

NetApp Storage Installation EngineerONTAP Professional Exam

Netapp NS0-185

Version Demo

Total Demo Questions: 13

Total Premium Questions: 136

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Planning	43
Topic 2, Installation	23
Topic 3, Configuration	51
Topic 4, Troubleshooting	19
Total	136

QUESTION NO: 1

An iSCSI LIF fails to migrate during a storage failover test.

In this scenario, why did the LIF not migrate?

- A. A failover port was not defined
- B. LIFs that serve block protocols do not failover**
- C. The failover policy was not configured
- D. LIF failover was not enabled

ANSWER: B

Explanation:

LIFs that serve block protocols do not failover is correct because an iSCSI LIF is a SAN data LIF, and ONTAP does not use automatic LIF failover to preserve SAN connectivity during a controller failover. Instead, SAN hosts are designed to use multipathing software and multiple configured paths to their LUNs. Each path normally uses a separate iSCSI LIF, preferably across both nodes in an HA pair. If a storage failover, controller issue, network-path loss, or other path disruption occurs, the host multipathing stack detects the unavailable path and directs I/O through an active alternate path. This behavior prevents the endpoint identity of a SAN LIF from moving between nodes in the way that NAS LIFs can move for client access continuity. A successful storage failover test for iSCSI should therefore validate that the host has redundant, healthy paths and that its multipathing configuration can continue I/O through the surviving path. NetApp documents that SAN LIFs do not fail over automatically and recommends configuring multiple paths for SAN access. See [NetApp ONTAP SAN Host Configuration documentation](#) and [NetApp LIF configuration documentation](#).

QUESTION NO: 2

A customer has created a volume on an aggregate that does not have a cloud tier attached. The customer now wants to apply a FabricPool tiering policy to move inactive blocks to object storage.

What must be done first?

- A. Attach a cloud tier to the aggregate that hosts the volume.**
- B. Create a SnapMirror relationship for the volume.
- C. Enable thin provisioning on the volume.
- D. Assign a new export policy to the volume.

ANSWER: A

Explanation:

Attach a cloud tier to the aggregate that hosts the volume is correct. A volume can use FabricPool tiering policies only when its local tier is FabricPool enabled through an attached cloud tier.

Creating a SnapMirror relationship, enabling thin provisioning, or changing the volume export policy does not provide an object-store capacity tier for FabricPool.

QUESTION NO: 3

Which tool is used to compare or check the dependencies of individual software and hardware components?

- A. Solution Builder**
- B. FMC_DC

C. Interoperability Matrix Tool (IMT)

D. Net2

ANSWER: C

Explanation:

The **Interoperability Matrix Tool (IMT)** is the NetApp resource used to verify compatibility and dependency requirements among hardware platforms, ONTAP releases, host operating systems, host utilities, protocols, switches, adapters, disk shelves, and other supported components. During a storage installation or upgrade, an engineer uses the IMT to select the intended configuration components and confirm that the precise combination is supported by NetApp. This validation is important because supportability is based on the complete configuration and on specific version combinations, rather than on whether each component is individually supported in isolation. The IMT results also identify applicable support notes, qualification limits, and any required versions or configuration conditions. Checking these dependencies before installation helps ensure that the deployed environment is supported and avoids interoperability issues during implementation. Access the tool directly through the [NetApp Interoperability Matrix Tool](#). NetApp's [ONTAP documentation](#) also directs administrators to validate supported configurations and related requirements before deploying or changing ONTAP systems.

QUESTION NO: 4

You attempt to create a data volume on aggr0 and receive a warning about severe performance or stability problems.

What is the reason for this warning?

- A. Controller failover and storage failover must occur in parallel.
- B. The performance load of a data volume should not share aggr0 with vol0.
- C. The performance load generated by vol0 cannot be shared.
- D. Controller failover and storage failover must occur at different times.

ANSWER: B

Explanation:

The performance load of a data volume should not share aggr0 with vol0. In ONTAP, aggr0 is normally the root aggregate and contains the node root volume, commonly named vol0. That root volume is essential to node operation because it holds configuration and system information needed by ONTAP. Consequently, the root aggregate should be reserved for root-volume and system use rather than used as a general-purpose location for application data volumes.

