

NetApp Certified Implementation Engineer - Data Protection

Netapp NS0-527

Version Demo

Total Demo Questions: 27

Total Premium Questions: 277

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, ONTAP Data Protection	195
Topic 2, SnapCenter	40
Topic 3, MetroCluster	18
Topic 4, Mix Questions	24
Total	277

QUESTION NO: 1

You decide to enable SnapMirror network compression on your NetApp clustered Data ONTAP environment. Which two statements are correct about enabling this feature?

(Choose two.)

- A. The data is compressed only while it traverses the network.
- B. The data transfer between the clusters must use the NDMP protocol.
- C. You may only enable SnapMirror network compression on one HA pair in a cluster, and traffic 1-1 is limited to routing through those nodes.
- D. The compression engine creates multiple transfer threads corresponding to the number of CPUs on the source storage system.

ANSWER: A D

Explanation:

SnapMirror network compression is a transport optimization: the source system compresses the replication stream before it is sent across the intercluster network, and the destination decompresses it as part of receiving the transfer. Consequently, it reduces WAN bandwidth consumption for eligible SnapMirror transfers without changing how the replicated data is stored on disk at either endpoint. This distinction is important when sizing replication links: the potential bandwidth saving depends on the compressibility of the transferred data, while storage efficiency on the source and destination remains governed by their configured storage-efficiency technologies.

The compression implementation is also designed to use available source-controller processing resources efficiently. Its compression engine can create parallel transfer/compression threads based on the CPUs available on the source storage system. Parallelism helps the system process replication data at a rate suitable for the network and workload, rather than treating compression as a single serial operation. NetApp exposes network compression as a SnapMirror relationship capability, including the `compression-enabled` setting in SnapMirror relationship creation and modification workflows. See the NetApp [snapmirror create command reference](#) and the [snapmirror modify command reference](#) for the relationship compression controls.

QUESTION NO: 2

What are two ways to determine if your SnapMirror relationships are up-to-date? (Choose two.)

- A. Use Active IQ and navigate to Cluster -> Volumes -> Protection.
- B. Use Active IQ Unified Manager and navigate to Protection -> Volume Relationships.
- C. Use ONTAP System Manager and navigate to Protection -> Volume Relationships on the destination cluster.
- D. Use ONTAP System Manager and navigate to Protection -> Volume Relationships on the source cluster.

ANSWER: B C

Explanation:

Active IQ Unified Manager provides a centralized view of data-protection relationships through *Protection > Volume Relationships*. This view reports the relationship state, transfer status, lag time, last transfer details, and health information. These indicators allow an administrator to verify that scheduled SnapMirror updates are completing and that the destination is not behind the source beyond the applicable recovery-point objective. It is especially useful when monitoring relationships across multiple clusters from one management interface.

ONTAP System Manager can also verify the current state of a SnapMirror relationship from the destination cluster through *Protection > Volume Relationships*. The destination owns the replicated copy and maintains the replication relationship information used to show whether the relationship is healthy and whether transfers are current. Reviewing the destination-side relationship provides the operational confirmation needed before recovery, failover, or other protection activities. NetApp documents relationship monitoring in both the Unified Manager protection interface and ONTAP System Manager

QUESTION NO: 3

A customer is running ONTAP 9.7 on a FAS2720 cluster that has 12 x 4 TB NL-SAS drives in the chassis with capacity running low. You decide to use Fabric Pool technology to solve this problem.

What does the customer need to do to accomplish this task?

- A. Add an SSD aggregate.
- B. Purchase a Data Protection Optimized (DPO) license.
- C. Use the CLI and add a capacity pool.
- D. Use the wizard and add a capacity pool.

ANSWER: A

Explanation:

Add an SSD aggregate. In ONTAP 9.7, a FabricPool-enabled aggregate on a FAS system requires an all-SSD aggregate as its local performance tier. The existing NL-SAS disks form an HDD-based local tier, so they cannot serve as the FabricPool local tier for this ONTAP release. The customer must therefore add supported SSD capacity and create an all-SSD aggregate before associating that aggregate with an object-store cloud tier. FabricPool then retains active or frequently accessed data in the SSD performance tier and can tier eligible inactive blocks to the configured capacity tier, helping relieve pressure on the local SSD tier while retaining transparent NAS or SAN access through ONTAP. The cloud tier must also be configured with a supported object-store provider or StorageGRID target and appropriate connectivity and credentials. This approach is specifically important for the stated ONTAP version: FabricPool local-tier support and capabilities have evolved over later ONTAP releases, so implementation requirements should be validated against the target release documentation. See NetApp's [FabricPool requirements](#) and the [FabricPool configuration documentation](#).

QUESTION NO: 4

You are planning to deploy MetroCluster IP. As part of the implementation strategy, you are asked to verify that you have the relevant hardware for the deployment.

Which two elements are required to deploy your solution? (Choose two.)

- A. storage controllers
- B. FC network switches
- C. cluster peering network
- D. MetroCluster IP Ethernet switches

ANSWER: A D

Explanation:

A MetroCluster IP deployment requires storage controllers and MetroCluster IP Ethernet switches. The storage controllers are the ONTAP systems at each site that host the MetroCluster configuration, including the replicated aggregates and mirrored data required for continuous availability across the two locations. Each MetroCluster IP site contains an HA pair of supported controllers, and the two sites together form the MetroCluster configuration.

