NetworkAppliance.NS0-593.v2026-01-16.q40

Exam Code:	NS0-593
Exam Name:	NetApp Certified Support Engineer ONTAP Specialist
Certification Provider:	Network Appliance
Free Question Number:	40
Version:	v2026-01-16
# of views:	126
# of Questions views:	400
https://www.freeqas.com/qa/NetworkAppliance/NS0-593/NetworkAppliance.NS0-593.v2026-01-16.q40.html	

NEW QUESTION: 1
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.

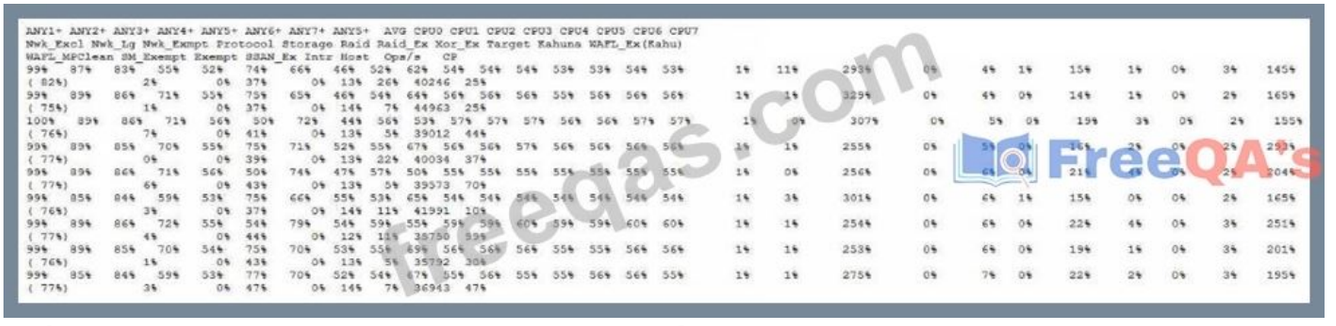

Referring to the exhibit, which statement is correct?
A. The customer should be advised to exclude certain workflows to reduce use.
B. You should immediately investigate further bygathering perfstat data and opening a support case.
C. High network exempt use could be a problem.
D. The CPU Is not a first-order monitoring metric for ONTAP.
Answer: B (LEAVE A REPLY)

NEW 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.

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

two.)

A. Enable the -use-secure parameter for the SVM DDNS services.

B. Enable the -is-enabledparameter for the SVM DDNS services

C. Remove the NFS protocol from the cifs_01 data LIF.

D. Disable the -vserver-fqdn parameter for the SVM DDNS services.

Answer: (SHOW ANSWER)

NEW QUESTION: 3

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 NetApp Support site credentials are correctly added to Cloud Manager.

B. Verify that the inbound and outbound rules allow AutoSupport on port 80.

C. Verify that your data LIFs allow traffic to mysupport .netapp.com on port 443.

D. Verify that your AWScredentials are correctly added to Cloud Manager.

Answer: C (LEAVE A REPLY)

NEW QUESTION: 4

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. Network I/O Control (NIOC)

B. vSphere Cluster Services (vCLS)

C. Storage I/O Control (SIOC)

D. Virtual machine ID (VMID)

Answer: B (LEAVE A REPLY)

NEW QUESTION: 5

Your customer wants to access a LUN on a FAS 8300 system from a VMware ESXi server through the FC protocol. They already created a new SVM, volume. LUN, and igroup for this purpose. The customer reports that the server's FC HBA port Is online, but the LUN does not show up.

Referring to the exhibit, what is the reason for this problem?

