



Free Questions for NS0-593 by go4braindumps

Shared by Hahn on 24-05-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

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?

Options:

- A- S3 license
- B- VE license
- C- TPM license
- D- FabricPool license

Answer:

D

Explanation:

FabricPool is an ONTAP feature that enables tiering of cold data from SSD aggregates to low-cost object storage, either on-premises or in the cloud¹. FabricPool requires a license to be installed on the cluster, and the license type depends on the cloud tier being used². In this scenario, the cloud tier is another ONTAP cluster (FAS8300), which is not supported by the new Cloud Tiering license that is used for most FabricPool configurations³. Therefore, the old FabricPool license that is retained for dark sites or MetroCluster systems using FabricPool Mirror must be used³. The FabricPool license defines the amount of capacity that can be tiered to the cloud tier, and it can be increased by add-on orders⁴. Reference:

1: FabricPool overview⁵

2: FabricPool requirements⁶

3: Install a FabricPool license²

4: ONTAP FabricPool (FP) Licensing Overview¹

Question 2

Question Type: MultipleChoice

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+	ANY5+	AVG	CPU0	CPU1	CPU2	CPU3	CPU4	CPU5	CPU6	CPU7													
Nwk_Excl	Nwk_Lg	Nwk_Exmpt	Protocol	Storage	Raid	Raid_Ex	Xor_Ex	Target	Kahuna	WAFI_Ex(Kahu)																			
WAPL_MPClean	SM	Exempt	Exempt	SSAN_Ex	Intr	Host	Ops/s	CP																					
99%	87%	83%	55%	52%	74%	66%	46%	52%	62%	54%	54%	53%	53%	54%	53%	1%	11%	293%	0%	4%	1%	15%	1%	0%	3%				
(82%)		2%	0%	37%	0%	13%	26%	40246	25%																				
99%	89%	86%	71%	55%	75%	65%	46%	54%	64%	56%	56%	56%	55%	56%	56%	1%	1%	329%	0%	4%	0%	14%	1%	0%	2%				
(75%)		1%	0%	37%	0%	14%	7%	44963	25%																				
100%	89%	86%	71%	56%	50%	72%	44%	56%	53%	57%	57%	57%	56%	56%	57%	1%	0%	307%	0%	5%	0%	19%	3%	0%	2%				
(76%)		7%	0%	41%	0%	13%	5%	39012	44%																				
99%	89%	85%	70%	55%	75%	71%	52%	55%	67%	56%	56%	57%	56%	56%	56%	1%	1%	255%	0%	5%	0%	16%	2%	0%	2%				
(77%)		0%	0%	39%	0%	13%	22%	40034	37%																				
99%	89%	86%	71%	56%	50%	74%	47%	57%	50%	55%	55%	55%	55%	55%	55%	1%	0%	256%	0%	6%	0%	21%	4%	0%	2%				
(77%)		6%	0%	43%	0%	13%	5%	39573	70%																				
99%	85%	84%	59%	53%	75%	66%	55%	53%	65%	54%	54%	54%	54%	54%	54%	1%	3%	301%	0%	6%	1%	15%	0%	0%	2%				
(76%)		3%	0%	37%	0%	14%	11%	41991	10%																				
99%	89%	86%	72%	55%	54%	79%	54%	59%	55%	59%	59%	60%	59%	59%	60%	1%	1%	254%	0%	6%	0%	22%	4%	0%	3%				
(77%)		4%	0%	44%	0%	12%	11%	35750	99%																				
99%	89%	85%	70%	54%	75%	70%	53%	55%	69%	56%	56%	56%	55%	55%	56%	1%	1%	253%	0%	6%	0%	19%	1%	0%	3%				
(76%)		1%	0%	43%	0%	13%	5%	35792	30%																				
99%	85%	84%	59%	53%	77%	70%	52%	54%	67%	55%	56%	55%	55%	56%	55%	1%	1%	275%	0%	7%	0%	22%	2%	0%	3%				
(77%)		3%	0%	47%	0%	14%	7%	36943	47%																				

Referring to the exhibit, which statement is correct?

Options:

- A- The customer should be advised to exclude certain workflows to reduce use.
- B- High network exempt use could be a problem.
- C- You should immediately investigate further by gathering perfstat data and opening a support case.
- D- The CPU Is not a first-order monitoring metric for ONTAP.

Answer:

D

Explanation:

= CPU utilization in ONTAP is not a linear measure of the system load, nor can it be used alone as a measure of the overall system utilization. ONTAP uses a Coarse Symmetric Multiprocessing (CSMP) design which partitions system functions into logical processing domains, each with its own scheduling rules and resource availability. Therefore, a high CPU utilization does not necessarily indicate a performance problem, unless it is accompanied by other contributing factors such as high latency, low throughput, or high queue depth. ONTAP has several mechanisms to optimize CPU usage and balance the workload across the cores, such as WAFL parallelization, exempt processing, and CPU pinning. The CPU utilization reported by the sysstat command is an average across all cores and domains, and does not reflect the actual CPU activity or availability for each domain. Therefore, the CPU is not a first-order monitoring metric for ONTAP, and other metrics such as latency, throughput, and queue depth should be considered first. Reference=What is CPU utilization in Data ONTAP: Scheduling and Monitoring?,How to measure CPU utilization,What are CPU as a compute resource and the CPU domains in ONTAP 9?,Monitoring CPU utilization before ONTAP upgrade

Question 3

Question Type: MultipleChoice

Your customer complains that u 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_iscsip_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_iscsip_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_iscsip_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_iscsip_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_iscsip_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_iscsip_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_iscsip_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?

