

3V0-21.23 Dumps

VMware vSphere 8.x Advanced Design

<https://www.certleader.com/3V0-21.23-dumps.html>



NEW QUESTION 1

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.)

- A. Use Storage DRS.
- B. Use VMDK anti-affinity rules.
- C. Use multiple datastores for heartbeat.
- D. Use a minimum of three storage arrays.
- E. Use VM to host DRS rules.

Answer: AC

NEW QUESTION 2

An architect is reviewing a physical storage design. The customer has specified that storage DRS will be used for ease of operational management for capacity and performance.

Which recommendation should the architect include in the design?

- A. Create smaller datastores to balance space with Storage DRS
- B. Use a larger number of storage profiles (varied disk speeds and RAID levels) to improve performance
- C. Create larger datastores to balance space with Storage DRS
- D. Create more datastores within each Storage DRS cluster to balance space and performance

Answer: D

NEW QUESTION 3

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

NEW QUESTION 4

An architect is designing the expansion of an existing vSphere 7 environment. The customer is requesting a design for a new cluster to support the anticipated future business growth. The requirements specified for the existing environment design must be considered when designing the new cluster.

The existing design has the following requirements:

- REQ01 The environment has an availability target of 99.5% for all infrastructure.
- REQ02 The recovery time objective (RTO) for Tier 1 virtual machines is one hour.
- REQ03 Windows and Linux virtual machines must reside on separate clusters.
- REQ04 Access to the management cluster within the environment must be controlled. Which of the listed requirements would be classified as a functional requirement?

- A. The environment has an availability target of 99.5% for all infrastructure
- B. The recovery time objective (RTO) for Tier 1 virtual machines is one hour
- C. Access to the management cluster within the environment must be controlled
- D. Windows and Linux virtual machines must reside on separate clusters

Answer: D

NEW QUESTION 5

An architect is designing a new greenfield environment that will install ESXi on local disks. There is a requirement to streamline initial and future installations of ESXi hosts.

Which configuration option should the architect recommend for installing ESXi hosts to meet these requirements?

- A. Installation with kick start script
- B. Auto Deploy with stateless caching mode
- C. Manual installation using boot from SAN
- D. Auto Deploy with stateful install mode

Answer: D

NEW QUESTION 6

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

NEW QUESTION 7

Which design decision must be included in a design to allow for the deployment of a minimum supported configuration of vCenter High Availability (HA)?

- A. A new subnet will be provisioned for vCenter HA services
- B. A vSphere cluster will consist of more than three nodes
- C. The deployed vCenter Server will be Tiny
- D. The vCenter HA network will support a latency of less than 50 ms

Answer: A

Explanation:

<https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.avail.doc/GUID-8FD87389-8CC9-429>

NEW QUESTION 8

An architect is designing a new backup solution for a vSphere platform that has been recently upgraded to vSphere 7. The architect wants the backup solution to perform the following:

- Full virtual machine image backup and restore
- Incremental virtual machine image backup and restore
- File level backup and restore within both Windows and Linux virtual machines
- LAN-free backup

Which functional requirement should the architect include in the design of the new backup solution?

- A. The backup solution must leverage the VMware Consolidated Backup (VCB) framework.
- B. The backup solution must leverage virtual machine snapshots.
- C. The backup solution must leverage VMware vSphere Storage APIs - Data Protection.
- D. The backup solution must leverage VMware vStorage APIs for Data Protection (VADP).

Answer: C

NEW QUESTION 9

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 support the concurrent running of 1,000 virtual machines
- The production environment must be delivered across two geographically dispersed data centers All virtual machines must be capable of running in either data center.
- The two data centers are currently connected to each other through a single but diversely routed, high bandwidth and low latency link.
- The link between the two data centers is capable of supporting a round-trip time (RTT) of 150 ms The existing server hardware standard document states that all virtual infrastructure hosts must be deployed using vSAN ReadyNodes
- The service owner has stated that it is critical to ensure the availability target of 99.9% All virtual machine backups must be completed using the existing backup service
- The recovery time objective (RTO) for the service is five minutes
- The recovery point objective (RPO) of the service is four hours

Which two elements represent risks to the successful delivery of this solution? (Choose two.)

