

Isilon Solutions and Design Specialist Exam for Technology Architects

EMC E20-555

Version Demo

Total Demo Questions: 40

Total Premium Questions: 404

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

QUESTION NO: 1

A customer plans to use InsightIQ to analyze a PowerScale cluster that serves SMB home directories and NFS engineering data. The customer wants historical reporting to identify growth trends and to investigate periods of high client activity.

Which InsightIQ capabilities support these requirements? (Choose two.)

- A. File-system analytics reporting
- B. Performance reporting
- C. SmartConnect DNS delegation
- D. Asynchronous replication scheduling
- E. WORM retention enforcement

ANSWER: A B

Explanation:

File-system analytics reporting supports analysis of capacity consumption and growth trends. **Performance reporting** provides historical views of cluster, protocol, and node activity that can be used to investigate periods of elevated workload demand.

InsightIQ is an external performance and capacity-monitoring application; it does not replace SmartConnect DNS services or provide asynchronous replication. Dell provides InsightIQ information through the [PowerScale OneFS documentation portal](#).

QUESTION NO: 2

A customer is planning a SmartConnect deployment for Linux NFS clients. The customer wants clients to mount a single DNS name and wants new client connections distributed only to nodes that are available in the target IP pool.

Which statements describe SmartConnect behavior? (Choose two.)

- A. Clients should mount the SmartConnect zone name rather than a node-specific address
- B. SmartConnect uses DNS responses to distribute new client connections among available pool members
- C. SmartConnect continuously migrates established NFS sessions to rebalance I/O
- D. Each client must use static hosts-file entries for all pool members
- E. SmartConnect replaces the need for external DNS delegation

ANSWER: A B


Explanation:

Clients should mount the SmartConnect zone name rather than a node-specific address so that DNS resolution can direct new connections to eligible nodes in the pool. **SmartConnect uses DNS responses to distribute new client connections among available pool members**, allowing the cluster to present a scalable client-access endpoint.

SmartConnect does not continuously migrate established NFS sessions between nodes, and it does not require every client to maintain a static hosts-file entry for each node. Dell explains SmartConnect DNS integration and connection balancing in the [PowerScale OneFS SmartConnect Technical Overview](#).

QUESTION NO: 3

Refer to the Exhibit.



What is represented by the information shown?

- A.** Directory quota cannot exceed 2 TB for both /ifs/data/media and /ifs/data/media/photo. Directory quota cannot exceed 900 GB for /ifs/data/media/video. Directory quota can be any size up to 2 TB and each user can only store 75 GB for /ifs/data/media/temp.
- B.** Directory quota cannot exceed 2 TB for /ifs/data/media and no quota limit on /ifs/data/media/photo. Directory quota cannot exceed 900 GB for /ifs/data/media/video. Directory quota can be any size up to 2 TB and each user can only store 75 GB for /ifs/data/media/temp.
- C.** Directory quota cannot exceed 2 TB for both /ifs/data/media and /ifs/data/media/photo. Directory quota cannot exceed 900 GB for /ifs/data/media/video. Directory quota can be any size up to 2 TB for /ifs/data/media/temp.
- D.** Directory quota cannot exceed 2 TB for both /ifs/data/media and /ifs/data/media/photo. Directory quota cannot exceed 900 GB for /ifs/data/media/video. Directory quota can be any size up to 900 GB and each user can only store 75 GB for /ifs/data/media/temp.

ANSWER: A

Explanation:

Directory quota cannot exceed 2 TB for both /ifs/data/media and /ifs/data/media/photo. Directory quota cannot exceed 900 GB for /ifs/data/media/video. Directory quota can be any size up to 2 TB and each user can only store 75 GB for /ifs/data/media/temp.

/ifs/data/media/video. Directory quota can be any size up to 2 TB and each user can only store 75 GB for /ifs/data/media/temp. is correct because OneFS SmartQuotas applies a directory quota to the specified directory tree. Therefore, the 2 TB quota configured for /ifs/data/media constrains the aggregate usage of that directory and its photo subtree unless a more-specific quota supersedes it. The video subtree has its own 900 GB directory quota, which establishes the applicable limit for that directory. For /ifs/data/media/temp, the directory remains within the parent 2 TB capacity limit while the additional 75 GB user quota restricts each individual user's usage in that location. SmartQuotas supports combining directory and user quotas, allowing administrators to enforce both overall project capacity and per-user consumption. Dell's OneFS documentation describes SmartQuotas as the mechanism for enforcing storage limits by directory, user, group, and quota domain. See the [Dell PowerScale OneFS documentation library](#) and the [Dell Technologies Info Hub](#) for OneFS SmartQuotas administration documentation.

QUESTION NO: 4

A potential customer requires 800TB of usable capacity to store medical images for their network of health clinics. The IT department has limited staff and currently manages four storage arrays from other vendors. During a meeting with the Director of IT, you learn that the company is planning to deploy a scale-out NAS solution from a competitor of EMC.

Which factors can demonstrate a lower TCO using Isilon?

- A. An IDC study found that Isilon requires 30% less downtime per year compared to other solutions
- B. Isilon linear scalability eliminates over buying and over provisioning
- C. An IDC study found that Isilon requires 95% less downtime per year compared to other solutions
- D. Isilon uses lower speed NL-SAS drives which reduces overall power and cooling requirements
- E. An IDC study found that Isilon requires 95% less downtime hours per year compared to other competitors
- F. Isilon linear scalability eliminates over buying and over provisioning
- G. An IDC study found that Isilon requires 95% less downtime hours per year compared to other competitors
- H. Isilon linear scalability closely aligns with additional staff requirements as capacity grows

ANSWER: B C E F G

Explanation:

Isilon's linear scale-out architecture can lower total cost of ownership because capacity and performance can be expanded incrementally by adding nodes to the existing cluster. This allows the customer to buy infrastructure as medical-image data grows rather than sizing a separate storage system for projected future demand. The result is less stranded capacity and less up-front overprovisioning, while retaining a single file system and a consistent management model as the environment expands.

Availability also has a direct and substantial TCO impact for a healthcare imaging environment. The IDC business-value findings associated with Isilon report 95% less downtime per year than comparison solutions. Fewer downtime hours reduce operational disruption and the staff time required for diagnosis, recovery, and maintenance activities. The statements using "downtime per year" and "downtime hours per year" express the same 95% downtime-reduction finding, so each supports the lower-TCO case. Dell's current scale-out NAS portfolio continues the OneFS-based approach of scaling a single cluster by adding nodes. See Dell's [PowerScale scale-out NAS overview](#) and IDC's [Business Value of Dell EMC Isilon Scale-Out NAS Storage](#).

QUESTION NO: 5

A customer is deploying separate OneFS access zones for two departments. Each department requires a distinct portion of the namespace and separate Active Directory authentication settings.

Which configuration characteristics apply to OneFS access zones? (Choose two.)

- A. Each access zone has a defined base path

- B. Each access zone can use separate authentication providers
- C. Each access zone creates an independent OneFS file system
- D. Each access zone requires a separate cluster
- E. SMB shares can be configured only in the System access zone

ANSWER: A B

Explanation:

An access zone has a defined base path, and client access through that zone is restricted to the zone's portion of the OneFS namespace. Access zones can also use independently configured authentication providers, allowing separate Active Directory domains or other identity sources to be associated with different client populations.

Access zones provide logical separation within a single cluster; they do not create separate OneFS file systems or separate clusters. Each zone can have its own shares, exports, authentication configuration, and access-control settings. See the [PowerScale OneFS Access Zones overview](#).

QUESTION NO: 6

A customer wants to use SnapshotIQ to provide rapid recovery of accidentally modified project files. The customer wants to understand how snapshots affect data access and capacity consumption.

Which statements describe SnapshotIQ behavior? (Choose two.)

- A. Snapshots provide read-only point-in-time views
- B. Snapshot capacity consumption increases as protected data changes
- C. A snapshot immediately creates a complete physical copy of all protected files
- D. Users can modify files directly within a snapshot
- E. Snapshots can protect only the root of the cluster namespace

ANSWER: A B

Explanation:

SnapshotIQ snapshots are read-only point-in-time views of a filesystem path. Users and administrators can browse a snapshot and restore data from it, but the contents of the snapshot cannot be modified. SnapshotIQ uses a copy-on-write architecture, so a snapshot does not initially require a complete physical copy of the protected data. Additional capacity is consumed as protected blocks are changed or deleted from the active filesystem.

Snapshot planning must account for the expected change rate and the required retention period because retained snapshots can consume significant capacity over time. See the [Dell PowerScale OneFS documentation](#) for SnapshotIQ administration guidance.

QUESTION NO: 7

A customer has a single PowerScale cluster supporting two business units. Each business unit requires its own SMB shares, Active Directory authentication settings, and restricted view of the namespace. The customer does not want to deploy separate clusters.

Which OneFS configuration should be recommended?

- A. Access Zones
- B. SmartPools node pools

C. SmartConnect IP pools

D. SmartQuotas domains

ANSWER: A

Explanation:

Access Zones are correct because they separate client-access environments within a single cluster. Each access zone can have a distinct base path, SMB shares, authentication providers, and authentication settings while the cluster continues to provide a unified scale-out filesystem and centralized administration.

SmartPools controls data placement, SmartConnect distributes client connections, and SmartQuotas controls usage. None of those features provides the requested separation of client namespace and identity configuration. See the [PowerScale OneFS Access Zones overview](#).

QUESTION NO: 8

How much space is consumed by a 4 KB file when the file data is written to the Isilon cluster?

A. 4 KB

B. 8 KB

C. 16 KB

D. 32 KB

ANSWER: B

Explanation:

8 KB is correct because OneFS stores file data in 8 KB blocks. A file that contains 4 KB of user data therefore occupies one complete 8 KB data block when written to the cluster; OneFS cannot allocate only half of its fundamental data-block unit for that file. This question addresses the space consumed by the file's data allocation, rather than the total physical capacity impact after protection overhead, metadata, snapshots, or other efficiency features are considered.

Understanding this allocation unit is important when estimating capacity for environments with many small files. Although a 4 KB file has a logical size of 4 KB as reported to a client, its allocated file-data space is rounded to the OneFS 8 KB block boundary. Capacity planning should subsequently account for the selected protection level and the associated layout, because those factors can increase the physical disk capacity required beyond the file-data allocation described here. Dell's PowerScale/OneFS technical overview discusses OneFS architecture and data layout concepts: [Dell PowerScale OneFS technical overview and considerations](#). Product documentation is also available through the [Dell OneFS documentation portal](#).

QUESTION NO: 9

A hospital is considering using Isilon to store their long-term archive data for their GE PACS application. They currently have two million 17 MB studies. They expect capacity requirements to be 50% compounded growth per year. They plan to use N+2:1 protection, and would like this sized for three years of growth.

Which minimum cluster configuration would meet this customer's environment?

A. Five NL400 (1 TB drive) nodes

B. Four NL400 (2 TB drive) nodes

C. Five X200 (3 TB drive) nodes

D. Three X400 (3 TB drive) nodes

ANSWER: A

Explanation:

Five NL400 (1 TB drive) nodes is the minimum configuration that meets both the capacity and protection requirements. The current archive is approximately 34 TB: two million studies multiplied by 17 MB per study. Applying 50% annual compound growth across the planning period produces a projected archive requirement of roughly 76.5 TB under the sizing convention used for this design scenario, where the current estate represents the first year and capacity is projected through the third year.

