

3V0-21.23 Dumps

VMware vSphere 8.x Advanced Design

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



NEW QUESTION 1

Which requirement would be classified as a functional requirement within the design documentation?

- A. The system must perform virtual machine backups through an API.
- B. Virtual machines must be patched within one month of the patch release date.
- C. Virtual machines must be restarted within 30 minutes of a host failure.
- D. The system must be able to scale to support 500 concurrent virtual machines.

Answer: C

NEW QUESTION 2

An architect is designing a new vSphere platform for a customer to meet the following requirements:

- The platform must be deployed into five physically separate sites.
- The sites are spread across multiple regions.
- Some sites require more than one vCenter Server.
- The platform must provide an administrator with the ability to access virtual infrastructure components across all sites from a single management tool instance.

Which single sign-on (SSO) design recommendation will meet these requirements?

- A. Use an SSO domain across all vCenter Server instances
- B. Use an SSO domain per region
- C. Use an SSO domain per vCenter Server instance
- D. Use an SSO domain per site

Answer: C

NEW QUESTION 3

During a transformation project kick-off meeting, an architect highlights specific areas on which to focus while developing the new conceptual design.

Which two of the listed statements are business requirements? (Choose two.)

- A. The project should use the existing storage devices within the data center
- B. Sites must support a network latency of less than 12 ms round-trip time (RTT)
- C. The solution must allow data replication between sites
- D. There is no budget specifically assigned for disaster recovery
- E. There must not be a single point of failure for the virtual infrastructure

Answer: CE

NEW QUESTION 4

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 5

An architect is designing a new vSphere platform to meet a list of requirements from the security team. Which two requirements would be classified as non-functional requirements? (Choose two.)

- A. Migration of virtual machines between hosts must be encrypted
- B. Log information must be verbose to support incident resolution
- C. Critical events generated within the platform must be logged to an external Syslog service
- D. Data integrity must be ensured
- E. A common content library must be maintained across all data centers

Answer: CD

NEW QUESTION 6

An architect is designing an environment for a retail customer. The design will use a single small vCenter Server Appliance and a cluster of eight ESXi hosts at a remote site. There is a single 10 GbE connected network at the remote site to support all management services. It is not possible to create additional management networks at the remote site. Virtual machine backups at the site will be dependent on the vCenter Server being available.

Which design decision should the architect make to maximize availability for backups?

- A. vCenter Server High Availability will be configured.
- B. The vCenter Server Appliance will be protected with vSphere Fault Tolerance.
- C. The cluster will be configured to use vSphere DRS in fully automated mode.
- D. The cluster will be configured with vSphere HA and set to restart virtual machines based on guest operating system heartbeat monitoring.

Answer: C

NEW QUESTION 7

In a meeting to discuss the minimum viable product (MVP) deployment of a new customer-facing application, the key stakeholder shares details of the application components and the application administrators share details of performance and integrity tests for the application.

The application will be made up of the following components:

➤ A web server

-Steps to confirm the web server is operating correctly will take 15 minutes after the application server is online.

➤ An application server

-Steps to confirm application server integrity will take 15 minutes after the database is online.

➤ A database server

-The database server will be managed by a database administrator, with an agreed service-level agreement (SLA) to restore and validate database services within one hour.

The existing VMware infrastructure offers a recovery point objective (RPO) of 5 minutes and recovery time objective (RTO) of 15 minutes through a combination of backups and replication.

In the event of an outage impacting all three application components, how long will it take for the application to recover and complete all checks?

- A. 15 minutes
- B. 60 minutes
- C. 105 minutes
- D. 90 minutes

Answer: C

Explanation:

15 restore VMs + 60 restore and test DB + 15 test app server + 15 test web server

NEW QUESTION 8

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 9

An architect is preparing a design for a company planning digital transformation. During the requirements gathering workshop, the following requirements (REQ) and constraints (CON) are identified:

- REQ01 The platform must host different types of workloads including applications that must be compliant with internal security standard.
- REQ02 The infrastructure must initially run 100 virtual machines.
- REQ03 Ten of the virtual machines must be compliant with internal security standard.
- CON01 The customer has already purchased the licenses as part of another project.
- CON02 The customer has five physical servers that must be reused.

Additionally, based on resource requirements, four physical servers will be enough to run all workloads. Which recommendation should the architect make to meet requirements while minimizing project costs?

- A. Use Network I/O Control to ensure the internal security zone has higher share value
- B. Purchase additional servers and plan separate, isolated clusters for workloads that must be compliant with internal security
- C. Use a single cluster and ensure that different security zones are separated at least with dedicated VLANs and firewall
- D. Use a single cluster and configure DRS anti-affinity rules to ensure internal security compliant virtual machines cannot migrate between ESXi hosts.

