ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

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

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

A user wishes to publish a VMware Cloud Foundation (VCF) Operations Orchestrator workflow to their VCF Automation project catalog, but Is blocked from publishing any workflows.

The following information has been provided:

- In the VCF Automation Organization portal, the user cannot see the Workflows option under Content Hub.
- The organization is not a Provider Consumption Organization.

Which are the two likely causes of this issue? (Choose two.)

A. An external VCF Operations Orchestrator is not integrated with their Organization.

B. The user is logged in with Project User rights.

C. The user is logged in the Project Advanced User rights

D. An embedded VCF Operations Orchestrator is not integrated with their Organization.

E. The user is logged in with Project Administrator rights.

Answer: A, D

Explanation:

In VMware Cloud Foundation 9.0, publishing a VCF Operations Orchestrator workflow to a VCF Automation project catalog requires that the Organization has a valid integration with VCF Operations Orchestrator. The question states that the user cannot see the Workflows option under Content Hub, and the organization is not a Provider Consumption Organization (PCO). According to the VCF 9.0 documentation, only organizations with VCF Operations Orchestrator integration are allowed to publish workflows into the catalog. Both embedded and external orchestrator integrations must be configured depending on the environment. If no orchestrator (embedded or external) is integrated with the organization, workflows cannot be listed or published. This aligns with the documented VCF Automation and VCF Operations Orchestrator design requirements, which

specify that workflow publishing is only available when the orchestrator instance is properly registered.

Additionally, user role permission issues could prevent workflow visibility, but the key blockers described in the scenario are the missing workflow section and the organization type. Because the organization is not a PCO, advanced provider features—including workflow publishing—are disabled unless a proper orchestrator integration exists. Therefore, the two most likely causes are:

A: An external VCF Operations Orchestrator is not integrated with their Organization.

D: An embedded VCF Operations Orchestrator is not integrated with their Organization.

These two conditions directly match the documented behavior in VMware Cloud Foundation 9.0.

Question: 2

An administrator wants to expand a VMware vSAN cluster in a workload domain by adding an unassigned host from the vSphere client. However, at the Host Selection screen no hosts are available and the following message displayed:

No unassigned hosts available with storage type VSAN. Commission hosts with physical NICs 0 & 1 to Add Host from UI.

How can the administrator commission hosts?

A. From the vSphere client by navigating to Supervisor Management.

B. From VCF Operations by navigating to Fleet Management.

C. From the SDDC manager by navigating to Workload Domains.

D. From the vSphere client by navigating to the Global Inventory.

Answer: C

Explanation:

In VMware Cloud Foundation 9.0, host commissioning is performed exclusively through SDDC Manager, not from the vSphere Client. When expanding a vSAN cluster inside a workload domain, all ESXi hosts must first be placed in an Unassigned state and then commissioned in SDDC Manager before they can appear in the "Add Host" wizard of the vSphere Client. The message in the problem—"No unassigned hosts available with storage type VSAN. Commission hosts with physical NICs 0 & 1 to Add Host from UI"—indicates that SDDC Manager has not yet commissioned any suitable hosts with the required NIC layout.

VCF 9.0 documentation states that for workload domain expansion, hosts must be commissioned under:

SDDC Manager → Workload Domains → (Select WLD) → Hosts → Commission Hosts.

This validates hardware, storage type (such as vSAN ESA or OSA), NIC placement, and ensures the host is compatible with the domain's configuration.

Options pointing to vSphere Client (A, D) or VCF Operations (B) do not perform the commissioning workflow. Therefore, the correct and verified answer is C, the only interface where host commissioning is officially supported.

Question: 3

An administrator is responsible for a VMware Cloud Foundation (VCF) fleet. The administrator has been tasked with commissioning four ESX hosts for a new workload domain that uses vSAN Express Storage Architecture (ESA) as the primary storage solution.

During the host validation stage in vSphere client, the process fails with the following errors:

esx-l.wld.vcf.local. Failed to validate vSAN HCL status.

esx-2.wld.vcf. local. Failed to validate vSAN HCL status.

esx-3.wld.vcf.local. Failed to validate vSAN HCL status.

esx~4.wid.vcf. local. Failed to validate vSAN HCL status.

What Is the cause of the errors?

A. The RAID controller in each ESX host is not configured to use RAID-O/Passthrough.

B. The ESX hosts are not using vSAN ESA certified storage devices.

C. The ESX hosts must have internet access to validate vSAN ESA compatibility.

D. The RAID controller in each ESX host needs to be reconfigured to use Tri-mode.

Answer: B

Explanation:

VMware Cloud Foundation 9.0 requires strict vSAN ESA hardware compatibility when creating a workload domain that uses vSAN Express Storage Architecture (ESA). During host validation, SDDC Manager and vSphere Client check whether each ESXi host meets ESA requirements, including CPU generation, storage controller type, and—most importantly—ESA-certified NVMe storage devices. The validation errors provided:

"Failed to validate vSAN HCL status" for every host

indicate that the hosts do not meet the vSAN ESA HCL requirements.

VCF 9.0 documentation states that ESA uses a next-generation log-structured filesystem requiring certified NVMe devices only, with no RAID controller dependencies. Unlike OSA, ESA eliminates disk groups, but it requires certified devices listed on the vSAN ESA HCL to pass host validation. If noncertified or unsupported NVMe/SAS devices are present, validation fails exactly as described.

Option A is incorrect because RAID pass-through settings apply to OSA, not ESA.

Option C is incorrect because ESA compatibility validation is performed ofline using the SDDC Manager BOM, not via internet lookup.

Option D is incorrect because ESA does not use tri-mode RAID controllers.

Therefore, the documented and verified cause is B: hosts are not using vSAN ESA certified storage devices.

Question: 4

An administrator has a vSphere 8.0 update 3 environment with the following configuration:

- A 3-node vSAN cluster
- A vSphere Standard Switch (VSS)
- Several standalone ESX hosts in the vCenter inventory

They want to convert this vSphere environment into a new VMware Cloud Foundation (VCF) 9.0 management domain.

Identify two changes they will need to make before converting this vSphere environment into a VMWare Cloud Foundation (VCF) Management domain? (Choose two.)

A. Remove the vSphere Standard Switch from the vCenter Inventory.

B. Upgrade vSphere 8.0 Update 3 to vSphere 9.0.

C. Configure a vSphere Distributed Switch.

D. Remove the standalone hosts from the vCenter inventory.

Answer: BC

Explanation:

To convert an existing vSphere environment into a VMware Cloud Foundation (VCF) 9.0 Management Domain, several prerequisites must be met as defined in the VCF 9.x documentation.

First, VCF 9.0 requires vSphere 9.0 as part of its Bill of Materials (BOM). The uploaded VCF 9.0 documentation confirms that VCF 9.0 is built on vSphere 9.0, vCenter 9.0, and NSX versions that align with the 9.x stack. A vSphere 8.0 Update 3 environment is not supported as a foundation for a VCF 9.0 management domain; therefore, the administrator must upgrade the entire vSphere platform to vSphere 9.0 before VCF deployment. (Reference: VCF 9.0 BOM — vSphere 9.0 is mandatory.)

Second, VCF management domain creation strictly requires vSphere Distributed Switches (vDS). VCF does not support vSphere Standard Switches (VSS) for any management domain hosts. The VCF 9.0 design and deployment guides state that all ESXi hosts intended for a management domain must use vDS for management, vSAN, and vMotion networking. Therefore, the existence of a VSS must be corrected by deploying and configuring a vSphere Distributed Switch and migrating host networking accordingly before Cloud Builder deployment.

Removing standalone hosts or removing a VSS from inventory is not required. Only the hosts selected for the management domain need to be prepared.

Thus, the required changes are:

✓ B. Upgrade vSphere 8.0 Update 3 to vSphere 9.0

✓ C. Configure a vSphere Distributed Switch

These are the only changes explicitly required by VCF 9.0 documentation.

Question: 5

An administrator determined that the VMware NSX admin password expired on their VMware NSX Edge Transport nodes. The administrator manually resets the password in the console of each Edge Transport node.

What additional action is required to synchronize the new password in VMWare Cloud Foundation (VCF) Operations?

A. In VCF Operations, rotate the admin password for each NSX Edge Transport node.

B. In VCF Operations, remediate the admin password for each NSX Edge Transport node.

C. In VCF Operations, sync the admin password for each NSX Edge Transport node.

D. In VCF Operations, update the admin password for each NSX Edge Transport node.

Answer: B

Explanation:

In VMware Cloud Foundation 9.0, password changes made manually on an NSX Edge Transport Node are not automatically synchronized with VCF Operations. VCF Operations maintains secure credential records for all managed components, including NSX Manager appliances and NSX Edge Transport Nodes. When credentials become stale—such as after a password expiration and manual reset—VCF Operations marks the credential object as out of sync and requires administrative remediation.

The official workflow described in VCF 9.0 Operations documentation states that administrators must use the "Remediate Password" function whenever a password was changed outside of VCF Operations, ensuring that the platform revalidates and updates the stored credentials used for monitoring, log collection, and automation tasks. Options such as "rotate," "sync," or "update" do not apply because rotation implies generating a new password managed by VCF, and "sync" does not overwrite the stored credential. Only remediation forces VCF Operations to re-validate and align credentials with the external system.

Therefore, after manually resetting the NSX Edge admin password, the administrator must perform password remediation in VCF Operations to restore operational consistency, making B the correct and verified answer.

Question: 6

An administrator Is responsible for managing a VMware Cloud Foundation (VCF) fleet. The administrator discovers intermittent performance issues with the supplemental storage (ISCSI) connected to VCF workload domain. The administrator discovers that the (iSCSI) target is reachable from most VMware ESX hosts, but some hosts consistently experience periods of slow I/O and connection drops.

Which two actions should the administrator take to diagnose and resolve this issue? (Choose two.)

A. Review the iSCSI target's configuration to ensure it's configured for maximum performance, including enabling

CHAP authentication.

B. Examine the iSCSI VMkernel port on all affected ESX hosts for TCP retransmissions and checksum ofload errors.

C. Update the network plugin on the ESX host to the latest version.

D. Ensure all ESX hosts have the VMkernel port MTU set to 1500.

E. Ensure all ESX hosts have the VMkernel port MTU set to 9000.

Answer: B, E

Explanation:

To diagnose and resolve the intermittent performance and connection drop issues with the supplemental iSCSI storage, the administrator should focus on network layer consistency and health, particularly regarding packet size (MTU) and delivery (TCP).

Examine the iSCSI VMkernel port for TCP retransmissions (Action B - Diagnose): "Intermittent" connection drops and slow I/O are classic symptoms of packet loss or fragmentation issues. By examining the ESXi network stats (e.g., using esxtop key n or viewing vSphere performance charts) for TCP retransmissions, the administrator can confirm if packets are being dropped or lost in transit. Checksum ofload errors can also indicate issues where the NIC hardware is incorrectly validating packets, causing the OS to drop them. This step identifies the root cause (packet loss/corruption).

Ensure all ESX hosts have the VMkernel port MTU set to 9000 (Action E - Resolve): For high-performance storage traffic like iSCSI in a VMware Cloud Foundation environment, it is best practice to use Jumbo Frames (MTU 9000) end-to-end (Host -> Switch -> Storage Array).

The symptom that some hosts are affected suggests configuration drift where those specific hosts might be set to a different MTU (e.g., 1500) or are mismatched with the physical network/target (which is likely set to 9000 for performance).

An MTU mismatch (e.g., Target sending 9000-byte frames to a Host/Switch expecting 1500) typically results in the "Do Not Fragment" (DF) bit causing packet drops, leading to the reported connection drops and retransmission delays. Ensuring a consistent MTU of 9000 across the fleet resolves this and aligns with VCF performance standards.

Note: Option A (CHAP) is for authentication security, not performance. Option C (Update network plugin) is a lifecycle task but less likely to be the immediate fix for "some hosts" having intermittent drops compared to the common issue of MTU mismatch. Option D (MTU 1500) would resolve drops if the physical network doesn't support Jumbo Frames, but would degrade performance, making E the preferred resolution for a "performance" storage tier.

Question: 7

An administrator has been tasked with expanding an existing VMware Cloud Foundation (VCF) workload domain by adding a new cluster. The VCF fleet has the following configuration:

- Three workload domains, including the management domain are configured.

- The management domain (WLD-01) and one of the workload domains (WLD-02) are running VCF 9.0.
- The other workload domain (WLD-03) is running VCF 5.2.1 and is an isolated workload domain.

