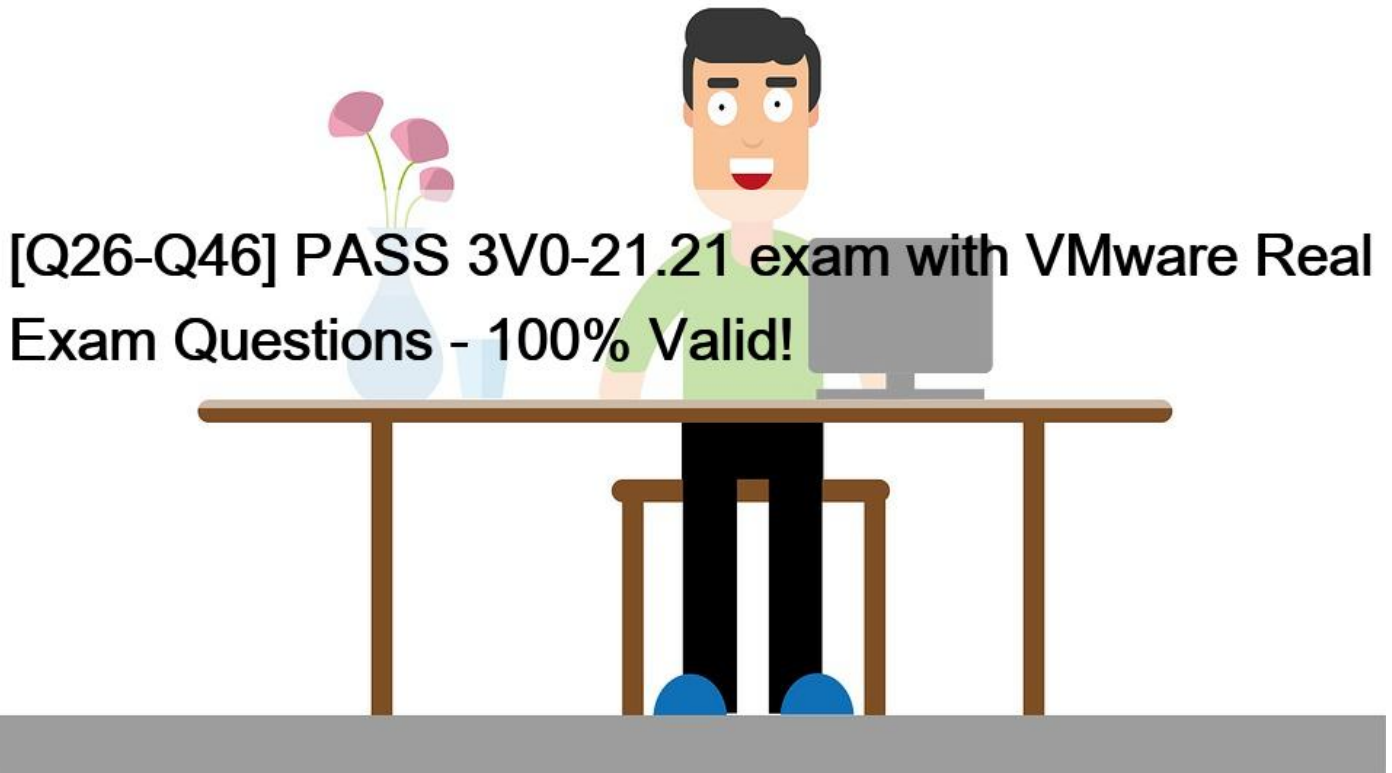


[Q26-Q46 PASS 3V0-21.21 exam with VMware Real Exam Questions - 100% Valid!



PASS 3V0-21.21 exam with VMware Real Exam Questions - 100% Valid!  
Actual 3V0-21.21 Exam Recently Updated Questions with Free Demo

#### NEW QUESTION 26

The storage team at an organization is planning to migrate from an older Fibre Channel storage environment to a new environment using IP-based storage.