- A. The use of only two data centers
- B. The network connectivity between data center sites
- C. The use of vSAN ReadyNodes
- D. The RTT on the link between the two data centers
- E. The use of the existing backup service

Answer: DE

NEW QUESTION 10

The Chief Operating Officer (COO) at an organization raises concerns that their virtual infrastructure environment is vulnerable. Recently, a security-related issue with a virtual machine caused all management services to become unavailable. No budget is available in the short term for additional platform investment. An architect is asked to review the current environment and make recommendations to mitigate concerns.

A virtualization administrator has provided the following details:

- There is a single four node cluster of ESXi servers
- There are two, Layer 2, physical network switches connecting resources
- The data center network is presented as a single /16 subnet

Given the information provided, which functional requirement should the architect include in the design to mitigate the COOs concerns?

- A. The virtual infrastructure environment must connect application virtual machines and management services to new physical network switches
- B. The virtual infrastructure environment must connect application virtual machines and management services to separate distributed virtual switches (DVS)
- C. The virtual infrastructure environment must connect application virtual machines and management services to separate VLANs
- D. The virtual infrastructure environment must connect management services to a vSphere standard switch (VSS)

Answer: D

NEW QUESTION 10

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 on-premises.
- 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: BE

NEW QUESTION 15

An architect is designing storage for a new vSphere environment to meet the following requirements: ➤ Asynchronous replication is required between two sites.

- The impact on the storage layer should not impact the performance of the compute layer.
- Each application tier will require different replication attributes.
- Virtual machine live migration across compute and storage must be supported.
- Virtual machine aware back up will be leveraged.
- Operational management overhead should be minimized.
- Operational automation should be supported.

Which storage design recommendations would meet the requirements?

- A. Two new Fibre Channel storage arrays will be deployed, one at either site. Each application tier will be initially provisioned a new LU
- B. Data replication will be offloaded to the new arrays.
- C. Two new vSphere clusters enabled with vSAN will be deployed, one at either site
- D. vSAN will be used to provide policy-based management for each application tier. vSphere Replication will be used to replicate the virtual machine data in an asynchronous configuration.
- E. Two new iSCSI storage arrays will be deployed, one at either site. Each application tier will be initially provisioned a new LU
- F. Data replication will be offloaded to the new arrays.
- G. Two new storage arrays will be deployed, one at either site. vSphere Volumes (vVOLS) will be used to provide policy-based management for each application tier
- H. Data replication will be offloaded to the new arrays.

Answer: B

NEW QUESTION 16

What is a benefit of using a scale-out method for handling vSphere cluster growth?

- A. An increase in the recovery time objective (RTO) for the cluster
- B. Faster to reach the limit of virtual machines per host
- C. An overall reduction in the license costs for the cluster
- D. Less potential impact to virtual machines during a single host failure

Answer: B

NEW QUESTION 20

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.)

- A. All virtual machines matching Virtual Machine Resource Profile 2 should be placed on Cluster One.
- B. All virtual machines matching Virtual Machine Resource Profile 1 should be placed on Cluster One.
- C. All virtual machines matching Virtual Machine Resource Profile 2 should be placed on Cluster Two.
- D. All virtual machines matching Virtual Machine Resource Profile 1 should be placed on Cluster Two.
- E. All virtual machines matching Virtual Machine Resource Profile 2 should be placed on Cluster Three.

