

Advanced Design VMware vSphere 7.x

VMware 3V0-21.21

Version Demo

Total Demo Questions: 11

Total Premium Questions: 111

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, IT Architectures, Technologies, Standards	2
Topic 2, VMware Products and Solutions	7
Topic 3, Planning and Designing	96
Topic 4, Installing, Configuring, and Setup	3
Topic 5, Performance-tuning, Optimization, and Upgrades	3
Total	111

QUESTION NO: 1

Following a recent acquisition, the architect learns that both companies use vSphere on-premise and will need to combine the data centers into one. The acquired company's licenses will not be renewed for cost-savings related to the acquisition. All consumed vSphere licenses must have active support to support line-of-business operations. The merged environment must maintain 25% spare capacity. The architect has a small budget remaining unallocated for hardware.

The architect has calculated that the current vSphere environment can absorb the acquired company's virtual machines but the cluster will run at 90% memory utilization and at 50% CPU utilization.

Which design decision can the architect make to incorporate the new company's virtual machines into the combined vSphere environment?

- A. Migrate the acquired company's virtual machines into the vSphere environment as it will currently fit.
- B. Use the current budget to add memory to the cluster to increase each ESXi host's capacity and add the new virtual machines.
- C. Purchase extra hosts to add to the cluster in anticipation of adding the acquired company's virtual machines.
- D. Purchase new licenses for some of the acquired company's ESXi hosts and add them to the cluster to hold the acquired company's virtual machines.

ANSWER: B

Explanation:

Use the current budget to add memory to the cluster to increase each ESXi host's capacity and add the new virtual machines. This directly addresses the constrained resource: after consolidation, memory utilization would be 90%, while CPU utilization would be only 50%. A requirement for 25% spare capacity means the design must retain at least one quarter of usable capacity as headroom, so memory consumption must be reduced to no more than approximately 75% of the available cluster memory capacity. Increasing memory in the existing ESXi hosts expands the cluster's usable memory capacity and can bring the projected utilization within that required headroom target. It is also aligned with the stated small remaining hardware budget because memory expansion is generally an incremental, targeted upgrade rather than a full host acquisition. The decision permits the acquired virtual machines to be migrated into the supported, actively licensed existing environment, avoiding dependence on licenses that will not be renewed. VMware describes how vSphere DRS balances workloads based on CPU and memory resource demand, making sufficient cluster memory capacity essential to maintaining placement and balancing flexibility: [vSphere Resource Management documentation](#). VMware's capacity-planning guidance also emphasizes accounting for resource utilization and growth when sizing clusters: [vSphere Capacity Planning](#).

QUESTION NO: 2

An architect is tasked with designing a greenfield VMware software-defined data center (SDDC) solution that will be used to deliver a private cloud service for a customer.

During the initial meeting with the service owner and business sponsor, the customer has provided the following information to help inform the design:

- The solution must initially support the concurrent running of 300 production and 600 development virtual machines.
- The production environment should be delivered across two geographically dispersed data centers.
- The development environment must be vSphere-based but does not have to be deployed onpremises.
- The two data centers are connected to each other through multiple diversely routed, high bandwidth and low latency links.
- The customer's server hardware standard document states that all virtual infrastructure hosts must be based on blade architecture only.
- The service owner has said that is important to ensure that neither the availability target of 99.5% nor the resource capacity is affected when the operations team completes maintenance activities, such as the monthly software patching and ad-hoc hardware break/fix.
- All virtual machine backups must be completed using the existing backup service.

- The recovery time objective (RTO) for the service is four hours.
- The recovery point objective (RPO) of the service is 24 hours.

Given the information from the customer, which two would be classified as assumptions within the design? (Choose two.)