Which two switch features or characteristics are appropriate for IP storage networks? (Choose two.)

- \* Fabric extending devices
- \* Spanning Tree Protocol (STP)
- \* 2:1 or greater bandwidth oversubscription for 10 GbE switches
- \* Non-blocking switch
- \* Deep or ultra buffered switches

<https://www.arista.com/en/solutions/ip-storage-network-infrastructures>

<https://www.dell.com/community/s/vjauj58549/attachments/vjauj58549/solutions-ch/488/1/h13104-dedicated-networks-for-ip-storage-wp.pdf>

### NEW QUESTION 27

The storage team at an organization is planning to migrate from an older Fibre Channel storage environment to a new environment using IP-based storage.

Which two switch features or characteristics are appropriate for IP storage networks? (Choose two.)

- \* Fabric extending devices
- \* Spanning Tree Protocol (STP)
- \* 2:1 or greater bandwidth oversubscription for 10 GbE switches
- \* Non-blocking switch
- \* Deep or ultra buffered switches

### NEW QUESTION 28

A architect is designing a new VMware software-defined data center (SDDC) using vSphere 7 to meet the following requirements:

The SDDC must be deployed at two locations: primary and secondary.

vSphere Replication must be used to replicate virtual machines between the two locations.

Site Recovery Manager must be used to orchestrate disaster recovery (DR) activities.

One single-sign on (SSO) domain must be used to authenticate access at both locations.

Which design decision should the architect make to meet these requirements?

- \* A vCenter Server Appliance will be deployed to each site. Unique SSO domains will be created per site.
- \* A vCenter Server will be installed on Windows virtual machines deployed to both sites.
- \* A vCenter Server Appliance will be deployed to each site.
- \* A vCenter Server Appliance will be deployed to the primary site only.

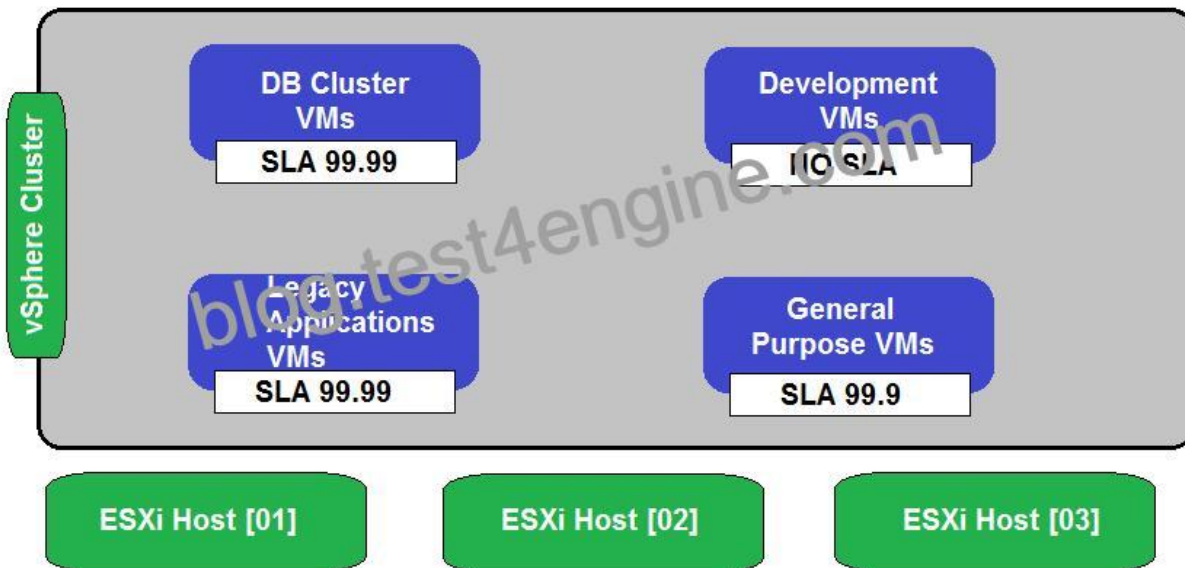
[https://docs.vmware.com/en/Site-Recovery-Manager/8.4/com.vmware.srm.install\\_config.doc/GUID-BB0C03E4-72BE-4C74-96C3-97AC6911B6B8.html](https://docs.vmware.com/en/Site-Recovery-Manager/8.4/com.vmware.srm.install_config.doc/GUID-BB0C03E4-72BE-4C74-96C3-97AC6911B6B8.html)

