PassQueen
a.k.a dumpsqueen.com

# Professional VMware vSphere 7.x

## VMware 2V0-21.20PSE

**Version Demo**

**Total Demo Questions: 13**

**Total Premium Questions: 130**

**Buy Premium PDF**

https://passqueen.com

support@passqueen.com

# Topic Break Down

| Topic | No. of Questions |
|---|---|
| Topic 1, Architectures and Technologies | 38 |
| Topic 2, VMware Products and Solutions | 16 |
| Topic 3, Installing, Configuring, and Setup | 19 |
| Topic 4, Performance-tuning, Optimization, and Upgrades | 18 |
| Topic 5, Troubleshooting and Repairing | 8 |
| Topic 6, Administrative and Operational Tasks | 31 |
| Total | 130 |

QUESTION NO: 1

Which two VMware features are supported in a virtual machine using Intel Software Guard

Extensions (Intel SGX) technology? (Choose two.)

**A.** Fault tolerance

**B.** vSphere vMotion

**C.** VM encryption

**D.** Virtual machine suspend and resume

**E.** Storage I/O Control

**ANSWER: C E**

**Explanation:**

VM encryption and Storage I/O Control are supported for a virtual machine configured to use Intel Software Guard Extensions (Intel SGX). Intel SGX exposes protected enclave capability to a guest while vSphere continues to provide compatible VM-level security and storage-quality-of-service functions. VM encryption protects the virtual machine's files, such as its configuration and virtual disks, through vSphere's key-management integration; this protection is distinct from the enclave protections that Intel SGX provides inside the guest. Therefore, using SGX does not prevent the VM from being protected with VM encryption.

Storage I/O Control is also supported because it operates at the datastore and storage-resource-management layer. It can continue to regulate and prioritize storage access for an SGX-enabled VM just as it does for other eligible workloads. VMware documents specific operational limitations for SGX-enabled VMs, while identifying VM encryption and Storage I/O Control as supported features. See the VMware vSphere Security documentation on Intel SGX support and limitations and the Intel SGX overview.

**QUESTION NO: 2**

An administrator is investigating the reason for a virtual machine restarting. In the vmware.log file, the administrator finds the following error:

```
svga | I120: WinBSOD: ( 1) 'A problem has been detected and Windows has
been shut down to prevent damage'
svga | I120: WinBSOD: ( 2) 'to your computer.'
svga | I120: WinBSOD: ( 4) 'If this is the first time you've seen this
Stop error screen,'
svga | I120: WinBSOD: ( 5) 'restart your computer. If this screen
appears again, follow'
svga | I120: WinBSOD: ( 6) 'these steps:'
svga | I120: WinBSOD: ( 8) 'Disable or uninstall any anti-virus, disk
defragmentation'
svga | I120: WinBSOD: ( 9) 'or backup utilities. Check your hard drive
configuration,'
svga | I120: WinBSOD: ( 10) 'and check for any updated drivers. Run
CHKDSK /F to check'
svga | I120: WinBSOD: ( 11) 'for hard drive corruption, and then restart
your computer.'
svga | I120: WinBSOD: ( 13) 'Technical information:'
svga | I120: WinBSOD: ( 15) '*** STOP: 0x00000024
(0x00000000001904FB, 0xFFFFF880069B2CE8, 0xFFFFF880069B2540,0'
svga | I120: WinBSOD: ( 16) 'xFFFFF880014F5FA2)
```

What is a likely cause of this virtual machine restart? (Choose the best answer.)

**A.** Guest operating system fault

**B.** User or API request to shut down guest operating system

**C.** High Availability

**D.** User or API request to restart guest operating system

**ANSWER: A**

**Explanation:**

Guest operating system fault is the likely cause. VMware records significant virtual-machine events in `vmware.log`, including conditions detected by the virtual CPU that cause the guest to stop unexpectedly and subsequently restart. A guest fault can arise from an operating-system kernel failure, a driver defect, corrupted guest components, or an application or workload that causes an unrecoverable error. In particular, fatal CPU exception conditions reported in the virtual-machine log indicate that the guest itself encountered a failure state; they are not evidence that the hypervisor independently initiated a planned power operation. The appropriate follow-up is to correlate the timestamp in `vmware.log` with the guest operating system's event, kernel, or crash logs, then investigate any dump files and recently changed guest drivers, updates, or software. VMware's guidance for troubleshooting unresponsive or failed guest systems also emphasizes examining both the virtual-machine log and guest-level diagnostic information to identify the originating fault. See VMware virtual machine log-file locations and troubleshooting information and vSphere Troubleshooting documentation.

