

Specialist - Technology Architect Midrange Storage Solutions Exam

EMC DES-1D12

Version Demo

Total Demo Questions: 18

Total Premium Questions: 189

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Dell EMC Unity Family	103
Topic 2, Dell EMC SC Series	25
Topic 3, Dell EMC VNX Series	2
Topic 4, Mix Questions	59
Total	189

QUESTION NO: 1 - (DRAG DROP)

DRAG DROP

In vCenter, what is the correct order of steps to create a VM in a storage container provisioned on a Dell EMC PowerStore system?

Select and Place:

Create vVol datastore	Step 1
Provision VMs based on vVol datastore	Step 2
Create VM Storage Policies	Step 3
Configure PowerStore connectivity	Step 4
Add PowerStore Cluster as VASA Storage Provider	Step 5

ANSWER:

Create vVol datastore	Configure PowerStore connectivity
Provision VMs based on vVol datastore	Add PowerStore Cluster as VASA Storage Provider
Create VM Storage Policies	Create VM Storage Policies
Configure PowerStore connectivity	Create vVol datastore
Add PowerStore Cluster as VASA Storage Provider	Provision VMs based on vVol datastore

Explanation:

Creating virtual machines with PowerStore vVol storage in vCenter follows the dependency chain between host access, VASA integration, policy-based storage control, and datastore consumption. First, configure PowerStore connectivity for the ESXi hosts. This establishes the supported storage path between the VMware environment and the PowerStore appliance, allowing the hosts to discover and use PowerStore storage resources. The connectivity configuration is the foundation for every later storage operation.

Next, add the PowerStore cluster as a VASA storage provider in vCenter. VASA is the VMware control-plane integration used by vVols. Through the registered provider, PowerStore advertises its available storage capabilities and reports the information vCenter needs to manage vVol-based storage. Dell describes this PowerStore and VMware integration in its [PowerStore Virtualization Integration white paper](#).

With capability information available through VASA, create VM storage policies. These policies express the storage requirements to be consumed by virtual machines, such as the service level or protection characteristics exposed by PowerStore. vCenter can then evaluate policy compatibility using the provider's published capabilities. After the policy layer is available, create the vVol datastore. The datastore represents the PowerStore vVol storage container made available to vCenter and is the location selected when deploying or configuring a VM.

Finally, provision VMs based on the vVol datastore. During provisioning, vCenter applies the selected VM storage policy and PowerStore creates and manages the individual virtual volumes that make up the VM's configuration, virtual disks, and related objects. This ordering ensures that both data access and policy-aware VASA management are ready before VM objects are deployed.

QUESTION NO: 2 - (SIMULATION)

Match each term with its description.

Term:	Description:
Throughput	Time taken for an IO to be executed
Latency	IOs in buffer waiting to be executed
IOPS	Amount of IO generated by an application
Queue depth	Sized by type and number of ports to handle traffic

ANSWER: See the explanation for the answer

Explanation:

The correct matches are:

QUESTION NO: 3

A customer has a mixed environment consisting of Dell Unity 350F Dell Unity XT 680F. and Dell PowerStore 5000X They want to have a single interface to monitor and report on all the arrays

Which software tool should the technical architect recommend?

- A. CloudIQ
- B. Unisphere Analyzer
- C. Unisphere Central
- D. PowerStore Manager

ANSWER: A

Explanation:

CloudIQ is the appropriate recommendation because it is Dell's cloud-based proactive monitoring and analytics application designed to provide a consolidated view across supported Dell infrastructure platforms. Both Dell Unity-family systems, including Unity 350F and Unity XT 680F, and Dell PowerStore systems can send telemetry to CloudIQ. This lets the customer use one interface to view system health, capacity consumption, performance trends, configuration information, alerts, and predictive analytics across the mixed storage estate.

CloudIQ is particularly suited to the stated monitoring and reporting requirement because it aggregates data from multiple arrays and presents fleet-level dashboards as well as array-specific details. It also provides historical reporting and health insights that help administrators identify capacity risks, performance anomalies, and operational issues before they affect workloads. Access is provided through a secure Dell-hosted portal, avoiding the need to deploy a separate on-premises management server solely to consolidate visibility. Dell describes CloudIQ as a proactive monitoring and predictive analytics application for supported Dell infrastructure, and its support documentation identifies Dell Unity and PowerStore among the supported storage platforms. See the [Dell CloudIQ overview](#) and the [Dell CloudIQ detailed review](#).

