DUMPSBOSS.

Dell PowerFlex Design 2023 Exam

EMC D-PWF-DS-23

Version Demo

Total Demo Questions: 6

Total Premium Questions: 61

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, PowerFlex Architecture	23
Topic 2, PowerFlex Design	16
Topic 3, PowerFlex Deployment	22
Total	61

QUESTION NO: 1

An architect is deploying an MDM cluster across three separate nodes to provide high availability and quorum protection. Which role should be assigned to the third MDM node after the Primary and Secondary roles have been assigned?

A. SDC

B. SDT

C. Tie-breaker

D. File Storage Node

ANSWER: C

Explanation:

The third node is assigned the tie-breaker role. In a three-node MDM configuration, the Primary MDM and Secondary MDM provide the management service and metadata resiliency, while the tie-breaker participates in quorum decisions. Placing the roles on separate nodes helps the cluster maintain a supported and resilient management-plane configuration.

QUESTION NO: 2

For what reason would an administrator choose to set the Force Clean SDS option to YES when adding devices to an SDS In the PowerFlex system?

A. Ensure that the device Is error-free and compatible with PowerFlex

B. Bypass restrictions and proceed with adding The device

C. Perform a clean check on the device before adding II

D. Overwrite existing data on the device during the addition process

ANSWER: D

Explanation:

“Overwrite existing data on the device during the addition process” is correct because Force Clean authorizes PowerFlex to clean a storage device before it is accepted by an SDS. A device being introduced into a PowerFlex storage pool can contain residual signatures, metadata, partition information, or data from an earlier deployment or another use. Enabling Force Clean instructs the add-device operation to remove the existing device content or metadata that would prevent the device from being initialized for PowerFlex use. This prepares the device for PowerFlex to write its own required layout and metadata as part of adding capacity to the SDS.

Administrators should use this setting deliberately, particularly for repurposed drives, because the cleaning operation is destructive to pre-existing information on the selected device. It is therefore appropriate only after confirming that the device is the intended target and that any required data has been retained elsewhere. This behavior supports a consistent and known device state before PowerFlex incorporates the device into its distributed storage configuration. Dell’s PowerFlex documentation resources and technical overview describe SDS device management and the preparation of storage devices for PowerFlex: Dell PowerFlex documentation and Dell PowerFlex Info Hub.

QUESTION NO: 3 - (SIMULATION)

Place the steps to set up remote replication on the Powerflex system in the correct order

ANSWER: See the explanation for the answer

Explanation:

Correct order for setting up PowerFlex remote replication:

1. Install the target PowerFlex system.
2. Exchange root certificates between the source and target systems.
3. Add the target system as a *Replication Peer.*
4. Configure the replication volumes.

The target environment must exist before it can receive replicated data. Certificate exchange establishes trusted secure communication, then the systems are paired as replication peers. Finally, configure the source and target volumes and replication settings.

QUESTION NO: 4

An administrator wants to configure SDTs to perform I/O with hosts that are going to use the NVMe over Fabric protocol What are three requirements to accomplish this task (Select 3)

A. The Storage port must be configured with the host role

B. The Discovery port must be configured with the storage role

C. The NVMe Qualified Name of the hosts is known

D. A minimum of two NVMe targets must be assigned to a Protection Domain.

E. The storage rote and hosts role must be configured to the same target IP addresses

ANSWER: A C D

Explanation:

Configuring SDTs for host I/O over NVMe over Fabrics requires that the storage-facing connectivity be prepared for host access, that each host can be uniquely identified, and that the NVMe target design provides availability within the protection domain. The requirement that *The Storage port must be configured with the host role* ensures the port is designated for the NVMe-oF host I/O path. The requirement that *The NVMe Qualified Name of the hosts is known* is essential because an NQN is the NVMe-defined unique identifier used to identify an NVMe host during access configuration and connectivity. Administrators use the host NQN when defining the host access relationship to NVMe namespaces/volumes. Finally, *A minimum of two NVMe targets must be assigned to a Protection Domain.* provides the target redundancy needed for highly available host access; hosts can maintain multiple paths to the storage service rather than relying on a single target. These configuration elements align the PowerFlex NVMe-oF target deployment with NVMe host identification and multipath-oriented availability practices. Dell describes PowerFlex architecture and deployment resources at Dell PowerFlex, while the NVMe specification defines NQN-based subsystem and host identification in the NVM Express Base Specification.

QUESTION NO: 5

An administrator wants to migrate a volume from one storage pool to another storage pool What two volume migrations are possible ?(Select 2)

A. From MG storage pool volume, non-zero padded and thick provisioned to FG storage pool volume, zero padded and thin provisioned

B. From MG storage pool volume, zero padded, and thick provisioned to FG storage pool volume, zero padded, and thin provisioned

C. From FG storage pool volume, zero padded, and thin provisioned to MG storage pool volume, non-zero padded and thick provisioned

D. From MG storage pool volume, non-zero padded, and thin provisioned to MG storage pool volume, zero padded, and thin provisioned

ANSWER: B D

Explanation:

PowerFlex volume migration can move a volume to a compatible destination storage pool without interrupting host access. A supported migration is from an MG storage pool volume that is zero padded and thick provisioned to an FG storage pool volume that is zero padded and thin provisioned. This supported conversion is useful when transitioning a workload to Fine Granularity storage while retaining zero padding and changing the destination provisioning method to thin provisioning. Fine Granularity pools provide more efficient allocation behavior and support the characteristics required by the destination volume in this migration scenario.

A migration from an MG storage pool volume that is non-zero padded and thin provisioned to an MG storage pool volume that is zero padded and thin provisioned is also supported. This is an MG-to-MG migration and preserves thin provisioning while allowing the destination volume to use zero padding. PowerFlex validates volume and pool compatibility during migration, and the migration workflow maintains the volume's host-facing availability while data is moved to the destination pool. These supported paths allow administrators to rebalance capacity, adopt a different pool configuration, and alter supported volume attributes as part of the move. See the Dell PowerFlex Family Information Hub and the Dell PowerFlex documentation library for the applicable configuration and migration guidance.

QUESTION NO: 6

An architect das configured a PowerFlex solution to use a tine granularity storage pool based on a customer's Initial request After validating the design against a LiveOptlcs output they modified the granularity of the configuration to medium What did the architect accomplish with this change'

A. Improved performance

B. Belter distribution of data blocks

C. Increased effective capacity

ANSWER: A

Explanation:

Improved performance is correct. In PowerFlex, storage-pool granularity determines the data-layout and allocation characteristics used by volumes in that pool. Moving a design from fine granularity to medium granularity changes it toward the layout intended for workloads that prioritize I/O performance and low latency. Medium-granularity pools use a larger allocation unit than fine-granularity pools, reducing allocation and metadata-processing overhead for typical larger-block and performance-sensitive I/O patterns. This makes medium granularity a suitable design choice when the workload analysis from LiveOptics indicates that responsiveness, throughput, or latency is more important than maximizing the space-efficiency features associated with fine granularity. The architect therefore used the validated workload profile to align the pool configuration with the customer's actual performance requirements. Granularity is selected when the storage pool is created, so this is an architectural design decision that ensures the appropriate pool type is deployed before volumes and application workloads are provisioned. Dell describes PowerFlex as software-defined block storage designed to deliver high and predictable performance at scale; see Dell PowerFlex family and the Dell PowerFlex documentation library.