One single-sign on (SSO) domain must be used to authenticate access at both locations; Install vCenter at Primary site, create SSO Domain. Install vCenter at Secondary site, join to SSO Domain.

### NEW QUESTION 29

Refer to the exhibit.

During a requirements gathering workshop, the customer shares the following diagram regarding their availability service-level agreements (SLAs):



The customer wants database application level availability to always take precedence. What should the architect recommend to meet the customer's requirement?

- \* Enable vSphere HA and add a VM Override with VM Restart Priority set to Highest.
- \* Enable Fault Tolerance.
- \* Enable Sphere HA and maintain the default settings.
- \* Enable vSphere HA and add a VM Override with VM Restart Priority set to Lowest.

### NEW QUESTION 30

An architect is designing a VMware solution for a customer to meet the following requirements:

The solution must use investments in existing storage array that supports both block and file storage.

The solution must support the ability to migrate workloads between hosts within a cluster.

The solution must support resource management priorities.

The solution must support the ability to connect virtual machines directly to LUNs.

The solution should use existing 32G fabric infrastructure.

There is no budget for additional physical hardware.

Which design decision should the architect make to meet these requirements?

- \* The ESXi hosts will leverage Fibre Channel (FC).
- \* The ESXi hosts will leverage iSCSI.
- \* The ESXi hosts will leverage Fibre Channel over Ethernet (FCoE).
- \* The ESXi hosts will leverage NFS.

### NEW QUESTION 31

During a requirements gathering workshop, the customer provides the following requirement that is pertinent to the design of a new vSphere environment:

The Maximum Tolerable Downtime (MTD) for all Tier 1 applications is one hour.

Which requirement classification is being gathered for the design documentation?

- \* Manageability
- \* Performance
- \* Availability
- \* Recoverability

MTD &#8211; Maximum Tolerable Downtime: Sum of the RTO and WRT, which is the total time required to recover from a disaster and start serving the business again. <https://vcdx133.com/2015/01/28/vcdx-availability-explained/>

### NEW QUESTION 32

An architect is considering placement of virtual machines within an existing VMware software-defined data center (SDDC).

During the discovery phase, the following information is documented:

Cluster One

Six ESXi hosts

vSphere HA with host failures cluster tolerates = 1

Proactive HA is enabled and set to automated

Fully Automated vSphere DRS

Transparent Page Sharing (TPS) is enabled

Cluster Two

Eight ESXi hosts

vSphere HA with host failures cluster tolerates = 1

Proactive HA is disabled

Partially Automated vSphere DRS

Transparent Page Sharing (TPS) is disabled

Cluster Three

Three ESXi hosts

vSphere HA with admission control is disabled

Proactive HA is not supported

Transparent Page Sharing (TPS) is disabled

Virtual Machine Resource Profile 1

Memory sharing techniques should not be used

Virtual machines should be automatically restarted in the event of host failure if resources are available Automated initial virtual machine placement Virtual Machine Resource Profile 2 Memory sharing techniques should not be used Virtual machines should be automatically restarted in the event of host failure regardless of available resources Automated initial virtual machine placement Which two recommendations should the architect make for placement of the virtual machines to meet resource profile requirements? (Choose two.)

