PassQueen
a.k.a dumpsqueen.com

NetApp Certified Data Administrator (ONTAP)

Netapp NS0-162

Version Demo

Total Demo Questions: 13

Total Premium Questions: 136

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Storage Administration	43
Topic 2, Data Protection	21
Topic 3, Security	20
Topic 4, Networking	29
Topic 5, High Availability	7
Topic 6, Performance	12
Topic 7, Troubleshooting	4
Total	136

QUESTION NO: 1

Your company has purchased two NetApp AFF A400 HA pairs. During deployment, how many root aggregates are created?

A. 6

B. 4

C. 2

D. 8

ANSWER: B

Explanation:

4 is correct. An AFF A400 HA pair contains two controller nodes, and ONTAP creates one root aggregate for each node during initial system configuration. Therefore, two HA pairs provide four nodes in total, resulting in four root aggregates. The root aggregate is node-specific: it contains the node's root volume and stores essential configuration and system information needed for ONTAP operation. It is not a shared aggregate for the HA pair, even though the nodes are configured as partners for takeover and giveback. Each controller maintains its own root aggregate so that it can continue to operate and retain its system configuration appropriately within the cluster architecture. This calculation is based on node count rather than the number of chassis or HA pairs: two pairs multiplied by two nodes per pair equals four nodes, and one root aggregate is created per node. NetApp documentation describes the root aggregate as the aggregate that contains a node's root volume and distinguishes it from data aggregates used to host user volumes. See the ONTAP disk and aggregate management documentation and the ONTAP aggregate and physical storage overview.

QUESTION NO: 2

Click the Exhibit button.

```
cluster1::> volume show -vserver svm2_cluster1 -volume data -fields space-guarantee,fractional-
reserve,autosize-mode,space-slo
vserver       volume  autosize-mode  space-slo   space-guarantee   fractional-reserve
------------  ------  -------------  ---------  ----------------  -------------------
svm2_cluster1 data    grow           none        none               0%

cluster1::> lun show -vserver svm2_cluster1 -volume data -fields space-reserve,space-allocation
vserver         path           space-reserve  space-allocation
-----------     -----------    -------------  ----------------
svm2_cluster1  /vol/data/lun  disabled       disabled
```

You have a 2-node NetApp FAS2750 ONTAP cluster. You create a new 20-GB LUN in a new 100-GB volume and write 10 GB of data to the LUN. No storage efficiencies are enabled for the volume or aggregate.

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

A. ONTAP reports that the volume is using 10 GB of its containing aggregate

B. ONTAP reports the volume as 20% full

C. ONTAP reports that the volume is using 100 GB of its containing aggregate

D. ONTAP reports the volume as 10% full

ANSWER: A B

Explanation:

ONTAP reports that the volume is using 10 GB of its containing aggregate because, with no storage efficiency features involved, the physical data written to the LUN is the relevant physical consumption in the aggregate. Only 10 GB of blocks

has been written, so the aggregate-side usage attributable to that data is 10 GB. This distinguishes physical consumption from the LUN's provisioned capacity.

ONTAP reports the volume as 20% full because the 20-GB LUN is space-reserved within the 100-GB volume. Space reservation ensures that sufficient volume capacity is set aside to satisfy the LUN's advertised size, independent of how much application data has currently been written. Therefore, the volume's reported fullness reflects 20 GB reserved out of 100 GB total capacity. The written-data amount and the LUN-reserved amount are separate capacity perspectives: physical blocks used in the aggregate versus capacity committed within the volume for SAN access. NetApp describes LUN space reservation and its capacity-guarantee behavior in its SAN space allocation documentation. Additional context on how ONTAP reports and manages volume space is available in the ONTAP volumes documentation.

