

Specialist - Implementation Engineer SC Series Exam

EMC DES-1721

Version Demo

Total Demo Questions: 18

Total Premium Questions: 187

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Storage Center Architecture	35
Topic 2, Storage Center Installation	58
Topic 3, Storage Center Configuration	60
Topic 4, Storage Center Administration	22
Topic 5, Storage Center Troubleshooting	12
Total	187

QUESTION NO: 1

You are on-site to install 2 new Brocade switches with a Storage Center. By default, every new Brocade switch IP address is set to 10.77.77.77. What is the command to check the IP address of a Brocade switch?

- A. cfgShow
- B. configipaddr
- C. ipaddrshow
- D. showconfig

ANSWER: C

Explanation:

`ipaddrshow` is the Fabric OS command used to display a Brocade switch's configured IP networking information. During initial deployment, this is especially important because two newly installed switches can have the same factory-default management address. Running the command from the switch CLI lets the implementation engineer verify the currently configured Ethernet management IP address before attempting further configuration or connecting both switches to the same management network.

The displayed information is used to confirm the active management-network settings and to ensure that each switch can subsequently be assigned and reached through a unique address. This check should be performed through the serial-console session or another known working connection when the switch is still using its default address. Brocade Fabric OS documents `ipaddrshow` as a command for viewing IP address configuration, while the associated IP configuration workflow uses the corresponding IP-address configuration command. See the [Broadcom Fabric OS Command Reference](#) and the [Broadcom Fabric OS Administration documentation](#).

QUESTION NO: 2

You are deploying a new SCv2020 system with four additional SCv100 expansion enclosures. What is the spare-to-managed drive ratio?

- A. 1 spare for every drive of that type per enclosure
- B. 1 spare for every 24 managed drives of that drive type
- C. 1 spare per enclosure regardless of the drive type
- D. 1 spare for every 20 managed drives of that drive type

ANSWER: D

Explanation:

The correct spare-to-managed-drive ratio is **1 spare for every 20 managed drives of that drive type**. SC Series Storage Center calculates spare capacity by drive type, ensuring that an appropriate replacement drive is available for a failed managed drive in the relevant drive class. The ratio is applied across the Storage Center rather than being calculated as an independent spare requirement for every individual expansion enclosure. Therefore, deploying an SCv2020 with four SCv100 expansion enclosures does not alter the fundamental 1:20 spare policy; the additional enclosures increase the total managed-drive population used in the calculation. This approach provides sufficient recovery capacity while avoiding unnecessary reservation of drives as spares. Capacity planning should account for the number of managed drives of each installed drive type, because each type requires appropriate spare coverage. Dell documentation for the SCv2000/SCv2020 family and its supported expansion enclosures is available from the [Dell SC Series manuals page](#). Dell's [SC Series support documentation](#) also provides product support and configuration resources.

QUESTION NO: 3

Where are Dell Storage Manager (DSM) automated reports stored for later access?

- A. DSM database
- B. DSM file system
- C. First created volume
- D. External document management system

ANSWER: B

Explanation:

DSM automated reports are stored in the **DSM file system** for later access. When an automated report is generated, DSM produces a report file and retains it in the file-system-based report repository associated with the DSM Data Collector installation. This lets administrators return to DSM later to view, download, or manage the generated report output without needing to recreate the report definition or rely on a storage volume as the report location. The automated-report configuration determines such items as report content, schedule, and retention behavior, while the generated output is maintained by the DSM/Data Collector host's file-system repository. This approach also makes report retention and available disk capacity on the DSM host relevant operational considerations for environments that generate frequent or large reports. Dell's Storage Manager documentation is available through the [Dell Compellent/SC Series manuals portal](#). Product documentation and support resources for Dell Storage Manager are also provided through the [Dell Support portal](#).

QUESTION NO: 4

What can be done in order to facilitate the WWN Zoning and port identification in Brocade switches?

- A. Create Zones using WWPNS
- B. Create Aliases for port numbers
- C. Create Aliases for the WWPNS
- D. Create Zones using port numbers

ANSWER: C

Explanation:

Create Aliases for the WWPNS is correct because Brocade Fabric OS aliases provide meaningful, administrator-defined names for device identifiers used in a zone configuration. A WWPN is the persistent worldwide identifier for an individual Fibre Channel port, such as an HBA port or a storage-array target port. Associating a descriptive alias with each WWPN lets an engineer build and review zones using names that identify the host, adapter, storage port, and purpose rather than relying on long hexadecimal values. This makes the zoning database easier to read, validate, document, and troubleshoot while retaining the benefits of WWN-based zoning. It also helps correlate a physical or logical device port with the corresponding zoning member during deployment and incident investigation. In Fabric OS, aliases can be created and then used as members of zones; the zone configuration ultimately resolves those alias definitions to the applicable WWPN values. This approach is a standard administrative practice for maintaining clear, portable SAN zoning configurations. Broadcom's Fabric OS administration documentation describes zoning objects, including aliases and their use in zone configuration management: [Fabric OS Administration Guide](#). Brocade Fibre Channel switching information is also available from [Broadcom Fibre Channel Networking](#).

