NetApp
NS0-593
NetApp certified support engineer - ONTAP specialist Exam

For More Information – Visit link below:
https://www.examsempire.com/

Product Version

1. **Up to Date products, reliable and verified.**
2. **Questions and Answers in PDF Format.**

Latest Version: 6.1

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 loot 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:
A B2B type consistency point (CP) occurs when a new CP is triggered before the previous CP is completed, due to the second memory buffer reaching a watermark. This can cause write latency to increase as user write operations are not replied until a write buffer frees up. However, this only affects the aggregate that is undergoing the B2B processing, and not the other aggregates on the same node. Therefore, the performance of the client operations on the data aggregates will not be affected by B2B processing on the root aggregate. Reference = What is the Back-to-Back (B2B) Consistency Point Scenario?, What are the different Consistency Point types and how are they measured in ONTAP 9?, What are the different Consistency Point types and how are they measured?

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.

```
vserver       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::*> vserver services dns dynamic-update show
Vserver         Is-Enabled Use-Secure Vserver 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 -vserver-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: B, D

Explanation:
To enable DDNS updates to work, two actions should be performed:
Remove the NFS protocol from the cifs_01 data LIF. This is because DDNS updates are only supported for LIFs that have only one data protocol enabled1. The cifs_01 LIF has both NFS and CIFS protocols enabled, which prevents it from registering its DNS record dynamically. By removing the NFS protocol from the cifs_01 LIF, it will become eligible for DDNS updates.
Enable the -is-enabled parameter for the SVM DDNS services. This is because the -is-enabled parameter controls whether the SVM sends DDNS updates to the DNS servers2. The exhibit shows that the –isenabled parameter is set to false for the svm1 SVM, which means that it does not send any DDNS updates. By enabling the -is-enabled parameter, the SVM will start sending DDNS updates for its eligible LIFs. Reference:
1: Configure dynamic DNS services3
2: Manage DNS/DDNS services with System Manager4

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
                                                  Drawer                             Usable
Physical
    Disk            Outage Reason  HA Shelf Bay  /Slot
Chan    Pool  Type     RPM     Size     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       4879     Node04    secondary
Node03    vldb     0        11       358      -         offline
Node03    vifmgr   11       11       4892     Node04    secondary
Node03    bcomd    11       11       62       Node04    secondary
Node03    crs      11       11       6        Node04    secondary
Node04    mgmt     11       11       4879     Node04    master
Node04    vldb     0        11       358      -         offline
Node04    vifmgr   11       11       4892     Node04    master
Node04    bcomd    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, pool0)
    RAID group /aggr2/plex0/rg0 (degraded, block checksums)

      RAID Disk Device          HA  SHELF BAY CHAN Pool Type  RPM  Used (MB/blks)    Phys
(MB/blks)
      --------- ------          ------------ ---- ---- ---- ----- --------------    ----------
----
      dparity   FAILED                N/A                        2538546/ -
      parity    3c.0.11         3c    0   11  SA:B   0  BSAS  7200 2538546/5198943744
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: D

Explanation:
The broken disk in Node03 is causing the cluster ring to be offline, which prevents the CIFS SVM from being accessible. The cluster ring is a distributed database that stores cluster configuration information and enables communication between cluster nodes. If the cluster ring is offline, the cluster cannot function properly and the CIFS SVM cannot serve data to clients. The other options are not relevant to the CIFS SVM connectivity issue. Reference = https://www.netapp.com/support-and-training/netapplearning-services/certifications/support-engineer/

https://mysupport.netapp.com/site/docs-and-kb

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: B

Explanation:
SnapMirror is a data replication technology that allows efficient and flexible data protection and disaster recovery for NetApp ONTAP storage systems1
SnapMirror transfers data between source and destination volumes using a network connection. SnapMirror can use storage efficiency features such as compression and deduplication to reduce the amount of data transferred and stored1
SnapMirror transfers are scheduled and controlled by policies that define the frequency, retention, and priority of the transfers. SnapMirror policies can also specify the network bandwidth limit for the transfers2
SnapMirror transfers are considered background tasks that run in the absence of user workload. SnapMirror transfers can consume CPU and disk resources on both source and destination systems, depending on the amount and type of data being replicated3
SnapMirror transfers can throttle up or down depending on the availability of system resources and network bandwidth. SnapMirror transfers will throttle up when there is no user workload, and throttle down when there is user workload. This is to ensure that SnapMirror transfers do not impact the performance of user operations3
Therefore, if a customer is concerned with high CPU and disk utilization on their SnapMirror destination system, the best answer is to explain that background tasks such as SnapMirror throttle up in the absence of user workload. This is normal and expected behavior, and it does not indicate a problem with the system3
Reference:
1: ONTAP 9 Data Protection - SnapMirror - The Open Group 2: ONTAP 9 Data Protection - SnapMirror Policies - The Open Group 3: SnapMirror storage efficiency configurations and behavior - Resolution Guide - NetApp Knowledge Base

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 I IF failover policy for each subnet inside a single broadcast domain.

Answer: C

Explanation:
A LIF (Logical Interface) is a logical entity that represents a network connection point on a node1.
A VIP LIF (Virtual IP LIF) is a LIF that can fail over across subnets within an IPspace2.
BGP (Border Gateway Protocol) is a routing protocol that enables VIP LIFs to advertise their IP addresses to external routers and to update the routing tables when a failover occurs3.
To connect a NetApp ONTAP cluster to a complex network that requires LIFs to fail over across subnets, you need to configure VIP LIFs using BGP on the cluster and on the external routers3.
This way, you can ensure that the network traffic is routed to the optimal node and port for each VIP LIF, and that the network connectivity is maintained in the event of a node or port failure3. Reference:
1: Logical Interfaces, ONTAP 9 Documentation Center
2: VIP LIFs, ONTAP 9 Documentation Center
3: Configuring BGP on a cluster, ONTAP 9 Documentation Center

Thank You for Trying Our Product

Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

Check our Customer Testimonials and ratings available on every product page.

Visit our website.

https://examsempire.com/