When attempting to perform the required steps using the vSphere Client UI the cluster cannot be added to the WLD-02 workload domain. What step should the administrator perform to complete the workload domain expansion?

A. Use the SDDC Manager UI to create the cluster in WLD-02.

B. Use the SDDC Manager API to create the cluster in WLD-03.

C. Use the vSphere Client UI to create the cluster in WLD-03.

D. Use the VCF Operations Fleet Manager UI to create the cluster in WLD-02.

Answer: D

Explanation:

VMware Cloud Foundation 9.0 introduces a major architectural redesign that replaces the traditional SDDC Manager–centric domain management model with a unified Fleet Management architecture implemented through VCF Operations Fleet Manager. In this model, each Workload Domain operates with its own vCenter, but Enhanced Linked Mode (ELM) is removed to improve isolation, reduce blast radius, and support multi-site scalability. As a result, administrators logged into the vSphere Client of the Management Domain can no longer manage or expand clusters in other Workload Domains, which explains why the vSphere UI blocks the attempted expansion of WLD-02.

Fleet Manager becomes the new authoritative control plane for lifecycle, topology, host commissioning, and workload domain expansion. Only Fleet Manager maintains the full global view necessary to orchestrate cluster addition operations across distributed vCenters and domains.
Because WLD-02 is running VCF 9.0 and is fully fleet-aware, its expansion must occur through VCF Operations Fleet Manager, not through the vSphere Client or legacy SDDC Manager workflows.

Options involving WLD-03 are invalid since that domain is running VCF 5.2.1, is isolated, and cannot participate in fleet-aware operations. SDDC Manager (A) is no longer the correct interface for VCF 9.0 domain expansion operations.

Question: 8

An administrator is responsible for managing a VMware Cloud Foundation (VCF) Fleet that is configured as follows:

- Single VCF instance with a single workload domain.
- The Workload Domain has a single 5-node VMware vSAN Express Storage Architecture (ESA) cluster.
- The vSAN Default Storage Policy is configured as RAID1.

The administrator is alerted to the fact that storage capacity is running low and, to improve space efficiency, attempts to change the vSAN storage policy on a number of large virtual machines to a 2 Failures - RAID-6 policy.

The policy change is immediately rejected.

What should the administrator do to reduce overall capacity usage while waiting for new storage devices to arrive?

A. Enable encryption on the vSAN Default Storage Policy.

B. Reconfigure the Virtual Machines to use a 1 Failure-RAID-5 Storage Policy.

C. Convert the Virtual Machines from thick provisioning to thin provisioning.

D. Enable compression on the vSAN Default Storage Policy.

Answer: C

Explanation:

In VMware Cloud Foundation 9.0 with vSAN ESA, storage policies must match the capabilities of the existing cluster. The scenario describes a 5-node vSAN ESA cluster where the vSAN Default Storage Policy is RAID-1 (FTT=1). The administrator attempts to apply a 2 Failures – RAID-6 policy, which ESA supports only on clusters with at least 7 nodes. Because the cluster has only five nodes, the policy fails immediately—this is expected and documented in the vSAN ESA design specifications.

Since RAID-6 is not an option and capacity is low, the administrator must look for a method to reclaim storage usage without requiring additional nodes or unsupported policy changes. Converting VMs from thick provisioning to thin provisioning is a safe and effective mitigation approach. Thin provisioning reduces consumed space by allowing disks to grow only as needed, immediately recovering unused blocks. This is a standard vSAN-supported method to temporarily alleviate capacity pressure.

Enabling encryption (A) or compression (D) does not reduce capacity usage retroactively and may actually increase overhead. Using RAID-5 (B) is also not possible because RAID-5 requires at least 6 ESA-enabled hosts.

Question: 9

An administrator has successfully mounted an NFS datastore as supplemental storage for a VMware Cloud Foundation (VCF) workload domain cluster. However, users report that data cannot be written to the datastore.

The administrator confirms the following:

- The NFS share is visible in the vSphere Client.
- Connectivity to the NFS server from the Virtual Machine.

What action should the administrator take next to troubleshoot the issue?

A. Verify the NFS server is listed in the VMware Hardware Compatibility Guide.

B. Reboot the ESX host to clear any file locks.

C. Verify that the NFS server permissions are not set to read-only for the ESX host.

D. Verify the MTU size configuration on the NFS VMkernel port group.

Answer: C

Explanation:

In VMware Cloud Foundation 9.0, supplemental storage such as NFS is fully supported for workload domains when configured correctly. When an NFS datastore mounts successfully in vSphere but users cannot write data, the issue almost always lies in the export permissions on the NFS server. vSphere will allow mounting a read-only NFS export, but write operations will fail silently at the VM or guest OS level.

VCF documentation confirms that ESXi requires explicit read/write export permissions, typically configured per-host or by IP subnet, on the NFS server. Even if network connectivity and VM-level access appear healthy, incorrect server-side permissions prevent ESXi from executing write operations.

Option A is incorrect because NFS servers are not validated by the HCL for write capability.
Option B (rebooting the host) is unnecessary and unrelated to permission enforcement.
Option D (MTU mismatch) may cause performance issues, not write-access failures.

Thus, the next troubleshooting step is to verify that the ESXi hosts have read/write access on the NFS share, making C the correct answer.

Question: 10

An administrator needs to confirm which account initiates tasks from VMware Cloud Foundation (VCF) Operations. As a test, a virtual machine (VM) is powered on/off through VCF Operations.

In the vCenter task pane, what account would be the initiator of the task?

A. The credentials of the logged in user.

B. The service account between VCF Operations and vCenter.

C. The service account between vCenter and SDDC Manager.

D. The administrator@vsphere.local account.

Answer: B

Explanation:

When VMware Cloud Foundation Operations performs actions on vCenter—such as powering on or off a VM—the tasks are initiated through an integration service account, not the identity of the user logged into the VCF Operations UI. VCF Operations connects to vCenter using a configured collector or integration credential, typically a service account defined during initial setup.

VCF documentation clarifies that all automated or orchestrated tasks originating from VCF Operations use this trusted

account to ensure consistent auditing, RBAC enforcement, and operational isolation from user identities. Therefore, in the vCenter task pane, the "Initiated By" field always reflects the VCF Operations → vCenter service account, even if the end-user triggered the action indirectly.

Option A is incorrect because the logged-in user does not directly interface with vCenter.
Option C refers to SDDC Manager's integration account, which is unrelated to VCF Operations workflows.
Option D (administrator@vsphere.local) appears only when vCenter's built-in admin performs the action.

Question: 11

In VMware Cloud Foundation (VCF) Automation an administrator is troubleshooting an issue with a newly created Organization. When the Organization administrator attempts to create a Namespace, they receive an error "Failed to list VPC after selecting a region.

The administrator logs into the NSX Manager for the Region and does not see an NSX Project for the Organization.
What could cause these symptoms?

A. The Provider Administrator hasn't set up the Organization's Networking Configuration for the selected Region.

B. The Organization Administrator hasn't created a Project in the selected Region.

C. The Provider Administrator hasn't granted the Organization Administrator role to the First User.

D. The Organization Administrator hasn't created a VPC in the selected Region.

Answer: A

Explanation:

In VMware Cloud Foundation 9.0 Automation, every Organization requires a properly configured Networking Configuration for each Region in which it operates. This configuration step —performed by the Provider Administrator— creates the NSX Project corresponding to the Organization, enabling Namespace creation, VPC visibility, and workload provisioning.

The error "Failed to list VPC after selecting a region" combined with the absence of an NSX Project in NSX Manager is a direct indicator that the Organization's Networking Configuration was never initialized. VCF Automation automatically creates the NSX Project only when the Provider Admin completes this step.

Option B is invalid because the Organization Administrator cannot create NSX Projects manually; they are system-generated during networking setup.

Option C is incorrect because role assignment affects administrative permissions, not NSX project creation.

Option D is also incorrect—the Organization Admin cannot create a VPC until the NSX Project exists.

Question: 12

An administrator has been tasked with the deletion of a workload domain within a VMware Cloud Foundation (VCF) instance. The following information has been provided:

- There are two workload domains and a management domain within the VCF instance.
- There is a single vSphere cluster within the workload domain to be deleted.
- There are no user created Virtual Machines in the workload domain cluster.

When performing the deletion in VCF Operations, the task fails at the Gather input for deletion of NSX component stage. The administrator checks the details of the failed task and notices the cause of the error is stated as Cannot read the array length because "<locall9>" is null.

What could be the possible cause of this error message?

A. The NSX Edge Cluster Deployment Removal Tool was run against the workload domain.

B. The NSX Edge cluster for the workload domain was deleted using NSX Manager.

C. The NSX Manager is shared between the workload domains.

D. The Network Pools associated with the workload domain were deleted using the vSphere client.

Answer: B

Explanation:

In VMware Cloud Foundation, deletion of a workload domain requires that VCF Operations can correctly discover and process the NSX components attached to that domain. The workload domain delete workflow explicitly includes removal of the NSX Manager and NSX Edge components associated with the domain, unless those NSX components are shared.

In earlier and current VCF guidance, VMware state that NSX Edge clusters for a workload domain must be removed using the documented/VCF-aware method (for example, using the NSX Edge removal process referenced in KB 78635, not by deleting objects directly in NSX Manager). If an administrator deletes the NSX Edge cluster directly in NSX Manager, the VCF inventory and orchestration logic still “believes” the Edge cluster exists. When the workload domain delete workflow reaches the stage “Gather input for deletion of NSX component”, it queries NSX / internal state for Edge cluster data. Because the underlying object has been manually removed, the returned structure is null, which results in an internal “Cannot read the array length because "<locall9>" is null” style error.

Using the NSX Edge Cluster Deployment Removal Tool as per documentation keeps VCF and NSX in sync and is the supported path, so option A is not the likely cause. Network pools and shared NSX Manager configurations do not match the specific NSX-component array/null condition described.

Question: 13

An administrator has created an alarm for an object in VMware Cloud Foundation (VCF) Operations. The alert does

not show up In the alert pane despite being configured on the object.

Parameters:

- Symptom definition: Read Latency (ms) is higher than 1 ms.
- Alert definition: Alert is triggered as soon as the latency is higher than the 1 ms defined in the symptom definition.
- Object type: Virtual Machine.

What is the reason the alert does not show up in the alert view?

A. The administrator is missing the privileges to view alerts for this object.

B. The metric used in the symptom definition does not apply to this object type.

C. The alert is not enabled in the policy.

D. This type of alert must be forwarded from VMware Cloud Foundation Operations for Logs.

Answer: C

Explanation:

In VMware Cloud Foundation 9.0, VCF Operations (vROps-based) uses policies to control which alerts, symptoms, and metrics are evaluated for a given object. Creating an alert definition and symptom alone is not sufficient; the alert must be associated with and enabled in a policy that is

actively applied to the target object (in this case, a Virtual Machine). The documentation shows that when you create an alert definition, there is an explicit Policies step, where you select the policy (for example, the default policy) so that the alert becomes active for objects governed by that policy.

The metric "Read Latency (ms)" is valid for virtual-machine–related objects: VCF Operations documents Read Latency metrics at the VM disk and VM–datastore link level (for Disk and Datastore metrics on Virtual Machines). Therefore, option B (metric not applicable) is incorrect. No requirement exists that such a performance alert must be forwarded from VCF Operations for Logs (D); log-based alerts are a separate alert type.

If the alert definition is not enabled in the effective policy for that VM, VCF Operations will not evaluate the symptom or generate the alert, and it will not appear in the alert pane—even though the definition technically exists. This matches option C exactly.

Question: 14

An administrator configures a new VMware Cloud Foundation (VCF) instance in a remote site using a vSAN Express Storage Architecture (ESA) for the workload domain cluster. vSAN ESA is configured with Auto-Policy Management and is designed to tolerate a single failure. The cluster experiences a hardware failure and on investigation, the administrator

discovers that the affected objects did not re-protect and remain in a "Reduced availability with no rebuild" state.

How can the administrator explain why the vSAN objects did not rebuild as expected?

A. The storage devices are not certified for vSAN.

B. The number of ESX hosts doesn't support rebuilds during an outage.

C. The storage policy needs to be modified to support forced provisioning.

D. The existing disk groups need to be expanded to support additional capacity.

Answer: B

Explanation:

In VMware Cloud Foundation 9.0, using vSAN Express Storage Architecture (ESA) with Auto-Policy Management, the system automatically selects the correct storage policy based on the cluster size and desired failure protection. When the administrator configures tolerance for a single failure (FTT=1 using RAID-1 mirroring), vSAN ESA requires sufficient remaining hosts during a failure event to reprotect objects.

