NetworkAppliance.NS0-593.v2024-09-13.q42

Exam Code:	NS0-593
Exam Name:	NetApp Certified Support Engineer ONTAP Specialist
Certification Provider:	Network Appliance
Free Question Number:	42
Version:	v2024-09-13
# of views:	334
# of Questions views:	420
https://www.dumpsdb.com/dumps/NetworkAppliance/NS0-593/NetworkAppliance.NS0-593.v2024-09-13.q42	

NEW QUESTION: 1

```
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?

A. Enter the storage failover giveback command on Site B.

B. Enter the metrocluster switchback command on Site A.

C. Enter the metroclustor switchback command on Site B.

D. Enter the storage failover giveback command on Site A.

Answer: (SHOW ANSWER)

NEW QUESTION: 2

You notice poor performance on your FlexGroup and execute the system node run -node * flexgroup show command for more Information. You notice the "Urge" column has non-zero values.

In this scenario, which statement is true?

A. The data placement is uneven.

B. The constituent volumes are out of Inodes.

C. The aggregate is completely full.

D. The constituent volumes are completely full.

Answer: A (LEAVE A REPLY)

= The "Urge" column in the flexgroup show command indicates the urgency of data rebalancing for each constituent volume. A non-zero value means that the data placement is uneven across the volumes, which can affect the performance of the FlexGroup. The target value is the desired percentage of data for each volume, and the difference value is the deviation from the target. The FlexGroup tries to balance the data placement by moving data between the volumes, but this process can be slow or interrupted by other factors. Therefore, it is recommended to monitor the "Urge" column and take corrective actions if the values are high or persistent. Reference = What do "flexgroup show" target and urgency and other columns mean?, NetApp ONTAP FlexGroup volumes - Best practices and implementation guide

NEW QUESTION: 3
You have a 4-node NetApp ONTAP 9.8 cluster with an AFF A400 HA pair and a FAS8300 HA pair with 16 TB NL-SAS drives. You are asked to automatically tier 150 TB of Snapshot copy data from the AFF A400 aggregates to the FAS8300.
In this scenario, which ONTAP license must be added to the cluster to accomplish this task?
A. VE license
B. FabricPool license
C. S3 license
D. TPM license
Answer: (SHOW ANSWER)

NEW QUESTION: 4
Which two automation methods does NetApp ONTAP Select support? (Choose two.)
A. REST
B. Anslble
C. Docker
D. PHP
Answer: A,B (LEAVE A REPLY)
= ONTAP Select supports REST and Ansible as automation methods. REST is the primary foundation for the automation technologies, and Ansible is a popular tool for automating ONTAP Select deployments using roles and playbooks. Reference = ONTAP Automation documentation, Automate ONTAP Select deployments with Ansible

NEW QUESTION: 5
You have a 4-node NetApp ONTAP 9.8 cluster with an AFF A400 HA pair and a FAS8300 HA pair with 16 TB NL-SAS drives. You are asked to automatically tier 150 TB of Snapshot copy data from the AFF A400 aggregates to the FAS8300.
In this scenario, which ONTAP license must be added to the cluster to accomplish this task?
A. S3 license
B. VE license
C. TPM license
D. FabricPool license
Answer: (SHOW ANSWER)
FabricPool is an ONTAP feature that enables tiering of cold data from SSD aggregates to low-cost object storage, either on-premises or in the cloud1. FabricPool requires a license to be installed on the cluster, and the license type depends on the

cloud tier being used2. In this scenario, the cloud tier is another ONTAP cluster (FAS8300), which is not supported by the new Cloud Tiering license that is used for most FabricPool configurations3. Therefore, the old FabricPool license that is retained for dark sites or MetroCluster systems using FabricPool Mirror must be used3. The FabricPool license defines the amount of capacity that can be tiered to the cloud tier, and it can be increased by add-on orders4. Reference:
1: FabricPool overview5
2: FabricPool requirements6
3: Install a FabricPool license2
4: ONTAP FabricPool (FP) Licensing Overview1

NEW QUESTION: 6
A storage administrator reports that a monitoring toot is reporting that the storage controller reads between 90% to 93% CPU use. You run the sysstat -m command against the node in question.

