NetApp NS0-165 Dumps

Shared by Alexander on 20-11-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A customer with a NetApp ONTAP based system has experienced a ransomware attack in the past. They need a way to protect critical volume snapshots so that even if an admin account is compromised, the attacker cannot delete or modify those snapshots. Which two ONTAP features protect against snapshot deletion by an administrator account? (Choose two.)

Options:

A- tamperproof snapshots
B- multi-admin verification
C- consistency group snapshots
D- Autonomous Ransomware Protection

Answer:

A, B

Explanation:

Securing backup recovery points against sophisticated ransomware attacks and rogue or compromised administrative credentials requires defense-in-depth protections built directly into the storage operating system. NetApp ONTAP provides two specialized security features engineered specifically to prevent the unauthorized deletion of Snapshot copies:

Tamperproof Snapshots (Statement A): Also known as SnapLock-compliant snapshot locking, this feature applies a WORM (Write Once, Read Many) retention lock directly to volume snapshots. Once locked, the Snapshot copy cannot be deleted, modified, or truncated by anyone---including the storage admin account or cluster superusers---until the mandatory retention period expires.

Multi-Admin Verification / MAV (Statement B): MAV establishes a 'four-eyes' security authorization framework inspired by military two-man rule principles. When sensitive, destructive commands (such as volume snapshot delete or volume delete) are protected by an active MAV rule, an administrator cannot execute the command unilaterally. Instead, the request enters an approval queue requiring explicit cryptographic or credentials-based authorization from a designated quorum of independent approvers before ONTAP executes the operation.

Consistency group snapshots create crash-consistent, point-in-time image groups across multiple volumes for databases, but they can be deleted by administrators if not locked. Autonomous Ransomware Protection (ARP) detects abnormal file encryption heuristics and triggers automated snapshots, but it does not inherently prevent a compromised administrator from issuing deletion commands.

Question 2

Question Type: MultipleChoice

After a SnapMirror relationship is initialized, hourly NetApp ONTAP Snapshot copies on the source volume do not appear on the destination.

What is causing this behavior?

Options:

A- SnapMirror throttle
B- tamperproof Snapshot copies
C- cluster peering
D- SnapMirror labels

Answer:

D

Explanation:

In NetApp ONTAP, modern data replication relationships utilize the unified Extended Data Protection (XDP) SnapMirror engine. Under XDP SnapMirror and SnapVault architectures, the transfer of specific point-in-time Snapshot copies from a source volume to a secondary destination volume is strictly controlled by SnapMirror labels.

When a Snapshot policy creates local Snapshot copies on the source volume (e.g., hourly, daily, or weekly schedules), each scheduled snapshot can be tagged with an optional metadata string called a snapmirror-label. On the destination cluster, the applied SnapMirror policy contains rules that explicitly define which SnapMirror labels should be transferred and how many retention copies to preserve. If the hourly Snapshot copies on the source volume lack a SnapMirror label, or if their label does not match any rule defined within the destination volume's SnapMirror policy, the replication engine skips those snapshots entirely during scheduled transfers. A SnapMirror throttle restricts transfer bandwidth (KBps) but does not selectively filter snapshots. Tamperproof snapshots protect copies from deletion, and a cluster peering failure would block all replication transfers rather than selectively omitting hourly snapshots.

==========

Question 3

Question Type: MultipleChoice

A storage administrator needs to ensure that critical-workload performance is not degraded by competing workloads in a NetApp ONTAP cluster.

Which two actions accomplish this task? (Choose two)

Options:

A- Set a throughput ceiling for a critical workload.
B- Set a throughput ceiling on a competing workload.
C- Set a throughput floor on a competing workload.
D- Set a throughput floor for a critical workload.

Answer:

B, D

Explanation:

NetApp ONTAP Storage Quality of Service (QoS) allows administrators to manage and prioritize system resources (IOPS and MBps bandwidth) across multi-tenant environments and mixed workloads, preventing 'noisy neighbor' problems from impacting business-critical applications. To protect a critical workload from being degraded by competing, lower-priority workloads, an administrator implements two coordinated QoS policies:

Set a throughput floor for a critical workload (Statement D): An administrator assigns an adaptive or fixed QoS minimum throughput limit (floor) to the critical workload. Supported on All-Flash FAS (AFF) and All-SAN Array (ASA) platforms, a QoS floor ensures that the critical workload receives guaranteed minimum I/O performance and prioritized controller scheduling, even during periods of heavy cluster congestion.

Set a throughput ceiling on a competing workload (Statement B): An administrator applies a maximum throughput limit (ceiling) to aggressive, non-critical background workloads (such as bulk data ingest, reporting scans, or dev/test environments). The ceiling throttles the competing workload if it attempts to exceed its allocated bandwidth or IOPS, protecting shared storage resources.