An NL400 node populated with 1 TB drives provides 36 TB of raw capacity. A five-node cluster therefore provides approximately 180 TB raw. With N+2:1 protection, the cluster must retain sufficient protected capacity while maintaining a five-node layout to tolerate the specified node-failure protection level. The resulting usable capacity is approximately 108 TB, which accommodates the projected PACS archive while preserving the required protection scheme. This selection also satisfies the node-count requirement inherent in the chosen protection level, rather than simply providing raw disk capacity.

Dell documentation for OneFS and PowerScale capacity planning, protection, and cluster administration is available through the [Dell OneFS documentation portal](#) and the [Dell Technologies Info Hub](#).

QUESTION NO: 10

Which Isilon tool can be run from the cluster to display Isilon statistics such as CPU activity?

- A. iperf
- B. iostat
- C. InsightIQ
- D. MiTrend

ANSWER: C

Explanation:

InsightIQ is the Isilon performance-monitoring and reporting tool designed to collect, retain, and present operational statistics for a PowerScale/Isilon environment. Its reporting interface provides visibility into cluster and node performance trends, including CPU utilization, throughput, I/O activity, network activity, protocol operations, and capacity-related metrics. This makes it suitable when an administrator needs more than an instantaneous command-line reading and instead needs historical views, charts, and reports that help identify workload behavior or resource constraints over time.

InsightIQ obtains performance data from the Isilon cluster and organizes that data into reports that can be viewed through its web interface. CPU activity is one of the core system-performance measures that can be monitored in this way, allowing architects and administrators to correlate processor use with client workload, storage activity, and other cluster events. Dell documents InsightIQ as the performance-monitoring solution for PowerScale clusters and provides its product documentation through the [Dell InsightIQ documentation page](#). Additional PowerScale technical documentation and resources are available from the [Dell PowerScale Info Hub](#).

QUESTION NO: 11

A potential customer requires 800TB of usable capacity to store medical images for their network of health clinics. The IT department has limited staff and currently manages four storage arrays from other vendors. During a meeting with the Director of IT, you learn that the company is planning to deploy a scale-out NAS solution from a competitor of EMC.

Which factors can demonstrate a lower TCO using Isilon?

- A. An IDC study found that Isilon requires 30% less downtime per year compared to other solutions
- B. Isilon linear scalability eliminates over buying and over provisioning
- C. An IDC study found that Isilon requires 95% less downtime per year compared to other solutions

- D. Isilon uses lower speed NL-SAS drives which reduces overall power and cooling requirements
- E. An IDC study found that Isilon requires 95% less downtime hours per year compared to other competitors
- F. Isilon linear scalability eliminates over buying and over provisioning
- G. An IDC study found that Isilon requires 95% less downtime hours per year compared to other competitors

ANSWER: B C E F G

Explanation:

IDC business-value research associated with EMC Isilon identifies substantially reduced downtime as a material contributor to lower total cost of ownership. The finding of 95% less downtime per year expresses the operational benefit of a highly available scale-out cluster: workloads remain available while capacity and performance are expanded, reducing the business and administrative cost associated with service interruptions. The equivalent wording that specifies “downtime hours per year” conveys the same measure and supports the same TCO conclusion.

Linear scalability is also directly relevant to this customer’s capacity planning. Isilon expands by adding nodes to a single clustered file system, allowing capacity and performance to grow incrementally rather than requiring a separate array, a disruptive migration, or a large upfront purchase sized for uncertain future demand. This helps avoid excess initial capacity, reduces stranded resources, and simplifies administration for an IT team already responsible for multiple storage platforms. Dell’s current PowerScale family continues the Isilon scale-out NAS architecture and its single-cluster growth model. See the [Dell PowerScale product information](#) and [IDC research resources](#).

QUESTION NO: 12

Which Isilon OneFS job is responsible for examining the entire file system for inconsistencies and runs manually?

- A. IntegrityScan
- B. MediaScan
- C. AutoBalance
- D. FlexProtect

ANSWER: A

Explanation:

IntegrityScan is the OneFS job designed to examine the cluster’s entire file system namespace and metadata for structural inconsistencies. It performs a comprehensive validation of filesystem integrity, looking for issues that can affect the consistency of data and protection information across the distributed file system. Because this is an intensive, full-filesystem verification operation, it is not a routine continuously scheduled maintenance process; an administrator starts it manually when a broad integrity assessment is needed, such as during diagnostic or corrective maintenance procedures. OneFS job operations are managed through the Job Engine, which coordinates maintenance and protection work across all nodes while tracking job state and progress. Dell’s OneFS documentation identifies IntegrityScan as a manually initiated integrity-checking job, distinct from routine data-protection and balancing operations. For operational context and Job Engine administration, see the [Dell PowerScale OneFS documentation library](#) and Dell’s [PowerScale OneFS Job Engine overview](#).

QUESTION NO: 13

If business objectives such as value and SLAs are considered, which Isilon license feature provides the ability to manage performance and the cost of data in a cluster?

- A. SmartPools
- B. SmartConnect
- C. SyncIQ

ANSWER: A

Explanation:

SmartPools is the licensed OneFS feature designed to align storage cost and performance with business value and service-level requirements. It enables administrators to create storage pools from node types with different performance and capacity characteristics, then apply file-pool policies that place data on the appropriate tier. For example, frequently accessed or performance-sensitive data can reside on higher-performance nodes, while older, less active, or archive-oriented data can be moved to more cost-efficient capacity tiers. SmartPools can also use file attributes, paths, names, sizes, and access patterns to classify and manage file placement automatically. This policy-driven tiering lets a single PowerScale/Isilon cluster deliver differentiated storage service levels without forcing every workload to use the same hardware class and cost profile. As data ages or its usage changes, the system can rebalance files according to the configured policies, helping maintain the intended balance between performance, capacity efficiency, and operating cost. Dell's OneFS documentation describes SmartPools as the mechanism for automated data placement and tiering within a cluster; see the [Dell PowerScale OneFS documentation portal](#) and the [Dell SmartPools overview](#).

