

Specialist - Systems Administrator, VxRail Appliance Exam

EMC DES-6332

Version Demo

Total Demo Questions: 17

Total Premium Questions: 172

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, VxRail Appliance Concepts	64
Topic 2, VxRail Appliance Installation	19
Topic 3, VxRail Appliance Administration	77
Topic 4, VxRail Appliance Troubleshooting	12
Total	172

QUESTION NO: 1

A virtual machine running on a VxRail cluster must be migrated to a different vSAN datastore while it remains powered on. The ESXi host does not need to change.

Which VMware feature should be used?

- A. vMotion
- B. Storage vMotion
- C. vSphere High Availability
- D. Distributed Resource Scheduler

ANSWER: B

Explanation:

Storage vMotion is correct. Storage vMotion migrates the files of a powered-on virtual machine from one datastore to another while the VM continues to run on its current ESXi host. The operation relocates virtual disks, configuration files, and other VM files according to the selected migration options.

vMotion, by contrast, moves the running compute state of a VM between ESXi hosts. A combined migration can move both compute and storage when required, but Storage vMotion is the appropriate feature when only the datastore location must change. See the [vSphere Virtual Machine Migration documentation](#).

QUESTION NO: 2

The vSphere Client is showing several failures on node 3. Support recommends opening an SSH session to troubleshoot. The connection fails to establish.

Which path should be used to enable SSH on the host from vSphere Client?

- A. Host > Monitor > Services > Enable SSH
- B. Cluster > Configure > VxRail > Hosts > Enable SSH
- C. Host > Configure > System > Services
- D. Cluster > Monitor > VxRail > Hosts > Services

ANSWER: C

Explanation:

Host > Configure > System > Services is the correct vSphere Client navigation path because SSH access is controlled through the ESXi host's service configuration. After selecting the affected host, opening its Configure tab, and navigating to System > Services, an administrator can locate the *TSM-SSH* service and start it. The service's startup policy can also be adjusted there when appropriate for the troubleshooting requirement. Once TSM-SSH is running, the host can accept an SSH connection, assuming network reachability and any relevant firewall controls permit TCP port 22.

This host-level location is important in a VxRail environment because VxRail nodes remain ESXi hosts managed through vCenter; SSH is not enabled through a VxRail cluster host-management workflow. Enabling SSH should generally be limited to the approved troubleshooting window and stopped again after the support activity is complete, consistent with VMware security guidance. VMware documents starting and configuring ESXi services through the host's Configure tab in the [vSphere ESXi configuration documentation](#). Additional guidance on using and securing ESXi Shell and SSH is available in the [Broadcom knowledge base](#).

QUESTION NO: 3

A VxRail administrator wants to configure an alarm that is triggered when an ESXi host loses connectivity to the vSAN network.

Which vSphere inventory object should be selected when creating the alarm?

- A. The VxRail cluster
- B. The vCenter Server Appliance VM
- C. The vSAN datastore folder
- D. An individual virtual machine

ANSWER: A

Explanation:

The **VxRail cluster** is correct. A vSAN network connectivity condition is evaluated in the context of the vSAN cluster, whose hosts participate in the distributed datastore and vSAN network. Creating the alarm at the cluster level allows the condition to be monitored across the participating hosts.

Cluster-level alarms are appropriate for conditions affecting shared cluster services, such as vSAN health, HA, DRS, and resource usage. Host-level alarms are more appropriate for conditions that apply only to an individual ESXi host. See the [vSphere Monitoring and Performance documentation](#).

QUESTION NO: 4

A company has bought VxRail V Series appliances with PCIe network cards. How can the PCIe NIC ports be configured in VxRail for data traffic?

- A. Create new Standard or DV switch and add these NICs as uplinks to carry both vMotion and VM network traffic
- B. During initial VxRail setup the NIC ports would be configured automatically to carry VM network traffic
- C. Create new Standard or DV switch and add these NICs as uplinks to carry VM network traffic
- D. During initial VxRail setup the NIC ports would be configured automatically to carry both vMotion and VM network traffic

ANSWER: C

Explanation:

Create new Standard or DV switch and add these NICs as uplinks to carry VM network traffic is correct. On VxRail V Series appliances, the PCIe NIC ports are available as additional physical network capacity after the appliance has been deployed. An administrator can create a separate vSphere Standard Switch or a vSphere Distributed Switch and assign the PCIe NIC ports as physical uplinks. Port groups on that switch can then provide connectivity for virtual-machine data networks, separating application or tenant traffic from the VxRail system networking used during deployment.