A. The FC service has not been configured on the SVM.

B. The zoning on the FC switches Is Incorrect.

C. The LUN Is not mapped to the correct SCSI ID.

D. The esxserver igroup contains incorrect IQNs.

Answer: A (LEAVE A REPLY)

To access a LUN on a FAS 8300 system from a VMware ESXi server through the FC protocol, the customer must configure the FC service on the SVM that owns the LUN. The FC service enables the SVM to act as an FC target and communicate with the FC initiators on the host. Without the FC service, the LUN will not be visible to the host, even if the LUN is mapped to an igroup and the FC LIFs are up. The exhibit shows that the FC service is not configured on the SVM, as the output of the command vserver fcp initiator show -vserver SVM1 is empty. Therefore, the reason for the problem is that the FC service has not been configured on the SVM. Reference = Configure an SVM for FC, Create an FC protocol service, Single IQN iSCSI session with ESXi on ONTAP when igroup has two IQNs

NEW QUESTION: 6

An administrator receives the following error message:

```
Mon Dec 23 00:20:36 EST [nodeA: wafl_exempt08: wafl.cp.toolong:error]: Aggregate
fas_01_DATA_AGGR experienced a long CP.
```

What are two causes for this error? (Choose two.)

A. A disk is falling.

B. There Is excessive SSD load causing the wear leveling to become unbalanced.

C. There is excessive SATA HDD load.

D. An SSD disk is performing garbage collection to create a dense data layout.

Answer: (SHOW ANSWER)

NEW QUESTION: 7

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 beenmounted.

C. The CIFS service is not in workgroup mode.

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

Answer: B (LEAVE A REPLY)

NEW QUESTION: 8

Your customer calls you because one application is not able to access a NetApp ONTAP S3 bucket. While reviewing the EMS log on the cluster, you see the following message:

```
Thu Jun 24 15:02:47 -0400 [node-01: kernel: nblade.css.s3.AccessDenied:error]: Access
is denied for user 'anonymous user' (Vserver 50), from client IP 10.10.10.10 accessing
resource '/offload-target'.
```

The same credentials work in another application.

In this scenario, what would cause this problem?

A. The application is trying to access using the HTTP protocol.

B. The application is using v2 signatures.

C. The application is using v4 signatures.

D. The application is trying to log in as an anonymous user.

Answer: D (LEAVE A REPLY)

NEW QUESTION: 9

Your customer added a new DS4246 shelf to their FAS2750 single-node system and created a new aggregate

on the new shelf. Two weeks later, they log into Active IQ and discover the Medium Impact error shown below.

Shelves with both connections via the same SAS path detected

```
Expanders on channel 0a:
Level    1: WWN 500a098005f0ed3f, ID  0, Serial Number ' SHJHU1527000148', Product
'DS424IOM6E      ', Rev '0172', Slot A
Level    2: WWN 500a098005b053bf, ID  1, Serial Number ' SHFHU1516001218', Product
'DS224IOM6       ', Rev '0191', Slot A
Level    3: WWN 500a09800772e3bf, ID  2, Serial Number ' SHJHU1653000119', Product
'DS424IOM6       ', Rev '0191', Slot A
Level    4: WWN 500a098005af77bf, ID  1, Serial Number ' SHFHU1516001218', Product
'DS224IOM6       ', Rev '0191', Slot B
Level    5: WWN 500a09800773b8bf, ID  2, Serial Number ' SHJHU1653000119', Product
'DS424IOM6       ', Rev '0191', Slot B

Expanders on channel 0b:
Level    1: WWN 500a09800773b0bf, ID  2, Serial Number ' SHJHU1653000119', Product
'DS424IOM6       ', Rev '0191', Slot B
Level    2: WWN 500a098005af77bf, ID  1, Serial Number ' SHFHU1516001218', Product
'DS224IOM6       ', Rev '0191', Slot B
Level    3: WWN 500a09802772e3bf, ID  2, Serial Number ' SHJHU1653000119', Product
'DS424IOM6       ', Rev '0191', Slot A
Level    4: WWN 500a098005b053bf, ID  1, Serial Number ' SHFHU1516001218', Product
'DS224IOM6       ', Rev '0191', Slot A
Level    5: WWN 500a098005f0ed3f, ID  0, Serial Number ' SHJHU1527000148', Product
```

Referring to the exhibit, which statement is correct?