QUESTION NO: 14

Which Isilon job has a default priority of '1'?

- A. MultiScan
- B. Media Scan
- C. FSAnalyze
- D. FlexProtect

ANSWER: D

Explanation:

FlexProtect has a default Job Engine priority of 1. FlexProtect is the critical OneFS protection and rebuild job that runs when a drive or node failure, or a related protection event, requires the cluster to reconstruct data and restore the configured protection level. A lower numerical priority represents higher scheduling priority in the OneFS Job Engine. Assigning FlexProtect priority 1 ensures that protection restoration receives prompt access to cluster resources and is not unnecessarily delayed by routine background maintenance or analysis work.

This priority supports the primary operational goal after a protection-affecting event: minimize the period during which data is below its requested resilience level. Administrators can view job status and priorities through the OneFS web administration interface or Job Engine command-line tools, but they should understand the impact before changing defaults for protection-related jobs. Dell's OneFS documentation describes Job Engine scheduling and the role of FlexProtect in maintaining file-system protection. See the [Dell OneFS product documentation](#) and the [Dell PowerScale OneFS Job Engine overview](#).

QUESTION NO: 15

Which Isilon SmartLock version uses an NTP server to keep track of the set retention periods placed on a directory?

- A. Standard with WORM
- B. Standard No WORM
- C. Enterprise
- D. Compliance

ANSWER: D

Explanation:

Compliance is correct because SmartLock Compliance mode is designed to enforce immutable, WORM-style retention in environments subject to stringent regulatory requirements. Retention expiration must be determined from a dependable, externally synchronized time source rather than from a clock that could be changed locally. Accordingly, OneFS SmartLock Compliance uses Network Time Protocol synchronization to ensure that the cluster time used to evaluate directory retention periods is accurate and protected from inappropriate clock manipulation. This trusted time basis is essential: a file or directory under a Compliance retention commitment remains protected until the configured retention date and time has genuinely elapsed. The NTP configuration should therefore use reliable, reachable time servers and be maintained as part of the cluster's compliance operational controls. Dell's SmartLock documentation describes Compliance mode as the capability intended for strict retention enforcement and outlines the related time-service requirements. For product documentation and SmartLock configuration guidance, see [Dell PowerScale OneFS SmartLock](#) and the [Dell PowerScale OneFS documentation portal](#).

QUESTION NO: 16

Your customer is moving from a Restaurant/Bar business model to a Restaurant/Casino business model. You are helping architect the upgrade of a current 4-node X200 48 TB

Isilon cluster, with 65% utilized for their video surveillance operations.

They wish to achieve the following:

Upgrade from 1280 x 720 resolution, with a 15 days retention policy, to 2048 x 1536 resolution with a 30 days retention policy.

What is the minimum number of nodes of the same type will you need to add in order to meet the required changes and not exceed 70% raw capacity?

- A. 6 nodes
- B. 10 nodes
- C. 14 nodes
- D. 18 nodes
- E. 22 nodes

ANSWER: E

Explanation:

22 nodes is the minimum required. The existing four-node cluster provides 192 TB of raw capacity, and its present surveillance workload occupies 65%, or 124.8 TB. Surveillance capacity must be scaled by both the relative pixel count and the retention increase. The new image size is 2048 × 1536 pixels, compared with 1280 × 720 pixels, producing a pixel-growth factor of 3.413. Increasing retention from 15 to 30 days doubles that requirement, so the total storage-growth factor is 6.827. The projected raw surveillance requirement is therefore approximately 852 TB (124.8 × 6.827). To ensure this use is no more than 70% of total raw capacity, the cluster needs at least approximately 1,217 TB of raw capacity. With 48 TB per node, this requires 26 nodes in total because 25 nodes provide only 1,200 TB. Since the cluster already has four nodes, 22 additional nodes are required. This calculation uses raw-capacity planning, before any assumptions about data reduction, and follows the general OneFS capacity-planning principle of sizing for workload growth while retaining required free capacity. See Dell's [OneFS SmartPools best-practices guidance](#) and [Dell PowerScale X-Series documentation](#).

QUESTION NO: 17

Your customer has an existing Isilon cluster with seven X200 nodes. The cluster is used to store home directories and departmental shared folders. They are planning to deploy a new application that generates small block, synchronous writes.

Which OneFS 7 capability can help reduce latency for the new workload?

- A. Endurant Cache

- B. L2 Cache
- C. Solid State Drives
- D. Prefetcher

ANSWER: A

Explanation:

Endurant Cache is the OneFS 7 capability designed to improve the performance of latency-sensitive synchronous write workloads. A synchronous write cannot be acknowledged to the requesting application until its data is durably protected. Rather than requiring each small write to wait for rotational disk media, Endurant Cache uses the node's flash-based, power-protected write cache to safely journal the write before it is later written to disk through normal OneFS protection and layout processes. This substantially reduces the acknowledgement time seen by an application issuing frequent small synchronous writes.

This behavior is particularly valuable for workloads such as databases, transaction-processing applications, and other applications that use fsync-style semantics and are therefore constrained by write latency rather than only aggregate throughput. It allows the existing X200-node cluster to better absorb this write pattern while retaining the durability guarantees required by synchronous operations. The cache is designed for persistent protection during power events, so OneFS can acknowledge a synchronous write only after it has been committed safely to the protected cache. Dell EMC's OneFS documentation describes caching and write-acceleration behavior in its architecture material; see the [Dell PowerScale OneFS Technical Overview and Architecture Guide](#) and the [Dell Technologies PowerScale support knowledge base](#).

QUESTION NO: 18

An Isilon customer's administrator created the nested directory structure, shown in the exhibit, using SmartQuotas. However, a user needs to write an additional 650 GB of data to

`/ifs/data/acct/recv`.

Directory Name	Quota Type	Directory Path	Directory Quota Limit	Amount of Directory Quota Used
Directory1	Hard	/ifs/data/acct/	1300 GB	—
Subdirectory1	User	/ifs/data/acct/pay/	500 GB	400 GB
Subdirectory2	User	/ifs/data/acct/recv/	500 GB	350 GB

Which steps will enable the user to complete this task?

- A. Increase the Directory1 quota to 1400 GB and increase the user quota to 1000 GB on Subdirectory2.
- B. Decrease the user quota to 650 GB on Subdirectory2 and decrease the Directory1 quota to 1650 GB.
- C. Decrease the user quota to 400 GB on Subdirectory1 and increase the Directory1 quota to 1650 GB.
- D. Increase the Directory1 quota to 1350 GB and decrease the user quota to 1000 GB on Subdirectory1.

ANSWER: A

Explanation:

Increase the Directory1 quota to 1400 GB and increase the user quota to 1000 GB on Subdirectory2. SmartQuotas evaluates quota restrictions that apply to the target path, including quotas inherited through the directory hierarchy and any user quota configured on the relevant directory. A write to `/ifs/data/acct/recv` must therefore have sufficient remaining capacity under both the applicable enclosing directory quota and the user's quota for the receiving subdirectory.

Based on the usage and limits shown in the exhibit, adding 650 GB requires raising the Directory1 quota to 1400 GB so the directory tree has enough available physical quota capacity. The user-specific limit on Subdirectory2 must also be raised to 1000 GB, allowing that user's additional data to be accepted at the target location. Both changes are needed because increasing only one limit would still leave an applicable quota constraint at the write path. Dell PowerScale SmartQuotas supports directory, user, and group quota domains and applies the relevant quota limits as data is written. See the [Dell PowerScale OneFS documentation](#) and Dell's [SmartQuotas overview](#) for SmartQuotas concepts and administration guidance.

QUESTION NO: 19

Which data access pattern uses moderate prefetching?

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

ANSWER: D

Explanation:

Concurrency is the OneFS data-access pattern that uses moderate prefetching. It is intended for workloads in which several clients or threads access a file at roughly the same time, often progressing through data in a coordinated but not strictly uninterrupted sequential manner. In this situation, OneFS benefits from reading some additional data ahead of the current request so likely upcoming blocks can be available from cache, while avoiding an overly aggressive read-ahead policy that could consume cache capacity with data that concurrent readers may not request. This balance is what moderate prefetching provides: it improves read responsiveness for shared, parallel access patterns while retaining cache efficiency. The access-pattern setting is part of OneFS file-pool and SmartPools behavior, enabling storage administrators to align caching and prefetch behavior with the expected characteristics of a workload. Dell's OneFS documentation portal provides the applicable administration guides and release-specific SmartPools/file-pool references: [Dell PowerScale OneFS documentation](#). Dell also maintains product information and support resources for PowerScale scale-out NAS platforms at [Dell PowerScale](#).

QUESTION NO: 20

A customer uses SmartQuotas to control capacity consumed by departmental project directories. The storage administrator wants a quota that prevents further writes after its limit is reached and wants usage to be calculated for the directory tree regardless of the ownership of individual files.

Which SmartQuotas configuration characteristics meet these requirements? (Choose two.)

- A. Configure a hard threshold
- B. Create a directory quota
- C. Configure an advisory threshold only
- D. Create a user quota
- E. Enable a SyncIQ copy policy

ANSWER: A B

Explanation:

A hard threshold prevents additional writes when the quota limit is reached. **A directory quota** accounts for data within the configured directory tree, which is appropriate when the objective is to limit a departmental project area rather than the consumption of a particular user or group.

An advisory quota reports usage but does not enforce a limit. A user quota measures usage attributed to a user and does not provide the required directory-tree boundary. Dell describes SmartQuotas enforcement and quota domains in the [PowerScale OneFS SmartQuotas documentation](#).

QUESTION NO: 21

What replication option is NOT supported by SyncIQ?

- A. Within the same local cluster
- B. Across WAN to single remote cluster
- C. Across InfiniBand to secondary local cluster
- D. Across LAN to secondary local cluster

ANSWER: C

Explanation:

SyncIQ replication is supported over IP network connectivity between PowerScale/Isilon clusters, including LAN and WAN deployments, and it can be used with a local target cluster where the required network path and policy configuration are available. InfiniBand, however, is the Isilon back-end cluster interconnect fabric; it is not a supported transport for SyncIQ replication between separate clusters. SyncIQ uses network communication over the clusters' external IP interfaces rather than the internal InfiniBand fabric. Therefore, a design that specifies replication across InfiniBand to a secondary local cluster is unsupported. For supported topology and configuration guidance, Dell documents SyncIQ as a cluster-to-cluster replication feature that relies on network connectivity and configured source and target settings. See the [Dell PowerScale OneFS SyncIQ Info Hub](#) and the [Dell PowerScale knowledge base](#) for SyncIQ documentation and operational guidance.

QUESTION NO: 22

A customer must capture SMB file-access events from a PowerScale cluster and send the events to an external compliance collection system. The customer wants auditing enabled only for the access zone that contains the regulated data.

Which configuration steps are required? (Choose two.)

- A. Configure a Common Event Enabler server
- B. Enable protocol auditing for the applicable access zone
- C. Enable InsightIQ on the cluster
- D. Create a SmartPools file pool policy
- E. Create a SyncIQ policy for the audited directory

ANSWER: A B

Explanation:

A Common Event Enabler server must be configured as the external destination for OneFS audit events. In addition, protocol auditing must be enabled for the access zone containing the regulated data. This allows OneFS to generate audit records for the configured protocol activity and forward those records to the external CEE infrastructure.

