Download NetApp NS0-593 Exam Dumps Free

Shared by Shaffer on 17-06-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

When an administrator tries to create a share for an existing volume named voll, the process fails with an error.

```
cluster1::> vserver cifs share create -vserver svm1 -share-name vol1 -
path /vol1                                                                    Error:
command failed: The specified path "/vol1" does not exist in the namespace belonging to Vserver
"svm1".

cluster1::> vserver cifs share show
Vserver        Share         Path              Properties Comment  ACL
-------------- ------------- ----------------- ---------- -------- -----------
svm1           admin$        /                 browsable  -        -
svm1           c$            /                 oplocks    -        BUILTIN\Administrators /
Full Control
                                               browsable
                                               changenotify
                                               show-previous-versions
svm1           ipc$          /                 browsable  -        -

3 entries were displayed.

cluster1::> vserver cifs show
             Server          Status    Domain/Workgroup Authentication
Vserver      Name            Admin     Name             Style
------------ --------------- --------- ---------------- --------------
svm1         SVM1            up        DEMO             domain

cluster1::*> volume show -vserver svm1 -volume vol1

                                Vserver Name: svm1
                                 Volume Name: vol1
                              Aggregate Name: cluster1_01_SSD_1
    List of Aggregates for FlexGroup Constituents: cluster1_01_SSD_1
                            Encryption Type: none
                List of Nodes Hosting the Volume: cluster1-01
                                 Volume Size: 20MB
                                Name Ordinal: base
                           Volume Data Set ID: 1028
                    Volume Master Data Set ID: 2162375168
                                Volume State: online
                                Volume Style: flex
                       Extended Volume Style: flexvol
                     FlexCache Endpoint Type: none
                     Is Cluster-Mode Volume: true
                      Is Constituent Volume: false
                               Export Policy: default
                                     User ID: -
                                    Group ID: -
                              Security Style: -
                             UNIX Permissions: ------------
                                Junction Path: -
                         Junction Path Source: -
                              Junction Active: -
                       Junction Parent Volume: -
                              Junction Active: true
                       Junction Parent Volume: svm1_root
                          Vserver Root Volume: false
                                      Comment:
                               Available Size: 18.76MB
                              Filesystem Size: 20MB
                    Total User-Visible Size: 19MB
                                  Used Size: 244KB
                              Used Percentage: 1%
                                          ...
                           Volume Tiering Policy: none
                Volume Tiering Minimum Cooling Days: -
                Performance Tier Inactive User Data: -
          Performance Tier Inactive User Data Percent: -
```

Referring to the exhibit, what Is the reason for the error?

Options:

A- The volume must have a type of DP.
B- The volume has not been mounted.
C- The CIFS service is not authenticating properly with the domain controller.
D- The CIFS service is not in workgroup mode.

Answer:

B

Explanation:

The error message indicates that the specified path ''/vol1'' does not exist in the namespace belonging to Vserver ''svm1''. This means that the volume ''vol1'' has not been mounted to the Vserver's namespace, which is required for creating a share. The volume type, the CIFS service status, and the CIFS service mode are not relevant to the error.Reference=https://www.netapp.com/support-and-training/netapp-learning-services/certifications/support-engineer/

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

Question 2

Question Type: MultipleChoice

A customer wants to connect a NetApp AFF A700 system to a 40GbE switch. The controllers have a 10/40G Ethernet card in slot 4 for this purpose. The link comes up fine on node 2, but it will not come up on node 1. You look at the AutoSupport data for the nodes in question and see the output shown in the exhibit.

```
Node1:
        slot 4: 10/40 Gigabit Ethernet Controller XL710 QSFP+
                e4a MAC Address:    d0:39:ea:43:18:4f (auto-unknown-fd-down)
                e4b MAC Address:    d0:39:ea:43:18:50 (auto-unknown-fd-down)
                e4c MAC Address:    d0:39:ea:43:18:51 (auto-unknown-fd-down)
                e4d MAC Address:    d0:39:ea:43:18:52 (auto-unknown-fd-down)
                Device Type:        XL710 B1
                Firmware Version:   fw 5.20 nvm 5.04 etid 800045ac
                Part Number:        111-02590
                Hardware Revision:  B0
                Serial Number:      032009000403
        slot 4: 10/40 Gigabit Ethernet Controller XL710 QSFP+
                e4e MAC Address:    d0:39:ea:43:18:53 (auto-unknown-fd-down)
                e4f MAC Address:    d0:39:ea:43:18:54 (auto-unknown-fd-down)
                e4g MAC Address:    d0:39:ea:43:18:55 (auto-unknown-fd-down)
                e4h MAC Address:    d0:39:ea:43:18:56 (auto-unknown-fd-down)
                Device Type:        XL710 B1
                Firmware Version:   fw 5.20 nvm 5.04 etid 800045ac
                Part Number:        111-02590
                Hardware Revision:  B0
                Serial Number:      032009000403

Node2:
        slot 4: 10/40 Gigabit Ethernet Controller XL710 QSFP+
                e4e MAC Address:    d0:39:ea:3f:b0:4e (auto-40g_sr4-fd-up)
                e4a MAC Address:    d0:39:ea:3f:b0:4a (auto-40g_sr4-fd-up)
```

