Free Questions for NS0-593 by dumpshq

Shared by Burns on 22-07-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Recently, a CIFS SVM was deployed and is working. The customer wants to use the Dynamic DNS (DDNS) capability available in NetApp ONTAP to easily advertise both data UFs to their clients. Currently. DNS is only responding with one data LIF. DDNS is enabled on the domain controllers.

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


cluster1::*> vserver services dns dynamic-update show
Vserver         Is-Enabled Use-Secure Vserver FQDN             TTL
--------------- ---------- ---------- ------------------------ -------
svm1            false      false      svm1.demo.net            24h
```

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

Options:

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

B- Remove the NFS protocol from the cifs_01 data LIF.

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

D- Enable the -is-enabled parameter for the SVM DDNS services

Answer:

B, D

Explanation:

To enable DDNS updates to work, two actions should be performed:

Remove the NFS protocol from the cifs_01 data LIF.This is because DDNS updates are only supported for LIFs that have only one data protocol enabled1. The cifs_01 LIF has both NFS and CIFS protocols enabled, which prevents it from registering its DNS record dynamically. By removing the NFS protocol from the cifs_01 LIF, it will become eligible for DDNS updates.

Enable the -is-enabled parameter for the SVM DDNS services.This is because the -is-enabled parameter controls whether the SVM sends DDNS updates to the DNS servers2. The exhibit shows that the -is-enabled parameter is set to false for the svm1 SVM, which means that it does not send any DDNS updates. By enabling the -is-enabled parameter, the SVM will start sending DDNS updates for its eligible LIFs.Reference:

1: Configure dynamic DNS services3

2: Manage DNS/DDNS services with System Manager4

Question 2

Question Type: MultipleChoice

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 loot aggregate.

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

Options:

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:

A B2B type consistency point (CP) occurs when a new CP is triggered before the previous CP is completed, due to the second memory buffer reaching a watermark. This can cause write latency to increase as user write operations are not replied until a write buffer frees up. However, this only affects the aggregate that is undergoing the B2B processing, and not the other aggregates on the same node. Therefore, the performance of the client operations on the data aggregates will not be affected by B2B processing on the root aggregate.Reference=What is the Back-to-Back (B2B) Consistency Point Scenario?,What are the different Consistency Point types and how are they measured in ONTAP 9?,What are the different Consistency Point types and how are they measured?

Question 3

Question Type: MultipleChoice

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

In this scenario, what is causing this problem?

Options:

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

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

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

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

Answer:

A

Explanation:

A speed mismatch occurs when the cluster ports on different nodes have different speeds, such as 10 Gbps and 40 Gbps1.

A speed mismatch can cause packet loss, retransmits, and latency on the cluster network, especially when data is accessed indirectly, meaning that the node serving the data is different from the node hosting the LIF2.

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

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

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

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

1: Cluster network cabling, ONTAP 9 Documentation Center

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

3: AFF A300 Tech Specs, NetApp

4: AFF A700 Tech Specs, NetApp

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

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

Question 4

Question Type: **MultipleChoice**

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

How would your customer prevent such problems in the future?

Options:

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

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

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

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

Answer:

A

Explanation:

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

https://docs.netapp.com/us-en/ontap-systems-upgrade/upgrade-arl-auto-app/resume_snapmirror_operations.html

Question 5

Question Type: MultipleChoice

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?

Options:

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:

D

Explanation:

The mgwd user space process is responsible for managing the cluster configuration and providing the management interface for ONTAP. If this process crashes, it may generate a core dump file in the /mroot/etc/mgwd/core directory of the node. To access this directory, you need to log in to the systemshell, which is a Linux shell that runs on the node. The systemshell can be accessed by logging in to the nodeshell using the node console and then running the systemshell command. The nodeshell is a custom shell that

provides access to ONTAP commands and utilities. The service processor and the clustershell are not relevant for collecting core dumps from the mgwd process.Reference=ONTAP 9 Cluster Administration, page 29-30;ONTAP 9 Troubleshooting Guide, page 23-24.

