documentation; platform support information is available on the Cisco UCS 6300 Series Fabric Interconnect support page.

QUESTION NO: 38

Refer to the exhibit.

An engineer is configuring a VMM domain integration, but the operation fails to complete. Which two actions resolve the issue? (Choose two.)

A. Use a different device package.

B. Configure the route from virtual machine managerB management network to vCenter.

C. Update the VMware vCenter version to be compatible with the Cisco ACI release.

D. Change the VMware vCenter access credentials.

E. Change VMware vCenter IP address.

ANSWER: B C

Explanation:

Configuring a route from the virtual machine managerB management network to vCenter restores the required IP reachability between the management components involved in the VMM integration. Cisco ACI must be able to establish management-plane communication with the VMware vCenter server to authenticate, retrieve inventory information, and create or update the distributed virtual switch configuration. The path must support the required HTTPS communication and return traffic; therefore, adding a valid route is a direct resolution when the management network lacks a route to vCenter.

Updating the VMware vCenter version to be compatible with the Cisco ACI release is also required because ACI VMM integration depends on the VMware APIs and vSphere versions supported by the installed APIC software release. A vCenter release outside the APIC compatibility matrix can prevent successful controller registration or subsequent VMM-domain provisioning even when IP connectivity is present. Before integrating, verify the exact APIC release and installed patches, then select a vCenter version listed as supported for that release. Cisco publishes supported software combinations and related interoperability information in its APIC device support and compatibility documentation. Cisco's APIC Virtualization Guide also describes VMware VMM integration requirements.