



Exam Code: NS0-593

Exam Name: NetApp certified support engineer - ONTAP specialist

Website: <https://VCEup.com/>

Team-Support: <https://VCEplus.io/>



Question No: 1

When you review performance data for a NetApp ONTAP cluster node, there are back-to-back (B2B) type consistency points (CPs) found occurring on the root aggregate.

In this scenario, how will performance of the client operations on the data aggregates be affected?

- A. During B2B processing, clients will be unable to write data.
- B. Data aggregates will not be affected by B2B processing on another aggregate.
- C. During B2B processing, all I/O to the node is stopped.
- D. During B2B processing, clients will be unable to read data.

Answer: B

Explanation:

Question No: 2

Recently, a CIFS SVM was deployed and is working. The customer wants to use the Dynamic DNS

(DDNS) capability available in NetApp ONTAP to easily advertise both data UFs to their clients.

Currently, DNS is only responding with one data LIF. DDNS is enabled on the domain controllers.

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

cluster1:~> vserver services dns dynamic-update show				
Vserver	Is-Enabled	Use-Secure	Vserver FQDN	TTL
svml	false	false	svml.demo.net	24h

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

(Choose two.)

- A. Disable the -vserver-fqdn parameter for the SVM DDNS services.
- B. Remove the NFS protocol from the cifs_01 data LIF.
- C. Enable the -use-secure parameter for the SVM DDNS services.
- D. Enable the -is-enabled parameter for the SVM DDNS services

Answer: A, D

Explanation:

Question No: 3

A customer is calling you to troubleshoot why users are unable to connect to their CIFS SVM.

```
ClusterB::*> storage disk show -broken

Original Owner: Node03
Checksum Compatibility: block

Physical
Disk      Outage Reason  HA Shelf Bay /Slot  Drawer  Usable
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      4875  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      4875  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)
Flex /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)
-----
parity  FAILED
parity  3c.0.11  3c  0  11  SA:B  0  BSAS  7200  2538546/5198943744
2543634/5209362816
data    3c.0.12  3c  0  12  SA:B  0  BSAS  7200  2538546/5198943744
2543634/5209362816
data    3c.0.13  3c  0  13  SA:B  0  BSAS  7200  2538546/5198943744
2543634/5209362816
data    3c.0.14  3c  0  14  SA:B  0  BSAS  7200  2538546/5198943744
2543634/5209362816
```

Referring to the Information shown in the exhibit, what is the source of the problem?

- A. The v1db database is offline.
- B. The aggregate aggr2 has a failed disk.
- C. The databases on Node03 must be switched from secondary to master.
- D. The broken disk in Node03 is the source of the problem.

Answer: C

Explanation:

Question No: 4

You have a customer who is concerned with high CPU and disk utilization on their SnapMirror destination system. They are worried about high CPU and disk usage without any user operations.

In this situation, what should you tell the customer?

- A. Suggest that the customer manually cancel any scanners on the destination to reduce CPU usage.
- B. Explain that background tasks such as SnapMirror throttle up in the absence of user workload.
- C. Suggest that the customer throttle their SnapMirror relationships to reduce resource consumption.

D. Explain that only user workload should use the CPU and Investigate further.

Answer: A

Explanation:

Question No: 5

You are attempting to connect a NetApp ONTAP cluster to a very complex network that requires LIFs to fail over across subnets.

How would you accomplish this task?

A. Configure an equal number of UFs on each subnet.

B. Configure VIP LIFs using OSPF.

C. Configure VIP LIFs using BGP.

D. Configure a LIF failover policy for each subnet inside a single broadcast domain.

Answer: C

Explanation:

Question No: 6

Which two statements about NetApp Cloud Volumes ONTAP licenses are true? (Choose two.)