Question 6

Question Type: MultipleChoice

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)

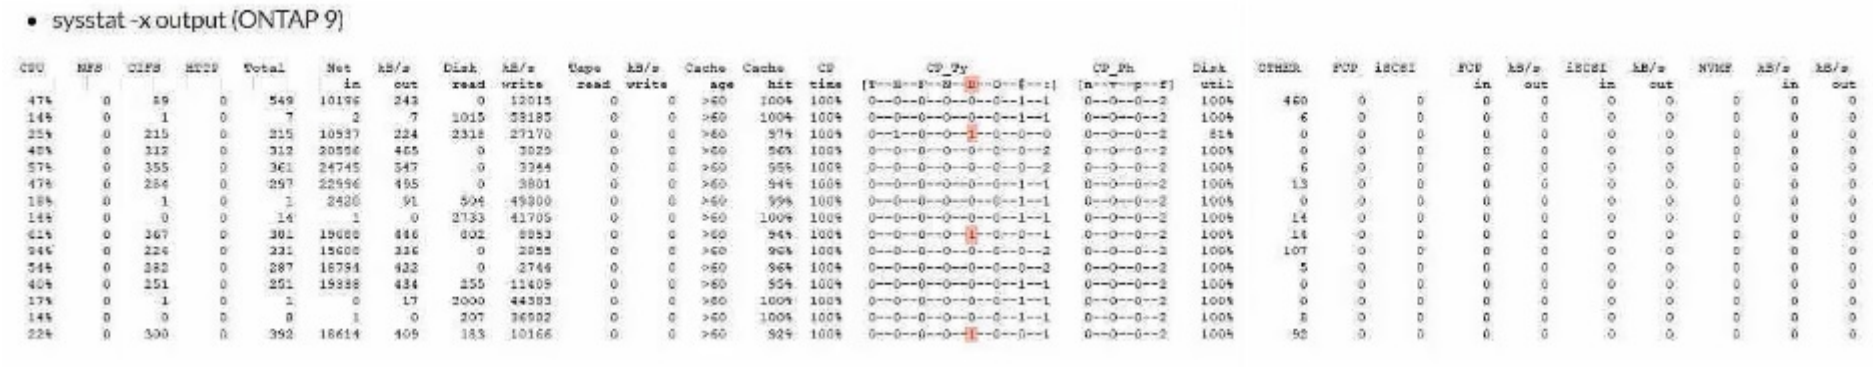

CPU	NFS	CIFS	HTTP	Total	Net in	Net out	Disk read	Disk write	Tape read	Tape write	Cache age	Cache hit	CP time	CP_Ty	CP_Ph	Disk util	OTHER	FCP	iSCSI	FCP in	FCP out	iSCSI in	iSCSI out	NVMF	NVMF in	NVMF out
47%	0	59	0	549	10196	243	0	12015	0	0	>60	100%	100%	0--0--0--0--0--0--1--1	0--0--0--2	100%	460	0	0	0	0	0	0	0	0	0
14%	0	1	0	7	2	7	1015	53185	0	0	>60	100%	100%	0--0--0--0--0--0--1--1	0--0--0--2	100%	6	0	0	0	0	0	0	0	0	0
25%	0	215	0	215	10537	224	2318	27170	0	0	>60	97%	100%	0--1--0--0--1--0--0--0	0--0--0--2	81%	0	0	0	0	0	0	0	0	0	0
40%	0	312	0	312	20556	465	0	3029	0	0	>60	96%	100%	0--0--0--0--0--0--0--2	0--0--0--2	100%	0	0	0	0	0	0	0	0	0	0
57%	0	355	0	361	24745	547	0	3264	0	0	>60	95%	100%	0--0--0--0--0--0--0--2	0--0--0--2	100%	6	0	0	0	0	0	0	0	0	0
47%	0	284	0	297	22996	495	0	3801	0	0	>60	94%	100%	0--0--0--0--0--0--1--1	0--0--0--2	100%	13	0	0	0	0	0	0	0	0	0
18%	0	1	0	1	2420	91	504	49300	0	0	>60	99%	100%	0--0--0--0--0--0--1--1	0--0--0--2	100%	0	0	0	0	0	0	0	0	0	0
14%	0	0	0	14	1	0	2733	41705	0	0	>60	100%	100%	0--0--0--0--0--0--1--1	0--0--0--2	100%	14	0	0	0	0	0	0	0	0	0
61%	0	367	0	381	19688	846	602	8853	0	0	>60	94%	100%	0--0--0--0--1--0--0--1	0--0--0--2	100%	14	0	0	0	0	0	0	0	0	0
54%	0	224	0	231	15600	336	0	2055	0	0	>60	90%	100%	0--0--0--0--0--0--0--2	0--0--0--2	100%	107	0	0	0	0	0	0	0	0	0
54%	0	282	0	287	18794	422	0	2746	0	0	>60	96%	100%	0--0--0--0--0--0--0--2	0--0--0--2	100%	5	0	0	0	0	0	0	0	0	0
40%	0	251	0	251	19388	434	255	11405	0	0	>60	95%	100%	0--0--0--0--0--0--1--1	0--0--0--2	100%	0	0	0	0	0	0	0	0	0	0
17%	0	1	0	1	0	17	2000	44383	0	0	>60	100%	100%	0--0--0--0--0--0--1--1	0--0--0--2	100%	0	0	0	0	0	0	0	0	0	0
14%	0	0	0	8	1	0	207	36992	0	0	>60	100%	100%	0--0--0--0--0--0--1--1	0--0--0--2	100%	8	0	0	0	0	0	0	0	0	0
22%	0	300	0	392	18614	409	183	10166	0	0	>60	92%	100%	0--0--0--0--1--0--0--1	0--0--0--2	100%	92	0	0	0	0	0	0	0	0	0

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

Options:

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:

A type B consistency point (CP) is triggered when the NVRAM buffer is full and needs to be flushed to disk.A type B CP indicates that the write workload is higher than the disk throughput, and the system is experiencing back pressure1

A switchless cluster is a cluster configuration that does not use external switches for cluster interconnect and management network.A switchless cluster has lower bandwidth and redundancy than a switched cluster, and is limited to two nodes2

The exhibit shows the output of the sysstat -x command, which displays the system performance statistics in extended mode. The output shows that the system has high CPU utilization, high disk utilization, high NVRAM utilization, and several type B CPs.These are signs of performance issues and resource contention3

The best corrective action to address these consistency points is to add an additional shelf of twenty-four 1.8 TB disks.This will increase the disk capacity and throughput, and reduce the disk utilization and the frequency of type B CPs4

Creating additional data LIFs will not address the consistency points, because the data LIFs are used for data access protocols, not for NVRAM flushing5

Converting the 2-node switchless cluster to a 2-node switched cluster will not address the consistency points, because the cluster interconnect and management network are not related to the disk performance6

Replacing the twenty-four 1.8 TB disks with twelve 4 TB disks will not address the consistency points, because the disk throughput will not increase, and the disk utilization will remain high7

1: Where can I learn more about Consistency Points?- NetApp Knowledge Base2: ONTAP 9 - Cluster and SVM Peering Express Guide - The Open Group3: ONTAP 9 - Commands: Manual Page Reference - The Open Group4: ONTAP 9 - Hardware Universe - The Open Group5: ONTAP 9 - Network Management Guide - The Open Group6: ONTAP 9 - Clustered Data ONTAP Concepts Guide - The Open Group7: ONTAP 9 - Logical Storage Management Guide - The Open Group

Question 7

Question Type: **MultipleChoice**

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?

Options:

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:

A

Explanation:

A Twinax cable is a type of copper cable that is used to connect cluster ports to cluster switches1.