QUESTION NO: 5

In an SC8000 environment, what is a reason why Legacy Port Mode should be implemented on the front-end legacy Mode?

- A. Simplify iSCSI connectivity
- B. Improve front-end redundancy

C. Support AIX servers without NPIV

D. Increase front-end performance

ANSWER: C

Explanation:

Support AIX servers without NPIV is correct because Legacy Port Mode is intended for Fibre Channel host environments whose HBAs, drivers, switches, or host operating systems cannot use NPIV (N_Port ID Virtualization). In the normal virtual-port implementation, an SC Series system uses NPIV to present virtual front-end WWPNs. Those virtual identities can move between physical front-end ports during controller failover, allowing the host to retain a consistent target identity while the storage system maintains high availability.

Some AIX configurations, particularly those using adapters or SAN designs without NPIV support, cannot discover or log in to those virtual NPIV-based target ports. Enabling Legacy Port Mode makes the physical front-end port identities available for host connectivity instead of requiring NPIV virtual ports. This provides the compatibility needed for these AIX hosts to attach to an SC8000 system. The setting should be selected as a host-compatibility requirement during front-end connectivity planning, with zoning and multipath configuration designed around the physical port presentation.

Dell's SC Series documentation and support resources describe front-end connectivity configuration and the platform documentation applicable to SC8000 systems: [Dell SC8000 manuals and documentation](#) and [Dell Storage Center OS best-practices resources](#).

QUESTION NO: 6

You are implementing a front-end connection in Legacy Port Mode on an SC4020 in a SAN with two iSCSI switches. What is the minimum number of port(s), per switch, that can be used by each controller?

A. One port on each Ethernet switch

B. Two ports on each Ethernet switch

C. Four ports on each Ethernet switch

D. Eight ports on each Ethernet switch

ANSWER: B

Explanation:

Two ports on each Ethernet switch is the minimum supported front-end connectivity per SC4020 controller when Legacy Port Mode is used with two iSCSI switches. This results in four front-end iSCSI connections per controller overall: two connections to one switch and two connections to the other. The design provides redundant paths across the independent Ethernet fabrics and allows the host multipathing configuration to maintain access if a path, port, or switch-side connection is unavailable. Legacy Port Mode uses distinct physical-port identities and target addressing rather than the virtual-port behavior available in Virtual Port Mode, so the physical port layout must meet the Storage Center redundancy requirements. For an SC4020, distributing two controller ports to each switch also supports balanced use of the available iSCSI paths while preserving switch-level resiliency. Dell's SC4020 deployment documentation describes the supported cabling and front-end connectivity requirements for the platform, and Dell's Storage Center documentation resources provide the associated configuration guidance. See the [Dell SC4020 Storage System Deployment Guide](#) and [Dell Storage Center documentation resources](#).

QUESTION NO: 7

Which SC Series array feature allows all front-end ports to act as active data paths?

A. Dual Fabric

B. ALUA

- C. Virtual Port Mode
- D. Legacy Port Mode

ANSWER: C

Explanation:

Virtual Port Mode is the SC Series feature that presents front-end connectivity so that all relevant front-end ports can participate as active host data paths. Rather than tying a host's preferred access path to a single controller-owned physical port, Virtual Port Mode uses virtual front-end ports that can be accessed through ports on both controllers. This design enables multipathed hosts to use available paths actively while the SC Series array retains the ability to handle controller or path changes transparently. The result is active/active-style host connectivity, improved path utilization, and nondisruptive continuation of access when a controller or physical port condition requires a path transition. Host multipathing software remains important: it discovers the presented paths and applies the supported path-selection and failover behavior. Virtual Port Mode is therefore the appropriate SC Series configuration when the requirement is for all front-end ports to function as active data paths. Dell's SC Series product documentation and support resources are available through the [Dell Support documentation page for SC Series systems](#); Dell also provides SC Series technical and support information through its [SC Series Storage Support Matrix](#).

QUESTION NO: 8 - (SIMULATION)

What is the correct sequence of steps to enable MPIO on a Microsoft Windows 2012 Server?