What is the cause of the customer's problem?

Options:

A- The cable is plugged in upside down.
B- The incorrect cable is used.
C- The ports are configured for 4x10Gbit instead of 40Gbit.
D- The port has an incorrect SFP inserted.

Answer:

D

Explanation:

= The cause of the customer's problem is that the port has an incorrect SFP inserted.According to the AutoSupport data, the node 1 has a device type of XL710 BQ, which is a 40GbE adapter that requires a QSFP+ transceiver1.However, the part number of the adapter is 111-02590, which is a 10GbE adapter that requires a SFP+ transceiver2. This mismatch prevents the link from coming up on node 1.The customer should replace the adapter with the correct one or use a compatible transceiver for the existing adapter.Reference=1:NetApp Hardware Universe - X1144A2:NetApp Hardware Universe - X1147A

Question 3

Question Type: MultipleChoice

Which two throughput elements are controlled by host applications? (Choose two.)

Options:

A- operation size
B- Jitter
C- concurrency
D- latency

Answer:

A, C

Explanation:

The two throughput elements that are controlled by host applications are operation size and concurrency. Operation size refers to the amount of data that is transferred in a single I/O request, and concurrency refers to the number of I/O requests that are issued simultaneously. Both of these factors affect the throughput of a host application, as they determine how much data can be sent or received within a given time period.Larger operation sizes and higher concurrency can increase the throughput, but they may also increase the latency and resource consumption12.Reference:1: NetApp Performance Analysis on Data ONTAP - Chapter 3: Host-Side Performance2: NetApp Performance Analysis on Data ONTAP - Chapter 4: Storage System Performance

Question 4

Question Type: MultipleChoice

After a normal power down of both nodes for building maintenance, Node01 of a 2-node cluster cannot be powered back up; however, all disk shelves are powered.

Which action should be performed to bring the cluster online and allow Node02 to serve data?

Options:

A- Recreate the cluster with the system configuration recovery cluster recreate -from node command.
B- Reboot the node With the system node reboot -node Node02 -bypass-optimization true

command.
C- Perform a takeover with the storage failover takeover -ofnode Node01 -option force command.
D- Reinitialize the cluster with option 4a from the boot menu.

Answer:

C

Explanation:

= The correct action to bring the cluster online and allow Node02 to serve data is to perform a takeover with the storage failover takeover -ofnode Node01 -option force command. This command will force Node02 to take over the resources of Node01 and serve the data from both nodes. This is necessary because Node01 is not responding and cannot initiate a graceful takeover.The other options are not correct because they will either destroy the existing cluster configuration (A and D) or reboot the node without taking over the resources of the other node (B).Reference=1Halt or reboot a node without initiating takeover in a two-node cluster - NetApp Documentation2Solved: Graceful shut down - NetApp Community

Question 5

Question Type: MultipleChoice

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.

