



BrainDumps
Collection

NetApp

NS0-592 Exam

NetApp Certified Support Engineer - ONTAP Specialist

Thank you for Downloading NS0-592 exam PDF Demo

You can also try our NS0-592 practice exam software

Download Free Demo

<https://www.braindumpscollection.com/NS0-592.html>

DEMO
VERSION

(LIMITED CONTENT)

Questions
& Answers

Version: 4.0

Question: 1

A customer mentions that their intercluster SnapMirror replication operations consistently take much longer than normal. While troubleshooting, you want to initiate a test SnapMirror session to collect performance data.

How would you accomplish this task?

- A. Run the network test-path command to simulate a SnapMirror connection and record the latency/throughput data.
- B. View the active network connections to the node to see if there a large number of active connections.
- C. Use the ifstat command to collect network performance data for the physical interface.
- D. Run the network statistics lif show command for the intercluster LIF to simulate a network connection.

Answer: A

Question: 2

The root aggregate on your single-node ONTAP cluster failed. You have the configuration backups saved on your FTP server?

In this scenario, which two steps are part of the recovery procedure? (Choose two.)

- A. Perform a Netboot using the configuration backup.
- B. Create a new aggregate and set its ha_policy option to sfo.
- C. Boot into maintenance mode.
- D. Create a new root aggregate and set its ha_policy option to cfo.

Answer: C,D

Question: 3

When does ONTAP 9.5 overwrite a block that contains existing data?

- A. after a user edits a file

- B. after the block has been marked free
- C. when file permissions are changed
- D. when a file is copied in a LUN

Answer: D

Question: 4

Your customer has initiated several volume move commands to redistribute the workload in their 8-node cluster. To verify that this process is not affecting their client access, they were monitoring the network switches that provide client access, but they do not see the network load of the volume move command.

In this scenario, what is the reason for the behavior?

- A. The volume move command uses the HA iWarp interconnect, not the client access switches.
- B. The volume move command uses copy offload protocol, so it does not show up when monitored.
- C. The volume move command uses the intercluster LIFs, not the data LIFs.
- D. The volume move command uses the cluster interconnect switches, not the client access switches.

Answer: B

Question: 5

A customer reports that a takeover has been disabled on an HA pair.

What would be a cause for this to happen?

- A. The HA interconnect is not functioning properly.
- B. There is a disk running too hot.
- C. The interface for the hardware assist has become unavailable.
- D. The cluster network is not functioning properly.

Answer: A

Question: 6

Click the exhibit button.

CPU	NFS	CIFS	HTTP	Total	Net in	kB/s out	HDD read	kB/s write	SSD read	kB/s write	Tape read	kB/s write	Cache age	Cache hit	CP time	CP ty	HDD util	SDD util	OTHER	FCP	iSCSI	FCP in	FCP out	Kb/s in	Iscsi in	Kb/s out
20%	89	31	0	127	52	93	8	12	0	0	0	0	31s	100%	0%	-										
4%	0%	0%	1	0	0	0	0	0	0	0	0	0	0	31s	100%	0%	-									
1%	0%	0%	5	0	0	7	77	4	12	0	0	0	0	31s	100%	0%	-									
1%	0%	0%	12	0	0	16	12	99	0	0	0	0	0	31s	100%	0%	-									
1%	0%	0%	0	0	0	10	14	8	4	12	0	0	0	31s	100%	0%	-									
2%	12	18	0	34	17	12	920	2344	0	0	0	0	0	31s	100%	13%	T	4%	0%	4	0	0	0	0	0	0
0%	23	39	0	62	16	15	4	4	0	0	0	0	0	31s	100%	0%	-									
1%	0%	0%	55	0	0	66	30	17	8	24	0	0	0	31s	100%	0%	-									
1%	0%	0%	33	0	113	52	33	0	0	0	0	0	0	31s	100%	0%	-									
0%	0%	0%	27	0	70	41	21	0	0	0	0	0	0	31s	100%	0%	-									
1%	0%	0%	9	0	41	20	16	16	24	0	0	0	0	31s	99%	0%	-									
1%	0%	0%	7	0	0	0	0	0	0	0	0	0	0	31s	100%	0%	-									
1%	0%	0%	0	0	75	70	25	4	4	0	0	0	0	31s	100%	0%	-									
0%	0%	0%	0	0	0	0	0	0	0	0	0	0	0													

```
smb2:IODSSNA002C1:session_setup_ops:30/s
smb2:IODSSNA002C1:session_setup_latency:5859166.83us
smb2:IODSSNA002C1:session_setup_latency_histogram.<20us:2
smb2:IODSSNA002C1:session_setup_latency_histogram.<40us:0
smb2:IODSSNA002C1:session_setup_latency_histogram.<60us:0
smb2:IODSSNA002C1:session_setup_latency_histogram.<80us:0
smb2:IODSSNA002C1:session_setup_latency_histogram.<600us:0
smb2:IODSSNA002C1:session_setup_latency_histogram.<800us:0
smb2:IODSSNA002C1:session_setup_latency_histogram.<1ms:0
smb2:IODSSNA002C1:session_setup_latency_histogram.<2ms:1
smb2:IODSSNA002C1:session_setup_latency_histogram.<4ms:0
smb2:IODSSNA002C1:session_setup_latency_histogram.<6ms:0
smb2:IODSSNA002C1:session_setup_latency_histogram.<8ms:0
smb2:IODSSNA002C1:session_setup_latency_histogram.<10ms:0
smb2:IODSSNA002C1:session_setup_latency_histogram.<12ms:0
smb2:IODSSNA002C1:session_setup_latency_histogram.<14ms:0
smb2:IODSSNA002C1:session_setup_latency_histogram.<16ms:0
smb2:IODSSNA002C1:session_setup_latency_histogram.<18ms:1
smb2:IODSSNA002C1:session_setup_latency_histogram.<20ms:2
smb2:IODSSNA002C1:session_setup_latency_histogram.<40ms:0
smb2:IODSSNA002C1:session_setup_latency_histogram.<60ms:0
smb2:IODSSNA002C1:session_setup_latency_histogram.<80ms:0
smb2:IODSSNA002C1:session_setup_latency_histogram.<100ms:0
smb2:IODSSNA002C1:session_setup_latency_histogram.<200ms:0
smb2:IODSSNA002C1:session_setup_latency_histogram.<400ms:0
smb2:IODSSNA002C1:session_setup_latency_histogram.<600ms:0
smb2:IODSSNA002C1:session_setup_latency_histogram.<800ms:87
smb2:IODSSNA002C1:session_setup_latency_histogram.<1s:264
smb2:IODSSNA002C1:session_setup_latency_histogram.<2s:118
smb2:IODSSNA002C1:session_setup_latency_histogram.<4s:482
smb2:IODSSNA002C1:session_setup_latency_histogram.<6s:256
smb2:IODSSNA002C1:session_setup_latency_histogram.<8s:230
smb2:IODSSNA002C1:session_setup_latency_histogram.<10s:300
smb2:IODSSNA002C1:session_setup_latency_histogram.<20s:0
```

Many end users are complaining that logging in their systems takes a long time. When the login to the desktop finishes, everything is fine. The home directories are on NetApp storage. You do not see any significant disk or CPU contention on the controller, so you gather output statistics from the controller and focus on SMB.

Referring to the exhibit, which two actions would help diagnose the problem? (Choose two.)

A. Engage your Windows team to make sure that the domain controllers are not overloaded.

- B. Engage your desktop team to collect network traces from the clients.
- C. Collect more statistics to determine whether the controller itself is causing latency.
- D. Engage your network team to make sure that there are no issues between the storage controller and the domain controller.

Answer: C,D

Question: 7

Click the exhibit button.

```
Fri Sep 21 06:38:52 CEST [cluster1: spmd: spm.ntpd.process.exit:error]: The ntpd with
ID 50230 exited as a result of normal exit (20). This server will attempt to restart.
Fri Sep 21 06:38:52 CEST [cluster1: spmd: spm.ntpd.process.exit:error]: The ntpd with
ID 50231 exited as a result of normal exit (20). This server will attempt to restart.
Fri Sep 21 06:38:52 CEST [cluster1: spmd: spm.ntpd.process.exit:error]: The ntpd with
ID 50233 exited as a result of normal exit (20). This server will attempt to restart.
Fri Sep 21 06:38:52 CEST [cluster1: spmd: spm.ntpd.process.exit:error]: The ntpd with
ID 50233 exited as a result of normal exit (20). This server will attempt to restart.
Fri Sep 21 06:38:52 CEST [cluster1: spmd: spm.ntpd.process.exit:error]: The ntpd with
ID 50234 exited as a result of normal exit (20). This server will attempt to restart.
Fri Sep 21 06:38:52 CEST [cluster1: spmd: spm.ntpd.process.exit:error]: The ntpd with
ID 50235 exited as a result of normal exit (20). This server will attempt to restart.
Fri Sep 21 06:38:52 CEST [cluster1: spmd: spm.ntpd.process.exit:error]: The ntpd with
ID 50236 exited as a result of normal exit (20). This server will attempt to restart.
Fri Sep 21 06:38:52 CEST [cluster1: spmd: spm.ntpd.process.exit:error]: The ntpd with
ID 50237 exited as a result of normal exit (20). This server will attempt to restart.
Fri Sep 21 06:38:52 CEST [cluster1: spmd: spm.ntpd.process.exit:error]: Call home for
EXITS THRESHOLD OF A PROCESS MANAGED BY SPM IS REACHED
```

A customer reports that their system recently began logging an error.

Referring to the exhibit, which two statements are correct? (Choose two.)

- A. The bcomd application is responsible for monitoring the observed behavior.
- B. A user is attempting to restart the ntpd daemon.
- C. The spmd application failed to restart the ntpd daemon.
- D. The ntpd daemon is responsible for time services.

Answer: B,D

Question: 8

Click the exhibit button.

```
cluster1::> -vserver VSEXX -igroup esx8
igroup show
Vserver Name: VSEXX
Igroup Name: esx8
Protocol: fcp
OS Type: vmware
Portset Binding -
Igroup:
Igroup UUID 27da5172-7b8f-11e2-bbf9-123478563412
ALUA: true
Initiators 20:00:00:25:b5:01:a0:08 (logged in)
            20:00:00:25:b5:01:a0:11 (not logged in)
            20:00:00:25:b5:01:a0:08 (logged in)
            20:00:00:25:b5:01:a0:11 (not logged in)
```

A customer has an FC SAN with brocade switches. Each server attached to the SAN has four FC ports. The customer describes that one of their servers has only four paths to the LUNs configured on the cluster. There should be eight paths on all the other servers. The switch administrator has verified the zone configuration.

Referring to the exhibit, which two actions would you perform to troubleshoot the problem? (Choose two.)

- A. Review the LIF configuration on the cluster.
- B. Review the driver version on the server.
- C. Review that the host side FC link ports are enabled.
- D. Review the igroup configuration on the cluster.

Answer: A,B

Question: 9

End users are complaining of high latency when accessing an ONTAP 9 cluster.

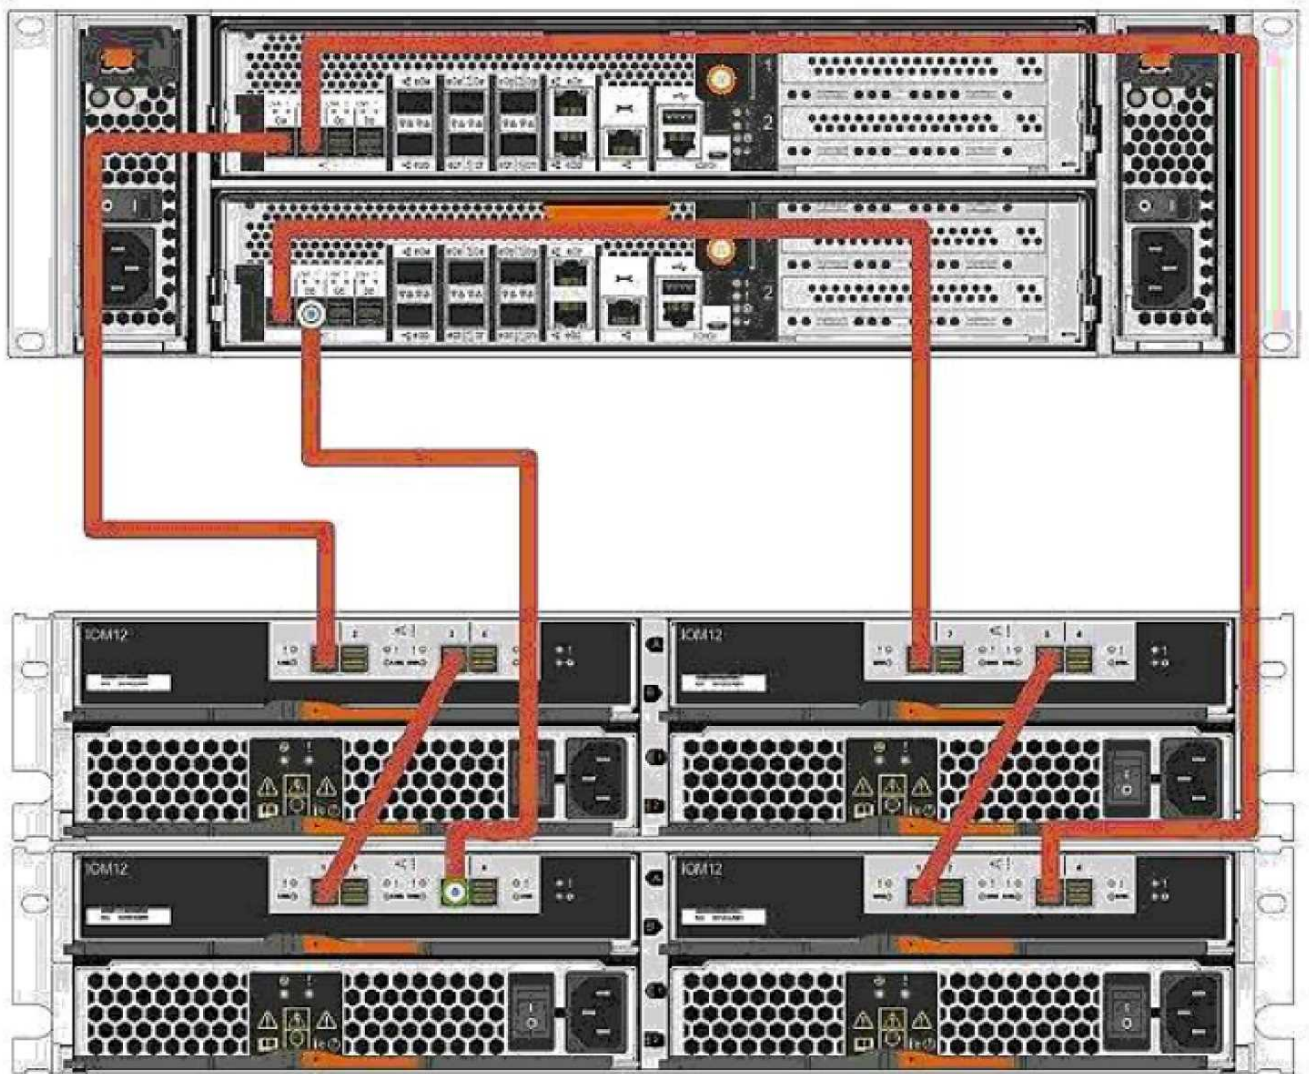
Which command helps you to determine whether there is a storage component that is contributing to this problem?

- A. qos statistics performance show
- B. qos statistics latency show
- C. sysstat -c 30 -M 1
- D. wafitop show -v -cpu -l 10 -c 5-n 6

Answer: A

Question: 10

Click the exhibit button.



You have just cabled a new shelf to a FAS8200. Running the system node `run -node * sysconfig` - a command indicates that multipath HA is configured. However, after running Config Advisor, you receive an error that your cabling is incorrect.

As indicated in the exhibit, which two actions will correct the cabling? (Choose two.)

- A. Move the 0b SAS connection on the bottom controller to 0d on the bottom controller.
- B. Move the 0b SAS connection on the top controller to 0d on the top controller.
- C. Move the 0b SAS connection on the top controller to 0d on the bottom controller.

D. Move the 0b SAS connection on the bottom controller to 0d on the top controller.

Answer: A,B

Question: 11

Click the exhibit button.

```
cluster01::> job show -id 18 -fields progress
```

```
id vserver  
progress
```

```
--  
-----  
-----
```

```
18 cluster01 "Cutover Deferred, Waiting for user intervention(10.30MB Sent)::Volume  
move job preparing transfer"
```

A volume move operation has been initiated to move a volume named MirrScr from one aggregate to another. The operation failed with the error that is shown in the exhibit.

In this scenario, which statement is correct?

- A. Manually trigger the cutover process in order to restart the job.
- B. The volume move cannot be completed at this time because of high use in this node.
- C. There are too many snapshots on the source volume.
- D. Cancel and restart the volume move operation.

Answer: A

Question: 12

Click the exhibit button.

===== SYSCONFIG-A =====

NetApp Release 8.3: Mon Mar 9 19:20:57 PDT 2015

System ID: 333333333 (NETAPPCLUS-01); partner ID: 444444444 (NETAPPCLUS-02)

System Serial Number: 8675309 (NETAPP1CLUS-01)

System Rev: B0

System Storage Configuration: Multi-Path HA

System ACP Connectivity: Full Connectivity

All-Flash Optimized: false

Backplane Part Number: 111-01459

Backplane Rev: D1

Backplane Serial Number: 021431000055

slot 0: System Board 2.8 GHz (System Board XX B0)

Model Name: FAS8080

slot 6: SAS Host Adapter 6c (PMC-Sierra PM8001 rev. C, SAS, <UP>)

MFG Part Number: USI 6086-000010-07 rev. MP V4.13

Part number: 111-00341+G1

Serial number: ST063EBJ11V

Date Code: 20141118

Firmware rev: 01.11.07.00

Base WWN: 5:00e004:000307b:50

Phy State: [0] Enabled, 6.0 Gb/s

[1] Enabled, 6.0 Gb/s

[2] Enabled, 6.0 Gb/s

[3] Enabled, 6.0 Gb/s

QSFP Vendor: Molex Inc.

QSFP Part Number: 112-00266+A0

QSFP Type: Optical (1490nm DFB) 2m ID:01

QSFP Serial Number: 93AG14035F0385

10.0 : NETAPP X425_HCBEP1T2A10 NA01 1142.3GB 520B/sect (K2K1KUTG)

Shelf 10: DS2246 Firmware rev. IOM6 A: 0187 IOM6 B: 0191

slot 12: SAS Host Adapter 12b (PMC-Sierra PM8001 rev. C, SAS, <UP>)

MFG Part Number: USI 6086-000010-07 rev. MP V4.13

Part number: 111-00341+G1

Serial number: ST063EBJ0YU

Date Code: 20141118

Firmware rev: 01.11.07.00

Base WWN: 5:00e004:000307b:3c

Phy State: [4] Enabled, 6.0 Gb/s

[5] Enabled, 6.0 Gb/s

[6] Enabled, 6.0 Gb/s

[7] Enabled, 6.0 Gb/s

QSFP Vendor: Molex Inc.

QSFP Part Number: 112-00266+A0

QSFP Type: Optical (1490nm DFB) 2m ID:01

QSFP Serial Number: 93AG13382F0001

10.0 : NETAPP X425_HCBEP1T2A10 NA01 1142.3GB 520B/sect (K2K1KUTG)

Shelf 10: DS2246 Firmware rev. IOM6 A: 0187 IOM6 B: 0191

===== ENVIRONMENT =====

Expander phy state on 12b.10 (NETAPP :DS224IOM6 SHFMS1449000979):

	Pwr	Invalid	Disprt	Loss	Phy	CRC	Phy				
	Pwr	Ccl	Rsv	Rls	DWord	Error	Sync	Reset	Error	Change	SATA
PHY/BAY STATE	ON	Cnt	Cnt	Cnt	Count	Count	Count	Prob	Cnt	Cnt	Affiliation
0/P0:0 6.0 Gb/s	---	---	---	---	0	0	0	0	0	45	-----
1/P0:1 6.0 Gb/s	---	---	---	---	0	0	0	0	0	45	-----
2/P0:2 6.0 Gb/s	---	---	---	---	0	0	0	0	0	45	-----
3/P0:3 6.0 Gb/s	---	---	---	---	0	0	0	0	0	45	-----
4/P1:0 6.0 Gb/s	---	---	---	---	0	0	0	0	0	2	-----
5/P1:1 6.0 Gb/s	---	---	---	---	0	0	0	0	0	2	-----
6/P1:2 6.0 Gb/s	---	---	---	---	0	0	0	0	0	2	-----
7/P1:3 6.0 Gb/s	---	---	---	---	0	0	0	0	0	2	-----

===== SYSCONFIG-A =====

```

NetApp Release 8.3: Mon Mar  9 19:20:57 PDT 2015
System ID: 44444444 (NETAPPCLUS-02); partner ID: 333333333 (NETAPP1CLUS-01)
System Serial Number: 9035768 (NETAPP1CLUS-02)
System Rev: B0
System Storage Configuration: Multi-Path HA
System ACP Connectivity: Partial Connectivity
All-Flash Optimized: false
Backplane Part Number: 111-01459
Backplane Rev: D1
Backplane Serial Number: 021432003770
slot 0: System Board 2.8 GHz (System Board XX B0)
      Model Name:      FAS8080
slot 6: SAS Host Adapter 6c (PMC-Sierra PM8001 rev. C, SAS, <UP>)
MFG Part Number:      USI 6086-000010-07 rev. MP V4.13
Part number:          111-00341+G1
Serial number:         ST063EBJ07F
Date Code:             20141118
Firmware rev:          01.11.07.00
Base WWN:              5:00e004:000307b:f0
Phy State:             [0] Enabled, 6.0 Gb/s
                      [1] Enabled, 6.0 Gb/s
                      [2] Enabled, 6.0 Gb/s
                      [3] Enabled, 6.0 Gb/s
QSFP Vendor:           Molex Inc.
QSFP Part Number:      112-00268+A0
QSFP Type:             Optical (1490nm DFB) 5m ID:01
QSFP Serial Number:    93AG14476F0081
10.0 : NETAPP X425_HCBEP1T2A10 NA01 1142.3GB 520B/sect (KZK1KUTG)
Shelf 10: DS2246 Firmware rev. IOM6 A: 0187 IOM6 B: 0191
slot 12: SAS Host Adapter 12b (PMC-Sierra PM8001 rev. C, SAS, <UP>)
MFG Part Number:      USI 6086-000010-07 rev. MP V4.13
Part number:          111-00341+G1
Serial number:         ST063EBJ0DH
Date Code:             20141118
Firmware rev:          01.11.07.00
Base WWN:              5:00e004:000307b:ac
Phy State:             [4] Enabled, Rate unknown
                      [5] Enabled, Rate unknown
                      [6] Enabled, 6.0 Gb/s
                      [7] Enabled, Rate unknown
QSFP Vendor:           Molex Inc.
QSFP Part Number:      112-00268+A0
QSFP Type:             Optical (1490nm DFB) 5m ID:01
QSFP Serial Number:    93AG14476F0137
10.0 : NETAPP X425_HCBEP1T2A10 NA01 1142.3GB 520B/sect (KZK1KUTG)
Shelf 10: DS2246 Firmware rev. IOM6 A: 0187 IOM6 B: 0191

```

===== ENVIRONMENT =====

```

Expander phy state on 12b.10 (NETAPP :DS224IOM6 SHFMS1449000979):
      Pwr      Invald Disprt Loss  Phy    CRC    Phy
      Pwr Ccl Rsv Rls DWord  Error  Sync  Reset  Error  Change SATA
PHY/BAY STATE      ON  Cnt Cnt Cnt Count  Count  Count  Prob  Cnt  Cnt  Affiliation
-----
0/P0:0 Disabled
1/P0:1 Disabled
2/P0:2 6.0 Gb/s    --- --- --- --- 0    0    0    0    0    39  -----
3/P0:3 Disabled
4/P1:0 6.0 Gb/s    --- --- --- --- 0    0    0    0    0    2  -----
5/P1:1 6.0 Gb/s    --- --- --- --- 0    0    0    0    0    2  -----
6/P1:2 6.0 Gb/s    --- --- --- --- 0    0    0    0    0    2  -----
7/P1:3 6.0 Gb/s    --- --- --- --- 0    0    0    0    0    2  -----

```

You are troubleshooting a system in which the customer reports that aggregates on Stack 6c/12b are having slow response times from system NETAPP1CLUS-02.

Referring to the exhibit, which two statements are correct? (Choose two.)

- A. The aggregates have only one PHY running at six gigabits per second on 6c.
- B. PHY 2 is running at six gigabits per second; on system NETAPP1CLUS-02, reseal IOM6 Module A on Shelf 10.
- C. The shelf module firmware is downrev; on system NETAPP1CLUS-02, replace IOM6 Module A on Shelf 10.
- D. The aggregates have only PHY at six gigabits per second on 12b.

Answer: A,C

Question: 13

What is a major difference between NFSv3 and NFSv4.x?

- A. NFSv4.x is stateful, and NFSv3 is stateless.
- B. NFSv4.x requires an auxiliary protocol for status monitoring, and NFSv3 does not.
- C. NFSv4.x is stateless, and NFSv3 is stateful.
- D. NFSv4.x requires an auxiliary protocol for mounting, and NFSv3 does not.

Answer: A

Question: 14

A customer is attempting to move from NFSv3 to NFSv4.x and reports that all attempts to mount using NFSv3 continue to be successful, however, all of their NFSv4.x mount attempts fail with an "access denied" error.

Which two statements would cause this error? (Choose two.)

- A. The NFS server does not have NFSv4 enabled.
- B. The SVM's root volume allows execute access for all users.
- C. The export policy applied to the data volume does not allow clients to access it using NFSv4.
- D. The export policy applied to the SVM's root volume does not allow the NFSv4 clients to access it.

Answer: B,C

Question: 15

In a 4-node cluster, nodes 1 and 2 are an HA pair, and nodes 3 and 4 are an HA pair. The following error appears in the log.

12/15/2018 15:48:47 node-2CRITICALclam.node.ooq: Node (name=node- 2, ID= 1001) is out of "CLAM quorum" (reason =quorum update).

In this scenario, which two statements are correct? (Choose two.)

- A. All LUNs are offline in the cluster.
- B. All LUNs are online in the cluster.
- C. Node 1 takes over node 2.
- D. Node 2 takes over node 1.

Answer: A,C

Question: 16

Click the exhibit button.

```
ONTAP_CLUSTER:> vservers cifs create -cifs-server svml -domain example.com -ou OU=Secure

In order to create an Active Directory machine account for the CIFS server, you must supply the
name and password of a Windows account with sufficient
privileges to add computers to the "OU=Secure" container within the "EXAMPLE.COM" domain.

Enter the user name: administrator

Enter the password:

Error: Machine account creation procedure failed
[ 61] Loaded the preliminary configuration.
[ 64] Successfully connected to ip 10.216.29.40, port 88 using
      TCP
[ 71] Requested service not found in Active Directory
      (KRB5KDC_ERR_S_PRINCIPAL_UNKNOWN)
[ 71] Failed to initiate Kerberos authentication. Trying NTLM.
[ 72] Successfully connected to ip 10.216.29.40, port 389 using
      TCP
**[ 77] FAILURE: Specified OU 'OU=Secure' does not exist in
**      EXAMPLE.COM

Error: command failed: Failed to create the Active Directory machine account "SVM1". Reason:
SecD Error: ou not found.
```

The storage administrator for the company example.com is attempting to join an SVM to their Active Directory as SVM1. The computer account must also be joined to the Secure organizational unit. Attempts to perform this action using the CLI fail with the error message shown in the exhibit.

What is the reason for this error?

- A. The requested OU Secure does not exist in the Active Directory and must be created by the Active Directory administrator.
- B. The requested OU Secure has incorrect principal and errored with

KRB5KDC_ERR_S_PRINCIPAL_UNKNOWN.

