

DUMPSBOSS.

VMware Cloud Foundation 5.2 Administrator

VMware 2V0-11.24

Version Demo

Total Demo Questions: 8

Total Premium Questions: 81

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, IT Architectures, Technologies, Standards	1
Topic 2, VMware by Broadcom Solution	16
Topic 3, Plan and Design	3
Topic 4, Install, Configure, Setup	19
Topic 5, Performance-tuning, Optimization	3
Topic 6, Troubleshooting and Repair	17
Topic 7, Administrative and Operational Tasks	22
Total	81

QUESTION NO: 1

What is the primary purpose of configuring a Virtual Distributed Switch (VDS) in a vSphere environment?

- A. To automate the deployment of virtual machines.
- B. To provide centralized management of network configurations across multiple ESXi hosts.
- C. To manage storage policies for datastores.
- D. To enable VM snapshots across multiple datastores.

ANSWER: B

Explanation:

To provide centralized management of network configurations across multiple ESXi hosts is correct because a vSphere Distributed Switch is a single logical switch spanning the hosts associated with it. Administrators configure distributed port groups, VLAN settings, uplink assignment, traffic shaping, NIC teaming, security policies, and other networking controls centrally through vCenter Server. The configuration is then applied consistently to the participating ESXi hosts, rather than requiring equivalent settings to be created and maintained separately on each host's standard switch. This centralized model reduces repetitive administration, helps avoid configuration drift, and supports consistent VM networking as workloads migrate between hosts with vMotion. A distributed switch also provides advanced networking capabilities and monitoring features that are managed at the vCenter level, while each host retains the operational switch components needed to forward traffic. VMware describes the vSphere Distributed Switch as providing centralized provisioning, administration, and monitoring of virtual networking across associated hosts. See the [VMware vSphere Distributed Switch documentation](#) and the [vSphere Networking guide](#).

QUESTION NO: 2

An administrator needs to ensure that their VMware Cloud Foundation-based private cloud solution can support multi-tenancy for different departments within their organization.

What two configurations and components are necessary to achieve this? (Choose two.)

- A. Configure vSAN storage policies for each department.
- B. Use vSphere Resource Pools to allocate resources to different departments.
- C. Deploy Aria Automation to enable self-service provisioning for tenants.
- D. Use NSX to create isolated network segments for each tenant.

ANSWER: B D

Explanation:

Use vSphere Resource Pools to allocate resources to different departments and **Use NSX to create isolated network segments for each tenant** provide the core compute-resource and network-isolation capabilities needed for departmental multi-tenancy. A vSphere resource pool creates a logical partition of cluster CPU and memory capacity. Administrators can assign reservations, limits, and shares so that each department receives an appropriate level of service and workloads in one department cannot consume all available compute resources needed by another. Resource pools can also be combined with permissions to delegate management within the intended scope.

NSX supplies the network-segmentation layer. Logical segments and security policies can separate each department's workload traffic while allowing administrators to define only the controlled connectivity required between departments or to shared services. Distributed firewall policies can be applied close to workloads, supporting consistent isolation even as virtual machines are provisioned or migrated within the environment. Together, resource pools address governed consumption of shared compute capacity, while NSX segments and protects the shared network fabric—two foundational controls for a multi-tenant private cloud. VMware documents resource-pool allocation controls in the [vSphere Resource Management documentation](#) and describes logical network segmentation in the [VMware NSX Administration documentation](#).

QUESTION NO: 3

An administrator is tasked with deploying a new application that requires both virtual machines and containers.

Which two components of VMware Cloud Foundation would be utilized to provide a unified platform for managing these workloads? (Choose two.)

- A. Aria Operations
- B. IaaS Control Plane
- C. HCX
- D. NSX
- E. Aria Suite Lifecycle

ANSWER: B D

Explanation:

IaaS Control Plane is the vSphere-integrated Kubernetes platform used in VMware Cloud Foundation to bring container and virtual-machine workloads under a common operational model. It enables administrators to provide Kubernetes namespaces and clusters while retaining vSphere-based governance, resource management, storage, and virtual-machine capabilities. This makes it the core workload platform when an application needs both traditional virtual machines and containerized services.

NSX complements that platform by delivering consistent software-defined networking and security for both workload types. It provides the networking services required by Kubernetes environments, including connectivity, load balancing capabilities where deployed, network segmentation, and distributed security controls. Using NSX alongside the IaaS Control Plane lets administrators apply consistent network policy and isolate application components regardless of whether those components execute in virtual machines or containers. Together, these components provide the compute, Kubernetes, networking, and security foundation needed to operate mixed workloads in VMware Cloud Foundation. VMware documents the vSphere Kubernetes capabilities in the [vSphere with Tanzu Configuration and Management documentation](#) and describes NSX integration within [VMware Cloud Foundation documentation](#).

QUESTION NO: 4

Which option should be chosen when placing a vSAN host into maintenance mode to ensure all data remains accessible during maintenance, while minimizing the time required to complete the operation?

- A. Quick Migration
- B. No Data Migration
- C. Ensure Accessibility
- D. Full Data Migration

ANSWER: C

Explanation:

Ensure Accessibility is the appropriate vSAN maintenance-mode data-evacuation option when the goal is to keep virtual-machine objects available while avoiding an unnecessary full evacuation of the host. Before placing the host into maintenance mode, vSAN evaluates the components residing on that host and evacuates only the components whose remaining placement would cause an object to become inaccessible. Components that can remain on the host without affecting object availability are not migrated. This targeted approach preserves access to vSAN data for running workloads while generally requiring less data movement, less temporary capacity, and less time than evacuating every component from the host. It is particularly suitable for routine, temporary host maintenance where the host is expected to return to the vSAN cluster after the work is complete. The operation still depends on sufficient healthy hosts and available capacity to create any

component copies required to maintain accessibility. VMware documents the maintenance-mode choices and their behavior in the [vSAN Administration documentation](#); see also the guidance in [vSAN maintenance mode options](#).

QUESTION NO: 5

A cloud administrator recently deployed a new VI workload domain. As part of the initial VI domain creation, the administrator created a new SSO domain. However, they reconsidered and now want the domain to use the same SSO domain as the management domain.

How can the VI domain be changed from a dedicated SSO domain to sharing the SSO domain with the management domain?

- A.** From the SDDC Manager console, use the rsutil command to add the VI workload domain to the management domain SSO ring.
- B.** The vSphere SSO domain cannot be changed once it has been deployed and a new VI workload domain that is part of the Management SSO domain must be created, and workloads migrated.
- C.** From the VI Workload Domain vCenter Server instance, use the rsutil command to join the management domain SSO ring.
- D.** From the Management Domain vCenter Server instance, use the rsutil command to add the VI workload domain to the management domain SSO ring.

ANSWER: B

Explanation:

The vSphere SSO domain cannot be changed once it has been deployed and a new VI workload domain that is part of the Management SSO domain must be created, and workloads migrated. In VMware Cloud Foundation, the SSO choice made during VI workload-domain deployment establishes the vCenter Server identity-domain relationship and the workload domain's SSO-ring membership. This is a foundational deployment configuration, not a post-deployment setting that can be converted from an isolated SSO domain into the management domain's shared SSO domain. Consequently, the supported approach is to deploy a replacement VI workload domain and select the management domain's existing SSO domain during its creation. The administrator can then migrate applicable workloads, networking dependencies, and operational configurations to that newly deployed domain, validating service operation before retiring the original workload domain. This approach preserves supported lifecycle management through SDDC Manager and avoids making unsupported identity-topology changes directly to deployed vCenter Server components. VMware Cloud Foundation documentation describes workload-domain deployment and the selection of an existing or new vSphere SSO domain as part of that deployment workflow. See the [VMware Cloud Foundation documentation](#) and the [VMware Cloud Foundation documentation portal](#).

QUESTION NO: 6

VMware Aria Operations reports sustained high memory use for several virtual machines on one ESXi host. The administrator needs to determine whether the host is under actual memory contention rather than observing normal guest operating system cache usage.

Which two metrics provide the strongest evidence of host memory contention? (Choose two.)