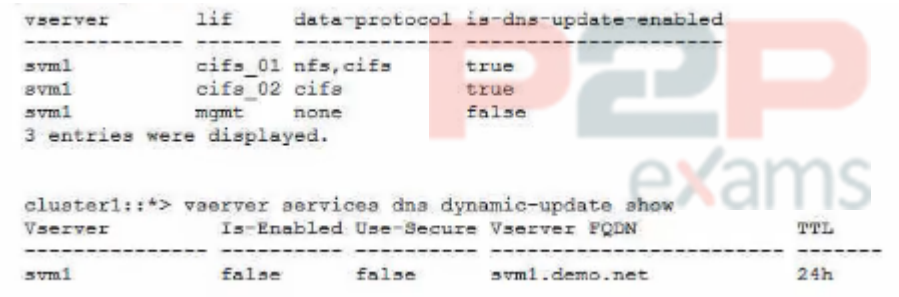

```
vserver       lif      data-protocol is-dns-update-enabled
------------- -------- ------------- ---------------------
svm1          cifs_01  nfs,cifs      true
svm1          cifs_02  cifs          true
svm1          mgmt     none          false
3 entries were displayed.

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.)

Options:

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 -is-enabled 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 6

Question Type: MultipleChoice

A customer with an FC MetroCluster configuration that is running NetApp ONTAP 9.8 software experienced a site outage and wants to know why the MetroClusler configuration did not switch over to the other site.

What are two reasons for this problem? (Choose two.)

Options:

A- There was a failed PSU in the switchover site.
B- The ISLs were down between the sites.
C- The site had unmirrored aggregates.
D- An ISL link was not operating at 16Gbps.

Answer:

B, C

Explanation:

A MetroCluster configuration requires the ISLs between the sites to be up and running for the switchover to occur. If the ISLs are down, the MetroCluster configuration cannot communicate with the other site and cannot perform the switchover. Additionally, a MetroCluster configuration requires all aggregates to be mirrored across the sites for data protection and availability. If a site has unmirrored aggregates, the switchover will fail because the data on those aggregates will not be accessible from the other site.Reference=Configuring the MetroCluster software in ONTAP,NetApp MetroCluster FC for ONTAP 9.8

Question 7

Question Type: MultipleChoice

Refer to the exhibit.