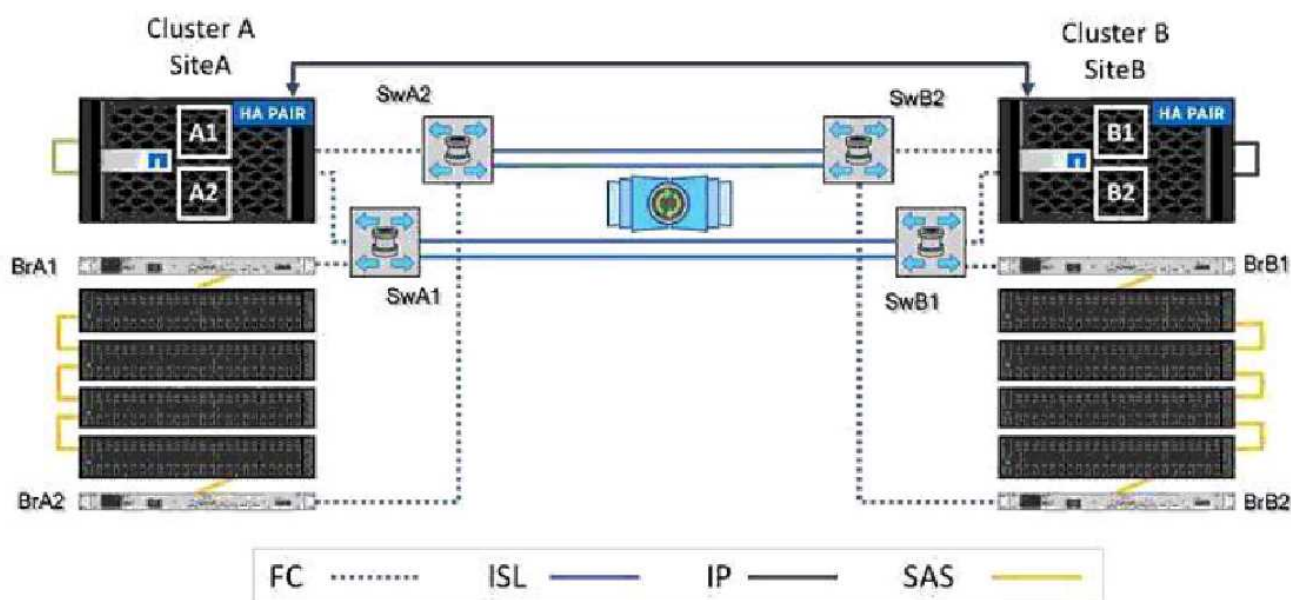
C. The administrator account password used in vserver cifs create was not entered correctly.

D. The syntax is incorrect; it should be vserver cifs create -cifs-server svm1 -domain example.com -ou=Secure, dc=example, dc=com.

Answer: B

Question: 17

Click the exhibit button.



You are investigating an AutoSupport Disk Redundancy Failed alert. Cluster A lost connectivity to storage bridge BrB1, while cluster B lost connectivity to storage bridge BrA1.

Referring to the exhibit, which three components need investigation? (Choose three.)

- A. storage bridges BrB1
- B. FC switch SwA1
- C. FC switch SwB1
- D. FC switch SwA2
- E. ISL cables

Answer: B,C,E

Question: 18

Two weeks ago, your customer experienced a data outage, which caused their MCC to switch over. When the power was restored, they found that 14 disks over three RAID groups had failed, which caused two data aggregates to lose their mirror plexes.

In this situation, before you perform a switchback, which three actions should be taken? (Choose three.)

- A. Delete the failed plexes.
- B. Initiate the aggregate healing phases.
- C. Adjust bridge configuration for the replaced disks.
- D. Mirror the unmirrored aggregates
- E. Assign all disks on the remote site to the surviving nodes.

Answer: B,C,D

Question: 19

After you replaced the motherboard, the node stopped in the “waiting for the cluster application to come online” state.

Which two actions should you perform to troubleshoot this problem? (Choose two.)

- A. Check the status of the HA with the interconnect-up field of storage failover show command.
- B. Verify the setting of the cluster interconnect switch.
- C. Verify that the physical links on the cluster ports are up.
- D. Check the status of the cluster with the cluster ring show command.

Answer: A,D

Question: 20

For restore purposes, a storage administrator on an ONTAP system created a FlexClone of a volume that contains a LUN from a specific snapshot. This LUN cannot be seen by a Windows client after a disk rescan is performed by using Disk Management.

In this scenario, which statement is correct?

- A. The cloned LUN must be brought online.
- B. The cloned LUN needs to be mapped to an initiator group.
- C. The cloned volume must be split from the source volume before the LAN can be used.
- D. The cloned LUN must be renamed to a different name than the name of the source LUN.

Answer: C

Question: 21

Click the exhibit button.

```
cl01::metrocluster interconnect> adapter show
```

Node	Adapter name	Adapter Type	Link Status	IP Address	Port Number
cl01-01	fcvi_device_0	FC-VI	Up	5.0.0.5	1a
cl01-01	fcvi_device_1	FC-VI	Up	6.0.0.6	1b
cl01-01	mlx4_0	IB	Down	192.0.5.56	ib0a
cl01-01	mlx4_0	IB	Down	192.0.5.57	ib0b

4 entries were displayed

Your customer is calling about their MetroCluster environment. They have provided you with the output of the metrocluster interconnect adapter show command.

Referring to the exhibit, which statement is correct?

- A. For all MetroCluster configurations, this is expected output.
- B. The InfiniBand cables should be connected in this configuration.
- C. The InfiniBand interconnect is not used in MetroCluster.
- D. This is a 2-node MetroCluster.

Answer: A

Question: 22

When you issue a volume clone create command, which entity in WAFL is copied?

- A. the vol_info block
- B. the buftree
- C. the data block
- D. the inode file

Answer: D

Question: 23

A customer called you after a failure in their Microsoft Hyper-V environment on the "SVM1" SVM. The customer wants to enable nondisruptive operations (NDO) to prevent future failures.

In this scenario, which two actions must be performed? (Choose two.)

- A. Verify that the continuously-available option is set on the CIFS share.
- B. Verify that SMB2 is enabled.
- C. Verify that SMB3 is enabled.
- D. Verify that the branchcache option is set on the CIFS share.

Answer: A,B

Question: 24

A customer has requested your assistance with replacing a DS224C disk shelf from their cluster nondisruptively.

Which two actions should be performed before physically replacing the shelf? (Choose two.)

- A. Verify that the system is configured as multipath high availability.
- B. Upgrade the IOMs in the shelf before starting the removal process.
- C. Verify that the drives in the shelves are spares and ownership is removed.
- D. Zero the disks before performing the physical shelf removal.

Answer: A,C

Question: 25

A customer's FAS8200 node1 went down an hour ago after a WAFL inconsistency error. All data is being served from partner node2, but the storage failover show command output is not showing a node1 as "Waiting for giveback".

What is the corrective action for this situation?

- A. Re-create the root aggregate on node1.
- B. Boot the system and run the wafliron command on the inconsistent data aggregate.
- C. From the LOADER prompt, run the wafliron command on the system's root aggregate.
- D. Boot the system and run the WAFL_check command on the inconsistent data aggregate.

Answer: A

Question: 26

Click the exhibit button.

```
netapp-cluster::> statistics volume show
```

```
netapp-cluster : 11/2/2018 13:34:36
```

Volume	Vserver	*Total Ops	Read Ops	Write Ops	Other Ops	Read (Bps)	Write (Bps)	Latency (us)
-----	-----	-----	-----	-----	-----	-----	-----	-----
nfs_volume	SVM04	2077	0	49	45	6567	2197980	194

A customer is reporting that numbers in the statistics volume show command output appear to be incorrect.

Referring to the exhibit, which statement is correct?

- A. The problem is due to a difference in operation size.
- B. Perfstat should be used to explain the behavior.
- C. There are operation types that are not listed in the output.
- D. A case with NetApp should be raised to find the root cause.

Answer: A

Question: 27

You are investigating a CIFS issue at a customer site, but it is unclear which SMB version is being used between the client and the server. The customer has provided you with a packet trace output.

In this scenario, which packet exchange contains information to indicate which SMB version was agreed upon by the client and the server?

- A. Ioctl1 Request FSCTL_DFS_GET_REFERRALS Ioctl1 Response FSCTL_DFS_GET_REFERRALS
- B. Session Setup Request Session Setup Response
- C. Negotiate Protocol Request Negotiate Protocol Response
- D. Tree Connect Request Tree Connect Request Response

Answer: D

Question: 28

Click the exhibit button.

The server supports Continuous Availability (persistent handles) and the request to create the handle succeeded. However, the server did not grant persistence. You should verify that the Resume Key Filter is running on the server and is attached to the target volume.

A customer reports that when they are failing over Hyper-V VMs, the process fails, and the server reports an error that indicates that there is something wrong with the storage.

Referring to the exhibit, which three actions would be required in ONTAP to solve this problem? (Choose three.)

- A. Verify that the -symlink-properties parameter is enabled on the share.
- B. Verify that the access-based-enumeration share parameter is configured.
- C. Verify that the continuously -available share parameter is configured.
- D. Verify that -offline-files parameter is set to manual on the share.
- E. Verify that the oplocks share property is configured.

Answer: B,C,E

Question: 29

Click the exhibit button.

```
cluster1::> aggr show
```

```
Info: Node cluster1-02 that hosts aggregate aggr0_cluster1_02 is offline
      Node cluster1-02 that hosts aggregate aggr2 is offline
```

Aggregate	Size	Available	Used%	State	#Vols	Nodes	RAID status
aggr0_cluster1_01	24.30GB	1.18GB	95%	online	1	cluster1-01	raid_dp, normal
aggr0_cluster1_02	-	-	-	unknown	-	cluster1_02	-
aggr1	255.8GB	1.18GB	99%	online	7	cluster1_01	raid_dp, normal
aggr2	-	-	-	unknown	-	cluster1_02	-

4 entries were displayed.

```
cluster1::> storage failover show
```

Node	Partner	Takeover Possible	State Description
cluster1_01	cluster1_02	true	Connected to cluster1_02
cluster1_02	cluster1_01	true	Connected to cluster1_01

2 entries were displayed.

You execute a few commands and many fail, as shown in the exhibit.

What could cause the error?

- A. The mgwd user space process on cluster1-02 failed.
- B. The aggregates were manually brought offline.

