Boost up Your Certification Score

NetApp
NS0-594
NetApp Certified Support Engineer - ONTAP Specialist (NCSE ONTAP Specialist)



For More Information – Visit link below:

https://www.examsboost.com/

Product Version

✓ Up to Date products, reliable and verified.
✓ Questions and Answers in PDF Format.

Visit us at: https://www.examsboost.com/test/ns0-594

Latest Version: 6.1

Question: 1

An SVM's root volume is protected by load-sharing mirrors that are updated by a daily schedule. This morning an administrator created a volume and mounted it directly under the SVM root. The mirror relationship history shows that the load-sharing mirror set has not been updated since before the volume was mounted.
Several hours later, NFS clients that list the SVM root do not see the new junction, and fresh mounts of the new path from clients that have never mounted it before fail with "no such file or directory". In another SVM without load-sharing mirrors, new volumes appear to clients immediately.
What is causing the behavior?

A. The export policy on the new volume denies these clients access to the new path
B. The SVM root is read from load-sharing mirrors not yet updated since the junction was created
C. The new volume's security style differs from the root volume's, which hides its junction from clients
D. The clients' NFS attribute cache hides the new directory entry until each of them remounts the SVM root

Answer: B

Explanation:
With load-sharing mirrors on the root volume, clients read the namespace root from the mirrors, so a new junction appears only after the mirror set is updated.

Question: 2

How does ONTAP Select support disaster-recovery operations in a virtualized environment?

A. By providing hardware-assisted failover only
B. Through rapid deployment of ONTAP on hypervisors for recovery
C. By removing the need for SnapMirror replication
D. By requiring dedicated HA storage arrays

Answer: B

Explanation:
Because ONTAP Select is a software instance, it can be stood up quickly on commodity hypervisor hosts at a recovery site, enabling DR without purpose-built storage hardware.

Question: 3

While reviewing a busy node, an engineer notices the system is running persistent back-to-back consistency point events with almost no idle gap between them.
What does this pattern most directly indicate?

A. A network outage is interrupting client sessions to the node

B. The system cannot finish one consistency point before the next begins
C. A name-service failure is delaying user authentication requests
D. A Snapshot schedule error is preventing scheduled copies

Answer: B

Explanation:
Back-to-back CPs mean the write path never drains, which points at a CPU or disk bottleneck in flushing writes to storage.

Question: 4

Which condition can cause an ONTAP node panic?

A. WAFL metadata inconsistency
B. Incorrect CIFS share permissions
C. Failed NDMP backup
D. Offline volume

Answer: A

Explanation:
A detected WAFL metadata inconsistency can force the node to panic to protect data integrity, which is why the event is captured in a core for RCA.

Question: 5

A customer reports that while Windows updates were being installed on their external Vscan servers, clients could not access any files on a CIFS SVM. Access returned to normal once the updates finished, and the customer wants to prevent this during the next patch window.
What are two ways to accomplish this?
(Choose two.)

A. Enable the -scan-mandatory flag in the Vscan policy.
B. Update the Vscan servers one at a time.
C. Modify the CIFS shares to be continuously available (CA) shares.
D. Disable the -scan-mandatory flag in the Vscan policy.

Answer: B,D

Explanation:
Patching the Vscan servers in a rolling fashion keeps at least one scanner available so CIFS access continues throughout the maintenance window.
With mandatory scanning disabled, ONTAP allows file access to continue even when no Vscan server is reachable, avoiding the access outage.

Question: 6

A customer moved a workload from a single-node Cloud Volumes ONTAP system to a Cloud Volumes ONTAP HA pair whose nodes are in two different availability zones. The instance type and cloud disk type are unchanged.
Read latency is the same as before, but write latency has risen. CPU is not saturated and no throttling is in place.
What explains the higher write latency?

A. Each write is mirrored to the partner in the other zone before it is acknowledged, adding inter-zone round-trip time
B. The mediator inspects each write before it is committed, which adds a network hop to every write the pair acknowledges
C. HA deployments place the aggregates on a slower cloud disk tier by default
D. Clients now connect through a floating IP that sends their writes through a NAT gateway

Answer: B

Explanation:
The pair mirrors each write to the partner before acknowledging it, and across zones that mirror travels a longer network path that reads never use.

Question: 7

A node panic message references a WAFL inconsistency.
What is the correct support action?

A. Immediately reinitialize the aggregate
B. Run wafliron in maintenance mode without review
C. Preserve the core file and open a support case
D. Delete and recreate affected volumes

Answer: C

Explanation:
Preserving the core for root cause analysis and engaging NetApp Support is the correct, non-destructive action.

Question: 8

After a node panic, what is the first recommended action?

A. Review EMS and panic string
B. Immediately reboot the node
C. Delete core files
D. Run wafliron

Answer: A

Explanation:

Examining the EMS log and the panic string is the first RCA step, because it identifies the fault before any disruptive recovery is attempted.

Question: 9

A Linux administrator working as root on an NFS client gets "Permission denied" when creating files in an exported directory that is owned by root with mode 755. On the same client, an ordinary user can create files in a directory that the user owns.
Two further observations: files that root creates in a world-writable directory on the same export appear on the storage owned by the anonymous UID rather than by 0, and root on the same client can write normally to a volume in another SVM.
The export policy rule that matches the client allows read-write access with the sys security flavor, and the volume uses UNIX security style.
What is the most likely cause?

A. Name services cannot resolve UID 0, so ONTAP rejects all requests from root
B. The rule's read-write setting leaves out the client's IP subnet
C. The rule's superuser setting maps root to anonymous
D. NFSv3 cannot carry root credentials to an ONTAP server at all

Answer: C

Explanation:
Without superuser access for the client's flavor, root is mapped to the anonymous user, which cannot write to a root-owned 755 directory.

Question: 10

After you replaced the motherboard, the node stopped in the waiting for the cluster application to come online state.
Which two actions should you perform to troubleshoot this problem?
(Choose two.)

A. Check the status of the HA with the interconnect-up field of the storage failover show command.
B. Verify the setting of the cluster interconnect switch.
C. Verify that the physical links on the cluster ports are up.
D. Check the status of the cluster with the cluster ring show command.

Answer: C,D

Explanation:
The cluster applications come online only after the node reaches its peers over the cluster network, so the cluster ports on the replacement board must have physical link.
cluster ring show reveals the replicated database (RDB) ring and quorum state, which determines whether the cluster application can come online on the rejoining node.

Visit us at: https://www.examsboost.com/test/ns0-594

Thank You for Trying Our Product

For More Information – Visit link below:

https://www.examsboost.com/

15 USD Discount Coupon Code:

G74JA8UF

FEATURES

- ✓ 90 Days Free Updates
- ✓ Money Back Pass Guarantee
- ✓ Instant Download or Email Attachment
- ✓ 24/7 Live Chat Support
- ✓ PDF file could be used at any Platform
- ✓ 50,000 Happy Customer

