Number: NS0-093
Passing Score: 800
Time Limit: 120
File Version: 3.0

Exam Code: NS0-093

Exam Name: NetApp Accredited Hardware Support Engineer

Exam A

QUESTION 1
You have created a case with NetApp Support for an issue with a DS4246 shelf on an ONTAP 9.12.1 system. They have requested that you provide shelf logs.
What action do you need to take to collect the shelf logs?

A. Provide the output of the nodeshell command rdfile/etc/log/shelflog.
B. Invoke an autosupport of type all using Active IQ Unified Manager.
C. Invoke a diagnostic AutoSupport with the subsystem storage.
D. Invoke a diagnostic autosupport with the subsystem log_files.

Correct Answer: C
Section:
Explanation:
To collect shelf logs for a DS4246 shelf in an ONTAP 9.12.1 system, you must invoke a diagnostic AutoSupport specifically targeting the storage subsystem. This action ensures that detailed storage logs, including shelf logs, are included in the AutoSupport bundle.
Steps to Collect Shelf Logs:
Use the following command:
bash
Copy code
system node autosupport invoke -node <node_name> -type diagnostic -subsystem storage
Replace <node_name> with the name of the node experiencing the issue.
This command generates an AutoSupport message that includes logs related to storage subsystems, such as disk shelves and adapters.
Provide the AutoSupport case number to NetApp Support for further analysis.
Why Other Options Are Incorrect:
A . Provide the output of the nodeshell command rdfile /etc/log/shelflog:
While this command allows manual reading of shelf logs, it is not a recommended or comprehensive approach for collecting logs for NetApp Support cases.
B . Invoke an AutoSupport of type all using Active IQ Unified Manager:
This action generates a generic AutoSupport bundle, which may not include detailed shelf logs unless explicitly targeted.
D . Invoke a diagnostic AutoSupport with the subsystem log_files:
The log_files subsystem targets general system logs, not storage-specific logs like shelf logs.
'ONTAP 9 AutoSupport and Diagnostics Guide' outlines the use of the subsystem storage option for collecting shelf logs.
The 'Troubleshooting Storage Subsystems' documentation specifies diagnostic AutoSupport as the method for gathering shelf-related logs.

QUESTION 2
What are two valid commands that can be used to trigger an AutoSupport? (Choose two.)

A. ::> autosupport history show-upload-details --node <nodename>
B. ::> system node coredump upload --node <nodename>
C. ::> autosupport invoke --node <nodename> --type all
D. ::> autosupport invoke-core-upload --node <nodename>

Correct Answer: C, D
Section:
Explanation:
To trigger an AutoSupport message in ONTAP, the following commands are valid:

1. ::> autosupport invoke --node <nodename> --type all
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.
2. ::> autosupport invoke-core-upload --node <nodename>
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.
Why Other Options Are Incorrect:
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.
'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.

QUESTION 3
You are replacing a boot device on a FAS8300 system that is running ONTAP 9.10P6 software. You attach a USB memory stick to the external USB port on the storage controller but cannot access the memory stick.
What step needs to be performed to access the boot device?

A. Set the port to ''enabled'' with setenv.
B. You need to use ONTAP 9.11 or later software.
C. Add the boot device before the BIOS is loaded.
D. The external USB port is not activated on NetApp systems.

Correct Answer: A
Section:
Explanation:
When replacing a boot device on a FAS8300 system and using a USB memory stick for recovery or installation, the external USB port must be explicitly enabled. This is done through the setenv command in the boot environment.
Steps to Enable the External USB Port:
Reboot the system and interrupt the boot process to access the bootloader prompt.
At the bootloader prompt, use the following command:
arduino
Copy code
setenv usbport_enabled true
Save the configuration and proceed with the boot process.
Why Other Options Are Incorrect:
B . You need to use ONTAP 9.11 or later software:
ONTAP 9.10P6 fully supports external USB recovery. There is no need to upgrade to ONTAP 9.11 for this functionality.
C . Add the boot device before the BIOS is loaded:
While the USB device must be inserted during the boot process, this alone will not enable access unless the port is enabled via setenv.
D . The external USB port is not activated on NetApp systems:
This is incorrect. The external USB port is supported but must be explicitly enabled in the bootloader environment.
NetApp Hardware Installation Guide for FAS8300 systems outlines the steps for enabling the USB port during recovery.

'ONTAP Boot Troubleshooting Guide' specifies the use of the setenv command to activate USB ports.

QUESTION 4
When you add a new disk to an ONTAP 9.1 system, you see an error that the disk ''has raid label with version (16), which is not within the currently supported range (14-15).'' What is one possible cause of this error?

A. The disk is a non-zeroed spare.
B. The disk firmware is too new for ONTAP software.
C. The disk needs to be assigned to the new system.
D. The disk was in a system with a newer version of ONTAP software

Correct Answer: D
Section:
Explanation:
Explanation of RAID Label Versions:
Disks in ONTAP systems contain metadata known as RAID labels. These labels store critical information such as ownership, aggregate membership, and versioning.
The error indicates a mismatch between the RAID label version on the disk and the supported versions of the ONTAP system.
Why Option D Is Correct:
If a disk has a RAID label created by a newer version of ONTAP (e.g., version 16) and is then inserted into a system running an older version of ONTAP (e.g., supporting versions 14-15), the label will not be recognized.
This mismatch causes ONTAP to reject the disk.
Resolution:
Upgrade the ONTAP version to match the RAID label on the disk.
Alternatively, zero the disk to reset its RAID label, but this will erase all data on the disk.
NetApp Reference Documentation:
The 'ONTAP Disk Management Guide' and 'ONTAP Compatibility Matrix' explain RAID label versions and compatibility issues when moving disks between ONTAP systems.

QUESTION 5
Panic_Message: PCI Error NMI from device(s):ErrSrcID(CorrSrc(0xf00),UCorrSrc(0x18)), RPT(0,3,0):Qlogic FC 16G adapter in slot 1 on Controller.
In which two sections of AutoSupport can you find information to analyze the following panic? (Choose two.)

A. HA-RASTRACE.TGZ
B. ALL-COREDUMP.XML
C. SSRAM-LOG
D. PCI-HIERARCHY.XML

Correct Answer: A, C
Section:
Explanation:
To analyze the provided panic error, the two sections of AutoSupport that are essential for investigation are:
1. HA-RASTRACE.TGZ
What it is: HA-RASTRACE.TGZ contains HA (High Availability) system trace logs. It records hardware diagnostics, error traces, and the HA system's response to hardware events. These logs are critical when analyzing hardware-related panics, including those caused by PCI errors.
Why it's relevant to the panic: In the given panic message, the NMI (Non-Maskable Interrupt) error originates from a Qlogic FC 16G adapter. HA-RASTRACE.TGZ will provide detailed diagnostics, including the error reporting from the HA interconnect and other hardware diagnostics. Specifically, it may include information about how the system detected the PCI fault and any actions taken to protect the system state.
How to analyze:
Extract the HA-RASTRACE.TGZ file from the AutoSupport bundle.
Review hardware-related trace messages for entries associated with the PCI bus or the Qlogic FC adapter.
Look for specific error codes or keywords like PCI Error, NMI, or Qlogic.
NetApp's 'AutoSupport Logs and Diagnostics Guide' highlights HA-RASTRACE.TGZ as a primary resource for debugging hardware faults.
The 'Panic Troubleshooting Guide' for ONTAP systems specifies HA-RASTRACE as a key source for identifying NMI-related errors.

2. SSRAM-LOG
What it is:
SSRAM-LOG records low-level hardware error details, including PCI device register states and uncorrectable memory errors. It is particularly useful for analyzing errors originating in peripheral hardware like network or storage adapters connected via PCI.
Why it's relevant to the panic: The panic message explicitly references a PCI Error NMI caused by a Qlogic FC adapter. SSRAM-LOG captures detailed state information for PCI devices, which can help identify whether the fault originated in the adapter hardware, the PCI bus, or another related component.
How to analyze:
Extract the SSRAM-LOG from the AutoSupport bundle.
Search for PCI-related errors, including the specific error source IDs (e.g., ErrSrcID(CorrSrc(0xf00),UCorrSrc(0x18))).
Review the log entries to confirm the root cause of the NMI.
The 'Hardware Diagnostics and Troubleshooting Guide for ONTAP' lists SSRAM-LOG as a key file for debugging PCI errors.
NetApp's documentation on PCI diagnostics emphasizes the use of SSRAM-LOG for validating hardware-level faults.

QUESTION 6
Which two commands confirm whether an aggregate is WAFL inconsistent? (Choose two.)

A. wafiron show <aggregate>
B. node run --node <node> sysconfig --r
C. storage aggregate show
D. node run --node <node> sysconfig --a

Correct Answer: A, B
Section:
Explanation:
To determine whether an aggregate is WAFL (Write Anywhere File Layout) inconsistent, the following two commands can be used:
1. wafiron show
What it does: This command directly checks the WAFL consistency status of the specified aggregate. If an aggregate is WAFL inconsistent, it will report the inconsistency in the output.
How to use:
Run the command: wafiron show (replace with the name of the aggregate).
Look for indications of WAFL inconsistency in the output.
Why it's relevant: The wafiron utility is specifically designed to provide WAFL status and diagnostics. It is the most accurate and direct way to confirm whether an aggregate is inconsistent.
'WAFL Troubleshooting Guide' from NetApp highlights the wafiron show command as a primary tool for checking aggregate consistency.
2. node run --node <node> sysconfig --r
What it does:
This command displays RAID information for all aggregates on the specified node. If an aggregate is WAFL inconsistent, it will be explicitly mentioned in the output.
How to use:
Run the command: node run --node <node> sysconfig --r.
Check the output for the phrase 'WAFL inconsistent' under the corresponding aggregate.
Why it's relevant: This command provides additional context, such as the RAID group details, which can help understand whether the inconsistency is isolated or part of a larger issue.
'ONTAP CLI Commands Guide' specifies sysconfig --r as a method to verify aggregate status, including WAFL consistency.
Why Other Options Are Incorrect:
C . storage aggregate show:
This command displays aggregate configuration and usage information but does not report WAFL inconsistency.
D . node run --node <node> sysconfig --a:
While this command shows detailed hardware configuration information, it does not include WAFL consistency status for aggregates.

QUESTION 7
What is the default amount of time that a volume is available for recovery from the volume recovery queue following a volume deletion?

A. 12 hours
B. 48 hours
C. 72 hours
D. 24 hours

Correct Answer: A
Section:
Explanation:
When a volume is deleted in a NetApp ONTAP system, it is placed in the Volume Recovery Queue. By default, the volume remains in this recovery queue for 12 hours before being permanently deleted. This allows administrators to recover mistakenly deleted volumes within the set retention period.
Explanation of Default Behavior:
Volume Recovery Queue:
This is a feature in NetApp ONTAP that acts as a safety mechanism, providing a grace period for recovering deleted volumes.
The default retention period for volumes in the recovery queue is 12 hours, as confirmed by the NetApp KB: 'How to use the Volume Recovery Queue.'
How to Recover a Deleted Volume:
Administrators can recover a deleted volume as long as it remains in the recovery queue and the retention period has not expired.
Use the ONTAP CLI command:
arduino
Copy code
cluster::> volume recovery-queue recover -vserver <vserver_name> -volume <volume_name>
This command restores the volume back to its original state.
How to Check the Volume Recovery Queue:
To view volumes in the recovery queue and their expiration times, use:
arduino
Copy code
cluster::> volume recovery-queue show
Changing the Default Retention Period:
While the default period is 12 hours, it can be adjusted by administrators to fit specific organizational requirements. This is done via system settings or policies.
Why the Other Options Are Incorrect:
B . 48 hours: Incorrect. The default retention period is not 48 hours; it is 12 hours by default.
C . 72 hours: Incorrect. While a custom configuration could allow this, it is not the default.
D . 24 hours: Incorrect. Although this was previously thought to be the default, NetApp documentation explicitly states it is 12 hours.
NetApp Knowledge Base Article: 'How to use the Volume Recovery Queue'.
NetApp ONTAP Documentation: Volume Recovery and Data Management Procedures.

QUESTION 8
Which two statements regarding drive 1.2.3.L1 are true? (Choose two.)

A. The drive is in shelf 2.
B. The drive is in bay 3.
C. The drive is in bay 2.
D. The drive is in shelf 1.

Correct Answer: A, B
Section:
Explanation:
The identifier 1.2.3.L1 follows the NetApp disk naming convention, which specifies the location of the drive in the system. Here is the breakdown of the identifier:
1: This indicates the stack ID or loop ID. It represents the stack number in the disk shelf configuration.
2: This indicates the shelf ID. In this case, the drive is located in shelf 2.

3: This indicates the bay ID or slot number within the shelf. The drive is in bay 3.
L1: This represents the logical port or logical disk identifier.
How to Interpret Disk Identifier 1.2.3.L1:
The shelf ID is 2, so the drive is in shelf 2 (A is correct).
The bay ID is 3, so the drive is in bay 3 (B is correct).
Why Other Options Are Incorrect:
C . The drive is in bay 2: The bay ID is explicitly specified as 3, not 2.
D . The drive is in shelf 1: The shelf ID is clearly given as 2, not 1.
NetApp Hardware Universe documentation provides details on disk naming conventions.
The 'ONTAP Disk Management Guide' includes a full explanation of disk IDs and their interpretation.

QUESTION 9
Where is a kernel core file stored on a FAS9000 system that is running ONTAP 9.12.1 software?

A. on the partner root aggregate
B. on the root aggregate
C. on the mailbox disk
D. on the boot device

Correct Answer: B
Section:
Explanation:
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.
Key Details:
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.
Why Other Options Are Incorrect:
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.

QUESTION 10
In the latest MANAGEMENT LOG AutoSupport message, you try to inspect the ENVIRONMENT section but find it empty.
In which section of AutoSupport can you find the reason?

A. AUTOSUPPORT-BUDGET.XML
B. HEADERS
C. AUTOSUPPORT-HISTORY.XML
D. MANIFEST.XML

Correct Answer: A
Section:
Explanation:
If the ENVIRONMENT section of the latest MANAGEMENT LOG AutoSupport message is empty, the reason can typically be found in the AUTOSUPPORT-BUDGET.XML file. This file contains information about AutoSupport resource allocation, including what sections were processed and any limits that were hit.
Key Details:
AUTOSUPPORT-BUDGET.XML:

This file provides a summary of the resources (budget) allocated for different AutoSupport sections.
If the ENVIRONMENT section is missing or empty, the AUTOSUPPORT-BUDGET.XML file will indicate whether it was skipped due to resource constraints or configuration limits.
Why Other Sections Do Not Apply:
B . HEADERS: This section only contains metadata about the AutoSupport message, such as timestamps and node details. It does not explain missing sections.
C . AUTOSUPPORT-HISTORY.XML: This file tracks the history of AutoSupport messages but does not provide information about missing sections.
'ONTAP AutoSupport Troubleshooting Guide' explains the role of the AUTOSUPPORT-BUDGET.XML file in diagnosing missing or incomplete AutoSupport sections.
MANIFEST.XML: This file lists the contents of the AutoSupport bundle but does not provide details on why a specific section is empty.

QUESTION 11
Which type of core file is generated when a node panics?

A. mgwd core
B. user space core
C. sync core
D. kernel core

Correct Answer: D
Section:
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 12
When you plan an ONTAP upgrade, which NetApp tool generates a detailed upgrade plan?

A. ONTAP System Manager
B. Active IQ Unified Manager
C. Upgrade Advisor
D. Active IQ Config Avisor

Correct Answer: C
Section:
Explanation:
The Upgrade Advisor is a NetApp tool that generates a detailed, step-by-step plan for upgrading an ONTAP system. This tool is available through the Active IQ portal and helps ensure a smooth and risk-free upgrade process.
Key Features of Upgrade Advisor:
Provides a tailored upgrade plan based on the current ONTAP version, cluster configuration, and desired target version.
Identifies potential risks, compatibility issues, and pre-requisite tasks for the upgrade.
Offers detailed instructions for each stage of the upgrade process.

Why Other Options Are Incorrect:
A . ONTAP System Manager:
While System Manager can be used to initiate upgrades, it does not generate a detailed upgrade plan.
B . Active IQ Unified Manager:
Unified Manager focuses on monitoring and management but does not provide upgrade plans.
D . Active IQ Config Advisor:
Config Advisor checks for best practices and configuration issues but is not used for generating upgrade plans.
NetApp's 'ONTAP Upgrade Guide' emphasizes the use of Upgrade Advisor for planning upgrades.
Active IQ documentation provides detailed instructions on accessing and using the Upgrade Advisor tool.

QUESTION 13
Which of the following scenarios could result in a NetApp WAFL inconsistency in a RAID DP aggregate?

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

Correct Answer: A
Section:
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 14
Which two statements are true about an IOM 12 module? (Choose two.)

A. It has two SAS ports.
B. It has four SAS ports.
C. It does not have an Ethernet port for alternate control path (ACP).
D. It has an Ethernet port for alternate control path (ACP).

Correct Answer: B, D
Section:
Explanation:
Overview of IOM 12 Module:

The IOM 12 module is used in NetApp storage shelves for SAS connectivity.
Key Features of IOM 12:
SAS Ports: The IOM 12 module has four SAS ports (two IN and two OUT) to support daisy-chaining of shelves and provide redundancy.
ACP (Alternate Control Path): The IOM 12 includes an Ethernet port for ACP, which is used for out-of-band management and monitoring of the storage shelves.
Elimination of Other Options:
Option A is incorrect because the module has four SAS ports, not two.
Option C is incorrect because the module does include an Ethernet port for ACP.
NetApp Reference Documentation:
'NetApp Hardware Universe' lists the specifications of the IOM 12 module, including its SAS and ACP capabilities.
The 'ONTAP Shelf Installation Guide' discusses ACP and its role in shelf management.

QUESTION 15
What is the recommended value for disk and CPU use when you plan an upgrade?

A. less than 50%
B. less than 90%
C. less than 85%
D. less than 70%

Correct Answer: D
Section:
Explanation:
Upgrade Considerations for Disk and CPU Utilization:
During an ONTAP upgrade, it is critical to ensure the system has sufficient resources to handle the upgrade process without impacting normal operations.
Recommended Threshold:
NetApp recommends that both disk and CPU utilization should be below 70% before initiating an upgrade. This ensures that there is enough headroom for the upgrade operations and avoids performance degradation.
Steps to Verify Utilization:
Use the system node show -fields cpu command to check CPU usage.
Use the storage aggregate show -fields used command to check aggregate disk utilization.
NetApp Reference Documentation:
'ONTAP Upgrade and Maintenance Guide' specifies the 70% threshold for disk and CPU usage during upgrade planning.
The 'ONTAP Performance Management Guide' provides methods for monitoring system resource utilization.

QUESTION 16
Which two tools can be used to recover an inconsistent aggregate? (Choose two.)

A. file check
B. wafl_check
C. wafl snapiron
D. wafliron

Correct Answer: B, D
Section:
Explanation:
To recover an inconsistent aggregate, the following tools can be used:
1. wafl_check
What it does: This tool is used to perform a consistency check on WAFL metadata. It identifies and attempts to fix WAFL inconsistencies in aggregates.
When to use: Run wafl_check after identifying WAFL inconsistencies to repair minor metadata issues.
2. wafliron

What it does: This tool repairs WAFL inconsistencies by reconstructing metadata. It is more powerful than wafl_check and should only be run under NetApp Support guidance, as improper use can result in data loss.
When to use: Use wafliron for severe WAFL inconsistencies that cannot be resolved by wafl_check.
Why Other Options Are Incorrect:
A . file check:
This is not a valid NetApp tool.
C . wafl snapiron:
While similar in name, snapiron is used for snapshot recovery, not aggregate recovery.
'ONTAP Aggregate Troubleshooting Guide' details the usage of wafl_check and wafliron.
NetApp Support documentation provides guidelines for recovering inconsistent aggregates.

QUESTION 17
Which two tools can you use to invoke AutoSupport? (Choose two.)

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

Correct Answer: B, C
Section:
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 18
A node has panicked with a PCI/NMI error. Giveback has not been performed.
Which two commands should you run to collect the logs to determine the cause? (Choose two.)

A. pelog --a --g=2
B. show pci --v
C. rdfile /mroot/etc/log/SSRAM
D. event log show

Correct Answer: A, C
Section:
Explanation:

To diagnose a PCI/NMI error and collect logs, use the following commands:
1. pelog --a --g=2
What it does: This command collects PCI error logs, including detailed information about PCI devices and the errors that caused the panic.
How to use: Run the command from the nodeshell to capture the required PCI log entries.
2. rdfile /mroot/etc/log/SSRAM
What it does: This command reads the SSRAM log file, which contains low-level error information related to PCI and other hardware subsystems.
How to use: Run the command to view the log entries directly for detailed troubleshooting.
Why Other Options Are Incorrect:
B . show pci --v:
While this command displays PCI device information, it does not provide detailed error logs.
D . event log show:
This displays event log entries but does not contain the specific PCI or NMI-related logs required for diagnosing the panic.
'ONTAP Hardware Troubleshooting Guide' lists pelog and SSRAM as tools for analyzing PCI errors.
'ONTAP Panic Analysis Guide' emphasizes the importance of collecting detailed hardware logs.

QUESTION 19
Which two steps are required to replace a drawer in a DS460c shelf? (Choose two.)

A. Shut down both nodes.
B. Disconnect the cable chains from the chassis.
C. Power off the shelf.
D. Evacuate all drives in the drawer.

Correct Answer: B, D
Section:
Explanation:
To replace a drawer in a DS460c shelf, the following steps must be taken:
1. Disconnect the cable chains from the chassis
Why this is required: Cable chains connect the drawer to the shelf and must be disconnected to safely remove the drawer.
2. Evacuate all drives in the drawer
Why this is required: Drives must be removed to avoid damage during the drawer replacement process and to reduce the weight of the drawer for safe handling.
Why Other Options Are Incorrect:
A . Shut down both nodes:
This is unnecessary because DS460c shelves support online replacement, and the system can remain operational.
C . Power off the shelf:
This is also unnecessary. DS460c shelves are hot-swappable, meaning they do not require the shelf to be powered down.
'DS460c Hardware Service Guide' explains the procedure for replacing a drawer.
NetApp's 'Field Replacement Guide' for DS460c shelves emphasizes online and hot-swappable replacements.

QUESTION 20
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 the corrected for best practices?

A. Use port 2a and 9a.
B. Use port 2a and 9b.
C. Use port 2a and 2c.
D. Use port 2b and 9d.

Correct Answer: A

Section:
Explanation:
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.

QUESTION 21
You are reviewing the output of disk show and one of the disks is reporting a container type of ''unknown''.
What is causing this status?

A. The disk is not owned by a member of the high-availability (HA) pair.
B. The disk is failed.
C. The disk is in the maintenance center.
D. The disk does not have an owner.

Correct Answer: D
Section:
Explanation:
Understanding 'Container Type: Unknown' in Disk Show Output:
The 'unknown' container type typically indicates that the disk is not properly configured or recognized by ONTAP.
This status often occurs when a disk does not have an owner assigned.
Root Cause:
For a disk to be used in an ONTAP system, it must be owned by a member of the high-availability (HA) pair.
If no ownership is assigned, the disk will not be initialized, resulting in an 'unknown' container type.
Steps to Resolve:
Use the disk assign command to manually assign ownership of the disk.
Example: storage disk assign -disk <disk_name> -owner <node_name>
NetApp Reference Documentation:
'ONTAP Disk Management Guide' explicitly states that unowned disks report 'unknown' container type until they are assigned to a node.
This is further detailed in the 'ONTAP Troubleshooting Guide' under disk configuration issues.

QUESTION 22
Which LOADER prompt command ensures that POST is done on boot?

A. setenv POST=true
B. bye
C. boot_diag
D. boot_ontap

Correct Answer: A
Section:
Explanation:

To ensure that POST (Power-On Self-Test) runs on boot, the setenv POST=true command is used at the LOADER prompt. This command enables the system to perform POST diagnostics before proceeding with the boot process.
Key Details:
POST Purpose: POST checks system hardware components (such as memory, disk, and controllers) for faults before loading the ONTAP kernel.
How to Use:
At the LOADER prompt, type:
arduino
Copy code
setenv POST=true
Save the configuration and reboot the system.
Why Other Options Are Incorrect:
B . bye:
This command restarts the system but does not ensure that POST runs on boot.
C . boot_diag:
This command boots the system into diagnostic mode but is not directly related to enabling POST on boot.
D . boot_ontap:
This command boots ONTAP but skips POST if it is not explicitly enabled.
NetApp 'ONTAP System Boot and Recovery Guide' describes setenv POST=true for enabling POST diagnostics.

QUESTION 23
In a SAS stack of shelves, what is the topology of the connection between expander and disk?

A. arbritrated loop
B. point-to-point
C. loop
D. ring

Correct Answer: B
Section:
Explanation:
In a SAS stack of shelves, the connection between the expander and the disk uses a point-to-point topology.
Key Details:
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.
Why Other Options Are Incorrect:
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.
'NetApp SAS Shelf and Disk Configuration Guide' specifies point-to-point communication between expanders and disks in SAS environments.