Download NetApp NS0-093 Exam Dumps Free

Shared by Bradshaw on 17-06-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Which three commands can be used to manually generate a kernel core file? (Choose three.)

Options:

A- SP> system power cycle
B- ::> reboot --node <node_name> --dump true
C- ::> halt --dump true --node <node_name>
D- ::> system node panic --node <node_name>
E- SP> system core
E- SP> system core:
This command displays core file information but does not generate a new core dump.

Answer:

B, C, D

Explanation:

To manually generate a kernel core file in ONTAP, the following commands can be used:

1. ::> reboot --node <node_name> --dump true

What it does: Reboots the specified node and generates a core dump before rebooting.

Example Usage:

reboot --node <node_name> --dump true

2. ::> halt --dump true --node <node_name>

What it does: Halts the specified node and generates a core dump before shutting down.

Example Usage:

halt --dump true --node <node_name>

3. ::> system node panic --node <node_name>

What it does: Forces a panic on the specified node, which triggers a kernel core dump.

Example Usage:

system node panic --node <node_name>

Why Other Options Are Incorrect:

A . SP> system power cycle:

This command reboots the system from the Service Processor but does not generate a core dump.

NetApp 'Kernel Core Dump Guide' explains the use of reboot, halt, and panic commands to trigger core dumps.

ONTAP CLI Reference includes the syntax for these commands.

Question 2

Question Type: MultipleChoice

What are two valid options for uploading a core file from a node that is running ONTAP 9.12.1 software to NetApp for analysis? (Choose two.)

Options:

A- system node autosupport invoke -diagnostic
B- Download via CIFS and upload to upload.netapp.com.
C- Download via Service Processor Infrastructure (SPI) and upload to upload.netapp.com.
D- system node autosupport invoke-core-upload

Answer:

B, D

Explanation:

Options for Uploading Core Files:

Core files are diagnostic dumps created during system failures for analysis by NetApp Support.

They can be uploaded via multiple methods, depending on system configuration and access:

Option B (CIFS Download):

Core files can be downloaded from the node using a CIFS share and then manually uploaded to upload.netapp.com.

This method is useful if automated processes are unavailable or connectivity is limited.

Option D (Autosupport Invoke-Core-Upload):

The command system node autosupport invoke-core-upload automates the process of uploading the core file to NetApp.

It uses the configured Autosupport mechanism to transfer the file to NetApp Support for analysis.

NetApp Reference Documentation:

'ONTAP Autosupport Guide' and 'ONTAP Troubleshooting Guide' provide instructions for manually and automatically uploading core files.

Question 3

Question Type: MultipleChoice

What should be the first step taken after detecting a NetApp WAFL inconsistency?

Options:

A- Run wafliron.
B- Stabilize the system.
C- Force offline the affected aggregate.
D- Run a dump to null.

Answer:

B

Explanation:

When a WAFL inconsistency is detected, the first step is to stabilize the system before attempting any repairs or recovery.

Steps to Handle WAFL Inconsistency:

Stabilize the System:

Ensure that the system is not experiencing ongoing hardware or software issues.

Avoid making changes to the affected aggregate or volume until the system is stable.

Assess the Impact:

Determine which aggregates or volumes are affected and the severity of the inconsistency.

Why Other Options Are Incorrect:

A . Run wafliron:

Running wafliron should only be done after stabilizing the system and under guidance from NetApp Support.

C . Force offline the affected aggregate:

Forcing an aggregate offline without proper assessment can lead to data loss.

D . Run a dump to null:

Dumping data is not relevant for resolving a WAFL inconsistency.

'NetApp WAFL Troubleshooting Guide' specifies that stabilization is the first step after detecting inconsistencies.

NetApp Support documentation provides guidelines for handling WAFL inconsistencies safely.

Question 4

Question Type: MultipleChoice

You are using wafliron on an aggregate named homedir on a production system.

When is the data on homedir available?

Options:

A- when the administrator manually onlines the aggregate
B- when wafliron is started
C- after Phase 4 of wafliron is complete
D- after the mounting phase of wafliron is complete

Answer:

C

Explanation:

When using wafliron to repair WAFL inconsistencies on an aggregate, the aggregate becomes

available after Phase 4 of the wafliron process is complete.