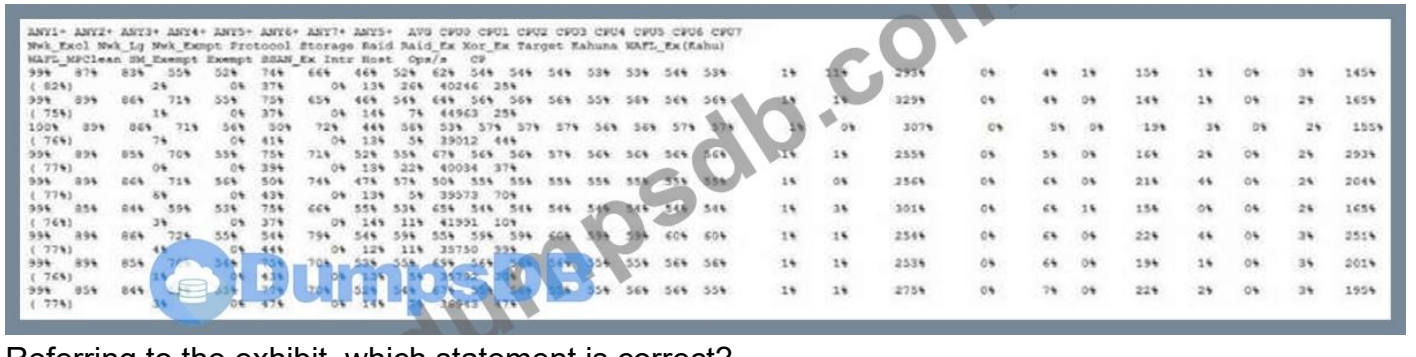

Referring to the exhibit, which statement is correct?
A. The customer should be advised to exclude certain workflows to reduce use.
B. High network exempt use could be a problem.
C. You should immediately investigate further by gathering perfstat data and opening a support case.
D. The CPU Is not a first-order monitoring metric for ONTAP.
Answer: D (LEAVE A REPLY)
= CPU utilization in ONTAP is not a linear measure of the system load, nor can it be used alone as a measure of the overall system utilization. ONTAP uses a Coarse Symmetric Multiprocessing (CSMP) design which partitions system functions into logical processing domains, each with its own scheduling rules and resource availability. Therefore, a high CPU utilization does not necessarily indicate a performance problem, unless it is accompanied by other contributing factors such as high latency, low throughput, or high queue depth. ONTAP has several mechanisms to optimize CPU usage and balance the workload across the cores, such as WAFL parallelization, exempt processing, and CPU pinning. The CPU utilization reported by the sysstat command is an average across all cores and domains, and does not reflect the actual CPU activity or availability for each domain. Therefore, the CPU is not a first-order monitoring metric for ONTAP, and other metrics such as latency, throughput, and queue depth should be considered first. Reference = What is CPU utilization in Data ONTAP: Scheduling and Monitoring?, How to measure CPU utilization, What are CPU as a compute resource and the CPU domains in ONTAP 9?, Monitoring CPU utilization before ONTAP upgrade

NEW QUESTION: 7
You have a 2-node NetApp FAS2750 switchless cluster with twenty-four 1.8 TB disks that is experiencing performance issues. Upon investigation, you discover severaltype B consistency points.

Referring to the exhibit, which corrective action would address these consistency points?

A. Add an additional shelf oftwenty-four 1.8 TB disks.

B. Create additional data LIFs.

C. Replace the twenty-four 1.8 TB disks with twelve 4 TB disks.

D. Convert the 2-node switchless cluster to a 2-node switched cluster.

Answer: A (LEAVE A REPLY)

NEW QUESTION: 8

After expanding a two-node AFF A300 cluster with two AFF A700 nodes, you observe latencies when data is accessed indirectly. The system node run -node -command netstat command shows retransmits and packet drops on the LIFs. The AFF A700 nodes with 40 Gbps and the AFF A300 connect with 10 Gbps to the cluster interconnect. You are using Cisco Nexus 3132Q cluster interconnect switches. You just updated to ONTAP 9.8 software.

In this scenario, what is causing this problem?

A. There is an NXOS firmware mismatch, verify in the Interoperability Matrix (1MTJ for a supportedversion.

B. The 10 Gbps to 10 Gbps cluster Interconnect is creating a speed mismatch.

C. AFF A300 and AFF A700 node mixing Isnot supported in ONTAP 9.8 software.

D. The AFF A300 Is using an adapter card for the cluster Interconnect, but only onboard ports are supported.

Answer: (SHOW ANSWER)

NEW QUESTION: 9

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.)

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 (LEAVE A REPLY)

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

NEW QUESTION: 10
You have a customer complaining of long build times from their NetApp ONTAP-based datastores. They provided you packet traces from the controller and client. Analysis of these traces shows an average service response time of 1 ms. QoS output confirms the same. The client traces are reporting an average of 15 ms in the same time period.
In this situation, what would be your next step?
A. The cluster is responding slowly and requires further investigation using performance archives.
B. The client that reports high latency should be investigated.
C. The cluster interconnects should be investigated.
D. A sync core should be triggered.
Answer: (SHOW ANSWER)
The question describes a scenario where the controller and client have a significant difference in their reported latency for the same datastores.