```
MCC-A::> metrocluster show

Configuration: IP-fabric

Cluster                   Entry Name          State
------------------------- ------------------- --------------------
Local: MCC-A
                          Configuration State  configured
                                        Mode  switchover
                          AUSO Failure Domain  auso-disabled
Remote: MCC-B
                          Configuration State  configured
                                        Mode  waiting-for-switchback
                          AUSO Failure Domain  auso-disabled
```

Referring to the exhibit, what do you need to do to return the MetroCluster to a normal state?

Options:

A- Enter the metroclustor switchback command on Site B.
B- Enter the metrocluster switchback command on Site A.
C- Enter the storage failover giveback command on Site B.
D- Enter the storage failover giveback command on Site A.

Answer:

B

Explanation:

The question refers to a MetroCluster configuration, which is a disaster recovery solution that uses two physically separated, mirrored clusters1.

The exhibit shows a MetroCluster switchover scenario, where Site A has experienced a disaster and Site B has taken over the tasks of Site A2.

To return the MetroCluster to a normal state, you need to perform a MetroCluster switchback operation, which reverses the switchover and activates the original sync-source storage virtual machines (SVMs) on Site A3.

To perform a MetroCluster switchback, you need to enter themetrocluster switchbackcommand on the cluster that was the source of the switchover, which is Site A in this case3.

The other options are not correct, because:

A)Entering themetrocluster switchbackcommand on Site B will not work, as Site B is the destination of the switchover, not the source3.

C)Entering thestorage failover givebackcommand on Site B will not work, as this command is used for local HA failover within a cluster, not for MetroCluster switchover between clusters4.

D)Entering thestorage failover givebackcommand on Site A will not work, as this command is used for local HA failover within a cluster, not for MetroCluster switchover between clusters4.Reference:

Understanding MetroCluster data protection and disaster recovery - NetApp

Perform IP MetroCluster switchover and switchback - NetApp

Performing a switchback - NetApp

High-availability configuration - NetApp

Question 8

Question Type: MultipleChoice

You are optimizing your Cloud Volumes ONTAP. For interoperability and consolidation purposes, you need to know the storage limits.

In this scenario, which source should be reviewed?

Options:

A- NetApp Cloud Volumes ONTAP Release Notes
B- Hardware Universe
C- Interoperability Matrix Tool (IMT)
D- Cloud Central

Answer:

A

Explanation:

= The NetApp Cloud Volumes ONTAP Release Notes provide the most up-to-date and accurate information about the storage limits for different cloud platforms, licenses, and configurations. The release notes also include details about supported features, known issues, and limitations for each version of Cloud Volumes ONTAP.Reference=Storage limits - NetApp,Storage limits | NetApp Documentation,Storage limits in AWS | NetApp Documentation,Storage limits in Azure | NetApp Documentation

Question 9

Question Type: MultipleChoice

A customer reports that while Installing Windows updates on their external Vscan servers, clients could not access any files on their CIF-S SVM. The problem disappeared after the update process was completed. The customer wants to prevent this issue from happening during the next patch window.

In this scenario, what are two ways to accomplish this task? (Choose two.)

Options:

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:

Updating the Vscan servers one at a time can prevent the clients from losing access to the files on the CIFS SVM, as there will always be at least one Vscan server available for scanning. Disabling the -scan-mandatory flag in the Vscan policy can also avoid the issue, as it allows the clients to access the files even if there are no Vscan servers online, but this may compromise the security of the data. Enabling the -scan-mandatory flag or modifying the CIFS shares to be CA shares will not help, as they do not affect the availability of the Vscan servers or the scanning process.Reference=Vscan server installation and configuration,Update the Vscan On-Access policy configuration for an SVM

Question 10

Question Type: MultipleChoice

You ate connecting a Windows Server 2019 host to ISCSI LUNs on a 4-node AFF A250 cluster running NetApp ONIAP 9.8 software. The host is unable to connect to the target interface, and the ONTAP event log has the message shown below.

ISCSI: network interface identifier disabled for use; incoming connection discarded

In this scenario, what is the cause of this connection issue?

Options:

A- iSNS is not enabled on the storage virtual machine.
B- The WWPN was not added to the initiator group.
C- The iSCSI service is not enabled on the interface.
D- CHAP is not configured correctly for the specified initiator.

Answer:

C

Explanation:

The iSCSI service must be enabled on the network interfaces that are used for iSCSI communication. If the iSCSI service is disabled on an interface, the incoming connection requests from the iSCSI initiator will be discarded and the event log will show the message ''ISCSI: network interface identifier disabled for use; incoming connection discarded''. To enable the iSCSI service on an interface, you can use the System Manager or the CLI commandnetwork interface modify -

vserver vserver_name -lif lif_name -data-protocol iscsi.Reference=Configure the iSCSI protocol with System Manager, [ONTAP 9 Network Management Guide]

Question 11

Question Type: MultipleChoice

Refer to the exhibit.

```
cluster1::> qos statistics volume latency show -vserver svm1 -volume vol1
Workload           ID    Latency    Network    Cluster       Data       Disk     QoS Max     QoS
Min      NVRAM
---------------- ------ ---------- ---------- ---------- ---------- ---------- ---------- ------
---- ----------
-total-             -
     4.71ms   185.00us     7.00us   468.00us     4.04ms        0ms        0ms    11.00us
vol1             32140    22.22ms    86.00us   125.00us   904.00us    21.11ms        0ms        0ms      1.00us
-total-             -
     4.58ms   253.00us     7.00us   991.00us     3.32ms        0ms        0ms    11.00us
vol1             32140    29.69ms    90.00us   162.00us     2.61ms    26.82ms        0ms        0ms      1.00us
-total-             -
     2.55ms   229.00us     8.00us   589.00us     1.71ms        0ms        0ms    13.00us
vol1             32140    18.84ms    61.00us   314.00us   136.00us    18.32ms        0ms        0ms      1.00us
-total-             -
     4.33ms   198.00us     9.00us  1221.00us     2.88ms        0ms        0ms    13.00us
vol1             32140   177.16ms   113.00us   226.00us   108.03ms    68.79ms        0ms        0ms      2.00us
```

Referring to the exhibit, from which delay center is most of the latency for vol1?

Options:

A- Disk
B- Network
C- Data
D- NVRAM

Answer:

C

Explanation:

. Data. However, I cannot guarantee the accuracy or completeness of this answer, and you should verify it with the NetApp Support Engineer ONTAP Specialist documents and learning resources.

Here are some references that might help you:

1:qos statistics volume latency show - NetApp2:How to use and understand qos statistics commands to monitor volume latency in real time - NetApp Knowledge Base