A. Having an Essentials package enables you to convert to another licensing option.

B. Capacity-based license packages support only single-node configurations.

C. BYOL licenses are purchased directly from NetApp.

D. AWS Marketplace contracts cannot be mixed with BYOL.

Answer: B, C

Explanation:

Question No: 7

Your customer complains that a host will constantly report losing a connection to the iSCSI target and then report that the session was reestablished.

```
Tue Sep 04 16:18:26 CDT [PRD-NTAP-01: iswti_iscsi_thread: iscsi.notice:notice]: ISCSI: New
session from initiator iqn.1991-05.com.microsoft:prd-app-01.netapp.com at IP addr 172.20.10.80
Tue Sep 04 16:18:26 CDT [PRD-NTAP-01: iswti_iscsi_thread: iscsi.warning:warning]: ISCSI: New
session request from initiator iqn.1991-05.com.microsoft:prd-app-01.netapp.com, a session from
this initiator already exists.
Tue Sep 04 16:19:02 CDT [PRD-NTAP-01: iswti_iscsi_thread: iscsi.notice:notice]: ISCSI: New
session from initiator iqn.1991-05.com.microsoft:prd-app-01.netapp.com at IP addr 10.10.20.60
Tue Sep 04 16:20:09 CDT [PRD-NTAP-01: iswti_iscsi_thread: iscsi.warning:warning]: ISCSI: New
session request from initiator iqn.1991-05.com.microsoft:prd-app-01.netapp.com, a session from
this initiator already exists.
Tue Sep 04 16:20:11 CDT [PRD-NTAP-01: iswti_iscsi_thread: iscsi.notice:notice]: ISCSI: New
session from initiator iqn.1991-05.com.microsoft:prd-app-01.netapp.com at IP addr 10.10.20.60
Tue Sep 04 16:20:11 CDT [PRD-NTAP-01: iswti_iscsi_thread: iscsi.warning:warning]: ISCSI: New
session request from initiator iqn.1991-05.com.microsoft:prd-app-01.netapp.com, a session from
this initiator already exists.
Tue Sep 04 16:20:45 CDT [PRD-NTAP-01: iswti_iscsi_thread: iscsi.notice:notice]: ISCSI: New
session from initiator iqn.1991-05.com.microsoft:prd-app-01.netapp.com at IP addr 172.20.10.80
```

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 second host with an IP address of 10.10.20.60 have the same IQN.

- B. 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 in different initiator groups.
- 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 are accessing the same LUN.
- D. A host with an IP address of 172.20.10.80 and a second IP address of 10.10.20.60 is accessing different LUNs.

Answer: C

Explanation:

Question No: 8

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

ANY1	ANY2	ANY3	ANY4	ANY5	ANY6	ANY7	ANY8	Avg	CPU0	CPU1	CPU2	CPU3	CPU4	CPU5	CPU6	CPU7
NwK_Ecol	NwK_Ig	NwK_Empd	Protocol	Storage	Raid	Raid_Ex	Xor_Ex	Target	Kahuna	NATL_Ex	(Kahui)					
99%	97%	83%	50%	52%	74%	66%	46%	52%	62%	54%	54%	53%	53%	54%	53%	1%
(82%)		2%	0%	37%	0%	13%	26%	40246	25%							1%
1%	0%	37%	0%	14%	7%	48963	25%									1%
100%	89%	86%	71%	56%	50%	72%	44%	56%	53%	57%	57%	57%	56%	56%	57%	57%
(75%)		7%	0%	41%	0%	13%	5%	39012	44%							1%
99%	99%	85%	70%	55%	75%	71%	52%	55%	67%	56%	56%	57%	56%	56%	56%	1%
(71%)		0%	0%	39%	0%	13%	22%	40034	37%							1%
99%	99%	86%	71%	56%	50%	74%	47%	57%	50%	55%	55%	55%	55%	55%	55%	1%
(77%)		6%	0%	43%	0%	13%	5%	39573	70%							1%
99%	95%	84%	59%	53%	78%	62%	55%	53%	65%	54%	54%	54%	54%	54%	54%	1%
(76%)		3%	0%	37%	0%	14%	11%	42391	10%							1%
99%	99%	86%	72%	55%	54%	79%	54%	59%	55%	55%	59%	60%	59%	59%	60%	1%
(77%)		4%	0%	44%	0%	12%	11%	35750	59%							1%
99%	99%	85%	70%	54%	75%	70%	53%	55%	69%	56%	56%	56%	55%	55%	56%	1%
(76%)		1%	0%	43%	0%	13%	5%	35792	30%							1%
99%	95%	84%	59%	53%	77%	70%	52%	54%	67%	55%	56%	55%	55%	56%	56%	1%
(77%)		3%	0%	47%	0%	14%	7%	36943	47%							1%

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: C

Explanation:

Question No: 9

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: C

Explanation:

Question No: 10

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

How would your customer prevent such problems in the future?

