

## Pass Your Next NS0-093 Certification Exam Easily & Hassle Free [Q22-Q36]



## Pass Your Next NS0-093 Certification Exam Easily & Hassle Free [Q22-Q36]

### Pass Your Next NS0-093 Certification Exam Easily & Hassle Free Free Network Appliance NS0-093 Exam Question Practice Exams

Network Appliance NS0-093 (NetApp Accredited Hardware Support Engineer) certification exam is a comprehensive test that evaluates your knowledge of NetApp hardware products and their support processes. If you are planning to pursue a career as a NetApp hardware support engineer, this certification will demonstrate your proficiency and credibility in the field. NS0-093 exam requires a thorough understanding of NetApp storage systems, hardware components, installation, and support methodologies.

**Q22.** Your customer notices that their AFF A220 high-availability (HA) system is showing Single-Path HA.



```
NetApp Release 9.6P5: Fri Dec 13 13:23:47 EST 2019
System ID: 0123456789 (clusterA_01); partner ID:9876543210(clusterA_02)
System Serial Number: 20004123456 (ClusterA_01)
System Rev: E3
System Storage Configuration: Single-Path HA
```

How do you make the system into multipath HA?

- \* 0a and 0b on both controllers should be connected.
- \* 0a and 0c on both controllers should be connected.
- \* 0c and 0d on both controllers should be connected.
- \* 0a and 0d on both controllers should be connected

In a NetApp high-availability (HA) configuration, multipath HA ensures redundancy by providing two independent paths from each controller to the connected storage. If the system is currently showing Single- Path HA, it means only one path is operational or connected, which is a potential single point of failure.

To convert the configuration to multipath HA, additional cabling must be configured to establish the second path. For the AFF A220 system, the correct connections are as follows:

Explanation of the Correct answer:

\* AFF A220 Port Configuration for HA:

\* Each controller in an AFF A220 has ports labeled 0a, 0b, 0c, and 0d for connectivity.

\* In a multipath HA configuration, each controller must have two independent paths to the connected storage via the partner controller's disk shelves.

\* Correct Cabling for Multipath HA:

\* Port 0a on both controllers must be connected to the corresponding port (e.g., disk shelf stack or Inter-Switch Link).

\* Port 0c on both controllers must also be connected to complete the redundant path.

\* These connections ensure that there are two independent data paths between the controllers and the storage system, providing fault tolerance in case one path fails.

\* Single-Path HA to Multipath HA:

\* If the system currently shows Single-Path HA, it is likely that only one set of ports (e.g., 0a and

0b) are cabled.

\* Adding the additional connections on 0a and 0c establishes the second path required for multipath HA.

Why the Other Options Are Incorrect:

\* A. 0a and 0b on both controllers should be connected:

\* This option only connects ports 0a and 0b, which does not establish a multipath HA configuration. Port 0c is also required for redundancy.

\* C. 0c and 0d on both controllers should be connected:

\* Connecting only 0c and 0d would leave ports 0a disconnected, failing to establish a proper multipath HA setup. Both 0a and 0c are required.

\* D. 0a and 0d on both controllers should be connected:

\* Connecting 0a and 0d is not the standard configuration for multipath HA in an AFF A220 system. Port 0c is part of the required configuration.

How to Verify the Configuration:

\* Use the following ONTAP command to verify the HA status:

```
cluster::> storage failover show -fields multipath-ha
```

\* This will display whether the system is in multipath HA or single-path HA mode.

\* Check the port connectivity:

```
cluster::> network port show -node <node_name>
```

References:

\* NetApp Hardware Installation and Setup Guide for AFF A220.

\* NetApp Knowledge Base: [Configuring Multipath HA for AFF Systems](#);

\* ONTAP 9.x System Administration Guide: [High Availability Configurations](#); [Understanding Single-Path HA vs. Multipath HA](#):

\* Single-Path HA indicates that only one path exists between the controllers and the connected shelves, leading to limited redundancy.

\* Multipath HA requires both controllers to have redundant SAS connections to the shelves.