- A. The backup service will store data in a secure facility
- B. The backup service has sufficient capacity for the new requirements
- C. The customer will update their hardware standard to support rack mount servers
- D. All virtual machines will be deployed with the same resource profile for production and development
- E. The clusters will have a minimum redundancy of N+1

ANSWER: A B

Explanation:


“The backup service will store data in a secure facility” is an assumption because the stated requirement only mandates use of the existing backup service; it does not establish the physical security, location, or protection characteristics of that service’s backup repository. The architect must treat that capability as presumed until it is validated with the backup-service owner and relevant security controls or agreements.

“The backup service has sufficient capacity for the new requirements” is also an assumption. The requirements identify the number of virtual machines, require use of the existing service, and define recovery objectives, but they provide no confirmed backup-storage capacity, ingest throughput, retention policy, backup-window availability, or restore performance. Those characteristics must be validated to establish that the service can protect the additional workloads and support recovery within the stated objectives. VMware’s vSphere availability guidance emphasizes planning recovery and protection capabilities against service requirements, while VMware backup integration relies on the backup product and infrastructure to supply sufficient repository and operational capacity. See the [vSphere Availability documentation](#) and the [vSphere API documentation](#) for VMware platform context.

QUESTION NO: 3

Refer to the exhibit.

During a requirements gathering workshop, the customer shares the following about their existing ESXi host virtual networking infrastructure:



The customer confirms that:

- Each ESXi host has approximately 200 virtual machines.
- They want to maximize the number of concurrent virtual machine migrations.
- When placing a host in maintenance mode, it takes a long time to evacuate the virtual machines.

Which two recommendations should the architect make in order to help the customer overcome their challenge? (Choose two.)

- A. Configure the network to use MTU for the VMotion VMKernel to 1,600 bytes
- B. Configure the network to use MTU for the VMotion VMKernel to 9,000 bytes
- C. Create an additional standard switch with pNIC3 to use for vMotion
- D. Use the 3 pNICs and bundle them in a link aggregation group (LAG) configuration
- E. Use 10 GbE NICs instead of 1 GbE

ANSWER: B E

Explanation:

Configuring the vMotion VMkernel network for an MTU of 9,000 bytes and using 10 GbE NICs directly address the throughput constraint that delays evacuation of hosts with approximately 200 virtual machines. vMotion transfers a virtual machine's memory state over the VMkernel network; therefore, the available physical-network bandwidth is a primary factor in how quickly migrations can complete and how many can make useful progress concurrently. Moving from 1 GbE to 10 GbE substantially increases available capacity for the vMotion traffic, reducing the time required to evacuate a host during maintenance.

Jumbo frames, configured end-to-end at 9,000-byte MTU across the ESXi VMkernel adapter, physical-switch ports, and all intervening network paths, reduce packet-processing overhead by carrying more payload per frame. VMware documents jumbo frames as an MTU of 9,000 bytes and notes that they can improve network efficiency where the complete path supports them. Combined with 10 GbE connectivity, this provides a higher-capacity, more efficient vMotion path suited to frequent or concurrent migrations. See VMware's [Jumbo Frames and VMware ESXi](#) guidance and the [vMotion requirements](#).

QUESTION NO: 4

An architect is designing a vSphere cluster that will host virtual machines with different business owners. The customer requires that each business owner can view and manage only its own virtual machines, while the underlying cluster remains shared.

Which design decision should the architect make?

- A. Create separate virtual machine folders and assign permissions to the required business-owner groups.
- B. Create a separate datastore cluster for each business owner.
- C. Apply a vSphere tag for each business owner.
- D. Create a separate vSphere cluster for each business owner.

ANSWER: A

Explanation:

Create **separate virtual machine folders and assign permissions to the required business-owner groups**. Folders provide inventory organization and can be used as permission boundaries in vCenter Server. Assigning appropriate roles and privileges to each group on its folder allows delegated administration of the virtual machines within that folder without granting visibility or management privileges to virtual machines belonging to other owners.