This approach preserves the VxRail-managed network configuration and allows the added adapters to be deliberately assigned according to the organization's virtual-machine network design. It also gives administrators the flexibility to use either vSphere switch type, subject to their vSphere licensing and operational model. VMware documents that physical adapters can be assigned as uplinks to standard switches and distributed switches, and Dell's VxRail documentation provides the supported networking guidance for VxRail cluster operations. See [VMware vSphere Networking documentation](#) and [Dell VxRail documentation](#).

QUESTION NO: 5 - (DRAG DROP)

DRAG DROP

What is the correct sequence of steps to deploy a second VDS for VxRail network traffic?

Select and Place:

Cable VxRail NIC ports to ToR switches	Step 1
Configure the new VDS	Step 2
Create a VDS with the required number of uplinks	Step 3
Configure ToR switches to support VxRail network traffic on switch ports that are connected to the VxRail NIC ports	Step 4
Create port groups for the VxRail networks that must be migrated to the new VDS	Step 5

ANSWER:

Cable VxRail NIC ports to ToR switches	Configure ToR switches to support VxRail network traffic on switch ports that are connected to the VxRail NIC ports
Configure the new VDS	Cable VxRail NIC ports to ToR switches
Create a VDS with the required number of uplinks	Create a VDS with the required number of uplinks
Configure ToR switches to support VxRail network traffic on switch ports that are connected to the VxRail NIC ports	Configure the new VDS
Create port groups for the VxRail networks that must be migrated to the new VDS	Create port groups for the VxRail networks that must be migrated to the new VDS

Explanation:

Deploying a second vSphere Distributed Switch for VxRail traffic starts with preparing the physical network that will carry the new traffic path. The ToR switch ports connected to the intended VxRail NIC ports need the required network configuration first, including the appropriate VLAN availability and any required uplink, trunking, or link-aggregation settings. This establishes a ready and consistent physical transport layer for the additional VDS.

Once the ToR configuration is ready, the applicable VxRail NIC ports can be cabled to those prepared switch ports. The new VDS is then created with the required number of uplinks, so its logical uplink design matches the available physical NIC connectivity. After the VDS exists, it is configured for its intended VxRail network use, including the needed host association and uplink behavior. Finally, port groups are created for the VxRail networks that will move to the new VDS. These port groups provide the destination logical networks required for a controlled migration of VxRail traffic.

This ordering avoids attempting to use VDS uplinks before the connected physical network is ready, while also ensuring that port groups are created on an already-created and configured distributed switch. Dell's VxRail documentation resources are available through the [Dell VxRail documentation portal](#), and VMware's networking documentation describes the vSphere Distributed Switch model at [vSphere Distributed Switch documentation](#).

QUESTION NO: 6

A VxRail has been configured with the VxRail-deployed vCenter Server option. Which service VMs can be optionally hosted on the VxRail cluster?

- A. VxRail Manager, vCenter Server, and Platform Services Controller
- B. VxRail Manager and Secure Remote Services/VE
- C. vCenter Server, Platform Services Controller, and Secure Remote Services/VE
- D. vRealize Log Insight and Secure Remote Services/VE

ANSWER: C

Explanation:

With the VxRail-deployed vCenter Server deployment model, the VxRail cluster can host the vCenter Server Appliance and, in versions that use the external Platform Services Controller architecture, a Platform Services Controller virtual machine. These management components are deployed as part of the VxRail implementation rather than requiring customer-provided external infrastructure. The cluster can also host a Secure Remote Services/VE virtual appliance when secure remote connectivity is enabled. Secure Remote Services/VE provides the secure support-connectivity capability used to facilitate remote monitoring and support operations through Dell's secure remote-services infrastructure.

This grouping identifies the service virtual machines that may be placed on the VxRail cluster for this deployment choice: vCenter Server, Platform Services Controller, and Secure Remote Services/VE. The exact vCenter and Platform Services Controller topology varies by VxRail and VMware release; modern vCenter versions use an embedded Platform Services Controller, while older supported designs could deploy it separately. Dell's VxRail documentation and support resources describe the appliance architecture, deployment models, and supported management-service components. See the [Dell VxRail product documentation](#) and the [VxRail support matrix](#).

QUESTION NO: 7

As a VxRail administrator you have been asked to configure vSAN fault domains After selecting the VxRail cluster, where would you navigate to in vCenter to accomplish this task?