Cabling for Multipath HA in AFF A220 Systems:

- \* To achieve multipath HA, you must connect 0c and 0d on both controllers to the appropriate stack ports on the shelves.
- \* This ensures dual paths for each controller to access the storage shelves.

NetApp Reference Documentation:

- \* Refer to the [ONTAP Multipath HA Configuration Guide](#); and the [NetApp Hardware Installation Guide](#); for AFF A220 systems for cabling diagrams and best practices.

**Q23.** Which type of AutoSupport message would you expect to see triggered automatically when ONTAP software detects a NetApp WAFL inconsistency on an aggregate?

- \* MEDIUM ERROR DURING RECONSTRUCTION
- \* CHECKSUM ERROR
- \* WAFL INCONSISTENT USER DATA BLOCK
- \* WAFL INCONSISTENT BLOCK

When ONTAP detects a WAFL inconsistency in an aggregate, it automatically generates an AutoSupport message with the description WAFL INCONSISTENT BLOCK.

- \* WAFL INCONSISTENT BLOCK:
  - \* This error indicates that WAFL metadata or user data blocks have been found to be inconsistent.
  - \* ONTAP triggers an automatic AutoSupport message to notify administrators and NetApp Support.

Key Details:

- \* A. MEDIUM ERROR DURING RECONSTRUCTION:
  - \* This error occurs during disk reconstruction, not due to WAFL inconsistencies.
- \* B. CHECKSUM ERROR:
  - \* A checksum error indicates a disk-level data integrity issue, not a WAFL inconsistency.
- \* C. WAFL INCONSISTENT USER DATA BLOCK:
  - \* While related, this is not the specific AutoSupport message triggered by ONTAP.

Why Other Options Are Incorrect:

- \* NetApp [ONTAP WAFL Troubleshooting Guide](#); describes AutoSupport messages related to WAFL inconsistencies.

References:

**Q24.** In a SAS stack of shelves, what is the topology of the connection between expander and disk?

- \* arbitrated loop
- \* point-to-point
- \* loop
- \* ring

In a SAS stack of shelves, the connection between the expander and the disk uses a point-to-point topology.

\* Point-to-Point:

\* Each SAS disk in a shelf connects directly to the expander using a dedicated channel. This ensures that communication between the disk and expander is independent of other disks, improving performance and reliability.

\* Why SAS Uses Point-to-Point:

\* SAS (Serial Attached SCSI) eliminates the shared bandwidth limitations of traditional bus architectures (e.g., arbitrated loop or ring) by dedicating a connection to each device.

Key Details:

\* A. arbitrated loop:

\* Arbitrated loop is a topology used in Fibre Channel systems, not SAS.

\* C. loop:

\* SAS does not use loop-based communication; this is typical of older technologies like SCSI Parallel Interface (SPI).

\* D. ring:

\* Ring topology is not used in SAS stacks.

Why Other Options Are Incorrect:

\* [NetApp SAS Shelf and Disk Configuration Guide](#); specifies point-to-point communication between expanders and disks in SAS environments.

References:

**Q25.** What are two valid commands that can be used to trigger an AutoSupport? (Choose two.)

\* `::> autosupport history show-upload-details -node <nodename>`

\* `::> system node coredump upload -node <nodename>`

\* `::> autosupport invoke -node <nodename> -type all`

\* `::> autosupport invoke-core-upload -node <nodename>`

To trigger an AutoSupport message in ONTAP, the following commands are valid:

\* **What it does:** This command manually triggers a complete AutoSupport message of type `all`; This includes logs and system information from all subsystems.

\* **How to use:**

\* Run the command: `autosupport invoke -node <nodename> -type all`

\* Replace `<nodename>` with the name of the node for which you want to generate the AutoSupport message.

\* **Why it's relevant:** This is the primary method for triggering a full AutoSupport message manually. It is commonly used

during troubleshooting to provide comprehensive system data to NetApp Support.

1. ::> autosupport invoke -node <nodename> -type all

\* What it does: This command is specifically used to upload core files (e.g., kernel or user space cores) from a node to NetApp Support for analysis.

\* How to use:

\* Run the command: autosupport invoke-core-upload -node <nodename>.

\* Replace <nodename> with the name of the node for which you want to upload core files.

\* Why it's relevant: If there is a system panic or other critical issue, this command ensures that core files are included in the AutoSupport message for detailed analysis.

2. ::> autosupport invoke-core-upload -node <nodename>

\* A. ::> autosupport history show-upload-details -node <nodename>:

\* This command displays the history of AutoSupport uploads but does not trigger a new AutoSupport.

\* B. ::> system node coredump upload -node <nodename>:

\* This command uploads coredumps directly to a support server but does not trigger an AutoSupport message.

Why Other Options Are Incorrect:

\* [ONTAP 9 AutoSupport Configuration Guide](#); confirms autosupport invoke as a valid command to trigger AutoSupport messages.

\* [ONTAP CLI Reference Manual](#); specifies autosupport invoke-core-upload for core file uploads.

References:

**Q26.** During an ONTAP upgrade, nodeA1 fails to boot. You notice that it is loading the previous version of ONTAP software. What step can be taken to boot the node successfully?

\* From the loader prompt of nodeA1, use setenv GX\_PRIMARY\_KERNEL\_URL to point to the correct version of ONTAP software.

\* From the loader prompt of nodeA1, run boot\_backup to run the correct version of ONTAP software.

\* From the cluster shell of the partner, run system node image modify to point to the correct version of ONTAP software.

\* From the cluster shell of nodeA1, run system node image modify to point to the correct version of ONTAP software.

When a node fails to boot and loads the previous version of ONTAP, it usually indicates that the primary boot location is incorrect. To resolve this:

\* Enter the LOADER Prompt:

\* Interrupt the boot process to access the LOADER prompt.

\* Set the Primary Kernel URL:

\* Use the following command:

Steps to Boot the Correct ONTAP Version:setenv GX\_PRIMARY\_KERNEL\_URL

<URL\_to\_correct\_ONTAP\_version>

\* Boot the Node:

\* Run the boot command to load the correct version.

\* B. From the loader prompt of nodeA1, run boot\_backup:

\* Booting the backup image may result in an outdated version of ONTAP, not the desired version.

\* C. From the clustershell of the partner, run system node image modify:

\* Modifications to the ONTAP image must be done from the LOADER prompt, not the clustershell.

\* D. From the clustershell of nodeA1, run system node image modify:

\* NodeA1 cannot load ONTAP, so this command cannot be run.

Why Other Options Are Incorrect:

\* [ONTAP System Boot Guide](#); explains how to resolve boot issues using the LOADER prompt.

\* NetApp documentation on environment variables includes the use of GX\_PRIMARY\_KERNEL\_URL.

References:

**Q27.** Which statement best describes the functionality of RAID Error Propagation (REP)?

- \* An unrecoverable error on a user data block does not trigger a NetApp WAFL inconsistency.
- \* An unrecoverable error on a metadata block does not trigger a NetApp WAFL inconsistency.
- \* A recoverable error on a user data block triggers a NetApp WAFL inconsistency.
- \* An unrecoverable error on a user data block triggers a NetApp WAFL inconsistency.

Functionality of RAID Error Propagation (REP):

\* RAID Error Propagation (REP) ensures that errors on the disk layer, such as unrecoverable read errors, do not propagate as inconsistencies into the WAFL (Write Anywhere File Layout) metadata layer.

\* WAFL is designed to manage data blocks, and if an error occurs on a user data block, WAFL can identify and handle it gracefully without flagging an inconsistency.

Why Option A Is Correct:

\* REP ensures that errors on a user data block are isolated and do not trigger WAFL inconsistencies.

Instead, the error is handled at the RAID layer or higher without corrupting metadata.

\* This design provides a robust mechanism to protect against data integrity issues caused by disk errors.

NetApp Reference Documentation:

\* Refer to the [ONTAP Storage Management Guide](#); and the [NetApp RAID Guide](#), which explain how RAID and WAFL interoperate and how REP is implemented.

**Q28.** You have upgraded your cluster from ONTAP 9.5 software to ONTAP 9.7 software. After the upgrade, you notice that the upgrade causes issues. You decide to revert to ONTAP 9.5 software.