**QUESTION NO: 3**

An administrator is tasked with upgrading an existing vSphere environment to version 7.0.

The current configuration is using an external Platform Services Controller.

Which two statements are true regarding this upgrade? (Choose two.)

**A.** The external Platform Services Controller needs to be converged into an embedded vCenter Server Appliance.

**B.** The external Platform Services Controller needs to be decommissioned following the deployment of the new vCenter Server Appliance.

**C.** The external Platform Services Controller backup can be restored to a new embedded vCenter Server Appliance.

**D.** The external Platform Services Controller needs to be upgraded after the vCenter Server Appliance.

**E.** The external Platform Services Controller needs to be powered down while a new vCenter Server Appliance is deployed.

**ANSWER: A B**

**Explanation:**

vSphere 7.0 no longer supports a deployment topology that uses an external Platform Services Controller. Therefore, an environment using an external Platform Services Controller must be converted to the embedded deployment model as part of its transition to vSphere 7.0. The external Platform Services Controller needs to be converged into an embedded vCenter Server Appliance because convergence moves the required Platform Services Controller services, including Single Sign-On domain data and associated infrastructure services, into the embedded appliance topology supported by vSphere 7.0.

Following successful deployment and convergence to the new embedded vCenter Server Appliance, the former external Platform Services Controller needs to be decommissioned. Decommissioning is an essential completion activity: the embedded appliance becomes the active provider of the migrated services, and retiring the obsolete external node prevents an unsupported or confusing residual topology. Administrators should validate the new embedded vCenter Server Appliance, confirm service health and inventory access, and then perform the documented decommissioning procedure. VMware documents the supported upgrade requirement and workflow for external Platform Services Controller deployments in its vCenter Server Upgrade documentation.

**QUESTION NO: 4**

Which two performance tools can be used to monitor virtual machines outside the guest operating system? (Choose two.)

**A.** vCenter Server performance charts

**B.** Task Manager

**C.** top

**D.** esxtop

**E.** Perfmon

**ANSWER: A D**

**Explanation:**

vCenter Server performance charts and esxtop can both monitor a virtual machine from the vSphere infrastructure layer, without requiring a monitoring utility to run inside the guest operating system. vCenter Server performance charts present collected performance-counter data for inventory objects, including virtual machines. They provide visibility into key resource metrics such as CPU usage, memory usage, disk activity, network throughput, latency, and resource contention, allowing an administrator to investigate VM performance centrally through the vSphere Client.

esxtop is an ESXi-host command-line performance-monitoring utility. It displays real-time statistics for host resources and for the virtual machines executing on that host. Its VM-focused views expose CPU scheduling and ready time, memory consumption and reclamation, virtual-disk I/O and latency, and virtual-network activity. Because esxtop runs in the ESXi environment rather than within a VM, it is particularly valuable when guest access is unavailable or when the administrator needs to correlate a VM's behavior with host-level resource usage. VMware documents performance-chart monitoring in the vSphere Monitoring and Performance documentation and provides guidance for using esxtop performance statistics.

**QUESTION NO: 5**

An administrator wants to create a master copy of an existing virtual machine (VM) named

“TestApp” and deploy VMs from this master copy whenever a user requests one.

Which method can the administrator use to create a master copy of TestApp? (Choose the best answer.)