Answer: BD

NEW QUESTION 25

An architect is finalizing the design for a new vSphere platform based on the following information:

- All Windows virtual machines will be hosted on a dedicated cluster for licensing purposes.
- All Linux virtual machines will be hosted on a dedicated cluster for licensing purposes. All management virtual machines will be hosted on a dedicated cluster.
- A total of ten physical sites will be used to host virtual machines.
- In the event of one physical datacenter becoming unavailable, the manageability of the virtual infrastructure in the remaining data centers should not be impacted.
- Access to configure the management virtual machines via vCenter Server must be controlled through the management Active Directory domain.
- Access to configure the Windows and Linux virtual machines must be controlled through the resource Active Directory domain.
- The management and resource Active Directory domains are part of separate Active Directory forests and do not have any trusts between them.
- The design will use Active Directory with Integrated Windows Authentication.

How should the architect document the vCenter Server configuration for this design?

- A. Deploy a vCenter server for the management cluster. Deploy a vCenter Server for all remaining cluster
- B. Create a shared SSO domain for each physical site.
- C. Deploy a vCenter Server for the management cluster. Deploy a vCenter Server for all remaining cluster
- D. Create a shared SSO domain across all physical sites.
- E. Deploy a vCenter Server for the management cluster with a dedicated SSO domain. Deploy a vCenter Server for all remaining clusters and use a dedicated SSO domain for each physical site.
- F. Deploy a vCenter Server for the management cluster with a dedicated SSO domain. Deploy a vCenter Server for all remaining clusters and use a dedicated SSO domain into a single physical site.

Answer: B

NEW QUESTION 27

An architect is designing a solution for an environment with two types of resource profiles that must be virtualized. The first type consists of Tier 1 virtual machines that are disk I/O intensive, but do NOT require high CPU or memory. The second type consists of Tier 2 virtual machines that require a lower CPU and memory allocation and have minimal disk I/O.

Which design recommendation should the architect make for distributing the resource profiles?

- A. Separate the two resource profiles into two cluster
- B. The Tier 1 cluster will have fast storage while the Tier 2 cluster will not.
- C. Run both resource profiles on the same cluster with the same host hardware platform.
- D. Separate the two resource profiles into two cluster
- E. The Tier 2 cluster will have faster CPU and more memory while the Tier 1 cluster will have slower CPU and less memory but more disk space.
- F. Run both resource profiles on the same cluster with host hardware that has fast CPU, large amounts of memory, and the fastest storage platform.

Answer: D

NEW QUESTION 28

An architect is designing a solution based on the following information:

- Each ESXi host has a single physical NIC with two 10 Gbps ports.

- There is a performance-based service-level agreement (SLA) that guarantees 15 Gbps bandwidth for production virtual machines at all times.
- There is no budget to purchase additional hardware.
- The hardware replacement SLA is based on a delivery agreement of two business days.

Which recommendation for the configuration of vSphere High Availability (HA) should the architect include in the design?