- C. The cluster-02 node is down.
- D. The bcomd user space process on cluster1-02 failed.

Answer: C

Question: 30

Your customer added a new DS2246 shelf to an existing stack on their FAS8060A system. The customer then noticed that the new shelf's disks appear with the shelf serial number instead of the shelf ID in the sysconfig -r command output.

What is the reason for this problem?

- A. The SAS is cabled incorrectly.
- B. There is failed IOM module in the new shelf.
- C. The ACP is cabled incorrectly.
- D. The new shelf has the same ID as another shelf.

Answer: A

Question: 31

A customer performs a takeover and giveback of their controllers and notices that several of their NFS clients lose access to their volumes for a short period of time during the SFO process. The customer reviewed the LIF failover policies, and the policies are correct.

What is the reason for this behavior?

- A. The clients are using NFSv4 to access the volumes.
- B. The export policies are incorrect and must be updated with node's HA partner LIF information.
- C. A firewall is blocking the necessary network ports and needs to be updated to allow NFS communication.
- D. NFSv4 file locks need to be broken manually before initiating an SFO operation.

Answer: A

Question: 32

Click the exhibit button.

```
Cluster::> vserver services name-service ns-switch show -vserver vs1
```

Vserver	Database	Source	Order
-----	-----	-----	-----
vs1	hosts	files,	dns
vs1	group	files,	ldap
vs1	passwd	files,	ldap
vs1	netgroup	files	
vs1	namemap	files	

5 entries were displayed.

```
Cluster::*> vserver services name-service ldap show
```

Vserver	Client Configuration
-----	-----
vs1	ad

A customer has the ns-switch shown in the exhibit.

Which two statements are true in this scenario? (Choose two.)

- A. User names not present in the local database are looked up in LDAP.
- B. User mapping decisions are always based on the LDAP configuration.
- C. Host names are always resolved using the configured DNS server.
- D. Groups can be looked up both in the local database and in LDAP.

Answer: C,D

Question: 33

A non-maskable interrupt event occurs. The system panics and then reboots.

In this scenario, what would have caused the panic?

- A. a secd process
- B. a watchdog timer
- C. an ICMP blackhole detector
- D. a volume location distributed database

Answer: B

Question: 34

A customer reports that one their SVMs has two default gateways. The customer wants to know whether this behavior is expected or if one of gateways should be deleted.

In this scenario, what do you tell the customer?

- A. The configuration is acceptable if the customer configures different metrics on all default routes.
- B. The configuration is acceptable only if both default gateways are in the same subnet.
- C. This situation is a misconfiguration, and the customer should delete one of the default gateways because there can be only gateway for each SVM.
- D. The configuration is acceptable. The customer can achieve load balancing between the default gateways by setting the same metric on all default routes.

Answer: A

Question: 35

You migrated from a FAS3250 system to a new A200 system. After the migration, you notice that data compaction fails to be enabled on the migrated volumes.

In this scenario, which statement is correct?

- A. You need to schedule background deduplication on all volumes.
- B. You need a license to have full efficiency on AFF systems.
- C. Data compaction is not supported on the A200 system.
- D. The volumes are thick-provisioned.

Answer: D

Question: 36

Your customer states that the storage efficiency on their FAS 8200 is far below what they had expected for a VMware NFS datastore. They enabled deduplication three days ago. The job ran every night for only 10-20 minutes, and the space savings now is less than 5%.

What is causing the space savings to be so low?

- A. They have enabled aggregate-level inline data compaction, so the savings are at the aggregate level.
- B. They have enabled aggregate-level inline deduplication, so the savings are at the aggregate level.
- C. They have not enabled volume-level compression.
- D. They have not run a baseline deduplication to scan old data.

Answer: C

Question: 37

Click the exhibit button.

```

disk          ut%  xfers  ureads--chain-usecs  writes--chain-usecs  creads--chain-usecs
greads--chain-usecs  gwrites--chain-usecs
/Aggr_SATA1/plex0/rq0:
0a.11.19      14  23.05   0.00   ....   . 11.11  62.42  199  11.94  58.82   156  0.00   ....   . 0.00   ....   .
0a.13.20      14  23.35   0.00   ....   . 11.40  60.89  213  11.96  58.81   160  0.00   ....   . 0.00   ....   .
0a.10.20      21  51.49  21.46  5.32 1781  9.13  28.54  470  20.91  12.12   725  0.00   ....   . 0.00   ....   .
0a.12.19      19  48.67  20.80  5.59 1644  8.34  27.18  492  19.53  12.45   675  0.00   ....   . 0.00   ....   .
0a.11.20      19  50.59  22.74  5.00 1597  8.26  27.94  456  19.59  11.65   668  0.00   ....   . 0.00   ....   .
0a.13.21      21  52.94  24.38  4.73 1912  8.57  27.95  476  19.99  11.76   737  0.00   ....   . 0.00   ....   .
0a.10.21      20  48.75  19.32  5.76 1568  8.71  28.67  429  20.72  12.00   662  0.00   ....   . 0.00   ....   .
0a.12.20      21  49.45  20.28  5.49 1702  9.02  28.07  492  20.15  12.32   686  0.00   ....   . 0.00   ....   .
0a.11.21      80  49.55  21.07  5.48 9799  8.53  28.23  932  19.95  11.61   982  0.00   ....   . 0.00   ....   .
0a.13.22      20  48.86  20.15  5.59 1734  8.79  28.43  455  19.91  12.17   736  0.00   ....   . 0.00   ....   .
0a.10.22      20  51.13  22.84  5.15 1788  8.69  27.24  469  19.59  11.63   712  0.00   ....   . 0.00   ....   .
0a.12.21      20  49.39  20.14  5.36 1640  8.94  30.22  438  20.32  12.48   669  0.00   ....   . 0.00   ....   .
0a.11.22      20  50.18  21.62  5.53 1664  8.55  28.12  480  20.01  11.75   730  0.00   ....   . 0.00   ....   .
0a.13.23      21  55.77  26.58  4.78 1778  8.95  28.90  486  20.23  12.81   663  0.00   ....   . 0.00   ....   .
0a.10.23      20  48.96  20.19  5.92 1750  8.78  28.77  454  19.99  11.96   686  0.00   ....   . 0.00   ....   .
0a.12.22      20  48.46  19.95  5.59 1755  8.71  28.92  481  19.80  12.33   668  0.00   ....   . 0.00   ....   .
0a.11.23      20  48.83  20.89  5.67 1771  8.44  29.17  437  19.50  12.00   704  0.00   ....   . 0.00   ....   .

```

Your customer has severe latency on one of their SATA aggregates. You run the `statit -b` and `statit -e` commands to obtain detailed disk statistics.

Referring to the output shown in the exhibit, which disk would be the problem?

- A. 0a.11.19
- B. 0a.11.21
- C. 0a.13.21
- D. 0a.11.23

Answer: A

Question: 38

Click the exhibit button.

```
cluster01::> broadcast-domain show -dbroadcast-domain Storage
```

IPspace Name	Broadcast Domain Name	MTU	Port List	Update Status	Details
Default	Storage	1500	cluster01-04:a0a-40 cluster01-04:eOM cluster01-05:a0a-40 cluster01-05:eOM cluster01-06:a0a-40 cluster01-06:eOM cluster01-03:a0a-40 cluster01-03:eOM	complete complete complete complete complete complete complete complete	

A customer notifies you that Active IQ is showing a warning that a broadcast domain in their environment has been configured with ports from multiple VLANs.

Referring to the exhibit, which two statements are correct? (Choose two.)

- A. You should move the eOM ports to a Broadcast Domain that is dedicated to only management traffic.
- B. A VLAN tag cannot be assigned to the eOM port, so the configuration is acceptable.
- C. You should configure eOM so that is tagged on VLAN 40.
- D. You should change the LIF service policy to net-intercluster to prevent the LIFs from failing over to the eOM ports.

Answer: C,D

Question: 39

You want to determine which SMB versions that a CIFS server can use when communicating with external servers.

Which two CIFS server settings should be reviewed to accomplish this task? (Choose two.)

- A. smb2-enabled
- B. smb1-enabled-for-dc-connections
- C. smb2-enabled-for-dc-connections
- D. smb1-enabled

Answer: C,D

Question: 40

A customer complains that the copy offload does not work when copying large file from a Windows 8 client between the CIFS shares that are hosted on the same ONTAP cluster.

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

- A. Deduplication is not enabled on the destination volume.
- B. The client should be using Windows 10 for ODX support.
- C. The source and target volumes reside on different cluster nodes.
- D. SMB 3.0 disabled on the CIFS SVM.

Answer: A,D

Question: 41

Click the exhibit button.

```
*****  
**  SYSTEM MESSAGES  **  
*****
```

```
The root volume is different than one previously used by this node. This node is not  
fully operational. Contact support personnel for the root volume recovery procedures.
```

When you log in to a node, a warning is displayed as shown in the exhibit.