Which revert step is allowed?

- \* Revert to ONTAP 9.5 software with the netboot procedure.
- \* Revert from ONTAP 9.7 software to ONTAP 9.5 software by using the boot backup procedure.
- \* Revert directly from ONTAP 9.7 software to ONTAP 9.5 software.
- \* Revert from ONTAP 9.7 software to ONTAP 8.6 software and then to ONTAP 9.5 software.

Reverting to a Previous ONTAP Version:

\* Reverting ONTAP software requires specific procedures and is only supported if the ONTAP version being reverted to is compatible with the current cluster configuration.

Why the Netboot Procedure Is Used:

- \* The netboot process loads the desired ONTAP version directly into memory from a network location and reinitializes the system.
- \* This method is used when a full software revert is required and is often necessary for downgrades to a version that cannot be directly reverted to from the running ONTAP version.

Key Notes:

- \* A direct revert is not supported from ONTAP 9.7 to ONTAP 9.5 using standard methods due to system metadata changes.
- \* Always verify compatibility using the [ONTAP Software Version Interoperability Matrix](#).

NetApp Reference Documentation:

\* Refer to the [ONTAP Revert/Downgrade Guide](#); for specific instructions on using the netboot procedure.

**Q29.** In preparation for an upgrade to ONTAP 9.6P8 software, you have uploaded the ONTAP image

96P8\_q\_nodar\_image.tgz onto your local web server with address 192.168.10.10.

Which is correct command syntax to download this image to the cluster image repository?

\* ClusterA::> storage firmware download -node local -package-url

http://192.168.10.10/96P8\_q\_nodar\_image.tgz

\* SP ClusterA-01> cluster image package get -url

http://192.168.10.10/96P8\_q\_nodar\_image.tgz

\* ClusterA::> cluster image package get -url



[http://192.168.10.10/96P8\\_q\\_nodar\\_image.tgz](http://192.168.10.10/96P8_q_nodar_image.tgz)

\* ClusterA::> cluster image package get -url

[http://netapp.com/ontap9/downloads/96P8\\_q\\_nodar\\_image.tgz](http://netapp.com/ontap9/downloads/96P8_q_nodar_image.tgz)

Command Syntax for Downloading ONTAP Image:

\* The cluster image package get command is used to download the ONTAP image from a specified URL to the cluster's image repository.

\* The syntax requires the -url parameter followed by the full URL of the image file.

Why Option C Is Correct:

\* This command syntax directly downloads the image from the local web server at <http://192.168.10.10/> to the ONTAP cluster repository.

\* Other options either refer to incorrect commands or unsupported URLs.

NetApp Reference Documentation:

\* The [ONTAP Software Upgrade Guide](#); provides the exact command and steps for downloading ONTAP images via HTTP or FTP.

**Q30.** What are the two steps to unfail a disk in maintenance mode? (Choose two.)

\* Reseat the disk.

\* Remove it from the Failed Disk Registry.

\* Use the `disk unfail` command.

\* Use the `disk remove` command.

Steps to Unfail a Disk in Maintenance Mode:

\* When a disk is marked as `failed`; it is added to the Failed Disk Registry and marked as unusable by the system.

Unfailing a disk involves removing it from this registry and resetting its state.

Step 1: Remove the Disk from the Failed Disk Registry

\* Use the command: `registry remove <disk>` to delete the disk from the Failed Disk Registry. This clears its `failed` status.

Step 2: Use the Disk Unfail Command

\* Execute the `disk unfail <disk>` command to reset the disk status and make it available for use.

Key Notes:

\* After unfailing, the disk will be treated as a spare and must be reassigned or re-added to an aggregate.

\* If the disk remains failed after these steps, physical hardware issues may need to be addressed.

NetApp Reference Documentation:

\* The [ONTAP Maintenance Mode Guide](#); and [Disk Management Guide](#); provide detailed instructions for handling failed disks and using the disk unfail command.

**Q31.** On an AFF A700 system, a SAS stack is connected to SAS ports 2a and 2b. The system has an additional 4- port SAS card in slot 9.

How should the cabling be corrected for best practices?

