PassQueen
a.k.a dumpsqueen.com

NetApp Certified Support Engineer ONTAP Specialist

Netapp NS0-593

Version Demo

Total Demo Questions: 9

Total Premium Questions: 95

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, ONTAP Software	80
Topic 2, ONTAP Hardware	15
Total	95

QUESTION NO: 1

Which two statements about NetApp Cloud Volumes ONTAP licenses ate 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: A C

Explanation:

Having an Essentials package enables you to convert to another licensing option. The Essentials capacity-based package is designed for flexibility: it supports different Cloud Volumes ONTAP deployment configurations and lets customers change to another capacity-based licensing package when their capacity, performance, or feature requirements change. This enables licensing to be aligned with changing workload needs without rebuilding the working environment.

BYOL licenses are purchased directly from NetApp. Bring Your Own License licensing is obtained from NetApp or an authorized NetApp sales channel rather than as an hourly pay-as-you-go charge through the cloud provider marketplace. After purchase, the license is associated with a serial number and provides licensed Cloud Volumes ONTAP capacity for use through BlueXP. NetApp licensing documentation describes BYOL as a customer-purchased license and distinguishes it from cloud-marketplace pay-as-you-go licensing. See the NetApp Cloud Volumes ONTAP licensing overview and the NetApp documentation for managing capacity-based licenses.

QUESTION NO: 2

A customer reports that writes to an NFS volume have failed with an out-of-space error. The volume is configured with autosize enabled in grow mode, the aggregate has sufficient free space, and the current volume size is equal to the configured maximum autosize size.

In this scenario, what should you do to allow the volume to grow automatically?

A. Increase the maximum autosize size for the volume.

B. Disable Snapshot copy reserve on the volume.

C. Move the volume to a different aggregate.

D. Set the volume guarantee to volume.

ANSWER: A

Explanation:

Increase the maximum autosize size for the volume. In grow mode, ONTAP automatically increases a volume only until it reaches the value configured by the `-maximum-size` parameter. When the current volume size equals that limit, autosize cannot grow the volume further, even if the hosting aggregate has available capacity.

After increasing the maximum size to an appropriate value, ONTAP can resume automatic growth when the volume reaches its configured growth threshold. The administrator should also confirm that aggregate free space remains sufficient for the new maximum volume size. See the volume autosize modify command reference and the ONTAP automatic volume size management documentation.

QUESTION NO: 3

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

What is the cause of the customer's problem?

A. The cable is plugged in upside down.

B. The incorrect cable is used.

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

D. The port has an incorrect SFP inserted.

ANSWER: C

Explanation:

The ports are configured for 4x10Gbit instead of 40Gbit. A 10/40GbE adapter port can operate either as one native 40GbE interface or as four breakout 10GbE interfaces, depending on its port configuration. For a direct 40GbE connection to the switch, the port on each controller must be configured and presented as a 40GbE port. The AutoSupport comparison indicates that node 2 has the expected 40GbE port personality, while node 1 is instead operating in the four-lane 10GbE breakout mode. In that mode, the physical interface and lane layout do not match the single 40GbE switch connection, so the link cannot establish even though the corresponding connection on node 2 works. Support should confirm the administrative and operational characteristics of the affected network port, then configure node 1 for the required 40GbE mode and verify the compatible switch-side configuration. NetApp documents port attribute administration through the network port modify command reference and recommends reviewing current port properties with the network port show command reference.

QUESTION NO: 4

```
MCC-A::> metrocluster show

Configuration: IP-fabric

Cluster                   Entry Name          State
------------------------- ------------------- ---------------------
Local: MCC-A
                            Configuration State  configured
                                   Mode  switchover
                            AUSO Failure Domain  auso-disabled
Remote: MCC-B
                            Configuration State  configured
                                   Mode  waiting-for-switchback
                            AUSO Failure Domain  auso-disabled
```

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

A. Enter the 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:

Enter the `metrocluster switchback` command on Site A. The exhibit indicates that Site A is the cluster that performed the MetroCluster switchover and is currently serving the affected workloads. After the partner site has been repaired, brought online, and its aggregates have been healed, a MetroCluster switchback returns aggregate ownership and normal workload operation to the recovered site. This completes the MetroCluster recovery process and restores the standard configuration in

which each site owns and serves its usual aggregates. The command must be issued from the surviving cluster that initiated and currently maintains the switchover state—in this case, Site A. Before switchback, administrators should confirm that MetroCluster healing has completed successfully and that the configuration is ready for switchback, because switchback is the operation that transfers service back to the recovered disaster-site cluster. NetApp documents the command syntax and its role in returning a MetroCluster configuration from switchover to normal operation in the metrocluster switchback command reference. For the broader recovery sequence, including healing and readiness checks before switchback, see NetApp's MetroCluster healing phase documentation.