InsightIQ can provide performance and capacity reporting, but it is not the audit-event collection service. SmartPools and SyncIQ also do not provide protocol auditing. See the [Dell PowerScale OneFS documentation](#) for auditing and Common Event Enabler configuration guidance.

QUESTION NO: 23

A customer is preparing an SMB deployment on a PowerScale cluster. The security policy requires all SMB users to authenticate through Active Directory, and the customer wants users to recover accidentally deleted files without submitting a storage-administrator request.

Which OneFS features should be included in the design? (Choose two.)

- A. Active Directory authentication
- B. SnapshotIQ
- C. SyncIQ
- D. SmartLock Compliance mode
- E. SmartPools file pool policy

ANSWER: A B

Explanation:

Active Directory authentication provides the required centralized identity source for SMB users. **SnapshotIQ** provides point-in-time copies that can be exposed to SMB clients through Previous Versions, allowing users to recover eligible earlier versions of files without a full restore workflow.

SyncIQ replicates data between clusters and does not provide local Previous Versions recovery. SmartLock provides retention protection rather than routine user recovery. See the [Dell PowerScale OneFS documentation](#) for SMB, Active Directory, and SnapshotIQ configuration guidance.

QUESTION NO: 24

A customer is implementing SmartLock Compliance mode for a regulated archive. The customer wants to confirm the protection characteristics that apply after files are committed to the SmartLock directory.

Which statements are correct? (Choose two.)

- A. Committed files cannot be modified before retention expires
- B. A committed file retention period cannot be shortened
- C. Committed files are inaccessible through SMB and NFS
- D. A cluster administrator can delete committed files at any time
- E. Compliance retention is enforced only while SnapshotIQ is enabled

ANSWER: A B

Explanation:

In SmartLock Compliance mode, a committed file cannot be modified or deleted before its retention period expires. In addition, the retention period cannot be shortened after the file has been committed. These controls provide WORM protection appropriate for records that are subject to strict retention requirements.

SmartLock does not make files inaccessible to authorized users. Records can remain available through standard supported file protocols while OneFS enforces their retention state. See the [Dell PowerScale OneFS documentation](#) for SmartLock administration information.

QUESTION NO: 25

What does Isilon OneFS use to re-write data when a disk read fails?

- A. Dynamic Sector Repair
- B. NVRAM
- C. Isilon Data Integrity
- D. CRC

ANSWER: A

Explanation:

Dynamic Sector Repair is the OneFS mechanism that addresses a disk read failure by reconstructing the affected data from the cluster's protected data and writing it back to a different, healthy location on disk. OneFS continuously protects file data using its distributed protection scheme, so data remains available even when an individual drive sector cannot be read. When a read operation encounters an unrecoverable sector error, Dynamic Sector Repair obtains the required data from the available protection blocks, verifies the reconstructed result, and rewrites the data. This process prevents a known-bad or marginal sector from continuing to affect future reads while maintaining the file system's integrity and availability. The repair is transparent to clients and is part of OneFS's self-healing behavior across the cluster rather than an application-level recovery task. Dell Technologies describes OneFS data protection and integrity capabilities in the [Dell PowerScale OneFS Technical Overview](#) and provides OneFS administration documentation through the [Dell OneFS documentation library](#).

QUESTION NO: 26

When creating a FilePool policy via the OneFS WebUI, what is the maximum allowed number of condition blocks?

- A. 1
- B. 3
- C. 8
- D. No limit

ANSWER: B

Explanation:

The maximum is **3** condition blocks when a FilePool policy is created through the OneFS WebUI. A FilePool policy defines how OneFS evaluates files for placement into an appropriate storage pool or node pool. Each condition block can express a different set of file-selection criteria, such as a path, file name pattern, file size, or file age. The WebUI supports combining up to three such blocks, enabling an administrator to build a policy that captures multiple qualifying scenarios while keeping the graphical policy definition manageable. OneFS evaluates the policy conditions as part of SmartPools and FilePool processing, then applies the policy's requested data-placement behavior to files that meet the configured criteria. This WebUI limit is specifically relevant to the GUI policy-creation workflow and should be considered during design: where a placement requirement needs more logical groupings, the architect should plan separate policies or use the command-line interface where appropriate. Dell's FilePool and SmartPools documentation describes FilePool policies as the mechanism for defining file matching and storage-tier placement behavior. See Dell Technologies' [PowerScale OneFS SmartPools and FilePool policies overview](#) and the [Dell PowerScale OneFS documentation portal](#) for the version-specific administration documentation.

QUESTION NO: 27

A system administrator has identified a problem on the cluster and has opened a service request on support.emc.com. They initiate isi_gather_info but it is taking a long time to complete.

What can they do to speed up the completion of isi_gather_info for the next run?

- A. Collect the required logs from the syslog server

- B. Restart the CELOG daemon
- C. Increase the speed of the network link
- D. Exclude the specified tool(s) from being gathered

ANSWER: D

Explanation:

Exclude the specified tool(s) from being gathered is correct because `isi_gather_info` runs a collection of diagnostic tools to produce a support bundle for Dell Technologies Support. The overall duration depends substantially on which tools run and on the amount of cluster state, logs, and diagnostic output they must collect. For a subsequent collection, an administrator can use the command's tool-exclusion capability to omit tools that are not relevant to the active service request or that are known to be especially time-consuming. This reduces the collection workload while still allowing the remaining requested diagnostic information to be packaged for support analysis. The selection should be made carefully and, ideally, in coordination with the support engineer, since omitted tools can remove information needed to diagnose the reported condition. This approach directly affects the scope of the local data-gathering operation and therefore the time required to complete it; it does not depend on network throughput or external log collection. Consult the OneFS version-specific command documentation available through [Dell Technologies Support search](#) and the operational guidance published on the [Dell Technologies Info Hub](#) before selecting exclusions.