- A. Quiesce the SnapMirror relationships before upgrading the ONTAP software.
- B. Verify that the cabling of the hardware port that is responsible for SnapMirror transfers is correct.
- C. Modify the schedules of the SnapMirror relationships after upgrading the ONTAP software.
- D. Delete the SnapMirror relationships and create them new after upgrading the ONTAP software.

Answer: B

Explanation:

Question No: 11

A customer enabled NFSv4.0 on an SVM and changed the client mount from NFSv3 to NFSv4.

Afterwards, the customer found that the directory owner was changed from root to nobody.

In this scenario, which statement is true?

- A. The customer did not configure name services on the SVM.
- B. The clients must be restarted to start using NFSv4.
- C. The export policy is not configured properly.
- D. The ID mapping domains do not match between the client and server.

Answer: C

Explanation:

Question No: 12

You have a NetApp ONTAP cluster consisting of four NetApp FAS8200 controllers with two NetApp CN1610 cluster switches running ONIAP 9.8 software. You are receiving several alert messages stating that the cluster network has degraded. After troubleshooting, you determine that the errors are being generated from Node 2, interface e0b.

In this scenario, what should you do first to solve this problem?

- A. Replace the Twinax cable between Node 2, Interface e0b. and the NetApp CN1610 switch.
- B. Replace the motherboard on Node 2.
- C. Replace both NetApp CN1610 switches.
- D. Replace the NetApp CN1610 switch that connects to Node 2, interface e0b.

Answer: B

Explanation:

Question No: 13

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


```
cluster1::*> network interface show -vserver svml_cluster1
Vserver      Logical      Status      Network      Current-      Current Is
Interface    Admin/Oper   Address/Mask Node          Port         Home
-----
svml_cluster1
cifs_01      up/up        192.168.0.131/24 cluster1-01   e0f          true
cifs_02      up/up        192.168.0.132/24 cluster1-02   e0f          true
jumbo_01     up/up        192.168.0.201/24 cluster1-01   e0g          true
jumbo_02     up/up        192.168.0.202/24 cluster1-02   e0g          true
mgmt         up/up        192.168.0.130/24 cluster1-01   e0c          true
5 entries were displayed.

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

Referring to the exhibit, which ONTAP command will help isolate a possible MTU mismatch?

A.

```
network ping -lif jumbo_01 -vserver svml_cluster1 -destination 192.168.0.210 -
disallow-fragmentation false -packet-size 1000
```

B.

```
network ping -lif jumbo_01 -vserver svml_cluster1 -destination 192.168.0.210 -
disallow-fragmentation true -packet-size 9000
```

C.

```
network ping -lif jumbo_01 -vserver svml_cluster1 -destination 192.168.0.210 -
disallow-fragmentation true -packet-size 1000
```

D.

```
network ping -lif jumbo_01 -vserver svml_cluster1 -destination 192.168.0.210 -
disallow-fragmentation true -packet-size 5000
```

Answer: B

Explanation:

Question No: 14

You have a customer complaining of long build times from their NetApp ONTAP-based datastores.

They provided you packet traces from the controller and client. Analysis of these traces shows an average service response time of 1 ms. QoS output confirms the same. The client traces are reporting an average of 15 ms in the same time period.

In this situation, what would be your next step?

A. The cluster is responding slowly and requires further investigation using performance archives.

B. The client that reports high latency should be investigated.

C. The cluster interconnects should be investigated.

D. A sync core should be triggered.

Answer: B

Explanation:

Question No: 15

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.