QUESTION NO: 5

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.

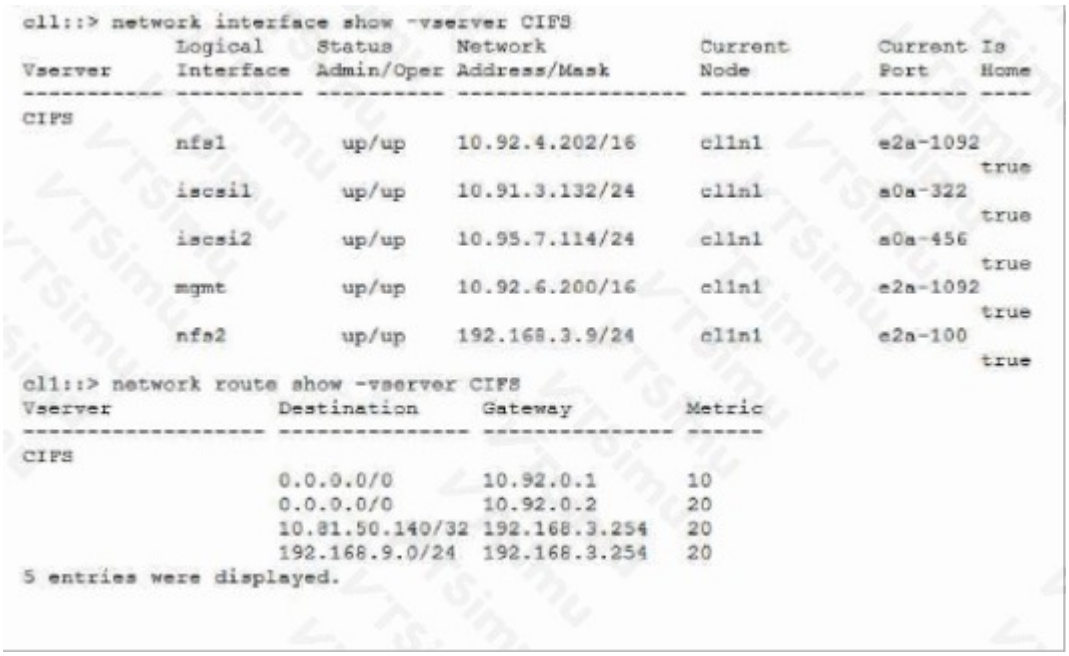

```
cl1::> network interface show -vserver CIFS
            Logical    Status     Network            Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
----------- ---------- ---------- ------------------ ------------- ------- ----
CIFS
         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

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

A. a0a-322

B. e2a-100

C. e2a-1092

D. a0a-456

ANSWER: C

Explanation:

e2a-1092 is the required interface because the exhibit associates the 192.168.9.0/24 client network with VLAN ID 1092. In ONTAP, a VLAN interface is represented by the parent physical port followed by a hyphen and the VLAN ID. Therefore, e2a-1092 is the logical VLAN port carrying traffic that is tagged for VLAN 1092 on physical port e2a. A packet trace must be started on the port or VLAN interface through which the traffic of interest enters or leaves the node. Tracing on e2a-1092 captures frames for the relevant client VLAN, allowing the administrator to inspect communication to and from clients in the specified subnet without mixing it with traffic on other VLANs carried by the same physical link. This is especially important on a trunk port, where multiple tagged networks can share one Ethernet interface. NetApp documents packet tracing as a node-and-port operation, and its VLAN documentation describes the parent-port/VLAN-ID naming convention used by interfaces such as e2a-1092. See the ONTAP network packet trace start command reference and ONTAP VLAN configuration documentation.

QUESTION NO: 6

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 LI

F.
DDNS is enabled on the domain controllers.

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

A. 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:

Dynamic DNS must first be enabled for the SVM, so enabling the `-is-enabled` parameter for the SVM DDNS services is required. This permits ONTAP to send DNS update requests when relevant SVM network information changes. For a CIFS workload in which clients need to locate individual data LIFs, disabling the `-vserver-fqdn` parameter is also appropriate. When this setting is disabled, ONTAP registers the individual LIF DNS records instead of limiting registration to the SVM fully qualified domain name. This allows DNS to return records for the data LIFs and supports client access to the available CIFS data interfaces. The DDNS configuration is applied at the SVM level with the `vserver services name-service dns dynamic-update` command family, which includes both settings. NetApp documents the procedure and the behavior of SVM dynamic DNS updates in its Configure dynamic DNS services guide. The command parameters are also documented in the ONTAP CLI reference for dynamic-update modify.

QUESTION NO: 7

Your customer has mounted an NFS SVM from a Linux client unci 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 svm1_cluster1
             Logical    Status     Network            Current       Current Is
Vserver      Interface  Admin/Oper Address/Mask       Node          Port    Home
-----------  ---------- ---------- ------------------ ------------- ------- ----
svm1_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                                       Update
Name    Domain Name    MTU  Port List                   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 svm1_cluster1 -destination 192.168.0.210 -
disallow-fragmentation false -packet-size 1000
```