Select "Server Roles" without selecting a role to install	STEP 1
Launch Server Manager and select "Add Roles and Features"	STEP 2
Select "Multipath I/O" within the Features window	STEP 3
Confirm the settings and click "Install"	STEP 4
Select "Role-based" or "Feature-based" installation	STEP 5

ANSWER: See the explanation for the answer

Explanation:

Correct sequence to enable/install MPIO on Windows Server 2012:

1. **Step 1:** Launch **Server Manager** and select **Add Roles and Features**.
2. **Step 2:** Select **Role-based or feature-based installation**.
3. **Step 3:** Select **Server Roles** without selecting a role to install.
4. **Step 4:** Select **Multipath I/O** in the **Features** window.
5. **Step 5:** Confirm the settings and click **Install**.

After installation, reboot the server if prompted. MPIO can then be configured from the MPIO control-panel applet, including claiming the relevant storage-device hardware IDs.

QUESTION NO: 9

Three VMware ESXi hosts require concurrent access to the same SC Series volume for a shared datastore.

Which object should be created before mapping the volume?

- A. Server cluster
- B. Storage Profile
- C. Replay Profile
- D. Fault Domain

ANSWER: A

Explanation:

A server cluster should be created for hosts that require shared access to the same volume. The server cluster object represents the collection of host server objects and allows the shared volume to be mapped consistently to every cluster member.

Mapping a shared datastore separately without using the appropriate cluster configuration can result in inconsistent host access and LUN presentation. The server cluster also provides the correct management context for a volume intended for clustered application or virtualization use.

QUESTION NO: 10

When configuring an SC Series array, why should a 4 MB data page size be selected?

- A. When systems have large files with infrequent snapshots
- B. This is the default data page size
- C. This is an invalid data page size and cannot be selected
- D. When applications have high performance needs

ANSWER: A

Explanation:

The correct selection is *When systems have large files with infrequent snapshots*. In Dell EMC SC Series storage, a data page is the allocation unit used by the Storage Center operating system to manage and move data across tiers. Selecting a larger 4 MB page size is appropriate for workloads dominated by large, sequentially accessed files, such as media repositories, large engineering files, backup data, or archival content. These workloads can benefit from fewer metadata entries and reduced page-management overhead because each allocation tracks a larger extent of data.

The snapshot qualification is important because SC Series snapshots are page based. When snapshots are created frequently, smaller page sizes can provide more granular copy-on-write behavior and can reduce the amount of unchanged data included when a page is modified. Conversely, where snapshots are infrequent, that granularity is less valuable, and the 4 MB page size better aligns with large-file capacity and sequential-I/O characteristics. Page size is therefore selected during initial configuration according to the expected workload profile and snapshot behavior, since it affects how data is allocated and protected over the life of the Storage Center. See Dell's [SC Series Storage Center OS Administrator Guide](#) and the [Dell EMC SC Series Best Practices Guide](#).

QUESTION NO: 11

An administrator wants to configure an SC Series array to use Virtual Port Mode. What is a requirement to accomplish this task?

- A. I/O Controller cards must use a single-ID firmware version

- B. Activation license must be installed
- C. NPIV must be disabled on Fibre Channel switches
- D. Fibre Channel switches must use an arbitrated loop topology
- E. Fibre Channel switches must support NPIV, and NPIV must be enabled.

ANSWER: E

Explanation:

Fibre Channel switches must support NPIV and have NPIV enabled to configure an SC Series array for Virtual Port Mode. Virtual Port Mode allows the SC Series system to present virtual Fibre Channel identities through its physical front-end ports. NPIV (N_Port ID Virtualization) is the Fibre Channel fabric capability that permits multiple virtual N_Port identities, each with its own fabric-assigned address, to operate through a single physical switch connection. This capability is essential for the array's virtual ports to log in to the fabric and be recognized correctly by connected hosts and zoning configurations.

Before enabling the feature, the administrator should verify that every switch in the relevant SAN fabric supports NPIV, enable NPIV according to the switch vendor's procedure, and confirm that the fabric design and zoning accommodate the virtual WWPNs exposed by the array. Dell's SC Series documentation identifies NPIV support and enablement on the Fibre Channel switches as a Virtual Port Mode requirement. See Dell's [Requirements for Virtual Port Mode](#) documentation and the [Dell SC Series Best Practices Guide](#).

QUESTION NO: 12

A customer notices a volume is using an above average amount of Tier 1 RAID 10 space.

What can be done to force this data into a lower RAID set/Tier?