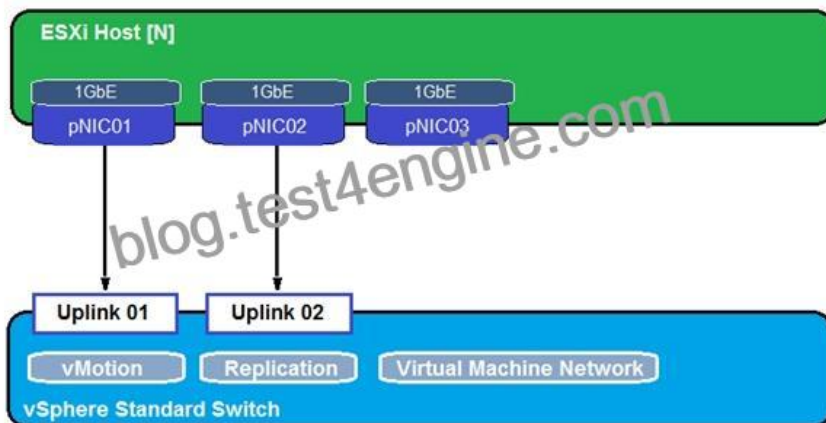
- \* All virtual machines matching Virtual Machine Resource Profile 2 should be placed on Cluster One.
- \* All virtual machines matching Virtual Machine Resource Profile 1 should be placed on Cluster One.
- \* All virtual machines matching Virtual Machine Resource Profile 2 should be placed on Cluster Two.
- \* All virtual machines matching Virtual Machine Resource Profile 1 should be placed on Cluster Two.
- \* All virtual machines matching Virtual Machine Resource Profile 2 should be placed on Cluster Three.

Reference:

<https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.resmgmt.doc/GUID-FEAC3A43-C57E-49A2-8303-B06DBC9054C5.html>

### NEW QUESTION 33

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.)

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

#### NEW QUESTION 34

An architect is tasked with reviewing the design of a VMware software-defined data center (SDDC) for a software development company. The platform is used to developing applications and services. It is important that the customer be able to accurately benchmark performance of developed applications.

The platform has recently commissioned new hosts to update the development cluster.

The development cluster host configuration is:

- \* 4 ESXi hosts with 2 sockets × 16 cores
- \* 512 GB RAM divided evenly between sockets
- \* There is no resource contention

The benchmarking cluster host configuration is:

- \* 8 ESXi hosts with 2 sockets × 8 cores
- \* 256 GB RAM divided evenly between sockets
- \* There is no resource contention

The customer is developing an application that includes a database virtual machine. The application developer states that the database virtual machine performs as required only when allocated 8 vCPUs 256 GB RAM. The database virtual machine performance meets the required levels when run from the development cluster.

Performance benchmarking for the database virtual machine yields highly variable results when run from the benchmarking cluster. The application cannot be released without reliable performance benchmarking data.

What is a possible reason for the difference in performance test results between the development and benchmarking clusters?

- \* The database tier breaches a single NUMA node boundary for the benchmarking cluster
- \* The database tier breaches a single NUMA node boundary for the development cluster
- \* The development cluster can support a lower %Ready time per vCPU
- \* The development cluster has more available RAM per host

#### NEW QUESTION 35

An architect is considering placement of virtual machines within an existing VMware software-defined data center (SDDC).

During the discovery phase, the following information is documented:

Cluster One

- Six ESXi hosts
- vSphere HA with host failures cluster tolerates = 1
- Proactive HA is enabled and set to automated
- Fully Automated vSphere DRS
- Transparent Page Sharing (TPS) is enabled

Cluster Two

- Eight ESXi hosts
- vSphere HA with host failures cluster tolerates = 1
- Proactive HA is disabled
- Partially Automated vSphere DRS
- Transparent Page Sharing (TPS) is disabled

Cluster Three

- Three ESXi hosts
- vSphere HA with admission control is disabled
- Proactive HA is not supported
- Transparent Page Sharing (TPS) is disabled

Virtual Machine Resource Profile 1

- Memory sharing techniques should not be used
- Virtual machines should be automatically restarted in the event of host failure if resources are available
- Automated initial virtual machine placement

Virtual Machine Resource Profile 2

- Memory sharing techniques can be used
- Virtual machines should be protected from any host hardware failures
- Automated initial virtual machine placement

Which two recommendations should the architect make for placement of the virtual machines to meet resource profile requirements?  
(Choose two.)