QUESTION NO: 3

Exhibit

```
cluster::> volume show -vserver SVM_A
Vserver   Volume       Aggregate    State      Type       Size  Available Used%
--------- ------------ ------------ ---------- ---- ---------- ---------- -----
SVM_A     SVM_A_root   aggr1        online     RW          1GB    970.7MB    0%
SVM_A     VOL_A        aggr1        online     RW        100GB     95.00GB    0%
2 entries were displayed.

cluster::> qos adaptive-policy-group show -policy-group apg-VOL-A

                    Name: apg-VOL-A
                 Vserver: SVM_A
                    Uuid: 402a3ae5-6ba3-11eb-ba3a-00a0988887c0
                      ID: 13773
          Expected IOPS: 1000IOPS/TB
               Peak IOPS: 2000IOPS/TB
   Absolute Minimum IOPS: 500IOPS
Expected IOPS Allocation: allocated-space
   Peak IOPS Allocation: used-space
              Block Size: ANY
    Number of Workloads: 1
```

Referring to the exhibit, what is the maximum possible IOPS for the VOL_A volume?

A. 1000 IOPS

B. 500 IOPS

C. 2000 IOPS

D. 100 IOPS

ANSWER: B

Explanation:

The maximum possible performance for VOL_A is **500 IOPS**. The exhibit indicates that VOL_A is governed by a Quality of Service policy that imposes a 500-IOPS maximum throughput limit. In ONTAP, a QoS maximum limit is an upper bound on the I/O operations that workloads assigned to the policy group can receive. Consequently, even if the aggregate, disks, controller, and network have additional available performance capacity, ONTAP throttles the volume workload as necessary to keep it at or below that configured ceiling. The applicable maximum is the limit associated with the volume's assigned QoS policy, rather than an estimate of available system performance or a minimum-performance value. QoS policy groups are used specifically to control workload resource consumption and provide predictable performance behavior for other workloads on the cluster. NetApp documents QoS maximum throughput as a policy setting that can be expressed in IOPS, with the policy applied to storage objects such as volumes. See the ONTAP documentation on QoS throughput limits and the qos policy-group create command reference.

QUESTION NO: 4

You need to configure data-at-rest encryption for your NetApp ONTAP 9.8 cluster. Your company does not have Key Management Interoperability Protocol (KMIP) services available but must require a passphrase to be entered when a node Is rebooted.

In this scenario, which two actions should be performed to satisfy these requirements? (Choose two.)

A. Configure an external key management server.

B. Enable onboard key management.

C. Enable cluster-wide FIPS-compliant mode.

D. Enable common criteria mode.

ANSWER: B D

Explanation:

Enable onboard key management because it provides an internal, cluster-based key management solution when an external Key Management Interoperability Protocol server is not available. The onboard key manager stores the keys used to protect NetApp Volume Encryption data within the ONTAP cluster, allowing encrypted volumes to be configured without deploying external key-management infrastructure. During onboarding, an administrator establishes the onboard key-management passphrase and securely records the recovery information.

Enable common criteria mode because this mode adds the required startup protection behavior for the onboard key manager. With Common Criteria mode enabled, ONTAP requires the onboard key-management passphrase to be supplied following node reboot before encrypted data can be made available. This meets the requirement that a passphrase be entered during reboot while retaining local, rather than KMIP-based, key management. The configuration therefore combines the cluster's onboard key repository with the Common Criteria security control that protects key access after restart. NetApp documents the onboard key-manager enable operation at security key-manager onboard enable and the Common Criteria configuration setting at security config modify.

QUESTION NO: 5

All SMB users are reporting that their network share performance is slow on a new NetApp FAS500f HA pair. You determine that the users have converged network adapters (CNAs) with flow control enabled and no option to disable.

In this scenario, which two flow control settings are recommended on the NetApp ONTAP physical ports? (Choose two.)

A. Disable flow control end-to-end on data network ports

B. Disable flow control end-to-end on cluster network ports

C. Enable flow control end-to-end on data network ports

D. Enable flow control end-to-end on cluster network ports

ANSWER: B C

Explanation:

Disable flow control end-to-end on cluster network ports and enable flow control end-to-end on data network ports are the recommended settings. ONTAP cluster interconnect traffic is highly sensitive to congestion and pause-frame behavior. NetApp recommends disabling Ethernet flow control throughout the cluster network path, including the ONTAP cluster ports and connected switch ports. This prevents pause-frame propagation from affecting the timely exchange of cluster-control traffic.