Options:

- 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:

A

Explanation:

[IQN stands for iSCSI Qualified Name, which is a unique identifier for an iSCSI initiator or target1.](#)

[ONTAP uses IQN to authenticate and authorize iSCSI sessions2.](#)

[If two hosts have the same IQN, they will cause a conflict and ONTAP will reject the new session request from the second host3.](#)

This will result in the host losing the connection to the iSCSI target and then reporting that the session was reestablished, as shown in the exhibit.

To avoid this problem, each host should have a unique IQN.Reference:

[iSCSI Qualified Name \(IQN\) - NetApp](#)

[iSCSI authentication and authorization - NetApp](#)

[Troubleshooting iSCSI issues - NetApp](#)

[\[Configuring iSCSI initiators - NetApp\]](#)

Question 4

Question Type: MultipleChoice

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

Options:

- 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:

A, C

Explanation:

Having an Essentials package enables you to convert to another licensing option. This is true because the Essentials package is a capacity-based licensing option that allows you to pay for Cloud Volumes ONTAP per TiB of capacity. You can switch to another capacity-based package or to a Keystone Flex Subscription at any time¹.

BYOL licenses are purchased directly from NetApp. This is true because BYOL stands for Bring Your Own License, which means you need to obtain a license from NetApp before deploying Cloud Volumes ONTAP systems in any cloud provider².

Capacity-based license packages support only single-node configurations. This is false because capacity-based license packages support both single-node and HA configurations¹.

AWS Marketplace contracts cannot be mixed with BYOL. This is false because you can mix AWS Marketplace contracts with BYOL as long as the license type matches the instance type³. Reference: ¹: Cloud Volumes ONTAP licensing | NetApp Documentation ²: Set up

Question 5

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 LIF 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 6

Question Type: MultipleChoice

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?

Options:

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

Answer:

B

Explanation:

SnapMirror is a data replication technology that allows efficient and flexible data protection and disaster recovery for NetApp ONTAP storage systems¹

SnapMirror transfers data between source and destination volumes using a network connection. SnapMirror can use storage efficiency features such as compression and deduplication to reduce the amount of data transferred and stored¹

SnapMirror transfers are scheduled and controlled by policies that define the frequency, retention, and priority of the transfers. SnapMirror policies can also specify the network bandwidth limit for the transfers²

SnapMirror transfers are considered background tasks that run in the absence of user workload. SnapMirror transfers can consume CPU and disk resources on both source and destination systems, depending on the amount and type of data being replicated³

SnapMirror transfers can throttle up or down depending on the availability of system resources and network bandwidth. SnapMirror transfers will throttle up when there is no user workload, and throttle down when there is user workload. This is to ensure that SnapMirror transfers do not impact the performance of user operations³

Therefore, if a customer is concerned with high CPU and disk utilization on their SnapMirror destination system, the best answer is to explain that background tasks such as SnapMirror throttle up in the absence of user workload. This is normal and expected behavior, and it does not indicate a problem with the system³

1: [ONTAP 9 Data Protection - SnapMirror - The Open Group](#)2: [ONTAP 9 Data Protection - SnapMirror Policies - The Open Group](#)3: [SnapMirror storage efficiency configurations and behavior - Resolution Guide - NetApp Knowledge Base](#)

Question 7

Question Type: MultipleChoice

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      4879  Node04 secondary
Node03    vldb      0        11      358   -      offline
Node03    vifmgr    11      11      4892  Node04 secondary
Node03    bcomd     11      11      62    Node04 secondary
Node03    crs       11      11      6     Node04 secondary
Node04    mgmt      11      11      4879  Node04 master
Node04    vldb      0        11      358   -      offline
Node04    vifmgr    11      11      4892  Node04 master
Node04    bcomd     11      11      62    Node04 master
Node04    crs       11      11      6     Node04 master
10 entries were displayed.

ClusterB:~> system node run -node Node04 -command aggr status -r aggr2
Aggregate aggr2 (online, raid_dp, degraded) (block checksums)
Plex /aggr2/plex0 (online, normal, active, pool0)
RAID group /aggr2/plex0/rg0 (degraded, block checksums)

RAID Disk Device          HA  SHELF BAY CHAN Pool Type  RPM  Used (MB/blks)  Phys
(MB/blks)
-----
----
  dparity  FAILED
  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?

Options:

- 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:

D

Explanation:

The broken disk in Node03 is causing the cluster ring to be offline, which prevents the CIFS SVM from being accessible. The cluster ring is a distributed database that stores cluster configuration information and enables communication between cluster nodes. If the cluster ring is offline, the cluster cannot function properly and the CIFS SVM cannot serve data to clients. The other options are not relevant to the CIFS SVM connectivity issue. Reference=<https://www.netapp.com/support-and-training/netapp-learning-services/certifications/support-engineer/>

<https://mysupport.netapp.com/site/docs-and-kb>

To Get Premium Files for NS0-593 Visit

<https://www.p2pexams.com/products/ns0-593>

For More Free Questions Visit

<https://www.p2pexams.com/netapp/pdf/ns0-593>