A minimum of 3 ESA-capable hosts is required for RAID-1, and re-protection after a failure requires enough hosts with available capacity to place new replica components. In small ESA clusters (e.g., 3 or 4 nodes), if one host fails, the remaining hosts may not meet the placement rules for automatic rebuild to restore compliance. ESA enforces strict placement rules to maintain consistent

performance and resilience; if vSAN determines that object layout compliance cannot be restored without violating these rules, it enters Reduced availability with no rebuild state.

This behavior is expected and documented: rebuilds cannot occur if the cluster does not have sufficient hosts or free capacity to recreate absent components. The administrator's ESA configuration behaved correctly given the cluster size limitation, making B the correct answer.

Question: 15

An administrator attempts to update the VMware vCenter root account password through VMware Cloud Foundation (VCF) Operations. The attempt fails with the following error message, "Failed to authenticate with the guest operating system using the supplied credentials." What is the cause of the failure?

A. The password does not meet policy requirements.

B. The password was previously updated on the vCenter directly.

C. vCenter is down.

D. The SSH service is not running.

Answer: B

Explanation:

VMware Cloud Foundation 9.0 Operations manages credentials for integrated components such as vCenter Server through its internal password vault. When administrators modify passwords directly on the component—such as manually changing the vCenter root password—VCF Operations is no longer able to authenticate using its stored credentials. As a result, any password rotation or update operation initiated through VCF Operations fails during the validation step.

The error "Failed to authenticate with the guest operating system using the supplied credentials" is a direct symptom of this condition. VCF Operations attempts to log in to vCenter using the previously stored credential, which no longer matches the actual root password. Documentation describes this as an "out-of-sync credential state," and the resolution is to perform password remediation to resynchronize VCF Operations with the system.

Option A (password complexity) is irrelevant because complexity is validated only after authentication.
Option C (vCenter down) would generate connectivity errors, not authentication errors.
Option D (SSH disabled) does not prevent password rotation because VCF Operations uses VMware Tools guest operations, not SSH, for authentication.

Question: 16

An administrator has been tasked with deploying a new workload domain consisting of six VMware ESX hosts with VMware vSAN into an existing VMware

Cloud Foundation (VCF) instance. After starting the deployment from VCF Operations, they discover that only four of the six hosts required are listed for

selection in the UI. The administrator checks the Unassigned Host Inventory view in the vSphere

Client and confirms that all six hosts are listed.

Which step should the administrator perform to identify why the two hosts are not available for selection?

A. Check that the management port group on the standard switch has been enabled for vSAN traffic.

B. Check that the failures to tolerate (FTT) setting for the workload domain is set to 0.

C. Check that all disk partitions have been deleted from the SSD drives of the hosts.

D. Check that the network pool the hosts have been associated with is enabled for vSAN.

Answer: D

Explanation:

When deploying a new workload domain in VMware Cloud Foundation (VCF), only ESXi hosts that fully meet all pre-requisites are displayed in the VCF Operations UI for selection. Although all six hosts appear in the Unassigned Host Inventory in vCenter, VCF performs additional validation before making them selectable for workload domain deployment.

One of the mandatory requirements for any vSAN-enabled workload domain is that the ESXi hosts must be associated with a Network Pool configured for vSAN traffic. A network pool defines the host network configuration (VLANs, MTU, NIC mapping) used during domain deployment.

If the two missing hosts are associated with a network pool that does not have vSAN traffic enabled, or are associated with no network pool at all, VCF will exclude them from the workload domain deployment wizard. This is documented behavior: VCF filters out hosts when required network intents—such as vSAN—are not present.

Other options are incorrect:

A . Management port group enabled for vSAN traffic — vSAN should never run on the management PG.

B . FTT setting — Has no effect on host visibility; applies only after deployment.

C . Disk partitions — Affects vSAN disk claim but does not prevent host selection in VCF.

Question: 17

A VMware NSX Edge node is present in the inventory but shows "Not Ready" status In NSX Manager UI. What should the administrator check first?

A. The NSX Edge has been added to an Edge cluster

B. The license key in NSX Manager UI

C. The NSX Edge node's uplink network configuration

D. The NSX Edge node's CPU reservation

Answer: C

Explanation:

The status "Node Not Ready" in the NSX Manager UI (specifically in the Configuration State column of the Edge Transport Nodes view) indicates that the NSX Manager has failed to push or validate the necessary configuration to the Edge VM.

Check Uplink Network Configuration (Option C): This is the most common cause for a "Node Not Ready" state during deployment or operation. For an Edge Node to be "Ready" (Success/Up), it must have a valid Transport Node configuration, which includes the Uplink Profile, IP Pool (for TEPs), and mapping to the Fastpath Interfaces (N-VDS). If the uplink configuration is missing, incorrect, or the management plane cannot communicate with the edge to apply it, the node remains in a "Not Ready" state.

Why not Option A? While an Edge must be in an Edge Cluster to be utilized by a Tier-0 Gateway, a standalone Edge Node should still report a status of "Success" (Configuration) and "Up" (Node Status) if it is healthy. Adding a "Not Ready" (unhealthy/unconfigured) node to a cluster will not fix the underlying configuration issue.

Why not Option D? Missing CPU reservations typically lead to a "Degraded" status or service crashes (Dataplane down), but "Node Not Ready" is the specific indicator of an incomplete or stalled configuration workflow, usually tied to the transport/uplink setup.

Question: 18

An administrator is asked to create a second provider gateway (provider gateway 02) in VMware Cloud Foundation (VCF) Automation Region-A.

After launching the Create Provider Gateway workflow in the VCF Automation Provider Management Portal, no Tier-0 Gateway is available for assignment.

How would you resolve this issue?

A. Create a new Region.

B. Log into the NSX Manager, create a new Tier-1 Gateway.

C. Log into the NSX Manager, create a new TO Gateway.

D. Retry the Create Provider Gateway workflow.

Answer: C

Explanation:

In VMware Cloud Foundation 9.0, a Provider Gateway in VCF Automation is always backed by an existing Tier-0 or Tier-0 VRF gateway in NSX. When the administrator launches the Create Provider Gateway workflow and no Tier-0 gateways appear for assignment, this indicates that VCF Automation cannot discover any valid Tier-0 gateways in the associated region.

The VMware Cloud Foundation 9.0 documentation explicitly states that before adding a Provider Gateway, an

administrator must first create an Active-Standby Tier-0 Gateway in NSX Manager. The Provider Gateway workflow only lists Tier-0 gateways that already exist and are properly configured in NSX. If none are present, the list will be empty.

From the documentation: “To add a provider gateway, first you must create an Active Standby tier-0 gateway in the NSX Manager associated with the region to back it.” . Provider gateways in VCF Automation are discovered from these preexisting Tier-0 gateways and cannot be created until they exist.

Creating a Tier-1 gateway (Option B) does not satisfy the requirement because Provider Gateways must map specifically to Tier-0, not Tier-1. Retrying the workflow (Option D) will not resolve the issue because the Tier-0 backing resource is missing. Creating a new region (Option A) is unnecessary unless required for other organizational reasons, and it still would not produce a Tier-0 gateway.

Therefore, the correct and verified solution is to log in to NSX Manager and create the required Tier-0 gateway, after which it will appear in the Provider Gateway creation workflow.

Question: 19

An administrator is troubleshooting an issue relating to VMware Cloud Foundation (VCF) Automation. While troubleshooting, the administrator realizes that debug-level information is not displayed in the VCF Automation Task Log.

How would the Administrator enable debug-level information in the Task Log?

A. Enable "display debug information" in the Administer > Settings section of the Organization Management portal.

B. Enable "display debug information" in the Administration > Feature Flag section of the Provider Management portal.

C. Enable "display debug information" in the Administration > Events and Tasks section of the Provider Management portal.

D. Enable "display debug information" in the Administration > General Settings section of the Provider Management portal.

Answer: B

Explanation:

In VMware Cloud Foundation (VCF) 9.0 Automation, the visibility of debug-level information in Task Logs is controlled centrally by the Provider Administrator through the Provider Management portal. Debug logging is not enabled by default because it exposes verbose operational details intended primarily for troubleshooting. According to the VCF Automation architecture and operations model, advanced logging capabilities—including debug output—are gated behind feature flags.

To enable debug-level information, the Provider Admin must navigate to:

Provider Management → Administration → Feature Flags → Display Debug Information

Once this flag is enabled, the system begins emitting additional diagnostic detail into Task Logs, improving insight into failures, orchestration flows, API calls, and service-to-service interactions. This aligns with VCF's multi-tenant design, where only the Provider tier has permission to modify global
settings that affect all Organizations.

Options A, C, and D are incorrect because Organization-level settings do not control system-wide logging, and the Events/Tasks or General Settings sections do not contain the mechanism for enabling debug output. Only the Feature Flag section controls this capability.

Question: 20

An administrator is preparing to import a vSphere environment into VMware Cloud Foundation (VCF) as a workload domain. The vSphere environment has the following configuration:

- vSphere version 8.0 update 3.
- Three-node vSAN cluster with a single OSA datastore.
- Two vSphere Distributed Switches (VDS).
- Three vmkernel adapters with DHCP assigned IP addresses.

What change must the administrator make before importing this environment?

A. Consolidate to a single vSphere Distributed Switch.

B. Upgrade vCenter and ESXi to vSphere 9.0.

C. Update the vmkernel adapters with statically assigned IPs.

D. Convert the vSAN datastore from OSA to ESA.

Answer: C

Explanation:

When importing an existing vSphere environment into VMware Cloud Foundation (VCF) as a workload domain, several strict prerequisites must be met. One of the key requirements documented in VCF 9.0 is that all VMkernel adapters (vmk ports) used for vSAN, vMotion, management, or other system traffic must have statically assigned IP addresses. DHCP-assigned VMkernel IPs are not supported for VCF workload domain bring-up or import operations.

In the provided scenario, the environment includes:

vSphere 8.0 U3

A 3-node vSAN OSA cluster

Two VDS switches

VMkernel adapters using DHCP

Before VCF can successfully validate and import the environment, the administrator must convert these VMkernel interfaces to static IP addressing. VCF uses IPAM assumptions and deterministic host networking configurations; DHCP introduces variability incompatible with automated lifecycle operations.

Option A (consolidating VDS) is unnecessary—VCF supports multiple VDS configurations during import.

Option B (upgrading to vSphere 9.0) is not required for import.

Option D (convert OSA to ESA) is impossible pre-import and not required—VCF supports OSA clusters.

Question: 21

An administrator has successfully created a new Organization for All Apps In VMware Cloud Foundation (VCF) Automation. When logging into the new organization using the first user account, only the Overview tab is visible.

What is a possible cause of this issue?

A. The first user account was assigned the Organization Auditor Role.

B. The first user account was assigned the Organization User Role.

C. The first user account was assigned a Custom Role.

D. The first user account was assigned the Organization Administrator Role.

Answer: B

Explanation:

This issue stems from an incorrect role assignment during the user creation process in VMware Cloud Director (VCF Automation).

Organization Administrator Role (Option D): This role grants full control, including visibility of the Administration tab (to manage users, groups, and settings), Data Centers, and Monitor tabs. If the user were an Admin, they would see all tabs.

Organization Auditor Role (Option A): This is a read-only role, but by definition, an Auditor can view anything an Organization Administrator can see (including the Administration settings), just without edit rights. Therefore, an Auditor would still see the Administration tab.

Organization User Role (Option B): This is a consumer-level role designed for deploying and managing vApps. By default,

this role does not have access to the Administration tab or high-level organization settings. If the organization is new and has no vApps or VDCs populated yet, a user with this role might see a very restricted view (effectively just a dashboard or "Overview") because they lack the rights to see the administrative configuration menus.

Conclusion: The fact that the "Administration" tab is missing (implied by "only Overview is visible") identifies the user as an Organization User (or a restricted Custom Role) rather than an Administrator or Auditor.

Question: 22

An administrator discovers that a VMware Cloud Foundation (VCF) workload domain four-node vSAN cluster is experiencing a network partition. The workload domain vCenter displays a "vSAN duster partition" warning. The performance across the cluster is degraded and the objects are showing as non-compliant.

