
Exam : **3V0-11.26**

Title : Advanced VMware Cloud
Foundation 9 Administrator

Version : DEMO

1.Domain: SDDC Manager & Architecture Operations

An administrator is commissioning four ESXi hosts in SDDC Manager to expand an existing VI Workload Domain cluster. During the host validation phase, the validation task fails with an error indicating that the host network configuration does not match the network pool specifications.

Which configuration requirement must be satisfied on the ESXi hosts prior to successful commissioning in SDDC Manager?

- A. The ESXi host must have an active VMkernel adapter (`vmk1`) pre-configured with the designated vSAN IP address.
- B. The ESXi management interface (`vmk0`) must be assigned an IP address within the range specified by the target Network Pool.
- C. The ESXi host must have standard vSphere Standard Switch (vSS) with two active uplinks configured and DNS forward/reverse lookups verified.
- D. The ESXi host must have an existing NSX Transport Node Profile applied prior to SDDC Manager discovery.

Answer: C

Explanation:

- Correct Option (C): When commissioning ESXi hosts in SDDC Manager, the hosts must be installed with supported ESXi builds, configured with a standard vSphere Standard Switch (`vSwitch0`) using active physical uplinks, and have valid forward and reverse FQDN resolution in DNS. SDDC Manager automatically manages the creation of vSphere Distributed Switches (VDS), VMkernel ports for vSAN/Geneve, and NSX encapsulation during the cluster creation or expansion workflows.
- Distractor Analysis:
 - Option A: VMkernel adapters for vSAN (`vmk1` or designated) and vSAN IPs are automatically generated and assigned by SDDC Manager during the cluster creation/expansion task from the selected Network Pool; pre-creating them causes validation failures.
 - Option B: The management IP (`vmk0`) is configured during ESXi host staging and must match the FQDN in DNS. Network Pools in SDDC Manager provide IP address ranges strictly for vSAN, vSphere vMotion, and NSX Overlay (Host TEPs), not ESXi host management.
 - Option D: ESXi hosts must not be configured as NSX Transport Nodes prior to commissioning; SDDC Manager automates NSX host preparation and Transport Node Profile attachment when the host is added to a workload domain cluster.

2.Domain: Lifecycle Management (LCM)

An administrator is preparing to apply an upgrade bundle to an existing VMware Cloud Foundation 5.x multi-cluster VI Workload Domain using the SDDC Manager UI. Before initiating the upgrade, the administrator runs the pre-check task, and the vCenter Server pre-check generates a warning regarding insufficient disk space on the `/storage/seat` partition.

What is the immediate and supported administrative action required to resolve this condition before continuing the LCM workflow?

- A. Increase the virtual disk size allocated to `/storage/seat` in vCenter Server Appliance Settings and reboot the appliance.
- B. Log into the vCenter Server Appliance shell and execute `vmon-cli --restart all` to purge the SEAT temporary cache.
- C. Use the SDDC Manager SOS utility with the `--force-lcm-bypass` argument to override the disk

warning.

D. Enable log rotation and purge historical event and task data using the vSphere Client management interface.

Answer: D

Explanation:

- Correct Option (D): The `/storage/seat` partition on the vCenter Server Appliance (VCSA) stores Statistics, Events, Alarms, and Tasks. During pre-checks, LCM verifies partition thresholds to guarantee database schema migration and rollbacks. The supported procedure to remediate high disk utilization on `/storage/seat` without breaking LCM state is to truncate or purge historical events, tasks, and statistical data via the vSphere Client or retention policy settings.
- Distractor Analysis:
 - Option A: While increasing virtual disks on VCSA is supported for general maintenance, resizing `/storage/seat` on a VCF-managed vCenter requires careful coordination with SDDC Manager inventory metadata and may mask runaway event logging issues. Purging historical records is the first recommended operational step.
 - Option B: The `vmon-cli` tool manages vCenter services; restarting services does not reclaim disk space on the persistent SEAT database partition.
 - Option C: The `sos` utility does not provide a `--force-lcm-bypass` parameter; bypassing safety pre-checks during core component upgrades is unsupported and risks bricking the appliance mid-migration.

3.Domain: Software-Defined Networking (NSX)

An administrator needs to deploy an NSX Edge Cluster for a newly created VI Workload Domain to provide north-south routing and centralized stateful services. The workload domain cluster utilizes a single vSphere Distributed Switch (VDS) with two 25 GbE uplinks.

When deploying the NSX Edge Nodes through the SDDC Manager interface, which network design requirement is mandatory for the Edge TEP communication?

- A. Edge TEPs must reside on the same VLAN and IP subnet as the Host TEPs.
- B. Edge TEPs must utilize a dedicated VLAN ID distinct from the Host TEP VLAN ID.
- C. Edge TEP interfaces must be configured with static IP addresses without a default gateway.
- D. Edge TEP interfaces must terminate on dedicated physical uplinks not shared with the VDS.