- \* All virtual machines matching Virtual Machine Resource Profile 2 should be placed on Cluster One.
- \* All virtual machines matching Virtual Machine Resource Profile 1 should be placed on Cluster One.
- \* All virtual machines matching Virtual Machine Resource Profile 2 should be placed on Cluster Two.
- \* All virtual machines matching Virtual Machine Resource Profile 1 should be placed on Cluster Two.
- \* All virtual machines matching Virtual Machine Resource Profile 2 should be placed on Cluster Three.

<https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.resmgmt.doc/GUID-FEAC3A43-C57E-49A2-8303-B06DBC9054C5.html> Profile 2 to Cluster 1 Fully Automated DRS allows Automated Initial VM Placement TPS is enabled to support Memory Sharing requirement (can be used) Profile 1 to Cluster 2 Partially Automated DRS still allows Automated Initial VM placement.

TPS is disabled to support Memory Sharing requirement (cannot be used)

### NEW QUESTION 36

An architect is designing a vSphere environment for a customer and learns that the customer has:

A single vSphere cluster

Two storage arrays with different RAID capabilities

Which two design decisions should the architect make to maximize data availability and data performance for this customer?  
(Choose two.)

- \* Use Storage DRS.
- \* Use VMDK anti-affinity rules.
- \* Use multiple datastores for heartbeat.
- \* Use a minimum of three storage arrays.
- \* Use VM to host DRS rules.

### NEW QUESTION 37

Which of the listed requirements would be classified as a recoverability non-functional requirement?

- \* The platform must be integrated with existing change control policies.
- \* The platform must be able to support a maximum tolerable downtime (MTD) of 30 minutes.
- \* Maintenance windows must be scheduled to take place monthly during an established overnight period.
- \* The platform must be available 24 hours a day, 7 days a week with the exception of scheduled downtime.

### NEW QUESTION 38

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 design diagram illustrating the configuration and specific attributes, such as IP addresses
  - \* A list of requirements, constraints, and risks
  - \* A solution architecture diagram with the components and data flow
  - \* An entity relationship diagram describing upstream and downstream dependencies for specific service components
- &#8220;The physical design is based on the logical design. The physical design includes specific hardware from specific vendors. This design also lists specific configurations for each of the components that are deployed&#8221;

### NEW QUESTION 39

An organization&#8217;s data scientists are executing a plan to use machine learning (ML). They must have access to graphical processing unit (GPU) capabilities to execute their computational models when needed. The solutions architect needs to design a solution to ensure that GPUs can be shared by multiple virtual machines.

Which two solutions should the architect recommend to meet these requirements? (Choose two.)

- \* NVIDIA vGPU
- \* AMD MxGPU
- \* vSphere DirectPath I/O
- \* vSGA
- \* vSphere Bitfusion

Explanation/Reference: <https://blogs.vmware.com/apps/files/2019/08/5521-VMW-GPU-MACHINE-LEARNING-GUIDE-USLET-WEB-20190812.pdf>

### NEW QUESTION 40

Refer to the exhibit.

During a requirements gathering workshop, the architect shares the following diagram:





What should the architect recommend for guaranteed throughput for each service?

