

# **Specialist - Implementation Engineer, Isilon Solutions Exam**

**EMC DES-1423**

**Version Demo**

**Total Demo Questions: 9**

**Total Premium Questions: 95**

**Buy Premium PDF**

**<https://passqueen.com>**

**[support@passqueen.com](mailto:support@passqueen.com)**

## Topic Break Down

| Topic                                 | No. of Questions |
|---------------------------------------|------------------|
| Topic 1, Isilon Solution Overview     | 7                |
| Topic 2, Isilon Cluster Installation  | 9                |
| Topic 3, Isilon Cluster Configuration | 28               |
| Topic 4, Isilon Data Management       | 44               |
| Topic 5, Isilon Troubleshooting       | 7                |
| Total                                 | 95               |

#### QUESTION NO: 1

What is the primary function of NVRAM in an Isilon node?

- A. Provides protected write caching for the node.
- B. Stores permanent copies of file metadata.
- C. Maintains the cluster configuration database.
- D. Provides L3 cache for client read operations.

**ANSWER: A**

#### Explanation:

NVRAM provides protected write caching for an Isilon node. Incoming write operations can be acknowledged after their intent is safely recorded in NVRAM, before the data is committed to disk. If a power interruption or node failure occurs, OneFS can recover the protected write information when the node restarts. This improves write performance while preserving data integrity.

NVRAM is not used as a general-purpose client cache or as permanent file storage. It is a hardware component used to protect in-flight write operations within the OneFS distributed file system. Dell documentation for the applicable hardware platform and OneFS release should be reviewed before performing hardware maintenance. See [Dell PowerScale OneFS documentation](#).

#### QUESTION NO: 2

What are the two ways a FlexProtect job can start?

- A. An at or batch command and Non-event driven
- B. Administrative and Event-driven
- C. An at or batch command and Alert-driven
- D. Administrative and Alert-driven

**ANSWER: B**

#### Explanation:

Administrative and Event-driven is correct. In the OneFS Job Engine, FlexProtect is the protection job that restripes data and restores the requested protection level after a drive or node-related protection event. It can be initiated administratively by an authorized operator, such as by starting the job through the OneFS command-line interface or web administration interface. It can also be initiated automatically by the Job Engine in response to an event that requires protection work. This event-driven behavior is essential to OneFS availability and data-protection operations because the cluster can react to a storage failure or topology change without waiting for an administrator to intervene.

The distinction concerns the initiation source: an administrative start is explicitly requested by an administrator, whereas an event-driven start is generated by OneFS processing an applicable cluster event. FlexProtect execution, prioritization, and status are managed through the Job Engine. Dell's OneFS documentation describes Job Engine operation and the FlexProtect job within its storage and data-protection guidance. See the [Dell PowerScale OneFS documentation portal](#) and the [Dell PowerScale OneFS Job Engine overview](#).

#### QUESTION NO: 3

What are the methods to join a node to an Isilon cluster?

- A. Configuration wizard, SSH connection, or CLI only

- B. Serial connection, Configuration wizard, or CLI only
- C. Front panel display, serial cable, or Configuration wizard only
- D. Front panel display, Configuration wizard, WebUI, or CLI only

**ANSWER: D**

**Explanation:**

Front panel display, Configuration wizard, WebUI, or CLI only is correct because OneFS supports adding a node to an existing cluster through both local-console and management-interface workflows. A node can be directed through its front-panel display during the node-joining process, which is especially useful when performing deployment work physically at the rack. Administrators can also use the Configuration wizard, which guides the cluster expansion process and collects the required cluster-join information. For an already operational cluster, OneFS administrative interfaces provide additional supported methods: the browser-based WebUI can be used to initiate and monitor node addition, while the command-line interface provides an equivalent administrative path for scripted, remote, or advanced operational procedures. These methods reflect the flexible deployment and expansion model of PowerScale OneFS, allowing implementation engineers to select the interface appropriate to the environment, access method, and operational workflow. Dell's OneFS documentation library contains the version-specific installation and administration guidance for adding and managing cluster nodes: [Dell OneFS documentation](#). The PowerScale product documentation portal also provides deployment and cluster-expansion resources: [Dell PowerScale OneFS documentation](#).

**QUESTION NO: 4**

A company's Isilon cluster has reached the capacity limit. The cluster consists of eight H500 nodes in a single node pool. Which action can be taken to resolve the capacity issue while new nodes are being delivered?

