NetApp

NS0-093 Exam

NetApp Accredited Hardware Support Engineer

Questions & Answers
Demo

Version: 4.0

Question: 1

In which two sections of AutoSupport can you find information to analyze the following panic? (Choose two.)
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

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

Answer: A, C

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.
Reference:
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.
Reference:
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: 2

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

Answer: A, B

Explanation:

To determine whether an aggregate is WAFL (Write Anywhere File Layout) inconsistent, the following two commands can be used:

1. wafiron show <aggregate>
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 <aggregate> (replace <aggregate> 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.
Reference:
"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.
Reference:
"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: 3

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

Answer: A

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.
Reference:
NetApp Knowledge Base Article: "How to use the Volume Recovery Queue".
NetApp ONTAP Documentation: Volume Recovery and Data Management Procedures.

Question: 4

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.

Answer: A, B

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.
Reference:
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: 5

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

Answer: B

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.
Reference:
"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.