- A. Stop I/O to the volume
- B. Change the storage profile of the volume
- C. Take a manual snapshot
- D. Move data to a different volume


ANSWER: C

Explanation:

Taking a manual snapshot is correct because it initiates a Data Progression opportunity for the SC Series volume. SC Series Data Progression places active data on the preferred high-performance tier and RAID level, while relocating less-active data to lower-cost tiers and RAID configurations according to the volume's configured storage profile and progression rules. A manually created snapshot causes the controller to evaluate and progress the volume's data, rather than waiting for the next scheduled snapshot or normal progression cycle. Consequently, blocks that no longer require Tier 1 RAID 10 performance can be migrated into the lower RAID set or storage tier permitted by the applicable storage profile. This is a nondisruptive controller-managed operation: host access to the volume remains available while the array performs the data movement in the background. The resulting placement remains policy-driven, so the array can preserve Tier 1 capacity for blocks that continue to demonstrate a need for higher performance. Dell SC Series documentation describes Data Progression as the automated mechanism for moving data among storage tiers and RAID levels, and identifies snapshot activity as part of its operation. See the [Dell SC Series documentation portal](#) and Dell's [SC Series storage overview](#).

QUESTION NO: 13

Refer to the Exhibit.



You are testing failover from a host perspective. The host has been configured in a “Failover Only MPIO” configuration. When running Iometer if the Data Path 2 cable from the host is disconnected, will the host survive the failover?

- A. No; Failover MPIO is configured for only Data Path 1
- B. No; Iometer is sending I/O only on one path
- C. Yes; the server has Data Path 1 connected to both controllers
- D. Yes; Failover MPIO should failover

ANSWER: D

Explanation:

Yes; Failover MPIO should failover is correct. A Failover Only MPIO policy presents multiple paths to the same storage LUN while using one preferred active path for normal I/O and retaining an alternate path for recovery. When the host loses connectivity through Data Path 2, the MPIO software detects the path failure and redirects outstanding and subsequent I/O through a surviving path to the same target device. This is precisely the resilience behavior being validated by disconnecting a host cable during an Iometer workload.

The application workload does not need to select or manage the replacement path itself. MPIO operates below the application layer, maintaining a single logical disk presentation to the host while handling path state changes through the configured multipathing policy. A brief pause can occur while the failed path is detected and I/O is retried, but the host should continue accessing the storage provided that an alternate physical path, fabric connection, and storage-controller path remain healthy. Microsoft describes MPIO as providing fault tolerance through multiple physical paths between the host and storage device; see [Microsoft MPIO overview](#). Dell Technologies support resources for SC Series documentation are available through the [SC Series documentation portal](#).

QUESTION NO: 14

What is the maximum number of Fault Domains on a SCv3000 Storage Center using SAS front-end?

- A. 1
- B. 2
- C. 3
- D. 4

ANSWER: D

Explanation:

The maximum number of Fault Domains supported by an SCv3000 Storage Center with a SAS front-end is **4**. In Dell Storage Center architecture, a fault domain is a grouping of storage enclosures arranged so that enclosure-level failures can be isolated and data protection can be maintained across independent hardware paths. The SCv3000 platform's SAS expansion topology supports up to four such domains. This limit is important during physical design because the enclosure cabling, SAS-chain layout, and disk placement policies must be planned within the supported fault-domain topology. Storage Center uses these domains as part of its resiliency design, helping ensure that protected data is distributed in a way that accounts for potential enclosure or path failures. When adding expansion capacity, administrators should therefore validate the current topology and remaining domain capacity before deployment, rather than treating each new enclosure as an unrestricted independent domain. Dell's SCv3000 product documentation and downloadable manuals are available from the [Dell SCv3000 documentation page](#). Additional Storage Center technical documentation is available through the [Dell Support portal](#).

QUESTION NO: 15

An SC Series administrator wants to use Remote Instant Replay to protect a volume at a second site. What must exist on both Storage Centers before a replication relationship can be configured?

- A. A replication connection between the Storage Centers
- B. A shared host server object on both Storage Centers
- C. A common iDRAC management subnet
- D. A single fault domain shared by both sites

ANSWER: A**Explanation:**

A replication connection between the Storage Centers is correct. Remote Instant Replay requires the source and destination Storage Centers to communicate through a configured replication connection. This connection defines the intersystem communication path used to transfer Replay data and supports the configuration of replication relationships.

A host volume mapping is not a replication prerequisite because replication occurs between Storage Centers, not through host paths. Likewise, a shared iDRAC network or a common fault domain does not establish the Storage Center replication connection.

QUESTION NO: 16

You are about to perform the post-setup tasks with the Storage-Center wizard. What information is required before proceeding?