For the data network, the CNA behavior is the determining factor in this scenario. Because the SMB clients' adapters have flow control enabled and cannot be configured to disable it, flow control must be configured consistently across the complete data path. Enabling it end-to-end means that the ONTAP data port and every applicable switch port use compatible flow-control settings. This avoids a mismatched configuration and allows pause handling to operate consistently while serving SMB workload traffic.

NetApp documents the separate recommendations for cluster and data ports in its ONTAP networking guidance: Configure network ports. Additional Ethernet and switch-configuration guidance is available in the NetApp switch configuration overview.

QUESTION NO: 6

You want to prepare your ONTAP cluster and your ESXi cluster to connect NFS datastores over a 10-GbE network using jumbo frames.

In this scenario, which three configurations would accomplish this task? (Choose three.)

A. Enable jumbo frames with an MTU of 1500 for your ESXi hosts

B. Enable jumbo frames with an MTU of 9000 for your ONTAP cluster

C. Enable jumbo frames with an MTU of 1500 for your ONTAP cluster

D. Enable jumbo frames with an MTU of 9000 for your ESXi hosts

E. Enable jumbo frames with an MTU of 9216 for your switches

ANSWER: B D E

Explanation:

Jumbo-frame support for NFS datastore connectivity must be configured consistently across the complete network path: the ESXi host VMkernel networking used for NFS, the ONTAP IP network hosting the NFS data LIFs, and every intervening Ethernet switch port. An MTU of 9000 on the ESXi hosts permits the VMkernel interfaces carrying NFS traffic to transmit and receive jumbo Ethernet frames. An MTU of 9000 on the ONTAP cluster's applicable broadcast domains, ports, and data LIF network path similarly enables ONTAP to use the same jumbo-frame payload size. The switches must support a frame size larger than the IP-interface MTU because Ethernet headers and encapsulation add bytes to the frame. Therefore, configuring a switch MTU of 9216 accommodates traffic generated by endpoints configured with a 9000-byte MTU. Consistency is essential: all devices and links in the NFS path must accept the required frame size, including any VLAN or port-channel path used by the storage network. NetApp documents setting the appropriate MTU for the broadcast domain and its member ports, while VMware documents configuring the MTU consistently on the physical network and VMkernel adapters for jumbo-frame use. See NetApp ONTAP broadcast-domain and MTU configuration and VMware ESXi jumbo-frame configuration.

QUESTION NO: 7

Exhibit.

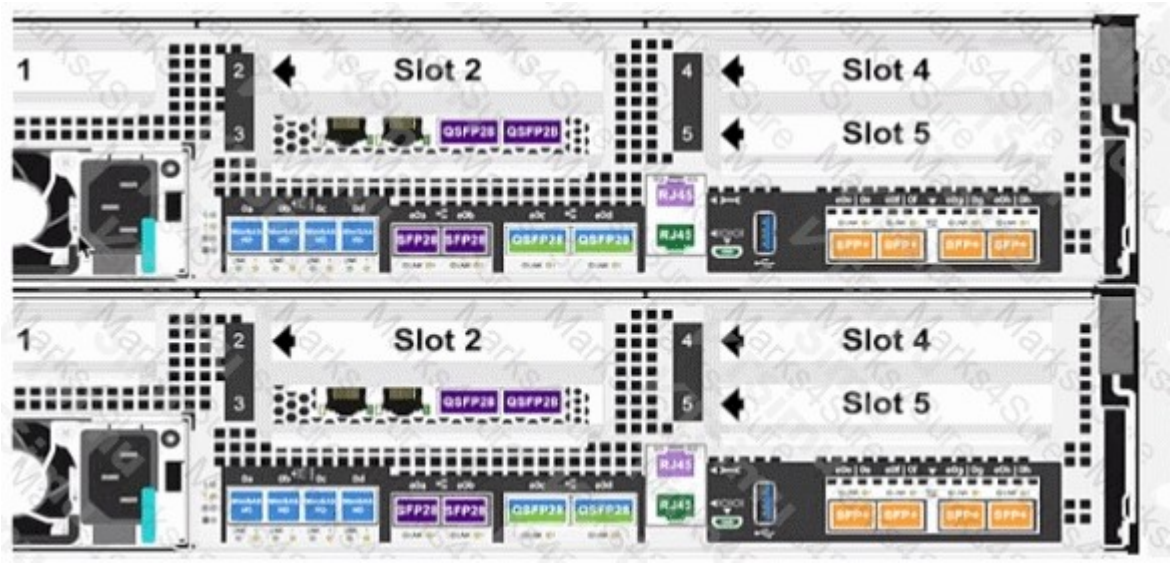

You are adding a NetApp AFF A400 system to an existing 4-node cluster and want to know which ports you would use to expand your cluster.