The controller's latency is 1 ms, which is within the normal range for ONTAP-based datastores1.
The client's latency is 15 ms, which is much higher than the controller's latency and could indicate a performance issue on the client side2.
Therefore, the next step is to investigate the client that reports high latency and identify the possible causes, such as network congestion, misconfiguration, resource contention, or application issues23.
The other options are not relevant or appropriate for this scenario, because:
A) The cluster is not responding slowly, as the controller's latency is low and QoS output confirms the same.
C) The cluster interconnects are not likely to be the cause of the latency difference, as they are used for communication between nodes within the cluster, not between the controller and the client4.
D) A sync core is a diagnostic tool that captures the state of the system at a given point in time, and is not a troubleshooting step for performance issues5. Reference:
ONTAP 9 Performance - Resolution Guide - NetApp Knowledge Base
Performance troubleshooting - NetApp
How to troubleshoot performance issues in Data ONTAP 8 7-mode
Cluster interconnect network - NetApp
How to generate a sync core on a node - NetApp

NEW QUESTION: 11
After a motherboard replacement on a NetApp AFF A300 in a SAN environment, the customer states that ports 0e and 0f are unable toconnect to the fabric. The ports report "offline".
What would you examine first to troubleshoot the issue?
A. system node hardware unified-connect show command output
B. storage port show command output
C. vserver fcp interface show command output
D. vserver fcp wwpn-alias show command output
Answer: B (LEAVE A REPLY)

NEW QUESTION: 12

You are trying to deploy a Connector in the AWScloud from NetApp Cloud Manager. The deployment fails and shows the message 'Insufficient permissions to deploy Cloud Connector". You have verified the AWS access key and the AWS secret key.
In this scenario, what Is the reason that the deployment failed?
A. The required Identity and Access Management (IAM) policies were not Installed.
B. The Connector can be deployed only inAWS GovCloud (US).
C. No AWS Marketplace subscription is associated with Cloud Manager.
D. The user lacks the permission to deploy within Cloud Manager.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 13
You created a new NetApp ONTAP FlexGroup volume spanning six nodes and 12 aggregates with a total size of 4 TB. You added millions of files to the FlexGroup volume with a flat directory structure totaling 2 TB, and you receive an out of apace error message on your host.
What would cause this error?
A. The maximum number of volume constituents has been reached in the ONTAP software.
B. All constituent volumes are full.
C. The inode limit is exceeded in the ONTAP software.
D. The maxdirsize Is exceeded in the ONTAP software.
Answer: D (LEAVE A REPLY)
The maxdirsize is the maximum size of a directory in a FlexVol or FlexGroup volume. It is determined by the number of inodes allocated to the directory. If the directory contains more files than the maxdirsize can accommodate, then the ONTAP software will return an out of space error message to the host, even if the volume has enough free space. This can happen when a FlexGroup volume has a flat directory structure with millions of files, as the maxdirsize is not automatically adjusted for FlexGroup volumes12. Reference: 1: FlexGroup volumes: Frequently asked questions | NetApp Documentation 2: How to increase the maxdirsize of a FlexVol volume - NetApp Knowledge Base

NEW QUESTION: 14
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 -vservser 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?

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: (SHOW ANSWER)
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

NEW QUESTION: 15
You are optimizing your Cloud Volumes ONTAP. For interoperability and consolidation purposes, you need to know the storagelimits.
In this scenario, which source should be reviewed?
A. Hardware Universe
B. Interoperability Matrix Tool (IMT)
C. NetApp Cloud Volumes ONTAP Release Notes
D. Cloud Central
Answer: C (LEAVE A REPLY)

NEW QUESTION: 16
You have a customer complaining of long build times from their NetApp ONTAP-based datastores. They provided you packet traces from the controller and client. Analysis of these traces shows an average service response time of 1 ms. QoS outputconfirms the same. The client traces are reporting an average of 15 ms in the same time period.
In this situation, what would be your next step?
A. The client that reports high latency should be investigated.
B. The cluster is responding slowly and requires further investigation using performance archives.
C. The cluster interconnects should be investigated.
D. A sync core should be triggered.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 17
You are troubleshooting a CIFS connection issue that is reported by some users. You decide to collect a packet trace.
In this scenario, after you generate the packet trace, where do you find the trace file?

A. /packet_traces of the node hosting the LIF
B. /etc/log/packet_traces of the node hosting the LIF
C. /etc/log/mlog/packet._traces Of all nodes
D. /vol0/paclcet_traces of the CIFS SVM
Answer: B (LEAVE A REPLY)
According to the NetApp Knowledge Base, packet traces are stored in the /etc/log/packet_traces directory of the node hosting the LIF. This applies to ONTAP 9.10 and above systems. For ONTAP 9.2 to 9.9 systems, the procedure to capture packet traces is similar, but the command syntax is slightly different. For ONTAP pre-9.2 systems, the procedure is different and requires the use of the tcpdump command. Reference = How to capture packet traces on ONTAP 9.10+ systems, How to capture packet traces (tcpdump) on ONTAP 9.2 to 9.9 systems, How to capture packet traces in ONTAP pre-9.2

NEW QUESTION: 18
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?