Setting a ceiling on a critical workload (Statement A) would artificially restrict its performance, and setting a floor on a competing workload (Statement C) would prioritize the competing workload over the critical application.

==========

Question 4

Question Type: MultipleChoice

Which two protocols are used to discover details about devices directly connected to a network? (Choose two)

Options:

A- NTP
B- LLDP
C- FTP
D- CDP

Answer:

B, D

Explanation:

In enterprise network administration and storage troubleshooting, discovering adjacent physical switches, intermediate routers, and switch port assignments is critical for verifying cabling topologies and cluster interconnect health. Two Layer 2 neighbor discovery protocols are natively supported and utilized in NetApp ONTAP:

Link Layer Discovery Protocol / LLDP (Statement B): LLDP is an open, vendor-neutral IEEE standard protocol (IEEE 802.1AB) that network devices use to advertise identity, capabilities, port descriptions, VLAN IDs, and system descriptions to directly connected neighbors across local Ethernet networks.

Cisco Discovery Protocol / CDP (Statement D): CDP is a proprietary Cisco Layer 2 discovery protocol that shares similar neighbor discovery capabilities when ONTAP systems are connected to Cisco Nexus or Catalyst switching fabrics.

In ONTAP, administrators can run the command network device-discovery show to inspect adjacent switch names, connected switch interface IDs, chassis capabilities, and MTU settings discovered via LLDP and CDP. Network Time Protocol (NTP) is utilized purely to synchronize system clocks across cluster nodes and directory services. File Transfer Protocol (FTP) is a legacy application-layer protocol for moving files across IP networks; neither protocol performs Layer 2 network neighbor discovery.

==========

Question 5

Question Type: MultipleChoice

Which NetApp ONTAP feature scans for identical blocks within a volume to reduce physical storage usage?

Options:

A- compression
B- deduplication
C- thin provisioning
D- compaction

Answer:

B

Explanation:

NetApp ONTAP includes an advanced suite of storage efficiency technologies embedded directly within the Write Anywhere File Layout (WAFL) filesystem architecture. NetApp deduplication is the specific storage efficiency feature engineered to search for and eliminate duplicate 4KB data blocks across a FlexVol or FlexGroup volume, or across an entire local tier (aggregate).

ONTAP deduplication operates by generating cryptographic digital fingerprints (hashes) for all written 4KB blocks and storing them in an internal fingerprint database. When duplicate fingerprints are detected---either in real time (inline deduplication) before writing to media or during post-process background scans (background deduplication)---ONTAP verifies the block contents. It then retains a single physical block on disk, updates WAFL metadata pointers so that all duplicate logical blocks point to that shared physical block, and releases the redundant physical space back to the storage pool. Compression reduces the physical footprint of individual blocks by encoding repetitive patterns using LZ-based algorithms. Compaction aggregates multiple sub-4KB data chunks into a single 4KB physical block, and thin provisioning allocates storage on demand rather than eliminating duplicate data.

==========

Question 6

Question Type: MultipleChoice

Which action is required before reconfiguring an existing NetApp ONTAP LUN to be converted to use the NVMe protocol?

Options:

A- Unmap the LUN from the igroup.
B- Take the LUN offline.
C- Take the data LIF offline.
D- Take the volume hosting the LUN offline.

Answer:

A

Explanation:

NetApp ONTAP provides an in-place conversion capability that allows administrators to convert an existing SAN LUN (traditionally accessed via Fibre Channel or iSCSI) directly into an NVMe namespace without requiring disruptive host-level data migrations or data copy operations.

Before initiating the LUN-to-namespace conversion command (lun modify -vserver <svm> -path <lun_path> -protocol nvme), ONTAP enforces strict configuration prerequisites. Crucially, the LUN must not be actively mapped to any traditional SCSI initiator group (igroup). If the LUN retains an active igroup mapping, the conversion operation will fail immediately and return an error to prevent inconsistent SCSI and NVMe target states. Therefore, the administrator must unmap the LUN from all igroups using the lun unmap command prior to conversion. Once unmapped and converted to an NVMe namespace, the storage entity can be assigned to an NVMe subsystem and mapped to the host's NQN. Taking the data LIF offline degrades network connectivity for other workloads, and taking the LUN or hosting volume offline is not the prerequisite enforced by the conversion command.

==========

Question 7

Question Type: MultipleChoice

Which ONTAP management UI lets administrators perform drag-and-drop operations?

Options:

A- Active IQ Unified Manager
B- NetApp Console
C- ONTAP System Manager
D- ONTAP CLI

Answer:

B

Explanation:

NetApp Console (the operational control plane interface within the unified NetApp BlueXP SaaS platform) provides an interactive, visual canvas designed to simplify hybrid and multicloud data management. Within the Console canvas, storage environments---including on-premises ONTAP AFF/FAS clusters, Cloud Volumes ONTAP instances in AWS/Azure/GCP, and public cloud object buckets---are represented as interactive visual working environments.