```

c11::> vservers fcp show -vservers SVM1
Vservers Name: SVM1
Target Name: 20:04:00:a0:98:a7:ca:18
Administrative Status: up

c11::> lun show -vservers SVM1
Vservers Path State Mapped Type Size
-----
SVM1 /vol/esx/datastore.lun online mapped vmware 900GB

c11::> lun igroup show -vservers SVM1
Vservers Igroup Protocol OS Type Initiators
-----
SVM1 esxserver fcp vmware 20:00:00:25:b5:12:1a:08
20:00:00:25:b5:12:1b:08

c11::> lun mapping show -vservers SVM1
Vservers Path Igroup LUN ID Protocol
-----
SVM1 /vol/esx/datastore.lun esxserver 0 fcp

c11::> vservers fcp initiator show -vservers SVM1
This table is currently empty.

```

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: C

Explanation:

Question No: 16

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

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

What should be your first action to correct this issue?

- A. Power cycle all the disk storage shelves that contain drives of the aggregate with the my_data_vol volume.
- B. Determine the root cause behind the inconsistency before attempting any recovery procedure.
- C. Use the storage takeover command on the storage controller that contains my_data_vol.
- D. Use the wafliron command against my_data_vol to solve the inconsistency on the volume.

Answer: A

Explanation:

Question No: 17

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

```
cluster1:~> vservers cifs share create -vsrvr svml -share-name voll -
path /voll
command failed: The specified path "/voll" does not exist in the namespace belonging to Vserver
"svml".
Errors:
```

```
cluster1:~> vservers cifs share show
Vserver      Share      Path      Properties Comment  ACL
-----
svml         admin$     /          browseable -      -
svml         c$         /          oplocks    -      BUILTIN\Administrators /
Full Control
svml         ipc$       /          browseable
              changenotify
              show-previous-versions
              browseable -      -
3 entries were displayed.
```

```
cluster1:~> vservers cifs show
Server      Status      Domain/Workgroup Authentication
Vserver     Name        Admin       Name        Style
-----
svml        SVML        up          DEMO         domain
```

```
cluster1:~> volume show -vsrvr svml -volume voll
Vserver Name: svml
Volume Name: voll
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: svml root
Vserver Root Volume: false
Comment:
Available Size: 18.76MB
Filesystem Size: 20MB
Total User-Visible Size: 19MB
Used Size: 244KB
Used Percentage: 1%
...
Volume Tiering Policy: none
Volume Tiering Minimum Cooling Days: -
Performance Tier Inactive User Data: -
Performance Tier Inactive User Data Percent: -
```

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

- A. The volume must have a type of DP.
- B. The volume has not been mounted.

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

D. The CIFS service is not in workgroup mode.

Answer: B

Explanation:

Question No: 18

The motherboard of Node-01 is being replaced. To perform this task, a takeover was Initiated from Node-02. Node-02 panics showing the string below.

PANIC; Permanent errors on all HA mailbox disks (while marshalling header) in SKprocess fnmbx_instanceWorker on release 9.5P8 (C) What has happened in this situation?

A. While in maintenance mode, the sldiag device show command was executed.

B. While in maintenance mode, the fcadmin device_map command was executed.

C. While in maintenance mode, the HA_config show command was executed.

D. While in maintenance mode, the mailbox destroy local command was executed.

Answer: B

Explanation:

Question No: 19

You are deploying NetApp ONTAP Select. When trying to add VMware vCenter to the inventory, you receive the message shown below.

Credential authentication failed for x.x.x.x with: HostConnectionFailed In this scenario, how would you correct this problem?

A. Change the vCenter password.

B. Change the IP address of the ONTAP Select Deploy VM.

C. Change the gateway address of the ONTAP Select Deploy VM.

D. Change the IP address of the vCenter server.

Answer: A

Explanation:

Question No: 20

You created a new NetApp ONTAP FlexGroup volume spanning six nodes and 12 aggregates with a total size of 4 TB. You added millions of files to the FlexGroup volume with a flat directory structure totaling 2 TB, and you receive an out of space error message on your host.

What would cause this error?

A. The maximum number of volume constituents has been reached in the ONTAP software.

B. All constituent volumes are full.

C. The inode limit is exceeded in the ONTAP software.

D. The maxdirsize is exceeded in the ONTAP software.

Answer: C

Explanation:

Question No: 21