B)

```
network ping -lif jumbo_01 -vserver svm1_cluster1 -destination 192.168.0.210 -
disallow-fragmentation true -packet-size 9000
```

C)

```
network ping -lif jumbo_01 -vserver svm1_cluster1 -destination 192.168.0.210 -
disallow-fragmentation true -packet-size 1000
```

D)

```
network ping -lif jumbo_01 -vserver svm1_cluster1 -destination 192.168.0.210 -
disallow-fragmentation true -packet-size 5000
```

A. Option A

B. `network ping -node -destination -df-bit true -packet-size 8972`

C. Option C

D. Option D

ANSWER: B

Explanation:

The appropriate command is `network ping -node <node_name> -destination <client_IP> -df-bit true -packet-size 8972`. A configured MTU of 9000 at the client and in the ONTAP broadcast domain does not prove that every physical link, switch port, VLAN, or intermediate network device supports jumbo Ethernet frames. The path must be tested end-to-end.

Setting `-df-bit true` requests that the probe packet not be fragmented. A packet size of 8972 bytes, together with the IPv4 and ICMP headers, produces a 9000-byte IP packet and is therefore appropriate for validating a jumbo-frame path configured for MTU 9000. If the ping succeeds from the relevant ONTAP node to the Linux client, it provides evidence that the route can carry the expected jumbo packet without fragmentation. If it fails while smaller probes work, the result helps isolate an MTU inconsistency or a device in the path that is not configured for jumbo frames.

Use the node that hosts, or can reach through the same network path as, the NFS data LIF so the test reflects the path used by the workload. NetApp documents the `network ping` parameters, including packet size and the do-not-fragment setting, in the ONTAP 9.8 network ping command reference. Additional ONTAP network troubleshooting guidance is available in the ONTAP network management documentation.

QUESTION NO: 8

An administrator receives the following error message:

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

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

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

B. A disk is falling.

C. There is excessive SATA HDD load.

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

ANSWER: A C

Explanation:

Excessive SSD load causing wear leveling to become unbalanced and excessive SATA HDD load can each produce the disk-performance condition shown in the error. SSDs perform internal wear-leveling activities to distribute program/erase cycles across flash media. Under sustained write pressure, these internal activities can add latency to host I/O, especially when the device has limited available flash blocks or is processing a heavy write workload. This can make disk responses slower and trigger an ONTAP alert associated with delayed disk I/O. SATA HDDs can similarly become overloaded when their outstanding I/O queue is saturated. Because spinning disks must service seeks and rotational latency, a heavily utilized SATA drive can take longer to complete commands; the delay can be reported by ONTAP as a disk-related performance event. Administrators should correlate the alert with aggregate, disk, and workload utilization, then reduce contention or redistribute workloads where appropriate. NetApp's disk and aggregate management documentation and performance monitoring documentation provide the operational context for investigating disk latency and workload pressure.

QUESTION NO: 9

You are configuring Fibre Channel access for a Linux host to a SAN SVM.

Which three objects must be configured to allow the host to access a LUN? (Choose three.)

A. an FC LIF

B. an initiator group

C. a LUN mapping

D. a Snapshot policy

E. an export policy

ANSWER: A B C

Explanation:

An FC LIF, an initiator group, and a LUN mapping are required for host access. The FC LIF provides a Fibre Channel endpoint on the SVM through which the host establishes fabric connectivity. The initiator group contains the host initiator WWPNs and defines the host access-control object. The LUN must then be mapped to that igroup so ONTAP can present it to the authorized initiators.

A Snapshot policy does not provide SAN connectivity or access authorization, and an export policy applies to NAS protocols rather than FC LUN access. NetApp documents these requirements in the ONTAP SAN provisioning documentation and the documentation for creating and mapping LUNs.