Phases of wafliron:

Phase 1: Initial scan to identify inconsistencies.

Phase 2: Corrects directory and inode structure issues.

Phase 3: Repairs blocks and metadata.

Phase 4: Completes final repairs and verification, after which the aggregate can be made available.

Why Other Options Are Incorrect:

A . when the administrator manually onlines the aggregate:

The aggregate cannot be manually brought online until wafliron completes Phase 4.

B . when wafliron is started:

Starting wafliron does not make the aggregate available; repairs need to be completed first.

D . after the mounting phase of wafliron is complete:

WAFLiron does not have a specific 'mounting phase.' Mounting happens after Phase 4 completes.

'NetApp WAFLiron Troubleshooting Guide' explains the availability of the aggregate after Phase 4.

NetApp Support documentation outlines the phases of wafliron and aggregate recovery.

Question 5

Question Type: MultipleChoice

Which two tools can you use to invoke AutoSupport? (Choose two.)

Options:

A- NetApp Cloud Insights
B- CLI
C- the NetApp Active IQ website
D- the SmartSolve tool

Answer:

B, C

Explanation:

To invoke AutoSupport in ONTAP, the following tools can be used:

1. CLI (Command Line Interface)

How to use: Run the command:

python

Copy code

autosupport invoke -node <nodename> -type all

This triggers AutoSupport to collect and send logs and system information.

2. NetApp Active IQ website

How to use: Log in to the Active IQ portal and use its interface to request an AutoSupport message from the connected ONTAP systems.

Why Other Options Are Incorrect:

A . NetApp Cloud Insights:

This tool is used for monitoring and performance analysis, not for triggering AutoSupport messages.

D . the SmartSolve tool:

SmartSolve is used for case resolution guidance but does not invoke AutoSupport.

'ONTAP AutoSupport Guide' provides instructions for invoking AutoSupport via CLI and Active IQ.

Question 6

Question Type: MultipleChoice

Which type of core file is generated when a node panics?

Options:

A- mgwd core
B- user space core

C- sync core
D- kernel core

Answer:

D

Explanation:

When a node panics in ONTAP, a kernel core file is generated. This core file contains information about the kernel's state at the time of the panic and is essential for debugging system crashes.

Key Details:

A kernel core file is produced during a node panic to capture information about the kernel, memory, and processes that led to the crash.

The core file is stored on the root aggregate by default and can be uploaded to NetApp Support using the autosupport invoke-core-upload command.

Why Other Options Are Incorrect:

A . mgwd core:

This is related to the Management Gateway daemon, which handles management traffic. It does not generate a core file during a panic.

B . user space core:

User space cores are generated for processes running in user space, not for kernel panics.

C . sync core:

Sync cores refer to synchronized cores for debugging but are not the primary type generated during a node panic.

'ONTAP Panic Troubleshooting Guide' specifies kernel core files as the output of a node panic.

'ONTAP Core File Management Guide' details the handling of kernel core files after a crash.

Question 7

Question Type: MultipleChoice

In maintenance mode, Which option best commands shows disk ownership?

Options:

A- sysconfig --v
B- disk show --v
C- aggr status
D- storage show disk --p

Answer:

B

Explanation:

In maintenance mode, the command disk show --v provides detailed information about disk ownership. This includes the current owner, previous owner, and disk health status.

How to Use:

Enter maintenance mode, and run:

disk show --v

The output will display ownership information for each disk, including its RAID group and current assignment.

Why Other Options Are Incorrect:

A . sysconfig --v:

This command provides system configuration details but does not show disk ownership.

C . aggr status:

This command shows the status of aggregates, not individual disks or their ownership.

D . storage show disk --p:

This command is not valid in maintenance mode.

NetApp 'Maintenance Mode Commands Guide' highlights disk show --v as the command to check disk ownership.

Question 8

Question Type: MultipleChoice

At what stage is a write acknowledged to a client?