- A.** High guest operating system memory usage.
- B.** Virtual machine ballooning activity.
- C.** Virtual machine CPU ready time.
- D.** Virtual machine swap activity.
- E.** Virtual machine network throughput.

ANSWER: B D

Explanation:

Ballooning and virtual-machine swapping are memory-reclamation mechanisms used when ESXi experiences memory pressure. Sustained ballooning indicates that ESXi is requesting memory reclamation through the VMware Tools balloon driver. Virtual-machine swap activity indicates a more severe condition in which ESXi must swap VM memory to disk.

High guest memory usage alone does not prove contention because operating systems commonly use available memory for file cache. CPU ready time and network throughput can identify other resource constraints, but they do not establish that the host is reclaiming memory.

QUESTION NO: 7 - (DRAG DROP)


DRAG DROP

Put the following steps in the correct order to update the driver/firmware in vSAN.

☐ Verify that the new driver and firmware versions are active and the host is functioning correctly.	Step 1
☐ Download the latest driver and firmware from the hardware vendor's website.	Step 2
☐ Apply the firmware update to the host using the vendor's update tool.	Step 3
☐ Plan for a maintenance window to perform the update.	Step 4
☐ Check the current driver and firmware versions using the vSphere Client.	Step 5
☐ Reboot the host to complete the update process.	Step 6
☐ Update the driver using the vSphere Lifecycle Manager (vLCM)	Step 7

ANSWER:

Answer:



Answer:

Explanation:

A safe vSAN driver and firmware update starts by planning a maintenance window. This ensures the host can be handled according to the cluster's availability requirements, including placing it into the appropriate maintenance state before any update activity affects workloads or vSAN data availability. Before changing anything, record the currently installed driver and firmware versions through the vSphere Client. That baseline is important for confirming the intended upgrade path and for validating the final state after the maintenance work is complete.

Next, obtain the approved driver and firmware packages from the server hardware vendor. Firmware is applied using the vendor's supported update utility because that tool is designed for the physical server components. The ESXi driver is then updated with vSphere Lifecycle Manager, which provides the VMware-managed lifecycle workflow for the host software component. After both updates have been staged or applied, rebooting the host completes activation of components that require a restart and ensures ESXi loads the new driver and the hardware initializes with the updated firmware.

The final validation confirms that the new driver and firmware versions are active and that the host is functioning correctly after it returns to service. This closes the maintenance workflow with an explicit operational check rather than assuming a successful installation means the updated components are in use. VMware's guidance for firmware updates and lifecycle management supports using vendor-supported firmware tools together with vSphere lifecycle processes: [Updating firmware on ESXi hosts](#) and [vSphere Lifecycle Manager documentation](#).

QUESTION NO: 8

An administrator successfully removed a host from a VMware vSphere cluster and decommissioned it in SDDC Manager.

What must be done before the host can be recommissioned and added back to the SDDC Manager inventory?

- A. Log in to the host console and select "Reset System Configuration" from the ESXi DCUI.
- B. Manually uninstall the NSX VIBs on the host.
- C. From the host console, use the `/sbin/generate-certificates` command to generate a new self-signed certificate.
- D. Reinstall ESXi on the host.

ANSWER: A

Explanation:

Log in to the host console and select “Reset System Configuration” from the ESXi DCUI. is the required preparation step. When a host is decommissioned from VMware Cloud Foundation, it must be returned to a clean, unconfigured state before SDDC Manager can commission it again. Resetting the ESXi system configuration removes the host's existing configuration and restores it to the state expected by the VMware Cloud Foundation host-commissioning workflow. This prevents prior cluster, networking, storage, and management configuration from interfering with the validation and automated configuration that SDDC Manager performs when the host is brought back into inventory. The reset is performed locally through the ESXi Direct Console User Interface, after which the host can be supplied with the required management connectivity and used in the applicable commissioning process. VMware's Cloud Foundation administration documentation describes the lifecycle operations for hosts managed by SDDC Manager, including removal and subsequent reuse through commissioning. The vSphere installation and setup documentation also describes using the ESXi console to reset a host's system configuration. See the [VMware Cloud Foundation 5.2 Administration Guide](#) and [vSphere ESXi Installation and Setup documentation](#).