



QUESTION & ANSWER

HIGHER QUALITY, BETTER SERVICE

Provide One Year Free Update!

<https://www.passquestion.com>

Exam : **D-PVM-OE-01**

Title : **Dell PowerMax Operate v2
Exam**

Version : **DEMO**

1.What are two characteristics of a SnapVX Linked Target?

- A. They are read writable (R/W).
- B. They support a maximum of 1024 snaps per volume
- B. They cannot send data to SRDF
- C. They can be made secure

Answer: AC

Explanation:

Step by Step Comprehensive Detailed

SnapVX Linked Targets: SnapVX is a snapshot technology used in Dell PowerMax storage arrays. A linked target is a volume that provides read/write access to a specific point-in-time copy (snapshot) of a source volume.

Read/Write Access: Unlike traditional snapshots, which are typically read-only, SnapVX linked targets allow modifications. This makes them suitable for use cases like testing, development, and data analysis where changes need to be made to a copy of the data without affecting the original source.

Secure Snapshots: SnapVX offers the capability to create "secure snapshots." These snapshots are write-protected and prevent any modifications or deletion, ensuring data integrity and protection against accidental or malicious changes.

Why other options are incorrect:

B. They support a maximum of 1024 snaps per volume: This is partially correct. While a source volume can have up to 1024 snapshots, this limit includes all types of snapshots (manual, automated, and legacy), not just linked targets.

B. They cannot send data to SRDF: This is incorrect. SnapVX linked targets can participate in SRDF (Symmetrix Remote Data Facility) replication, allowing for disaster recovery and data mobility.

Reference and documents of Dell's public documentation for PowerMax Operate v.2:

Dell Solutions Enabler 10.0.0 TimeFinder SnapVX CLI User Guide: This guide provides detailed information about SnapVX features and commands, including how to create, link, and manage snapshots. It confirms the read/write capability of linked targets and the ability to create secure snapshots. You can find this document on the Dell Support website by searching for "Solutions Enabler TimeFinder SnapVX CLI User Guide."

Dell PowerMax Family: Essentials and Best Practices Guide: This guide offers a comprehensive overview of PowerMax technologies, including SnapVX. It highlights the benefits of SnapVX linked targets for various use cases. You can find this document on the Dell Support website by searching for "PowerMax Family Essentials and Best Practices Guide."

2.Which three device types can be managed using Solutions Enabler and Unisphere?

- A. SRDF Thin Devices (RDF1 or RDF2)
- B. Thin BCV Devices (BCV+TDEV)
- C. Internal Thin Devices (Int+TDEV)
- D. Data Devices (TDATs)
- E. Thin Devices (TDEV)

Answer: ACE

Explanation:

Step by Step Comprehensive Detailed

Dell PowerMax storage arrays utilize different device types for various purposes. Solutions Enabler

(SYMCLI) and Unisphere for PowerMax are management tools that can interact with these device types. Here's a breakdown:

SRDF Thin Devices (RDF1 or RDF2): These devices are specifically used for SRDF (Symmetrix Remote Data Facility) replication. RDF1 devices represent the local copy of data in an SRDF relationship, while RDF2 devices represent the remote copy. Both Solutions Enabler and Unisphere can manage these devices to configure and monitor SRDF replication.

Internal Thin Devices (Int+TDEV): These are thin provisioned devices that reside within the PowerMax storage array. They are used for general storage purposes and can be managed by both Solutions Enabler and Unisphere for tasks like provisioning, allocating capacity, and monitoring performance.

Thin Devices (TDEV): This is a general term for thin provisioned devices in PowerMax. Thin provisioning allows for efficient storage utilization by allocating capacity on demand. 1 Both Solutions Enabler and Unisphere can manage these devices.

[https://www.n-able.com/blog/thin-provision-vs-thick-](https://www.n-able.com/blog/thin-provision-vs-thick-provision#:~:text=Thin%20provisioning%20allocates%20disk%20space,need%20at%20any%20given%20time.)

[provision#:~:text=Thin%20provisioning%20allocates%20disk%20space,need%20at%20any%20given %20time.](https://www.n-able.com/blog/thin-provision-vs-thick-provision#:~:text=Thin%20provisioning%20allocates%20disk%20space,need%20at%20any%20given%20time.)

