

Implementation EngineerData Protection Exam

Netapp NS0-528

Version Demo

Total Demo Questions: 6

Total Premium Questions: 67

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Data Protection Solution Planning	30
Topic 2, Data Protection Solution Implementation	32
Topic 3, Data Protection Solution Troubleshooting	5
Total	67

QUESTION NO: 1

Which two of the following need to be enabled to use tamper-proof Snapshot copies? (Choose two.)

- A. Compliance Clock
- B. three NTP servers
- C. SnapRestore license
- D. NetApp ONTAP One license

ANSWER: A D

Explanation:

Tamper-proof Snapshot copies use ONTAP SnapLock technology to provide WORM-style protection: after a Snapshot copy is committed with a retention period, it cannot be modified or deleted until that retention period expires. The **Compliance Clock** must be initialized because SnapLock relies on this secure, cluster-wide time source to calculate and enforce retention. Initializing it establishes the trusted clock foundation that prevents retention from being bypassed by changing ordinary system time. The **NetApp ONTAP One license** is also required because SnapLock functionality is included in the ONTAP One software suite. With the licensing entitlement in place and the Compliance Clock initialized, administrators can configure the applicable SnapLock settings and create protected Snapshot copies with enforceable retention. These prerequisites ensure that the immutability behavior is backed both by the licensed ONTAP feature set and by a protected time mechanism. NetApp documents SnapLock licensing and its inclusion in ONTAP One in the [ONTAP license management documentation](#). Details on establishing the secure retention time base are available in the [SnapLock Compliance Clock documentation](#).

QUESTION NO: 2

Your customer set up a SnapMirror synchronous relationship between two sites. They want to convert it to a SnapMirror active sync configuration for availability reasons.

What action must you take before you convert this relationship?

- A. Change the SnapMirror policy to Continuous.
- B. Destroy all volumes in the destination SVM.
- C. Unmap all LUNs on the destination volumes.
- D. Stop the destination SVM that is hosting the LUNs.

ANSWER: C

Explanation:

Unmap all LUNs on the destination volumes before converting the relationship. A SnapMirror synchronous destination is normally a protection copy whose LUNs should not remain presented to hosts. By contrast, SnapMirror active sync is designed to provide continuous availability for SAN workloads through coordinated access to replicated LUNs across the two clusters. Removing destination-side LUN mappings before the conversion prevents pre-existing host presentation from conflicting with the active-sync configuration and ensures that the LUN access path can be established in the supported manner after conversion.

This requirement is especially important in an environment using multipathing, because hosts must see and use the intended, consistently configured paths after the active-sync relationship is created. The conversion process therefore begins by reviewing all LUN mappings on each destination volume, removing those mappings, and then proceeding with the relationship conversion and supported host-path configuration. NetApp documents SnapMirror active sync requirements and conversion procedures in the [SnapMirror active sync documentation](#) and describes the relevant replication relationship operations in the [ONTAP SnapMirror CLI reference](#).

QUESTION NO: 3

An administrator needs to maintain 10 snapshots per volume on a NetApp AFF system and replicate to a destination with a snapshot retention of 120.

Which NetApp feature should the administrator use?

- A. SnapMirror active sync
- B. XCP
- C. MetroCluster
- D. SnapMirror

ANSWER: D

Explanation:

SnapMirror is the appropriate feature because it provides ONTAP's native volume replication and supports retaining a different number of Snapshot copies at the source and destination. The administrator can keep a short operational retention window of 10 Snapshot copies on each AFF source volume while using a SnapMirror policy at the destination to retain up to 120 replicated Snapshot copies. This arrangement supports longer-term recovery requirements without requiring the production AFF volume to hold the full destination retention set.

In ONTAP, SnapMirror policies contain rules that associate Snapshot labels with retention counts. A policy designed for backup retention, such as a mirror-vault or vault policy, can retain labeled Snapshot copies at the destination according to the specified count. Source Snapshot schedules and retention settings can therefore remain aligned with the 10-copy local operational requirement, while destination policy rules preserve the required 120 copies for recovery. SnapMirror relationships also transfer Snapshot copies incrementally after the baseline, making this an efficient and standard ONTAP data-protection design. See NetApp's [SnapMirror replication workflow documentation](#) and its [custom replication policy documentation](#) for policy-rule retention behavior.

QUESTION NO: 4

An administrator must continuously synchronize an on-premises SMB share to cloud object storage by using BlueXP copy and sync. The source and target are not ONTAP clusters and cannot use SnapMirror.

Which two requirements must be met to create the synchronization relationship? (Choose two.)

- A. Deploy a BlueXP copy and sync data broker.
- B. Provide network connectivity from the data broker to both endpoints.
- C. Create an ONTAP cluster peer relationship.
- D. Create one intercluster LIF on each node.

ANSWER: A B

Explanation:

A **data broker** is required because it performs the discovery, transfer, and synchronization work for BlueXP copy and sync. The data broker must also have network connectivity to both the source SMB share and the target object-storage service so that it can read source data and write the synchronized copy.

Cluster peering and intercluster LIFs are ONTAP requirements for SnapMirror replication. They are not required for a BlueXP copy and sync relationship between heterogeneous file and object-storage endpoints.

QUESTION NO: 5

A company is using SnapMirror SVM replication, with identity preserve set to false to protect an SMB-only dataset in case of a disaster.

What needs to be manually configured to allow SMB to have access to the data on the destination?

- A. SMB server
- B. home directories
- C. SMB shares
- D. name mappings

ANSWER: A

Explanation:

SMB server is correct. An SMB server is the SVM-level CIFS/SMB identity that provides the destination SVM's ability to participate in Active Directory and serve SMB clients. It includes the server name and the associated Active Directory machine-account relationship. When an SVM disaster-recovery relationship is configured with identity preservation disabled, the destination is intended to retain a distinct identity rather than assume the source SVM's identity. Consequently, SMB server identity and its Active Directory integration must be configured for the destination before it can provide SMB access after activation.

After the destination SVM is activated, administrators configure an SMB server appropriate for that destination environment, join it to the relevant Active Directory domain, and ensure the required network and name-service connectivity is available. This gives clients a valid SMB endpoint through which they can authenticate and access the replicated namespace and data. NetApp documents that identity preservation controls whether source SVM identity configuration is replicated as part of SVM disaster recovery; a non-identity-preserved destination requires configuration for serving data. See NetApp's [SVM disaster-recovery replication documentation](#) and its guidance for [creating an SMB server in an Active Directory domain](#).

QUESTION NO: 6

An administrator is deploying NetApp BlueXP backup and recovery for VMs to protect both Windows and Linux systems. The site has two instances of VMware Cloud on AWS that need to be backed up.

How will this goal be accomplished? (Choose two.)

- A. Pair each instance of VMware Cloud on AWS with a unique instance of BlueXP backup and recovery.
- B. Deploy one instance of BlueXP backup and recovery on a single Linux VM.
- C. Pair both instances of VMware Cloud on AWS to a single instance of BlueXP backup and recovery.
- D. Deploy two instances of BlueXP backup and recovery on separate Linux VMs.

ANSWER: C D

Explanation:

Both deployment models can meet the requirement to protect virtual machines running Windows and Linux across two VMware Cloud on AWS instances. "Pair both instances of VMware Cloud on AWS to a single instance of BlueXP backup and recovery" supports a centralized design: one BlueXP backup and recovery deployment can be paired with and manage backup operations for both VMware Cloud on AWS environments. This approach centralizes administration, policy management, monitoring, and backup control while allowing workloads in each connected environment to be protected.

"Deploy two instances of BlueXP backup and recovery on separate Linux VMs" is also valid when the administrator wants a distributed deployment. Each BlueXP backup and recovery instance can be deployed on its own Linux virtual machine and paired to a VMware Cloud on AWS instance. This model can be selected to provide administrative separation, isolate operations between environments, or align the backup architecture with site-specific operational requirements. In either design, protection is VM-based, so the guest operating system can be Windows or Linux; backup is managed at the VMware virtual-machine layer rather than requiring a separate backup deployment for each guest OS. NetApp describes BlueXP

backup and recovery capabilities at the [BlueXP backup and recovery product page](#); VMware Cloud on AWS environment concepts are documented in the [AWS VMware Cloud on AWS User Guide](#).