What could be causing the network partition?

A. IGMP snooping is disabled on the multicast group.

B. The VLAN was changed on the physical switch port.

C. Jumbo frames are configured on the vSphere distributed switch (VDS).

D. The vSAN Witness service was added to the vMotion network.

Answer: B

Explanation:

A vSAN cluster network partition occurs when vSAN nodes cannot communicate over the designated vSAN network. In VMware Cloud Foundation workload domains, the vSAN network relies on L2 adjacency, consistent VLAN configuration, and stable multicast/BUM behavior (in older versions). VCF 9.0 uses unicast-mode vSAN, so multicast-related issues (such as IGMP snooping configuration) are no longer relevant.

A network partition can occur when the VLAN ID on the physical switch port differs from the VLAN configured on the vSphere Distributed Switch (VDS) for the vSAN VMkernel adapters. The documentation emphasizes that consistent VLAN configuration across the physical and virtual network is required for proper vSAN cluster communication. If a switch port is reconfigured—

intentionally or accidentally—to use a different VLAN, the node becomes isolated from the rest of the vSAN cluster, causing:

"vSAN cluster partition" warnings in vCenter

degraded performance

objects marked as non-compliant

resyncs that cannot complete

Option A (IGMP snooping) does not apply because modern vSAN uses unicast, not multicast.

Option C (Jumbo frames) would cause packet loss only if inconsistently configured, but it does not cause a full network partition.

Option D (vSAN Witness on vMotion) is relevant only for stretched clusters and does not cause a partition in a standard four-node cluster.

Question: 23

After upgrading from VMware Cloud Foundation (VCF) 5.2 to VMware Cloud Foundation (VCF) 9.0 the administrator attempts to enable SSH access through the vCenter console to the newly upgraded VCF Ops instance and Is not able to. They attempt to log in through SSH as the root user and they are unable to. What needs to be done to enable SSH access to the VCF Ops instance?

A. Reset the root password.

B. Reboot the appliance and enable SSH.

C. Rollback to snapshot because the upgrade did not work as expected.

D. Use VCF Operations to remediate the password

Answer: D

Explanation:

In VMware Cloud Foundation (VCF) 9.0, the management of appliance credentials and lifecycle operations is centralized within the VCF Operations Fleet Manager (which subsumes the roles of the legacy SDDC Manager Life Cycle Management).

The Problem: The administrator is unable to log in as root via the console or SSH. This indicates a credential synchronization issue or account lockout, which prevents them from manually enabling SSH via the console (the traditional method).

The Solution (Remediate Password): The "Remediate Password" workflow in VCF Operations allows the administrator to reset and synchronize the root password for VCF components (like the VCF Ops

instance itself) directly from the management plane.

By navigating to Fleet Management > Passwords (or similar path in VCF 9.0), the administrator can select the affected instance and choose Remediate.

This process updates the password in the centralized database and on the appliance, restoring the ability to log in.

Once the root access is restored via remediation, the administrator can then proceed to enable SSH (either via the VCF

Operations settings UI or the console). Without the correct password (which "Remediate" fixes), SSH cannot be enabled.

Note: Options A and B (Reset/Reboot) are legacy manual steps that do not ensure the VCF inventory database is updated, potentially leading to further "configuration drift" or sync errors. Option C is unnecessary for a credential issue.

Question: 24

An administrator is attempting to log into the vCenter using the vSphere Client but receives an error stating "no healthy upstream" What are two possible causes for this? (Choose two.)

A. The vpxd service is not running.

B. The SSO Service is not running.

C. Port 443 is not opened between the local machine and the vCenter.

D. The administrator logged in with the root account.

E. The vmware-rbd-watchdog service is not running.

Answer: A, B

Explanation:

The vSphere Client “no healthy upstream” error is a classic indicator that one or more vCenter backend services are not running or responding, preventing the reverse proxy layer (envoy / nginx) from routing requests to the appropriate upstream services.

Two services in particular are known root causes:

A . vpxd service not running

vpxd is the core vCenter Server service responsible for inventory, host management, and client interaction. If vpxd is stopped, crashed, or restarting, the vSphere Client cannot communicate with backend APIs, resulting in the “no healthy upstream” condition.

B . SSO (vmware-stsd / identity service) not running

Authentication in vCenter depends on the SSO/Identity service. If SSO is unavailable, login sessions cannot be validated, and vCenter marks the upstream service as unhealthy.

Other options do not match the behavior:

C (Port 443 closed) would produce a connection failure, not the upstream error.

D (logging in with root) is fully supported and does not trigger this message.

E (vmware-rbd-watchdog) relates to backup/restore health, not core authentication/management planes.

Question: 25

An administrator attempts to configure a Microsoft Certificate Authority in VMware Cloud Foundation (VCF) Operations supplying a certificate template name of VMware. The attempt fails with error, "Certificate authorities update failed."

What is the possible cause of this failure?

A. The user account has only the "Enroll" permission on the certificate template.

B. The user account does not have the "Enroll" permission on the certificate template.

C. The user account does not have the "Read" and "Autoenroll" permission on the certificate template.

D. The user account has only the "Read" and "Enroll" permission on the certificate template.

Answer: A

Explanation:

To successfully configure a Microsoft Certificate Authority (CA) in VMware Cloud Foundation (VCF) Operations (formerly vRealize/Aria Operations), the service account used for the integration must have specific permissions on the Certificate Template (e.g., the "VMware" template).

Required Permissions: The VCF 9.0 and Aria Operations documentation explicitly states that the service account must be assigned Read and Enroll permissions on the target Certificate Template.

Read: This permission is critical for the "Discovery" and "Validation" phase. It allows VCF Operations to query the CA, list available templates, and read the template's properties (like Key Usage and Extended Key Usage) to ensure they meet the security requirements (e.g., Server Authentication, Non-Repudiation).

Enroll: This permission allows the account to actually submit a Certificate Signing Request (CSR) via

the interface and receive a signed certificate.

The Cause of Failure (Option A): If the user account is configured with only the "Enroll" permission, it effectively lacks the "Read" permission. Without "Read", VCF Operations cannot "see" or validate the template during the configuration wizard. The application attempts to fetch the template details, fails (because the template is invisible to it), and throws the error "Certificate authorities update failed."

Why other options are incorrect:

Option D (Read and Enroll): This is the correct and recommended configuration. If the user had these permissions, the operation would succeed (assuming other prereqs like Basic Auth are met).

Option C (Autoenroll): The Autoenroll permission is designed for Windows Group Policy-based background renewal. It is not required for the VCF Operations API-based integration, which relies on explicit "Enroll" calls.

Question: 26

An administrator is creating a new workload domain from VMware Cloud Foundation (VCF) Operations. They are blocked at the Hosts selection screen as no ESX hosts are available. They see the following message:

"No suitable hosts available to create a VI workload domain. Hosts must be unassigned, commissioned with at least one physical NIC and the same storage type as the VI workload domain, and the ESX version must be compatible with the lowest ESX version present in the management domain."

How can the administrator commission new hosts to enable the creation of the VI workload domain?

A. Using the Cloud Builder.

B. Using the vSphere client.

C. Using the VCF Installer.

D. Using VCF Operations.

Answer: D

Explanation:

In VMware Cloud Foundation 9.0, all host commissioning operations are performed through VCF Operations, not through vSphere Client, Cloud Builder, or the VCF Installer. Once VCF is deployed, Cloud Builder is no longer used, and the VCF Installer is for lifecycle and bundle management—not

for host workflows. The vSphere Client also cannot commission hosts because host commissioning is a foundational VCF workflow requiring hardware validation, storage type checks, NIC checks, HCL conformance, and version compatibility.

The error message provided:

“Hosts must be unassigned, commissioned with at least one physical NIC and the same storage type… and the ESX version must be compatible…”

is a standard VCF 9.0 validation message shown when no commissioned hosts matching the workload domain requirements exist. VMware documentation states that hosts must be commissioned under:

VCF Operations → Fleet Management → Hosts → Commission Host

Here, VCF validates:

Storage type (vSAN ESA, vSAN OSA, NFS, FC, etc.)

Network pool membership (matching the WLD plan)

ESXi version compatibility with the Management Domain baseline

NIC mapping and certifications

Until hosts are commissioned, they cannot appear in the workload domain creation wizard.

Thus, the correct method to commission hosts is D. Using VCF Operations.

Question: 27

An administrator is tasked to add a new host to a vSphere cluster that was created with VMware vSAN Express Storage Architecture (ESA) as its principal storage in an existing workload domain.

The administrator successfully commissions the new host with a VMware vMotion only network pool but is unable to add the host to the existing cluster.

What must the administrator do to be able to complete this task?

A. Decommission, reinstall ESX, and recommission the new host to the network pool for the existing vSAN ESA cluster.

B. Change the network pool associated to the new host to the network pool for the existing vSAN ESA cluster.

C. Manually configure the vSAN network on the new host within vCenter.

D. Reconfigure the currently associated network pool with a vSAN network.

Answer: B

Explanation:

In VCF 9.0, when adding a host to a vSAN ESA-enabled cluster, the host must be commissioned with a network pool that includes a vSAN network configuration. Network pools define host-level networking templates for VCF, including management, vSAN, vMotion, and overlay networks. A host commissioned with a vMotion-only network pool does not have the required vSAN ESA network interfaces (vmk + NIC mapping) to join an ESA cluster.

Because the administrator successfully commissioned the new host but only using a vMotion-only network pool, VCF correctly prevents the host from being added to the ESA cluster.

The required action is:

Reassociate the host with the correct network pool that includes the vSAN ESA network.

Option A (reinstall ESXi) is unnecessary; commissioning workflows can be redone.
Option C (manual vCenter configuration) is explicitly unsupported—VCF manages host networking. Option D (reconfiguring the existing pool) is not correct because the new host must be associated with the same network pool used by the existing ESA cluster, not change the pool definition itself.

Therefore, the precise and VMware-documented resolution is B.

Question: 28

An administrator is responsible for supporting a VMware Cloud Foundation (VCF) fleet and has been tasked with

deploying VMware Cloud Foundation (VCF) Operations for Logs. To complete this task, the administrator needs to configure a new ofline depot within VCF Operations fleet management.

The following information has been provided to the administrator to complete the task:

- Ofline Depot Type: Webserver
- Repository URL: http://10.138.148.160/depot/
- Username: depotuser
- Password: P@sswordl23!
- Accept imported certificate: True

When the administrator attempts to configure the depot, the following error message is presented:

Either the depot URL provided is partial or invalid or not reachable or download token is invalid. Check logs for more details.

The administrator completes the following troubleshooting steps:

- Confirms the Repository URL is valid by connecting to it through a web browser.
- Reviews the command used to create the depot:
  - `./vcf-download-tool binaries download —depot-store=/VCF —depot-download-token-file=<token_file_path> -vcf-version=9.0.0.0 —sku=VCF —coaponent=VRLI —type= INSTALL`
- Confirms that the downloaded folder and files were copied into the /depot shared folder on the web server hosting the repository

Which two actions must the administrator take to resolve the issue? (Choose two.)

A. Reconfigure the web server to share the /vcf/ folder containing the depot files.

B. When configuring the ofline depot, the Repository URL should be changed to http://10.138.148.160.

C. When configuring the ofline depot, the OflineDepotType should be changed to Local Path.

D. Reconfigure the Fleet Manager appliance to share the /data/ folder.

E. When configuring the ofline depot, the Repository URL should be changed to https://10.138.148.160/depot/.

Answer: A, E

Explanation:

To resolve the "partial or invalid or not reachable" error when configuring the VCF 9.0 Ofline Depot, the administrator must address two critical misconfigurations related to the protocol and the file path mapping:

Switch to HTTPS (Option E): VMware Cloud Foundation 9.0 enforces HTTPS by default for all depot connections to ensure security. The administrator's configuration uses http://, which the VCF Fleet Manager will reject (or fail to connect to)

unless the system has been explicitly modified via internal properties files to allow insecure transport. Changing the Repository URL to https://10.138.148.160/depot/ aligns with the default security requirements of the VCF 9.0 binaries download and validation process.

Reconfigure Web Server Pathing (Option A): The command --depot-store=/VCF instructs the download tool to create a repository structure rooted at /VCF. The administrator then copied this "downloaded folder" into the /depot folder on the web server, resulting in a nested path (e.g., /var/www/html/depot/VCF/...). However, the configured URL is .../depot/, which points to the parent directory where the required index.json or metadata files are not immediately visible. The administrator must reconfigure the web server (e.g., via DocumentRoot or Alias settings) to explicitly share the specific /vcf/ (or /VCF/) folder content at the target URL so the Fleet Manager can locate the manifest files.