```
MCC-A::> metrocluster show
```

Configuration: IP-fabric		
Cluster	Entry Name	State

Local: MCC-A	Configuration State	configured
	Mode	switchover
	AUSO Failure Domain	auto-disabled
Remote: MCC-B	Configuration State	configured
	Mode	waiting-for-switchback
	AUSO Failure Domain	auto-disabled

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

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

Answer: B

Explanation:

Question No: 22

Your customer mentions that they have accidentally destroyed both root aggregates in their twonode cluster.

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

- A. Rejoin the second node to the re-created cluster.
- B. Re-create the cluster from the local backup.
- C. Install ONTAP from a USB device.
- D. Re-create the cluster from the remote backup.

Answer: A, C

Explanation:

Question No: 23

You have recently discovered that NetApp ONTAP Cloud Manager is not sending AutoSupport messages to NetApp.

In this scenario, what would solve this Issue?

- A. Verify that your AWS credentials are correctly added to Cloud Manager.
- B. Verify that your NetApp Support site credentials are correctly added to Cloud Manager.
- C. Verify that your data LIFs allow traffic to mysupport .netapp.com on port 443.
- D. Verify that the inbound and outbound rules allow AutoSupport on port 80.

Answer: C

Explanation:

Question No: 24

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' (Veervert 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: A

Explanation:

Question No: 25

You Just used the CLI to create a NetApp ONTAP FlexGroup Volume on a NetApp Cloud Volumes ONTAP instance. After creation, you notice odd behavior in NetApp Cloud Manager.

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

- A. Cloud Volumes ONTAP does not support FlexGroup volumes.
- B. The Capacity Management Mode within Cloud Manager should be set to Manual.
- C. The Capacity Management Mode within Cloud Manager should be set to Automatic.
- D. Cloud Manager must be used to provision the FlexGroup volume.

Answer: D

Explanation:

Question No: 26

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.

• sysstat -x output (ONTAP 9)

CPU	MEM	CPU%	MEM%	Total	Net	MB/s	Clash	MB/s	Uops	MB/s	Cache	Clash	CU	CP_01	CP_02	SLA	CP02	CPD	CP02	CPD	MB/s	CP02	MB/s	CPD	MB/s	CPD
47%	0	1%	0	54%	101%	242	0	101%	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
14%	0	1	0	2	7	101%	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
23%	0	21%	0	21%	101%	224	231%	231%	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
47%	0	11%	0	31%	201%	44%	0	32%	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
57%	0	35%	0	36%	207%	38%	0	35%	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
47%	0	23%	0	26%	201%	15%	0	36%	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
13%	0	1	0	1	24%	91	0	24%	493%	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
14%	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
62%	0	26%	0	36%	190%	81%	0	32%	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
54%	0	22%	0	23%	150%	32%	0	25%	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
74%	0	24%	0	28%	170%	82%	0	27%	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
49%	0	25%	0	25%	180%	43%	0	25%	114%	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
57%	0	1	0	1	0	17	200%	443%	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
14%	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
23%	0	53%	0	30%	165%	40%	15%	101%	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

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

Explanation:

Question No: 27

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: A

Explanation:

Question No: 28

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

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

- A. Network I/O Control (NIOC)
- B. Virtual machine ID (VMID)
- C. Storage I/O Control (SIOC)
- D. vSphere Cluster Services (vCLS)

Answer: D

Explanation:

Question No: 29

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: C

Explanation:

Question No: 30

After users start reporting the inability to create new files in a CIFS share, you find EMS events for `wafldir.size.max` logged for the volume of the SVM to which the share points.

In this scenario, which action should you take to solve this issue?

A. Delete unneeded files from the directory.

B. Move the volume to a different aggregate.

C. Increase the maximum number of files for the volume.

D. Increase the maximum directory size for the volume.

Answer: C

Explanation:

Question No: 31

You have connection issues from a Linux host to an AFF A400 FC LUN. You are asked by support to run the NetApp `sanlun` command to provide troubleshooting information.

Which NetApp software package must be installed on the host?

A. Virtual Storage Console

B. Inventory Collect Tool

C. Host Utilities package

D. SnapDrive for UNIX

Answer: C

Explanation:

Question No: 32

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: A

Explanation:

Question No: 33

B. Jitter

C. concurrency

D. latency

Answer: B, D

Explanation:

Question No: 37

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