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 (LEAVE A REPLY)
= 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 - X1144A 2: NetApp Hardware Universe - X1147A

NEW QUESTION: 19

You receive the panic message shown in the exhibit.

```
Uncorrectable Machine Check Error at CPU0. MC0 Error: STATUS<0xb200000430000800>
(Val,UnCor,Enable,PCC,ErrCode(Src,NTO,Gen,Mem,L0)). MC5 Error: STATUS<0xf2000010c4300e0f>
(Val,OverF,UnCor,Enable,PCC,ErrCode(Gen,NTO,Gen,Gen,Gen)); Uncorrectable error at DIMM-1,
Channel 0, Serial: BA-00-1131-00098398|69002460-I01-NTA-T1?!, FERR(0x400), NERR(0x402), MERR
M10Err, Rank 3, Bank 6, CAS 0x1e8, RAS 0x1bcf Uncorrectable error at DIMM-1, Channel 0, Serial:
BA-00-1131-00098398|69002460-I01-NTA-T1?!, MERR M10Err, Rank 3, Bank 6, CAS 0x1e8, RAS 0x1bc.
```

In this scenario, which component should you troubleshoot first?

A. the CPU

B. the PCI card in slot 3

C. the MetroCluster FC-VI card in slot 6

D. the memory module in slot 1

Answer: A (LEAVE A REPLY)

NEW QUESTION: 20

You have a new VMware vSphere cluster with ESXI 7.0U2 hosts. The hosts are connected to a 4-node AFF A400 NetApp ONTAP 9.8 cluster with FC LUNs. You have a requirement to identify and follow I/Os from each VM on the shared FC LUN-backed datastores for troubleshooting purposes.

In this scenario, which VMware feature Is supported by ONTAP software to accomplish this task?

A. Network I/O Control (NIOC)

B. Virtual machine ID (VMID)

C. Storage I/O Control (SIOC)

D. vSphere Cluster Services (vCLS)

Answer: B (LEAVE A REPLY)

The VMware feature that is supported by ONTAP software to identify and follow I/Os from each VM on the shared FC LUN-backed datastores is Virtual machine ID (VMID). VMID is a unique identifier assigned to each VM by the vCenter Server. ONTAP software can use VMID to map the VMs to the FC LUNs and provide visibility into the I/O performance and latency of each VM. This feature is also known as Virtual Storage Console (VSC) for VMware vSphere12. Reference: 1: Virtual Storage Console for VMware vSphere | NetApp Documentation 2: How to use Virtual Storage Console (VSC) to monitor the performance of virtual machines - NetApp Knowledge Base

NEW QUESTION: 21

Your customer installed the shelf firmware for their NS224 shelf over a week ago, and the firmware has notupgraded on shelf 1 module B. The customer wants to know what the next steps would be to get the firmware upgraded after verifying that the shelf firmware is indeed loaded onto the system.

Which step would you perform to complete the firmware upgrade?

A. Reseat the NSM100 module.

B. Power cycle the shelf.
C. Reseat the PSU of the shelf.
D. Reseat the disk in Bay 0.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 22
Which two statements about NetApp Cloud Volumes ONTAP licenses ate true? (Choose two.)
A. Having an Essentials package enables you to convert to another licensing option.
B. BYOL licenses are purchased directly from NetApp.
C. Capacity-based license packages support only single-node configurations.
D. AWS Marketplace contracts cannot be mixed with BYOL.
Answer: B,C (LEAVE A REPLY)

NEW QUESTION: 23
A system panic due to an "L2 watchdog timeout hard reset" error occurred. You have found a FIFO message in the SP log.
Which FIFO message Is useful for Investigating this Issue?