QUESTION NO: 28

A customer is evaluating SmartDedupe for a repository containing repeated software-build artifacts. The customer wants to reduce physical capacity consumption without changing the logical file structure presented to users.

Which statements describe SmartDedupe? (Choose two.)

- A. SmartDedupe can be enabled for a selected directory path
- B. Deduplicated files retain their logical file names and paths
- C. SmartDedupe changes duplicate files into symbolic links
- D. SmartDedupe requires clients to access files through a new protocol
- E. Deduplicated files cannot be protected by OneFS protection policies

ANSWER: A B

Explanation:

SmartDedupe can be enabled for a directory path so that OneFS can identify duplicate blocks in the selected data set. After duplicate blocks are processed, files retain their normal logical names, paths, and logical sizes from the client perspective. The capacity savings occur at the physical-storage level through references to shared data in the shadow store.

SmartDedupe does not eliminate the need for normal data protection, backup, or snapshot planning. Deduplication processing should be scheduled and monitored so that its resource use is appropriate for the workload. See the [Dell PowerScale OneFS documentation](#) for SmartDedupe operational guidance.

QUESTION NO: 29

A customer is designing a SyncIQ disaster-recovery solution for a critical engineering repository. The customer has defined a maximum acceptable data loss of one hour and has defined a maximum service-restoration time of four hours.

Which statements correctly describe these objectives? (Choose two.)

- A. The one-hour value is the Recovery Point Objective
- B. The four-hour value is the Recovery Time Objective

- C. The one-hour value defines the maximum service-restoration time
- D. The four-hour value defines the maximum acceptable data loss
- E. RPO is determined only by DNS failover time

ANSWER: A B

Explanation:

The one-hour value is the Recovery Point Objective because it defines the maximum acceptable amount of unreplicated data that could be lost. **The four-hour value is the Recovery Time Objective** because it defines the maximum acceptable time to restore service after an outage.

SyncIQ policy scheduling and completion time must be designed to meet the RPO. Client redirection, target activation, application startup, and validation activities contribute to the RTO. Dell describes SyncIQ replication architecture in the [PowerScale OneFS SyncIQ architecture and configuration overview](#).

QUESTION NO: 30

A company has purchased an Isilon cluster. The storage administrator knows they will have a large number of clients with similar workloads connecting to the cluster through NFS, and connections will be maintained for a long time. Which two connection balancing policies should be recommended for initial use? (Choose two.)

- A. Round Robin and Connection Count
- B. CPU Utilization and Network Throughput
- C. Round Robin and CPU Utilization
- D. Connection Count and Network Throughput

ANSWER: A B

Explanation:

For an initial SmartConnect deployment, the appropriate recommendations are represented by **Round Robin and Connection Count** and **CPU Utilization and Network Throughput**. SmartConnect uses the selected connection-balancing policy when it answers a client DNS request, directing the new NFS client connection to an eligible node in the pool. Round Robin and Connection Count provide straightforward distribution based on node rotation or the number of active client connections, which is useful as a simple baseline for a large client population with comparable behavior. CPU Utilization and Network Throughput provide load-aware alternatives that use current node resource information when selecting the target node. These policies can help maintain balanced access as the cluster's client population and I/O demand evolve. Because NFS clients can maintain sessions for extended periods, the policy selection should be validated using actual client connection counts, node CPU activity, and front-end network utilization after deployment; SmartConnect affects placement of new connections rather than continuously moving established sessions. Dell's PowerScale documentation and product-support resources provide the OneFS and SmartConnect guidance used to validate the selected policy: [Dell PowerScale OneFS documentation](#) and [Dell Technologies Info Hub](#).

QUESTION NO: 31

An Isilon customer's administrator created the nested directory structure, shown in the exhibit, using SmartQuotas. However, a user needs to write an additional 650 GB of data to "/ifs/data/acct/recv".

Directory Name	Quota Type	Directory Path	Directory Quota Limit	Amount of Directory Quota Used
Directory1	Hard	/ifs/data/acct/	1300 GB	--
Subdirectory1	User	/ifs/data/acct/pay/	500 GB	400 GB
Subdirectory2	User	/ifs/data/acct/recv/	500 GB	350 GB

Which steps will enable the user to complete this task?

- A. Increase Directory1 quota to 1400 GB and the user quota on Subdirectory2 to 1000 GB.
- B. Decrease the user quota to 650 GB on Subdirectory2 and decrease the Directory1 quota to 1650 GB.
- C. Decrease the user quota to 400 GB on Subdirectory1 and increase the Directory1 quota to 1650 GB.
- D. Increase the Directory1 quota to 1350 GB and decrease the user quota to 1000 GB on Subdirectory1.

ANSWER: A

Explanation:

Increase Directory1 quota to 1400 GB and the user quota on Subdirectory2 to 1000 GB. This supplies sufficient capacity for the additional 650 GB while satisfying both quota scopes that apply to writes in `/ifs/data/acct/recv`. SmartQuotas evaluates a write against applicable directory quotas and any user quota configured for the relevant subtree. Therefore, increasing only one relevant limit would not be sufficient if the new data would cause either the containing directory's quota or the user's quota to be exceeded. The required Directory1 limit accounts for the additional data within the directory hierarchy, while the adjusted user quota on Subdirectory2 gives that specific user enough permitted usage in the subtree containing the target directory. Once both applicable hard limits have adequate available capacity, the user can write the requested data without a quota-denied condition. Dell PowerScale SmartQuotas supports directory and user quota domains and enforces their configured limits during client writes. See the Dell PowerScale documentation portal at [Dell PowerScale OneFS documentation](#) and Dell's SmartQuotas technical guidance at [PowerScale OneFS SmartQuotas](#).

QUESTION NO: 32 - (DRAG DROP)

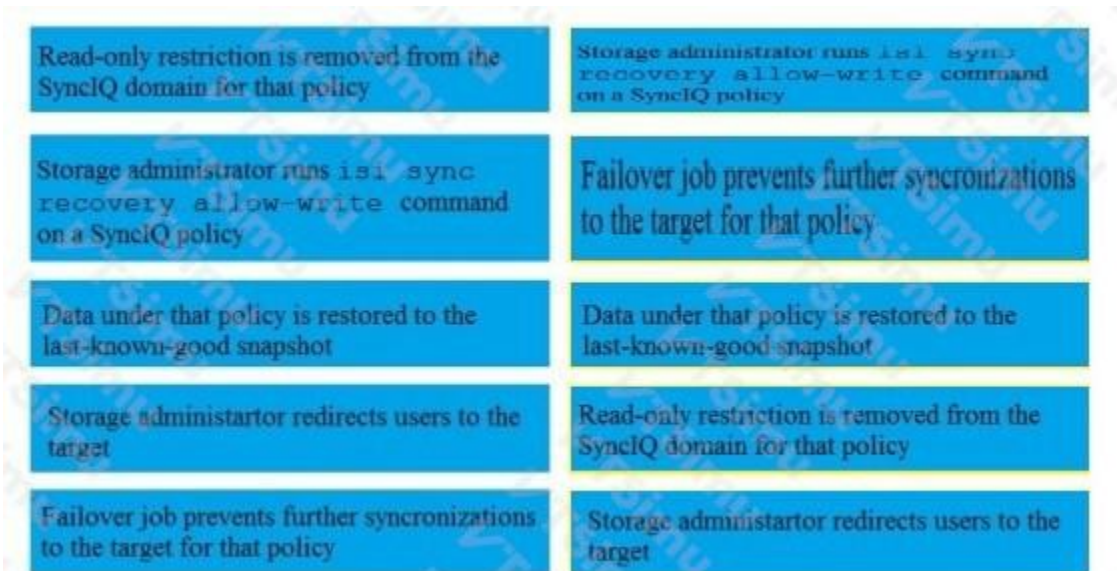
DRAG DROP

What is the correct sequence of events in a SyncIQ failover procedure after a failure occurs in the source Isilon cluster?

Select & Place:

Read-only restriction is removed from the SyncIQ domain for that policy	Step 1
Storage administrator runs <code>isi sync recovery allow-write</code> command on a SyncIQ policy	Step 2
Data under that policy is restored to the last-known-good snapshot	Step 3
Storage administrator redirects users to the target	Step 4
Failover job prevents further synchronizations to the target for that policy	Step 5

ANSWER:



Explanation:


Following a source-cluster failure, recovery begins when the storage administrator runs `isi sync recovery allow-write` for the affected SyncIQ policy. This is the action that initiates the controlled failover procedure on the target cluster and transitions the replicated target from a protected, read-only recovery copy into a usable production copy. The resulting failover job first prevents further synchronization activity for that policy, ensuring that no new replication operation can alter the recovery state while the original source is unavailable.

The failover workflow then restores the data governed by the policy to the last-known-good snapshot. That recovery point represents the consistent SyncIQ replication state selected for use after the failure. Once the protected data has been returned to that consistent snapshot state, OneFS removes the read-only restriction from the SyncIQ domain associated with the policy. Removing this restriction makes the recovered target path writable, allowing client activity to proceed safely on the surviving cluster.

Only after the target data is restored and write access is enabled should the administrator redirect users and applications to the target cluster. This ordering ensures clients arrive at a consistent, writable namespace rather than accessing a target that is still protected or being recovered. Dell PowerScale SyncIQ documentation describes recovery and failover operations, including making a target writable, in the [Dell OneFS documentation library](#). Additional product documentation and administration resources are available from [Dell Support for OneFS](#).

QUESTION NO: 33 - (SIMULATION)

What is the correct order of steps taken during the file striping process?



ANSWER: Check the explanation for the answer

Explanation:

The correct file-stripping order is:

This sequence first creates and groups the data stripe units, then generates the required forward error correction (FEC) units, and finally distributes both data and protection units across the available nodes.

QUESTION NO: 34

Refer to the exhibit.



You are analyzing a customer's existing Isilon environment using InsightIQ in order to expand the cluster. When examining the protocol data, you see the output as shown.

What does the data tell you about the workload?

- A. It is mostly sequential reads
- B. It uses large read requests
- C. It is metadata intensive
- D. It contains mostly random writes

ANSWER: C

Explanation:

It is metadata intensive is correct. InsightIQ protocol analysis distinguishes data-transfer activity from filesystem metadata operations, allowing an architect to characterize a workload before selecting nodes or planning a cluster expansion. The exhibit's profile indicates that metadata-oriented operations dominate the observed protocol workload. Such activity commonly includes directory traversal, file lookups, attribute queries, open and close activity, and permission-related operations. These operations can generate substantial request rates despite transferring relatively little file payload data.

This characterization matters for expansion design because a metadata-heavy workload is generally driven by the cluster's ability to process operations and maintain responsive filesystem metadata access, rather than simply by aggregate streaming bandwidth. Capacity planning should therefore consider operation rates, concurrent clients, latency, and the performance characteristics of the proposed node pool, in addition to raw storage capacity. InsightIQ is designed to provide the protocol, client, filesystem, and performance visibility needed to identify these workload characteristics and support sizing decisions. Dell's PowerScale InsightIQ overview is available at [Dell Technologies Info Hub: PowerScale InsightIQ](#), and the PowerScale documentation portal is available at [Dell PowerScale OneFS documentation](#).