[https://www.n-able.com/blog/thin-provision-vs-thick-](https://www.n-able.com/blog/thin-provision-vs-thick-provision#:~:text=Thin%20provisioning%20allocates%20disk%20space,need%20at%20any%20given%20time.)

[provision#:~:text=Thin%20provisioning%20allocates%20disk%20space,need%20at%20any%20given %20time.](https://www.n-able.com/blog/thin-provision-vs-thick-provision#:~:text=Thin%20provisioning%20allocates%20disk%20space,need%20at%20any%20given%20time.)

Why other options are incorrect:

B. Thin BCV Devices (BCV+TDEV): BCV (Business Continuance Volume) devices are used for creating point-in-time copies for disaster recovery. While Solutions Enabler can manage BCV devices, Unisphere for PowerMax has limited functionality for managing them directly.

D. Data Devices (TDATs): TDATs are physical devices within the PowerMax array. While Solutions Enabler can interact with TDATs at a lower level, Unisphere for PowerMax primarily focuses on managing logical devices and storage groups.

Reference and documents of Dell's public documentation for PowerMax Operate v.2:

Dell PowerMax Family: Essentials and Best Practices Guide: This guide provides an overview of PowerMax devices and their management. It mentions the different device types and how they are used in the PowerMax environment.

Dell Solutions Enabler 10.0.0 CLI User Guide: This guide provides detailed information about Solutions Enabler commands for managing various device types, including SRDF devices, thin devices, and internal devices.

Dell Unisphere for PowerMax 10.0.0 Online Help: The online help documentation for Unisphere for PowerMax explains how to manage different device types through the graphical user interface, including provisioning, monitoring, and configuring storage.

3.DRAG DROP

Place the steps in the correct order for performing a Non-Disruptive Migration

Steps

Initiate data synchronization from the source to the target and cutover to make source array paths inactive so target array services all host I/O.

Remove the configured array-to-array connectivity for the data migration pathway on both storage arrays.

Set up the migration environment to migrate applications from the source array to the target array.

Examine the specific storage for the applications on the source array and automatically provision equivalent storage on the target array.

Commit the necessary changes after the source to target data synchronization is complete and all application data has been migrated to the target.

Correct Order

Answer:

Steps

Initiate data synchronization from the source to the target and cutover to make source array paths inactive so target array services all host I/O.

Remove the configured array-to-array connectivity for the data migration pathway on both storage arrays.

Set up the migration environment to migrate applications from the source array to the target array.

Examine the specific storage for the applications on the source array and automatically provision equivalent storage on the target array.

Commit the necessary changes after the source to target data synchronization is complete and all application data has been migrated to the target.

Correct Order

Set up the migration environment to migrate applications from the source array to the target array.

Examine the specific storage for the applications on the source array and automatically provision equivalent storage on the target array.

Initiate data synchronization from the source to the target and cutover to make source array paths inactive so target array services all host I/O.

Commit the necessary changes after the source to target data synchronization is complete and all application data has been migrated to the target.

Remove the configured array-to-array connectivity for the data migration pathway on both storage arrays.

Explanation:

Set up the migration environment to migrate applications from the source array to the target array.
Examine the specific storage for the applications on the source array and automatically provision

equivalent storage on the target array.

Initiate data synchronization from the source to the target and cutover to make source array paths inactive so target array services all host I/O.

Commit the necessary changes after the source to target data synchronization is complete and all application data has been migrated to the target.

Remove the configured array-to-array connectivity for the data migration pathway on both storage arrays.

Correct Order:

Set up the migration environment to migrate applications from the source array to the target array.

Why: This is the initial setup phase, where you configure the necessary settings on both the source and target arrays to enable the migration.

This involves actions like:

Verifying compatibility between the source and target arrays.

Ensuring that the required licenses are in place (e.g., NDM license).

Configuring network connectivity (FC or iSCSI) for data transfer between the arrays.

Examine the specific storage for the applications on the source array and automatically provision equivalent storage on the target array.

Why: Before migrating data, you need to ensure that the target array has the appropriate storage capacity and configuration to host the applications.

How: NDM can often automatically provision equivalent storage on the target based on the source configuration. This includes creating storage groups, volumes, and masking views that mirror the source.

Initiate data synchronization from the source to the target and cutover to make source array paths inactive so target array services all host I/O.

Why: This is the core of the migration process. Data is copied from the source to the target while the application continues to run. Cutover is the final step where I/O is redirected to the target array.