Referring to the exhibit, which tool would you use to accomplish this task?

A. Active 1Q Unified Manager

B. NetApp Interoperability Matrix (IMT)

C. Active IQ Conflg Advisor

D. NetApp Hardware Universe (HWU)

ANSWER: D

Explanation:

NetApp Hardware Universe (HWU) is the appropriate tool because it provides platform-specific hardware documentation, including controller port layouts, port labels, supported protocols, and the designated roles for onboard ports. When adding an AFF A400 HA pair to an existing cluster, cluster connectivity must use ports supported and designated for the cluster network on that platform. Hardware Universe lets an administrator select the AFF A400 model and review the detailed port information needed to identify those cluster expansion ports accurately. It also provides cabling and platform reference material that helps ensure the new nodes are connected according to supported NetApp configurations. This is especially important during a cluster expansion because the physical interfaces, port speeds, and supported connection types must align with the existing cluster switch infrastructure and the AFF platform's documented requirements. Review the AFF A400 system information in NetApp Hardware Universe and use NetApp's AFF A400 installation and setup documentation to validate the physical cluster-network cabling before bringing the new nodes into the cluster.

QUESTION NO: 8

Exhibit.

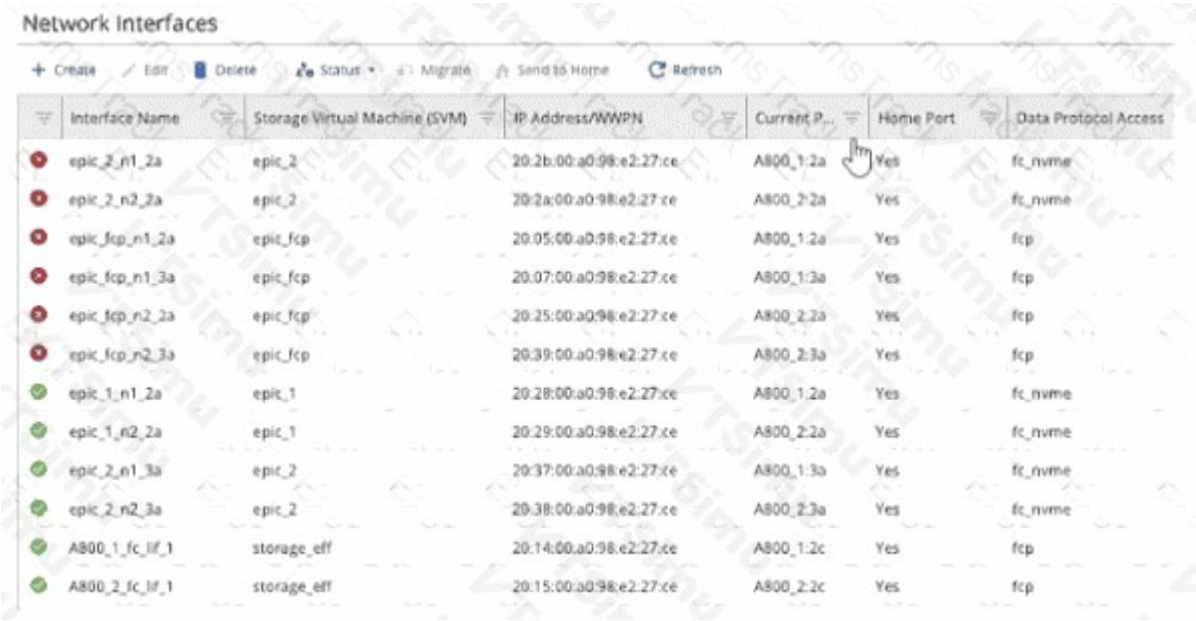

Referring to the exhibit, what would cause a customer to have an Issue with SVM eplc_fcp when mapping a namespace to a host?