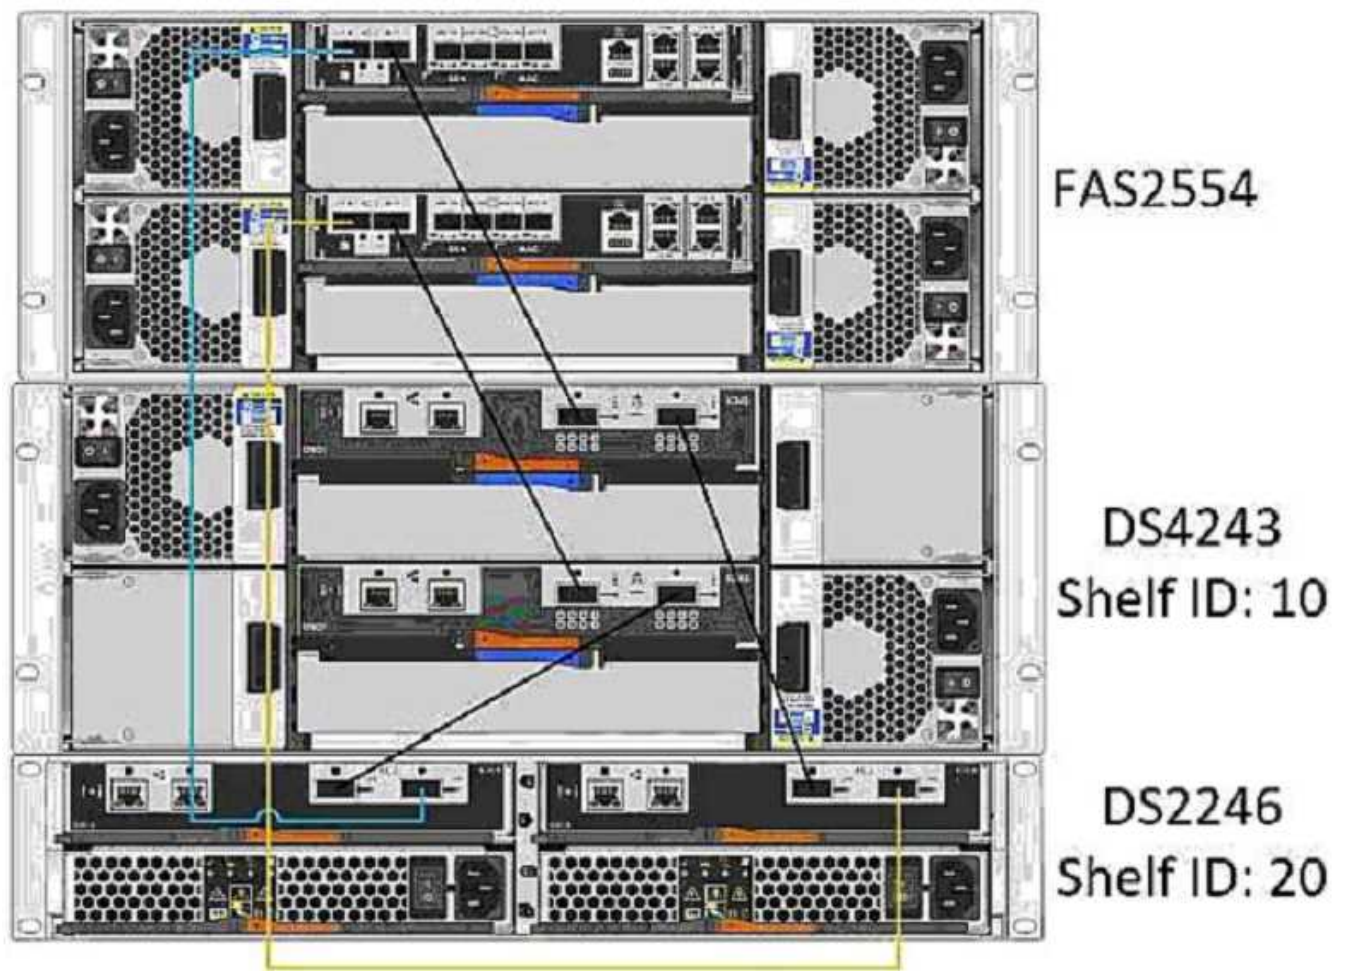
Which statement is correct in this scenario?

- A. The node root volume was moved to a new aggregate.
- B. The partner still hosts the CFO aggregate for the node.
- C. One or more RDB db files are damaged due to an internal software error.
- D. The node cannot communicate with the rest of the nodes in the cluster.

Answer: D

Question: 42

Click the exhibit button.



Your customer reports that their FAS2554A system reports multi-path HA when running the sysconfig command. However, Config Advisor reports an error about incorrect cabling.

What are the two reasons for the messages? (Choose two.)

- A. Shelves 10 and 20 have different IOM modules and must be in separate stacks.
- B. The FAS2554 system is cabled incorrectly to shelf ID 20.
- C. The DS2246 shelf is wrong; the shelf ID must be 11.
- D. Disk shelves 10 and 20 are cabled incorrectly to each other.

Answer: B,D

Question: 43

A field support engineer has just replaced the IOXM on NODE2 on a FAS8080 running ONTAP 9.1. After the replacement, the customer is unable to perform a giveback because the interconnect is down.

In this scenario, which statement is correct?

- A. Verify that there cross-links between the interconnects.
- B. Verify that the installed FC-VI cards are configured properly.
- C. Verify that the interconnect cables on the motherboard are connected correctly.
- D. Verify that the interconnect cables on the IOXM are connected correctly.

Answer: D

Question: 44

A customer reports that OnCommand Performance Manager shows that the workload on node-02 in a 4-node cluster has increased to a level that is higher than the threshold. The customer wants to avoid this issue becoming a performance problem.

Which two actions would result in achieving this goal? (Choose two.)

- A. Use the volume move command to move one or more volumes to another aggregate that is owned by another node that is less utilized.
- B. Use the storage aggregate relocation command to offload an aggregate to its partner.
- C. Use the storage aggregate scrub command to verify the parity on the busy aggregates.
- D. Physically move the shelf to another node that is less utilized.

Answer: A,D

Question: 45

Click the exhibit button.

```

...
slot 1: FC Host Adapter 1a (QLogic 2532 rev. 2, N-port, <UP>)
...
san31:5.126L0      : ATTO      FibreBridge6500N 1.50  FB6500N108390
san31:6.126L0      : ATTO      FibreBridge6500N 1.50  FB6500N108363
san51:5.126L0      : ATTO      FibreBridge6500N 1.50  FB6500N108406
san51:6.126L0      : ATTO      FibreBridge6500N 1.50  FB6500N108893
san51:5.shelf10    : DS2246    Firmware rev. IOM6 A: 0173  IOM6 B: 0173
san51:6.shelf31    : DS4246    Firmware rev. IOM6 A: 0173  IOM6 B: 0173
san31:5.shelf0     : DS2246    Firmware rev. IOM6 A: 0173  IOM6 B: 0173
san31:6.shelf20    : DS4246    Firmware rev. IOM6 A: 0173  IOM6 B: 0173
slot 1: FC Host Adapter 1b (QLogic 2532 rev. 2, N-port, <UP>)
...
san32:5.126L0      : ATTO      FibreBridge6500N 1.50  FB6500N107744
san52:5.126L0      : ATTO      FibreBridge6500N 1.50  FB6500N108403
san52:6.126L0      : ATTO      FibreBridge6500N 1.50  FB6500N108903
san52:5.shelf10    : DS2246    Firmware rev. IOM6 A: 0173  IOM6 B: 0173
san52:6.shelf31    : DS4246    Firmware rev. IOM6 A: 0173  IOM6 B: 0173
san32:5.shelf0     : DS2246    Firmware rev. IOM6 A: 0173  IOM6 B: 0173
slot 1: FC Host Adapter 1c (QLogic 2532 rev. 2, N-port, <UP>)
...
san31:5.126L0      : ATTO      FibreBridge6500N 1.50  FB6500N108390
san31:6.126L0      : ATTO      FibreBridge6500N 1.50  FB6500N108363
san51:5.126L0      : ATTO      FibreBridge6500N 1.50  FB6500N108406
san51:6.126L0      : ATTO      FibreBridge6500N 1.50  FB6500N108893
san51:5.shelf10    : DS2246    Firmware rev. IOM6 A: 0173  IOM6 B: 0173
san51:6.shelf31    : DS4246    Firmware rev. IOM6 A: 0173  IOM6 B: 0173
san31:5.shelf0     : DS2246    Firmware rev. IOM6 A: 0173  IOM6 B: 0173
san31:6.shelf20    : DS4246    Firmware rev. IOM6 A: 0173  IOM6 B: 0173
slot 1: FC Host Adapter 1d (QLogic 2532 rev. 2, N-port, <UP>)
...
san32:5.126L0      : ATTO      FibreBridge6500N 1.50  FB6500N107744
san52:5.126L0      : ATTO      FibreBridge6500N 1.50  FB6500N108403
san52:6.126L0      : ATTO      FibreBridge6500N 1.50  FB6500N108903
san52:5.shelf10    : DS2246    Firmware rev. IOM6 A: 0173  IOM6 B: 0173
san52:6.shelf31    : DS4246    Firmware rev. IOM6 A: 0173  IOM6 B: 0173
san32:5.shelf0     : DS2246    Firmware rev. IOM6 A: 0173  IOM6 B: 0173

```

A customer's system reported several errors. There was no disruption, but the customer is concerned about the system's health.

Referring to the exhibit, which statement is true?

- A. One of the switches has failed.
- B. One of the IOM modules on shelf 20 has failed.
- C. Storage redundancy is degraded.
- D. Port 1D on the system has failed.

Answer: C

Question: 46

You configured iSCSI multipathing on a Windows host to a 2-node ONTAP cluster so that you can attach to some LUNs. You see that the host logging in to both iSCSI LIFs in the ONTAP cluster, and you

notice that the igroup is correctly set. After you open Disk Management and rescan for disks, the NetApp LUNs do not show up.

In this scenario, what should be performed to enable the LUNs to be visible to the Windows host?

- A. Set up LUN mapping and try scanning again.
- B. Configure the SVM for FCP and try scanning again.
- C. Set up Selective LUN Mapping (SLM) and try scanning again.
- D. Configure the ONTAP nodes for HA and try scanning again.

Answer: C

Question: 47

A customer has a 2-node MetroCluster. They experienced a complete power loss at site B. The customer has lost connectivity to site B and now realizes that the DR partner at site A has not taken over the DR partner at site B.