Options:

A- when the write is present in the local node RAM and NVRAM
B- when the write has been flushed to disk
C- when the write is present in the NVRAM on the local node and its HA partner
D- when the write has been flushed from NVRAM to RAM

Answer:

C

Explanation:

Write Acknowledgment in NetApp ONTAP:

In a clustered ONTAP system, write requests are acknowledged to the client only after they are securely stored in NVRAM on both the local node and its HA (High Availability) partner.

This ensures redundancy and data protection in case of a node failure.

Why Other Options Are Incorrect:

A . when the write is present in the local node RAM and NVRAM:

Writes are not acknowledged until the HA partner also stores the data in its NVRAM.

B . when the write has been flushed to disk:

Writes are acknowledged before they are written to disk, as NVRAM ensures durability.

D . when the write has been flushed from NVRAM to RAM:

Data is not acknowledged based on RAM; NVRAM on both nodes is the requirement.

NetApp's 'ONTAP Write I/O Processing Guide' explains the role of NVRAM and HA in write acknowledgment.

'Data Protection in ONTAP' highlights the synchronization of NVRAM between HA partners.

Question 9

Question Type: MultipleChoice

Which option best scenarios could result in a NetApp WAFL inconsistency in a RAID DP aggregate?

Options:

A- two disks failing and a block error during reconstruction
B- rebooting a node during a disk reconstruction
C- two disks failing within seconds of each other
D- both party disks failing

Answer:

A

Explanation:

A NetApp WAFL (Write Anywhere File Layout) inconsistency in a RAID-DP aggregate could occur in the following scenarios:

1. Two disks failing and a block error during reconstruction

Why this causes inconsistency:

RAID-DP is designed to handle up to two concurrent disk failures. However, if a block error occurs during the reconstruction process (e.g., unreadable data on the surviving disks), the RAID group cannot rebuild the lost data, leading to WAFL inconsistencies.

2. Two disks failing within seconds of each other

Why this causes inconsistency:

If two disks in the same RAID group fail nearly simultaneously (before the RAID-DP can reconstruct data from the first failed disk), the system cannot recover the data, resulting in WAFL inconsistencies.

Why Other Options Are Incorrect:

B . rebooting a node during a disk reconstruction:

Rebooting a node does not cause WAFL inconsistency because ONTAP ensures that RAID reconstructions resume upon reboot without data loss.

D . both party disks failing:

This is not a valid RAID-DP term.

'WAFL and RAID-DP Operations Guide' explains failure scenarios that could cause inconsistencies.

NetApp's 'Troubleshooting RAID Groups and Aggregates' covers recovery procedures for double-disk failures and reconstruction errors.

Question 10

Question Type: MultipleChoice

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

Options:

A- MEDIUM ERROR DURING RECONSTRUCTION
B- CHECKSUM ERROR
C- WAFL INCONSISTENT USER DATA BLOCK
D- WAFL INCONSISTENT BLOCK

Answer:

D

Explanation:

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

Key Details:

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.

Why Other Options Are Incorrect:

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.

NetApp 'ONTAP WAFL Troubleshooting Guide' describes AutoSupport messages related to WAFL inconsistencies.

Question 11

Question Type: MultipleChoice

Which two of the following methods are valid ways to access a node which is not booting? (Choose two.)

Options:

A- node management port
B- cluster management port
C- Service Processor
D- serial console

Answer:

C, D

Explanation:

If a node is not booting, the following methods can be used to access the system for troubleshooting:

1. Service Processor (SP)

What it does: The SP provides out-of-band management access to the node, even if the node is not booting.

How to use:

Connect to the SP using SSH or a direct console connection.

Use SP commands to gather logs or perform diagnostics.

2. Serial Console

What it does: The serial console provides direct access to the node's bootloader or maintenance

mode.

How to use:

Connect to the serial port using a terminal emulator.

Use console commands to interact with the system.

Why Other Options Are Incorrect:

A . node management port:

The node management interface is only accessible if the node is booted and ONTAP is running.

B . cluster management port:

The cluster management interface requires the cluster to be operational, which is not possible if the node is not booting.

NetApp 'Service Processor and Serial Console Guide' provides instructions for accessing a non-booting node.