```
NOTICE token.node.out.of.quorum: All token references from node
(ID - 6e274936-ad68-11e6-9864-9daa7f8a8bb1) are dropped because the node went out of quorum.
DEBUG ctran.acn.received: CTRAN has received an ACN with value 0x00000002.
DEBUG qmm.ssg.grp.update
DEBUG qmm.acn.event.recv: hostname="cm2220-cn-02",
nvram_id="0x0", clam_id="1001", attribs="0x8"
DEBUG qmm.acn.event.recv: hostname="cm2220-cn-01",
nvram_id="0x0", clam_id="1000", attribs="0xa"
INFORMATIONAL clam.update.partner.state: CLAM on node (ID=1000)
updated failover state of partner (ID=1001) to nohb.
cm2220-cn-01 DEBUG clam.quorum.epoch: CLAM set the quorum epoch
on local node to 386.
INFORMATIONAL clam.received.quorum: Local node received a quorum update from Cluster node
(name=cm2220-cn-01, ID=1000).
EMERGENCY clam.node.oog: Node (name=cm2220-cn-02, ID=1001) is out of "CLAM quorum"
(reason=heartbeat failure).
EMERGENCY callhome.clam.node.oog: Call home for NODE(S) OUT OF CLUSTER QUORUM.
DEBUG clam.nodes.avail.change: The availability status of node
(name=cm2220-cn-02, ID=1001) changed from Available to Unavailable.
```

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: C

Explanation:

Question No: 38

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

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

A. -mail-hosts

B. -transport

C. -proxy-url

D. -to

Answer: A, C

Explanation:

Question No: 39

You want to assign storage access to an NVMe protocol host.

Which three objects must be configured to accomplish this task? (Choose three.)

- A. the NVMe subsystem
- B. the NVMe broadcast domain
- C. an NVMe LIF
- D. the NVMe namespace
- E. an NVMe LUN

Answer: A, B, C

Explanation:

Question No: 40

You receive the "Unable to connect to the management gateway server" error when trying to connect to a node management IP.

In this situation, how do you determine whether core dumps are generated for the mgwd user space process?

- A. Log in to the nodeshell using the node console and collect the core dumps from the nodeshell.
- B. Log in to the service processor and collect the core dumps from the service processor.
- C. Log in to the clustershell using the node console and collect the core dumps from the clustershell.
- D. Log in to the systemshell and collect core dumps from the systemshell.

Answer: B

Explanation:

Question No: 41

You are troubleshooting a CIFS connection issue that is reported by some users. You decide to collect a packet trace.

In this scenario, after you generate the packet trace, where do you find the trace file?

- A. /packet_traces of the node hosting the LIF
- B. /etc/log/packet_traces of the node hosting the LIF
- C. /etc/log/mlog/packet._traces Of all nodes
- D. /vol0/paclcet_traces of the CIFS SVM

Answer: C

Explanation:

Question No: 42

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

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

- A. There was a failed PSU in the switchover site.

- B. The ISLs were down between the sites.
- C. The site had unmirrored aggregates.
- D. An ISL link was not operating at 16Gbps.

Answer: A, C

Explanation:

Question No: 43

Cluster1::> qos statistics volume latency show -vserver svml -volume voll										
Workload	ID	Latency	Network	Cluster	Data	Disk	QoS Max	QoS		
Min	NVRAM									
-total-	-	-	-	-	-	-	-	-	-	-
4.71ms	185.00us	7.00us	469.00us	4.04ms	0ms	0ms	11.00us	0ms	1.00us	
voll	32140	22.22ms	86.00us	125.00us	504.00us	21.11ms	0ms	0ms	0ms	
-total-	-	-	-	-	-	-	-	-	-	-
4.58ms	253.00us	7.00us	991.00us	3.32ms	0ms	0ms	11.00us	0ms	1.00us	
voll	32140	29.69ms	90.00us	162.00us	2.61ms	26.82ms	0ms	0ms	0ms	
-total-	-	-	-	-	-	-	-	-	-	-
2.55ms	229.00us	8.00us	589.00us	1.71ms	0ms	0ms	13.00us	0ms	1.00us	
voll	32140	18.84ms	61.00us	314.00us	136.00us	18.32ms	0ms	0ms	0ms	
-total-	-	-	-	-	-	-	-	-	-	-
4.33ms	198.00us	9.00us	1221.00us	2.88ms	0ms	0ms	13.00us	0ms	2.00us	
voll	32140	177.16ms	113.00us	226.00us	108.03ms	68.79ms	0ms	0ms	0ms	

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

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

Answer: B

Explanation:

Question No: 44

Which two automation methods does NetApp ONTAP Select support? (Choose two.)

- A. REST
- B. Ansible
- C. Docker
- D. PHP

Answer: A, B

Explanation:

Question No: 45

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

iSCSI: network interface identifier disabled for use; incoming connection discarded In this scenario, what is the cause of this connection issue?

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

Answer: B

Explanation:

Question No: 46

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

Explanation:

Question No: 47

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

In this scenario, which source should be reviewed?

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

Answer: A

Explanation:

Question No: 48

You receive the panic message shown in the exhibit.