- A. Configure > vSAN > Configuration Assist > Fault Domains & Stretched Cluster
- B. Configure > vSAN > General > Fault Domains & Stretched Cluster
- C. Configure > vSAN > Disk Management > Fault Domains & Stretched Cluster
- D. Configure > vSAN > Fault Domains & Stretched Cluster

ANSWER: B

Explanation:

Configure > vSAN > General > Fault Domains & Stretched Cluster is the appropriate vCenter navigation path for configuring vSAN fault domains in the VxRail environment and vSphere client workflow covered by this exam. Fault domains enable an administrator to group vSAN hosts according to a shared physical failure boundary, such as a rack, power source, or top-of-rack switch. vSAN can then place replica components across those domains, helping ensure that a single infrastructure failure does not make all replicas of an object unavailable.

After the VxRail cluster is selected, the **Configure** tab exposes cluster configuration settings. The **vSAN** section contains the General configuration area, where the **Fault Domains & Stretched Cluster** control provides access to fault-domain

configuration and the related stretched-cluster workflow. This location is also where an administrator can create and manage fault domains before assigning hosts to the appropriate physical failure boundaries. Proper fault-domain design must reflect the actual hardware and network topology so that vSAN availability policies can deliver the intended resilience. VMware's vSAN administration documentation provides the configuration context for fault domains and stretched clusters: [vSAN Administration documentation](#). For VxRail operational guidance and lifecycle considerations, consult [Dell VxRail Support Matrix](#).

QUESTION NO: 8

Which components are managed by the vCenter server in the VxRail infrastructure?

- A. vSAN datastore, ESXi hosts, and vRealize Operations
- B. Virtual machines, vRealize LogInsight, and vSAN datastore
- C. vSAN cluster, ESXi hosts, and virtual networking
- D. Virtual networking, vRealize Operations, and virtual machines

ANSWER: C

Explanation:

vSAN cluster, ESXi hosts, and virtual networking is correct because vCenter Server is the central VMware management plane for the virtual infrastructure that underlies a VxRail cluster. It maintains the cluster inventory and configuration, manages the ESXi hosts that contribute compute and storage resources, and provides the interface through which administrators configure and monitor vSAN. Although vSAN uses local disks in the VxRail nodes, its datastore, policies, capacity, health, and cluster-level services are administered through vCenter Server.

vCenter Server also manages the virtual networking used by the hosts and workloads. This includes network objects and configuration such as port groups, VMkernel adapters, uplinks, and, where deployed, vSphere Distributed Switches. Centralizing these functions in vCenter allows consistent configuration across all nodes and enables cluster services such as vMotion, vSAN traffic, and VM networking. In VxRail, VxRail Manager integrates with vCenter Server and automates lifecycle and platform operations, while vCenter remains the authoritative interface for the VMware cluster, host, storage, and virtual-networking constructs. Dell's VxRail documentation describes this integrated management model, and VMware documents vCenter's cluster and host-management role in its vSphere documentation: [Dell VxRail documentation](#) and [VMware vSphere documentation](#).

QUESTION NO: 9

What is the limit of locally cached managed storage included in the CloudArray Virtual Edition license packaged with a VxRail system?

- A. 1 TB
- B. 5 TB
- C. 10 TB
- D. 15 TB

ANSWER: A

Explanation:

1 TB is the limit of locally cached managed storage included with the CloudArray Virtual Edition license packaged with the applicable VxRail system offering. CloudArray Virtual Edition provides a storage-gateway function: it presents storage locally while using cloud or object storage as the durable back end. Its local cache is used to retain active data and support efficient access, while the managed-storage entitlement establishes the capacity permitted by the bundled license. For this VxRail package, the included entitlement is specifically 1 TB rather than a larger capacity tier.

This question refers to the bundled CloudArray Virtual Edition licensing described for the VxRail release documentation, not to the physical capacity of the VxRail cluster, the amount of cloud-object capacity that may be available in a target, or a separately purchased CloudArray license expansion. Therefore, the relevant value is the licensed locally cached managed-storage allowance included with the system. Dell EMC's VxRail vCenter Server planning guidance identifies the CloudArray Virtual Edition entitlement associated with this package. See the [VxRail vCenter Server Planning Guide](#) and Dell's [data storage resources](#) for product and storage-platform context.

QUESTION NO: 10

A VxRail administrator is planning to add capacity disks to an existing vSAN disk group that uses the original storage architecture.

What is the maximum number of capacity devices supported in one disk group?

