

Netapp

NS0-593 Exam

NetApp Certified Support Engineer ONTAP Specialist

Question: 1

When you review performance data for a NetApp ONTAP cluster node, there are back-to-back (B2B) type consistency points (CPs) found occurring on the root aggregate.

In this scenario, how will performance of the client operations on the data aggregates be affected?

- A. During B2B processing, clients will be unable to write data.
- B. Data aggregates will not be affected by B2B processing on another aggregate.
- C. During B2B processing, all I/O to the node is stopped.
- D. During B2B processing, clients will be unable to read data.

Answer: B

Explanation:

Question: 2

Recently, a CIFS SVM was deployed and is working. The customer wants to use the Dynamic DNS (DDNS) capability available in NetApp ONTAP to easily advertise both data UFs to their clients. Currently, DNS is only responding with one data LIF. DDNS is enabled on the domain controllers.

```
vservers      lif      data-protocol  is-dns-update-enabled
-----
svm1          cifs_01  nfs,cifs      true
svm1          cifs_02  cifs          true
svm1          mgmt     none          false
3 entries were displayed.
```

```
cluster1::*> vservers services dns dynamic-update show
Vservers      Is-Enabled  Use-Secure  Vservers FQDN      TTL
-----
svm1          false      false      svm1.demo.net      24h
```

Referring to the exhibit, which two actions should be performed to enable DDNS updates to work? (Choose two.)

- A. Disable the -vservers-fqdn parameter for the SVM DDNS services.
- B. Remove the NFS protocol from the cifs_01 data LIF.
- C. Enable the -use-secure parameter for the SVM DDNS services.
- D. Enable the -is-enabled parameter for the SVM DDNS services

Answer: A, D

Explanation:

Question: 3

A customer is calling you to troubleshoot why users are unable to connect to their CIFS SVM.

```
ClusterB::*> storage disk show -broken

Original Owner: Node03
Checksum Compatibility: block

Physical                               Drawer          Usable
Disk
Chan  Pool  Type  Outage Reason  HA Shelf Bay /Slot  Size
-----
--
      1.0.2      failed      3b    0    2  -/-    B
FAILED BSAS  7200  1.62TB  1.62TB

ClusterB::*> cluster ring show
Node      UnitName Epoch  DB Epoch DB Trnxs Master  Online
-----
Node03    mgmt     11      11      4875  Node04 secondary
Node03    vldb     0        11      358   -      offline
Node03    vifmgr   11      11      4892  Node04 secondary
Node03    bccmd    11      11      62    Node04 secondary
Node03    crs      11      11      6     Node04 secondary
Node04    mgmt     11      11      4875  Node04 master
Node04    vldb     0        11      358   -      offline
Node04    vifmgr   11      11      4892  Node04 master
Node04    bccmd    11      11      62    Node04 master
Node04    crs      11      11      6     Node04 master
10 entries were displayed.

ClusterB::*> system node run -node Node04 -command aggr status -r aggr2
Aggregate aggr2 (online, raid_dp, degraded) (block checksums)
Plex /aggr2/plex0 (online, normal, active, pool10)
RAID group /aggr2/plex0/rg0 (degraded, block checksums)

RAID Disk Device          HA  SHELF BAY CHAN Pool Type  RPM  Used (MB/blks)  Phys
(MB/blks)
-----
-----
      parity  FAILED
      parity  3c.0.11      3c    0    11  SA:B  0  BSAS  7200  2538546/ -
2543634/5209362816
      data    3c.0.12      3c    0    12  SA:B  0  BSAS  7200  2538546/5198943744
2543634/5209362816
      data    3c.0.13      3c    0    13  SA:B  0  BSAS  7200  2538546/5198943744
2543634/5209362816
      data    3c.0.14      3c    0    14  SA:B  0  BSAS  7200  2538546/5198943744
2543634/5209362816
```

Referring to the Information shown in the exhibit, what is the source of the problem?

- A. The v1db database is offline.
- B. The aggregate aggr2 has a failed disk.
- C. The databases on Node03 must be switched from secondary to master.
- D. The broken disk in Node03 is the source of the problem.

Answer: C

Explanation:

Question: 4

You have a customer who is concerned with high CPU and disk utilization on their SnapMirror destination system. They are worried about high CPU and disk usage without any user operations. In this situation, what should you tell the customer?

- A. Suggest that the customer manually cancel any scanners on the destination to reduce CPU usage.

- B. Explain that background tasks such as SnapMirror throttle up in the absence of user workload.
- C. Suggest that the customer throttle their SnapMirror relationships to reduce resource consumption.
- D. Explain that only user workload should use the CPU and Investigate further.

Answer: A

Explanation:

Question: 5

You are attempting to connect a NetApp ONTAP cluster to a very complex network that requires LIFs to fail over across subnets.

How would you accomplish this task?

- A. Configure an equal number of UFs on each subnet.
- B. Configure VIP LIFs using OSPF.
- C. Configure VIP LIFs using BGP.
- D. Configure a LIF failover policy for each subnet inside a single broadcast domain.

Answer: C

Explanation: