

Dell Storage Networking Professional - version 2

Dell DC0-261

Version Demo

Total Demo Questions: 13

Total Premium Questions: 136

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

QUESTION NO: 1

You are selecting a disk to be used as a hot-spare replacement. Which disk IDs are exempt from use on an FC4700?

- A. 0:0 - 0:8
- B. 0:0 - 0:2
- C. 0:0 - 0:4
- D. Any disk can be used as a hot-spare replacement.

ANSWER: C

Explanation:

0:0 - 0:4 is correct because these are the first five disks in the FC4700 disk processor enclosure and are reserved as the array's vault disks. Vault disks hold critical system information, including write-cache vaulting data and operating/system metadata needed to preserve data integrity during an unexpected power event and to restore the storage system safely. Because these disks have this essential platform role, they are not eligible to be selected or assigned as hot spares. When choosing a hot spare for an FC4700, select an eligible non-vault disk with compatible characteristics for the protected drives, such as the appropriate drive technology, capacity, and speed class. This protects the availability of spare capacity while preserving the dedicated vault-disk layout required by the array. Dell's legacy CLARiiON documentation resources are available through the [Dell Support documentation page for the CLARiiON FC4700](#). Dell's [vault-drive identification knowledge-base article](#) also describes the purpose of vault drives in CLARiiON/VNX architectures.

QUESTION NO: 2

Which joins multiple switches in a Storage Area Network (SAN)?

- A. trunking
- B. SAN routing
- C. cascading
- D. zoning
- E. LUN masking

ANSWER: C

Explanation:

cascading is the practice of interconnecting multiple SAN switches through inter-switch links (ISLs) so they operate together as a larger switched fabric. This allows hosts and storage attached to different switches to communicate across the fabric, while also expanding the available port count beyond what a single switch can provide. In Fibre Channel SAN environments, the interconnected switches exchange fabric information and maintain connectivity information needed to forward traffic to devices on other switches. A cascaded topology can also be designed with redundant ISLs and multiple paths to improve availability and resilience. The term therefore directly describes joining switches into a common SAN fabric, rather than a host-access or storage-access control feature. Cisco's SAN switching documentation describes ISLs as the links used to connect Fibre Channel switches and form a fabric: [Cisco Nexus SAN Switching Configuration Guide](#). Broadcom's Fabric OS administration documentation likewise covers switch-to-switch fabric connectivity and ISL configuration: [Fabric OS Administration Guide](#).

QUESTION NO: 3

Along with what software is the Navisphere Host Agent installed?

- A. PowerPath
- B. Navisphere Manager
- C. CLI
- D. AccessLogix

ANSWER: C

Explanation:

The Navisphere Host Agent is installed along with the Navisphere CLI software. In the CLARiiON/Navisphere management architecture, the host agent runs on a connected host and communicates host configuration information to the storage system's management software. This information can include the host name, operating system details, initiators, and logical-unit visibility data. It also supports host registration and provides the storage array with the information needed to associate host connectivity correctly during configuration and management operations.

Installing the host agent as part of the CLI package is practical because both components are host-side administration tools. The CLI supplies command-line access for array-management tasks, while the host agent supplies the local host-to-array communication service used by Navisphere management functions. Dell EMC CLARiiON documentation and support resources identify Navisphere as the management environment for these legacy arrays and provide product documentation for supported management utilities. See the [Dell Support documentation page for CLARiiON CX4 systems](#) and the [Dell Technologies Support portal](#) for legacy Navisphere documentation and downloads.

QUESTION NO: 4 - (DRAG DROP)

DRAG DROP

Click the Task button. Place the steps for powering down a Dell / EMC storage array in the correct order.

Select and Place:

Place the steps for powering down a Dell / EMC storage array in the correct order.


Steps	
place here	power down all DAEs
place here	stop I/O, power down host servers
place here	power down switches
place here	power down DPE/SPE using the SPS switch

Order		
First	Second	Third
Fourth		

Done

ANSWER:

Place the steps for powering down a Dell / EMC storage array in the correct order.



Explanation:

The correct shutdown sequence preserves data integrity and ensures that dependent infrastructure remains available until it is no longer needed. First, stop all application I/O and power down the host servers. This prevents hosts from issuing new read or write requests while the storage environment is being taken offline. The shutdown should occur only after normal host and application procedures have completed, so cached writes and outstanding I/O have been safely handled. Next, power down the switches. Once the hosts are off and no longer require SAN connectivity, the switching fabric can be shut down without interrupting active storage access. The disk-array enclosures are then powered down. At that point, no hosts or fabric components need access to the drives, so the DAEs can be safely removed from service. Finally, power down the DPE/SPE by using the SPS switch. The DPE/SPE contains the storage processors and is the core controller component of the array; leaving it until last maintains orderly control of the storage system throughout the preceding shutdown actions. This dependency-based sequence—hosts and I/O first, network infrastructure next, drive enclosures after that, and the processor enclosure last—is the safe operational method for a complete planned array shutdown. Dell documentation provides product-specific power-down procedures and safety guidance through its [Dell Unity support documentation](#) and the [Dell VNX power-down knowledge-base guidance](#).

QUESTION NO: 5

LUN 3 in RAID Group 2 is no longer needed. When LUN 3 was deleted, RAID Group 2 disappeared. What is the most likely cause of this?

- A. The RAID Group is automatically deleted with the last LUN is destroyed.
- B. The Automatically Destroy After Last LUN Unbound checkbox was selected when the RAID Group was created.
- C. The Move Unbound RAID Groups to the Unbound Storage Group checkbox was selected when the RAID Group was created.
- D. Upon the deletion of a LUN, a system level defrag executes, which deletes any unoccupied RAID Groups.

ANSWER: B

Explanation:

The Automatically Destroy After Last LUN Unbound checkbox was selected when the RAID Group was created. On Dell EMC VNX/CLARiiON block-storage systems, a RAID Group can be configured with an automatic-destruction attribute. When this setting is enabled, destroying or unbinding the final LUN associated with that RAID Group causes the array to automatically destroy the now-empty RAID Group. Its disks are then released from that RAID Group and become available for subsequent storage provisioning operations. This behavior is intentional and is controlled by the RAID Group configuration selected at creation time; it prevents an empty RAID Group from remaining in the system after its last LUN has been removed. The observation that RAID Group 2 disappeared immediately after LUN 3, its final LUN, was deleted directly

matches this configured automatic-destruction behavior. Dell's VNX documentation and product support resources describe RAID Group and LUN provisioning attributes, including management of RAID Groups after LUNs are unbound or destroyed. See the [Dell EMC VNX Series documentation page](#) and the [Dell Support knowledge base](#) for VNX storage-management guidance.

QUESTION NO: 6

Which three statements are correct about CHAP authentication for iSCSI? (Choose three.)

- A. One-way CHAP authenticates the initiator to the target.
- B. Mutual CHAP authenticates both the initiator and the target.
- C. CHAP uses a shared secret for authentication.
- D. CHAP encrypts iSCSI payload data.
- E. CHAP eliminates the need for iSCSI target discovery.

ANSWER: A B C

Explanation:

CHAP provides an authentication mechanism for iSCSI sessions. One-way CHAP authenticates the initiator to the target, while mutual CHAP authenticates both the initiator and the target. CHAP uses a shared secret during its challenge-response process, so matching authentication credentials must be configured on the applicable initiator and target. CHAP authenticates the endpoints but does not encrypt iSCSI payload data; IPsec or another appropriate protection mechanism is needed when encryption in transit is required. The [CHAP RFC](#) describes the challenge-handshake authentication process.

QUESTION NO: 7

When is the RAID type of a RAID Group defined?