- \* Use port 2a and 9a.
- \* Use port 2a and 9b.
- \* Use port 2a and 2c.
- \* Use port 2b and 9d.

Best Practices for SAS Cabling in AFF A700 Systems:

- \* The AFF A700 system has built-in SAS ports (e.g., 2a and 2b) as well as additional SAS ports on optional SAS cards.
- \* To ensure high availability and redundancy, it is recommended to distribute SAS connections across multiple SAS ports from different controllers or slots.

Why Port 2a and 9a Are Recommended:

- \* Port 2a is a built-in SAS port on the AFF A700 system.
- \* Port 9a belongs to the additional SAS card in slot 9.
- \* By connecting the stack using 2a and 9a, you utilize different SAS domains (built-in controller ports and add-on card ports), providing both path redundancy and load balancing.

NetApp Reference Documentation:

- \* [NetApp Hardware Universe](#); and [ONTAP Hardware Installation Guide](#); highlight that SAS cabling for redundancy should leverage different ports, including those from separate SAS controllers or add-on cards.
- \* NetApp's best practice guidelines suggest avoiding connections to the same SAS controller or port group for critical stacks.

**Q32.** Where is a kernel core file stored on a FAS9000 system that is running ONTAP 9.12.1 software?

- \* on the partner root aggregate
- \* on the root aggregate
- \* on the mailbox disk
- \* on the boot device

On a FAS9000 system running ONTAP 9.12.1, the kernel core file is stored on the root aggregate. This is the default location where ONTAP writes kernel core files for system-level failures.

\* The root aggregate is the aggregate that contains the root volume for a given node in the cluster. This aggregate is used for critical system files and logs, including kernel core files.

\* When a kernel panic or other critical failure occurs, the core dump is written to the root aggregate for later analysis by NetApp Support.

Key Details:

- \* A. on the partner root aggregate: The partner root aggregate is not used for storing core files unless explicitly configured (which is not the default behavior).
- \* C. on the mailbox disk: The mailbox disk is used for cluster quorum and configuration information, not for storing core files.
- \* D. on the boot device: The boot device contains ONTAP software and boot files but does not store kernel core dumps.
- \* [ONTAP System Administration Guide](#); specifies that core files are stored on the root aggregate.
- \* [NetApp's Troubleshooting and Diagnostics Guide](#); confirms the default behavior for kernel core file storage.

Why Other Options Are Incorrect:References:

**Q33.** On an AFF A700 system, a SAS stack is connected to SAS ports 2a and 2b. The system has an additional 4- port SAS card in slot 9.

How should the cabling be corrected for best practices?

- \* Use port 2a and 9a.
- \* Use port 2b and 9d.
- \* Use port 2a and 9b.
- \* Use port 2a and 2c.

**Q34.** Which two factors can increase the chances of experiencing a NetApp WAFL inconsistency on an aggregate?

(Choose two.)

- \* ignoring NetApp Active IQ health risks
- \* constantly running at high load
- \* using outdated disk firmware
- \* upgrading frequently
- \* A. Ignoring NetApp Active IQ health risks

\* Explanation:NetApp Active IQ provides proactive monitoring and alerts for system health risks, including risks related to WAFL inconsistencies. Ignoring these alerts can lead to unresolved issues such as metadata corruption, hardware failures, or system misconfigurations, increasing the chances of WAFL inconsistencies.

\* Examples of Active IQ Risks:

- \* Disk health warnings.
- \* High write latency or metadata bottlenecks.
- \* Aggregates running out of space.

\* C. Using outdated disk firmware

\* Explanation:Outdated disk firmware can cause performance issues, drive failures, or incorrect handling of I/O operations, potentially leading to WAFL inconsistencies. Disk firmware updates often include fixes for critical issues that improve reliability and prevent data corruption.

\* Importance: Keeping firmware updated ensures compatibility with ONTAP and minimizes risks of disk-level errors that can propagate to WAFL.

Why the Other Options Are Incorrect:

\* B. Constantly running at high load:

\* Running at high load can impact performance but does not directly cause WAFL inconsistencies.

WAFL is designed to handle high I/O loads efficiently.

