



# QUESTION & ANSWER

HIGHER QUALITY, BETTER SERVICE

**Provide One Year Free Update!**

<https://www.passquestion.com>

**Exam** : **NS0-165**

**Title** : NetApp Data ONTAP  
Administrator Exam

**Version** : DEMO

1.A SAN Administrator is hardening the security of an iSCSI deployment. The security mandate requires Bidirectional (Mutual) CHAP authentication between the Windows Server hosts and the ONTAP storage array.

The administrator configures the ONTAP cluster:

```
cluster1::> vserver iscsi security create -vserver svm_san -initiator iqn.1991-05.com.microsoft:win-host1
-auth-type CHAP -user-name target_user -outbound-user-name initiator_user
```

When the Windows Server attempts to discover the target and log in, the connection is instantly rejected with an "Authentication Failure" error.

Based on the mechanics of Bidirectional iSCSI CHAP, which TWO of the following password/secret constraints and configuration rules must the administrator verify to resolve the login failure? (Choose 2.)

- A. Bidirectional CHAP mathematically disables Asymmetric Logical Unit Access (ALUA), forcing the administrator to manually pin the LUN to a specific physical port.
- B. In the Windows iSCSI Initiator configuration, the "Target secret" field must exactly match the password associated with the outbound-user-name (initiator\_user) configured on the ONTAP array.
- C. The inbound password (used by the host to authenticate to ONTAP) and the outbound password (used by ONTAP to authenticate back to the host) MUST be mathematically identical to satisfy the IPsec mutual trust requirement.
- D. In the Windows iSCSI Initiator configuration, the "Target secret" field must exactly match the password associated with target\_user on the ONTAP array.
- E. ONTAP strictly enforces that the inbound CHAP secret and the outbound CHAP secret MUST be different; using the same password for both directions poses a replay-attack vulnerability and is actively rejected by the WAFL security kernel.

**Answer:** B, E

2.An IT Manager receives an alert that a scheduled SnapMirror transfer failed during the night. The manager logs into ONTAP System Manager and navigates to the Protection dashboard. The GUI displays a generic "Transfer Failed" status indicator on the relationship, but the network team requires the exact error code and timestamp of the network timeout to correlate with a firewall change.

The manager opens an SSH session to the cluster to investigate further.

...

```
cluster1::> snapmirror show -destination-path svm_dest:vol_dest
Progress
Source      Destination Mirror Relationship Total      Last
Path      Type Path      State  Status    Progress Healthy Updated
-----
svm_src:vol DP   svm_dest:vol Snapmirrored Idle    -        false  -
...
```

Which TWO ONTAP CLI commands provide the granular, low-level diagnostic output required to identify the specific network timeout failure that is abstracted in the standard System Manager GUI? (Choose 2.)

- A. volume show -vserver svm\_dest -volume vol\_dest -fields state,status
- B. system node run -node cluster1-01 -command sysstat -c 1
- C. event log show -message-name snapmirror\* -severity Error
- D. network interface show -role cluster-management -status-admin up
- E. snapmirror show -destination-path svm\_dest:vol\_dest -instance

**Answer: C, E**

3.A NAS Administrator configures a new tree quota to restrict an existing qtree (qt\_projectX) to 500GB. However, a developer reports they are still able to write data well beyond this specified limit.

The administrator checks the active quota status via the CLI:

...

```
cluster1::> volume quota show -vserver svm_nas -volume vol_data
```

```
Vserver: svm_nas   Volume: vol_data   State: on
```

```
cluster1::> volume quota report -vserver svm_nas -volume vol_data
```

```
Vserver: svm_nas
```

```

          ----Disk----  ----Files-----  Quota
Volume  Tree   Type  ID    Used Limit  Used  Limit  Specifier
-----  -
vol_data qt_home tree   1     400GB 1TB   12000 -     qt_home

```

(Note: qt\_projectX is missing from the active report output)

...

Which TWO valid administrative actions will force ONTAP to compile and apply the newly created qt\_projectX quota rule to the active filesystem? (Choose 2.)

- A. Re-mount the NFS export on the client workstations to clear stale capacity caches from the Linux kernel.
- B. Reboot the physical storage controller hosting the volume to flush and rebuild the quota rule tables.
- C. Execute the volume quota resize -vserver svm\_nas -volume vol\_data command to hot-load the new rule into memory.
- D. Execute the volume quota off followed immediately by the volume quota on command for the specific volume.
- E. Temporarily migrate the volume to a different aggregate using volume move to trigger a backend quota recalculation.

**Answer: C, D**

4.A Systems Engineer manages a 2-node HA pair (node1 and node2). All data is secured using NVE, and the master keys are housed on an external KMIP server.

A severe network outage completely isolates the HA pair from the external KMIP server. As established, active I/O continues because the keys are cached in RAM.

However, during this exact network outage, node1 suffers a catastrophic hardware panic and reboots. node2 is perfectly healthy.

Based on the architectural interaction between HA takeovers, NVRAM caching, and external key management, what is the exact operational state of the encrypted data residing on node1's disks immediately following the panic?

- A. node1 boots but halts at the bootloader, unable to contact the KMIP server to unlock its root aggregate. node2 performs an HA takeover of node1's data aggregates. Since node2 remained operational, its RAM cache retains the master keys, enabling it to decrypt and seamlessly serve node1's data to clients despite the KMIP outage.
- B. Both nodes are mathematically locked out of the encrypted data. Following node1's panic event, node2 is forced to flush its entire secure RAM cache containing the master keys, thereby destroying its

key copies and halting all cluster I/O operations until the KMIP server network connectivity is restored.  
C. The cluster would drop into a degraded plaintext mode to preserve data availability during the dual failure scenario, which violates the FIPS 140-2 compliance boundary and is not a supported operational state in ONTAP for NVE-secured volumes.

D. node2 successfully takes over node1's aggregates; however, because the WAFL metadata was marked "dirty" due to node1's panic, node2 is required to re-authenticate with the KMIP server to verify the NVRAM journal integrity. This mandatory re-authentication fails due to the network outage, causing the takeover process to abort and the data aggregates to go offline.

**Answer: A**

5. An IT Manager is designing the performance monitoring strategy for a new AFF cluster. The cluster hosts two distinct environments: high-frequency trading (HFT) databases with strict contractual SLAs, and hundreds of standard departmental file shares with unpredictable user behaviors.

The manager must decide how to implement Active IQ Unified Manager (AIQUM) threshold policies to monitor these environments effectively.

Which of the following approaches represent valid, architectural trade-offs when designing AIQUM threshold policies for this mixed environment? (Select all that apply.)

A. Using static thresholds exclusively for all workloads reduces the time AIQUM takes to poll data from ONTAP because historical machine learning calculations on the cluster are bypassed.

B. Combining dynamic and static thresholds on the same HFT database volumes allows administrators to simultaneously catch unusual behavioral deviations and strict contractual SLA breaches.

C. Enabling global dynamic thresholds for the entire cluster consumes significantly less AIQUM server CPU and Memory resources than applying user-defined static thresholds to individual volumes.

D. Applying static thresholds specifically to the HFT databases ensures that alerts are strictly aligned with hard business SLAs (e.g., latency > 2ms), regardless of whether the database's historical baseline behavior is naturally higher.

E. Utilizing dynamic thresholds for the departmental file shares provides excellent anomaly detection without requiring the administrator to manually calculate and guess "normal" latency for unpredictable user workloads.

**Answer: B, D, E**