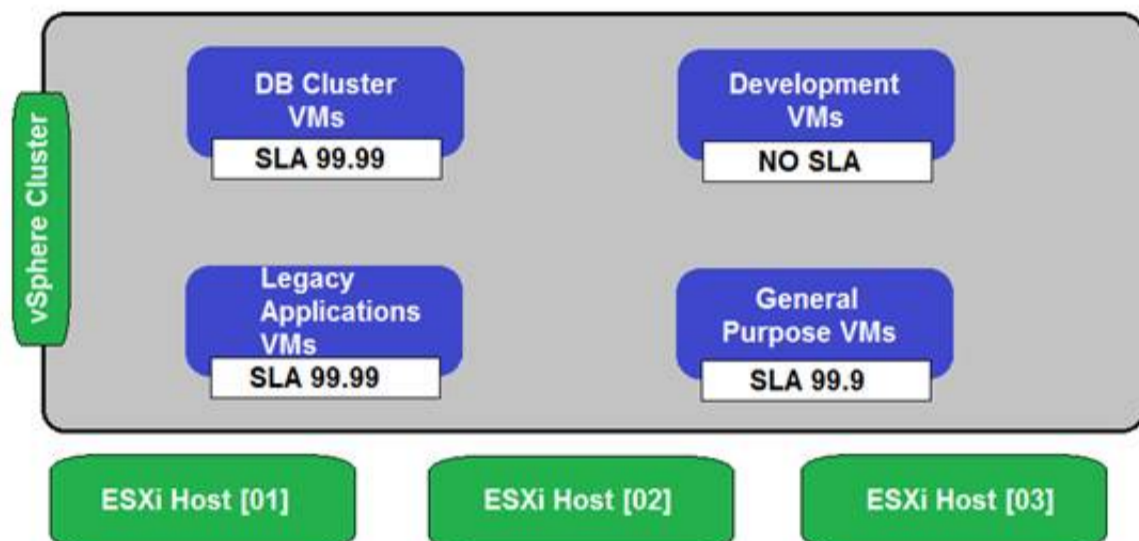
- A. Configure vSphere HAConfigure % based admission control Configure two isolation addresses Consider an OEM with NIC failure conditions in their Proactive HA plugin
- B. Configure vSphere HASet das.IgnoreRedundantNetWarning to trueConsider an OEM with NIC failure conditions in their Proactive HA plugin
- C. Configure vSphere HAConfigure two existing data stores for heartbeatConsider an OEM with NIC failure conditions in their Proactive HA plugin
- D. Configure Proactive HA Automation Level: Automated Remediation: Maintenance mode for all failuresConsider an OEM with NIC failure conditions in their Proactive HA plugin

Answer: A

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?

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

Answer: D

NEW QUESTION 30

An architect is designing a new vSphere cluster. The requirement is to provide a total of 96 CPU cores and 1.5 TB RAM across all hosts.

The following information has been provided:

Two different physical hardware profiles are available for the ESXi hosts in the cluster.

-Profile 1: 16 CPU cores and 256 GB RAM

-Profile 2: 32 CPU cores and 512 GB RAM

Profile 2 is twice as expensive to purchase as Profile 1.

Which two aspects should the architect consider when selecting the hardware profile? (Choose two.)

- A. The manufacturer and model of the CPUs in the hosts
- B. The amount of capacity available for failover of virtual machines within the cluster
- C. The downtime allowed for virtual machines that will be running within the cluster
- D. The cost to procure and maintain the hardware
- E. The number of virtual machines that will be running within the cluster

Answer: BE

NEW QUESTION 32

A customer defines a requirement to minimize the vMotion migration time during a maintenance period. The servers being used are equipped with eight 1 GbE network adapters.

Per the defined logical network configuration, there are two network adapters each used for:

- Management traffic
- vMotion traffic
- iSCSI traffic
- Virtual machine traffic

Which design decision should the architect make to meet the customer requirement?

- A. Use Network I/O Control to define a reservation for vMotion traffic.
- B. Implement Multi-NIC vMotion by adding additional vMotion VMkernels.
- C. Configure a dedicated TCP/IP stack for vMotion traffic.
- D. Combine vMotion and Management traffic to make use of four adapters.

Answer: A

NEW QUESTION 37

An architect is designing a vSphere environment for a customer based on the following information:

- The vSphere cluster will have three hosts only due to budget considerations.
- A database cluster (node majority) consisting of three virtual machines will be running on the vSphere cluster.

Which two recommendations can the architect make so that the customer achieves the highest level of application availability while taking into consideration operational resiliency? (Choose two.)

- A. Create VM-VM anti-affinity rules
- B. Set `das.respectvmvmantiaffinityrules` to false
- C. Create VM-Host anti-affinity rules
- D. Disable vSphere HA during maintenance
- E. Set `das.ignoreinsufficienthbdastore` to true

Answer: BC

NEW QUESTION 38

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 Protocol (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.)

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

Answer: ABC

NEW QUESTION 39

An architect decides to separate virtual desktops and application servers into separate vSphere clusters to meet security and management requirements.

What are two implications of this design decision? (Choose two.)