- A. 4
- B. 5
- C. 7
- D. 8

ANSWER: C

Explanation:

7 is correct. In the vSAN original storage architecture, each disk group contains one cache device and from one to seven capacity devices. The cache device provides the caching tier for the disk group, while the capacity devices contribute persistent storage capacity to the vSAN datastore.

When additional capacity is required beyond seven capacity devices, the administrator must create another disk group with its own cache device, subject to the VxRail hardware configuration and supported expansion options. See the [vSAN Planning and Deployment documentation](#) and [Dell VxRail Documentation](#).

QUESTION NO: 11

A VxRail administrator creates a vSAN storage policy that includes an Object Space Reservation value of 100%.

What is the effect of this setting?

- A. The full logical size of the object is reserved on the vSAN datastore
- B. The object is limited to 100 IOPS
- C. Two full object replicas are created
- D. Data is automatically encrypted at rest

ANSWER: A

Explanation:

The full logical size of the object is reserved on the vSAN datastore is correct. Object Space Reservation controls the amount of vSAN capacity reserved for an object before the object consumes that capacity with guest data. A value of 100% reserves capacity equal to the complete logical size of the virtual disk or other object.

This setting is commonly considered when an application requires thick-provisioned storage behavior or when the administrator must ensure capacity is available for the full object size. The reservation is applied in addition to capacity overhead required by the assigned availability policy, such as mirroring. See the [vSAN Administration documentation](#).

QUESTION NO: 12

How much capacity of vSphere Data Protection is provided with a VxRail appliance?

- A. 1 TB
- B. 8 TB
- C. 10 TB
- D. 16 TB

ANSWER: B

Explanation:

8 TB is the vSphere Data Protection capacity included with the VxRail appliance offering referenced by this question. VxRail integrated VMware software and data-protection capabilities as part of its appliance solution, and the bundled vSphere Data Protection entitlement provided 8 TB of backup capacity. This capacity is used by the vSphere Data Protection virtual appliance for protecting virtual machines and is distinct from the raw or usable storage capacity installed in the VxRail cluster. In practical terms, the included entitlement lets an administrator deploy and configure VDP for image-level virtual-machine backups without separately purchasing a smaller VDP capacity license. When planning protection, administrators should still account for retention requirements, change rates, deduplication, and the backup destination, since these factors determine whether the included capacity meets the environment's operational recovery objectives. The VxRail data-protection white paper identifies the included VDP capacity in its description of the appliance's data-protection capabilities. See the [Dell EMC Data Protection for VxRail white paper](#) and VMware's [vSphere Data Protection documentation](#).

QUESTION NO: 13 - (SIMULATION)

What is the order of steps to generate the SolVe procedure to add a capacity drive to a VxRail node?

ANSWER: See the explanation for the answer

Explanation:

Correct order:

1. Sign in to **Dell SolVe Online** with a Dell Support account.
2. Select **VxRail**, then select **VxRail Appliance**.
3. Open the VxRail **procedures** list and select **Add a Capacity Drive to a VxRail Node** (the drive-add/capacity-expansion procedure).
4. Provide the requested appliance/node identification and hardware information (for example, Service Tag and the capacity-drive details).
5. Select **Generate** to create/download the customized SolVe procedure, then use that generated procedure to perform the drive addition.

The Service Tag/node and drive information must be entered before generating the document so SolVe produces the procedure appropriate for the specific VxRail hardware configuration.

Reference: [Dell KB 000022572 — VxRail: Where do I locate VxRail Procedural Documentation?](#)

QUESTION NO: 14

Which tool is used for troubleshooting vSAN related issues when vCenter Server is unavailable?

- A. vSAN Proactive Tests
- B. Event Code Reference
- C. ESXCLI

ANSWER: C

Explanation:

ESXCLI is the appropriate tool because it can be run directly on an ESXi host through an SSH session or the ESXi Shell, without requiring vCenter Server to be available. The ESXCLI vSAN command namespace provides host-level visibility into vSAN configuration, cluster membership, network settings, disk groups, physical disks, and resynchronization activity. For example, an administrator can use commands such as `esxcli vsan cluster get` to verify the host's vSAN cluster state and `esxcli vsan storage list` to inspect vSAN storage devices. This makes ESXCLI especially valuable during a vCenter outage, when vSphere Client workflows and vCenter-dependent health views cannot be accessed.

In VxRail troubleshooting, Dell documentation specifically directs administrators to access affected ESXi hosts and use ESXCLI commands to validate vSAN cluster information when hosts appear disconnected or not responding in vCenter. VMware also documents the ESXCLI vSAN command set as a supported interface for retrieving and managing vSAN host information. These direct host commands provide the diagnostic data needed to determine whether the underlying vSAN services and membership remain healthy independently of vCenter availability. See [Dell VxRail troubleshooting guidance](#) and [VMware ESXCLI vSAN documentation](#).