- A. when the LUN is introduced to the OS
- B. when the first LUN is bound within the RAID Group
- C. when the RAID Group is first created
- D. when the RAID Group is placed into a Storage Group
- E. during the first LUN bind, when the RAID type is selected

ANSWER: B E

Explanation:

The RAID type is defined when the first LUN is bound within the RAID Group. In the traditional Dell EMC CLARiON/VNX RAID Group model, creating the RAID Group initially associates a set of physical disks; the RAID layout is selected as part of the first LUN-binding operation. That initial bind establishes the RAID Group's RAID configuration, which then governs the organization, protection characteristics, and usable capacity of the disks in that group. Subsequent LUNs bound to that same RAID Group use its already-established RAID type rather than redefining it.

"During the first LUN bind, when the RAID type is selected" is also correct because it states the same configuration point more explicitly. It is not a later host-presentation or Storage Group action: those steps control access and visibility, whereas LUN binding is the operation that commits the RAID layout for the RAID Group. Dell documentation and support resources for VNX terminology and RAID Group administration can be found through the [Dell Support search for VNX RAID Group LUN binding](#) and the [Dell VNX product documentation library](#).

QUESTION NO: 8

Which SP port must be available when setting up a MirrorView connection on a CX700?

- A. 2
- B. 0
- C. 3
- D. 1
- E. any unused port

ANSWER: C

Explanation:

3 is the required available Storage Processor port for configuring a MirrorView connection on a CLARiiON CX700. In this platform's supported port-assignment design, SP port 3 is reserved for the inter-array Fibre Channel connectivity used by MirrorView. The MirrorView configuration process relies on this dedicated port so that the storage processors can establish the required remote mirror connection between the participating CX-series arrays. Before creating the MirrorView connection, verify that port 3 is present, operational, and not allocated to a conflicting host-facing or other purpose. This reservation helps ensure the replication path has the expected controller-port configuration and supports reliable MirrorView operation. Dell's legacy CLARiiON CX700 product documentation is available from the [Dell Support documentation page for CLARiiON CX700](#). For access to platform documentation, advisories, and support resources, use the [Dell Technologies Support portal](#).

QUESTION NO: 9

Which three statements are correct about zoning with Brocade switches? (Choose three.)

- A. Zoning configuration information is stored on the hosts.
- B. Zone aliases are not required to define a zone.
- C. A zone can be included in more than one zone configuration.
- D. Each zone must be given a unique name.
- E. Each device can be included in only one zone.

ANSWER: B C D

Explanation:

Zone aliases are not required to define a zone because Brocade Fabric OS supports defining zone members directly by their device identifiers, such as a WWN, domain/port address, or FCID. Aliases are an optional administrative convenience that makes zoning configurations easier to read and maintain; they are not a prerequisite for creating a valid zone. A zone can be included in more than one zone configuration because a configuration is a named collection of zones. This enables administrators to build and save alternative zoning configurations, although only one configuration can be enabled in a fabric at a time. Each zone must be given a unique name because the zoning database uses the zone name as the object identifier within the fabric's defined configuration. Unique names allow Fabric OS to distinguish zones unambiguously when zones are added to configurations and when the active configuration is enabled. These behaviors are fundamental to Brocade zoning-object management and configuration activation. See the [Broadcom Fabric OS Zoning User Guide](#) and Dell's [Connectrix B-Series switches information hub](#) for Brocade zoning documentation and operational guidance.

QUESTION NO: 10

Which three statements are correct about RAID 6? (Choose three.)

- A. It uses distributed dual parity.
- B. It can tolerate the failure of two disks in the RAID set.
- C. It requires at least four disks.
- D. It provides data protection by mirroring every disk.
- E. It has no capacity overhead for data protection.

ANSWER: A B C

Explanation:

RAID 6 uses distributed dual parity across the member disks and can tolerate the failure of any two disks in the RAID set without data loss. A RAID 6 configuration requires at least four disks because capacity equivalent to two disks is used for parity protection. RAID 6 is commonly selected where additional fault tolerance is required, particularly for large-capacity drive configurations, although its parity calculations can introduce more write overhead than RAID types without parity.

QUESTION NO: 11

Which statement is correct about Navisphere Analyzer?