Creating separate clusters would provide a stronger resource and operational boundary, but it is unnecessary when the stated requirement is limited to inventory visibility and delegated management. Tags support classification and reporting but

do not independently enforce access permissions. A datastore cluster is a storage-management construct and does not provide virtual-machine administrative separation. See the [vSphere Security documentation](#).

QUESTION NO: 5

An architect is designing a cluster for an application that has a virtual machine with a 32 GB memory reservation. The cluster uses vSphere HA admission control based on dedicated failover hosts.

What must the architect ensure when selecting a failover host?

- A. It must have sufficient unreserved resources to satisfy protected virtual-machine reservations
- B. It must be configured as the DRS automation master
- C. It must use the same datastore as the vCenter Server Appliance
- D. It must have vSphere Fault Tolerance enabled

ANSWER: A

Explanation:

The selected failover host must have sufficient unreserved CPU and memory capacity to accommodate the largest protected virtual-machine reservations that might need to restart on it. In this scenario, the host must be able to satisfy the virtual machine's 32 GB memory reservation in addition to any other protected workloads that HA might need to restart according to the configured failover-host design.

A dedicated failover host is reserved for HA recovery and is not normally used for running workloads during normal cluster operation. Selecting a host solely by total installed memory is insufficient if reservations, resource usage, host compatibility, or other constraints prevent the protected virtual machine from being placed there. VMware documents vSphere HA admission-control policies and failover-host behavior in the [vSphere Availability documentation](#).

QUESTION NO: 6

The architect for a large enterprise is tasked with reviewing a proposed design created by a service partner.

Which design elements are expected to be detailed within the physical design section of the documentation?

- A. A design diagram illustrating the configuration and specific attributes, such as IP addresses
- B. A list of requirements, constraints, and risks
- C. A solution architecture diagram with the components and data flow
- D. An entity relationship diagram describing upstream and downstream dependencies for specific service components

ANSWER: A

Explanation:

A design diagram illustrating the configuration and specific attributes, such as IP addresses is correct because a physical design translates the approved logical architecture into an implementable, environment-specific configuration. It documents the concrete placement and connectivity of infrastructure components: physical hosts and clusters, network segments, VLANs, IP addressing, storage connectivity, rack or site placement where applicable, and the detailed configuration values needed by deployment and operations teams. A physical design diagram therefore identifies exactly how the solution is realized in the customer environment, rather than merely showing functional relationships.

For vSphere designs, this level of detail is particularly important because networking and storage implementation choices directly affect availability, performance, and manageability. Documenting attributes such as management, vMotion, vSAN, and virtual-machine network addressing or VLAN associations allows the implementation to be validated against the approved design and operational requirements. VMware design guidance distinguishes the detailed physical realization from

higher-level conceptual and logical design views. See VMware's [vSphere installation and setup documentation](#) and the [VMware Validated Design architecture and design guidance](#).

QUESTION NO: 7

A customer requests a review of its current vSphere platform design.

The following information is noted:

- There are three different workload profiles for the virtual machines:
 - Tier-1 virtual machines operate resource-intensive applications and require dedicated allocations for CPU and RAM.
 - Tier-2 virtual machines operate internet-facing applications and require access to externally facing networks.
 - Tier-3 virtual machines operate platform management tools such as vCenter Server and have different lifecycle management requirements.
- Tier-1, Tier-2 and Tier-3 virtual machines are all hosted on a single large vSphere cluster.
- The Chief Information Security Officer (CISO) has raised concerns that hosting externally facing applications alongside management tools does not meet internal compliance standards.
- The Operations team has raised concerns about Tier-1 virtual machines negatively impacting the performance of vCenter Server.
- The Operations lead has stated that management changes have consistently been rejected by application teams.

As a result of the review, which recommendation should the architect make regarding the design of this platform?