Question: 29

An administrator is troubleshooting a vSAN issue. As part of the initial investigation, the following observations were identified:

- vSAN cluster capacity is decreased.
- Some virtual machine components are marked as degraded.
- Component rebuild process started automatically.

What is the cause of this issue?

A. VM migration to another cluster is in progress.

B. vSAN license capacity is too small.

C. Too many virtual machines were created in the vSAN cluster.

D. Physical disk failure.

Answer: D

Explanation:

The symptoms described—reduced cluster capacity, degraded virtual machine components, and automatic component rebuild operations—are classic indicators of a vSAN disk failure or disk group degradation. vSAN continuously monitors the health of disks, disk groups, and network paths. When a physical disk or disk group becomes unavailable, vSAN will:

Mark affected components as degraded because the required number of replicas or witnesses cannot be maintained.

Trigger automatic repair/rebuild operations, provided there are enough healthy disks remaining in the cluster to satisfy the storage policy (e.g., FTT=1, RAID1/5/6).

Reduce available storage capacity because the failed device is removed from contributing to the vSAN datastore.

These behaviors align directly with documented vSAN failure-response logic, which states that component rebuilds begin

automatically after a disk failure, assuming the cluster still has adequate resources.

The other options do not match the symptoms:

A . VM migration to another cluster → does not reduce vSAN capacity nor trigger component rebuilds.

B . vSAN license capacity too small → restricts features, not component state or capacity changes.

C . Too many VMs created → may cause capacity pressure but does not mark components degraded or trigger automated rebuilds.

Only physical disk failure accurately explains all three observations simultaneously.

Question: 30

An administrator created a new VPC with an associated subnet, configured with a DHCP Server.

When attaching virtual machines to the VPC subnet, an IP address is assigned, but the DNS and NTP settings are not configured.

How can the administrator update the DHCP server configuration to set DNS and NTP?

A. Update the default VPC Service Profile to include the IP addresses for the DNS and NTP servers.

B. Change the DHCP Server mode from DHCP Server to DHCP Relay.

C. Enable DNS and NTP Passthrough on the DHCP Server.

D. Switch the DHCP Network mode from Distributed Connectivity to Centralized Connectivity.

Answer: A

Explanation:

In VMware Cloud Foundation 9.0 Automation, each VPC is governed by a VPC Service Profile, which defines the default network services applied to the VPC's DHCP server—this includes DNS servers, NTP servers, DHCP lease values, and other network attributes. When a subnet is associated with a VPC and DHCP is enabled, the DHCP service inherits its DNS and NTP configuration from the VPC Service Profile.

In the scenario, virtual machines attached to the new VPC subnet receive an IP address, but not DNS or NTP settings. This indicates that the DHCP server is functioning correctly, but its service profile lacks DNS and NTP configuration. Updating the default VPC Service Profile allows the administrator to specify DNS resolver addresses and NTP time sources, which will then automatically be pushed to all DHCP-enabled subnets under that VPC.

Option B (changing to DHCP Relay) is incorrect because relay mode does not configure DNS/NTP—it delegates DHCP to an external DHCP server.

Option C (enable DNS/NTP passthrough) is not a feature of NSX DHCP.

Option D (changing connectivity mode) affects routing and service placement, not DHCP options.

Question: 31

An administrator is troubleshooting a problem with NSX.

Which command can be used to validate installed NSX VIBs on the ESX host?

A. esxtop -b -d 2 -n 100

B. esxcli software vib list

C. nsxcli get version

D. esxcfg software list

Answer: B

Explanation:

When troubleshooting NSX on an ESXi host, VMware requires verification that NSX VIBs (vSphere Installation Bundles) are installed and in the correct state. VIBs are responsible for NSX datapath, control-plane modules, and kernel extensions on ESXi. The authoritative and documented method to list VIBs on an ESXi host is the command:

esxcli software vib list

This command displays all installed kernel modules, version numbers, NSX packages, and their installation status. For NSX-T (now part of VCF networking), administrators expect to see VIBs such as nsx-aggservice, nsx-bridge, nsx-esx-datapath, and others. If any required NSX VIBs are missing or inconsistent, the ESXi host will fail to join NSX transport nodes or will show "Not Ready."

Option A (esxtop) is for performance monitoring and does not show VIB information.

Option C (nsxcli get version) checks NSX version on Edge Nodes or host transport nodes but does not list VIBs.

Option D (esxcfg software list) is an outdated and invalid command.

Question: 32

An administrator is preparing to upgrade their VMware Cloud Foundation (VCF) management domain from VCF 5.0 to VCF 9.0.

After configuring the online depot, they see the SDDC Manager 9.0 upgrade bundle is available.

However, the 9.0 upgrade bundles for vCenter, ESX, and NSX are missing.

How can the administrator resolve this issue?

A. Upgrade the SDDC Manager to 9.0.

B. Use the VCF Download Tool to download the missing 9.0 upgrade bundles.

C. Upgrade the management domain from VCF 5.0 to VCF 5.2.

D. Use the VCF Ofline Bundle Transfer Utility (OBTU) to download the missing 9.0 upgrade bundles.

Answer: A

Explanation:

When upgrading from VCF 5.0 to VCF 9.0, the upgrade workflow requires that the SDDC Manager be upgraded first before any other component bundles (vCenter, ESX, NSX) become visible. This is explicitly stated in the VMware Cloud Foundation upgrade process: the upgrade bundles for the management domain components are dependent on the SDDC Manager version. The online depot will not present the 9.0 upgrade bundles for vCenter, ESX, or NSX until the SDDC Manager itself has reached the target major version (in this case, 9.0).

This is because SDDC Manager contains the updated Lifecycle Management (LCM) engine and updated bundle manifests, which are required to understand, download, and orchestrate the remaining component upgrades. Attempting to download the other bundles without upgrading SDDC Manager first is not supported.

Options B and D (download tools) are incorrect because the issue is not that the bundles are missing from the depot, but that SDDC Manager 5.x cannot interpret 9.0 component bundles. Option C (upgrade to 5.2 first) is also incorrect because the VCF 5.x → 9.x upgrade path is directly managed by the upgrade planner once SDDC Manager is upgraded.

Thus, the correct resolution is to upgrade the SDDC Manager to 9.0, after which the remaining component bundles will become available.

Question: 33

An administrator has observed that the vSphere Global Inventory is only available from the management domain vCenter. The Global Inventory is not available from the workload domain's vCenter.

Why is the "Global Inventory" missing from the workload domain's vCenter?

A. VCF SSO and vCenter Linking have not been configured.

B. Supervisor Management has not been enabled.

C. An inventory sync was not run following the workload domain creation.

D. An external VIDB instance has not been configured.

Answer: A

Explanation:

The Global Inventory List (GIL) is only available when multi-vCenter SSO domain linking is configured. In VMware Cloud Foundation, the management domain vCenter is deployed first and becomes the root vCenter for global inventory data. For workload domains, their vCenter Servers must be

registered into the same SSO domain and linked with the management-domain vCenter in order for the global inventory data (VMs, hosts, clusters, content libraries) to appear.

If a workload domain vCenter is not SSO-linked, it operates in its own identity domain, and therefore cannot access or present Global Inventory, resulting in exactly the symptom described: the management domain vCenter shows the GIL,

while the workload domain vCenter does not.

Option B (Supervisor Management) relates to vSphere with Tanzu and has no impact on Global Inventory.

Option C (inventory sync) is incorrect—there is no manual sync required; GIL relies entirely on SSO linking.

Option D (VIDB) is not related to vCenter linking or inventory visibility; it is used by VCF Identity Broker.

Therefore, the reason the Global Inventory is missing from the workload domain vCenter is that SSO/vCenter Linking has not been configured, which is required for federation across all VCF vCenters.

Question: 34

An administrator is responsible for managing a remote VMware Cloud Foundation (VCF) fleet with the following configuration:

- A single VCF instance with a single Workload Domain.
- The Workload Domain has a single VMware vSAN Express Storage Architecture (ESA) cluster.
- VCF is licensed using the disconnected mode.

The administrator discovers a notification in VCF Operations showing that the VCF licenses have expired. Which three steps should the administrator take to resolve the issue? (Choose three.)

A. Increase the license core count in SDDC Manager.

B. Restart SDDC Lifecycle Manager Service in the VCF Operations console.

C. Export the usage file from VCF Operations and upload to the VCF Business Services console.

D. Use the VCF Business Services console to export a new VCF license file.

E. Import the license file into VCF Operations and assign to the workload domain vCenter.

F. Import the license file into VCF Operations and assign to the SDDC Manager.

Answer: C, D, F

Explanation:

In VMware Cloud Foundation (VCF) 9.0 using disconnected mode licensing, VCF Operations does not automatically synchronize license status with VMware's cloud services. Instead, the administrator must periodically refresh the license file using a manual ofline workflow. When the VCF Operations console reports that licenses have expired, it means the license entitlement in the VCF Business Services portal is out of date, and therefore VCF Operations cannot validate the current usage.

The VMware-documented ofline licensing workflow requires the following steps:

Export the usage file from VCF Operations.

This usage file contains consumption details needed to generate a new ofline license.

→ C is correct.

Upload the usage file to the VCF Business Services console and generate a new ofline license file. In disconnected mode, the Business Services portal is the only mechanism to create updated license entitlements.
→ D is correct.

Import the updated VCF license file into VCF Operations, specifically assigning it to the SDDC Manager.
SDDC Manager is the system that validates and enforces licensing across workload domains, so the new license must be applied there—not only to a vCenter.
→ F is correct.

Options A and B do not affect license validation.
Option E is incorrect because workload-domain vCenter licensing is independent and not the root cause of VCF license expiration.

Question: 35

An administrator is planning to apply updates to a VMware vCenter instance.

What two actions can the administrator take to confirm the status of the vCenter services? (Choose two.)

A. Connect to the vSphere Client and review vCenter performance charts.

B. Connect to the vCenter appliance shell and run the services-control -status command.

C. Connect to the vCenter Server Management console and review the services statuses.

D. Connect to the vCenter appliance shell and run the vim-top command.

E. Connect to the ESX DCUI where the vCenter Appliance is running and run the services.sh script.

Answer: B, C

Explanation:

Before applying updates to a vCenter Server Appliance (VCSA), an administrator must validate that all vCenter services are healthy. VMware provides two supported and documented methods for checking vCenter service status:

1. Using the vCenter Appliance Shell

Running the command:

services-control --status

This command displays the status of all vCenter-related services (vmdird, vmcad, vpxd, vsan-health, etc.). It is the authoritative diagnostic tool embedded in the appliance for confirming whether services are running, stopped, or in a degraded state. This method is explicitly documented in vSphere 9.0 service management procedures.

This matches Option B.

2. Using the vCenter Server Management Interface (VAMI)

Accessed at:

https://<vcenter-fqdn>:5480

The VAMI console provides a graphical interface under Services, showing the real-time health, status, and start/stop controls for all vCenter services. VMware documentation instructs administrators to review service status here before performing upgrades or maintenance operations.

This matches Option C.

Incorrect Options Explained

A . vSphere performance charts → These show workload data, not service health.

D . vim-top command → Displays vSphere hosts' runtime metrics, not vCenter services.

E . Running services.sh on ESXi DCUI → vCenter does not run ESXi services; this script is for ESXi hosts only.

Question: 36

An administrator is tasked with replacing a VMware vCenter certificate in VMware Cloud Foundation (VCF) Operations with an external CA-signed certificate. The certificate import completes successfully but when running the certificate replacement task, it fails with the following error: Certificate replacement has failed...The Certificate Chain validation failed due to 'Signature does not match' What is the possible cause of this issue?

A. The Certificate Signing Request (CSR) included the IP address of the vCenter.

B. The external CA is not trusted by VCF Operations.

C. The external CA is not accessible to VCF Operations.

D. The server certificate was copied to the wrong field.

Answer: D

Explanation:

When replacing certificates in VMware Cloud Foundation (VCF) Operations, the system performs strict certificate chain validation. The error shown:

"Certificate chain validation failed due to 'Signature does not match'"

indicates that VCF Operations attempted to validate the presented certificate chain but detected that the server certificate did not correctly match the signing CA certificate. This occurs most commonly when the administrator pastes

the server certificate and CA root/intermediate certificates into the wrong fields during import.