Which two actions must be performed in this scenario? (Choose two.)

- A. On the DR partner at site A, perform the process to switchover.
- B. Ensure that site B is prevented from serving data or simply powered off (nodes and disk shelves).
- C. On the DR partner at site A, perform the process to heal the data aggregates.
- D. On the DR partner at site B, perform the process to heal the root aggregate.

Answer: A,D

Question: 48

A reboot was performed by a watchdog reset on the node.

In this scenario, which statement is true?

- A. The service processor no longer responds.
- B. In case of an L1 watchdog event, determine whether a panic message was recorded.
- C. In case of an L2 watchdog event, determine whether a core was dumped.
- D. The wafli_check command should be executed.

Answer: A

Question: 49

Which statement is correct about the wafliiron command?

- A. The wafiron command is used to correct parity inconsistencies.
- B. The wafiron command is used to correct inconsistent WAFL metadata.
- C. The wafiron command is used to correct WAFL i2p inconsistencies.
- D. The wafiron command is used to correct inconsistent user data blocks.

Answer: B

Question: 50

In a 4-node MetroCluster cluster, both clusters are alerting that the connection to an ATTO bridge is lost.

Which three components should be verified to determine what is causing this problem? (Choose three.)

- A. the SFPs that are connecting the ATTO bridge to the switch
- B. the dark fiber cables between sites
- C. the ATTO bridge
- D. the SFPs in the storage controllers
- E. the fiber cable that is connecting the ATTO bridge to the switch

Answer: A,C,E

Question: 51

A customer has created an interface group (igroup) named a0a, a VLAN named a0a-2500, and a data LIF associated with the VLAN interface that has been assigned an IP of 10.254.10.100. The customer is able to ping 10.254.10.100 from a Windows client without any issues, but data traffic is not working.

In this scenario, what is the solution to the problem?

- A. Disable flow control on the a0a interface.
- B. Configure an appropriate service policy for the data LIF.
- C. Correct the MTU mismatch between the broadcast domain configuration and the upstream network switches.
- D. Allow data traffic by using LIF firewall policies.

Answer: D

Question: 52

At a customer site, the storage failover giveback command fails when giving back the root aggregate.

What would be two reasons for this failure? (Choose two.)

- A. There are active NFS connections.
- B. There are failed disks.
- C. The HA interconnect is not operational.
- D. There are active SIS operations running.

Answer: B,C

Question: 53

The error message shown below is reported in the EMS log file.

“[sis.changelog.full:warning]: Change logging metafile on volume DATA1 is full and can not hold any more fingerprint entries.”

What are two ways to solve this problem? (Choose two.)

- A. Clear the changelog manually so that the operation can restart.
- B. Increase the size of the changelog to allow the operation to continue.
- C. Modify the storage efficiency policy so that deduplication runs more frequently.
- D. Modify the threshold that triggers a deduplication operation.

Answer: B,D

Question: 54

A customer reports that a LUN is offline because the volume is out of space even though the space-reserve parameter is configured to enabled on the LUN.

Which two statements are correct in this scenario? (Choose two.)

- A. The containing aggregate is thin-provisioned.
- B. The fractional reserve is set to 0%.
- C. Deduplication is inflating the LUN when it is brought online.
- D. The volume has snapshots.

Answer: A,D

Question: 55

Click the exhibit button.

```
cluster1::*> network interface show -vserver FS -fields data-protocol,curr-node,curr-port,is-
home
vserver lif    data-protocol curr-node    curr-port is-home
-----
FS        lif01 cifs          cluster1-01 e0d        true
FS        lif02 cifs          cluster1-02 e0d        true
2 entries were displayed.
```

```
cluster1::*> vsriver cifs share show -vserver FS -share-name Users
```

```

Vserver: FS
Share: Users
CIFS Server NetBIOS Name: FS1
Path: /Users
Share Properties: oplocks
                  browsable
                  changenotify
                  show-previous-versions
Symmlink Properties: symlinks
File Mode Creation Mask: -
Directory Mode Creation Mask: -
Share Comment: -
Share ACL: Everyone / Full Control
File Attribute Cache Lifetime: -
File Mode Creation Mask: -
Directory Mode Creation Mask: -
Share Comment: -
Share ACL: Everyone / Full Control
File Attribute Cache Lifetime: -
Volume Name: Users
Offline Files: manual
Vscan File-Operations Profile: standard
Maximum Tree Connections on Share: 4294967295
UNIX Group for File Create: -
```

```
cluster1::*> vsriver cifs options show -vserver FS
```

```

Vserver: FS

Client Session Timeout: 900
Copy Offload Enabled: true
Default Unix Group: -
Default Unix User: pcuser
Guest Unix User: -
Are Administrators mapped to 'root': true
Is Advanced Sparse File Support Enabled: true
Is Fsctl File Level Trim Enabled: true
Direct-Copy Copy Offload Enabled: true
Export Policies Enabled: false
Grant Unix Group Permissions to Others: false
Is Advertise DFS Enabled: false
Is Client Duplicate Session Detection Enabled: true
Is Client Version Reporting Enabled: true
Is DAC Enabled: false
Is Fake Open Support Enabled: true
Is Hide Dot Files Enabled: false
Is Large MTU Enabled: false
Is Local Auth Enabled: true
Is Local Users and Groups Enabled: true
Is Multichannel Enabled: false
Is NetBIOS over TCP (port 139) Enabled: true
Is NBNS over UDP (port 137) Enabled: false
Is Referral Enabled: true

...

SMB1 Enabled: false
Max Buffer Size for SMB1 Message: 65535
SMB2 Enabled: true
SMB3 Enabled: true
SMB3.1 Enabled: true
Map Null User to Windows User or Group: -
WINS Servers: -
Report Widelink as Reparse Point Versions: SMB1
Max Credits to Grant: 128
```

Referring to the exhibit, which change is required to enable the CIFS Autolocation function for the Users share?

- A. Edit the CIFS share's properties to remove the changenotify property.
- B. Re-create the LIFs to allow both CIFS and NFS protocols.
- C. Change the CIFS share's Symlink Properties parameter to symlinks.
- D. Modify the CIFS server's option to enable SMB multichannel.

Answer: C

Question: 56

A customer calls you about a performance issue that they experienced over the past few days. They want to know if data collection and root cause analysis (RCA) is possible for this issue. This issue appears to be a one time event.

Which tool would be used to satisfy the customer's request?

- A. perfstat
- B. perfmon
- C. OneCollect
- D. Performance Archive

Answer: C

Question: 57

A customer reports that copying data over iSCSI attached LUNs is taking significantly longer than usual. The customer asks you to help to identify the root cause of this problem.

Which two perfstat sections would you use to accomplish this task? (Choose two.)