A cluster port is a network port that is configured for cluster communication and data access2.

A cluster switch is a network switch that is used to interconnect the nodes in a cluster and provide redundancy and load balancing3.

A cluster network is a network that enables cluster communication and data access between the nodes in a cluster and external clients4.

A cluster network can be degraded due to various reasons, such as misconfiguration, malfunction, or excessive link errors on the cluster ports or the cluster switches.

Link errors are errors that occur on the physical layer of the network, such as CRC errors, length errors, alignment errors, or dropped packets.

Link errors can indicate a problem with the cable, the switch port, the network interface card (NIC), or the cable connector.

In this scenario, the alert messages state that the cluster network has degraded and the errors are being generated from Node 2, interface e0b.

The first step to solve this problem is to replace the Twinax cable between Node 2, interface e0b and the NetApp CN1610 switch, as this could be the source of the link errors.

Replacing the cable could resolve the issue and restore the cluster network to a healthy state.

If replacing the cable does not solve the problem, then other steps may be required, such as checking the switch port, the NIC, or the cable connector, or replacing the switch or the motherboard.Reference:

1: Cluster network cabling, ONTAP 9 Documentation Center

2: Cluster ports, ONTAP 9 Documentation Center

3: Cluster switches, ONTAP 9 Documentation Center

4: Cluster network, ONTAP 9 Documentation Center

[5]: How to troubleshoot CLUSTER NETWORK DEGRADED error messages, NetApp Knowledge Base

[6]: Cluster network degraded due to high CRC errors on cluster ports, NetApp Knowledge Base

Question 8

Question Type: MultipleChoice

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?

Options:

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

B- Configure VIP LIFs using OSPF.

C- Configure VIP LIFs using BGP.

D- Configure a I IF failover policy for each subnet inside a single broadcast domain.

Answer:

C

Explanation:

A LIF (Logical Interface) is a logical entity that represents a network connection point on a node1.

A VIP LIF (Virtual IP LIF) is a LIF that can fail over across subnets within an IPspace2.

BGP (Border Gateway Protocol) is a routing protocol that enables VIP LIFs to advertise their IP addresses to external routers and to update the routing tables when a failover occurs3.

To connect a NetApp ONTAP cluster to a complex network that requires LIFs to fail over across subnets, you need to configure VIP LIFs using BGP on the cluster and on the external routers3.

This way, you can ensure that the network traffic is routed to the optimal node and port for each VIP LIF, and that the network connectivity is maintained in the event of a node or port failure3.Reference:

1: Logical Interfaces, ONTAP 9 Documentation Center

2: VIP LIFs, ONTAP 9 Documentation Center

3: Configuring BGP on a cluster, ONTAP 9 Documentation Center

Question 9

Question Type: MultipleChoice

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

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

Options:

A- Enable the -scan-mandatory flag in the Vscan policy.

B- Update the Vscan servers one at a time.

C- Modify the CIFS shares to be continuously available (CA) shares.

D- Disable the -scan-mandatory flag in the Vscan policy.

Answer:

B, D

Explanation:

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

Vscan servers or the scanning process.Reference=Vscan server installation and configuration,Update the Vscan On-Access policy configuration for an SVM

Question 10

Question Type: MultipleChoice

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

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

Options:

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:

If both root aggregates are destroyed in a two-node cluster, the cluster will be inoperable and the data will be inaccessible. To recover from this situation, you need to perform the following actions:

Install ONTAP from a USB device on one of the nodes. This will create a new root aggregate and a new cluster on that node.

Rejoin the second node to the re-created cluster. This will also create a new root aggregate on the second node and synchronize it with the first node.

Restore the cluster configuration and data from a backup, if available.Reference=

ONTAP 9 Documentation Center

Storage System Recovery Troubleshooting

Recovering from a root aggregate failure