- A. Navisphere Analyzer does not require Navisphere Manager in the storage domain in order to run.
- B. The Analyzer Provider must be loaded onto each storage array you want to analyze.
- C. The Navisphere Analyzer UI does not need to be on the same storage array as the Management UI.
- D. Navisphere Analyzer can only monitor real-time performance data.

ANSWER: B

Explanation:

The Analyzer Provider must be loaded onto each storage array you want to analyze. Navisphere Analyzer collects, stores, and presents array performance statistics through a distributed architecture. The Analyzer Provider is the array-resident component responsible for gathering performance data from the specific storage processor and its resources, such as LUNs, RAID groups, ports, and cache-related activity. Because performance metrics originate on each individual array, an Analyzer Provider must be installed and enabled on every array whose statistics are to be included in analysis. This arrangement allows the Navisphere Analyzer interface to retrieve both current measurements and previously collected archival data for the selected array, supporting troubleshooting, capacity planning, and trend analysis. In practice, administrators deploy and manage this component as part of the Navisphere environment before attempting to view meaningful performance information for an array. Dell's legacy CLARiiON documentation and downloads, including Navisphere-related materials, are available through the [Dell Support documentation page for CLARiiON CX4 systems](#). Dell's broader [Support portal](#) also provides access to applicable Navisphere software, release notes, and product documentation by array model and software revision.

QUESTION NO: 12

Which two statements describe supported capabilities of MirrorView? (Choose two.)

- A. UsingMirrorView can guarantee that the source and target LUNs will be identical, because it uses synchronous writes.
- B. MirrorView can support bi-directional remote mirroring.
- C. You can useMirrorView to synchronously mirror data up to 100 Kilometers.
- D. BecauseMirrorView runs on the host, it consumes none of the storage array's processing power.

ANSWER: A B

Explanation:

Using MirrorView can guarantee that the source and target LUNs will be identical, because it uses synchronous writes. is correct in the context of MirrorView/S. MirrorView/S performs remote replication synchronously: a host write is not acknowledged until the write has been committed at both the primary and secondary storage systems. This write-order-preserving behavior maintains a consistent secondary image that is synchronized with the source while the mirror is operating normally. It is intended for applications that require a highly current remote copy and can accommodate the latency associated with synchronous replication.

MirrorView can support bi-directional remote mirroring. is also correct. MirrorView supports configurations in which each participating storage system can act as the source for one mirror and the destination for another. This enables reciprocal replication between sites, provided that the required LUNs and mirror relationships are configured independently and storage-system resources are appropriately planned. Dell's VNX documentation describes MirrorView as array-based remote replication with both synchronous and asynchronous modes, and covers the supported mirror relationship configurations. See the [Dell VNX product documentation](#) and Dell's [MirrorView overview knowledge article](#).

QUESTION NO: 13

Which four are valid Fibre Channel port types? (Choose four.)

- A. P_Port
- B. FL_Port
- C. F_Port
- D. N_Port
- E. E_Port
- F. X_Port

ANSWER: B C D E

Explanation:

FL_Port, **F_Port**, **N_Port**, and **E_Port** are the standard Fibre Channel port types represented in this question. An **N_Port** is a node-facing port on an end device, such as a host bus adapter or storage target, when it connects to the fabric. An **F_Port** is the corresponding fabric-switch port used for an **N_Port** attachment. Together, these port roles provide the normal point-to-point host or storage device connection into a switched Fibre Channel SAN.

An **FL_Port** extends fabric services to a Fibre Channel Arbitrated Loop, allowing loop-attached devices to access the switched fabric. An **E_Port** is an expansion port that connects one Fibre Channel switch to another through an inter-switch link; these links allow the fabric to scale beyond a single switch and exchange fabric-control information. Correct identification of these port roles is important for SAN design, switch configuration, and troubleshooting because the connected device type and topology determine the operational port role. Dell Connectrix Fibre Channel environments use these standard fabric concepts and port-role terminology. See Dell's [Connectrix B-Series information hub](#) and the [Fibre Channel Industry Association glossary](#) for Fibre Channel terminology and implementation resources.