

DUMPSBOSS.

Dell Information Storage and Management Foundations v2 Exam

EMC D-ISM-FN-01

Version Demo

Total Demo Questions: 8

Total Premium Questions: 82

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Data Storage and Infrastructure	16
Topic 2, Storage Systems	22
Topic 3, Storage Networking Technologies	19
Topic 4, Data Protection and Mobility	20
Topic 5, Cloud Computing	3
Topic 6, Mix Questions	2
Total	82

QUESTION NO: 1

Question:

A company is looking for a backup solution that can reduce its IT management overhead and offer scalability and flexibility. Which type of backup solution would meet these requirements?

- A. Image-based backup
- B. Agent-based backup
- C. Cloud-based backup
- D. Tape-based backup

ANSWER: C

Explanation:

Cloud-based backup is the best fit because the backup infrastructure is delivered and operated as a service by a cloud provider or backup-as-a-service provider. This substantially reduces the company's operational burden: the provider manages the underlying storage infrastructure, hardware lifecycle, platform availability, and capacity expansion, while IT staff configure protection policies, monitor jobs, and perform recovery activities. Cloud backup also provides elastic scalability. Rather than purchasing, installing, and forecasting sufficient on-premises backup hardware in advance, an organization can increase protected capacity as data volumes, workloads, or retention requirements grow. This consumption model supports flexibility for protecting distributed systems and retaining backup copies off site without building and maintaining a separate disaster-recovery location. Dell describes Dell APEX Backup Services as a fully managed, SaaS-based offering that helps protect data while simplifying management. Similarly, AWS explains that cloud storage provides scalable capacity without the need to provision and manage physical storage infrastructure. These characteristics directly meet the stated needs for lower IT management overhead, scalability, and flexibility. Reference: [Dell APEX Backup Services](#); [AWS: What is Cloud Storage?](#)

QUESTION NO: 2

Question:

An IT team needs to reduce or eliminate data frame loss during times of network congestion in their converged Ethernet network. What standard should they implement?

- A. QoS
- B. DCB
- C. NAT
- D. VLAN

ANSWER: B

Explanation:

Data Center Bridging (DCB) is the appropriate standard suite for a converged Ethernet environment that must reduce or prevent frame loss during congestion. DCB extends conventional Ethernet with capabilities designed for data-center traffic types, including storage and other workloads that benefit from predictable, low-loss transport. Its key congestion-handling mechanism is Priority-based Flow Control (PFC), standardized as IEEE 802.1Qbb. PFC allows a receiving device to send pause control for a specific traffic priority when its buffers are approaching exhaustion, rather than pausing all Ethernet traffic on the link. This selective, priority-level flow control supports lossless behavior for designated classes of traffic, such as Fibre Channel over Ethernet (FCoE), while allowing unrelated traffic classes to continue forwarding.

DCB also incorporates Enhanced Transmission Selection for bandwidth allocation among traffic classes and uses data-center bridging exchange mechanisms to help maintain consistent configuration across connected devices. Together, these capabilities enable converged networks to carry LAN, storage, and other traffic over shared Ethernet infrastructure while

applying the lossless treatment required by selected priorities. IEEE describes Priority-based Flow Control in [IEEE 802.1Qbb](#), and the broader bridging and traffic-class framework is maintained in the [IEEE 802.1Q standard](#).

QUESTION NO: 3

Question:

A company is implementing a new information security policy. They want to ensure that the policy aligns with its risk appetite, internal policies, and external regulations. Which framework can help the company achieve this goal?

- A. GRC
- B. GDPR
- C. COBIT
- D. ITIL

ANSWER: A

Explanation:

GRC is the correct choice because Governance, Risk, and Compliance provides an integrated organizational approach for directing and controlling activities, identifying and treating risk, and meeting mandatory obligations. In an information security policy initiative, governance establishes accountability, decision rights, and alignment with business objectives. Risk management connects the policy's controls and requirements to the organization's defined risk appetite, so leadership can determine whether residual risks are acceptable. Compliance then maps policy requirements and evidence to applicable legal, regulatory, contractual, and internal obligations.

This integrated view is especially useful because a security policy is not merely a technical document: it must be approved, enforced, measured, reviewed, and defensibly aligned to the organization's risk decisions. GRC practices support those lifecycle activities through common controls, risk registers, compliance assessments, reporting, and audit evidence. ISACA describes GRC as the coordinated capability to govern the enterprise, manage risk, and ensure compliance; NIST's cybersecurity guidance likewise emphasizes that organizational risk management should inform cybersecurity governance and policy decisions. See [ISACA's GRC overview](#) and [NIST Cybersecurity Framework](#).

QUESTION NO: 4

Question:

A host bus adapter is replaced in a Fibre Channel SAN that uses WWN zoning. The replacement adapter has a different World Wide Name but remains connected to the same switch port. What action is required before the host can access its assigned storage?

- A. Update the WWN zone membership with the new HBA WWN
- B. Move the replacement HBA to a different switch port
- C. Recreate the storage LUN on the array
- D. Convert the fabric to port zoning

ANSWER: A

Explanation:

Update the WWN zone membership with the new HBA WWN is correct. WWN zoning identifies zone members by their World Wide Names rather than their physical switch-port locations. A replacement HBA has a new WWN, so the active zoning configuration must include that new identifier.

Changing the switch port is not required merely because the HBA was replaced. Replacing the storage LUN or recreating the host file system does not address fabric access control. The new HBA must also be correctly configured on the host, but updating the zone membership is the required SAN zoning action.

QUESTION NO: 5

Question:

Which two benefits can storage virtualization provide in a data center? Select two.

- A. Improved storage utilization
- B. Elimination of all backup requirements
- C. Non-disruptive data migration
- D. Guaranteed identical performance for all applications
- E. Removal of all physical storage dependencies

ANSWER: A C

Explanation:

Improved storage utilization and **non-disruptive data migration** are benefits of storage virtualization. Virtualization abstracts physical storage resources and presents logical storage to hosts. This abstraction enables capacity from different physical resources to be pooled and allocated more efficiently.

It can also allow data to be moved between physical storage resources while preserving the logical storage presentation to hosts, reducing or avoiding application disruption. Storage virtualization does not eliminate the need for data protection, and it does not guarantee that all workloads receive the same performance.

QUESTION NO: 6

Question:

A company implements continuous data protection for a database that must be recoverable to points before a logical corruption event. Which two functions enable this capability? Select two.

- A. A splitter intercepts production writes for replication processing
- B. A journal retains a history of protected changes
- C. A RAID rebuild reconstructs failed drives
- D. A full backup copies all selected data at scheduled intervals
- E. Deduplication eliminates duplicate backup segments

ANSWER: A B

Explanation:

A CDP solution uses a **splitter to intercept production writes and send write information to the replication process** and a **journal to retain a history of changes**. The journaled write history enables recovery to an available point before the corruption occurred.

A RAID rebuild reconstructs data after a failed physical drive and does not provide granular point-in-time recovery. A full backup is a periodic copy rather than continuous write capture. Deduplication reduces the physical capacity required to store duplicate data, but it does not create a continuous recovery history.

QUESTION NO: 7 - (SIMULATION)

Match each OSD (Object Storage Device) feature with its description:

ANSWER: See the explanation for the answer

Explanation:

Correct matches:

Flexible data access method — Supports multiple methods/protocols for accessing objects, such as REST/SOAP APIs and file access protocols including NFS or CIFS.

Metadata-driven policy — Uses rich object metadata and policies to control object placement, management, retention, and protection according to service requirements.

Data protection — Provides object durability and availability using mechanisms such as replication and/or erasure coding.

Global namespace — Presents one unified namespace that hides the objects' physical storage locations, enabling access regardless of where data resides.

QUESTION NO: 8

Question:

A company needs to rebuild a failed disk in its RAID 5 array. Which two factors would significantly impact the rebuild time?

- A. Network bandwidth between the server and the storage array
- B. Write performance of the remaining disks
- C. Capacity of the failed disk
- D. Capacity of the cache memory

ANSWER: B C

Explanation:

Write performance of the remaining disks and capacity of the failed disk significantly affect a RAID 5 rebuild. During rebuilding, the RAID controller reads data and parity information from the surviving members, reconstructs the contents that belonged on the failed member, and writes the reconstructed blocks to the replacement drive. The sustained write capability of the remaining disks, together with their ability to service the rebuild workload while handling normal I/O, therefore directly determines how quickly reconstruction can proceed.

Capacity also matters because a larger failed member represents a larger address range that the array must reconstruct. Rebuilding more blocks requires more reads, parity calculations, and writes, so—under otherwise similar conditions—a larger-capacity drive takes longer to rebuild. In practical storage systems, rebuild priority and concurrent application activity can further influence the observed duration, but disk capacity and the performance of the surviving drives are fundamental determinants. Dell documentation describes rebuild as reconstructing data on a replacement physical disk from the remaining disks, while RAID references describe the parity-based reconstruction process used by RAID 5. See the [Dell PERC rebuild documentation](#) and the [SNIA RAID overview](#).