- A. Activate SmartPools
- B. Enable GNA
- C. Enable Spillover
- D. Release VHS drives

**ANSWER: D**

**Explanation:**

Release VHS drives is the appropriate temporary measure because Virtual Hot Spare (VHS) capacity is reserved on each node to provide drive-failure protection. Releasing that reserved VHS capacity makes additional disk space available to the OneFS file system without waiting for expansion nodes to arrive. In an eight-node H500 pool, the accumulated reserved capacity can provide enough short-term headroom to bring utilization back below the critical threshold and allow client writes to continue.

This action must be treated as a carefully controlled, temporary capacity-relief procedure. The administrator should first confirm cluster health, verify that no drive failures or rebuild activity are present, determine the amount of free space required, and monitor capacity closely until the delivered nodes are added and the protection configuration is restored. Releasing VHS capacity reduces the immediately available protection reserve, so it is suitable only when the operational need for short-term writable capacity outweighs that temporary reduction in drive-failure tolerance. Dell's OneFS documentation and PowerScale best-practice guidance cover capacity administration, hardware protection, and the operational considerations for maintaining adequate free space: [Dell PowerScale OneFS documentation](#) and [Dell PowerScale OneFS best practices](#).

**QUESTION NO: 5**

You are configuring the network on a Isilon source cluster. The administrator wants to dedicate a new subnet for replication. Which consideration must be addressed for SyncIQ to work properly?

- A. SyncIQ must run on a static IP pool.
- B. SnapshotIQ must be licensed first.
- C. 10GbE ports must be used for replication
- D. SyncIQ must run on a dynamic IP pool

**ANSWER: A**

**Explanation:**

**SyncIQ must run on a static IP pool.** When a subnet is dedicated to SyncIQ replication, its pool must use static address allocation. SyncIQ establishes and maintains replication communications between the source and target clusters, and it requires predictable, persistently assigned source-network addresses for the participating nodes. A static pool provides those stable endpoints and lets the replication network be explicitly separated from client-facing and management traffic.

During implementation, the administrator should create the dedicated subnet, configure appropriate routing and reachability to the target cluster, and assign a static pool with sufficient addresses for the nodes that will participate in replication. Network paths, firewall rules, and name-resolution settings must also allow the source and target clusters to communicate through the intended replication interfaces. This design supports reliable policy execution and failover behavior while preserving a clearly defined network path for SyncIQ traffic.

Dell PowerScale documentation and SyncIQ implementation guidance are available through the [Dell PowerScale OneFS documentation portal](#) and the [Dell Technologies Info Hub SyncIQ basic configuration guide](#).

**QUESTION NO: 6**

Which data access pattern makes the prefetch caching algorithm to be highly aggressive?

- A. Random
- B. Parallel
- C. Concurrency
- D. Streaming

**ANSWER: D**

**Explanation:**

**Streaming** is correct because OneFS identifies a streaming workload when a client reads a file sequentially, progressing through adjacent blocks in order. This pattern provides a strong, reliable indication that blocks immediately following the current read position will soon be requested. The prefetch caching algorithm can therefore act highly aggressively, reading ahead and placing upcoming data into cache before the client explicitly requests it. By overlapping anticipated disk access with the application's current processing, prefetch reduces storage-read latency and helps sustain high sequential-read throughput.

This behavior is particularly valuable for workloads such as media playback, large-file copying, backup reads, and sequential analytics, where data is normally consumed from beginning to end. OneFS uses the observed access pattern rather than merely the file's size to determine the appropriate prefetch behavior, allowing the cache to be used effectively for predictable sequential I/O. Dell EMC's SmartPools data-layout guidance describes streaming access as the workload pattern for which prefetch is most aggressive. See the [EMC Community SmartPools Data Layout Prefetch article](#) and Dell's [PowerScale OneFS documentation](#) for OneFS storage and caching documentation.

**QUESTION NO: 7 - (SIMULATION)**

What is the correct sequence of steps in a client read operation?