A. before NMI BBSP:C=00 L=00 FIFO: F3 F2 F1 F0 5 B2 0 0 0 80 20 E F E3 C0 C2 3 4 6 7 9 9 A
B. BIOS exit BBSP:C=2F L=20 FIFO: 2 C8 C9 C4 24 4 6 18 7 8 11
C. before Reset BBSP:C=00 L=00 FIFO:
D. ONTAP Shutdown BBSP:C=00 L=2F FIFO:
Answer: A (LEAVE A REPLY)
= The FIFO message before NMI is useful for investigating the issue because it shows the state of the system before the non-maskable interrupt (NMI) was triggered by the L2 watchdog timeout. The FIFO message contains information about the CPU registers, the stack pointer, the instruction pointer, and the last executed instructions. This can help identify the cause of the system hang or deadlock that led to the watchdog reset. The other FIFO messages are not useful because they show the state of the system after the reset or shutdown, which may not reflect the original problem. Reference = https://kb.netapp.com/onprem/ontap/hardware/Handling_L2_Watchdog_Resets_on_the_FAS8200_and_AFF_A300_platforms https://docs.netapp.com/us-en/ontap-metrocluster/install-ip/task_sw_config_restore_defaults.html

NEW QUESTION: 24
Your customer installed the shelf firmware for their NS224 shelf over a week ago, and the firmware has not upgraded on shelf 1 module B. The customer wants to know what the next steps would be to get the firmware upgraded after verifying that the shelf firmware is indeed loaded onto the system.
Which step would you perform to complete the firmware upgrade?
A. Reseat the NSM100 module.
B. Reseat the disk in Bay 0.
C. Power cycle the shelf.
D. Reseat the PSU of the shelf.
Answer: A (LEAVE A REPLY)
The question refers to a scenario where the shelf firmware for an NS224 shelf has not been upgraded on one of the NVMe shelf modules (NSM) after a week of installation.

The NSM is responsible for managing the communication between the drives and the I/O modules (IOM) in the shelf1.
The shelf firmware for the NSM is automatically updated when the NSM is inserted into the shelf or when the system is rebooted2.
If the automatic update does not work, the manual update process involves reseating the NSM, which means removing it from the shelf and inserting it back3.
Reseating the NSM triggers the firmware update and also resets the NSM's state3.
The other options are not correct, because:
B) Reseating the disk in Bay 0 will not affect the NSM firmware update, as the disk is not connected to the NSM1.
C) Power cycling the shelf will disrupt the I/O operations and may cause data loss or corruption4.
D) Reseating the PSU of the shelf will not affect the NSM firmware update, as the PSU is not connected to the NSM1.
Reference:
NS224 NVMe drive shelf overview - NetApp
Shelf firmware update process - NetApp
Module firmware upgrade stuck on NS224 shelf - NetApp Knowledge Base
Power cycle a disk shelf - NetApp

NEW QUESTION: 25
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 (LEAVE A REPLY)
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

NEW QUESTION: 26
Refer to the exhibit.

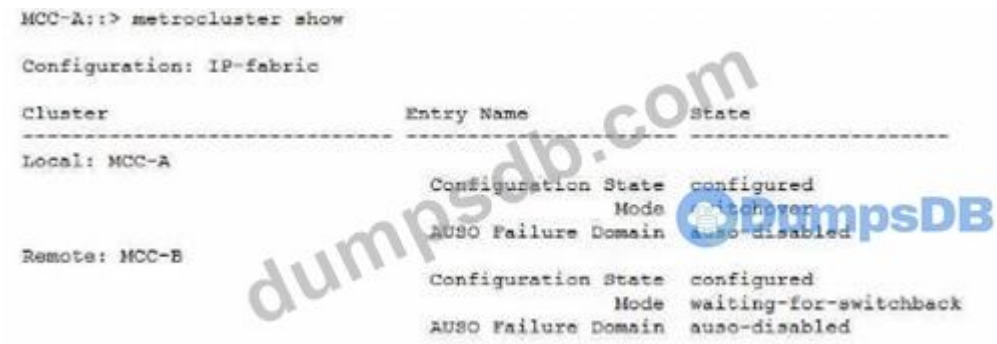

Referring to the exhibit, what do you need to do to return the MetroCluster to a normal state?

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: (SHOW ANSWER)

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 the metrocluster switchback command 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 the metrocluster switchback command on Site B will not work, as Site B is the destination of the switchover, not the source3.

C) Entering the storage failover giveback command 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 the storage failover giveback command 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

NEW QUESTION: 27

Your customer complains about missing volume snapshot copies on a SnapMIrror destination. While investigating this case, you notice an executed SnapMirror resync operation in the event logs of the system.

In this scenario, what Is the cause of this problem?

A. Newer Snapshot copies than the common snapshot that was chosen for resync were removed automatically at the destination.

B. Snapshot copies were manually removed from the destination after the resync.

C. NetApp ON TAP System Manager removed these snapshot copies due to a snapshot schedule.

D. The autodelete option removed Snapshot copies before the resync.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 28
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 (LEAVE A REPLY)
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

NEW QUESTION: 29
A user mentions that their home drive, that Is an export within a volume, is no longer allowing them to save files. The drive reports that it Is full, even though It shows that minimal data is written to it.
Which statement would explain this behavior?