A. This situation is normal for a FA57750 single-node system.

B. Active IQ Is not aware of the shelf configuration for a single-node system.

C. In a FAS2750 system, the internal shelf is only seen on one path, and the external shelves are seen on two paths.

D. The cables are cabled Incorrectly on the SAS stack.

Answer: A (LEAVE A REPLY)

NEW QUESTION: 10

Your customer complains that u host will constantly report losing a connection to the iSCSI target and then

report that the session was reestablished.

As shown in the exhibit, what is a cause of this flapping?

A. A host with an IP address of 172.20.10.80 and a secondhost with an IP address of 10.10.20.60 are in different initiator groups.

B. A host with an IP address of 172.20.10.80 and a secondIP address of 10.10.20.60 is accessing different LUNs.

C. A host with an IP address of 172.20.10.80 and a second host with an IP address of 10.10.20.60 have the same IQN.

D. A host with an IP address of 172.20.10.80 and a second host with an IP address of 10.10.20.60 are accessing the same LUN.

Answer: (SHOW ANSWER)

NEW QUESTION: 11

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 aggregate aggr2 has a failed disk.

B. The databases On Node03 must be Switched from secondary to master.

C. Thev1db database is offline.

D. The broken disk in Node03 is the source of the problem.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 12

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: 13
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 several type B consistency points.

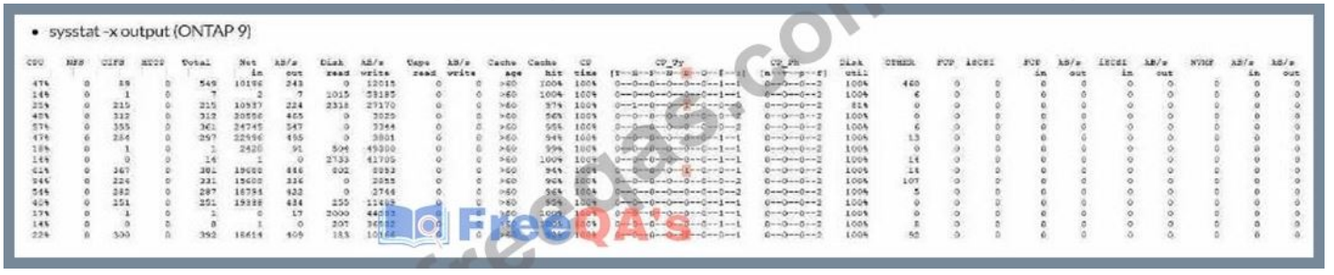

Referring to the exhibit, which corrective action would address these consistency points?
A. Create additional data LIFs.
B. Convert the 2-node switchless cluster to a 2-node switched cluster.
C. Add an additional shelf of twenty-four 1.8 TB disks.
D. Replace the twenty-four 1.8 TB disks with twelve 4 TB disks.
Answer: C (LEAVE A REPLY)
A type B consistency point (CP) is triggered when the NVRAM buffer is full and needs to be flushed to disk. A type B CP indicates that the write workload is higher than the disk throughput, and the system is experiencing back pressure1 A switchless cluster is a cluster configuration that does not use external switches for cluster interconnect and management network. A switchless cluster has lower bandwidth and redundancy than a switched cluster, and is limited to two nodes2 The exhibit shows the output of the sysstat -x command, which displays the system performance statistics in extended mode. The output shows that the

system has high CPU utilization, high disk utilization, high NVRAM utilization, and several type B CPs. These are signs of performance issues and resource contention3 The best corrective action to address these consistency points is to add an additional shelf of twenty-four 1.8 TB disks. This will increase the disk capacity and throughput, and reduce the disk utilization and the frequency of type B CPs4 Creating additional data LIFs will not address the consistency points, because the data LIFs are used for data access protocols, not for NVRAM flushing5 Converting the 2-node switchless cluster to a 2-node switched cluster will not address the consistency points, because the cluster interconnect and management network are not related to the disk performance6 Replacing the twenty-four 1.8 TB disks with twelve 4 TB disks will not address the consistency points, because the disk throughput will not increase, and the disk utilization will remain high7 Reference:
1: Where can I learn more about Consistency Points? - NetApp Knowledge Base 2: ONTAP 9 - Cluster and SVM Peering Express Guide - The Open Group 3: ONTAP 9 - Commands: Manual Page Reference - The Open Group 4: ONTAP 9 - Hardware Universe - The Open Group 5: ONTAP 9 - Network Management Guide - The Open Group 6: ONTAP 9 - Clustered Data ONTAP Concepts Guide - The Open Group 7: ONTAP 9 - Logical Storage Management Guide - The Open Group

NEW QUESTION: 14
Your customer added a new DS4246 shelf to their FAS2750 single-node system and created a new aggregate on the new shelf. Two weeks later, they log into Active IQ and discover the Medium Impact error shown below.
Shelves with both connections via the same SAS path detected

Referring to the exhibit, which statement is correct?
A. In a FAS2750 system, the internal shelf is only seen on one path, and the external shelves are seen on two paths.
B. The cables are cabled Incorrectly on the SAS stack.
C. This situation is normal for a FA57750 single-node system.
D. Active IQ Is not aware of the shelf configuration for a single-node system.
Answer: B (LEAVE A REPLY)

According to the SAS cabling rules and concepts for shelves with IOM12/IOM12B modules1, one of the goals is to provide a single, easily understood universal algorithm for all SAS products and configurations. One of the rules is that of the two paths to a single disk storage stack from a single node, one should connect to the top and one to the bottom of the storage stack. This ensures maximum possible resilience and minimizes the reliance on controller takeovers. The exhibit shows that both SAS connections from one controller go to the same shelf, which violates this rule and causes the error message. This situation is not normal for a FAS2750 single-node system and can reduce reliability and availability. The cables should be reconnected according to the SAS cabling rules and the controller-to-stack cabling worksheet template for quad-pathed connectivity2. Reference = SAS cabling rules and concepts - shelves with IOM12/IOM12B modules, SAS cabling issue: Both SAS connections from one controller go to the same shelf, Controller-to-stack cabling worksheet template for quad-pathed connectivity.