How:

Synchronization: Data is copied in the background.

Cutover: Once synchronization is complete, a brief cutover is performed. In a well-planned NDM, this cutover is designed to be within the I/O timeout limits of most applications.

Commit the necessary changes after the source to target data synchronization is complete and all application data has been migrated to the target.

Why: This step finalizes the migration and makes it permanent.

What it involves: The migration session is acknowledged and the configuration is finalized on the target array.

Remove the configured array-to-array connectivity for the data migration pathway on both storage arrays.

Why: After the migration is complete, the temporary connections used for data transfer between the arrays should be removed to free up resources and maintain a clean configuration.

What it involves: This typically means removing the FC zones or iSCSI settings that were configured specifically for the NDM process.

4.What are two characteristics of a SnapVX Clone?

- A. Maximum 1024 snaps per volume
- B. Can be made secure

- C. Restores directly to the source volume
- D. Is Read-only
- E. Can be made crash consistent

Answer: BE

Explanation:

Step by Step Comprehensive Detailed

SnapVX Clones: SnapVX clones are full, writable copies of a source volume created using the SnapVX snapshot technology. They are independent volumes that can be used for various purposes, such as testing, development, or data analysis.

Secure Snapshots: SnapVX offers the capability to create "secure snapshots." When a clone is derived from a secure snapshot, it inherits the same protection, making it immutable and preventing any modifications or deletion. This ensures data integrity and protection against accidental or malicious changes.

Crash Consistent: SnapVX clones can be made crash consistent. This means that the clone captures a point-in-time copy of the source volume that is consistent with a database or application crash. This is important for ensuring data integrity and recoverability in situations where the source volume experiences an unexpected outage.

Why other options are incorrect:

- A. Maximum 1024 snaps per volume: This limit applies to the source volume, not the clones themselves. Each clone is an independent volume.
- C. Restores directly to the source volume: Clones are independent copies and do not directly restore to the source volume. Data can be copied or moved from the clone to the source if needed.
- D. Is Read-only: SnapVX clones are fully writable copies, not read-only.

Reference and documents of Dell's public documentation for PowerMax Operate v.2:

Dell Solutions Enabler 10.0.0 TimeFinder SnapVX CLI User Guide: This guide provides detailed information about SnapVX features and commands, including how to create and manage clones. It confirms the ability to create secure clones and the option to make them crash consistent.

Dell PowerMax Family: Essentials and Best Practices Guide: This guide offers a comprehensive overview of PowerMax technologies, including SnapVX. It highlights the benefits of SnapVX clones for various use cases.

5. Your organization is planning to expand its data center capacity and must configure a new PowerMax 2500 system. The goal is to ensure high performance and scalability while maintaining redundancy. Which configuration option is specific to the PowerMax 2500 system?

- A. Supports up to 16 nodes
- B. Supports up to 1.5 PB of effective capacity
- C. Scales from one to a maximum of two node pairs
- D. Uses Storage Class Memory

Answer: C

Explanation:

Step by Step Comprehensive Detailed

The Dell PowerMax 2500 is a mid-range storage array designed for enterprise environments. It offers a balance of performance, capacity, and scalability.

Scalability: The PowerMax 2500 scales from a single node pair to a maximum of two node pairs. This

allows for increased performance and capacity as needed. Each node pair provides processing power, cache memory, and connectivity.

Why other options are incorrect:

A. Supports up to 16 nodes: This is incorrect. The PowerMax 8500, the higher-end model, supports up to 16 nodes.

B. Supports up to 15 PB of effective capacity: While the PowerMax 2500 offers significant capacity, its maximum effective capacity is lower than 15 PB.

D. Uses Storage Class Memory: Both the PowerMax 2500 and 8500 utilize Storage Class Memory (SCM) for enhanced performance.

Reference and documents of Dell's public documentation for PowerMax Operate v.2:

Dell PowerMax 2500 Hardware Information Guide: This guide provides detailed technical specifications for the PowerMax 2500, including its scalability options, node configurations, and capacity limits. You can find this document on the Dell Support website by searching for "PowerMax 2500 Hardware Information Guide."

Dell PowerMax Family: Essentials and Best Practices Guide: This guide offers a general overview of the PowerMax family, including the 2500 and 8500 models. It highlights the key differences in scalability and performance between the two models.