**ANSWER: See the explanation for the answer**

**Explanation:**

**Correct client read sequence:**

1. The client connects to a cluster node (typically selected through SmartConnect) and sends the file read request.
2. The receiving/protocol node resolves the file's metadata (LIN/inode) to determine the file layout and the nodes/disks holding the requested data blocks.
3. The protocol node reads locally held blocks and sends internal cluster requests to any peer nodes that own required blocks. Those nodes read the blocks from disk (or cache) and return them across the back-end network.
4. The protocol node assembles the requested data, as necessary, and sends the read response back to the client over the client-facing network.

**In short:** Client request → metadata/layout lookup → local/remote block reads → data returned through the connected protocol node to the client.

### QUESTION NO: 8 - (DRAG DROP)

**DRAG DROP**

What is the correct sequence of actions performed by OneFS to generate an access token?


**Select and Place:**

**Answer Area**

|                                             |        |
|---------------------------------------------|--------|
| Perform user mapping                        | Step 1 |
| Look up user in external directory services | Step 2 |
| Determine on-disk identity                  | Step 3 |
| Perform ID mapping                          | Step 4 |

**ANSWER:**

## Answer Area



### Explanation:

OneFS builds an access token in a deliberate order so that the token represents one consistent security identity across the cluster's supported authentication and file-access environments. The process starts when OneFS looks up the user in external directory services. This supplies the authenticated user's directory information, including the identity attributes and group membership data required to construct the token.

OneFS next performs user mapping. User mapping applies the configured rules that associate or rewrite a presented identity into its effective identity for access. Applying these rules at this point ensures that all later identity-related processing uses the mapped user identity rather than only the original name supplied during authentication.

Once the effective user identity is known, OneFS performs ID mapping. This creates or resolves the numeric and cross-protocol identity values required to support access to data through the cluster's multiprotocol environment. The access token can then carry the identity information necessary for authorization checks regardless of the client protocol used to reach the file system.

Finally, OneFS determines the on-disk identity. At this stage, OneFS has the directory lookup result, the mapped identity, and the required identity mappings available, allowing it to select the identity representation that is used for file-system ownership and permission evaluation. This sequence ensures that OneFS evaluates access with a complete and internally consistent token. Dell's PowerScale documentation resources are available through the [Dell PowerScale OneFS documentation portal](#), and further multiprotocol and identity-management guidance is published in the [Dell Technologies Info Hub](#).

## QUESTION NO: 9

A company plans to build two Isilon clusters in an active-active disaster recovery environment. The initial requirement is to have one access zone for their Finance department. Which recommended directory structure meets the requirement?

### A. Cluster\_A:

/ifs/cluster\_a/finance

/ifs/cluster\_b/finance Cluster\_B:

/ifs/cluster\_a/finance

/ifs/cluster\_b/finance

### B. Cluster\_A:

/ifs/finance

/ifs/finance

Cluster\_B:

/ifs/finance

/ifs/finance

### C. Cluster\_A:

/ifs/finance/cluster\_a

/ifs/finance/cluster\_b Cluster\_B: /ifs/finance/cluster\_a

/ifs/finance/cluster\_b

D. Cluster\_A:  
/ifs/cluster\_a  
/ifs/cluster\_b Cluster\_B:  
/ifs/cluster\_a  
/ifs/cluster\_b

**ANSWER: A**

**Explanation:**

The recommended structure is **Cluster\_A: /ifs/cluster\_a/finance; /ifs/cluster\_b/finance Cluster\_B: /ifs/cluster\_a/finance; /ifs/cluster\_b/finance**. An active-active SyncIQ disaster-recovery design requires each cluster to maintain an equivalent view of the data namespace while clearly separating the data that is authoritative locally from the data replicated from the peer cluster. Creating top-level, cluster-specific directories under `/ifs`, with the Finance directory beneath each cluster root, provides that separation and makes replication policies, failover procedures, and access-zone paths predictable.

Both clusters therefore contain the same pair of namespace roots: one for Cluster\_A-owned Finance data and one for Cluster\_B-owned Finance data. The Finance access zone can be configured consistently against the applicable Finance path, and client access can be redirected during a site outage without needing to redesign the directory hierarchy. This arrangement also scales cleanly when additional departments or access zones are introduced, because each can be placed beneath the corresponding cluster-specific root while preserving clear SyncIQ source and target boundaries. Dell's SyncIQ design guidance discusses active-active replication considerations, and the OneFS documentation describes access zones and their association with directory paths: [Dell PowerScale SyncIQ: A Detailed Review](#) and [Dell OneFS Web Administration Guide: Access Zones](#).