A. FC has not been licensed on the cluster.

B. NVMe Is not configured for the SVM.

C. FC switch zoning has not been done correctly.

D. NVMe is not licensed on the cluster.

ANSWER: B

Explanation:

NVMe Is not configured for the SVM. is the correct answer because an NVMe namespace can be presented to a host only through the ONTAP NVMe protocol configuration for the SVM. In ONTAP, namespace access is not established by directly mapping a namespace to an FC host in the same way that a LUN is mapped with an igroup. Instead, the administrator enables and configures NVMe on the SVM, creates an NVMe subsystem, adds the host NQN to that subsystem, and maps the namespace to the subsystem. The host then discovers the namespace through a supported NVMe transport, such as NVMe/FC or NVMe/TCP.

An SVM named `eplc_fcp` may indicate that it was created or configured for traditional FCP SAN access. That configuration alone does not enable NVMe namespace provisioning. Before namespace mapping can succeed, the SVM must have NVMe configured and the required NVMe subsystem and host NQN association must be available. NetApp documents that NVMe must be configured on the SVM before NVMe namespaces can be used, and that namespaces are mapped to NVMe subsystems for host access. See Configure an SVM for NVMe and Map an NVMe namespace to a subsystem.

QUESTION NO: 9

You are preparing an NFS client to access a volume in an SVM.

Which two ONTAP configurations are required to provide NFS access to the volume? (Choose two.)

A. An NFS data LIF

B. An export policy assigned to the volume

C. An iSCSI igroup mapped to the volume

D. An FC target port configured on every node

ANSWER: A B

Explanation:

An NFS data LIF is required because NFS clients need a network endpoint through which they can communicate with the SVM. The LIF must be configured on a reachable network and provide the NFS data service.

An export policy assigned to the volume is also required because ONTAP uses export policies and their rules to determine which NFS clients can access the volume and whether they have read-only or read-write permissions. The volume must also be mounted in the SVM namespace for client access through its junction path. See the ONTAP NFS configuration documentation.

QUESTION NO: 10

Exhibit.

```
Cluster1::> ifgrp show
         Port       Distribution                 Active
Node     IfGrp      Function     MAC Address     Ports   Ports
-------- ---------- ------------ --------------- ------- -------------------
Node1
         a0a        ip           02:a0:98:f1:2e:be partial e0c, e0d
Node2
         a1a        ip           02:a0:98:f1:21:88 full    e0c, e0d
2 entries were displayed.

Cluster1::> network port show -node Node1 -port e0c,e0d -fields mac,health-status
node                 port mac               health-status
-------------------- ---- ----------------- -------------
Node1   e0c    00:a0:98:f1:2e:b2   degraded
Node1   e0d    00:a0:98:f1:2e:b3   healthy
```

You have received an alert from Active IQ Unified Manager that interface eOc has gone offline. Network engineering asks for the corresponding MAC address for the port that has a problem.

Referring to the exhibit, which MAC address would you supply to the engineering team?

A. 02:a0:98:fl:2e:be

B. 00:a0:98:fl:2e:b2

C. 00:a0:98:fl:2e:b3

D. 02:a0:98:fl:21:88

ANSWER: C

Explanation:

The MAC address to provide is **00:a0:98:fl:2e:b3**. The exhibit associates each ONTAP network port name with its hardware MAC address. The affected interface is identified as e0c, so the required value is the MAC address listed on the same e0c port entry. This is the physical, globally administered address for that specific Ethernet port and is the value network engineering needs to locate the connected switch interface, validate CAM-table learning, and investigate a link-down condition.

ONTAP exposes per-port network information, including administrative and operational state, speed, MTU, and MAC address. Administrators can retrieve and verify this information from the cluster CLI with commands such as `network port show`, optionally filtering for the node and port involved. Matching the alert's port identifier exactly to the port record is important because MAC addresses are assigned per physical network port; a nearby address from the same node could identify a different cable and switch connection. NetApp documents how to display port-level network details in the ONTAP network port information documentation. For the relationship between ONTAP physical ports and network configuration, see the ONTAP physical ports documentation.