A. The mount is stale and uses a cached version of the volume.
B. The client system needs to remount the export to show the proper space.
C. Other users wrote to this user's home drive.
D. Other files within the volume are also owned by the user, exceeding the user quota.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 30
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?

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 (LEAVE A REPLY)

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

NEW QUESTION: 31

Your customer mentions that they have accidentally destroyed both rootaggregates in their two-node cluster.

In this scenario, what are two actions that must be performed? (Choose two.)

A. Rejoin the second node to the re-created cluster.

B. Re-create the cluster from the local backup.

C. Install ONTAP from a USB device.

D. Re-create the cluster from the remote backup.

Answer: A,C (LEAVE A REPLY)

NEW QUESTION: 32

You recently discovered the error message shown below in your ONTAP logs.

```
[wafl.incons.userdata.vol:error]: WAFL inconsistent: volume my_data_vol
```

What should be your first action to correct this Issue?

A. Power cycle all the disk storage shelves that contain drives of the aggregate with the my_data_vol volume.

B. Determine the root cause behind the inconsistency before attempting any recovery procedure.

C. Use the storage takeover command on the storage controller that contains my_data_vol.

D. Use the wafliron command against my_data_vol to solve the inconsistency on the volume.

Answer: B (LEAVE A REPLY)

= The error message indicates that the volume my_data_vol is WAFL inconsistent, which means that there is a discrepancy between the data blocks and the metadata in the file system. This can be caused by various factors, such as hardware failures, software bugs, power outages, or network disruptions. The first action to correct this issue is to determine the root cause behind the inconsistency before attempting any recovery procedure, as recommended by the NetApp documentation1. This is because some recovery procedures, such as wafliron or storage takeover, may not work or may cause further damage if the underlying cause is not resolved. For example, if the inconsistency is due to a faulty disk or shelf, running wafliron may not fix the problem and may even corrupt more data. Therefore, it is important to identify and isolate the cause of the inconsistency before taking any further steps. Reference = 1 Volume Showing WAFL Inconsistent - NetApp Knowledge Base

NEW QUESTION: 33
You have a new VMware vSphere cluster with ESXI 7.0U2 hosts. The hosts are connected to a 4-node AFF A400 NetApp ONTAP9.8 cluster with FC LUNs. You have a requirement to identify and follow I/Os from each VM on the shared FC LUN-backed datastores for troubleshooting purposes.
In this scenario, which VMware feature Is supported by ONTAP software to accomplish this task?
A. vSphere Cluster Services (vCLS)
B. Virtual machine ID (VMID)
C. Network I/O Control (NIOC)
D. Storage I/O Control (SIOC)
Answer: (SHOW ANSWER)

NEW QUESTION: 34
You have recently discovered that NetApp ONTAP Cloud Manager is not sending AutoSupport messages to NetApp.
In this scenario, what would solve this Issue?
A. Verify that your AWS credentials are correctly added to Cloud Manager.
B. Verify that your NetApp Support site credentials are correctly added to Cloud Manager.
C. Verify that your data LIFs allow traffic to mysupport .netapp.com on port 443.
D. Verify that the inbound and outbound rules allow AutoSupport on port 80.
Answer: (SHOW ANSWER)
= AutoSupport is enabled by default for Cloud Manager and ONTAP Cloud systems. It proactively monitors the health of your systems and sends messages to NetApp technical support. To enable AutoSupport, you must provide your NetApp Support site credentials to Cloud Manager. If your credentials are incorrect or expired, Cloud Manager will not be able to send AutoSupport messages to NetApp. Therefore, the solution for this issue is to verify that your NetApp Support site credentials are correctly added to Cloud Manager. Reference = Troubleshooting Cloud Manager and ONTAP Cloud, Set up AutoSupport

NEW QUESTION: 35
You have a 4-node NetApp ONTAP 9.8 cluster with an AFF A400 HA pair and a FAS8300 HA pair with 16 TB NL-SAS drives. You are asked to automatically tier 150 TB of Snapshot copy data from the AFF A400 aggregates to the FAS8300.
In this scenario, which ONTAP license must be added to the cluster to accomplish this task?
A. TPM license

B. VE license
C. S3 license
D. FabricPool license
Answer: (SHOW ANSWER)

NEW QUESTION: 36
Your customer has mounted an NFS SVM from a Linux client unci performance is very poor. The customer is certain that they have jumbo frames enabled. They have verified an MTU of 9000 on both the Linux client and the broadcast domain on the NetApp ONTAP 9.8 cluster.