QUESTION NO: 15

What is a consideration when using a VxRail Cluster with vSphere Standard Edition licensing?

- A. Administrative workload is decreased
- B. Workload balancing is automated
- C. vGPU is not supported
- D. Drive replacement is streamlined

ANSWER: C

Explanation:

vGPU is not supported is the relevant consideration for a VxRail cluster licensed with vSphere Standard Edition. NVIDIA virtual GPU functionality depends on VMware capabilities and licensing that are available with vSphere Enterprise Plus rather than vSphere Standard in the VxRail licensing model covered by this exam. A VxRail environment can contain GPU hardware, but assigning a physical NVIDIA GPU to virtual machines as shared, mediated virtual GPU resources requires the appropriate vSphere edition as well as compatible NVIDIA vGPU software and licensing.

This distinction matters during VxRail design: workloads such as virtual desktop infrastructure, graphics-intensive engineering applications, and machine-learning workloads may require GPU virtualization. The cluster must therefore be planned with the required Enterprise Plus licensing before deploying those workloads. This is a licensing and feature-entitlement consideration, not merely a hardware-selection decision. Dell's VxRail planning guidance identifies licensing implications that must be evaluated when choosing vSphere editions, while NVIDIA's deployment documentation describes the VMware vSphere requirements for vGPU environments. See the [Dell VxRail vCenter Server Planning Guide](#) and the [NVIDIA Virtual GPU Software User Guide](#).

QUESTION NO: 16

An application consisting of two VMs must be kept on two separate hosts. Which VMware feature would enable this?

- A. Fault Tolerance
- B. High Availability
- C. Storage vMotion

ANSWER: D

Explanation:

Distributed Resource Scheduler (DRS) VM-VM anti-affinity rule is the appropriate feature because it defines a placement constraint between the two virtual machines: they should run on different ESXi hosts. In vSphere, this is configured as a VM/VM anti-affinity rule, commonly expressed as “separate virtual machines.” DRS evaluates the rule when initially placing workloads and during load-balancing recommendations or migrations, helping maintain host separation while managing cluster resources. This is particularly useful when two application components need host-level isolation to reduce the effect of a single host failure or to meet an application availability requirement.

For the strongest enforcement, the rule can be configured as a mandatory rule, so DRS does not intentionally place both VMs together. Administrators should ensure that the cluster has sufficient eligible hosts and capacity to satisfy the rule, especially during maintenance or host-failure scenarios. VMware’s vSphere documentation describes VM/VM affinity rules as cluster rules that control whether specified VMs run together or separately, and identifies DRS as the service that applies these placement rules. See the [vSphere VM/VM affinity rules documentation](#) and the [vSphere DRS resource-management documentation](#).

QUESTION NO: 17

An environment has two VxRail clusters with the following hardware configuration:

- Cluster-A is using FTT = 2 and FTM = Erasure Coding
- Cluster-B is using FTT = 2 and FTM = Mirroring

A 100 GB VM is planned to be vMotioned from Cluster-A to Cluster-B.

How much more vSAN capacity is consumed in Cluster-B than in Cluster-A?

- A. 0 GB; consumption is unchanged
- B. 100 GB
- C. 150 GB
- D. 167 GB

ANSWER: C

Explanation:

150 GB is correct. With failures to tolerate set to 2, vSAN’s erasure-coding policy uses the RAID-6 configuration, which stores four data components and two parity components. This produces a 1.5× capacity overhead for protected data. Therefore, the 100 GB VM consumes 150 GB of vSAN capacity in Cluster-A.

For the same failures-to-tolerate value with mirroring, vSAN uses RAID-1 mirroring with three full replicas of the data: the primary replica plus two additional replicas. The 100 GB VM consequently consumes 300 GB in Cluster-B. Comparing the destination consumption of 300 GB with the source consumption of 150 GB gives an additional capacity requirement of 150 GB after the migration.

This calculation concerns the policy-level storage consumption for the VM’s data and assumes the stated logical VM size, without adding operational overheads such as witness components, metadata, snapshots, swap objects, or slack space. VMware documents the capacity multipliers for vSAN storage policies, including 1.5× for RAID-6 and 3× for RAID-1 at failures to tolerate = 2, in its [vSAN design and sizing guidance](#). See also the [VMware vSAN space-efficiency overview](#).