QUESTION NO: 4

What is the maximum capacity of a file system on a Dell Unity XT 380?

- A. 64 TB
- B. 32 TB

C. 256 TB

D. 128 TB

ANSWER: C

Explanation:

The maximum supported capacity for a single Dell Unity XT file system is **256 TB**. This is a per-file-system platform limit: it defines the largest logical file system that can be created and presented through Unity NAS services, rather than the total raw or usable capacity of the storage array. A Unity XT 380 can host multiple file systems, subject to the applicable system, pool, and software-limit considerations, but no individual file system can exceed 256 TB. The limit applies when planning NAS storage for NFS and SMB client access, including the sizing of the storage pool that supplies the file system. Administrators should therefore validate both the required file-system size and sufficient available pool capacity before provisioning or expanding a file system. Dell's Unity XT documentation and product support materials provide the authoritative platform specifications and configuration guidance for the Unity XT 380 family. See the [Dell Unity XT 380 documentation page](#) and the [Dell Technologies Info Hub for Dell Unity XT](#).

QUESTION NO: 5

Which Dell PowerStore policy object can combine local snapshot scheduling with asynchronous replication settings for a storage resource?

A. Protection policy

B. Host group

C. Storage policy

D. Mobility rule

ANSWER: A

Explanation:

Protection policy is correct. A PowerStore protection policy is assigned to a supported storage resource to automate data-protection operations. The policy can associate a snapshot rule for local point-in-time copies and a replication rule for asynchronous replication. This allows the administrator to apply a reusable protection configuration to volumes, volume groups, and other supported resources.

A snapshot rule defines snapshot frequency and retention behavior, while a replication rule defines the replication schedule and recovery point objective. Neither rule alone provides the combined policy assignment. Dell describes protection policies, snapshot rules, and replication rules in the [Dell PowerStore Protect Data documentation](#).

QUESTION NO: 6

What component or feature is provided by both the Dell EMC PowerStore 1000T array and the Dell EMC Unity XT 380 Series array?

A. Container Integration

B. Configured Block with Hypervisor

C. Unified System

D. NVMe disks

ANSWER: C

Explanation:

Unified System is provided by both platforms. The Dell EMC PowerStore 1000T is a PowerStore T model, designed to deliver unified storage services from the same appliance. It supports block storage, file storage through NAS services, and VMware virtual volumes, allowing organizations to consolidate multiple workload types on one system. The Dell EMC Unity XT 380 Series is likewise a unified storage platform, providing both block and file services in a single array. Its unified architecture can present block storage through protocols such as Fibre Channel and iSCSI while also providing NAS file services, depending on the configured connectivity and services. This shared capability is significant because it enables each array to support mixed application environments without requiring separate dedicated block-only and file-only storage systems. In practice, both platforms can therefore serve virtualized infrastructure, databases, and file-sharing workloads through one storage family and common management model. Dell's PowerStore documentation describes the unified block and file capabilities of PowerStore T appliances, while Dell Unity XT documentation identifies Unity XT as a unified storage family with integrated block and file functionality. See the [Dell PowerStore product documentation](#) and [Dell Unity family product documentation](#).

QUESTION NO: 7

A customer has a Dell EMC SC9000 array with six SC420 enclosures running SCOS 7.3. The enclosure configuration is as follows:

- Two enclosures have write-intensive SSDs
- Four enclosures have read-intensive SSDs

What is recommended as a general guideline for making spare drive assignments?

- A. 2 spare drives per shelf
- B. Dedicated spares are not required
- C. 2 spares per drive type per array
- D. 1 spare per drive type per shelf

ANSWER: D

Explanation:

1 spare per drive type per shelf is the recommended general guideline. In an SC Series system, spare capacity must be suitable for the failed drive it may replace, including the appropriate drive/media type. Maintaining a spare of each installed drive type in every enclosure ensures that a locally available compatible drive is ready if a disk fails. For this configuration, each enclosure containing write-intensive SSDs should have a write-intensive SSD spare, and each enclosure containing read-intensive SSDs should have a read-intensive SSD spare. This approach supports prompt reconstruction without relying on a compatible spare located in another enclosure and helps preserve expected availability and performance characteristics during a failure event. Spare planning must account for the physical distribution of drives, not merely the total number of disks in the array, because an enclosure can be affected independently and because the two SSD endurance classes are distinct drive types for capacity-management purposes. Dell's SC Series documentation and support resources provide the platform-specific guidance for drive, enclosure, and storage management procedures. See the [Dell EMC SC9000 documentation page](#) and Dell's [SC Series Storage Manager and Unisphere documentation resources](#).