Creating a data volume on aggr0 can introduce unpredictable workload contention for the disks and aggregate resources supporting vol0. Intensive client I/O, metadata activity, snapshot processing, or capacity pressure from the data volume can affect the responsiveness of the root volume and therefore interfere with important system functions. ONTAP displays the warning to ensure that the administrator consciously acknowledges this operational risk. The standard deployment practice is to create one or more separate data aggregates for user and application volumes, leaving the root aggregate isolated from ordinary production workloads. NetApp's aggregate documentation describes aggregates as the storage containers on which volumes are provisioned, while its volume administration guidance distinguishes system/root usage from normal data-volume management. See [NetApp ONTAP disk and aggregate management documentation](#) and [NetApp ONTAP logical storage management documentation](#).

QUESTION NO: 5

You completed the installation of a new ONTAP cluster. AutoSupport is configured to use HTTPS, and DNS resolution of the NetApp support destination succeeds. The customer firewall permits outbound HTTPS traffic only through an HTTP proxy.

What should you configure to allow AutoSupport messages to be delivered?

- A. Configure the AutoSupport proxy URL.

- B. Configure an SMTP mail host.
- C. Move the cluster management LIF to an intercluster port.
- D. Create a static route through the cluster peer network.

ANSWER: A

Explanation:

Configure the AutoSupport proxy URL is correct. ONTAP must be given the proxy-server information so that HTTPS AutoSupport traffic can be forwarded through the customer's required outbound proxy.

Configuring an SMTP mail host would support email transport rather than the required HTTPS path. Modifying the intercluster LIF or changing DNS does not configure ONTAP to use the proxy.

QUESTION NO: 6

You have completed a base installation and system registration of an AFF A250 system, but you notice that you have one failed drive. Neither un-failing nor reseating the drive repaired it.

In this scenario, which two actions should you take? (Choose two.)

- A. Ask the customer to open a technical case with NetApp Support.
- B. Ask the customer for a drive with the same capacity and replace it.
- C. Run Active IQ Config Advisor to collect all the logs on the system.
- D. Power off the system until the drive has been replaced.

ANSWER: A C

Explanation:

Ask the customer to open a technical case with NetApp Support once the drive remains failed after the standard recovery actions of reseating it and attempting to return it to service. A support case formally records the hardware failure, enables NetApp to validate the reported symptoms and system information, and initiates the appropriate replacement or return-material-authorization process under the customer's support entitlement. Before or while the case is being opened, run Active IQ Config Advisor to collect all the logs on the system. Config Advisor provides a support-ready assessment of the cluster configuration and can collect diagnostic information that helps establish the installed-system baseline and identify any configuration conditions relevant to service activity. Its output gives Support and the installation team useful evidence for efficiently handling the failed-drive condition. These steps preserve normal system availability while ensuring that replacement work is performed through the supported hardware-service process and with the relevant diagnostic data available. See the [Active IQ Config Advisor documentation](#) and NetApp's [Support site](#) for case-management resources.

QUESTION NO: 7

While you are introducing the customer to the NetApp Support website, they ask you to ensure that the new ONTAP storage array is covered by support before you complete the knowledge transfer procedure.

Referring to the exhibit, which two widgets would you select to accomplish this task?(Choose two.)



- A. Active IQ
- B. Product Support
- C. Systems
- D. Contracts and Warranties

ANSWER: C D

Explanation:

Systems and **Contracts and Warranties** are the appropriate Support Site widgets for confirming that a newly installed ONTAP storage array is entitled to support. The Systems widget lets the installer locate and review the registered storage system, typically by serial number or customer site. This establishes that the array is present in the customer's NetApp support inventory and that its system details are associated with the correct account.

After identifying the array, the Contracts and Warranties widget provides the entitlement information needed to verify coverage. It is used to review the support contract or warranty associated with the serial number, including coverage status and relevant service dates. Checking both areas before knowledge transfer gives the customer confidence that the system is registered and has active support entitlement, and it provides an opportunity to address registration or contract-association issues while the installation engineer is still engaged.

NetApp's Support Site is the central portal for managing systems, support cases, contracts, and warranty information: [NetApp Support Site](#). NetApp also provides support and entitlement resources through its customer support portal: [NetApp Support and Training](#).