Answer: B

Explanation:

- Correct Option (B): When deploying NSX Edge VMs on the same VDS and physical uplinks as the ESXi hosts (collapsed compute and edge design), Edge TEPs and Host TEPs cannot share the same VLAN ID. This is due to how Geneve encapsulation and packet forwarding operate on the hypervisor: if an Edge VM sends encapsulated Geneve traffic to a local Host TEP on the same VLAN, the local vSwitch drops the traffic to prevent loops. Therefore, Edge TEPs must use a distinct VLAN ID.
- Distractor Analysis:
 - Option A: Sharing the same VLAN and subnet between Edge TEPs and Host TEPs on a single VDS causes Geneve encapsulation drops and breaks BGP peering across TEP endpoints.
 - Option C: Edge TEPs require default gateways (or specific static routes) to allow Geneve encapsulation across routed layer-3 underlay networks.
 - Option D: VCF natively supports collapsed Edge deployments where Edge VM vNICs attach to trunk port groups on the existing VDS shared with system traffic.

4.Domain: Software-Defined Storage (vSAN)

An enterprise deploys a 4-host VI Workload Domain cluster utilizing vSAN Express Storage Architecture (ESA). The administrator is tasked with defining an appropriate Storage Policy-Based Management (SPBM) policy for business-critical virtual machines requiring high performance and resilience against a single host failure.

Which configuration represents the default and recommended failure tolerance method for this 4-host vSAN ESA cluster?

- A. RAID-1 (Mirroring) — 1 failure - RAID-1 (Mirroring)
- B. RAID-5 (Erasure Coding) — 1 failure - RAID-5 (Single Parity) with 4+1 scheme
- C. RAID-6 (Erasure Coding) — 2 failures - RAID-6 (Dual Parity) with 4+2 scheme
- D. RAID-5 (Erasure Coding) — 1 failure - RAID-5 (Single Parity) with 2+1 scheme

Answer: D

Explanation:

- Correct Option (D): In vSAN Express Storage Architecture (ESA), VMware introduced adaptive RAID-5 erasure coding. For clusters with 3 to 5 hosts, vSAN ESA automatically employs a 2+1 RAID-5 erasure coding scheme (consuming only 1.5x storage capacity compared to RAID-1's 2.0x while tolerating 1 failure). When the cluster scales to 6 or more hosts, vSAN ESA switches to a 4+1 scheme (1.25x overhead). In a 4-host cluster, 2+1 is the native supported RAID-5 configuration.
- Distractor Analysis:
 - Option A: RAID-1 is supported but is not the recommended default in vSAN ESA because ESA provides RAID-5 erasure coding with performance equal to or better than RAID-1, without the 100% capacity penalty.
 - Option B: The 4+1 RAID-5 scheme in vSAN ESA requires a minimum of 6 hosts (4 data + 1 parity + 1 spare/rebuild host).
 - Option C: RAID-6 (4+2 dual parity) requires a minimum of 7 hosts in vSAN ESA.

5.Domain: SDDC Manager & Architecture Operations

An administrator needs to rotate and replace the self-signed SSL/TLS certificates across all components in the Management Domain (SDDC Manager, vCenter Server, and NSX Manager) using an enterprise Microsoft Active Directory Certificate Services (AD CS) instance.

What is the correct operational procedure within SDDC Manager to accomplish this task?

- A. Manually generate OpenSSL private keys on each appliance, import the certificates into the appliance keystores, and click "Synchronize Certificates" in SDDC Manager.
- B. Configure the Microsoft CA authority under Security Settings in SDDC Manager, generate CSRs through SDDC Manager, submit them to the CA, and apply the issued certificates via the SDDC Manager UI.
- C. Use the vSphere Certificate Manager tool on vCenter to generate certificates for all SDDC components and push them via VMware Aria Suite Lifecycle.
- D. Log into the SDDC Manager appliance via SSH and run `/opt/vmware/vcf/commonsvcs/cert-manager --force-replace-all`.

Answer: B

Explanation:

- Correct Option (B): SDDC Manager features built-in Certificate Management automation. The

administrator connects SDDC Manager to the enterprise Certificate Authority (such as OpenSSL or Microsoft AD CS), generates Certificate Signing Requests (CSRs) directly from the SDDC Manager interface, obtains signed certificates, and instructs SDDC Manager to orchestrate certificate replacement across all endpoints (vCenter, NSX Manager, SDDC Manager) seamlessly.

- Distractor Analysis:

- Option A: Manually updating certificates directly on individual appliances bypasses SDDC Manager inventory, leading to trust store mismatches, broken API communications, and LCM failure during future updates.
- Option C: vSphere Certificate Manager only manages certificates local to vCenter Server and cannot coordinate certificate lifecycles for NSX Managers or SDDC Manager.
- Option D: There is no such command or directory path; direct filesystem manipulation of keystores without SDDC Manager orchestration breaks component trust.