```
cluster1::*> network interface show -vserver svm1_cluster1
            Logical      Status     Network            Current       Current Is
Vserver     Interface    Admin/Oper Address/Mask       Node          Port    Home
----------- ---------- ---------- ------------------ ------------- ------- ----
svm1_cluster1
            cifs_01      up/up    192.168.0.131/24   cluster1-01   e0f     true
            cifs_02      up/up    192.168.0.132/24   cluster1-02   e0f     true
            jumbo_01     up/up    192.168.0.201/24   cluster1-01   e0g     true
            jumbo_02     up/up    192.168.0.202/24   cluster1-02   e0g     true
            mgmt         up/up    192.168.0.130/24   cluster1-01   e0c     true
5 entries were displayed.

cluster1::*> network port broadcast-domain show -broadcast-domain Jumbo
IPspace Broadcast                                    Update
Name    Domain Name   MTU  Port List                 Status Details
------- ----------- ------ -------------------------- --------------
Default Jumbo         9000
                           cluster1-01:e0f            complete
                           cluster1-01:e0g            complete
                           cluster1-02:e0f            complete
                           cluster1-02:e0g            complete
```

Referring to the exhibit, which ONTAP command will help isolate a possible MTU mismatch?
A.
```
network ping -lif jumbo_01 -vserver svm1_cluster1 -destination 192.168.0.210 -
disallow-fragmentation true -packet-size 5000
```
B.
```
network ping -lif jumbo_01 -vserver svm1_cluster1 -destination 192.168.0.210 -
disallow-fragmentation true -packet-size 9000
```
C.
```
network ping -lif jumbo_01 -vserver svm1_cluster1 -destination 192.168.0.210 -
disallow-fragmentation false -packet-size 1000
```
D.
```
network ping -lif jumbo_01 -vserver svm1_cluster1 -destination 192.168.0.210 -
disallow-fragmentation true -packet-size 1000
```
Answer: B (LEAVE A REPLY)

NEW QUESTION: 37
You have recently discovered that NetApp ONTAP Cloud Manager is not sending AutoSupport messages to NetApp.
In this scenario, what would solve this Issue?
A. Verify that your AWScredentials are correctly added to Cloud Manager.
B. Verify that your data LIFs allow traffic to mysupport .netapp.com on port 443.
C. Verify that the inbound and outbound rules allow AutoSupport on port 80.
D. Verify that your NetApp Support site credentials are correctly added to Cloud Manager.

Answer: (SHOW ANSWER)

NEW QUESTION: 38

A storage administrator reports that a monitoring toot is reporting that the storage controller reads between 90% to 93% CPU use. You run the sysstat -m command against thenode in question.

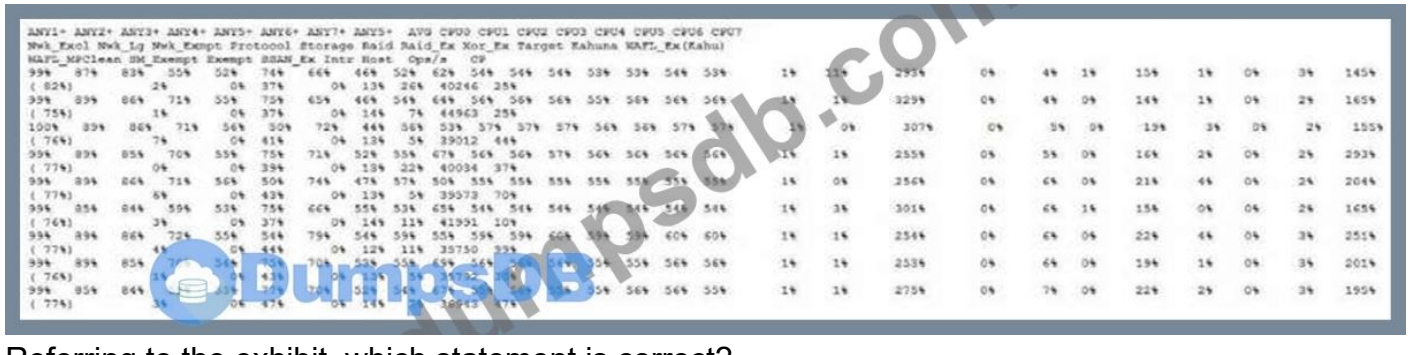

Referring to the exhibit, which statement is correct?

A. High network exempt use could be a problem.

B. The CPU Is not a first-order monitoring metric for ONTAP.

C. The customer should be advised to exclude certain workflows to reduce use.

D. You should immediately investigate further bygathering perfstat data and opening a support case.

Answer: D (LEAVE A REPLY)

NEW QUESTION: 39

Your customer Informs you about SnapMIrror problems after upgrading NetApp ONTAP software to a newer version. After investigating the event logs and the SnapMirror history, you see information about delayed updates of the SnapMirror relationships.

How would your customer prevent such problems in the future?