NEW QUESTION: 15

Refer to the exhibit.

```
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 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 (LEAVE A REPLY)

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: 16
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: (SHOW ANSWER)
= 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: 17
You want to assign storage access to an NVMe protocol host.
Which three objects must beconfigured to accomplish this task? (Choose three.)
A. an NVMe LIF
B. the NVMe subsystem

C. an NVMe LUN
D. the NVMe broadcast domain
E. the NVMe namespace
Answer: A,B,D (LEAVE A REPLY)

NEW QUESTION: 18
Your customer is running a NetApp AFF A800 system with NetAppONTAP 9.8 software and states that their NS224 shelf is not showing with the correct shelf ID. You analyze the data and the shelf shows an identification of "1.SHFHUXXXXXXXXX."
What is the cause of the reported Issue?
A. A new shelf was added to the controller.
B. The shelf was cabled Incorrectly.
C. The shelf has been power cycled.
D. The shelf had a disk replaced too quickly.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 19
Your customer Informs you about SnapMlrror 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. Delete the SnapMirror relationships and create them new after upgrading the ONTAP software.
B. Verify that the cabling of the hardware port that Is responsible for SnapMirror transfers Is correct.
C. Quiesce the SnapMirror relationships before upgrading the ONTAP software.
D. Modify the schedules of the SnapMirror relationships after upgrading the ONTAP software.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 20
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.

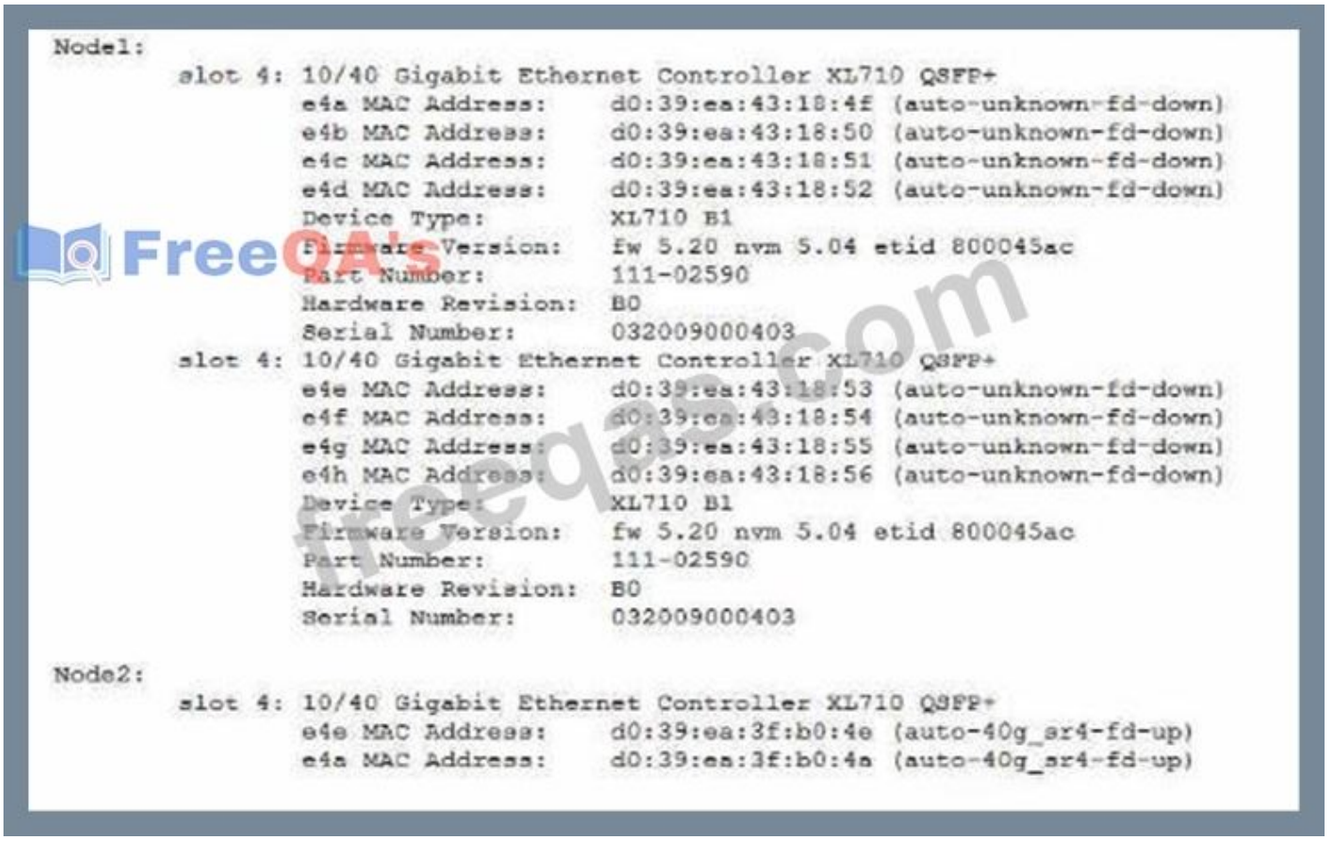

What is the cause of the customer's problem?

A. The cable is plugged in upside down.

B. The incorrect cable is used.

C. The port has an incorrect SFP inserted.

D. The ports are configured for 4x10Gbit instead of 40Gbit.

Answer: D (LEAVE A REPLY)

NEW QUESTION: 21

A customer's storage administrator Informs you about the deactivated Automatic Switchover (AUSO) feature on their MetroCluster IP environment.

What Information would you tell your customer in this scenario?

A. The AUSO feature cannot be activated due to a non-mirrored aggregate.

B. The AUSO feature is not available in Metrocluster IP installations by design.

C. A faulty link automatically deactivated the AUSO feature.

D. This is a configuration failure, you can activate that feature for the customer.

Answer: B (LEAVE A REPLY)

The AUSO feature is a MetroCluster functionality that enables an automatic switchover to the surviving site in the event of a disaster that affects one site1.

The AUSO feature is only available in MetroCluster FC configurations, which use Fibre Channel (FC) switches and FC-to-SAS bridges to connect the nodes and disk shelves across sites1.

MetroCluster IP configurations, which use Ethernet switches and network adapters to connect the nodes and disk shelves across sites, do not support the AUSO feature2.

Instead, MetroCluster IP configurations use the ONTAP Mediator service, which is a software component that monitors the health and connectivity of the MetroCluster nodes and initiates a Mediator-assisted unplanned switchover when a disaster occurs3.
The Mediator-assisted unplanned switchover is similar to the AUSO feature, but it requires the ONTAP Mediator service to be configured and running on a separate host3.
Therefore, you would tell your customer that the AUSO feature is not available in MetroCluster IP installations by design, and that they need to use the ONTAP Mediator service instead for disaster recovery.
Reference:
1: Understanding MetroCluster data protection and disaster recovery, ONTAP MetroCluster Documentation Center
2: Differences among the ONTAP MetroCluster configurations, ONTAP MetroCluster Documentation Center
3: Configure the ONTAP Mediator service from a MetroCluster IP configuration, ONTAP MetroCluster Documentation Center

NEW QUESTION: 22
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: D (LEAVE A REPLY)
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: 23
A 2-node cluster hosts only NAS data UFs. Allnodes lose access to the cluster network.

Referring to the exhibit, which statement is correct?

A. No takeover event will occur and NAS data UFs will not be accessible.

B. All nodes in the cluster reset and NAS data UFs are not accessible after reboot.

C. NAS data UFs are accessible on only the node that has Epsilon.

D. Both nodes race to takeover the other, one node will takeover the other node.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 24

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?

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 (LEAVE A REPLY)

= 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 = 1 Halt or reboot a node without initiating takeover in a two-node cluster - NetApp Documentation 2 Solved: Graceful shut down - NetApp Community

NEW QUESTION: 25

A 2-node cluster hosts only NAS data UFs. All nodes lose access to the cluster network.

Referring to the exhibit, which statement is correct?

A. Both nodes race to takeover the other, one node will takeover the other node.

B. No takeover event will occur and NAS data UFs will not be accessible.

C. All nodes in the cluster reset and NAS data UFs are not accessible after reboot.

D. NAS data UFs are accessible on only the node that has Epsilon.

Answer: D (LEAVE A REPLY)

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 Performance 2: NetApp Performance Analysis on Data ONTAP - Chapter 4: Storage System Performance

NEW QUESTION: 26

Your customer noticed in NetApp Active IQ that their NetApp Cloud Volumes ONTAP forAzure 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. -proxy-url

B. -mail-hosts

C. -transport

D. -to

Answer: A,B (LEAVE A REPLY)

NEW QUESTION: 27

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. Other users wrote to this user's home drive.
C. Other files within the volume are also owned by the user, exceeding the user quota.
D. The client system needs to remount the export to show the proper space.
Answer: C (LEAVE A REPLY)
= A user quota is a limit on the amount of disk space or the number of files that a user can consume on a volume. A user can be represented by multiple IDs, such as UNIX user ID (UID), Windows user ID (SID), or name mapping ID. A file owned by any of these IDs is subject to the restriction of the user quota. If a user has files in different directories or qtrees within the same volume, all of those files count towards the user's quota limit. Therefore, if other files within the volume are also owned by the user, they could exceed the user quota and prevent the user from saving more files to their home drive, even if the home drive itself shows minimal data usage. Reference = How quotas work with users and groups overview, ONTAP - How quotas work with users and groups

NEW QUESTION: 28
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. FabricPool license
C. TPM license
D. VE license
Answer: C (LEAVE A REPLY)

NEW QUESTION: 29
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. /etc/log/mlog/packet._traces Of all nodes
B. /packet_traces of the node hosting the LIF
C. /vol0/paclcet_traces of the CIFS SVM
D. /etc/log/packet_traces of the node hosting the LIF
Answer: A (LEAVE A REPLY)

NEW QUESTION: 30
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   504.00us    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?

A. Disk

B. Network

C. Data

D. NVRAM

Answer: C (LEAVE A REPLY)

. 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 - NetApp 2: How to use and understand qos statistics commands to monitor volume latency in real time - NetApp Knowledge Base

NEW QUESTION: 31

A user reports that a colleague saved a file called Test.txt from a UNIX system to a multiprotocol volume. When opening the file later from a Windows system, it was not the file that they wanted. The file that they wanted was named TEST~1.TXT.

Which statement explains this behavior?

A. UNIX name mapping updated the filename.

B. Case Insensitivity of SMB clients caused the file to be displayed with a different name.

C. Windows Volume Shadow Copy Service stored an older version of the file.

D. A Snapshot copy preserved two versions of the file.

Answer: (SHOW ANSWER)

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. Use the wafliron command against my_data_vol to solve the inconsistency on the 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. Power cycle all the disk storage shelves that contain drives of the aggregate with the my_data_vol volume.

Answer: (SHOW ANSWER)

NEW QUESTION: 33