- A. Separate Tier-1, Tier-2 and Tier-3 virtual machines using dedicated distributed virtual switches (DVS)
- B. Separate Tier-2 virtual machines onto a dedicated cluster
- C. Separate Tier-1, Tier-2 and Tier-3 virtual machines onto dedicated clusters
- D. Separate Tier-1, Tier-2 and Tier-3 virtual machines using resource pools and shares

ANSWER: C

Explanation:

Separate Tier-1, Tier-2 and Tier-3 virtual machines onto dedicated clusters is the appropriate recommendation because it creates clear operational, security, and resource boundaries aligned to each workload profile. A dedicated management cluster for the Tier-3 services isolates vCenter Server and other platform-management components from externally exposed workloads, supporting the stated compliance requirement and reducing the management-plane attack surface. It also gives the operations team an independent lifecycle-management domain: maintenance, upgrades, and configuration changes can be planned for management services without requiring agreement from application owners.

A dedicated Tier-1 cluster provides a separate compute and memory contention domain for resource-intensive workloads. This protects the management services from performance effects caused by Tier-1 demand and allows cluster capacity, availability settings, and admission-control decisions to be sized for Tier-1 requirements. A dedicated Tier-2 cluster similarly provides a distinct boundary for internet-facing services and their external network connectivity. VMware recognizes clusters as logical groupings of hosts that provide a common resource and availability domain, making cluster separation the suitable architectural control here. See the [vSphere 7 documentation](#) and VMware's [vSphere Availability technical guidance](#).

QUESTION NO: 8

The Chief Information Security Officer (CISO) for an organization is concerned about the security posture of the operating system images that are used for the provisioning of their Software-as-aService (SaaS) applications. The organization is in a growth period. The organization is opening a new data center to launch its next phase of new SaaS-based solutions.

The DevOps team currently creates encrypted virtual machine (VM) templates that are used for various operating systems and adds these to the vSphere inventory. The DevOps team already uses a published content library and has been granted a role with the ability to add and delete library items.

The following requirements have been noted:

- Impacts to the DevOps team's operational processes must be kept to a minimum. ▪ The DevOps team must be able to regularly check out a copy of the image for updates and check in a new version of the image.
- Images must be synchronized from the primary data center to the new data center.

Which three recommendations should the architect make to design a content library solution that will meet these requirements? (Choose three.)

- A. Clone virtual machines as VM templates to the published content library
- B. Create a subscribed library from the published library and synchronize Open Virtualization Format (OVF) templates on-demand
- C. Create a subscription and publish VM templates to a subscribed content library
- D. Create a subscribed library from the published library and synchronize Open Virtualization Format (OVF) templates automatically
- E. Clone virtual machines as Open Virtualization Format (OVF) templates to the published content library
- F. Update the role for the DevOps team with new privileges

ANSWER: A D F

Explanation:

The correct design is to clone virtual machines as VM templates to the published content library, create a subscribed library with automatic synchronization, and update the DevOps role with the required additional privileges. A vSphere VM template stored in a content library supports the required lifecycle: a user can check out the template to create an editable virtual machine, apply operating-system image updates, and check the VM back in as the updated template. This preserves a workflow that is close to the team's existing template-based process while retaining the managed library copy.

The published library in the primary data center is the source for distribution. A subscribed library at the new data center with automatic synchronization ensures the updated VM templates are copied to the subscriber as library items are published or changed. This is important for VM templates because their synchronization to subscribed libraries requires automatic synchronization rather than on-demand retrieval.

The existing ability to add and delete library items does not by itself authorize the checkout/check-in workflow. The DevOps role must be extended with the applicable content-library VM-template privileges so the team can check out templates, check them in after maintenance, and continue managing the published library. VMware documents content-library management and VM-template workflows in the [vSphere 7 documentation](#) and its [vSphere ESXi and vCenter Server documentation](#).

QUESTION NO: 9

Which two of the listed requirements would be classified as manageability non-functional requirements? (Choose two.)