A. Quiesce the SnapMirror relationships before upgrading the ONTAP software.

B. Verify that the cabling of the hardware port that Is responsible for SnapMirror transfers Is correct.

C. Modify the schedules of the SnapMirror relationships after upgrading the ONTAP software.

D. Delete the SnapMirror relationships and create them new after upgrading the ONTAP software.

Answer: A (LEAVE A REPLY)

Quiescing the SnapMirror relationships before upgrading the ONTAP software ensures that the data replication is completed and consistent across the source and destination volumes. This prevents any data loss or corruption due to the upgrade process. Quiescing also suspends the scheduled updates until the relationships are resumed, avoiding any delays or failures in the transfers. The other options are not effective in preventing the SnapMirror problems after the upgrade. Reference = https://docs.netapp.com/us-en/ontap/upgrade/concept_upgrade_requirements_for_snapmirror.html https://docs.netapp.com/us-en/ontap-systems-upgrade/upgrade-arl-auto-app/resume_snapmirror_operations.html

NEW QUESTION: 40

Your customer noticed in NetApp Active IQ that their NetApp Cloud Volumes ONTAP for Azure HA solution is no longer sending AutoSupport messages over HTTPS. A support ticket has been opened to find out why. No changes have been made to the Cloud Volumes ONTAP for Azure HA environment.

In this scenario, which two autosupport command parameters should be used to validate that AutoSupport Is working properly? {Choose two.)

A. -mail-hosts

B. -transport

C. -proxy-url

D. -to

Answer: B,C (LEAVE A REPLY)

= The -transport parameter specifies the protocol used to send AutoSupport messages, which should be HTTPS by default. The -proxy-url parameter specifies the proxy server used to send AutoSupport messages, which should be the Connector's IP address and port if the Cloud Volumes ONTAP nodes do not have outbound internet access. These two parameters can be used to check the AutoSupport configuration and connectivity status. Reference = Verify AutoSupport setup, Troubleshoot your AutoSupport configuration, High-availability pairs in Azure

NEW QUESTION: 41

After expanding a two-node AFF A300 cluster with two AFF A700 nodes, you observe latencies when data is accessed indirectly. The system node run -node -command netstat command shows retransmits and packet drops on the LIFs. The AFF A700 nodes with 40 Gbps and the AFF A300 connect with 10 Gbps to the cluster interconnect. You are using Cisco Nexus 3132Q cluster interconnect switches. You just updated to ONTAP 9.8 software.

In this scenario, what is causing this problem?

A. The 10 Gbps to 10 Gbps cluster Interconnect is creating a speed mismatch.

B. AFF A300 and AFF A700 node mixing Is not supported in ONTAP 9.8 software.

C. The AFF A300 Is using an adapter card for the cluster Interconnect, but only onboard ports are supported.

D. There is an NXOS firmware mismatch, verify in the Interoperability Matrix (1MTJ for a supported version.

Answer: A (LEAVE A REPLY)

A speed mismatch occurs when the cluster ports on different nodes have different speeds, such as 10 Gbps and 40 Gbps1. A speed mismatch can cause packet loss, retransmits, and latency on the cluster network, especially when data is accessed indirectly, meaning that the node serving the data is different from the node hosting the LIF2.

In this scenario, the AFF A300 nodes have 10 Gbps cluster ports, while the AFF A700 nodes have 40 Gbps cluster ports, creating a speed mismatch on the cluster interconnect34.

The Cisco Nexus 3132Q cluster interconnect switches support both 10 Gbps and 40 Gbps cluster ports, but they have a shallow buffer size of 9 MB per port, which can be easily overwhelmed by the bursty traffic generated by the speed mismatch.

To solve this problem, you need to either upgrade the cluster ports on the AFF A300 nodes to 40 Gbps, or downgrade the cluster ports on the AFF A700 nodes to 10 Gbps, to achieve a homogeneous cluster interconnect1.

Alternatively, you can use QoS policies to limit the bandwidth of the cluster ports to avoid oversubscription and congestion on the cluster interconnect. Reference:

1: Cluster network cabling, ONTAP 9 Documentation Center

2: Cluster Network Latency - Troubleshooting Guide, NetApp Knowledge Base

3: AFF A300 Tech Specs, NetApp

4: AFF A700 Tech Specs, NetApp

[5]: Why do network speed mismatches create problems with shallow buffered switches?, NetApp Knowledge Base

[6]: How to use QoS to limit the bandwidth of cluster ports, NetApp Knowledge Base

NEW QUESTION: 42

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?

A. The CIFS service is not in workgroup mode.

B. The volume must have a type of DP.

C. The volume has not beenmounted.

D. The CIFS service is not authenticating properly with the domain controller.

Answer: C (LEAVE A REPLY)