**A.** Deploy from an Open Virtual Machine Format (OVF) template

**B.** Convert a template to a VM

**C.** Clone the VM to a template

**D.** Create a new VM

**ANSWER: C**

**Explanation:**

Clone the VM to a template is correct because a vSphere template is specifically designed to serve as a reusable master image for deploying additional virtual machines. The administrator can clone the existing TestApp virtual machine directly into a template, preserving its configured guest operating system, installed application, virtual hardware settings, and other desired baseline configuration. After the template is created, it can be selected as the source when deploying a new VM through the vSphere Client. Each deployment produces an independent virtual machine based on the master copy, allowing requested instances to be created consistently and efficiently.

This workflow is appropriate when an existing VM has already been prepared and validated as the desired standard image. Before creating the template, administrators commonly ensure that the source VM is in the required state and that its operating-system customization requirements have been considered, so deployed VMs can receive unique settings such as host names and network identities. VMware documents that templates are master copies of VMs that can be used to create and provision new virtual machines, and that cloning can create a template from an existing VM. See the vSphere documentation on cloning virtual machines and templates and the vSphere documentation on VM templates.

## QUESTION NO: 6

In a vSphere High Availability (vSphere HA) cluster, which condition can be detected by datastore heartbeating? (Choose the best answer.)

**A.** Fault tolerance virtual machine (VM) failover

**B.** Network isolation

**C.** Virtual machine (VM) disk failure

**D.** Datastore All Paths Down (APD) event

## ANSWER: B

**Explanation:**

**Network isolation** is correct. vSphere HA primarily uses management-network heartbeats to determine whether hosts can communicate with one another. If those network heartbeats are lost, the remaining hosts must determine whether the affected host has actually failed or is still running but isolated from the management network. Datastore heartbeating provides this additional evidence: each HA host periodically updates a heartbeat region on selected shared datastores. When the host's network heartbeats are absent but its datastore heartbeat continues to be updated, vSphere HA can determine that the host is still active and has likely experienced network isolation rather than a complete host failure. This distinction helps HA make an appropriate availability decision and avoid unnecessarily restarting virtual machines that may still be running on the isolated host. Datastore heartbeating is therefore an important secondary communication mechanism for confirming host liveness when management-network communication is interrupted. VMware documents datastore heartbeats as part of HA's host-failure detection process in the vSphere Availability documentation and in the vSphere Availability Guide.

## QUESTION NO: 7

Which feature is a pre-requisite for configuring admission control on a cluster? (Choose the best answer.)

**A.** vSphere Enhanced vMotion Compatibility (EVC)

**B.** vSphere High Availability (HA)

**C.** vSphere Fault Tolerance

**D.** vSphere Distributed Resource Scheduler (DRS)

## ANSWER: B

**Explanation:**

vSphere High Availability (HA) is required because admission control is an HA cluster policy. HA monitors ESXi hosts and virtual machines and, after a host failure, attempts to restart the affected virtual machines on surviving hosts. Admission control supports this availability objective by reserving sufficient uncommitted cluster capacity for failover. When enabled, it can prevent a virtual machine from being powered on if doing so would leave the cluster unable to satisfy the configured host-failure tolerance policy. In vSphere 7, HA admission control can use dedicated failover hosts, a slot-based policy, or a cluster-resources-percentage policy, depending on the cluster's requirements. These settings are available as part of configuring vSphere HA for the cluster, so HA must first be enabled before admission control can be configured. This relationship ensures that capacity reservations are tied directly to the cluster's automated restart and failover behavior rather than being a general-purpose resource-management setting. VMware documents admission control under the vSphere Availability and HA configuration guidance: vSphere HA Admission Control and vSphere Availability documentation.

## QUESTION NO: 8

What are two supported ESXi boot options? (Choose two.)

**A.** NFS

**B.** iSCSI

**C.** vSAN

**D.** SAN

**E.** vSphere Virtual Volumes

**ANSWER: B D**

**Explanation:**

iSCSI and SAN are supported ESXi boot options. ESXi can boot from an iSCSI logical unit number when the server's network adapter, firmware, and boot configuration support iSCSI boot. The iSCSI boot LUN contains the ESXi system image, allowing the host to start across an IP storage network rather than from a local disk. VMware documents specific requirements for iSCSI boot, including correctly configured network connectivity, adapter/firmware support, and an accessible boot target.

