

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

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

- 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 IPv4 based network infrastructure.
- There is no budget for additional physical hardware.

Which two design decisions could the architect make to meet these requirements? (Choose two.)

- A. The ESXi hosts will leverage Fibre Channel (FC)
- B. The ESXi hosts will leverage NFS 3
- C. The ESXi hosts will leverage Fibre Channel over Ethernet (FCoE)
- D. The ESXi hosts will leverage iSCSI
- E. The ESXi hosts will leverage NFS 4.1

Answer: BD

Explanation:

<https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.storage.doc/GUID-8A929FE4-1207-4C> Starting from vSphere 7.0, VMware no longer supports software FCoE in production environments.

<https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.storage.doc/GUID-6B49866F-7005-40>

NEW QUESTION 4

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 5

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 6

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 7

An architect is designing a new vSphere environment to meet the following requirements:

- The environment must support 5,000 virtual machines.
- The environment will be built initially using 350 hosts.

Which vCenter Server appliance deployment size should the architect specify for the design?

- A. Large
- B. Small
- C. Tiny
- D. Medium

Answer: A

Explanation:

<https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vcenter.install.doc/GUID-077C7523-E0EA-492>

NEW QUESTION 8

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 9

Following a company merger, there are two data centers running vSphere environments. Both data centers are leveraging separate Layer 3 vMotion networks. Which requirement must be met in order to enable vMotion migration between these locations?

- A. The vMotion service must be configured on the Management VMkernel adapter
- B. A dedicated TCP/IP stack for vMotion with a dedicated gateway must be configured
- C. A stretched vMotion network must be configured between data centers
- D. Virtual machines must be powered off in order to migrate them between data centers

Answer: B

NEW QUESTION 10

An architect is designing a new greenfield environment with 600 ESXi hosts in an automated fashion. The engineering department already has a PXE Boot server, TFTP server, and DHCP server set up with an NFS mount for their current Linux servers.

The architect must be able to demonstrate and meet a security requirement to have all infrastructure processes separated.

Which recommendation should the architect make for the ESXi host deployment?

- A. Request an isolated network segment to use and dedicate it to Auto Deploy functions
- B. Ask the business to expand the engineering environment to service the virtual environment as well
- C. Request a common shared network with flexible security measures to accommodate different auto deployment options
- D. Deploy each ESXi host individually and document it to satisfy security requirements

Answer: A

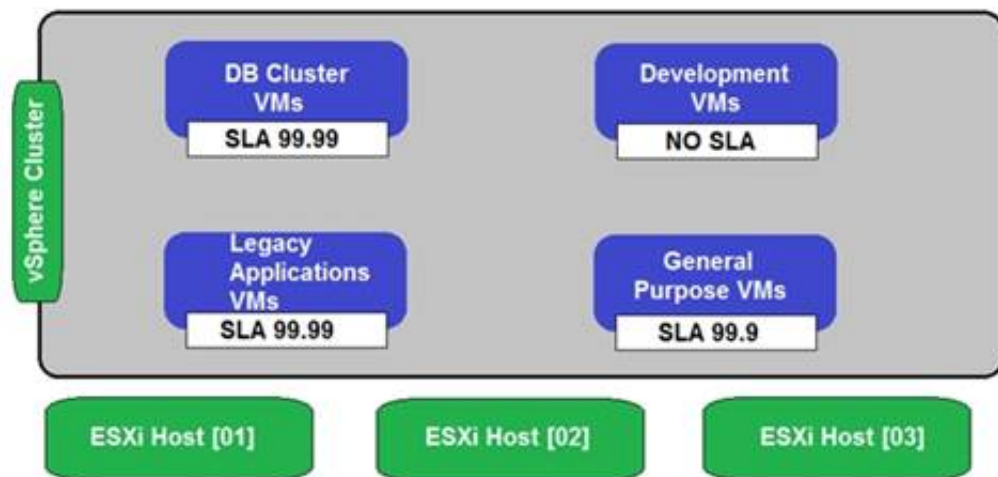
Explanation:

<https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.esxi.install.doc/GUID-8DAC6FEE-0441-4072>

NEW QUESTION 10

Refer to the exhibit.

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



The customer states that there is no application level availability for legacy applications.

Which recommendation could the architect make to meet the customer's high availability requirements for the legacy applications virtual machines?

- A. Enable vSphere HA and add a VM Override with VM Restart Priority set to Disabled
- B. Enable Fault Tolerance
- C. Achieve application availability with snapshots
- D. Enable vSphere HA and add a VM Override with VM Restart Priority set to Lowest

Answer: D

NEW QUESTION 12

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 13

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 15

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 18

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 22

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 26

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 29

During a requirements gathering workshop, the customer provides the following information:

- Each host has 2 × 10 GbE NIC
- EtherChannel is not currently configured
- No changes can be made to the physical network
- Network throughput must be prioritized for defined critical services

Which two recommendations should the architect make with regard to virtual networking? (Choose two.)

- A. Use Route Based on Physical NIC Load.
- B. Use Network I/O Control with Shares.
- C. Use Network I/O Control with Reservation.
- D. Use Link Aggregation Control Protocol (LACP).
- E. Use Network I/O Control with Limits.

Answer: AD

NEW QUESTION 33

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 36

An architect is designing a new VMware software-defined data center (SDDC) that will consist of 100 branch sites connected to a single VMware vCenter Server within the primary data center. To allow for the use of existing automation scripts, there is a requirement to replicate the names of the virtual distributed port groups across all sites. The procurement team purchases licensing and there is no further budget allocated.