VCF requires the certificate bundle to be uploaded in the correct format:

Server certificate → Server Certificate field

Intermediate certificates → Intermediate Chain field

Root certificate → Root CA field

If the chain order is wrong or the server certificate is mistakenly placed in an intermediate or root CA field, the cryptographic signature validation fails. This exact failure mode is documented in VMware certificate replacement workflows.

Option A is incorrect because including an IP address in a CSR does not invalidate chain signatures. Option B is incorrect because an untrusted CA produces a trust failure, not a signature mismatch. Option C is unrelated: accessibility is not required for certificate validation.

Question: 37

An administrator recently deployed a new three-node VMware vSAN Express Storage Architecture (ESA) cluster to an existing workload domain. After creating a number of Virtual Machines (VMs), the administrator discovers that storage is being consumed a lot quicker than expected.

While investigating the issue, the administrator discovers that the datastore default policy has been set to RAID-1 by Auto-Policy Management rather than the expected RAID-5.

What is a possible cause?

A. The RAID-5 policy is only supported on a vSAN ESA storage cluster.

B. The vSAN ESA cluster must have a minimum of four hosts.

C. The vSAN storage policy has Force Provisioning enabled.

D. The vSAN ESA cluster has Host Rebuild Reserved enabled.

Answer: B

Explanation:

In vSAN Express Storage Architecture (ESA), Auto-Policy Management determines which default storage policies can be used based on the number of hosts in the cluster. RAID-5 and RAID-6 policies require a minimum number of hosts to satisfy fault domain and component placement rules.

For vSAN ESA, the minimum hosts required are:

RAID-1 (FTT=1) → minimum 3 hosts

RAID-5 (FTT=1) → minimum 4 hosts

RAID-6 (FTT=2) → minimum 6 hosts

In this scenario, the administrator deployed a three-host ESA cluster. Since RAID-5 requires at least four ESA-capable hosts, vSAN Auto-Policy Management automatically falls back to RAID-1, the highest level of resilience possible with the available cluster size. This results in significantly higher storage consumption, which matches exactly what the administrator observed.

Option A is incorrect because RAID-5 is fully supported on ESA—but only with enough hosts.
Option C (Force Provisioning) does not change the default policy selected.
Option D (Host Rebuild Reserve) does not control RAID policy selection.

Question: 38

An administrator logs into the vSphere client to check the health of a cluster. An alert appears on the cluster stating, "vSphere HA host status".

The administrator toggles vSphere HA off and on and the following error appears on the host "A general system error occurred: Failed to start fdm service on host".

What is the cause of this issue?

A. The vmware-fdm service is disabled on the ESX host.

B. The vmware-fdm vib is missing from the ESX host.

C. vSphere HA Admission Control settings are not configured correctly.

D. vSphere HA startup policy is not configured correctly.

Answer: B

Explanation:

vSphere High Availability (HA) depends on the FDM agent (Fault Domain Manager) that runs on every ESXi host in the cluster. When an administrator enables HA on a cluster, vCenter automatically installs or updates the vmware-fdm VIB on each participating ESXi host. This VIB contains the HA agent binaries and is mandatory for HA services to start.
The error encountered:

"A general system error occurred: Failed to start fdm service on host"

is a classic and well-documented symptom of a missing or corrupted vmware-fdm VIB. When vSphere HA is toggled off and on, vCenter attempts to reinstall or restart the FDM agent; if the VIB is not present, HA cannot deploy successfully, and the FDM service fails to start.

Why the other answers are incorrect:

A. The vmware-fdm service is disabled

ESXi does not allow manual disabling of this system service in normal operations. If the service fails to start, the root

cause is usually the absence or corruption of the VIB—not a disabled service.

C. Admission Control settings not configured correctly

Admission Control errors affect VM failover capacity, not the ability to start FDM services.

D. HA startup policy not configured correctly

There is no per-host HA startup policy that prevents FDM from starting.

Question: 39

An administrator creates a tag for a virtual machine (VM) in VMware Cloud Foundation (VCF) Operations. When assigning the tag to the virtual machine In vCenter, the tag was not found.

What is the cause of this error?

A. The tag was not pushed to Custom Groups.

B. The vCenter version is incorrect.

C. The tag was not pushed to the vCenter instance.

D. VM Tools is not installed.

Answer: C

Explanation:

In VMware Cloud Foundation 9.0 Operations, tags created inside VCF Operations do not automatically appear in vCenter. Tags must be explicitly synchronized ("pushed") to the selected vCenter instance before they become usable for VM tagging within vCenter. This is because VCF Operations maintains its own metadata store for tags, super metrics, groups, and policies.

The correct workflow is:

Create the tag in VCF Operations.

Push (synchronize) the tag to the appropriate vCenter instance.

The tag then appears in vCenter's Tags & Custom Attributes section.

Administrators can then assign the tag to VMs.

If the push step is skipped, the tag exists only inside VCF Operations and cannot be referenced by vCenter, which is exactly the symptom described: tag not found when attempting to assign it to a VM.

Option A is incorrect because Custom Groups do not affect vCenter tag visibility.

Option B is incorrect because tag synchronization is not tied to a specific vCenter version as long as the vCenter is officially supported by VCF 9.x.

Option D is irrelevant—VMware Tools has nothing to do with tag visibility.

Question: 40

An administrator is responsible for managing a VMware Cloud Foundation (VCF) fleet. The following information has been provided about the VCF fleet configuration:

- The VCF fleet consists of a single VCF instance with a single management domain and a single workload domain.
- VCF Automation has a single Organization for VM Apps configured with a VCF Cloud Account for the workload domain.

The administrator has been tasked with creating a new Organization for All Apps to support the developers need to deploy Kubernetes-based applications in a new region in a workload domain.

The administrator attempts to create a new region through the VCF Automation Provider Portal but the VMware NSX manager for the workload domain does not appear on the list of available NSX managers.

What action must the administrator complete to resolve the issue?

A. Deploy an additional VCF workload domain cluster.

B. Trigger an inventory synch in VCF Operations fleet management.

C. Add the SDDC Manager integration for the VCF instance.

D. Deploy a new VCF workload domain.

Answer: C

Explanation:

In VMware Cloud Foundation 9.0 Automation, the Provider Portal must have full visibility into the underlying VCF inventory—including NSX Managers, clusters, regions, vCenters, and SDDC Manager objects—before new regions can be created for Kubernetes-based deployments (All Apps Orgs).

The issue described:

"The NSX Manager for the workload domain does not appear in the list of available NSX Managers"

occurs when SDDC Manager is not integrated into VCF Automation. Without this integration, VCF Automation cannot discover workload domains or their associated NSX Managers. As a result, when attempting to create a new region, the NSX Manager list is empty.

The required action is:

Add the SDDC Manager integration under VCF Automation → Provider Portal → Integrations.

This integration enables Automation to pull:

NSX Manager inventory

vCenter endpoints

Workload domain topology

Cluster details

Only after this integration is complete will the NSX Manager appear and allow region creation.

Option A and D (deploying new WLD or cluster) are unnecessary—inventory access is the problem, not resources.
Option B (triggering inventory sync) cannot work because no SDDC Manager integration exists.

Question: 41

An administrator attempts to add a new user (provideradmin05) within the VMware Cloud Foundation (VCF) Automation Provider Management Portal, however provideradmin05 cannot be found for import.

The following information is provided:

- The existing VCF Fleet uses VMware Identity Broker (VIDB) for single sign-on.
- VIDB uses Active Directory as the identity provider.
- A group named VCFA_ProviderAdmins was created in Active Directory, populated with the appropriate user accounts and synchronized with VIDB.
- Five days later provideradmin05 was added to VCFA_ProviderAdmins.

What will resolve this issue?

A. In VCF Operations, manually resync the directory.
B. In the VCF Automation Provider Management Portal, enable the Advanced Rights Bundle Mode.
C. In VCF Operations, disable VCF SSO for VCF Automation.
D. In the VCF Automation Provider Management Portal, import provideradmin05 as an LDAP user.

Answer: A

Explanation:

VMware Cloud Foundation (VCF) 9.x uses VMware Identity Broker (VIDB) as the central identity provider for the entire VCF fleet. VIDB synchronizes user and group metadata from the connected enterprise identity source, in this case Active Directory. When a user is added to an AD group after the group was already synced into VIDB, VIDB does not automatically resync group membership on demand unless a directory synchronization is performed.

In this scenario, the group VCFA_ProviderAdmins was synchronized five days earlier. When the new user provideradmin05 was later added to the AD group, VIDB—and therefore the VCF Automation Provider Management Portal—does not recognize that new user until a manual directory resynchronization occurs from VCF Operations.

This operation forces VIDB to:

Requery Active Directory

Update group membership information

Repopulate available users for import into VCF Automation

Options B and D are incorrect because they do not influence Identity Broker directory synchronization. Option C (disable VCF SSO) would break authentication and is not a valid solution.

Question: 42

An administrator is managing a VMware Cloud Foundation (VCF) environment. They receive a request from the developers to enable vDefend - Distributed Firewall. However, they noticed It cannot be enabled due to a missing license.

Where must the new license be applied?

A. SDDC Manager.

B. NSX Manager.

C. VCF Automation.

D. VCF Operations.

Answer: B

Explanation:

vDefend – Distributed Firewall is a security capability delivered by NSX within VMware Cloud Foundation. Although VCF components such as SDDC Manager, VCF Operations, and VCF Automation rely on licensing frameworks, the enforcement and activation of NSX features—including Distributed Firewall—occur entirely within NSX Manager.

To enable vDefend (Distributed Firewall), NSX Manager must detect a valid NSX license that includes security features. Without applying the correct license directly to NSX Manager:

The Distributed Firewall feature remains locked

vDefend cannot be enabled in workload domains

Security rules and micro-segmentation capability remain unavailable

VCF does not apply NSX security licensing at the SDDC Manager, VCF Automation, or VCF Operations layers. Instead, NSX Manager handles all feature entitlement checks internally.

Therefore, the new license must be installed directly in NSX Manager, under:

System → Licensing → NSX → Add License

Options A, C, and D are incorrect because none of those components control NSX feature activation.

Question: 43

An administrator logs into the VMware NSX Manager UI and discovers a time sync issue that has been reported in the VMWare Cloud Foundation (VCF) installer.

The administrator performs the following steps:

1. Validates that the NTP server IP addresses are present in the NTP configuration on the VCF Installer.
2. Validates that the DNS records are correctly set for the FQDN and IP address of the two NTP servers.
3. Validates that the NTP servers can be pinged by name and IP address from the VCF Installer.
4. Validates that the time between the NTP servers and the VCF Installer is synchronized successfully.

What additional step should the administrator perform to help identify the cause of the error?

A. Confirm that the ESX hosts have been configured to use host time synchronization.

B. Confirm that the NTP service has an allowed rule in the iptables on the VCF Installer.

C. Confirm that the NTP server details have been specified in the deployment parameter workbook using the required FQDN format.

D. Confirm that the time on the ESX hosts allocated for the management domain is synchronized with the same NTP servers as the VCF Installer.

Answer: D

Explanation:

During VMware Cloud Foundation bring-up, time synchronization across all management components is mandatory. The VCF Installer, ESXi hosts, NSX Manager nodes, and vCenter must all sync to the same NTP servers. If even one host or component has a time skew exceeding VMware's allowed limits, VCF will report time sync errors during bring-up or post-deployment.

The administrator validated NTP configuration, DNS resolution, ping connectivity, and time sync only on the VCF Installer appliance, but did not verify the ESXi hosts' time synchronization. NSX Manager obtains its time reference from the underlying ESXi host during deployment, so if the ESXi hosts are not synchronized with the same NTP sources, NSX Manager will drift, triggering the exact error described.

Option B (iptables) does not apply—the VCF Installer does not block outbound NTP by default. Option C refers to workbook formatting, which would fail earlier in deployment—not after NSX Manager is running.

Option A is incorrect because ESXi should never use "host time sync"; NTP must be used.

Question: 44

An administrator is automating the deployment of a new VMware Cloud Foundation (VCF) fleet using VCF Installer. The VCF fleet must include VCF Automation being deployed in a simple deployment model.

The administrator creates a JSON file, but during the installation attempt the VCF Installer returns an error indicating that the JSON validation has failed.

What is the cause of the errors?