- A. There will be an increase in management overhead.
- B. Identical hardware must be procured for all hosts.
- C. There will be a reduction in performance.
- D. The patching cycles will affect both clusters at the same time.
- E. There will be additional licensing and cost requirements for both clusters.

Answer: DE

NEW QUESTION 44

An architect is reviewing a physical storage design. The customer has specified that a new active-passive based storage array will be used to provide storage for the vSphere clusters.

Which configuration should for the architect recommended?

- A. `VMW_SATP_LOCAL`
- B. `VMW_PSP_MRU`
- C. `VMW_SATP_DEFAULT_AA`
- D. `VMW_PSP_FIXED`

Answer: B

NEW QUESTION 47

An architect is tasked with designing a new VMware software-defined data center (SDDC) solution for an online retail customer who has a primary and secondary data center as well as 10 distribution hubs.

The customer has provided the following business requirements to help inform the design:

- The solution must support the running of up to 1,000 concurrent virtual machines across the primary and secondary data center.
- The solution must support the running of up to 20 concurrent virtual machines in each distribution hub.
- The solution must support the separation of management and lines-of-business application virtual machines.
- All management components (including directory services, backup, automation, operations and logging) must be deployed to the primary data center.
- All virtual infrastructure components must have redundancy of N+1.
- The solution should support a monthly uptime target of 99.9%.
- The recovery time objective (RTO) for the solution must be four hours.
- The recovery point objective (RPO) for the solution must be 24 hours.

Given the information from the customer, which assumption should the architect include in the design?

- A. All business application virtual machines can be deployed into a single cluster within the primary data center.
- B. Each distribution hub should be configured with a backup device.
- C. The wide area network has sufficient bandwidth to support centralized management.
- D. Each cluster will have a minimum of four hosts.

Answer: B

NEW QUESTION 50

A new vSphere platform is being created. The platform will host virtual machines that will run management services and line-of-business applications. What should the architect consider when designing the number and type of clusters required?

- A. Maximum tolerable downtime
- B. Predicted platform growth
- C. Auditing requirements for the virtual machines
- D. The level of isolation required between virtual machine classifications

Answer: D

NEW QUESTION 53

An architect is designing a VMware software-defined data center (SDDC) solution based on the following customer requirements:

- > The solution must initially support 1,000 virtual machines
- > The solution must scale to support the concurrent running of up to 5,000 virtual machines
- > The production environment should be delivered across two data centers
- > The solution should have a maximum tolerable downtime (MTD) of four hours
- > The solution should have a monthly service availability target of 99.8%

Which two assumptions could the architect make based on the information from the customer to help size the solution? (Choose two.)

- A. The number of vSphere hosts in a cluster
- B. The average resource utilization of a virtual machine
- C. The size (CPU/RAM/storage) of the average virtual machine
- D. The guest operating system for each virtual machine
- E. The size (CPU/RAM/storage) of the vSphere hosts

Answer: AE

NEW QUESTION 55

An architect is designing a new vSphere environment with the following resources:

- > 600 vCPU
- > 5,760 GB RAM

Average resource usage is:

- > 60 vCPU
- > 1,152 GB RAM

The design must meet the following requirements:

- > The environment has the ability to burst by 25%.
- > Each host can schedule 36 vCPUs and has 512 GB RAM.
- > Management overhead is 20%.

What is the minimum number of hosts required to meet the design requirements?

- A. Three
- B. Five
- C. Four
- D. Two

Answer: D

NEW QUESTION 56

An architect is planning the physical server configuration for a vSAN-based infrastructure.

Which operations mode should a RAID controller support to minimize potential server downtime during physical disk failures?

- A. RAID controller with Passthru mode
- B. RAID controller with RAID 5 mode
- C. RAID controller with RAID 10 mode
- D. RAID controller with RAID 6 mode

Answer: D

NEW QUESTION 60

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your 3V0-21.23 Exam with Our Prep Materials Via below:

<https://www.certleader.com/3V0-21.23-dumps.html>