SAN boot is also supported and describes booting an ESXi host from storage presented through a storage-area network rather than a directly attached local boot device. This includes supported SAN-based implementations such as Fibre Channel, Fibre Channel over Ethernet, and iSCSI, subject to the host adapter, storage array, multipathing, and VMware compatibility requirements. Centralizing boot media on SAN storage can simplify hardware replacement and operational management because the ESXi installation resides on shared storage accessible to the host during startup. VMware's vSphere Storage documentation covers SAN boot configuration and operational considerations, while Broadcom's knowledge base summarizes key design considerations for ESXi boot from SAN storage. See vSphere 7 Storage documentation and Considerations for booting ESXi from SAN.

**QUESTION NO: 9**

Which two configuration tasks must be performed from the vCenter Server Management

Interface? (Choose two.)

**A.** Add administrator accounts.

**B.** Add an Network Time Protocol (NTP) server.

**C.** Add a vCenter Server license.

**D.** Enable SSH login.

**E.** Change the logging level.

**ANSWER: B C**

**Explanation:**

Add an Network Time Protocol (NTP) server and Add a vCenter Server license are performed through the vCenter Server Management Interface (VAMI), which is accessed on port 5480 of the vCenter Server Appliance. The appliance's time configuration is managed there because the VAMI provides the appliance-level time settings used to select a time-synchronization mode and specify NTP servers. Accurate and consistent time is essential for vCenter services, authentication, certificates, logging, and cluster operations.

The VAMI also provides the appliance-level licensing workflow for the vCenter Server instance. An administrator can review the currently assigned license and assign or replace the vCenter Server license key from its licensing area. This is distinct from managing inventory and broader vSphere licensing administration in the vSphere Client; the task named in the question concerns licensing the vCenter Server Appliance itself. VMware documents appliance administration functions, including time configuration and licensing, in the vCenter Server Configuration guide. See the vCenter Server Configuration documentation and the vSphere Installation and Setup documentation.

## QUESTION NO: 10

An administrator wants to configure vCenter Server High Availability so that each node is located in a different data center.

Which requirement must be met? (Choose the best answer.)

**A.** Each node must be configured with at least three network adapters.

**B.** Network latency between the nodes must be less than 10 milliseconds.

**C.** All nodes must be located on a vSAN datastore in each data center.

**D.** A vCenter Server license must be assigned to each node.

## ANSWER: B

**Explanation:**

Network latency between the nodes must be less than 10 milliseconds. vCenter Server High Availability uses an active, passive, and witness node, with continuous replication from the active node to the passive node over the dedicated vCenter HA network. When the nodes are distributed between data centers, the inter-site network must still provide sufficiently low latency for this replication and for reliable failover monitoring. VMware specifies a round-trip latency of less than 10 ms between vCenter HA nodes. Meeting this threshold helps maintain the replication stream, prevents health-state instability caused by delayed heartbeats, and allows the passive node to remain suitably current for takeover if the active node fails. The requirement applies to the network path used by the vCenter HA traffic, so administrators should validate the actual end-to-end latency between the sites before deploying the cluster. The nodes can be placed on separate ESXi hosts and in separate locations for resiliency, but the placement does not remove the vCenter HA network-performance requirement. See VMware's vCenter Server High Availability requirements documentation at vCenter Server High Availability Requirements and the vCenter Server Configuration documentation.

## QUESTION NO: 11

Which two statements are true regarding the use of firmware update packages in vSphere

Lifecycle Manager? (Choose two.)

**A.** A hardware support manager must be registered as a vCenter Server extension.

**B.** The firmware and drivers add-on are distributed through the official VMware online depot.

**C.** The firmware and drivers add-on can be used on clusters that are managed with baselines.

**D.** The firmware and drivers add-on can add or remove components from the ESXi base image.

**E.** The firmware and drivers add-on is provided by a third-party vendor.

## ANSWER: A D

**Explanation:**

A hardware support manager must be registered as a vCenter Server extension. The Hardware Support Manager integration is the mechanism through which vSphere Lifecycle Manager communicates with supported hardware-management software for firmware lifecycle operations. Registering the Hardware Support Manager as an extension establishes the required vCenter Server integration and enables Lifecycle Manager to obtain the firmware information and desired-state data needed for image-based cluster management.