MetroCluster IP Ethernet switches provide the dedicated IP switching fabric that carries the MetroCluster back-end traffic between controller modules and across sites. These switches must be supported and configured for MetroCluster IP; they provide the redundant network paths used for remote nonvolatile memory replication and storage replication. Although MetroCluster IP uses Ethernet connectivity, the required switching infrastructure is identified as MetroCluster IP switches

rather than generic iSCSI switches. NetApp's architecture documentation describes the controllers and IP switches as core components of an IP MetroCluster configuration. Review the supported platform, switch, and cabling requirements before deployment in the [NetApp MetroCluster IP configuration components documentation](#) and the [MetroCluster IP installation considerations](#).

QUESTION NO: 5

Using ONTAP System Manager, you are asked to replicate all the volumes contained in a Storage VM to a remote ONTAP 9.7 system for quick disaster recovery.

Which two steps need to be completed to accomplish this task? (Choose two.)

- A. Create a SnapMirror relationship for each volume.
- B. Create a cluster peering relationship.
- C. Create an SVM-DR relationship.
- D. Create an SVM peering relationship.

ANSWER: B C

Explanation:

Creating a cluster peering relationship and creating an SVM-DR relationship are the required high-level steps for replicating the contents of an entire Storage VM to a disaster-recovery cluster. Cluster peering establishes the trusted intercluster connectivity needed for SnapMirror replication traffic between the source and destination ONTAP clusters. It also provides the foundation on which the Storage VM disaster-recovery configuration is created.

An SVM-DR relationship is specifically designed to protect a Storage VM as a unit. It replicates the eligible volumes and Storage VM configuration required to recover service on the destination cluster, rather than requiring a separate volume-level relationship for every volume. In ONTAP System Manager, creating the SVM-DR SnapMirror relationship includes selecting the source Storage VM and destination Storage VM and configuring the destination relationship for disaster recovery.

This approach is appropriate when the objective is rapid recovery of all protected volumes in the Storage VM at a remote site. NetApp documents the SVM disaster-recovery workflow and its cluster-peering prerequisite in the [ONTAP SVM replication relationship documentation](#). Additional details about cluster peering are available in the [ONTAP cluster peering documentation](#).

QUESTION NO: 6

You are deploying a mission-critical business application for a healthcare customer that, according to government policies and regulations, requires zero data loss.

In this scenario, which NetApp solution enables you to set up data replication at a volume level to satisfy your business needs?

- A. MetroCluster
- B. SnapMirror-Synchronous
- C. SnapShot Copies
- D. SyncMirror

ANSWER: B

Explanation:

SnapMirror-Synchronous is the appropriate solution because it provides synchronous replication between ONTAP source and destination volumes. For each protected write, ONTAP confirms that the data has been committed at both the primary and secondary sites before acknowledging completion to the host. This behavior delivers a recovery point objective of zero (RPO=0) while the synchronous relationship is operating normally, which is essential for an application whose policy requires no acknowledged data loss.

The protection is configured as a SnapMirror relationship for individual FlexVol volumes, allowing the replication scope to match the application's data layout and protection requirements. SnapMirror Synchronous supports both strict and sync policies. A strict policy preserves zero-RPO semantics by causing host I/O to stop if synchronous replication cannot be maintained; this is commonly selected when the requirement is absolute zero data loss. The destination volume maintains a consistent, current replica that can be used for recovery if the source site becomes unavailable.

NetApp documents synchronous SnapMirror as a volume-level replication technology intended for workloads requiring zero RPO. See the [ONTAP SnapMirror Synchronous deployment documentation](#) and the [ONTAP SnapMirror relationship configuration documentation](#).

QUESTION NO: 7

Your customer wants to implement SVM DR, but the customer is not certain whether the identity-preserve parameter should be enabled. The customer explains that the source and destination clusters are in different network subnets.

In this scenario, which SVM DR solution should the customer use?

- A. SVM DR using the identity-preserve parameter disabled and the default SnapMirror policy.
- B. SVM DR using the identity-preserve parameter enabled and then administratively take the data LIFs offline on the DR cluster.
- C. SVM DR with the identity-preserve parameter enabled and a SnapMirror policy with the discard-configs network parameter.
- D. SVM DR with the identity-preserve parameter enabled and the default SnapMirror policy.

ANSWER: C

Explanation:

SVM DR with the identity-preserve parameter enabled and a SnapMirror policy with the discard-configs network parameter is the appropriate design when the source and destination clusters use different network subnets. Identity preservation replicates the SVM's identity-related configuration so that the destination can assume the protected SVM's role during a disaster-recovery activation. However, network configuration requires special treatment when the DR site cannot use the same IP subnets, broadcast domains, or routing environment as the production site. A custom SnapMirror policy that discards the `network` configuration category prevents source network objects, such as LIF networking configuration, from being applied at the destination. This lets the destination cluster retain or use networking that is valid for its own site while still receiving the protected SVM data and applicable SVM configuration. The result is an identity-preserving SVM DR relationship that is usable across dissimilar network environments without attempting to bring online addresses from the source subnet at the DR site. NetApp documents the use of the `discard-configs` policy setting for excluding specific SVM configuration categories and describes SVM disaster-recovery relationship configuration in the [ONTAP SVM DR documentation](#) and [SVM DR destination management guidance](#).