```
Uncorrectable Machine Check Error at CPU0. MC0 Error: STATUS<0xb200000430000800>
{Val,UnCor,Enable,PCC,ErrCode(Src,NTQ,Gen,Mem,I0)}. MCS Error: STATUS<0xf2000010c4300e0f>
{Val,OverF,UnCor,Enable,PCC,ErrCode(Gen,NTQ,Gen,Gen,Gen)}: Uncorrectable error at DIMM-1,
Channel 0, Serial: BA-00-1131-00098398169002460-I01-NTA-T1?!, FERR(0x400), NERR(0x402), MERR
M10Err, Rank 3, Bank 6, CAS 0x1e8, RAS 0x1bcf Uncorrectable error at DIMM-1, Channel 0, Serial:
BA-00-1131-00098398169002460-I01-NTA-T1?!, MERR M10Err, Rank 3, Bank 6, CAS 0x1e8, RAS 0x1bc.
```

In this scenario, which component should you troubleshoot first?

- A. the PCI card in slot 3

- B. the MetroCluster FC-VI card in slot 6
- C. the memory module in slot 1
- D. the CPU

Answer: D

Explanation:

Question No: 49

You are trying to deploy a Connector in the AWS cloud from NetApp Cloud Manager. The deployment fails and shows the message 'Insufficient permissions to deploy Cloud Connector'. You have verified the AWS access key and the AWS secret key.

In this scenario, what is the reason that the deployment failed?

- A. No AWS Marketplace subscription is associated with Cloud Manager.
- B. The required Identity and Access Management (IAM) policies were not installed.
- C. The user lacks the permission to deploy within Cloud Manager.
- D. The Connector can be deployed only in AWS GovCloud (US).

Answer: C

Explanation:

Question No: 50

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

```
===== SAS-EXPANDER-MAP =====
Expanders on channel 0a:
Level 1: WWN 500a098005f0ed3f, ID 0, Serial Number 'SHJHU1527000148', Product
'DS424ICM6E', Rev '0172', Slot A
Level 2: WWN 500a098005b053bf, ID 1, Serial Number 'SHFHU1516001218', Product
'DS224ICM6', Rev '0191', Slot A
Level 3: WWN 500a09800772a3bf, ID 2, Serial Number 'SHJHU1653000119', Product
'DS424ICM6', Rev '0191', Slot A
Level 4: WWN 500a098005af77bf, ID 1, Serial Number 'SHFHU1516001218', Product
'DS224ICM6', Rev '0191', Slot B
Level 5: WWN 500a09800773b8bf, ID 2, Serial Number 'SHJHU1653000119', Product
'DS424ICM6', Rev '0191', Slot B