QUESTION NO: 8

You installed a new 2-node FAS500f ONTAP 9.8 switchless cluster that serves NAS volumes to both NFSv4.1 and SMB clients. The host operating systems have full flow control enabled on their data ports.

In this scenario, which two flow control settings would you verify? (Choose two.)

- A. Verify that the data network ports have full flow control disabled.
- B. Verify that the cluster network ports have full flow control enabled.
- C. Verify that the data network ports have full flow control enabled.

D. Verify that the cluster network ports have full flow control disabled.

ANSWER: C D

Explanation:

Verify that the data network ports have full flow control enabled because flow control should be configured consistently across the end-to-end NAS data path. Since the NFSv4.1 and SMB hosts use full flow control on their data ports, the corresponding ONTAP data ports should also use full flow control. This allows Ethernet pause-frame handling in both directions, helping manage transient congestion and preserving reliable, high-throughput NAS connectivity. ONTAP supports configuring a physical port's flow-control mode, including full flow control, through the network port configuration.

Verify that the cluster network ports have full flow control disabled because cluster interconnect traffic requires low-latency communication between the two controllers. The switchless cluster connections are dedicated to internal ONTAP cluster communication rather than client NAS traffic, and they should use no flow control. This is the expected configuration for cluster ports and avoids pause-frame behavior on the cluster interconnect.

These settings separate the flow-control policy for client-facing data connectivity from the policy required for dedicated cluster networking. See the ONTAP [network port modify command reference](#) for the flow-control setting and NetApp's [flow-control configuration guidance](#).

QUESTION NO: 9

You were dispatched to install a cluster made up of FAS8060 and FAS8080 controllers. There are three FAS8060 controllers and three FAS8080 controllers.

Which solution would provide the largest possible cluster?

- A. a cluster made up of two nodes
- B. a cluster made up of four nodes
- C. a cluster made up of five nodes
- D. a cluster made up of six nodes

ANSWER: B

Explanation:

a cluster made up of four nodes provides the largest supported permanent configuration from the available controllers. An ONTAP cluster is constructed from high-availability pairs, and the two nodes in a normal HA pair must be the same controller model. Therefore, two FAS8060 controllers can form one HA pair and two FAS8080 controllers can form another HA pair. These two pairs can then join the same ONTAP cluster, creating a four-node mixed-platform cluster.

Although ONTAP supports a cluster containing different supported platform models, that flexibility applies between HA pairs; it does not allow the remaining individual FAS8060 and FAS8080 controllers to be added as standalone nodes in a production cluster. Each must have a compatible HA partner. With three controllers of each model, one controller from each model remains unpaired after the two homogeneous HA pairs have been formed. Consequently, the largest valid cluster uses four nodes.

NetApp describes an HA pair as the fundamental availability unit for ONTAP storage and documents the HA architecture at [ONTAP HA pair overview](#). Platform support and permitted configurations should also be verified for the relevant ONTAP release in the [NetApp Hardware Universe](#).

QUESTION NO: 10

You are attempting to add a new ONTAP 9.7 node to your ONTAP 9.8 NetApp cluster. After attempting to add the node, the operation fails as shown in the exhibit.

```
ns0184-98::> cluster add-node-status
```

Node Name	Cluster IP	Status	Error Reason
ns0184-98-01	169.254.111.185	success	
ns0184-98-02	169.254.124.154	success	
ns0184-97	168.254.242.195	failure	Failed to start cluster join on node: 169.254.242.195: Cluster join operation cannot be performed at this time: Cannot join node to the cluster because this is a mixed-version cluster. All nodes in the cluster should be running the same software version before a node is added. To attempt the join anyway, use the "-allow-mixed-version-join" parameter. Resolve the issue, then try the command again.

3 entries were displayed.

What are two nondisruptive ways to add the new node to your ONTAP cluster? (Choose two.)

- A. Revert the cluster from ONTAP 9.8 to ONTAP 9.7.
- B. Re-initialize all the nodes in the ONTAP cluster with ONTAP 9.7.
- C. Upgrade the new node from ONTAP 9.7 to ONTAP 9.8.
- D. Add the node using the `-allow-mixed-version-join` parameter.