QUESTION NO: 8

A company just installed SnapCenter to help manage their ONTAP 9.7 environment. The database administrator wants to add new volumes with LUNs for a 10 TB Oracle database.

Which two tools would accomplish this task? (Choose two.)

- A. SnapCenter Plug-in for Oracle Database
- B. OnCommand Insight

C. ONTAP System Manager

D. Active IQ Unified Manager

ANSWER: A C

Explanation:

SnapCenter Plug-in for Oracle Database is appropriate because SnapCenter's Oracle plug-in integrates application-aware Oracle database management with ONTAP storage operations. It can be used to provision and manage storage associated with Oracle workloads, allowing the database administrator to use SnapCenter workflows while ensuring that the resulting LUN-based storage is properly associated with the Oracle environment. This is particularly useful for a large database because SnapCenter maintains application context and provides a consistent platform for subsequent protection, cloning, and recovery operations.

ONTAP System Manager is also appropriate because it is the ONTAP management interface used to administer storage objects directly. An ONTAP administrator can create the required volume, configure SAN settings, create LUNs within the volume, and map those LUNs to the appropriate initiator group for the Oracle host. System Manager therefore provides the underlying ONTAP storage provisioning capability, while the Oracle plug-in provides database-aware management through SnapCenter. Together, these tools support both the storage configuration and the Oracle application workflow required for the new database capacity. See the [SnapCenter Plug-in for Oracle Database documentation](#) and the [ONTAP documentation for creating volumes and configuring LUNs](#).

QUESTION NO: 9

Click the Exhibit button.

Supported Adapter Cards - FAS9000 (Single Chassis HA Pair (FAS9000A)) 9.5 ONTAP

Priority	Category	Bus Type	Mktg Part No.	Mfg Part No	Description	Optional Module	Supported Speed(s)	Min ONTAP	Max Qty	Priority Slot Assignment
1	NVRAM	IO Module	X93140A	111-03190	NVRAM10 32GB	Not supported		9.1RC2, 9.2RC1, 9.3RC1, 9.4RC1, 9.5RC1.	1	6
2	Networking	IO Module	X91440A	111-02590	2p 40GbE NIC Cu Op	OSPF Optional	ETH: 10/40 Gbps	9.1RC2, 9.2RC1, 9.3RC1, 9.4RC1, 9.5RC1.	8	4, 6, 9, 1, 10, 2, 11, 3, 7, 5
3	Block Access	IO Module	X91135A	111-03789	4p 32Gb FC SFP+Op	SFP+ Included	FC: 8/16/32 Gbps FC-NVMe: 16/32 Gbps	9.3RC1, 9.4RC1, 9.5RC1.	8	9, 1, 10, 2, 11, 3, 7, 8
4	Block Access	IO Module	X91134A	111-03431	2 p 32Gb FC Op	SFP+ Included	FC: 8/16/32 Gbps FC-NVMe: 16/32 Gbps	9.1RC2, 9.2RC1, 9.3RC1, 9.4RC1, 9.5RC1.	8	9, 1, 10, 2, 11, 3, 7, 5, 8
5	Performance Acceleration	IO Module	X9170A	111-00421	Flash Cache,1.0TB NVMe, Module	Not supported		9.1RC2, 9.2RC1, 9.3RC1, 9.4RC1, 9.5RC1.	2	61, 62
6	Performance Acceleration	IO Module	X9172A	108-00423	Flash Cache,4.0TB NVMe, Module	Not supported		9.1RC2, 9.2RC1, 9.3RC1, 9.4RC1, 9.5RC1.	2	61, 62
7	Block Access	IO Module	X91143A	111-02397	4p 16Gb 10Gb UTA2 Cu Op	SFP+ Optional	ETH: 1/10 Gbps FC: 4/8/16 Gbps	9.1RC2, 9.2RC1, 9.3RC1, 9.4RC1, 9.5RC1.	8	9, 1, 10, 2, 11, 3, 7, 5, 8
8	Networking	IO Module	X91143A	111-02397	4p 16Gb 10Gb UTA2 Cu Op	SFP+ Optional	ETH: 1/10 Gbps FC: 4/8/16 Gbps	9.1RC2, 9.2RC1, 9.3RC1, 9.4RC1, 9.5RC1.	8	9, 1, 10, 2, 11, 3, 7, 5, 8
9	Networking	IO Module	X91120A	111-02399	4p 10GBASE-T NIC	Not supported	ETH: 1/10 Gbps	9.1RC2, 9.2RC1, 9.3RC1, 9.4RC1, 9.5RC1.	8	9, 1, 10, 2, 11, 3, 7, 5, 8
10	Storage	IO Module	X91135A	111-03789	4p 32Gb FC SFP+ Op	SFP+ Included	FC: 8/16/32 Gbps FC-NVMe: 16/32 Gbps	9.3RC1, 9.4RC1, 9.5RC1.	8	9, 1, 10, 2, 11, 3, 7, 8
11	Storage	IO Module	X91143A	111-02397	4p 16 Gb 10Gb UTA2 Cu Op	SFP+ Optional	ETH: 1/10 Gbps FC: 4/8/16 Gbps	9.1RC2, 9.2RC1, 9.3RC1, 9.4RC1, 9.5RC1.	8	9, 1, 10, 2, 11, 3, 7, 5, 8

A customer needs to add an additional X91143A card into a FAS9000 to allow mirroring between the production system and the disaster recovery FAS9000 platform. Slots 1, 4, 6, 8, and 9 have cards in them.