\* D. Upgrading frequently:

\* Frequent upgrades, when performed correctly, do not increase the chances of WAFL inconsistencies. However, improperly executed upgrades or skipping recommended steps might introduce risks.

References:

\* NetApp Knowledge Base: Articles on Active IQ health checks and disk firmware management.

\* ONTAP System Administration Guide: Recommendations for preventing WAFL inconsistencies.

**Q35.** Which command can you use to confirm successful completion of an ONTAP upgrade?

\* system controller sp upgrade show

\* cluster image show-update-progress

\* job show -jobtype upgrade

\* system node upgrade-revert upgrade-task -node <nodename> -version <ontap Version>

To confirm the successful completion of an ONTAP upgrade, the cluster image show-update-progress command provides the necessary details.

\* Command:

Key Command: cluster image show-update-progress

\* Displays the status of the cluster upgrade, including which nodes have successfully upgraded and whether the process has completed.

\* A. system controller sp upgrade show:

\* This command displays information about SP (Service Processor) firmware upgrades, not ONTAP upgrades.

\* C. job show -jobtype upgrade:

\* While this command shows upgrade jobs, it does not confirm the completion of the cluster-wide ONTAP upgrade.

\* D. system node upgrade-revert upgrade-task -node <nodename> -version <ontap Version>:

\* This command reverts an upgrade task, not confirms completion.

Why Other Options Are Incorrect:

\* NetApp &#8220;ONTAP Upgrade Guide&#8221; specifies the cluster image show-update-progress command as the primary method to verify upgrade completion.

References:

**Q36.** While performing a health check on a cluster, you notice the following entries in the cluster event log:

```
Thu Jun 18 12:12:09 PDT [nodeA2: disk_server_0: disk.ioMediumError:notice]:  
Medium error on disk 0b.64.18L2: op  
0x88:00000003c1ce73d8:00000200 sector 16136434984 SCST: medium error -  
Unrecovered read error - If the disk is in a RAID group, the subsystem will  
attempt to reconstruct unreadable data (3 11 ff 0) (6955) Disk 0b.64.18L2  
Shelf 64 Bay 18 [NETAPP X381_HLBRE10TSDB NA01]  
S/N [7PGVHKSG] UID [5000CCA2:51CC0C1B:00000000: 00000000: 00000000: 00000000:  
00000000: 00000000: 00000000: 00000000]
```

Referring to the exhibit, which of the following actions do you take?

- \* Sanitize the disk.
- \* Reseat the disk.
- \* Review the current firmware and known issues.
- \* Use diskcopy to copy the disk to a working spare.

Analyzing the Error in the Event Log:

- \* The log entry indicates a medium error on a disk, suggesting an unrecoverable read issue.
- \* While the RAID subsystem attempts to reconstruct unreadable data, this error might indicate a firmware issue or a compatibility problem with the disk.

Why Reviewing Firmware and Known Issues Is Important:

- \* Medium errors can sometimes result from outdated or incompatible firmware.
- \* By reviewing firmware release notes and known issues for the disk model (NETAPP X381\_HLBRE10TSDB in this case), you can identify if this is a known issue and resolve it by updating the firmware.

Other Options:

- \* Sanitize the disk (Option A): Not relevant here, as sanitization is used for secure data erasure.
- \* Reseat the disk (Option B): Useful for addressing hardware seating issues, but not the first step here.
- \* Diskcopy to a spare (Option D): This is a last-resort recovery step and not the primary action.

NetApp Reference Documentation:

\* &#8220;ONTAP Disk Management Guide&#8221; and &#8220;Disk Firmware Release Notes&#8221; detail how to handle medium errors and update firmware.

Becoming a NetApp Accredited Hardware Support Engineer can open up many career opportunities in the IT industry. NetApp storage systems are widely used in businesses of all sizes, and there is a high demand for hardware support engineers who can provide expert support and maintenance for these systems. In addition, the NS0-093 certification is recognized by many employers as a valuable credential that demonstrates a candidate's knowledge and skills in this area.

**Ace NS0-093 Certification with 63 Actual Questions:** <https://www.validbraindumps.com/NS0-093-exam-prep.html>