QUESTION NO: 8

When sizing a Dell EMC Unity Hybrid array, what are considerations for growth?

- A. Port utilization and load balancing
- B. Cache utilization and hardware bottlenecks
- C. Latency and queue depth
- D. General capacity and active data

ANSWER: D

Explanation:

General capacity and active data is correct because sizing growth for a Dell EMC Unity Hybrid array requires more than projecting total raw or usable capacity. The design must account for both overall capacity growth and the amount and location of active data over time. In a hybrid system, the active or “hot” portion of the dataset has a direct effect on the flash tier, since FAST VP data relocation places the most frequently accessed 256 MB slices on the most appropriate tier. As the environment grows, architects should anticipate whether the active working set will grow, whether hot data will move to new applications or volumes, and whether additional flash capacity will be needed to maintain expected performance. Longer-term growth planning should also consider expected consumption trends and expansion options so that capacity and performance can scale together without disrupting service levels. Dell’s Unity documentation describes how FAST VP automatically relocates data among tiers according to activity, making active-data growth a key sizing consideration for hybrid configurations. See the [Dell Unity FAST VP overview](#) and the [Dell Unity support documentation](#).

QUESTION NO: 9

How many hierarchical snapshots are supported by Dell EMC Unity?

- A.** 32 levels deep for LUNs
16 levels deep for file systems
- B.** 16 levels deep for LUNs
8 levels deep for file systems
- C.** 16 levels deep for LUNs
16 levels deep for file systems

ANSWER: B

Explanation:

16 levels deep for LUNs

8 levels deep for file systems is correct. Dell EMC Unity supports hierarchical snapshots, meaning that a snapshot can itself be used as the source for another snapshot. This capability is useful when preserving multiple recovery points in a dependent snapshot tree while retaining space efficiency through Unity’s redirect-on-write snapshot architecture.

The supported hierarchy depth differs by storage object type. For block storage, Unity supports a snapshot hierarchy up to 16 levels deep for LUNs. For NAS/file storage, it supports a hierarchy up to 8 levels deep for file systems. These limits are platform design limits and should be considered when planning protection workflows that create snapshots from existing snapshots, including application-consistent recovery designs and retention schemes. Administrators should also account for snapshot capacity consumption and monitor the storage pool to ensure sufficient space is available as changed data accumulates across the snapshot hierarchy.

Dell’s Unity product documentation and technical resources describe Unity snapshot functionality, operational limits, and data-protection planning considerations. See the [Dell EMC Unity Family Information Hub](#) and the [Dell Unity Family product documentation](#).

QUESTION NO: 10

What are the two biggest factors that determine IOPS when using SSDs in a Dell EMC SC Series array?

- A.** Current latency and queue depth
- B.** Application block size and overall throughput
- C.** Multi-tenant arrays and tier configuration
- D.** Read/write percentage and peak workload

ANSWER: B

Explanation:

Application block size and overall throughput are the primary determinants of IOPS because IOPS is a rate of individual input/output operations, whereas throughput is the amount of data transferred per unit of time. The relationship can be expressed conceptually as $IOPS = \text{throughput} \div \text{average I/O size}$. Consequently, a workload transferring a given amount of data with small application blocks produces far more I/O operations than one transferring the same amount with large blocks. For example, at a comparable data rate, an 8 KiB workload generates approximately eight times as many operations as a 64 KiB workload. In an SC Series SSD design, performance sizing must therefore begin with the application's observed or expected block-size profile and its required aggregate throughput, rather than treating an SSD as having one fixed IOPS value. These workload characteristics establish the operation rate the storage system must sustain and are essential inputs for estimating SSD performance and capacity requirements. Dell's SC Series documentation describes performance planning and workload considerations for Storage Center systems, while Dell's storage best-practices material emphasizes basing configuration decisions on measured application I/O characteristics. See the [Dell Storage Center documentation](#) and [Dell Storage Center best practices](#).

QUESTION NO: 11

A customer must protect file storage resources through asynchronous replication to other Dell Unity XT systems in different locations.

What is the maximum number of remote systems that can be configured?

- A. 4
- B. 24
- C. 16
- D. 8

ANSWER: A

Explanation:

4 is the maximum number of remote Dell Unity XT systems that can be configured for replication. In a Dell Unity XT replication topology, a remote system is a peer Unity XT array registered with the local system so that replication sessions can be created and managed between the arrays. This limit applies to the number of distinct remote systems configured on a Unity XT system; it is not a limit on the number of individual file-system replication sessions. A customer can therefore use up to four geographically separate Unity XT peer systems to design asynchronous file-replication protection and recovery workflows, subject to the applicable Unity OE release, platform, and replication-session limits. Asynchronous replication transfers updates on a scheduled basis, supporting recovery at a remote site without requiring each write to be acknowledged by the destination before the host operation completes. Dell's Unity XT documentation contains the replication planning and system-limit guidance used when designing these configurations. See the [Dell Unity XT documentation library](#) and the [Dell Unity XT Info Hub](#) for current configuration and replication information.

QUESTION NO: 12

When configuring Fibre Channel host connectivity to a Dell Unity XT LUN, which paths are considered optimized for host I/O?

- A. Paths through the Storage Processor that owns the LUN
- B. Paths through either Storage Processor, regardless of LUN ownership
- C. Paths using only the management Ethernet ports
- D. Paths through the Storage Processor with the lowest CPU utilization

ANSWER: A

Explanation:

Paths through the Storage Processor that owns the LUN is correct. Dell Unity XT uses Asymmetric Logical Unit Access (ALUA) to identify the preferred paths for a block resource. The Storage Processor that owns a LUN provides the optimized paths to that LUN. Hosts and supported multipathing software should direct normal I/O through these paths to avoid unnecessary inter-SP processing.

Paths through the peer Storage Processor remain available for high availability, but they are non-optimized for a LUN owned by the other SP. If I/O is received on a non-optimized path, Unity can process the request through its internal inter-SP communication architecture. Correct host multipathing configuration recognizes ALUA path states and preferentially uses optimized paths. See the [Dell Unity XT documentation portal](#) for host-connectivity and multipathing guidance.

QUESTION NO: 13

How many PCI-E lanes are provided for Slot 0 and Slot 1 in Dell EMC Unity XT arrays?

- A. 8 for Slot 0; 8 for Slot 1
- B. 16 for Slot 0; 16 for Slot 1
- C. 8 for Slot 0; 16 for Slot 1
- D. 16 for Slot 0; 8 for Slot 1

ANSWER: D

Explanation:

Dell EMC Unity XT arrays provide 16 PCIe lanes to Slot 0 and 8 PCIe lanes to Slot 1 on each storage processor. Thus, “16 for Slot 0; 8 for Slot 1” is correct. The differing lane widths are an intentional part of the Unity XT storage-processor hardware design: Slot 0 has a x16 PCIe electrical connection, while Slot 1 has a x8 electrical connection. PCIe lane count is important because it defines the maximum host-interface bandwidth available to an I/O module installed in that slot. When planning connectivity expansion, architects should account for the slot’s electrical width as well as the supported I/O-module type, host protocol, and array model. This ensures that the selected module is installed in a supported location and that its available throughput is understood in the context of the platform hardware. Dell’s Unity XT technical documentation identifies the PCIe lane allocation and slot layout as part of the platform’s hardware specifications. See the [Dell EMC Unity XT technical guide](#) and the [Dell Unity XT documentation portal](#) for hardware documentation and supported configuration information.

QUESTION NO: 14

What is a consideration when sizing a Dell EMC SC Series system for capacity?

- A. Read/Write ratio
- B. IOPS
- C. Page size
- D. RAID overhead

ANSWER: D

Explanation:

RAID overhead is a key consideration when sizing an Dell EMC SC Series system for capacity because the raw capacity installed in drives is not the same as the usable capacity available to hosts. SC Series storage uses RAID protection to provide fault tolerance, and the selected RAID type consumes a portion of physical capacity for parity or mirrored data. The

amount consumed varies with the protection scheme and disk configuration, so capacity planning must calculate usable capacity after this overhead rather than relying on raw-drive totals. A sound sizing exercise also allows for the expected data footprint, growth, snapshots, replication where applicable, and operational free-space requirements; however, RAID overhead is the direct capacity-efficiency factor identified here. Dell's SC Series documentation is available through the [Dell Support documentation portal for Compellent/SC Series products](#). Dell's storage resource library also provides product documentation and planning resources at the [Dell Technologies data storage page](#).

QUESTION NO: 15

A storage administrator is trying to create a new hybrid pool on their Dell Unity XT 680. They want to add

- 18x 3.84 TB SSD drives
- 72x 1200 GB 10K SAS drives
- 64x 12 TB NL-SAS drives

However, this action is not being permitted by the Unity software. What is the reason for this?

- A. 3.84 TB SSD are only supported in all flash pools
- B. There are too many disks in the SAS tier
- C. 12 TB NL-SAS drives are only permitted in archive pools
- D. The pool must be created with only a single drive type

ANSWER: A

Explanation:

"3.84 TB SSD are only supported in all flash pools" is correct. Dell Unity XT validates the drive configuration during pool creation against the platform's supported drive and pool-layout rules. In this configuration, the 3.84 TB SSD type cannot be used as the Flash tier of a hybrid pool. Those drives are supported when the pool is configured as an all-flash pool, in which every drive in the pool is an SSD. Because the proposed pool also includes 10K SAS and 12 TB NL-SAS drives, Unity classifies it as a hybrid pool and rejects the configuration before the pool can be created. The administrator must either create an all-flash pool using the supported SSD configuration or select an SSD type supported for the Flash tier of a hybrid pool, subject to the Unity OE version and installed drive-support rules. Pool and drive support can vary by Unity OE release, so the installed system's release-specific documentation and compatibility information should be consulted before changing the design. Dell's Unity XT documentation is available from the [Dell Unity XT 680 support documentation page](#), and Dell provides current platform information at the [Dell Unity XT family page](#).

QUESTION NO: 16

What is the benefit of performing a Dell Unity OE software upgrade using the supported upgrade workflow?

- A. It enables a nondisruptive upgrade of the Storage Processors
- B. It converts a hybrid array into an all-flash array
- C. It automatically replaces failed disk drives
- D. It requires all host access to be permanently removed

ANSWER: A

Explanation:

It enables a nondisruptive upgrade of the Storage Processors is correct. Dell Unity OE upgrades are designed to be nondisruptive when the array, host connectivity, and multipathing configuration meet the applicable requirements. The

upgrade process updates the Storage Processors in a controlled sequence so that host I/O can continue through the available SP and paths while the peer SP is updated.

This behavior depends on properly configured redundant host paths and supported multipathing. An upgrade does not require recreating pools, deleting storage resources, or manually moving every LUN to another array. Dell provides release-specific upgrade requirements and procedures through the [Dell Unity Family documentation portal](#).

QUESTION NO: 17

Which Dell EMC Unity storage array models support NVMe?

- A. 380, 880, and 880F
- B. 680F, 880, and 880F only
- C. 680, 680F, 880, and 880F
- D. 380, 380F, and 480F
- E. 380F, 480F, 680F, and 880F

ANSWER: E

Explanation:

Dell EMC Unity XT NVMe support applies to the all-flash Unity XT platforms: 380F, 480F, 680F, and 880F. These systems support NVMe SSD configurations, enabling the lower-latency, higher-parallelism characteristics of NVMe flash media for workloads requiring high transaction performance and responsive storage. The model suffix "F" identifies an all-flash Unity XT system; the hybrid models are not included in the supported NVMe model list. Therefore, the complete and accurate answer must name all four supported all-flash models: 380F, 480F, 680F, and 880F. This is consistent with Dell's Unity XT platform documentation and product specifications, which distinguish the all-flash configurations and their supported drive technologies. Consult Dell's [Unity XT product information](#) and the [Dell Unity XT documentation portal](#) for the applicable hardware and software release documentation.

QUESTION NO: 18

How are vVol datastores mounted in Dell EMC PowerStore storage during the creation of vVol datastores in the vSphere environment shared by ESXi hosts?

- A. Volume Groups
- B. Storage containers
- C. Storage Groups
- D. VMFS datastores

ANSWER: B

Explanation:

Storage containers is correct. In a VMware vVol environment, the storage array presents capacity to vSphere as a storage container through the vASA provider. vCenter Server discovers this container and mounts it as a vVol datastore, which ESXi hosts can then use for virtual-machine storage. The vVol datastore is therefore the vSphere-side representation of array-provided storage capacity rather than a conventional VMFS filesystem created on an individual block volume. PowerStore supports vVols by registering its vASA provider with vCenter and exposing a storage container for the PowerStore cluster. During vVol datastore creation, an administrator selects the discovered PowerStore storage container; vSphere creates and mounts the corresponding vVol datastore for the participating hosts. This integration also lets PowerStore provide policy-based storage capabilities to VMware through VMware Storage Policy Based Management. Dell describes PowerStore vVol configuration and storage-container presentation in its [PowerStore VMware Virtual Volumes documentation](#). VMware also

explains that a vVol datastore represents a storage container supplied by a storage provider in the [vSphere Storage documentation](#).