- A. stats lun
- B. lun stats -vserver
- C. statit
- D. stats volume

Answer: B,C

Question: 58

Click the exhibit button.

```

disk          ut%  xfers  ureads--chain-usecs  writes--chain-usecs  cpreads-chain-usecs
greads--chain-usecs  gwrites-chain-usecs
/aggr0_600sas/plex0/rg0:
0a.00.0       4    8.30    0.28    1.00    0    4.10    36.27    164    3.92    39.52    150    0.00    ....    .    0.00    ....    .
3b.03.0       5    8.67    0.28    1.00    0    4.57    32.78    175    3.82    38.93    121    0.00    ....    .    0.00    ....    .
0a.00.1      48   97.79   90.89   7.34   1240   2.05   13.41   1142   4.85   10.58   727    0.00    ....    .    0.00    ....    .
3b.06.11     48   87.63   82.13   7.64   1229   1.96   30.38   649    3.54   11.71   575    0.00    ....    .    0.00    ....    .
0a.00.2     49   91.64   85.86   7.39   1145   1.77   37.21   304    4.01    8.44   694    0.00    ....    .    0.00    ....    .
3b.03.2     52   96.58   91.73   8.18   1132   1.77   26.53   589    3.08   13.27   1021   0.00    ....    .    0.00    ....    .
0a.00.3     51   94.90   88.47   8.30   1159   2.24   17.96   647    4.20   14.47   522    0.00    ....    .    0.00    ....    .
0a.03.3     50   93.50   87.54   7.66   1143   1.96   25.86   378    4.01   11.81   957    0.00    ....    .    0.00    ....    .
0a.00.4     50   98.91   93.60   6.73   1242   1.96   26.29   694    3.36   13.78   613    0.00    ....    .    0.00    ....    .
3b.03.4     50   97.70   92.10   6.51   1349   1.77   29.63   549    3.82    9.10   1185   0.00    ....    .    0.00    ....    .
0a.00.5     47   83.90   77.93   8.68   1081   1.86   34.55   405    4.10   10.68   700    0.00    ....    .    0.00    ....    .
0a.03.5     52   95.09   89.87   7.51   1255   1.68   28.00   466    3.54   19.45   451    0.00    ....    .    0.00    ....    .
0a.00.6     50   91.45   85.95   7.34   1257   1.77   22.21   557    3.73   15.03   539   0.00    ....    .    0.00    ....    .
3b.03.6     51   92.29   87.26   8.32   1179   1.77   28.68   442    3.26   14.71   619   0.00    ....    .    0.00    ....    .
0a.00.7     48   92.38   87.35   7.66   1107   1.58   26.65   380    3.45   13.59   799   0.00    ....    .    0.00    ....    .
0a.03.7     51   97.79   91.92   8.02   1161   1.96   21.29   454    3.92   11.43   725   0.00    ....    .    0.00    ....    .
3b.05.18     33   59.48   53.60   4.05   2151   1.86   29.60   404    4.01   11.30   580   0.00    ....    .    0.00    ....    .
3b.06.9      33   60.22   54.63   3.91   1907   2.33   38.84   286    3.26   13.91   466   0.00    ....    .    0.00    ....    .
3b.05.20     33   57.52   51.55   4.86   1855   1.96   36.43   390    4.01   13.60   579   0.00    ....    .    0.00    ....    .
3b.05.21     34   60.59   55.00   4.43   1984   1.96   33.67   368    3.64   14.72   451   0.00    ....    .    0.00    ....    .
3b.05.22     33   55.37   49.22   5.51   1641   2.24   39.96   303    3.92   14.60   721   0.00    ....    .    0.00    ....    .
3b.05.23     32   63.11   57.33   4.04   1892   1.96   33.10   367    3.82   11.07   780   0.00    ....    .    0.00    ....    .
/aggr0_600sas/plex0/rg1:
0a.00.8       3    5.50    0.00    ....    .    2.52   61.00   124    2.98   53.50   61    0.00    ....    .    0.00    ....    .
3b.03.8       2    5.50    0.00    ....    .    2.52   61.00   107    2.98   53.50   67    0.00    ....    .    0.00    ....    .
0a.00.9     50   87.63   80.92   7.55   1353   2.24   20.17   593    4.47   11.56   969   0.00    ....    .    0.00    ....    .
3b.04.23     51   96.02   89.87   7.88   1210   1.77   20.37   581    4.38   10.91   655   0.00    ....    .    0.00    ....    .
0a.00.10     52   99.10   93.78   6.93   1405   1.68   21.72   691    3.64   11.69   732   0.00    ....    .    0.00    ....    .
3b.03.10     52
101.33   95.74   6.89   1409   1.86   19.85   683   3.73    6.60   1170   0.00    ....    .    0.00    ....    .
0a.00.11     52   97.88   92.10   6.98   1271   1.68   17.17   874    4.10    7.32   1140   0.00    ....    .    0.00    ....    .
0a.03.11     47   88.38   81.66   8.34   1087   1.96   16.62   722    4.75    7.06   1522   0.00    ....    .    0.00    ....    .
0a.00.12     50   96.58   90.43   8.09   1136   2.14   16.48   760    4.01   11.37   630   0.00    ....    .    0.00    ....    .
3b.03.12     47   89.12   82.87   8.04   1117   2.14   17.91   755    4.10    8.05   887   0.00    ....    .    0.00    ....    .
0a.00.13     50   90.24   83.25   7.96   1209   2.14   18.61   551    4.85    9.15   592   0.00    ....    .    0.00    ....    .
0a.03.13     52   91.54   85.11   8.14   1319   2.05   25.82   572    4.38   10.96   786   0.00    ....    .    0.00    ....    .
0a.00.14     52
104.22   98.16   7.81   1194   1.86   16.05   654   4.20    8.20   791   0.00    ....    .    0.00    ....    .
3b.03.14     49   91.27   85.67    7.75   1127   1.77   20.32   806    3.82   11.27   680   0.00    ....    .    0.00    ....    .
0a.00.15     53
106.09   99.47   7.03   1303   2.14   21.52   685   4.47    7.85   899   0.00    ....    .    0.00    ....    .
0a.03.15     48   89.49   83.43   7.96   1238   2.05   16.77   783    4.01    9.91   742   0.00    ....    .    0.00    ....    .
3b.06.1      25   46.24   40.09   5.60   1454   2.14   29.17   463    4.01   15.12   418   0.00    ....    .    0.00    ....    .
0a.03.9      24   42.23   35.61   4.79   1721   2.42   27.50   302    4.20   16.73   331   0.00    ....    .    0.00    ....    .
3b.06.3      25   48.76   42.79   4.83   1582   2.14   38.35   290    3.82   19.44   442   0.00    ....    .    0.00    ....    .
3b.03.22     22   45.68   39.53   5.41   1257   2.24   28.54   349    3.92   13.00   546   0.00    ....    .    0.00    ....    .
3b.06.5      24   47.54   41.48   5.31   1280   2.24   35.83   291    3.82   18.00   314   0.00    ....    .    0.00    ....    .
3b.06.7      26   49.04   41.86   5.14   1502   2.33   35.08   383    4.85   15.35   434   0.00    ....    .    0.00    ....    .

```

A customer reports that storage latency has recently increased substantially. No changes have been introduced, and the traffic did not change.

Referring to the output shown in the exhibit, which corrective step should be taken?

- A. Add more spindles.
- B. Throttle the client traffic.
- C. Perform the volume reallocation.
- D. Move some volumes off the affected aggregate.

Answer: B

Question: 59

Click the exhibit button.

----- SAS-EXPANDER-MAP -----

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	' SHJHU1516001218' , 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	SHJHU1516001218' , 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	500a09800773b8bf,	ID	2, Serial Number	' SHJHU1653000119' , Product
'DS424IOM6			', Rev '0191', Slot B			
Level	2:	WWN	500a098005af77bf,	ID	1, Serial Number	' SHJHU1516001218' , Product
'DS224IOM6			', Rev '0191', Slot B			
Level	3:	WWN	500a09800772e3bf,	ID	2, Serial Number	' SHJHU1653000119' , Product 'DS424IOM6 ', Rev '0191', Slot A
Level	4:	WWN	500a098005b053bf,	ID	1, Serial Number	' SHJHU1516001218' , Product 'DS224IOM6 ', Rev '0191', Slot A
Level	5:	WWN	500a098005f0ed3f,	ID	0, Serial Number	' SHJHU1527000148' , Product 'DS424IOM6E ', Rev '0172', Slot A

Your customer added a new DS4246 shelf to their FAS system. They created an aggregate on the new shelf. Two weeks later, they log in to Active IQ and discover the Medium Impact error shown below.

“Both SAS connections from one controller go to the same shelf.”

Referring to the exhibit, which statement is correct?

- A. You are unable to obtain a “Cable Visualization” overview from Active IQ.
- B. A faulty cable is causing the problem.
- C. A faulty IOM is causing the problem.
- D. You can correct this problem without downtime on the FAS system.

Answer: D

Question: 60

Click the exhibit button.


```

offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<20us:0
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<40us:0
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<60us:0
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<80us:0
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<100us:0
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<200us:0
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<400us:0
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<600us:1
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<800us:156
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<1ms:6546
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<2ms:20591
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<4ms:40357
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<6ms:9414
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<8ms:2411
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<10ms:1452
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<12ms:387
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<14ms:237
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<16ms:353
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<18ms:273
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<20ms:132
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<40ms:0
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<60ms:0
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<80ms:0
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<100ms:0
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<200ms:0
offbox_vscan_server:vsl:10.23.23.54:nb-sfs-01:scan_latency_histogram.<400ms:0

```

A storage administrator reports that when they enabled Vscan, they began receiving reports of CIFS users having very slow access to all their files. The storage administrator recovered from the issue by disabling Vscan on their SVM. The storage administrator was able to collect perfstat output during the time of the reported slow access, as shown in the exhibit.

In this situation, what is the problem?

- A. The problem is on the SVM, as shown in the scan_latency output.
- B. Vscan server 192.168.100.144 was slower than Vscan server 10.23.23.54 in completing Vscan requests.
- C. Vscan server 10.23.23.54 was slower than Vscan server 192.168.100.144 in completing Vscan requests.
- D. The problem was with both Vscan servers 192.168.100.144 and 10.23.23.54.

Answer: B

Question: 61

A new cluster running ONTAP 9.5 has just been installed and you want to create a FlexGroup volume. Which action will accomplish this task?

- A. Use the storage pool create command.
- B. Use the volume create command.
- C. Use the Deploy utility to provision the FlexGroup volume.
- D. Use the volume size command.

Answer: A

Question: 62

An administrator wants to keep an automatic daily up-to-date copy of a volume at a disaster recovery location. What should the administrator do to achieve this goal?

- A. Perform a volume relocation to the remote location.
- B. Configure a SnapMirror relationship to the remote location.
- C. Configure a SyncMirror relationship to the remote location.
- D. Perform an NDMP dump to the remote location.

Answer: B

Question: 63

You are running a fabric-attached MetroCluster configuration with ONTAP 9.5 software. There is a maintenance window for the upcoming weekend. Plexes in your mirrored aggregates between your

two data centers have not been synchronizing for two days.

When the resynchronizations begin, they are not as quick as expected, and the process takes more than a day.

You have been modifying the number of concurrent I/O streams using the storage raid-options modify raid.resync.num_concurrent_ios_per_rg command, but the resynchronization speed stays slow.

You have four separate ISLs between your data centers using Brocade 6510 switches. Hardware problems are not an issue.

Which two actions should you take to prevent this problem in the future?

(Choose two.)

A. Update to at least ONTAP 9.4 software.

B. Restart the resynchronizations after modifying storage raid-options modify raid.resync.num_concurrent_ios_per_rg command.

C. Avoid changes on your aggregates during the maintenance window.

D. Take the corresponding plexes offline before maintenance to avoid a base-level synchronization.

Answer: AB

Question: 64

Which two statements are true about a data volume? (Choose two.)

A. A data volume can be accessed by a LIF on any node.

B. A data volume can be moved to any aggregate with space for it in the cluster.

C. A data volume can be accessed by a user connected to any SVM.

D. An SVM root or root load sharing mirror volume must be present on every node.

Answer: AB

Question: 65

You have a 2-node ONTAP Select HA deployment. In this scenario, which application manages the quorum, if one node fails?

A. ONTAP Mediator

B. ONTAP Deploy VM

C. vCenter Server

D. Tiebreaker

Answer: B

Thank You for trying NS0-592 PDF Demo

To try our NS0-592 practice exam software visit link below

<https://www.braindumpscollection.com/NS0-592.html>

Start Your NS0-592 Preparation

Use Coupon “20OFF” for extra 20% discount on the purchase of Practice Test Software. Test your NS0-592 preparation with actual exam questions.