A. VCF components binaries are not downloaded.

B. Second IP address for VCF Automation is not specified.

C. NSX Manager size was defined as large.

D. A separate distributed switch was defined for vSAN traffic.

Answer: B

Explanation:

In VCF 9.0, when deploying VCF Automation using the VCF Installer in a Simple Deployment Model, the appliance requires two IP addresses:

Primary IP – Management interface

Secondary IP – Required for service separation and internal routing for Automation services

VMware's JSON schema for VCF Installer enforces this requirement. If the second IP is missing, incorrectly formatted, or placed under the wrong JSON section, the installer validation will fail immediately with a JSON schema error before deployment begins.

This is one of the most common causes of validation failure for VCF Automation deployment.

Option A (component binaries missing) produces a bundle download error, not JSON schema failure.

Option C (NSX Manager size = large) is allowed and does not break JSON validation.

Option D (separate vDS for vSAN) is allowed if defined correctly and also does not cause JSON schema failure.

Question: 45

An administrator is troubleshooting network connectivity issues on a VMware ESX host configured with a dedicated VMware vSAN vSphere Distributed Switch (vDS) port group. The VMware vSAN vDS port group has two physical adapters and two uplinks assigned. After a failure of the active physical adapter, the vSAN vDS connection over the vSAN network was lost.

What is the cause of the issue?

A. The vSAN storage policies are misconfigured.

B. VLAN tagging is not correctly configured on the vDS.

C. A physical adapter is set to "Not Used" in the vDS configuration.

D. The vDS failover policy does not allow fallback.

Answer: C

Explanation:

In vSAN ESA or OSA networking configured through a dedicated vSphere Distributed Switch (vDS), each vSAN vmkernel port must have at least one Active physical uplink available at all times. The scenario describes a vDS with two physical adapters and two uplinks, but after failure of the active uplink, vSAN traffic was lost. This only occurs when the second physical NIC is not actually assigned to the vSAN port group—typically because its uplink is set to "Unused".

In such a misconfiguration:

vSAN traffic only uses the single active uplink.
When that uplink fails, vSAN has no failover path, causing immediate connectivity loss.

Option A (storage policies) does not affect network uplink behavior.
Option B (VLAN tagging) could cause connectivity failure but would not suddenly break only after an uplink failure.
Option D (failover policy not allowing fallback) affects recovery order, not immediate redundancy.

Question: 46

An Administrator has been tasked with creating a new VMware Cloud Foundation (VCF) Automation Region named Region-2. The following information has been provided:

- The current environment has two workload domains named WLD1 and WLD2.
- The workload domains share one NSX Local Manager deployment.
- A VCF Automation Region named region-1 exists that uses the shared NSX Local Manager deployment.

When creating the second Region in VCF Automation, the administrator sees "No results" when attempting to select a NSX Local Manager for the Region. What should the Administrator do to resolve this issue?

A. Add an additional NSX Edge Cluster In WLD1.

B. Deploy a third workload domain that includes a new, dedicated NSX Local Manager deployment.

C. Deploy an additional vSphere cluster in WLD1.

D. Ensure that that the NSX Manager is deployed in HA mode.

Answer: B

Explanation:

In VMware Cloud Foundation (VCF) Automation, each Automation Region must be associated with a dedicated NSX Local Manager. A single NSX Local Manager instance cannot be reused across multiple Automation Regions.

In the provided scenario:

The existing environment has WLD1 and WLD2, both sharing one NSX Local Manager.

Region-1 in VCF Automation already consumes this shared NSX Local Manager.

When creating Region-2, the interface shows “No results” when selecting an NSX Local Manager.

This behavior matches documented VCF Automation constraints: an NSX Local Manager can only be mapped to a single Automation Region. Once it is consumed by one region, it is not available for any additional region.

To create a second region (Region-2), a new NSX Local Manager instance must exist in the environment. The only supported method to obtain a new NSX Local Manager is to deploy a new workload domain, because NSX Local Manager is deployed as part of every VI Workload Domain.

Thus, the administrator must deploy a new (third) workload domain, which includes its own NSX Local Manager package, allowing Region-2 to be created successfully.

Question: 47

An administrator is attempting to import a certificate chain In VMware Cloud Foundation (VCF) Operations by uploading a certificate file. The validation fails with an error stating, "The provided certificate content is invalid.'

What is a possible cause for this error?

A. The certificate is not PEM-encoded.

B. The certificate chain order is invalid.

C. The certificate chain is missing the root CA.

D. The certificate chain does not include the private key.

Answer: A

Explanation:

VCF Operations enforces strict certificate format validation when importing certificate chains. According to VMware Cloud Foundation 9.x certificate management requirements, all uploaded certificates must be PEM-encoded. A PEM certificate must contain:

ASCII-encoded content

Proper headers such as:

---BEGIN CERTIFICATE ------

---END CERTIFICATE---------

If the certificate is encoded in DER, PFX, PKCS#12, or any non-PEM format, VCF Operations will reject the upload with

the error:

"The provided certificate content is invalid."

This matches the behavior described in the question.

Option B (chain order invalid) and Option C (missing root CA) can cause validation issues only after the certificate file is successfully parsed. The error described indicates the file itself cannot be parsed, which directly points to encoding.

Option D (missing private key) is incorrect because certificate chain uploads must NOT include a private key — private keys are only used during CSR signing and are handled separately by the system.

Question: 48

An administrator is adding a vSphere Supervisor using VMware NSX classic to an existing VMware Cloud Foundation (VCF) cluster using Distributed Connectivity. When attempting to enable the vSphere Supervisor for the domain the cluster shows up as incompatible with the reason:

No valid edge cluster for VDS 50 Ob 4d 9a cb 32 62 4d - 76 78 6b 92 cd 87 c4 5a

Why is the cluster showing up as incompatible?

A. The WCPReady tag has not been been assigned to the NSX Edge Cluster.

B. The NSX Edge transport nodes have been deployed as large.

C. vSphere Supervisor requires Central Connectivity.

D. AVI load balancing has not been enabled for the NSX Edge Cluster.

Answer: A

Explanation:

A Comprehensive and Detailed Explanation: When enabling vSphere Supervisor with NSX Classic (using the traditional NSX-T Data Center networking stack rather than the newer NSX VPC mode), the vSphere Workload Management wizard filters the list of available NSX Edge Clusters to ensure they are explicitly designated for use with Kubernetes workloads.

The "WCPReady" Tag Requirement: The primary mechanism vCenter uses to identify a valid, compatible Edge Cluster for Workload Management is a specific tag on the NSX Edge Cluster object. This tag must be WCPReady (case-sensitive).

Symptoms: If this tag is missing—which often happens if the Edge Cluster was created manually in NSX Manager rather

than through the SDDC Manager automation—the validation process will fail to find any usable clusters. This results in the specific error message: "No valid edge cluster for VDS [UUID]", or simply an empty list of compatible clusters in the wizard.

Resolution: The administrator must log in to the NSX Manager, navigate to System > Fabric > Nodes > Edge Clusters, select the target cluster, and manually add the tag WCPReady (often with the scope "Created for", though the tag itself is the critical filter).

Why other options are incorrect:

B: Large Edge nodes are actually a requirement for vSphere Supervisor (Small/Medium are typically unsupported for this role), so deploying them as Large would make the cluster compatible, not incompatible.

C: vSphere Supervisor fully supports Distributed Connectivity (connecting directly to the VDS), so Central Connectivity is not a hard requirement causing this specific error.

D: While AVI (NSX Advanced Load Balancer) is a supported load balancer, the "No valid edge cluster" error occurs during the Edge Cluster discovery phase, preceding the load balancer configuration.

Question: 49

A VMware Cloud Foundation (VCF) administrator cannot deploy Virtual Machines (VMs) to a compute cluster.

The administrator discovers that the vCLS VMs on the problematic cluster are powered off and cannot be powered on.

What action can the administrator take to enable deployment of VMs?

A. Delete all resource pools in the affected cluster.

B. Disable HA on the affected cluster.

C. Enable retreat mode on the affected cluster.

D. Set DRS Automation level to fully automated.

Answer: C

Explanation:

In vSphere 7+ and VCF-managed clusters, the vSphere Cluster Services (vCLS) VMs must remain powered on for DRS, cluster health, and policy enforcement to function. If the vCLS VMs cannot power on, no workloads—including new VMs—can be

deployed to the cluster because vSphere considers the cluster unhealthy.

A common cause is insufficient resources (CPU/memory), datastore issues, or policy conflicts preventing vCLS VMs from starting. VMware provides Retreat Mode as a troubleshooting mechanism to temporarily disable vCLS, allowing the administrator to deploy VMs and correct underlying issues. Enabling retreat mode:

Removes vCLS from the cluster

Restores ability to deploy VMs

Allows remediation of storage/placement issues

Can later be disabled to restore DRS health

Option A (deleting resource pools) does not restore vCLS VM power state.

Option B (disabling HA) does not affect vCLS behavior.

Option D (setting DRS automation level) does not correct vCLS placement problems.

Question: 50

A user attempts to deploy a catalog item into a vSphere Namespace in a VMware Cloud Foundation (VCF) Automation Organization for All Apps. The catalog item will not deploy into zone3.

The following information is provided:

- The vSphere Supervisor has three zones (zonel, zone2, zone3).
- The user has successfully deployed the catalog item into zonel and zone2 of the vSphere Namespace.

What is the cause of this issue?

A. The user does not have Project Advanced User role for the vSphere Namespace.

B. The vSphere Namespace is assigned the default large vSphere Namespace Class.

C. The vSphere Namespace does not include zone3.

D. The user does not have the Project User role for the vSphere Namespace.

Answer: C

Explanation:

In VMware Cloud Foundation (VCF) Automation for All Apps, a vSphere Namespace can span multiple Supervisor Zones. However, workloads—including catalog item deployments—can only be deployed into zones that are explicitly assigned to that Namespace. The user in the scenario successfully deploys into zone1 and zone2, which confirms that those zones are correctly associated with the Namespace.

The failure to deploy into zone3, while deployments into the other zones work, strongly indicates that zone3 is not

part of the Namespace configuration.

This behavior matches how Supervisor Zones function:

A zone must be added to the Namespace in Supervisor configuration.

If the zone is not associated, VCF Automation will not present it as an eligible deployment location, and deployment into that zone fails.

Option A and D (project roles) are incorrect because insufficient permissions would prevent deployment into any zone, not a single missing zone.

Option B (Namespace Class) is irrelevant because Namespace Classes define resource limits, not which Supervisor Zones the Namespace is mapped to.

Question: 51

An administrator is attempting to activate a new vSphere Supervisor for use with VMware Cloud Foundation (VCF) Automation on a newly deployed cluster. In the VMware vSphere client, when going through the vSphere Supervisor activation having selected VCF Networking with VPC, the
Virtual Private Cloud (VPC) Connectivity Profile dropdown is empty on the workload network page. The administrator verified that a Virtual Private Cloud (VPC) Connectivity Profile exists in NSX.

What is the cause of the issue?

A. The TO gateway is in active/active mode.

B. The vSphere Supervisor control plane is set to high-availability.

C. The selected NSX Project is the Default Project.

D. The default VPC has not been created.

Answer: C

Explanation:

When activating a vSphere Supervisor using VCF Networking with VPC, the Supervisor Workload Network must use a VPC Connectivity Profile. These profiles are scoped to an NSX Project, and cannot be consumed from the Default Project.

VCF Automation requires that:

A custom NSX Project be used for VPC networking integrations.

The Default Project cannot host Connectivity Profiles or VPC constructs intended for Supervisor activation.

Even though the administrator verified that a VPC Connectivity Profile exists in NSX, the Supervisor wizard will not

display it if:

The VPC Connectivity Profile belongs to a different project, or

The current selection is the Default Project, which blocks visibility.

This exact behavior—empty VPC Connectivity Profile dropdown—is documented when attempting Supervisor activation under the Default NSX Project.

Option A (T0 active/active) affects North-South routing but does not hide VPC profiles.

Option B (Supervisor HA mode) does not impact network profile selection.

Option D (missing default VPC) is incorrect because the wizard is complaining about availability of Connectivity Profiles, not VPC instances.

Question: 52

An administrator is creating an additional Organization for All Apps within VMware Cloud Foundation (VCF) Automation.

After logging into the VCF Automation Provider Management Portal UI, the administrator is only able to create new Organizations for All Apps.

What action can the administrator take to resolve the issue and complete the task?