Referring to the exhibit, in which slot should you install the card?

A. 5

B. 10

C. 3

D. 7

ANSWER: D

Explanation:

The X91143A card should be installed in slot 7. For a FAS9000, adapter placement must follow the supported PCIe-slot assignment rules for the controller and the adapter type. These rules are not simply a matter of selecting any physically empty slot: they preserve the required PCIe connectivity, bandwidth, and supported I/O layout for the system. Given the occupied slots listed and the placement guidance shown in the exhibit, slot 7 is the appropriate available location for the additional adapter used for the mirroring network connection.

Before installation, the engineer should also verify the precise controller configuration, adapter part number, supported ONTAP release, and any port-use requirements in NetApp Hardware Universe. Hardware Universe is the authoritative source for supported adapter configurations and slot-placement information for a specific platform. After physically installing the adapter, confirm that ONTAP recognizes it, validate cabling and link status, and configure the applicable intercluster or data-protection network connectivity before establishing or resuming the mirror relationship.

See [NetApp Hardware Universe](#) and the [FAS9000 documentation](#) for platform-specific installation and configuration guidance.

QUESTION NO: 10

You are configuring a SnapLock Compliance volume for records that must not be deleted before their retention periods expire.

Which two statements are correct? (Choose two.)

- A. WORM files cannot be deleted before their retention periods expire.
- B. Privileged delete is not supported.
- C. Any cluster administrator can shorten a file retention period.
- D. Ordinary clients can modify a committed WORM file.

ANSWER: A B

Explanation:

Files committed to a SnapLock Compliance volume cannot be deleted before their WORM retention periods expire. This protection applies even to highly privileged ONTAP administrators, helping organizations meet strict regulatory and records-retention requirements.

SnapLock Compliance does not support privileged delete. Privileged delete is a SnapLock Enterprise capability that allows specifically authorized users to remove WORM files before retention expiry under controlled circumstances. Compliance volumes are designed to provide the stricter retention model. See the [SnapLock overview](#) and the [SnapLock WORM file deletion documentation](#).

QUESTION NO: 11

Which two configurations would be used to protect volumes in a production ONTAP system at two different disaster recovery (DR) locations? (Choose two.)

- A. SnapMirror in a fan-out topology
- B. SVM DR in a cascade topology
- C. SnapMirror in a cascade topology
- D. SVM DR in a fan-out topology

ANSWER: A B

Explanation:

SnapMirror in a fan-out topology provides two independently maintained destination copies from the same production source volume. The production volume can therefore replicate directly to a volume at each DR site, allowing either site to be selected for recovery according to the organization's recovery plan. This topology is well suited when both DR locations need current, source-based replication rather than one site depending on the other site's replication path.

SVM DR in a cascade topology provides multistage protection for an entire storage virtual machine, including its data volumes and relevant SVM configuration. The production SVM is replicated to the first DR location, and that DR SVM serves as the source for replication to the second DR location. This establishes recoverable SVM-level copies at two geographically separate locations while retaining an ordered replication path. It is appropriate when the recovery design requires the SVM's configuration and workload data to travel through an intermediate DR site before being protected at a further site.

These ONTAP data-protection topologies enable protection at two DR locations using supported SnapMirror-based replication designs. See the NetApp documentation for [supported data-protection deployment configurations](#) and [SVM disaster-recovery replication](#).

QUESTION NO: 12

Which two cloud providers allow you to provision Cloud Volumes ONTAP as a SnapMirror destination? (Choose two.)

- A. Azure
- B. DigitalOcean
- C. AWS
- D. RackSpace


ANSWER: A C

Explanation:

Azure and AWS are correct because Cloud Volumes ONTAP can operate as an ONTAP SnapMirror destination in both of these public-cloud environments. A source ONTAP system replicates volume Snapshot copies to a destination Cloud Volumes ONTAP system using native SnapMirror relationships, enabling disaster recovery, data migration, and cloud-based secondary storage while retaining ONTAP replication behavior. The destination is provisioned in the selected cloud provider through NetApp BlueXP and configured with the required connectivity, intercluster networking, peer relationships, and destination volumes. NetApp's Cloud Volumes ONTAP documentation identifies both Amazon Web Services and Microsoft Azure as supported deployment providers, and its replication documentation describes replication between on-premises ONTAP and Cloud Volumes ONTAP as well as between Cloud Volumes ONTAP systems. Therefore, either provider can host the Cloud Volumes ONTAP destination required by the question. See the [Cloud Volumes ONTAP documentation](#) and the [BlueXP replication documentation](#) for deployment and SnapMirror replication guidance.

QUESTION NO: 13

Click the Exhibit button.



You have a SnapProtect v11 environment as shown in the exhibit. What are the two default Snapshot copies created automatically after creating a Storage Policy.

(Choose two.)