Expanders on channel 0b:
Level 1: WWN 500a09800773b8bf, ID 2, Serial Number 'SHJHU1653000119', Product
'DS424ICM6', Rev '0191', Slot B
Level 2: WWN 500a098005af77bf, ID 1, Serial Number 'SHFHU1516001218', Product
'DS224ICM6', Rev '0191', Slot B
Level 3: WWN 500a09800772a3bf, ID 2, Serial Number 'SHJHU1653000119', Product
'DS424ICM6', Rev '0191', Slot A
Level 4: WWN 500a098005b053bf, ID 1, Serial Number 'SHFHU1516001218', Product
'DS224ICM6', Rev '0191', Slot A
Level 5: WWN 500a098005f0ed3f, ID 0, Serial Number 'SHJHU1527000148', Product
'DS424ICM6E', Rev '0172', Slot A
```

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: C

Explanation:

Question No: 51

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 aggregate is completely full.
- B. The constituent volumes are out of Inodes.
- C. The data placement is uneven.
- D. The constituent volumes are completely full.

Answer: A

Explanation:

Question No: 52

After a motherboard replacement on a NetApp AFF A300 in a SAN environment, the customer states that ports 0e and 0f are unable to connect to the fabric. The ports report "offline".

What would you examine first to troubleshoot the issue?

- A. vserver fcp wwpn-alias show command output
- B. system node hardware unified-connect show command output
- C. storage port show command output
- D. vserver fcp interface show command output

Answer: C

Explanation:

Question No: 53

A customer reports that some of their client computers are unable to access a FlexGroup over CIFS.

This problem involves only their older Windows XP clients and some printers and scanners.

In this scenario, which statement is correct?

- A. FlexGroup requires CIFS shares to be located in qtrees, which these clients do not support.
- B. FlexGroup requires CIFS shares to use the Continuously Available Feature, which these clients do not support.
- C. These systems use the SMQI protocol which is an unsupported feature in FlexGroup
- D. These systems use Offloaded Data Transfer, which is an unsupported feature in FlexGroup.

Answer: B

Explanation:

Question No: 54

Your customer has deployed a two petabyte NetApp ONTAP FlexGroup volume across their 4-node ONIAP 9.8 cluster. They plan to store over three billion files. They want to prevent file 10 conflicts since files are placed into the FlexGroup volume.

In this scenario, which two NFS SVM parameters should be enabled? (Choose two.)

- A. -v3-64bit-identifiers
- B. -v4-fsid-change
- C. -v3-fsid-change
- D. -v4-64bit-identifiers

Answer: A, C

Explanation:

Question No: 55

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.

```
cli1:~> network interface show -vserver CIFS
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
CIFS	nfs1	up/up	10.92.4.202/16	clln1	e2a-1092	
	iscsi1	up/up	10.91.3.132/24	clln1	a0a-322	true
	iscsi2	up/up	10.95.7.114/24	clln1	a0a-456	true
	mgmt	up/up	10.92.6.200/16	clln1	e2a-1092	true
	nfs2	up/up	192.168.3.9/24	clln1	e2a-100	true


```
cli1:~> 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: C

Explanation:

Question No: 56

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

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

- A. Newer Snapshot copies than the common snapshot that was chosen for resync were removed automatically at the destination.
- B. The autodelete option removed Snapshot copies before the resync.
- C. Snapshot copies were manually removed from the destination after the resync.
- D. NetApp ON TAP System Manager removed these snapshot copies due to a snapshot schedule.

Answer: C

Explanation:

Question No: 57

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

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

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

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

What is the cause of the customer's problem?

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

Answer: C

Explanation:

Question No: 58

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. A Snapshot copy preserved two versions of the file.
- C. Windows Volume Shadow Copy Service stored an older version of the file.
- D. Case Insensitivity of SMB clients caused the file to be displayed with a different name.

Answer: B

Explanation:

Question No: 59

Your customer is running a NetApp AFF A800 system with NetApp ONTAP 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.SHFHUXXXXXXXXXX." What is the cause of the reported Issue?

- A. The shelf has been power cycled.
- B. A new shelf was added to the controller.
- C. The shelf had a disk replaced too quickly.
- D. The shelf was cabled Incorrectly.

Answer: B

Explanation:

Question No: 60

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

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

- A. Enable the -scan-mandatory flag in the Vscan policy.
- B. Update the Vscan servers one at a time.
- C. Modify the CIFS shares to be continuously available (CA) shares.
- D. Disable the -scan-mandatory flag in the Vscan policy.

Answer: A, C

Explanation: