

DUMPS ARENA

Dell VxRail Design

EMC D-VXR-DS-00

Version Demo

Total Demo Questions: 6

Total Premium Questions: 64

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

What are two reasons to provide minimum ceiling clearance above a VxRail rack? (Choose two.)

- A. Room for forklift
- B. Sprinkler system
- C. Service access
- D. Anti-static protection
- E. Airflow management

ANSWER: B E

Explanation:

“Sprinkler system” and “Airflow management” are the two reasons for maintaining the required unobstructed space above a VxRail rack. Data-center fire-suppression design depends on sprinkler discharge reaching the protected area without rack tops, cable pathways, or other infrastructure interfering with the intended spray pattern. Maintaining adequate vertical clearance helps support effective operation of the room’s fire-protection system and alignment with the facility’s applicable fire-code requirements.

Ceiling clearance also supports the thermal design of the rack and room. VxRail nodes produce heat that must be removed through a controlled airflow path. Space above the rack allows hot exhaust air to return to the room cooling system and reduces the risk that overhead obstructions will impede exhaust flow, trap heat, or contribute to recirculation. This is particularly important when the room uses hot-aisle/cold-aisle containment or overhead return-air pathways. Dell’s VxRail planning guidance requires customers to prepare a suitable installation environment, including appropriate power, cooling, and physical-site conditions; the final rack layout should therefore be coordinated with both HVAC and fire-protection requirements. See the [Dell VxRail documentation library](#) and the [ASHRAE Thermal Guidelines for Data Processing Environments](#).

QUESTION NO: 2

A customer requires a VxRail stretched cluster that can continue serving protected workloads after the loss of either data site.

Which two design decisions support this requirement? (Choose two.)

- A. Deploy the vSAN witness in an independent third failure domain.
- B. Assign Site disaster tolerance=Site mirroring – stretched cluster to protected workloads.
- C. Deploy the vSAN witness in the preferred data site with the majority of the hosts.
- D. Assign Site disaster tolerance=None – standard cluster to protected workloads.
- E. Use RAID 5 erasure coding instead of site mirroring for all protected workloads.

ANSWER: A B

Explanation:

Deploying the vSAN witness in an independent third failure domain and assigning Site disaster tolerance=Site mirroring – stretched cluster to protected workloads are required design decisions. The witness provides the quorum component outside the two data sites, while site mirroring maintains a complete data replica in each data-site fault domain.

Locating the witness at either data site weakens failure-domain separation. A standard-cluster policy does not provide site-level mirroring. RAID 5 alone provides local erasure coding and does not replace the site-mirroring requirement for protection from a complete site failure.

QUESTION NO: 3

What are the two types of port groups of vSphere Distributed Switches? (Choose two.)

- A. Virtual machine port groups
- B. Uplink port groups
- C. VMkernel port groups
- D. Distributed port groups

ANSWER: B D

Explanation:

A vSphere Distributed Switch provides two distinct port-group types: distributed port groups and uplink port groups. Distributed port groups define the network connectivity and policies applied to virtual machine network adapters and VMkernel adapters. They are configured centrally on the distributed switch, with their settings propagated consistently to participating ESXi hosts. This central management is particularly important in VxRail environments, where hosts must use consistent networking configurations for services such as management, vMotion, vSAN, and virtual-machine traffic.

Uplink port groups represent the physical-network-facing connections of the distributed switch. They contain uplink ports that are mapped to physical NICs on each ESXi host, allowing the distributed switch to forward traffic from distributed port groups onto the physical network. The number of uplinks is defined for the switch and can be associated with host physical adapters as appropriate. VMware's vSphere networking documentation explicitly identifies distributed port groups and uplink port groups as the port-group types available on a vSphere Distributed Switch. See the [vSphere Distributed Switch documentation](#) and the [vSphere port group documentation](#).

QUESTION NO: 4

What are the minimum network bandwidths required for vSAN?

- A. 10 GbE for OSA, 25 GbE for ESA
- B. 1 GbE for OSE, 10 GbE for ESA
- C. 10 GbE for OSA, 10 GbE for ESA

ANSWER: A

Explanation:

10 GbE for OSA, 25 GbE for ESA is correct for VxRail vSAN networking. The Original Storage Architecture (OSA) supports a minimum 10 GbE vSAN network, which provides the bandwidth needed for storage traffic such as replication, rebuilds, resynchronization, and normal virtual-machine I/O. The Express Storage Architecture (ESA) is designed around higher-performance NVMe-based storage and a log-structured file system, and it requires a minimum 25 GbE network. This higher baseline is necessary to prevent the network from becoming a bottleneck as ESA processes storage traffic at substantially greater throughput and lower latency. These values apply to the vSAN network used for interhost storage communication; a resilient design should also provide redundant physical networking, with appropriate VLAN and switch configuration, to maintain availability if a link or adapter fails. VMware documents the 25 GbE minimum for ESA and the 10 GbE minimum for OSA in its vSAN planning guidance. See the [VMware vSAN Network Design documentation](#) and Dell's [VxRail Support Matrix](#) for platform-specific compatibility and design requirements.

QUESTION NO: 5

A VxRail architect must design a new VxRail cluster using the VxRail Sizing Tool. The new cluster will replace an existing 3 tier infrastructure and must support Hadoop and Tableau workloads.

Which purpose and use case should be selected in the VxRail sizing tool?

- A. Purpose: GreenfieldUse Case: In Memory database
- B. Purpose: Datacenter ConsolidationUse Case: Relational database
- C. Purpose: GreenfieldUse Case: Big data
- D. Purpose: VxRail Technology RefreshUse Case: Big data

ANSWER: D

Explanation:

Purpose: VxRail Technology Refresh

Use Case: Big data is the appropriate selection. Replacing an existing three-tier infrastructure is a technology-refresh initiative: the customer is modernizing an established environment by moving its compute, storage, and virtualization infrastructure to a VxRail hyperconverged platform. Selecting VxRail Technology Refresh allows the sizing exercise to reflect an existing-workload migration or replacement scenario rather than a net-new deployment.

Hadoop is a distributed data-processing and analytics framework commonly used for large-scale data workloads. Tableau is a business-intelligence and visualization platform that can consume and present analytics results from enterprise data sources, including big-data platforms. Together, these requirements align with the Big data use case in the VxRail Sizing Tool. The use-case selection helps frame workload characteristics and sizing assumptions for the intended application category, while the architect still collects detailed resource requirements, capacity needs, availability requirements, and projected growth before finalizing the VxRail configuration.

Dell positions VxRail as an integrated HCI system for modernizing data-center infrastructure and supporting a broad range of applications. See the [Dell VxRail overview](#) and the [VxRail Support Matrix](#) for current product and compatibility information.

QUESTION NO: 6

What is a result of using the VxRail Sizing Tool to size a new cluster?

- A. The lowest price solution is the recommended configuration.
- B. The highest price solution is the recommended configuration.
- C. Single workloads are only considered when sizing the new cluster.
- D. A recommended cluster configuration is produced based on single or combined workload profiles.

ANSWER: D

Explanation:

Using the VxRail Sizing Tool results in a recommended VxRail cluster configuration derived from the resource requirements of the workload profile or profiles entered into the tool. The sizing workflow captures workload characteristics such as virtual-machine quantities, compute requirements, memory, storage capacity, and performance-related needs, then uses those inputs to determine an appropriate node configuration and cluster size. This enables a design to account for a consolidated environment containing multiple workloads, rather than restricting the assessment to one workload at a time. The recommendation is therefore a technical sizing outcome intended to meet the stated workload demand and provide a suitable VxRail configuration; it is not inherently a recommendation of either the lowest- or highest-priced configuration. Dell describes VxRail as an integrated HCI system designed to simplify infrastructure deployment and lifecycle operations, while sizing is performed against workload and capacity requirements. See the [Dell VxRail product overview](#) and the [VxRail Support Matrix](#) for Dell VxRail platform information and supported configuration guidance.