- \* Use explicit failover order with pNIC0 as Active for ESXi Management and VM Network Use explicit failover order with pNIC1 as Active for backup network Use explicit failover order with pNIC2 as Active for vMotion Use explicit failover order with pNIC3 as Active for replication
- \* Use the Route Based on IP Hash for ESXi management and VM network Use the Route Based on IP Hash for backup network Use the Route Based on the Originating Virtual Port for vMotion Use failover with pNIC3 as Active for replication
- \* Create a link aggregation group (LAG) for vDS\_01

Use the Route Based on Physical NIC Load for vMotion Use the Route Based on Physical NIC Load for replication

- \* Use the Route Based on IP Hash for ESXi management and VM network Use failover with pNIC1 as Active for backup network Create a link aggregation group (LAG) for vDS\_02

&#8221; The problem is that in A VM network and management is together. In this scenario backup and VM network should be together. From the load point of view, it makes sense, as backup can saturate 100% of NIC but it is not service.&#8221;

#### NEW QUESTION 41

A new real-time financial service application is being developed by the engineering team at a financial firm and will be released as a public Software-as-a-Service (SaaS) offering. The solutions architect has designed and deployed a new vSphere environment and the supporting network infrastructure for hosting all public services. ESXi hosts are configured to use Precision Time Proalhazi@Itoocol (PTP) and a local stratum-1 network time server.

Application provisioning and scaling will be managed by VMware vRealize Automation and can be run on Microsoft Windows or multiple distributions of Linux.

Which three recommendations should the architect include in the design to ensure that the service maintain timekeeping within an accuracy of one second? (Choose three.)

- \* Use Microsoft Windows Server as the guest operating system.
- \* Configure the chrony time-sync agent on each virtual machine guest operating system.
- \* Set the virtual hardware device to use Host System Time (NTP) for each virtual machine running the application.
- \* Add a precision clock virtual device to each virtual machine running the application.
- \* Use a Linux distribution as the guest operating system.
- \* Add a virtual watchdog timer (VWDT) device to each virtual machine running the application.

<https://blogs.vmware.com/apps/2021/04/lets-be-precise-enabling-and-configuring-precision-time-protocol-in-vsphere.html>

<https://core.vmware.com/blog/introducing-ptpvmw-new-linux-driver-achieve-high-time-synchronization-accuracy>

<https://blogs.vmware.com/apps/2020/09/ensuring-accurate-time-keeping-in-virtualized-active-directory-infrastructure.html>

#### NEW QUESTION 42

Which two of the listed requirements would be classified as performance non-functional requirements?

(Choose two.)

- \* The vSphere platform must be able to provide a recovery time objective of 30 minutes
- \* The vSphere platform must be able to provide a minimum throughput of 400 MB/s
- \* The vSphere platform must be able to provide N+1 redundancy
- \* The vSphere platform must be able to provide a maximum read latency of 15 ms
- \* The vSphere platform must be able to provide a service-level agreement (SLA) of 99,9%

Explanation/Reference: <https://technicloud.com/category/vmware/>

#### NEW QUESTION 43

Which of the listed requirements would be classified as a recoverability non-functional requirement?

- \* The platform must be integrated with existing change control policies.
- \* The platform must be able to support a maximum tolerable downtime (MTD) of 30 minutes.
- \* Maintenance windows must be scheduled to take place monthly during an established overnight period.
- \* The platform must be available 24 hours a day, 7 days a week with the exception of scheduled downtime.

#### NEW QUESTION 44

During a requirements gathering workshop, the customer's Chief Information Security Office (CISO) provides the following requirements that are pertinent to the design of a new vSphere environment:

- \* All operating system critical patches must be installed within 24 hours of release.
- \* All virtual machine templates must be updated every three months in line with company policy.

Which requirement classification is being gathered for the design documentation?

- \* Security
- \* Manageability
- \* Recoverability
- \* Availability

Explanation

#### NEW QUESTION 45

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

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

### NEW QUESTION 46

A customer has six hosts available in a cluster. When running at full capacity, all virtual machines can be run on two hosts.

How many hosts can the customer place into maintenance mode at the same time while still providing N+2 resiliency to the cluster?

- \* Two
- \* Three
- \* One
- \* None

**3V0-21.21 Free Sample Questions to Practice One Year Update:**

[https://www.test4engine.com/3V0-21.21\\_exam-latest-braindumps.html](https://www.test4engine.com/3V0-21.21_exam-latest-braindumps.html)