- A. Customer contact information for the Storage Center
- B. Management IPv4 address for the Storage Center
- C. Email address for the default administrator
- D. NTP server IPv4 address for the Storage Center

ANSWER: B**Explanation:**

A management IPv4 address for the Storage Center is required because it establishes the management-network identity through which administrators and supporting management applications communicate with the controller. The Storage Center wizard uses this address as part of the initial configuration needed to make the system reachable for ongoing administration, monitoring, and subsequent configuration tasks. It must be planned in advance as an available address on the management subnet, together with the applicable subnet mask and gateway information, so that the Storage Center can communicate

beyond its local network when required. After the management interface is configured, administrators can connect to the Storage Center using the appropriate Dell management tools and complete the remaining storage, connectivity, alerting, and operational configuration. Dell's SC Series documentation describes initial configuration and management networking as core setup activities for a Storage Center. See the [Dell SC Series Storage Center Setup Guide](#) and Dell's [SC Series support documentation](#) for setup and administration resources.

QUESTION NO: 17

A customer is running the Dell Storage Manager (DSM) Suite as a software stack on Microsoft Windows. To reduce the Microsoft licensing costs, the customer wants to switch to the DSM Virtual Appliance.

Which recommended deployment should the customer use?

- A. Configure as a Primary Data Collector
- B. Configure as a Remote Data Collector
- C. Migrate the DSM database manually
- D. Migrate from the existing Data Collector

ANSWER: D

Explanation:

Migrate from the existing Data Collector is the recommended deployment when replacing an established DSM installation on Microsoft Windows with the DSM Virtual Appliance. This deployment path is designed to move the current Data Collector environment into the appliance while retaining the operational information that DSM depends on, including its configuration and historical database data. It provides continuity for the Storage Center systems and other resources already under management, rather than requiring the administrator to build an entirely new management environment and reconfigure discovery, monitoring, alerts, users, and connections.

For this scenario, the customer's objective is specifically to replace the Windows-hosted DSM software stack and eliminate the associated Microsoft licensing requirement. Migrating the existing Data Collector to the virtual appliance directly accomplishes that objective while preserving the existing DSM deployment. The migration procedure should be planned with suitable backups and maintenance scheduling, and the appliance must meet the supported virtual-machine requirements for the relevant DSM release. Dell publishes DSM documentation and product support materials through the [Dell Storage Manager documentation portal](#). Current Dell storage product support resources are also available from the [Dell Support portal](#).

QUESTION NO: 18 - (DRAG DROP)

DRAG DROP

Match the Snapshot type with the description related to Snapshot creation.

Select & Place:

Type:**Description:**

VMware Snapshot	Replay Manager uses VSS to halt IO for associated volumes when creating a snapshot
Standard Snapshot	Halts all IO to associated volumes when creating a snapshot
SQL or Exchange Snapshot	Takes snapshots in series for associated volumes when creating a snapshot
Parallel Snapshot	Takes snapshots simultaneously for associated volumes when creating a snapshot
Consistent Snapshot	DSITV instructs Data Instant Replay to take a snapshot of the Dell SC volume

ANSWER:**Type:****Description:**

VMware Snapshot	SQL or Exchange Snapshot
Standard Snapshot	Standard Snapshot
SQL or Exchange Snapshot	Consistent Snapshot
Parallel Snapshot	Parallel Snapshot
Consistent Snapshot	VMware Snapshot

Explanation:

SC Series snapshot creation types are selected according to the application or consistency behavior required for the volumes involved. A Standard Snapshot halts I/O to the associated volumes while the snapshot is created, providing a stable point-in-time copy. A Consistent Snapshot creates the snapshots for associated volumes in series, which coordinates the group of volumes during creation. A Parallel Snapshot instead creates snapshots for the associated volumes simultaneously, reducing the time needed to establish the group recovery point. These behaviors are important when volumes have application or dependency relationships and a coordinated recovery image is needed.

Application-integrated snapshot workflows use the type that corresponds to the integration mechanism. For SQL Server or Exchange, Replay Manager uses Microsoft Volume Shadow Copy Service to quiesce the relevant I/O while it creates the snapshot. This provides the application-aware snapshot behavior expected for these Microsoft workloads. For VMware environments, Dell Storage Integration Tools for VMware instructs Data Instant Replay to take a snapshot of the Dell SC volume. This ties the SC volume snapshot operation to the VMware integration workflow. Dell's SC Series documentation

describes Replay Manager, Data Instant Replay, and snapshot operations as part of the SC storage and data-protection feature set; see the [Dell Compellent SC Series documentation library](#) and Dell's [data protection overview](#).