Answer: C

NEW QUESTION 10

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?

- A. Manageability
- B. Performance
- C. Availability
- D. Recoverability

Answer: C

NEW QUESTION 10

An architect is tasked with planning the design of a new vSphere environment. When commissioned, this environment will be used to migrate an existing set of virtual machines.

An inventory of the existing infrastructure, including configured vCPU, RAM and storage sizes has been provided.

In order for each virtual machine to be migrated, which two data sources with peak and average utilization data are required for sizing? (Choose two.)

- A. %Ready
- B. Disk Write latency
- C. CPU
- D. Ballooned memory
- E. IOPS

Answer: BE

NEW QUESTION 12

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 15

An organization'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.)

- A. NVIDIA vGPU
- B. AMD MxGPU
- C. vSphere DirectPath I/O
- D. vSGA
- E. vSphere Bitfusion

Answer: AE

Explanation:

<https://blogs.vmware.com/apps/2018/07/using-gpus-with-virtual-machines-on-vsphere-part-1-overview.html>

NEW QUESTION 16

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 17

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: AC

NEW QUESTION 21

As part of a new hybrid cloud initiative for a large financial company, the customer technical team is presenting an overview of the current state of the infrastructure and their vision for a new solution.

The project team captures notes during the presentation and adds them to the discovery documentation. Which of the listed statements is a design constraint?

- A. The applications are created in-house with in-guest recovery protection
- B. The maximum tolerable data loss is 10 minutes
- C. The two data center locations have a network latency of 8 ms round-trip time (RTT)
- D. The existing storage is out of maintenance

Answer: D

NEW QUESTION 25

An architect is tasked with expanding an existing VMware software-defined data center (SDDC) solution so that it can be used to deliver a virtual desktop infrastructure (VDI) service off-shore development activities.

The production environment is currently delivered across two geographically dispersed data centers. The two data centers are currently connected to each other through multiple diversely routed, high bandwidth and low latency links. The current operations management components are deployed to a dedicated management cluster that is configured with N+1 redundancy. The current VMware software-defined data center (SDDC) has a monthly availability target of 99.5%, which includes all management components.

The customer requires that the new solution scale to support the concurrent running of 500 persistent virtual desktops. The virtual desktops must not share the same virtual infrastructure as existing virtual machines, but can be managed using the same VMware operations management components. Any new VDI service management components must be installed into the management cluster. There is no requirement to back up the virtual desktops because all relevant user data is stored centrally. The VDI service is providing business critical services and must have an availability target of 99.9%.

Given the information from the customer, which two assumptions would the architect include in the design? (Choose two.)

- A. The existing virtual infrastructure has sufficient capacity to host the new VDI workloads
- B. The existing operations monitoring tools have sufficient capacity to monitor the new VDI services
- C. The existing management cluster has enough available capacity to host any VDI service management component
- D. The management cluster has N+1 redundancy
- E. The VDI service has a higher service-level agreement (SLA) than the operations management SLA

Answer: BD

NEW QUESTION 26

An architect is tasked with designing a new VMware software-defined data center (SDDC) using VMware vSAN. The architect uses a storage assessment tool to determine the storage requirements for the new vSAN cluster. The new SDDC is going to be deployed into the existing data center and must be connected to a shared core network switch.

The architect decides to use vSAN ReadyNodes with the following configuration:

- Two disk groups with:
- Write Intensive NVMe 800 GB drive for cache
- Four 3.84 TB Mixed Use NVMe for capacity
- Four 10 GbE ports

Which element represents a risk that should be included in this design?

- A. The number of 10 GbE capable ports in the vSAN ReadyNode
- B. The use of vSAN ReadyNodes
- C. The existing network is 10 GbE capable
- D. The use of NVMe drives for cache and capacity

Answer: C

NEW QUESTION 30

A VMware Service Provider is tasked with delivering a solution for continuous availability for a subset of Tier 1 virtual machines (VMs) and vApps running in their vSAN environment. The VMs make up a

mission-critical application and there can be no data loss in the event of an outage at their primary data center. In the event of a regional outage, they have established a 10-minute recovery point objective (RPO). Failover/failback to the third site must be automated.

They have the following in place:

- Two local data centers (primary and secondary) connected with 100 Gb dedicated fiber
- 2ms round-trip time (RTT) latency between the sites A third data center located on another power grid
- 70ms latency between the primary and secondary data centers
- Matching storage arrays at all locations

Which two solutions could be used to meet the requirements? (Choose two.)

- A. Site Recovery Manager
- B. Snapshots
- C. vSAN Metro Cluster

- D. vSphere Data Protection
- E. vStorage APIs for Array Integration (VAAI)

Answer: BC

NEW QUESTION 31

.....

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>