A signature architectural feature of NetApp Console is its intuitive drag-and-drop workflow engine. Administrators can drag one storage working environment directly onto another to instantly launch automated data service wizards. This includes establishing SnapMirror asynchronous replication pipelines, configuring cloud tiering policies to object storage, setting up backup and recovery schedules, and initiating Cloud Sync data transfers. In contrast, ONTAP System Manager is a localized, cluster-specific dashboard that uses standard web forms, tables, and drop-down menus rather than a multi-system drag-and-drop canvas. Active IQ Unified Manager is an alerting and reporting dashboard, and the ONTAP CLI is a text-based terminal interface.

==========

Question 8

Question Type: MultipleChoice

An administrator has presented a thin-provisioned iSCSI LUN to an ESXi host. The administrator deleted data from the datastore in ESXi, but no blocks were freed on the LUN.

Which two steps should you use to resolve this issue? (Choose two)

Options:

A- Use the vmfs unmap command.
B- Rescan the storage from the host.
C- Set the space allocation to enabled.
D- Set the volume to thin-provisioned.

Answer:

A, C

Explanation:

When virtual machines or files are deleted from a VMware ESXi VMFS datastore backed by a thin-provisioned ONTAP LUN, the underlying storage system is unaware that blocks have been discarded unless SCSI UNMAP primitives are exchanged. In thin-provisioned SAN environments, resolving this block reclamation failure requires two essential steps:

Set the space allocation to enabled (Statement C): In ONTAP, the LUN must have space allocation turned on (lun modify -vserver <svm> -path -space-allocation enabled). This setting instructs the ONTAP SCSI target to support and process the T10 SCSI UNMAP primitive, enabling the storage array to deallocate and free underlying 4KB blocks in the volume. If space allocation is disabled, ONTAP ignores deallocation requests.

Use the vmfs unmap command (Statement A): On the VMware ESXi host, the administrator must initiate space reclamation across the VMFS datastore using the command esxcli storage vmfs unmap (or rely on automated VMFS6 reclamation). This triggers ESXi to send SCSI UNMAP commands to the LUN, identifying deleted blocks so ONTAP can free physical storage.

Rescanning the storage adapter discovers newly mapped LUNs or capacity expansions, but does not trigger block reclamation. The volume is already stated as thin-provisioned, so setting the volume type does not resolve host-to-array UNMAP communication.

==========

Question 9

Question Type: MultipleChoice

What must a NetApp ONTAP administrator create in a new SVM so that an NFS client can access a FlexVol volume?

Options:

A- a qtree
B- an export policy rule
C- a POSIX-style access control list
D- an initiator group

Answer:

B

Explanation:

In NetApp ONTAP, client access to Network Attached Storage (NAS) file resources over the Network File System (NFS) protocol is strictly governed by export policies and their associated export rules. An export policy functions as a network firewall and authorization gatekeeper between the client IP network and the ONTAP storage namespace.

When a new Storage Virtual Machine (SVM) and a new FlexVol volume are created, client mount requests will be rejected by default unless the volume is associated with an export policy containing a matching export rule. The export policy rule defines:

The client match criteria (specific IP address, subnet CIDR, or hostname domain).

The allowed NFS protocol versions (NFSv3, NFSv4.0, NFSv4.1).

The access permissions granted (read-only, read-write, or superuser/root access).

Security flavor requirements (sys/AUTH_SYS, krb5, krb5i, krb5p).

Without an explicit export policy rule matching the NFS client's IP address, the ONTAP storage node returns an 'Access Denied' or 'Permission Denied' RPC mount error. A qtree is an optional subdirectory boundary used for quotas. POSIX ACLs control file-level permissions within the filesystem after the mount succeeds, and initiator groups (igroups) are used strictly for SAN block protocols (iSCSI/FC).

==========

Question 10

Question Type: MultipleChoice

Which two default, automatic HA capabilities operate together to reduce and prevent client outages in a NetApp ONTAP FAS cluster? (Choose two)

Options:

A- automatic giveback
B- automatic disk assignment
C- hardware-assisted takeover
D- MetroCluster switchover

Answer:

A, C

Explanation:

In NetApp ONTAP FAS and AFF platforms, High Availability (HA) pairs are architected to deliver continuous data access during hardware faults, software panics, or node reboots. Two foundational, default automatic capabilities work in tandem to minimize client disruption:

Hardware-Assisted Takeover (Statement C): Traditional failover mechanisms rely on heartbeat keep-alives transmitted across the cluster interconnect. If a node suffers a catastrophic controller reboot or power failure, heartbeat loss detection can take several seconds. Hardware-assisted takeover leverages the Service Processor (SP) or Baseboard Management Controller (BMC) to monitor low-level hardware states. If a hardware fault occurs, the SP instantly alerts the surviving partner over the management network, initiating an immediate, accelerated takeover in hundreds of milliseconds and preventing client timeouts.

Automatic Giveback (Statement A): Once the failed or rebooted node successfully reboots, passes Power-On Self-Tests (POST), and reaches a stable state, ONTAP automatically executes an orderly, nondisruptive giveback of storage aggregates to their home owner, restoring normal balanced operations without administrative intervention.

Automatic disk assignment manages drive ownership for unassigned disks but does not arbitrate controller failover outages. MetroCluster switchover is an explicit disaster-recovery failover mechanism used in geographically distributed MetroCluster environments, not a standard local HA pair capability.

==========

Question 11

Question Type: MultipleChoice

Which feature is enabled by default on newly created storage units after upgrading an existing

cluster to NetApp ONTAP 9.18.1 software?

Options:

A- QoS policy
B- ARP/AI
C- SnapMirror
D- encryption

Answer:

B

Explanation:

NetApp ONTAP 9.18.1 introduces substantial enhancements to native cybersecurity, prominent among which is the expansion and default enablement of Autonomous Ransomware Protection with Artificial Intelligence (ARP/AI). While earlier implementations of Autonomous Ransomware Protection required manual provisioning and extended heuristic learning modes, ONTAP 9.18.1 embeds advanced AI-driven behavioral models directly into the file and block storage subsystems.

On newly deployed systems or existing clusters upgraded to NetApp ONTAP 9.18.1 (notably across modern All-Flash and All-SAN Array r2 architectures), ARP/AI is automatically enabled by default on all newly created storage units and volumes following a standard upgrade grace period. ARP/AI continuously analyzes read/write I/O patterns, file entropy metrics, and burst volume mutations in real time to spot zero-day ransomware attacks. Upon detecting abnormal encryption patterns, ARP/AI immediately generates high-priority EMS alerts and automatically takes an un-deletable, clean snapshot to guarantee immediate recovery without data loss. Storage Quality of Service (QoS) policies and SnapMirror data protection require deliberate administrative provisioning and policy assignment. Drive and volume-level encryption depend on configured key managers and storage efficiency templates rather than universal automatic default enablement on all newly provisioned units.

==========

Question 12

Question Type: MultipleChoice

Refer to the exhibit:

[Exhibit Details: Cluster A is expanded by adding a new HA pair. Peering relationships exist between Cluster A and Cluster B across IPspace AB, and between Cluster A and Cluster C across IPspace AC. Following expansion, data replication fails from the new HA pair in Cluster A to both Cluster B and Cluster C.]

After the customer expands Cluster A by adding a new HA pair, they cannot replicate data from the new HA pair to Cluster B or Cluster C.

Which two actions will resolve the problem? (Choose two)

Options:

A- Add two management LIFs to IPspace AB on Cluster A.
B- Add two intercluster LIFs to IPspace AB on Cluster A.
C- Add two intercluster LIFs to IPspace AC on Cluster A.
D- Add two management LIFs to IPspace AC on Cluster A.

Answer:

B, C

Explanation:

In NetApp ONTAP, SnapMirror and SnapVault intercluster replication workflows depend entirely on dedicated intercluster LIFs to establish data transport sessions between peered clusters. A foundational architectural rule of ONTAP cluster peering is full-mesh node connectivity: every individual node in a cluster participating in an intercluster relationship must host at least one active intercluster LIF in each IPspace designated for cluster peering.

When Cluster A is expanded by adding a new High Availability (HA) pair (two physical nodes), the newly added nodes do not automatically receive intercluster LIFs. If volumes hosted on the new HA pair are replicated to destination clusters, replication transfers fail because the new nodes lack endpoints to communicate with Cluster B and Cluster C. Furthermore, because Cluster A uses multi-tenant network isolation via dedicated IPspaces (IPspace AB for Cluster B peering, and IPspace AC for Cluster C peering), each new node in Cluster A must have an intercluster LIF provisioned within each IPspace:

Creating two intercluster LIFs (one per new node) in IPspace AB restores connectivity to Cluster B (Statement B).

Creating two intercluster LIFs (one per new node) in IPspace AC restores connectivity to Cluster C (Statement C).

Management LIFs handle administrative CLI, REST, and System Manager traffic; they play no role in data replication paths.

==========

To Get Premium Files for NS0-165 Visit
https://www.p2pexams.com/products/ns0-165

For More Free Questions Visit
https://www.p2pexams.com/netapp/pdf/ns0-165