A. Create the new Organization for VM Apps using the VCF Automation API.

B. Delete the existing Organization for VM Apps using the VCF Automation API.

C. Delete any existing Organizations for All Apps from the Provider Management Portal UI.

D. Enable the creation of new Organization for VM Apps feature in the Provider Management Portal UI.

Answer: D

Explanation:

In VMware Cloud Foundation (VCF) 9.0 Automation, Provider Administrators manage which types of Organizations can be created:

VM Apps Organizations

All Apps Organizations

These capabilities are controlled by Feature Flags within the VCF Automation Provider Management Portal. If the administrator logs in and only sees the ability to create All Apps Organizations, it means that the feature flag enabling VM Apps Organization creation has not been turned on.

VCF Automation requires the Provider Admin to explicitly enable creation of VM Apps Organizations, because doing so

exposes VM-centric consumption models and allows the environment to differentiate between VM-only and hybrid (VM + Kubernetes) application deployments.

Therefore, the administrator simply needs to navigate to:

Provider Management Portal → Administration → Feature Flags → Enable "Create VM Apps Organizations"

Option A (creating via API) is unnecessary—the UI will support it once the feature is enabled.

Option B (deleting existing VM Apps orgs) has no effect on feature availability.

Option C (deleting All Apps orgs) is unrelated and would not unlock VM Apps org creation.

Question: 53

Through the VMware NSX Manager user interface, the administrator has identified an issue with BGP peering. Which command on the NSX Edge Transport Node provides more information about the issue?

A. get edge-cluster status

B. get logical-routers

C. get edge-cluster history state

D. get log-file routing follow

Answer: D

Explanation:

When troubleshooting BGP peering issues on an NSX Edge Transport Node, VMware documentation directs administrators to examine routing logs, because BGP failures are often caused by adjacency negotiation errors, authentication mismatches, keepalive/hold timer issues, or route-policy failures.

The NSX Edge CLI command:

get log-file routing follow

streams real-time routing logs, including BGP daemon logs (bfdd, routed, wdog) and provides detailed insight into:

BGP session establishment and teardown

Keepalive and hold timer exchanges

Neighbor state transitions

Route advertisement or rejection

Authentication mismatches

MTU or connectivity issues on TEP / uplinks

This is the only command in the list that exposes diagnostic-level BGP information needed to troubleshoot peering.

Option A (edge-cluster status) shows cluster membership only.
Option B (get logical-routers) shows logical router configuration, not BGP logs.
Option C (edge-cluster history state) is unrelated to routing.

Question: 54

An administrator has successfully deployed and configured the Application Monitoring Telegraf Agent to 30 virtual machines through VMware Cloud Foundation (VCF) Operations.

After 24 hours, the administrator is alerted to the fact that no additional data has been collected since the agents were deployed on the virtual machines.

What could be the possible cause of the issue?

A. There is a time synchronization issue between the Telegraf Agent and the Cloud Proxy.

B. The Service Discovery Management Pack has not been configured.

C. Application monitoring has been configured to use a single Cloud Proxy rather than a Collector Group.

D. There is a compatibility issue between the version of Virtual Machine Hardware and VMware Tools.

Answer: A

Explanation:

Application Monitoring in VCF Operations uses Telegraf agents running inside virtual machines.
These agents forward metrics to the Cloud Proxy, which then sends them to the Operations analytics cluster. One of the most common reasons an agent stops reporting data—especially exactly 24 hours after deployment—is clock drift or time mismatch between the VM (running the Telegraf agent) and the Cloud Proxy.

VCF Operations enforces strict timestamp validation. If the timestamps from the agent are outside the acceptable drift window, the Cloud Proxy rejects incoming data as invalid. In this case, the Telegraf agents appear installed and functional, but no new metrics are received by the analytics engine.

This is a well-known issue documented in VMware Aria/VCF Operations agent-based monitoring, where:

Agents send metrics with local system time.

Cloud Proxy enforces time validation to prevent corrupt metric ingestion.

A drift >5 minutes commonly results in zero data collection despite healthy connectivity.

Options B and C cannot stop data flow after exactly 24 hours; they would prevent initial collection. Option D (virtual hardware/tools compatibility) affects VM operations but not Telegraf metric timestamp validation.

Question: 55

An administrator has identified that the VMware NSX Admin account is locked out. The administrator is unable to login to the NSX Manager UI using this account.

How could the administrator resolve this issue?

A. SSH into NSX Manager as Admin and remove API and CLI password lockouts.

B. Login into vCenter and increasing the password age policy.

C. Login to SDDC Manager and rotate admin account password.

D. Console into NSX Manager as root and clear API and CLI password lockouts.

Answer: D

Explanation:

When an NSX Admin account becomes locked in NSX Manager, this occurs due to failed login attempts exceeding the lockout threshold for either:

CLI access,

API access, or

UI login, which is tied to API authentication.

Once locked, the only supported method to recover the NSX admin account is to log in to the NSX Manager console as the root user and manually clear the lockout counters. This is documented in NSX Manager password-recovery procedures and is the standard administrative recovery action.

The root console provides access to:

clear account-lockout admin

or the equivalent reset methods within NSX Manager.

Why the other options are incorrect:

A. SSH into NSX Manager as Admin

Impossible — the admin account is locked and cannot be used to SSH.

B. Change password age policy in vCenter

NSX Manager accounts are not governed by vCenter password policy.

C. Rotate admin password in SDDC Manager

SDDC Manager rotates NSX passwords when unlocked; it cannot unlock a locked account.

Question: 56

An administrator logs into the VMware NSX Manage UI and observes a "Remote Logging Not Configured" alarm for each NSX Management node. What is a possible reason for this issue?

A. Update the NSX Edge Cluster Profile to configure a remote logging server.

B. Update the NSX Uplink Profile to configure a remote logging server.

C. Update the NSX Node Profile to configure a remote logging server.

D. Update the NSX Configuration Profile to configure a remote logging server.

Answer: C

Explanation:

The "Remote Logging Not Configured" alarm in NSX Manager is a system-health alert indicating that one or more Transport Nodes (Edges or Hypervisors) or Management Nodes do not have a Syslog server defined.

NSX Node Profiles: In VMware NSX (and by extension VCF), the standard method to apply consistent administrative settings—such as Syslog Servers, NTP settings, and Core Dump configurations—across a fleet of nodes is to use an NSX Node Profile.

Configuration Path: The administrator should navigate to System > Fabric > Profiles > Node Profiles. Here, they can create or edit a profile that specifies the remote syslog server's IP/FQDN, port, and protocol.

Application: Once the Node Profile is applied to the NSX Management Cluster or Edge Clusters, the configuration is pushed to all respective appliances, clearing the alarm.

Why not A/B: Edge Cluster Profiles manage networking/BFD settings; Uplink Profiles manage NIC teaming and MTU.

Question: 57

An administrator has received reports of high CPU ready times on several Virtual Machines (VMs) running within a VMware Cloud Foundation (VCF) with error" F.iain and has been tasked with collecting detailed metrics for all running Virtual Machines from each ESX host.
com.vmware.esx.setungs_aaemon.sonware.scan_spec.".

What is the cause of this error?

A. The cluster was not assigned a default image.

B. A component was removed from the image.

C. A vendor add-on was not provided in the image.

D. A device driver Is incompatible with the included firmware.

Answer: B

Explanation:

The error message com.vmware.esx.settings.daemon.software.scan_spec (reconstructed from the typo setungs_aaemon.sonware.scan_spec) is a specific failure generated by the vSphere Lifecycle Manager (vLCM) compliance engine on the ESXi host.

Cause - Removed Component: This error is a documented known issue in vSphere 8.x/VCF 5.x+, specifically occurring when an administrator has removed a component from the Cluster Image that the system validates as "required" or "structural" for the host's hardware configuration (common with hosts using DPUs or specific OEM add-ons).

The Scenario: The mention of "High CPU ready times" likely implies the administrator attempted to streamline the host image by removing a perceived "bloatware" component (like a vendor monitoring agent or unused driver) to improve performance.

The Result: When vLCM attempts to build the "Scan Specification" (scan_spec) to validate the host against this modified image, the internal struct validation fails because the removed component creates an invalid dependency state, throwing the Invalid field software_spec... or scan_spec exception.

Question: 58

An administrator has received reports of high CPU ready times on several Virtual Machines (VMs) running within a VMware Cloud Foundation (VCF) workload domain and has been tasked with collecting detailed metrics for all running Virtual Machines from each ESX host.

Which command line utility will enable the administrator to collect the required metrics?

A. vimtop

B. esxcli

C. vim-cmd

D. esxtop

Answer: D

Explanation:

To collect detailed per-VM CPU metrics—especially CPU Ready (%RDY)—the correct command-line utility on an ESXi host is esxtop. This tool provides real-time, low-level performance data for CPU, memory, disk, and network usage, and is the authoritative method for diagnosing CPU contention issues in VMware environments.

When troubleshooting high CPU Ready times, esxtop allows administrators to:

View CPU contention at the VM level

Inspect co-stop, wait, and scheduling delays

Monitor NUMA distribution and pCPU saturation

Capture historical performance snapshots using batch mode

The other options do not provide the necessary VM-level CPU scheduling metrics:

A . vimtop: Only available on vCenter Server Appliance (VCSA), not ESXi; does not show VM CPU ready.

B . esxcli: Used for configuration and health checks; not for real-time CPU metrics.

C . vim-cmd: Used to manage VMs via vSphere API bindings; not a performance monitoring tool.

Question: 59

An administrator is attempting to troubleshoot why the vSAN witness node cannot form a stretched cluster with the vSAN data nodes. The administrator can successfully ping the vSAN data node from the vSAN witness using the following command:

vmkping -I <witness-vmk#> <vsan-IPaddress> -s <1472> -d

What could be the possible cause of the issue?

A. Port 12321 is not opened bidirectionally between all nodes.

B. Port 443 is not opened bidirectionally between all nodes.

C. The customer does not have any virtual machines in the vSAN Cluster.

D. Jumbo Frames have not been enabled on the Witness Network.

Answer: A

Explanation:

In a vSAN Stretched Cluster, communication between the witness node and data nodes requires several specific TCP/UDP ports. The ability to successfully execute:

vmkping -I <witness-vmk> <vsan-IP> -s 1472 -d

confirms that:

L2/L3 connectivity is present

MTU is correctly configured

ICMP traffic flows without fragmentation

However, vmkping alone does not verify vSAN control-plane communication.

For the vSAN Witness to properly form a cluster, TCP port 12321 must be open bidirectionally between:

Witness → Data nodes

Data nodes → Witness

Port 12321 is required for:

vSAN cluster membership

Witness traffic

vSAN object health/state synchronization

If this port is blocked by firewall policy or misconfigured network ACLs, the nodes can ping each other, but vSAN witness traffic will fail, preventing the stretched cluster from forming.

Why the other options are incorrect:

B . Port 443 — Required for management, not cluster formation.

C . No VMs in cluster — Has no impact on witness formation.

D . Jumbo frames not enabled — Already ruled out by the successful 1472-byte vmkping with DF bit.

Question: 60

The administrator has to change the DRS automation level in preparation to upgrade the vCenter. When making this change through VCF Operations, the following error occurs: 'Internal Error: Failed to retrieve vim client'.

What is the possible cause of this error?

A. DRS Automation is already set on the vSphere Client.

B. The vCenter is overloaded with API requests from VCF Operations.

C. Connectivity issue between vCenter and VCF Operations.

D. Insufficient licensing for the advanced vCenter features.

Answer: C

Explanation:

The error:

“Internal Error: Failed to retrieve vim client”

occurs when VCF Operations cannot establish a functional API session with vCenter. The vim client is the internal vSphere API client library used by VCF Operations to perform cluster actions such as modifying DRS settings, powering on/off workloads, or retrieving inventory.

When this error appears, VMware documentation identifies these common root causes:

Loss of connectivity between VCF Operations and vCenter

DNS resolution issues

Network interruption

Stale or expired authentication tokens

Credential mismatch

If the vCenter password was changed manually, VCF Operations may be unable to authenticate.

vCenter services restarting or unavailable

If vCenter backend services (vpxd, sts, etc.) are unstable, VCF Operations cannot establish a vim session.

Option A is incorrect—DRS automation state in the vSphere Client does not cause vim client retrieval errors.
Option B (vCenter overloaded by API requests) would cause timeouts, not a vim client initialization failure.
Option D (insufficient licensing) affects feature use, not API connectivity.