

DUMPS ARENA

Dell VxRail Operate 2023

EMC D-VXR-OE-23

Version Demo

Total Demo Questions: 7

Total Premium Questions: 71

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, VxRail Concepts and Features	26
Topic 2, VxRail Management	18
Topic 3, VxRail Maintenance	17
Topic 4, VxRail Troubleshooting	10
Total	71

QUESTION NO: 1

What are two considerations of deploying VxRail Cluster with an existing customer-supplied vCenter Server? (Select 2)

- A. Requires an internal DNS server
- B. Supports Enhanced Link Mode with other VxRail-managed vCenter Servers
- C. Supports vCenter HA
- D. Uses customer-defined procedures for vCenter LCM

ANSWER: B D

Explanation:

“Supports Enhanced Link Mode with other VxRail-managed vCenter Servers” is a valid consideration because a customer-supplied vCenter Server can participate in an Enhanced Link Mode configuration when its vCenter Single Sign-On domain and topology meet VMware requirements. Enhanced Link Mode provides a common inventory view and enables administrators to use one vSphere Client session to manage linked vCenter Server systems. This can be useful where the customer needs centralized visibility across VxRail environments and other vSphere resources while retaining an existing vCenter Server deployment. VMware describes the capabilities and configuration concepts for Enhanced Link Mode in its [vCenter Server Configuration documentation](#).

“Uses customer-defined procedures for vCenter LCM” is also correct. When VxRail is deployed with a customer-supplied vCenter Server, the customer owns operational responsibility for that vCenter instance, including its lifecycle planning, maintenance windows, upgrades, backups, and compatibility validation. VxRail lifecycle management continues to orchestrate supported VxRail infrastructure software and hardware updates, but it does not assume ownership of the externally managed vCenter lifecycle. Dell documents the customer-managed vCenter deployment model and associated operational considerations in the [Dell VxRail System Tech Book](#).

QUESTION NO: 2 - (SIMULATION)

An existing 4-node VxRail standard cluster has to be scaled out with 2 additional nodes and configured as a stretched cluster. An administrator must prepare a vSAN storage policy that should have the following:

- vSAN Policy Name: vSAN

Stretched Cluster

- Data protection across the sites

- Local protection: RAID-1 with

FTT=1

- Thick Provisioning

An administrator should be able to provision a VM even if the currently available cluster resources cannot satisfy the current policy.

Use the simulator to create the required vSAN storage policy.

ANSWER: See the explanation for the answer

Explanation:

Create the VM Storage Policy with the name: vSAN Stretched Cluster.

In the policy wizard, select the **vSAN** storage rules and configure the availability and provisioning rules as follows:

Site disaster tolerance / data protection across sites: Dual site mirroring (stretched-cluster protection).

Primary level of failures to tolerate (PFTT): 1 failure - RAID-1 (Mirroring), if this field is presented separately.

Secondary level of failures to tolerate (SFTT), the local-site protection: 1 failure - RAID-1 (Mirroring).

Object space reservation / provisioning type: Thick provisioning (equivalent to 100% object-space reservation).

Force provisioning: Yes (enabled). This permits provisioning even when the currently available vSAN resources cannot satisfy the policy.

Complete the wizard with **Finish**. The essential stretched-cluster configuration is dual-site mirroring with local RAID-1, FTT=1, plus thick and forced provisioning.

QUESTION NO: 3

An administrator wants to get information about a specific host and runs this code:

```
GET/v1/hosts/{sn}
```

An API code 404 is returned

What is the reason for the error code?

- A. The administrator was not authenticated
- B. The provided {sn} was not valid
- C. There was a server error
- D. The API is designed to return all hosts in a cluster

ANSWER: B

Explanation:

The provided {sn} was not valid. The `GET /v1/hosts/{sn}` endpoint requests one particular host, using its serial number as the resource identifier. A successful response depends on replacing {sn} with the serial number of a host that is known to the VxRail API service and accessible within the relevant VxRail environment. When no matching host resource is found for that identifier, the service returns HTTP status 404 Not Found. In practical terms, the administrator should verify that the placeholder was replaced rather than sent literally, confirm the serial number was copied accurately, and ensure it identifies a host currently represented by the API inventory. HTTP defines 404 as the status used when the origin server cannot find a current representation for the requested target resource. This is consistent with a host lookup whose supplied serial-number path parameter does not resolve to an existing host. See the [MDN reference for HTTP 404 Not Found](#) and Dell Technologies' [Developer Portal](#) for VxRail API resources and documentation access.

QUESTION NO: 4

An administrator is validating jumbo-frame connectivity between two VxRail hosts on the vSAN network. The vSAN VMkernel adapter is vmk1, and the physical network is configured for an MTU of 9000.

Which command should be used from an ESXi host to validate that an MTU-sized packet can traverse the vSAN network without fragmentation?

- A. `vmkping -I vmk1 -d -s 8972`
- B. `vmkping -I vmk0 -d -s 8972`
- C. `ping -s 9000`
- D. `vmkping -I vmk1 -s 9000`

ANSWER: A

Explanation:

`vmkping -I vmk1 -d -s 8972 <remote-vSAN-IP>` is correct. The `-I vmk1` option sends the test through the VMkernel adapter used for vSAN traffic. The `-d` option sets the Don't Fragment bit, so the test fails if the path cannot support the requested packet size. For an MTU of 9000, an IPv4 payload size of 8972 bytes plus the 28-byte IP and ICMP headers validates the full MTU.

A standard `ping` test does not validate jumbo-frame support, and testing through `vmk0` validates the management network rather than the vSAN network.

QUESTION NO: 5

A VxRail cluster is deployed with a customer-supplied external vCenter Server. The administrator must determine which lifecycle responsibilities remain coordinated through VxRail Manager.

Which two operations are managed through VxRail Manager? (Select 2)

- A. Updating the VxRail cluster by using a supported composite bundle
- B. Adding a compatible VxRail node through the VxRail expansion workflow
- C. Upgrading the customer-supplied external vCenter Server
- D. Backing up the customer-supplied external vCenter Server
- E. Changing the external vCenter Single Sign-On domain configuration

ANSWER: A B

Explanation:

VxRail Manager coordinates supported VxRail cluster lifecycle updates by using VxRail composite bundles, and it manages the VxRail workflow for adding compatible VxRail nodes to an existing cluster. These operations preserve the VxRail appliance configuration and supported hardware and software interoperability.

With a customer-supplied external vCenter Server, the customer uses customer-defined procedures for vCenter lifecycle management and backup. VxRail Manager does not assume ownership of the externally managed vCenter Server.

QUESTION NO: 6 - (SIMULATION)

VxRail administrator got the following warning while on node physical view.

"Health monitoring is currently disabled and information may not be accurate.

Dell recommends that health monitoring is enabled."

Use the simulator to enable the relevant service and resolve the issue.

ANSWER: See the explanation for the answer

Explanation:

Enable the ESXi CIM Server service (CIM health monitoring) on the affected VxRail node(s). The Physical View warning is displayed when the host hardware health/CIM monitoring service is disabled.

1. In the vSphere Client, select the VxRail ESXi host reporting the warning.
2. Select **Configure > System > Services**.
3. Select **CIM Server** (the `sfcbd-watchdog` service).

4. Click **Start**.
5. Use **Edit Startup Policy** and set it to **Start and stop with host** (or the simulator's equivalent automatic startup option), then save.
6. Repeat for other VxRail hosts if CIM Server is disabled on them.

After the CIM Server is running, return to the VxRail **Physical View** and refresh it. Hardware health monitoring can then collect node health information and the warning is resolved.

QUESTION NO: 7

Using the VxRail plug-in, where can iDRAC settings be configured?

- A. Cluster > Configure > VxRail > iDRAC Configuration
- B. Host > Configure > Hardware > iDRAC Configuration
- C. Cluster > Configure > VxRail > Networking
- D. Host > Configure > VxRail > iDRAC Configuration

ANSWER: D

Explanation:

Host > Configure > VxRail > iDRAC Configuration is the correct location because iDRAC configuration is managed on a per-host basis in the VxRail vCenter plug-in. An iDRAC belongs to an individual physical VxRail node and provides out-of-band hardware management for that node, including its remote-access networking and related controller settings. Selecting the applicable host first ensures that changes are directed to the iDRAC associated with that specific appliance rather than to a cluster-level virtual infrastructure setting.

Within the selected host's configuration view, the VxRail section exposes VxRail-specific management workflows, including the iDRAC Configuration page. This integration enables an administrator to review and update supported iDRAC settings through vCenter while remaining within the VxRail management experience. Appropriate administrative privileges are required, and administrators should follow their organization's network and security procedures before changing remote-management connectivity. Dell's VxRail product documentation is available through the [Dell VxRail documentation portal](#); Dell also provides VxRail operational guidance through the [Dell Technologies Info Hub](#).