ANSWER: C D

Explanation:

Upgrading the new node from ONTAP 9.7 to ONTAP 9.8 is a nondisruptive approach because it brings the node to the same ONTAP release as the existing cluster before the join operation. Matching the cluster version is the standard and preferred state for a newly added node, and the software upgrade can be performed on the unjoined node without affecting data availability or service on the current cluster.

Adding the node using the `-allow-mixed-version-join` parameter is also a supported nondisruptive approach for a controlled, temporary mixed-version condition. The parameter explicitly permits the cluster join when the joining node has a compatible but different ONTAP version. This is useful when an expansion or upgrade workflow requires the node to join before its ONTAP image is aligned. The mixed-version state should be treated as temporary: after the node joins, the environment should be brought back to a uniform ONTAP version according to the supported upgrade path. Neither method requires taking the existing cluster offline or reinitializing its nodes. NetApp documents the mixed-version requirements and the cluster join command parameters in the [mixed-version cluster requirements](#) and the [cluster add-node command reference](#).

QUESTION NO: 11

After installation, the customer requests training on system administration.

Which two responses are appropriate? (Choose two.)

A. Visit the NetApp MySupport site and use Express Guides.

B. Use the Knowledge Base for specific issues.

C. Use NetApp SmartSolve for all questions.

D. Open a support case for training questions.

ANSWER: A B

Explanation:

Visit the NetApp MySupport site and use Express Guides. is appropriate because MySupport provides customers with access to NetApp documentation and support resources, while Express Guides provide concise, task-focused instructions for deploying, configuring, and administering supported NetApp solutions. These guides help a new administrator build practical operational knowledge after the installation handoff and provide a useful starting point for common administration workflows. NetApp's documentation portal is available at [NetApp Documentation](#).

Use the Knowledge Base for specific issues. is also appropriate because the Knowledge Base is intended for targeted technical research. An administrator can search it for documented symptoms, configuration guidance, known issues, troubleshooting procedures, and applicable technical articles when a particular operational question arises. This complements structured documentation by providing issue-specific information that is useful during day-to-day system administration. Customers with appropriate support access can use the NetApp Knowledge Base through [NetApp Knowledge Base](#). Together, these resources support both foundational post-installation learning and efficient self-service research for specific administrative needs.

QUESTION NO: 12

You are preparing to physically install a new AFF A400 cluster into a customer's racks. All equipment has been delivered.

Which three tools do you need to complete this installation? (Choose three.)

A. additional network cables

B. Phillips #2 screwdriver

C. laptop computer

D. slide rail kit

E. cable management tray

ANSWER: B

Explanation:

Phillips #2 screwdriver is the correct choice from the options provided. NetApp's AFF A400 installation instructions identify a Phillips #2 screwdriver as a required tool for physically installing the system, including securing the rail hardware and chassis in the rack. This is a conventional hand tool needed during the mechanical installation process, so it directly fits the question's request for a tool. NetApp distinguishes required installation tools from system hardware and accessories that are shipped with, or connected to, the storage system. A proper preinstallation check should also confirm that the rack location satisfies the system's space, airflow, weight, power, and cable-routing requirements before equipment is mounted. For the complete rack-installation sequence, including rail installation and chassis mounting, follow the AFF A400 documentation and the applicable cabinet or rack guidance. See the [NetApp AFF A400 detailed installation guide](#) and the [NetApp AFF A400 quick installation guide](#).

QUESTION NO: 13

You have completed the installation and initial configuration of a new ONTAP cluster. Before handing the system over to the customer, you want to validate both the installed configuration and the ability to deliver support telemetry.

Which two actions should you perform? (Choose two.)

- A. Run an Active IQ Config Advisor assessment.
- B. Generate a test AutoSupport message.
- C. Wait for a scheduled AutoSupport message without checking transmission status.
- D. Disable AutoSupport after the installation is complete.
- E. Perform an unplanned controller takeover as the only validation step.

ANSWER: A B

Explanation:

Running an Active IQ Config Advisor assessment validates the installed configuration against NetApp rules and best practices. It can identify cabling, HA, shelf, and configuration conditions that should be corrected before handover.

Generating a test AutoSupport message verifies that the cluster can collect and send support telemetry through its configured delivery method. This is preferable to waiting for the next scheduled AutoSupport message because it provides an immediate installation-time validation.