- A. ESXi clusters must scale when compute resources are sustained above 70% for five business days
- B. vSphere Fault Tolerance must be supported to improve application uptime
- C. ESXi host updates must be installed within one week of release
- D. The vSphere environment must support administrator password rotation
- E. ESXi clusters must scale to 500 concurrent virtual machines

ANSWER: C D

Explanation:

Manageability non-functional requirements define how effectively administrators can operate, maintain, secure, and govern the vSphere environment throughout its lifecycle. “ESXi host updates must be installed within one week of release” is a manageability requirement because it establishes an operational maintenance target for patching hosts. A design must account for the processes and tooling needed to assess updates, place hosts into maintenance mode, remediate them in a controlled sequence, and verify compliance within the required timeframe. VMware vSphere Lifecycle Manager provides the lifecycle-management capabilities used to manage ESXi image and update remediation: [VMware vSphere Lifecycle Manager documentation](#).

“The vSphere environment must support administrator password rotation” is also manageability-focused because credential rotation is an ongoing administrative security operation. The requirement affects how administrator accounts are managed, how access is maintained without operational disruption, and how authentication and authorization controls are governed over time. vCenter Server supports identity and access management through roles, privileges, users, and authentication sources, all of which are central operational design considerations: [vSphere Authentication and Authorization documentation](#).

QUESTION NO: 10

An architect is designing a vSphere Lifecycle Manager image-managed cluster. The customer requires a consistent ESXi image across all hosts and wants remediation to include the supported firmware and driver components supplied by the hardware vendor.

Which two components should the architect include in the desired image? (Choose two.)

- A. An ESXi base image
- B. A vendor add-on
- C. A vSphere Distributed Switch backup
- D. A vCenter Server file-based backup
- E. A virtual-machine storage policy

ANSWER: A B

Explanation:

An ESXi base image is required because it defines the ESXi version for the image-managed cluster. vSphere Lifecycle Manager uses the desired image to evaluate host compliance and remediate hosts toward the configured ESXi release.

A vendor add-on should also be included when supported by the server vendor because it can provide the vendor-specific firmware and driver components that are intended to work with the selected ESXi base image. This supports consistent lifecycle management across the cluster and reduces the risk of unmanaged differences between hosts. VMware documents image-based cluster management and vendor add-ons in the [vSphere Lifecycle Manager documentation](#).

QUESTION NO: 11

During a requirements gathering workshop to design a physical to virtual migration, the customer provides the following information:

- There is no physical firewall in the data center with no anticipated plans for a future network refresh.
- Leveraging the virtual infrastructure to mitigate the lack of network security must be addressed in the design.
- All physical servers to be migrated exist on the same VLAN.

Which recommendation should the architect make to address the customer requirement with regard to virtual networking?

- A. Split the virtual machines into several VLANs Use tag actions
- B. Create port groups with different names and same VLAN IDs Enable traffic shaping for ingress and egress traffic

C. Enable traffic filtering and marking Use allow or drop actions

D. Disable traffic filtering and marking Use tag actions

ANSWER: C

Explanation:

Enable traffic filtering and marking Use allow or drop actions is the appropriate recommendation because a vSphere Distributed Switch can enforce distributed firewall-like policy directly at the virtual networking layer. Traffic filtering rules are applied to virtual-machine network traffic on the host, allowing the design to control communication between migrated workloads even when those workloads originate from, or remain connected to, the same VLAN. Rules can match traffic characteristics such as source and destination IP addresses, protocols, and ports, then explicitly allow or drop matching traffic. This provides workload-level segmentation without requiring an upstream physical firewall or a network refresh.

For this migration, the architect can define an allow-list policy for required application flows and use drop rules to prevent unnecessary east-west communication between server tiers. Because enforcement occurs at the distributed switch port level, policy follows the virtual-machine network connection across ESXi hosts, providing a consistent control point during and after the physical-to-virtual migration. VMware documents traffic filtering and marking as a vSphere Distributed Switch capability and identifies allow and drop as supported rule actions. See the [vSphere Networking documentation](#) and the [VMware guidance for configuring traffic filtering and marking policies](#).