The firmware and drivers add-on can add or remove components from the ESXi base image. In an image-managed cluster, Lifecycle Manager constructs the desired image from the selected ESXi base image and additional image elements. A firmware and drivers add-on can declare component-level changes required by the hardware platform, including components that must be added to, or removed from, the base ESXi image. Lifecycle Manager then validates that combined desired image and remediates hosts toward the validated image state. This integration allows firmware requirements and the

associated driver/component requirements to be handled together as part of host lifecycle management. See the VMware vSphere Lifecycle Manager documentation on vSphere Lifecycle Manager and firmware updates with Lifecycle Manager.

## QUESTION NO: 12

Which solution can be used to automatically deploy a fully configured VMware softwaredefined data center

(SDDC)? (Choose the best answer.)

**A.** vCenter Server

**B.** vRealize Suite

**C.** vSphere Lifecycle Manager

**D.** VMware Cloud Foundation

## ANSWER: D

**Explanation:**

VMware Cloud Foundation is the solution designed to deploy and operate a complete, standardized software-defined data center as an integrated platform. It combines the required SDDC components—including vSphere compute virtualization, vSAN storage virtualization, NSX network virtualization, and management services—into workload domains with validated architecture and configuration. Its Cloud Builder deployment workflow automates the initial bring-up process, converting prepared physical hosts into a fully configured management domain. This creates the foundational SDDC environment with consistently deployed infrastructure services rather than requiring each product to be installed and configured separately.

After initial deployment, VMware Cloud Foundation also provides lifecycle management capabilities for the integrated stack, helping maintain supported component versions and configurations across the environment. This full-stack automation and lifecycle approach is the defining use case for VMware Cloud Foundation: delivering an operational SDDC platform from infrastructure bring-up through ongoing management. VMware describes Cloud Foundation as its integrated hybrid-cloud platform for deploying and managing a complete cloud infrastructure stack. See the VMware Cloud Foundation product page and the VMware Cloud Foundation documentation.

## QUESTION NO: 13

An administrator is tasked with upgrading an existing vSphere environment to version 7.0. The environment consists of mounted NFS v3 datastores. The new design requires configuration of NFS v4.1 datastores.

Which statement is true regarding this upgrade? (Choose the best answer.)

**A.** The NFS v3 datastores can be unmounted and remounted using NFS v4.1.

**B.** A new VMKernel port for NFS v4.1 needs to be created before the upgrade.

**C.** Following first boot, NFS v3 datastores are automatically converted to NFS v4.1 by ESXi 7.0

**D.** ESXi 7.0 does NOT support NFS v4.1.

**E.** The NFS v3 datastore must be replaced by a newly mounted NFS v4.1 datastore, and workloads must be migrated to it.

## ANSWER: E

**Explanation:**

The NFS v3 datastore must be replaced by a newly mounted NFS v4.1 datastore, and workloads must be migrated to it. NFS protocol selection is made when ESXi mounts the datastore; it is not an in-place property that can be changed for an existing vSphere datastore. Consequently, moving from NFS v3 to NFS v4.1 requires presenting an export that is available through NFS v4.1, mounting it in vSphere as a new NFS 4.1 datastore, and then migrating the virtual machines and their files to that datastore. The migration can be performed with Storage vMotion when the environment and licensing support it,

allowing workloads to remain powered on, or by using an appropriate offline migration process. This design approach also permits validation of NFS v4.1 connectivity, authentication configuration where applicable, and datastore accessibility from every ESXi host before production workloads are moved. VMware documents NFS datastores as distinct datastore types and describes the configuration of NFS 4.1 storage separately from NFS 3 storage in the vSphere Storage documentation. VMware migration guidance is also available in the vSphere virtual-machine migration documentation.