QUESTION NO: 11

You need to make a point-in-time copy of a volume available to clients for testing without consuming capacity for a full copy of the volume.

Which ONTAP feature satisfies this requirement?

A. FlexClone

B. FlexCache

C. SnapMirror

D. SnapLock

ANSWER: A

Explanation:

FlexClone is correct because a FlexClone volume is a writable, space-efficient clone of a parent volume or Snapshot copy. Initially, the clone shares the parent volume blocks, so creating it does not require a complete physical copy of the source data. Additional capacity is consumed only as blocks diverge because of writes to the clone or parent.

FlexClone volumes are useful for test and development, analytics, application validation, and other workflows that require a writable copy of production data. See the ONTAP documentation for creating FlexClone volumes.

QUESTION NO: 12

You must configure external key management for NetApp Volume Encryption (NVE) by using Key Management Interoperability Protocol (KMIP) servers.

Which two actions are required in this scenario? (Choose two.)

A. Install the certificate authority certificate for the KMIP server.

B. Install a client certificate for the ONTAP cluster.

C. Create an iSCSI initiator group for the KMIP server.

D. Configure a SnapMirror relationship to the KMIP server.

ANSWER: A B

Explanation:

You must **install the certificate authority certificate for the KMIP server** so that ONTAP can validate the identity of the external key-management server. This establishes trust for the TLS connection used to protect KMIP communication.

You must also **install a client certificate for the ONTAP cluster**. The client certificate allows the KMIP server to authenticate the cluster when ONTAP requests or stores encryption keys. After the required certificates are installed, the external key manager can be enabled and associated with the KMIP server addresses. See the NetApp documentation for external key management.

QUESTION NO: 13

Click the Exhibit button.

```
cluster::> network interface show -vserver SVM_A -fields data-protocol
vserver      lif                  data-protocol  address
------------ -------------------- -------------- ---------------
SVM_A        lif-SVM-A            nfs,cifs       10.0.1.250

cluster::> volume show -fields vserver,volume,aggregate,policy,junction-path -vserver SVM_A
vserver  volume     aggregate  policy          junction-path
-------  --------   ---------  -------------   -------------
SVM_A    SVM_A_root  aggr1     policy-vol_a    /
SVM_A    VOL_A       aggr2     policy-vol_a    -
2 entries were displayed.

cluster::> export-policy rule show -vserver SVM_A -policyname policy-vol_a
             Policy          Rule    Access      Client              RO
Vserver      Name            Index   Protocol    Match               Rule
------------ --------------  ------  ----------- ------------------  --------
SVM_A        policy-vol_a    1        any        10.0.1.0/24         any
```

A Linux host with the 10.0.1.24 IP address is unable to mount the VOL_A volume using NFSv3.

Referring to the exhibit, what is the problem?

A. The volume is in read-only

B. The logical interface does not allow the NFSv3 protocol

C. The host IP address does not match the export policy rule

D. The volume is not mounted in the namespace

ANSWER: C

Explanation:

The host IP address does not match the export policy rule is correct. In ONTAP, NFS access to a volume is controlled by the export policy assigned to that volume or its effective junction path. Each export-policy rule includes a client match value that identifies the clients permitted to access the export, such as an individual IP address, a subnet expressed in CIDR notation, a hostname, or a netgroup. For the Linux client at 10.0.1.24 to mount VOL_A with NFSv3, its source address must match the client specification in an applicable export-policy rule and that rule must permit the requested access. The exhibit's rule does not include an address range or client definition that matches 10.0.1.24, so ONTAP denies the mount request at export-policy evaluation. Correcting the client match—for example, by adding the host address or using the appropriate subnet CIDR range—allows the client to be evaluated under the intended NFS access settings. NetApp documents export policies

as the mechanism used to control client access to volumes and explains the client-match criteria used by export rules. See Create an export policy and Add a rule to an export policy.