- A. Primary (Tape)
- B. Primary (Classic)
- C. Primary (NetApp)
- D. primary (Snap)

ANSWER: B D

Explanation:

When a SnapProtect/Commvault storage policy is created for an IntelliSnap-enabled workflow, the system creates the associated default copy definitions needed to protect data through both snapshot and conventional backup mechanisms. **Primary (Classic)** is the default conventional, or classic, primary copy. It represents the copy used for traditional backup data management and can be associated with standard backup storage resources and retention settings. **primary (Snap)** is the default snapshot copy used for snapshot-based protection, including snapshots managed through the storage array integration. These two copy types are created as complementary parts of the storage-policy design: the snapshot copy supports rapid local recovery and snapshot operations, while the classic copy provides the conventional copy framework used by the policy. This default pairing allows the policy to manage snapshot protection alongside classic backup processing without requiring an administrator to manually define both initial copies. Commvault documentation describes storage policies as the mechanism that defines data-copy management and identifies snapshot copies as the storage-policy copy type used for IntelliSnap operations. See [Commvault: Creating a Storage Policy](#) and [Commvault: IntelliSnap](#).

QUESTION NO: 14

For DR purposes, data must be replicated between on-premises clusters and multiple public clouds. In this scenario, which solution would simplify management?

- A. Cloud Volumes ONTAP
- B. Cloud Manager
- C. Cloud Sync
- D. Cloud Insights

ANSWER: B**Explanation:**

Cloud Manager is correct because it provides a centralized management plane for NetApp storage resources across on-premises ONTAP environments and supported public-cloud environments. For a disaster-recovery design requiring replication between clusters and more than one cloud, it simplifies operational management by presenting the relevant working environments and data services through one interface. Administrators can use it to discover or deploy environments, configure and monitor replication relationships, and manage the associated cloud storage resources without treating each location as an entirely separate management domain.

Cloud Manager was the product name used for the NetApp cloud management platform at the time represented by this exam content; the platform is now called BlueXP. Its replication service uses ONTAP SnapMirror technology to support asynchronous replication between compatible ONTAP systems, including on-premises systems and Cloud Volumes ONTAP deployments in cloud providers. This centralized approach is particularly useful for DR because it improves visibility of protection relationships and makes it easier to manage recovery copies distributed across hybrid and multicloud infrastructure. NetApp documents the current BlueXP management platform at [BlueXP documentation](#) and describes its replication capability at [BlueXP replication documentation](#).

QUESTION NO: 15

You have an environment with multiple database servers that are performing scheduled nightly dumps. These dumps vary in both the time to complete and the size of data produced. You need to ensure that there is always enough room in the volume for the data and that a nightly SnapMirror update does not take place until the database dump process is complete.

In this scenario, which NetApp tool ensures that this process is completed every night?

- A. OnCommand Unified Manager
- B. OnCommand Insight
- C. OnCommand Workflow Automation
- D. OnCommand System Manager

ANSWER: C**Explanation:**

OnCommand Workflow Automation is designed to orchestrate and automate multi-step storage workflows that involve dependencies, conditional logic, and scheduling. In this scenario, a workflow can check or provision sufficient capacity for the nightly database dump, invoke or coordinate the dump-related steps, and delay the SnapMirror update until the dump has completed successfully. This is important because both the duration of the dump and the resulting data size are variable; a fixed-time SnapMirror schedule alone cannot reliably account for either condition.

Workflow Automation uses predefined and custom workflows built from NetApp commands, PowerShell commands, and scripts. Administrators can define execution order and conditions so that subsequent operations occur only after prerequisite operations finish. The completed workflow can then be scheduled to run nightly, providing repeatable execution and an auditable automation process across the database and storage environment. NetApp describes Workflow Automation as a solution for automating storage-management tasks and creating workflows that coordinate ONTAP operations. See the [NetApp OnCommand Workflow Automation documentation](#) and the [ONTAP SnapMirror documentation](#).

QUESTION NO: 16

You do not want to impact your ongoing SnapMirror relationship between your primary and secondary site. Which step must be taken prior to creating a FlexClone at the disaster recovery site?

- A. Cascade the SnapMirror destination volume to another SnapMirror destination.
- B. Create a Snapshot copy of the SnapMirror destination volume.

- C. Cascade the SnapMirror destination volume to another SnapVault destination.
- D. Create a Snapshot copy of the SnapMirror source volume and wait for it to transfer.

ANSWER: D

Explanation:

Create a Snapshot copy of the SnapMirror source volume and wait for it to transfer. This ensures that the Snapshot copy used as the FlexClone parent exists on both the primary source and the disaster-recovery destination as part of the SnapMirror replication process. The transferred Snapshot copy is a common Snapshot copy for the relationship, so the destination has a consistent, replicated point-in-time image from which the FlexClone can be created.

Using a Snapshot copy that was created at the source and successfully transferred preserves the normal SnapMirror update workflow. The FlexClone at the DR site is then created as a separate writable volume while the destination volume remains a protected read-only replica and the established replication relationship can continue. Operationally, the administrator should create the source Snapshot copy, run or wait for the scheduled SnapMirror update, verify that the transfer has completed, and then select that transferred Snapshot copy as the clone parent at the destination. NetApp documents FlexClone volume creation and Snapshot-based volume operations in the [ONTAP logical storage management documentation](#); SnapMirror replication behavior and update operations are covered in the [ONTAP data protection documentation](#).

QUESTION NO: 17

A FAS2750 cluster administrator experiences network issues during replication updates that pass over a low- bandwidth WAN link to their DR site. They also replicate locally to another cluster for testing and development. The administrator wants to enable network compression only on the existing WAN Link SnapMirror jobs.

In this scenario, which statement is correct?

- A. Network compression must be enabled when initializing the replica, the current jobs need to be deleted and new ones created.
- B. Network compression is a global option and it cannot be used.
- C. The current DR transfers must be aborted and the option enabled on each job, then resume the transfer.
- D. Network compression is possible but not on a FAS2750.

ANSWER: C

Explanation:

The current DR transfers must be aborted and the option enabled on each job, then resume the transfer. Network compression is configured as a property of an individual SnapMirror relationship, rather than as a cluster-wide setting. This allows the administrator to enable compression only for the relationships traversing the constrained WAN path while retaining the existing, uncompressed local replication relationships used for testing and development. The setting is changed with the `snapmirror modify` command by setting `-is-network-compression-enabled true` for the applicable destination paths. A relationship cannot be modified while a SnapMirror transfer is actively running, so any active update transfer must first be aborted. After the relationship property has been updated, the administrator can resume replication; a new baseline initialization is not required merely to enable network compression. This provides a targeted way to reduce WAN bandwidth consumption, with the expected trade-off of additional CPU processing for compression and decompression. NetApp documents the relationship-level network-compression parameter in the [snapmirror modify command reference](#); transfer interruption and restart behavior is covered in the [ONTAP data protection documentation](#).

QUESTION NO: 18

A customer wants to assign a user to a build-in SnapCenter RBAC Role that will allow them to review and start existing SnapMirror processes only.

In this scenario, which Role will allow the customer to accomplish this task?

- A. App Backup and Clone Admin
- B. Infrastructure Admin
- C. Backup and Clone Viewer
- D. SnapCenter Admin

ANSWER: A

Explanation:

The **App Backup and Clone Admin** role is the appropriate built-in SnapCenter RBAC role for a user who must work with existing data-protection activity, including reviewing and initiating the backup-related operations that drive SnapMirror updates. SnapCenter treats secondary replication as part of an application data-protection workflow: after a protected backup is created, SnapCenter can update the associated SnapMirror relationship according to the configured policy and schedule. A user assigned this role has the application-level backup and clone operational permissions needed to work with those existing protection workflows without being assigned the unrestricted administrative scope intended for full SnapCenter administration.

In practice, this role supports delegated operations for application owners or backup operators. The administrator can also scope the user's access to the relevant resources, resource groups, policies, and storage assets, so the delegated user can operate only the existing protection configuration that they are responsible for. This aligns with the least-privilege approach recommended when assigning SnapCenter roles. NetApp documents the built-in roles and the way RBAC roles and asset assignments control what a user can perform in the [SnapCenter role-based access control documentation](#).

QUESTION NO: 19

Click the Exhibit button.

The screenshot shows the 'Refine Search Criteria' interface in SnapCenter. It features a top navigation bar with 'Sort', 'Policies & Guidelines', 'Generate URL', 'Filters', and 'Clear' buttons. Below this, there are several search criteria panels, each with a title and a list of components with checkboxes. The 'SnapCenter Server' panel has 'SnapCenter Server 4.1.1' checked. The 'Host OS (OS)' panel has 'Microsoft Windows Server 2019 DCE 64-bit, x64' checked. The 'SnapCenter Host Plugin' panel has 'SnapCenter Host Plug-in 4.1.1 for Microsoft Windows' checked. The 'SnapCenter Application Plugin' panel has 'SnapCenter Custom Plug-in 4.1.1 for Microsoft SQL Server' checked. The 'Virtual Storage Console (VMware)' panel has 'NetApp Virtual Storage Console (VMware) 6.2.1' checked. The 'Browser' panel has 'Google Chrome 64.0' checked. The 'SnapCenter Hypervisor Plugin' panel has 'SnapCenter Plug-in 4.1.1 for VMware vSphere' checked.

Referring to the exhibit, what two Search Criteria check boxes need to be used to provide a supported solution for protecting Windows Server 2019 DCE 64-bit, x64 host using local NetApp LUNs with SnapCenter Server 4.1.1? (Choose two.)

- A. SnapCenter Application Plugin
- B. SnapCenter Host Plugin
- C. Browser
- D. Host OS (OS)

ANSWER: B D

Explanation:

Use **SnapCenter Host Plugin** and **Host OS (OS)** as the search criteria. The supported configuration must be validated from the perspective of the host-based SnapCenter component that performs application and host operations, as well as the operating system on which that component is installed. Selecting the SnapCenter host-plugin criterion identifies the relevant compatibility entries for the plug-in deployed to the Windows host. Selecting the host operating-system criterion then narrows those entries to the Windows Server 2019 Datacenter Edition 64-bit (x64) platform specified in the scenario.

This combination is necessary because SnapCenter support is determined by the complete interoperability relationship among the SnapCenter release, the host plug-in, the host operating system, and the connected NetApp storage configuration. It ensures that the Windows host and its SnapCenter plug-in are evaluated as a supported pairing for protection of data residing on local NetApp LUNs. NetApp directs administrators to use the Interoperability Matrix Tool to validate supported product, host, and plug-in combinations before implementation. See the [NetApp Interoperability Matrix Tool](#) and the [SnapCenter 4.1.1 plug-in installation and setup documentation](#).

QUESTION NO: 20

Click the Exhibit button.

Your customer wants to protect a volume and replicate it to a secondary volume. The customer has tried to configure this relationship using OnCommand System Manager. Upon reviewing this newly configured relationship, you discover that the configuration does not satisfy the customer's stated requirements.

Referring to the exhibit, what should you do to ensure that the volume is protected?

- A. Deletethe relationship and recreate it.
- B. Resynchronize the relationship.
- C. Initialize the relationship.
- D. Change the schedule to hourly.

ANSWER: A

Explanation:

Deletethe relationship and recreate it. is the appropriate action because the configured SnapMirror relationship shown in the exhibit does not meet the required protection design. A replication relationship is defined by core settings such as its source and destination volumes, relationship type, replication policy, and direction. Where those foundational characteristics were selected incorrectly during creation, the reliable corrective method is to remove the incorrect relationship and create a new one with the intended protection configuration. The new relationship can then be initialized, establishing the baseline transfer from the primary volume to the secondary volume and allowing later scheduled updates to protect subsequent changes.

When creating the replacement, the administrator should confirm that the destination volume is selected correctly, the required policy is applied, and the appropriate replication schedule is configured before initialization. This produces a relationship aligned with the customer's stated recovery and replication requirements rather than attempting to continue with an incorrectly defined relationship. NetApp describes the required relationship-creation choices and initialization workflow in its [ONTAP replication relationship documentation](#). NetApp also documents the removal process in [Delete a replication relationship](#).

QUESTION NO: 21

A storage administrator is reviewing NAS volume spaceallocation and is concerned that the Snapshot copies that are recently created to protect the volumes are not storing the data from the volumes. One volume contains over 20 GB of data, although the latest SnapShot copy contains only 198 KB of data.

In this scenario, which statement is true?

- A. ONTAP references metadata when creating Snapshot copies instead of copying data blocks.
- B. The Snapshot policy is defined to create the Snapshot copy with only the most recent data changes.
- C. ONTAP uses data from all Snapshot copies when restoring of the volume is necessary.
- D. The Snapshot copies are created by using a background process that takes time to build.

ANSWER: A

Explanation:

ONTAP references metadata when creating Snapshot copies instead of copying data blocks. A Snapshot copy is a read-only, point-in-time image of a volume that is created almost instantaneously. Rather than making a second physical copy of all 20 GB of active file data, ONTAP preserves the volume's block pointers and filesystem metadata as they existed at the instant the Snapshot copy was created. Consequently, the initial Snapshot copy can consume very little space—such as the 198 KB shown—because its primary overhead is metadata, not a duplicate of every data block in the volume.

The active filesystem and its Snapshot copies share unchanged blocks. Snapshot space usage increases only after data is changed or deleted in the active file system: ONTAP must retain the older blocks that are still referenced by the Snapshot copy, while the active filesystem writes new blocks for the changed data. This block-sharing architecture is what enables efficient point-in-time recovery without immediately duplicating the full volume contents. NetApp documents that Snapshot copies are low-overhead because they use pointers to existing data blocks and consume storage as blocks diverge after creation. See the [ONTAP documentation on local Snapshot copies](#) and the [Snapshot copy concepts documentation](#).

QUESTION NO: 22

Which two cloud providers allow you to provision Cloud Volumes ONTAP as a SnapMirror destination? (Choose two.)

- A. Azure
- B. DigitalOcean
- C. AWS
- D. RackSpace

ANSWER: A C

Explanation:

Azure and AWS are correct because Cloud Volumes ONTAP is available in both Microsoft Azure and Amazon Web Services and supports NetApp SnapMirror replication. A Cloud Volumes ONTAP system can be deployed as the destination for SnapMirror relationships originating from an ONTAP system, enabling replication of volumes to cloud storage for disaster recovery, migration, or data-protection use cases. In this configuration, the source system transfers Snapshot copy updates to the destination volume hosted by Cloud Volumes ONTAP, while SnapMirror maintains the replication relationship and recovery point objectives. BlueXP supports creating and managing Cloud Volumes ONTAP working environments in AWS and Azure, including the networking and storage configuration required for replication connectivity. NetApp documents SnapMirror as a supported replication technology between ONTAP environments and Cloud Volumes ONTAP, and lists AWS and Azure among the supported cloud providers for Cloud Volumes ONTAP. See the [Cloud Volumes ONTAP supported configurations](#) and [BlueXP replication documentation](#).

QUESTION NO: 23

You must retain regulated records on a SnapLock volume. A legal investigation requires specific WORM files to remain undeletable after their normal retention periods expire.

Which SnapLock capability satisfies this requirement?

- A. Legal hold

- B. Snapshot reserve
- C. SnapMirror quiesce
- D. Volume autosize

ANSWER: A

Explanation:

A **legal hold** prevents deletion of specified WORM files after their retention periods expire. The legal hold remains in effect until an authorized administrator explicitly releases it, allowing the organization to preserve records for the duration of an investigation, audit, or litigation matter.

Legal hold is separate from the file retention period. It does not alter the original retention setting; instead, it prevents eligible deletion while the hold is active. NetApp documents this capability in [Place files under legal hold](#).

QUESTION NO: 24

You are copying your data and replicating it off site. What is the called?

- A. a backup
- B. a demonstration
- C. a version
- D. a Snapshot


ANSWER: A

Explanation:

a backup is correct because a backup is a separate copy of data maintained so that it can be restored after data loss, corruption, deletion, ransomware, or a site-level failure. Keeping that copy off site is a key backup practice: it separates the recovery copy from the primary location and therefore improves resilience if the production system or its physical site becomes unavailable. In NetApp environments, backup workflows can protect ONTAP data by creating and retaining copies in a separate target, including object storage or another location, with policies controlling retention and recovery. This provides an independent recovery source and supports restoration when the original data is no longer usable. NetApp describes BlueXP backup and recovery as a service that protects ONTAP volume data by creating backup copies in object storage, including options for restoring protected data when needed. See the [NetApp BlueXP backup and recovery overview](#) and the [ONTAP data protection and disaster recovery documentation](#) for related protection concepts.

QUESTION NO: 25

Click the Exhibit button.



You must set up SVM replication for the first time using CLI. This replication will be used for DR purposes. The SVM configuration does not need to be replicated at this time. In this configuration, you will be reusing a default policy.

Referring to the exhibit, what are two missing steps from the SVM replication workflow? (Choose two.)

- A. Exclude specific LIFs and protocols from the replication policy.
- B. Create a relationship with the id-preserve option set to false.
- C. Initialize the newly created SnapMirror relationship.
- D. Initialize the relationship with the id-preserve option set to true.

ANSWER: B C

Explanation:

Because the requirement is to protect the SVM's data for disaster recovery without replicating its configuration, the SnapMirror SVM relationship must be created with identity preservation disabled. The `-identity-preserve false` setting causes the destination to retain its own SVM identity and configuration rather than receiving the source SVM configuration. This is the appropriate choice when only the data protection relationship is needed at this time. A default SnapMirror policy can be specified or allowed as applicable when the relationship is created.

After the relationship is created, it must be initialized. Initialization performs the first baseline transfer, creating the common Snapshot copy and copying the source volumes' protected data to the destination SVM. Until this operation completes, the newly defined relationship does not provide a usable DR baseline. The initialization operation is issued separately with the `snapmirror initialize` command against the destination relationship path.

NetApp's SVM replication procedure describes creating an SVM DR relationship with the required identity-preservation behavior and then initializing the relationship. See the [NetApp SVM replication documentation](#) and the [snapmirror create command reference](#).

QUESTION NO: 26

Using ONTAP System Manager, you are asked to replicate all the volumes contained in a Storage VM to a remote ONTAP 9.7 system for quick disaster recovery.

Which two steps need to be completed to accomplish this task? (Choose two.)

- A. Create a SnapMirror relationship for each volume.
- B. Create a cluster peering relationship.
- C. Create an SVM-DR relationship.
- D. Create an SVM peering relationship.

ANSWER: B C

Explanation:

To protect every eligible volume in a Storage VM as a single disaster-recovery unit, first create a cluster peering relationship. Cluster peering establishes the trusted intercluster connectivity needed for SnapMirror replication traffic between the source and remote ONTAP clusters. The intercluster LIFs and peer relationship provide the communication foundation on which the replication configuration operates.

Next, create an SVM-DR relationship. SVM disaster recovery uses SnapMirror to replicate the Storage VM configuration and its data volumes to a destination Storage VM, providing a consistent failover scope that is much faster to manage than configuring independent protection relationships volume by volume. When configuring SVM DR in ONTAP System Manager, the necessary SVM peer relationship is created as part of the SVM-DR relationship workflow; it is not a separate required step in this scenario. After initialization, scheduled SnapMirror updates keep the destination synchronized and ready for activation during a disaster.

NetApp documents the cluster peer prerequisite and SVM disaster-recovery setup in its [SVM replication relationship documentation](#) and explains the underlying [cluster peering configuration](#).

QUESTION NO: 27

Which NetApp technology helps to protect against ransomware attacks?

- A. SnapLock
- B. SyncMirror
- C. NetApp Volume Encryption
- D. NetApp Aggregate Encryption

ANSWER: A

Explanation:

SnapLock helps protect against ransomware by providing write-once, read-many (WORM) data retention on ONTAP. When files are committed to a SnapLock volume, they cannot be modified, renamed, or deleted until their assigned retention period expires, including by highly privileged users. This immutability creates a protected recovery copy that ransomware cannot encrypt or remove, allowing administrators to restore trustworthy data after an attack.

SnapLock supports compliance-oriented and enterprise-oriented retention use cases, enabling organizations to apply retention controls that match regulatory or operational recovery requirements. It can be used alongside ONTAP Snapshot copies and replication technologies as part of a layered cyber-resilience design: Snapshot copies provide rapid point-in-time recovery, while SnapLock can preserve recovery data against deletion or alteration for a specified period. NetApp documents SnapLock as an ONTAP WORM solution and includes it in its ransomware-protection capabilities. See the [NetApp ONTAP SnapLock documentation](#) and NetApp's [ransomware protection overview](#).