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

- A. A new vCenter Server will be deployed for each branch site
- B. A new host and cluster folder will be created for each branch site
- C. The automation script will be updated to reflect unique naming for each site
- D. A new virtual data center will be created for each branch site

Answer: B

NEW QUESTION 41

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 44

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

- A. The application must be hosted with redundancy levels of N+1 or better.
- B. Penetration testing must be executed quarterly with a pass rate of 80% or higher.
- C. The application must be capable of handling 200 transactions per second.
- D. Administrators must monitor the network traffic of the desired systems.

Answer: C

NEW QUESTION 46

An architect is creating a network design for a new vSphere environment.

Based on customer requirements, the environment must support the following types of traffic:

- Management

- > vMotion
- > vSAN
- > Fault Tolerance
- > Virtual machine traffic, which cannot be impacted by other types of traffic

Which design recommendation can the architect make for a resilient infrastructure with vSphere network service tiering?

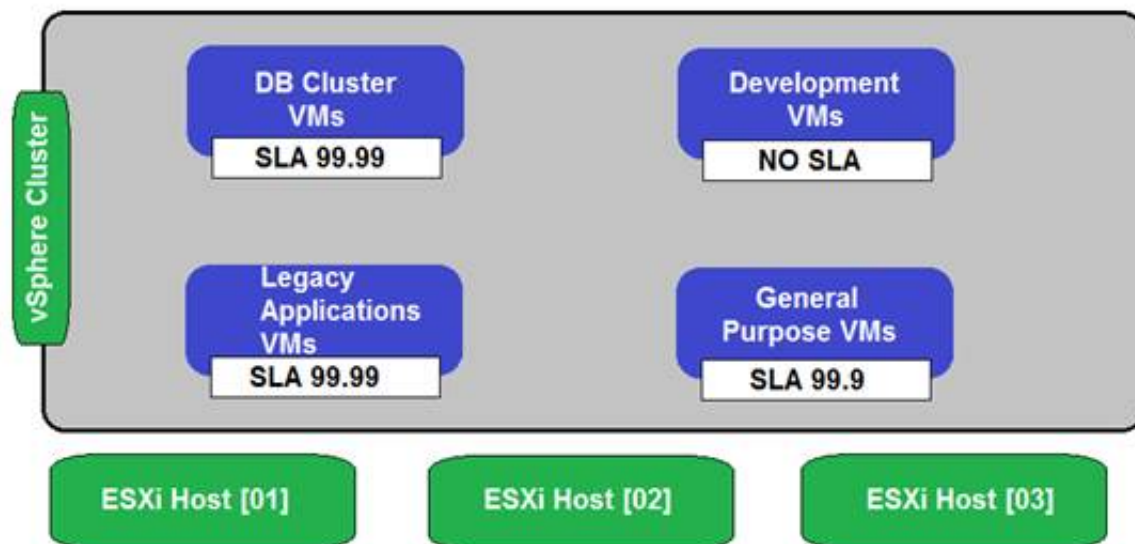
- A. Use different logical networks to ensure traffic is isolated with separate VLANs
- B. Use Network I/O Control and ensure appropriate share value is defined for different types of traffic giving priority to the virtual machines traffic
- C. Use two dedicated virtual switches with a single adapter each, dedicating one virtual switch for Management, vMotion, vSAN and Fault Tolerance traffic, and the second one for virtual machine traffic
- D. Use a NIC teaming policy based on the physical NIC load

Answer: A

NEW QUESTION 49

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 50

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?

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

Answer: C

NEW QUESTION 55

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 56

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 60

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: A

NEW QUESTION 63

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 66

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 70

During a requirements gathering workshop, the customer provides the following requirement (REQ) and constraints (CON):

- > REQ01: The customer is looking for a way to limit database virtual machine (VM) placement to save on CPU licensing costs.
- > CON01: There is a single cluster with no budget to scale.
- > CON02: All virtual machines must run on the consolidated cluster.

Which two design decisions should the architect make to meet the customer requirement? (Choose two.)

- A. The solution must use VM-VM anti-affinity rules
- B. The solution must use vSphere DRS in manual mode
- C. The solution must use a vRealize Orchestrator workflow for VM placement
- D. The solution must use VM-Host affinity rules
- E. The solution must use vSphere VM and host DRS groups

Answer: DE

NEW QUESTION 72

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 75

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?

- A. Two
- B. Three
- C. One
- D. None

Answer: A

NEW QUESTION 80

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

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

Answer: BD

NEW QUESTION 85

A customer requires the use of data encryption to ensure data is not accessible when a drive is removed from the primary storage platform. However, there is also a requirement to use deduplication and compression against all workloads in order to conserve space.

Which solution meets the customer requirements?

- A. Data-in-transit encryption
- B. OS-level encryption
- C. Encrypted backups
- D. Array-based encryption

Answer: D

NEW QUESTION 88

Application owners require support of a Microsoft Windows Server Failover Cluster (WSFC).

Their current environment consists of the following components:

- vSphere 7.0 and vSAN 7.0
- External array supporting NFS 3.0/4.1, Server Message Block (SMB) 2.1
- 10 GbE storage connectivity for all devices

The solution architect is tasked with coming up with a solution to meet this requirement while utilizing their existing investments.

Which two recommendations could the architect make? (Choose two.)

- A. Use vSAN native support for WSFC
- B. Use NFS 4.1 shares for quorum and shared disk
- C. Use raw device mapping (RDM)
- D. Use the SMB 2.1 protocol for sharing disks
- E. Run WSFC on vSAN iSCSI Target Service

Answer: AE

Explanation:

<https://blogs.vmware.com/virtualblocks/2018/04/18/vsan-6-7-introducing-wsfc-support-vsan>

NEW QUESTION 89

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 92

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 97

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 100

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 102

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 104

.....

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>