```
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    504.00us    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?

A. Disk

B. Network

C. Data

D. NVRAM

Answer: B (LEAVE A REPLY)

NEW QUESTION: 34

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.

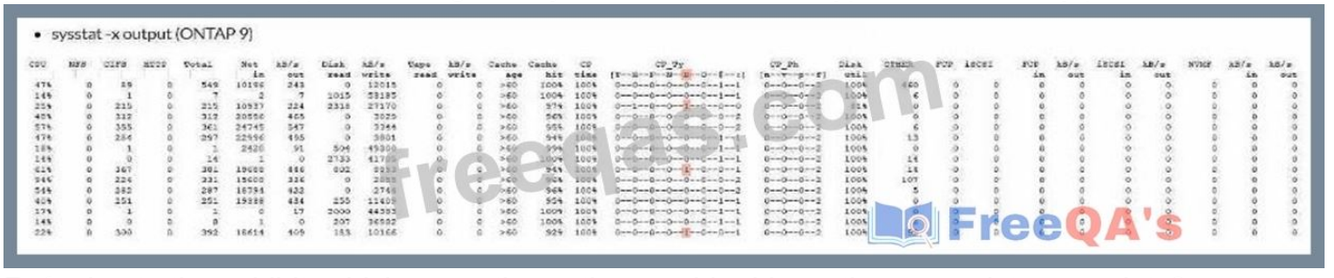

Referring to the exhibit, which corrective action would address these consistency points?

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

B. Create additional data LIFs.

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

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

Answer: (SHOW ANSWER)

NEW QUESTION: 35
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. ONTAP Shutdown BBSP:C=00 L=2F FIFO:
C. before Reset BBSP:C=00 L=00 FIFO:
D. BIOS exit BBSP:C=2F L=20 FIFO: 2 C8 C9 C4 24 4 6 18 7 8 11
Answer: D (LEAVE A REPLY)

NEW QUESTION: 36
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: (SHOW ANSWER)
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: 37

Your customer calls you because one application is not able to access a NetApp ONTAP S3 bucket. While reviewing the EMS log on the cluster, you see the following message:

```
Thu Jun 24 15:02:47 -0400 [node-01: kernel: nblade.css.s3.AccessDenied:error]: Access
is denied for user 'anonymous user' (Vserver 50), from client IP 10.10.10.10 accessing
resource '/offload-target'.
```

The same credentials work in another application.
In this scenario, what would cause this problem?
A. The application is trying to log in as an anonymous user.
B. The application is using v2 signatures.
C. The application is using v4 signatures.
D. The application is trying to access using the HTTP protocol.
Answer: B (LEAVE A REPLY)
The error message indicates that the access is denied for user 'anonymous user' (Vserver 50), from client IP 10.10.10.10 accessing resource '/offload-target'. This means that the application is not providing a valid user name or access key to authenticate with the ONTAP S3 server. According to the NetApp documentation1, ONTAP S3 server supports both v2 and v4 signatures, but v2 signatures are deprecated and not recommended. Therefore, the application should use v4 signatures to access the ONTAP S3 bucket. Additionally, the application should specify the bucket name in the host name, not in the resource path, as shown in the following example2:
https://bucket-name.s3-server-name.com/object-name
Reference = 1 ONTAP S3 Server Overview - NetApp Documentation 2 Accessing an ONTAP S3 bucket - NetApp Documentation

NEW QUESTION: 38
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 anout of apace error message on your host.
What would cause this error?
A. The maxdirsize Is exceeded in the ONTAP software.
B. The inode limit is exceeded in the ONTAP software.
C. The maximum number of volume constituents has been reached in the ONTAP software.
D. All constituent volumes are full.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 39
A customer calls you to troubleshoot a network issue. The customer wants to create a packet trace for all clients on the 192.168.9.0/24 subnet.

```
          Logical     Status     Network            Current       Current Is
Vserver   Interface   Admin/Oper Address/Mask       Node          Port    Home
--------- ----------- ---------- ------------------ ------------- ------- ----
CIFS
          nfs1        up/up      10.92.4.202/16     cl1n1         e2a-1092
                                                                          true
          iscsi1      up/up      10.91.3.132/24     cl1n1         a0a-322
                                                                          true
          iscsi2      up/up      10.95.7.114/24     cl1n1         a0a-456
                                                                          true
          mgmt        up/up      10.92.6.200/16     cl1n1         e2a-1092
                                                                          true
          nfs2        up/up      192.168.3.9/24     cl1n1         e2a-100
                                                                          true
cl1::> network route show -vserver CIFS
Vserver             Destination     Gateway         Metric
------------------- --------------- --------------- ------
CIFS
                    0.0.0.0/0       10.92.0.1       10
                    0.0.0.0/0       10.92.0.2       20
                    10.81.50.140/32 192.168.3.254   20
                    192.168.9.0/24  192.168.3.254   20
5 entries were displayed.
```

Referring to the exhibit, on which interface do you have to run the packet trace?

A. a0a-322

B. e2a-100

C. e2a-1092

D. a0a-456

Answer: B (LEAVE A REPLY)

To create a packet trace for all clients on the 192.168.9.0/24 subnet, you have to run the packet trace on the interface that has an IP address in the same subnet. According to the exhibit, the only interface that matches this criterion is e2a-100, which has an IP address of 192.168.3.1/24. The other interfaces have IP addresses in different subnets, such as 10.92.4.202/16, 10.91.3.132/24, 10.95.7.114/24, and 10.92.6.200/16.

Therefore, the correct answer is B. e2a-100. 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 download captured packet traces from ONTAP

NEW QUESTION: 40

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.

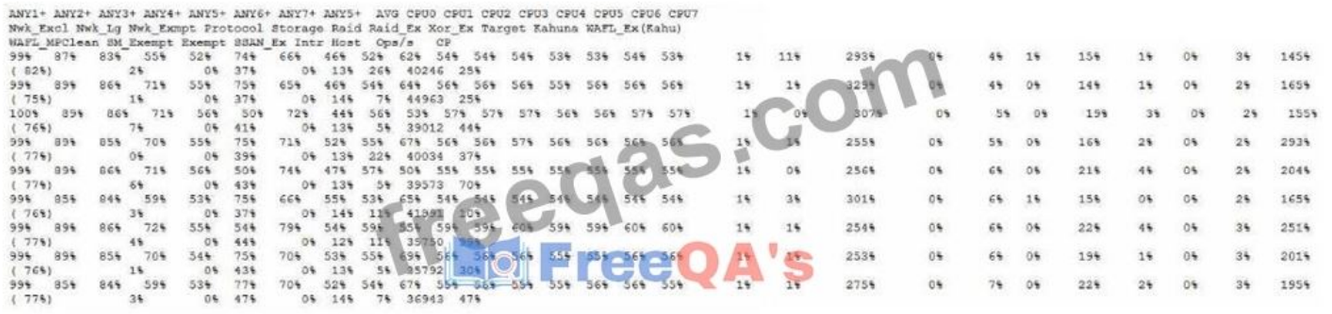

Referring to the exhibit, which statement is correct?

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