DUMPSBOSS.

Dell Data Protection Management Foundations v2 Exam

DELL EMC D-DP-FN-01

Version Demo

Total Demo Questions: 12

Total Premium Questions: 122

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Data Protection Concepts	114
Topic 2, Data Protection Solutions	7
Topic 3, Mix Questions	1
Total	122

QUESTION NO: 1

An organization is preparing a disaster recovery plan for a critical order-processing application.

Which two activities should be included to demonstrate that the plan can meet its recovery objectives? (Choose two.)

A. Document application recovery dependencies

B. Perform regular recovery tests

C. Increase backup frequency without testing restores

D. Retain all recovery copies at the production site

E. Use production credentials for all recovery activities

ANSWER: A B

Explanation:

Document application recovery dependencies and **Perform regular recovery tests** are correct. A recovery plan must identify dependencies such as databases, identity services, network connectivity, storage, and application startup order. Recovering a virtual machine alone does not establish that the application service is usable.

Regular tests validate that documented procedures, recovery copies, staffing, and infrastructure can meet the defined recovery time objective and recovery point objective. Merely increasing backup frequency does not prove that the application can be recovered, and retaining only local copies does not protect against a site-level disaster.

QUESTION NO: 2

A business application must resume service within three hours after an outage. The business can tolerate the loss of no more than 20 minutes of recently committed data.

Which recovery objectives meet these requirements?

A. RTO of 20 minutes and RPO of three hours

B. RTO of three hours and RPO of 20 minutes

C. RTO of six hours and RPO of 20 minutes

D. RTO of three hours and RPO of one hour

ANSWER: B

Explanation:

An RTO of three hours and an RPO of 20 minutes is correct. The Recovery Time Objective (RTO) defines the maximum acceptable time to restore service following an interruption. The Recovery Point Objective (RPO) defines the maximum acceptable amount of data loss measured in time.

Reversing the values would allow the service to remain unavailable for too long and would set an unnecessarily restrictive data-loss target. Longer RTO or RPO values do not meet the stated business requirements.

QUESTION NO: 3

A company is defining protection policies for customer records, application logs, and temporary processing files.

Which two actions support policy-driven data lifecycle management? (Choose two.)

A. Classify data according to business and regulatory value

B. Assign retention periods based on defined requirements

C. Apply one retention period to every dataset

D. Delete data solely because it is inactive

E. Retain all temporary processing files indefinitely

ANSWER: A B

Explanation:

Classify data according to business and regulatory value and **Assign retention periods based on defined requirements** are correct. Classification identifies the sensitivity, importance, and operational value of data. Retention rules then determine how long each class must be retained before it can be disposed of, archived, or moved to another tier.

Applying one retention period to every dataset ignores differing business and compliance obligations. Deleting data solely because it is inactive can violate retention requirements, while keeping every temporary file indefinitely increases storage cost and risk without supporting a defined lifecycle objective.

QUESTION NO: 4

A microservice-based application experiences transient faults due to network instability.

What is the best practice for handling these transient faults?

A. Switching to a different network provider

B. Implementing retry logic with exponential backoff

C. Logging the faults and alerting the support team

D. Increasing the timeout period for requests

ANSWER: B

Explanation:

Implementing retry logic with exponential backoff is the appropriate resiliency practice because transient network faults are temporary conditions that commonly clear without administrative intervention. A retry enables the application to repeat an operation that failed because of a short-lived connectivity interruption, packet loss, throttling event, or temporary downstream-service unavailability. Exponential backoff increases the delay between successive attempts, allowing the affected network path or service time to recover while reducing the risk that many immediate retries will amplify congestion or overload a struggling dependency.

In a microservice architecture, retries should be bounded by a maximum number of attempts and integrated with sensible timeout and cancellation policies. Where applicable, randomized jitter should also be added to retry delays so that many clients do not retry at the same moment. Operations should be designed to be idempotent, or protected with an idempotency mechanism, so a retry cannot unintentionally create duplicate work when the original request completed but its response was lost. This pattern supports service continuity while maintaining controlled behavior during temporary failures. Microsoft's guidance on transient-fault handling describes retry strategies and exponential backoff: Transient fault handling. AWS also documents exponential backoff and jitter as retry best practices: Retry with backoff pattern.

QUESTION NO: 5

A storage system is configured with erasure coding, divided up into 10 data segments and 2 coding segments. Each segment is written to different drives.

What is the maximum number of drive failures that can be withstood without losing the data in this configuration?

A. 2

B. 3

C. 4

D. 5

ANSWER: A

Explanation:

2 is correct. In an erasure-coded layout expressed as 10+2, each protected stripe consists of 10 data segments and 2 coding segments, for 12 segments total. The coding segments contain parity information calculated from the data segments. This information enables the storage system to mathematically reconstruct data from missing segments, provided no more segments are unavailable than the coding protection can cover.

Because the configuration contains two coding segments, it can sustain the loss of any two segments in the stripe. Since the question specifies that every segment is placed on a different drive, two segment losses correspond to two drive failures. The remaining 10 segments provide sufficient information to reconstruct the complete original data. This is the core resilience property of a 10+2 erasure-coding scheme: two segments of redundancy provide tolerance for two simultaneous drive or segment failures without data loss.

Erasure coding is commonly described in terms of data and parity fragments, with the number of parity fragments determining the supported fragment-loss tolerance. Dell Technologies also presents erasure coding as a capacity-efficient data-protection method in its storage documentation. See Dell PowerStore erasure coding documentation and Erasure code overview.

QUESTION NO: 6

Which router in the Virtual Router Redundancy Protocol (VRRP) group is elected as the primary router?

A. The first router in the VRRP group

B. The router with the highest priority

C. The router with the highest bandwidth link

D. The router with the most active connections

ANSWER: B

Explanation:

The router with the highest priority is elected as the VRRP primary router, formally called the Master router. VRRP lets multiple physical routers share a virtual router identity, including a virtual IP address that end hosts use as their default gateway. The Master router is responsible for forwarding packets sent to that virtual address and periodically advertises its status to the Backup routers. VRRP election behavior is controlled primarily by the configured priority: a higher numerical priority is preferred. This allows network administrators to intentionally select the normal active gateway by assigning it a higher priority than its peers. The default VRRP priority is 100; the router that owns the virtual IP address has priority 255 and takes precedence. If routers have the same priority, VRRP uses the numerically highest primary IP address as the tie-breaker. When the current Master becomes unavailable, a suitable Backup router takes over, maintaining gateway availability for attached hosts. This priority-based role selection is defined by the VRRP standard and is independent of traffic load, link bandwidth, connection count, or the order in which routers joined the group. See RFC 5798, Virtual Router Redundancy Protocol Version 3 and Cisco VRRP configuration guide.

QUESTION NO: 7

A data protection management application provides administrators with access to backup policies, retention settings, and recovery operations. The organization wants to reduce the risk associated with compromised administrator credentials.

Which two controls should be implemented? (Choose two.)

A. Require multi-factor authentication for administrator access

B. Apply role-based access control with least-privilege assignments

C. Use a shared administrator account for all backup operators

D. Encrypt backup data at rest

E. Enable audit logging only

ANSWER: A B

Explanation:

Multi-factor authentication adds an independent authentication factor, so possession of a password alone is insufficient to access the management application. **Role-based access control with least-privilege assignments** limits each administrator to only the functions required for their responsibilities, reducing the scope of changes possible through a compromised account.

Audit logging is important for detection and investigation, but it does not prevent account misuse. Encrypting backup data protects stored information but does not control access to the management interface. Using a shared administrator account reduces accountability and makes it more difficult to identify the individual responsible for an action.

QUESTION NO: 8

Data unavailability has occurred in your organization due to natural disasters.

Which two actions should the company have taken prior to the disaster to safeguard their data? (Choose two.)

A. Collaborate with disaster response agencies

B. Implement data encryption

C. Implement a robust disaster recovery plan

D. Establish geographically distributed backups

E. Minimize reliance on on-site storage

ANSWER: C D

Explanation:

Implement a robust disaster recovery plan and **Establish geographically distributed backups** are the complementary measures required to preserve data availability following a natural disaster. A disaster recovery plan documents the people, processes, recovery priorities, recovery point objectives, recovery time objectives, and validated procedures needed to restore applications and data after the primary environment is unavailable. Crucially, this plan should be tested regularly so that recovery teams can execute it reliably under the pressure of an actual incident.

Geographically distributed backups protect against the loss of a complete site or regional infrastructure. Maintaining backup copies in a separate location, such as a remote data center or cloud location, means that a flood, fire, hurricane, or similar event affecting the production site does not also destroy the recoverable data. This follows the core resilience principle of separating backup copies from the systems and location being protected. Used together, an off-site backup strategy provides a surviving recovery copy, while the disaster recovery plan provides the established method for restoring business services from that copy. Dell describes data protection and cyber resilience capabilities at Dell Data Protection, and NIST provides contingency-planning guidance, including recovery strategy and plan testing, at NIST SP 800-34 Rev. 1.

QUESTION NO: 9

A remote office has a constrained WAN connection to a centralized PowerProtect DD system. The office wants to reduce the volume of backup data sent over the WAN.

Which two characteristics make source-based deduplication suitable for this environment? (Choose two.)

A. Duplicate segments are identified before transmission

B. Only unique segments are transferred across the WAN

C. All backup data is first copied to the central site before deduplication

D. It guarantees that no source-side processing resources are used

E. It removes the need to retain recovery copies at a backup target

ANSWER: A B

Explanation:

Duplicate segments are identified before transmission and **Only unique segments are transferred across the WAN** are correct. Source-based deduplication processes backup data near the protected source and avoids transmitting segments already represented on the target. This reduces WAN consumption and can help remote backup jobs complete within limited network windows.

The approach does not eliminate the need for a central backup target or recovery policies. It may use source-side processing resources, so it is not selected because it guarantees the lowest client CPU use.

QUESTION NO: 10

How does block-level deduplication with variable-length segments improve storage efficiency?

A. Adjusts segments based on data changes

B. Uses object IDs to store and retrieve data

C. Stores multiple copies of the same data segment

D. Breaks files down into fixed-length segments

ANSWER: A

Explanation:

Adjusts segments based on data changes is correct because variable-length, content-defined segmentation identifies chunk boundaries from patterns in the data itself instead of imposing rigid byte-offset boundaries. This makes deduplication resilient when a protected file changes slightly between backup versions. For example, after an insertion or deletion, the segmentation process can quickly establish matching boundaries again beyond the changed area. The unchanged portions therefore produce the same fingerprints as previously stored segments and can be referenced rather than written again. Only genuinely new or changed segments require additional capacity.

This behavior improves deduplication ratios for evolving datasets, such as databases, virtual-machine images, and frequently revised files, where small changes may occur within otherwise highly similar data streams. By retaining recognition of unchanged content despite shifted byte positions, content-defined segmentation reduces the amount of unique physical storage consumed across full and incremental backup data. Dell PowerProtect DD systems use deduplication and compression to reduce the disk capacity required for protected data; the effectiveness of that reduction depends on accurately detecting repeated content across backup versions. See Dell's PowerProtect DD Series Appliances Architecture Guide and the PowerProtect DD Operating System documentation.

QUESTION NO: 11

A ransomware attack has compromised production servers and an administrator account used to manage backup policies.

Which two measures most directly reduce the risk that attackers can destroy recovery copies? (Choose two.)

A. Maintain immutable recovery copies

B. Require multi-factor authentication for administrative access

C. Increase backup frequency without protecting backup retention

D. Store all backups only on the production network

E. Disable audit logging to reduce management overhead

ANSWER: A B

Explanation:

Maintain immutable recovery copies and **Require multi-factor authentication for administrative access** are correct. Immutability prevents protected copies from being changed or deleted during their configured retention period, limiting an attacker's ability to destroy recovery data. Multi-factor authentication reduces the likelihood that a stolen password alone can be used to access the data protection management plane.

Increasing the frequency of backups can reduce data loss but does not prevent deletion of the copies. Storing all backups on the same production network increases exposure, and disabling audit logging removes evidence needed to detect and investigate unauthorized activity.

QUESTION NO: 12

Which type of virtual machine (VM) is created from a snapshot of the parent VM in a clone operation?

A. A snapshot clone

B. A replicated clone

C. A full clone

D. A linked clone

ANSWER: D

Explanation:

A linked clone is created using a snapshot of an existing parent virtual machine. Rather than duplicating every block of the parent VM's virtual disk, the linked clone retains a relationship to the parent's base disk and snapshot, while writing its own changes to separate delta disks. This design allows the clone to be provisioned quickly and uses substantially less storage at creation time than copying an entire virtual disk. The snapshot provides the point-in-time state from which the linked clone is derived, and the clone reads unchanged data through its dependency chain while maintaining only its unique changes independently. This approach is particularly useful where many similar VMs must be deployed efficiently, such as lab, test, training, and virtual desktop environments. VMware documents that linked clones are created from a snapshot of a parent virtual machine and share the parent's virtual disks; therefore, appropriate lifecycle management of the parent VM and its snapshot is essential. Dell data-protection administrators should recognize this relationship when planning protection, retention, and recovery operations because the linked clone's disk chain depends on that source snapshot state. See VMware's virtual-machine cloning documentation and Horizon linked-clone documentation.