



**Free Questions for 1Z0-499 by certscare**

**Shared by May on 06-06-2022**

**For More Free Questions and Preparation Resources**

**Check the Links on Last Page**

# Question 1

---

**Question Type:** MultipleChoice

---

What happens to synchronous writes after they land in DRAM?

## Options:

---

- A- They are acknowledged as complete to the client immediately.
- B- They are acknowledged only after they are persistently stored on disk or flash.
- C- They are acknowledged only after they are persistently stored in DRAM on the opposite cluster node.
- D- They are acknowledged only after they are persistently stored on disk or flash from the opposite cluster node.

## Answer:

---

B

## Explanation:

---

Synchronous writes to the Oracle ZFS Storage Appliance are not acknowledged immediately upon landing in DRAM. Instead, they are acknowledged only once they are persistently stored on disk or flash.

References:

<http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/o14-001-architecture-overview-zfsa-2099942.pdf>

## Question 2

---

**Question Type:** MultipleChoice

---

Which three RAID levels are not recommended when configuring a storage pool for OLTP (random) workload and why?

**Options:**

---

- A-** RAIDZ3, due to bad performance because of the lowest IOPS among RAID levels
- B-** RAID1, due to great availability but high cost
- C-** RAIDZ2, due to not very good performance because of low IOPS
- D-** RAID0, due to good performance but no redundancy

**Answer:**

---

A, C, D

## Explanation:

---

A: Triple parity RAID, wide stripes. RAID in which each stripe has three disks for parity. This is the highest capacity option apart from Striped Data. Resilvering data after one or more drive failures can take significantly longer due to the wide stripes and low random I/O performance.

C: Double parity RAID is a higher capacity option than the mirroring options and is intended either for high-throughput sequential-access workloads (such as backup) or for storing large amounts of data with low random-read component.

D: RAID0 does not provide redundancy.

References:

[https://www.doag.org/formes/pubfiles/2262661/docs/Konferenz/2010/vortraege/Infrastruktur%20\(inkl.%20SUN\)/390-2010-K-INF-Vogel-S7000\\_Einsatz-Praesentation.pdf](https://www.doag.org/formes/pubfiles/2262661/docs/Konferenz/2010/vortraege/Infrastruktur%20(inkl.%20SUN)/390-2010-K-INF-Vogel-S7000_Einsatz-Praesentation.pdf) p. 30

[https://docs.oracle.com/cd/E51475\\_01/html/E52872/goden.html](https://docs.oracle.com/cd/E51475_01/html/E52872/goden.html)

## Question 3

---

**Question Type:** MultipleChoice

---

Which hardware failure cannot be prevented by the No Single Point of Failure feature on a ZFS Storage Appliance?

### **Options:**

---

- A- disk backplane failure
- B- JBOC failure
- C- controller failure
- D- power supply failure

### **Answer:**

---

B

### **Explanation:**

---

Adding a single JBOD to a double parity RAID-Z NSPF config makes it impossible to preserve NSPF characteristics. However, you can still add the JBOD and create RAID stripes within the JBOD, sacrificing NSPF in the process.

References:

[https://docs.oracle.com/cd/E51475\\_01/html/E52872/godge.html](https://docs.oracle.com/cd/E51475_01/html/E52872/godge.html)

## Question 4

---

**Question Type:** MultipleChoice

---

Which three Oracle solutions are supported by ZFS Storage Appliances?

### Options:

---

- A- InfiniBand (IB) connectivity for Exadata Backup
- B- storage expansion of Oracle Private Cloud Appliance
- C- Quality of Service Plus (QoS Plus) for auto-tiering
- D- Oracle Intelligent Storage Protocol (OISP) for Oracle Database

### Answer:

---

A, B, D

### Explanation:

---

When equipped with native QDR InfiniBand and 10Gb Ethernet connectivity options, the ZFS Storage Appliance is ideal for reliably backing up Oracle Exadata.

The Oracle Private Cloud Appliance includes compute nodes, management nodes, virtual networking, and integrated Oracle ZFS Storage ZS3-ES for internal storage.

The new OISP allows the Oracle Database NFSv4 client to pass ODM optimization information to the NFSv4 server of the ZFS Storage Appliance. The ZFS Storage Appliance takes advantage of the ODM optimization information to simplify database configuration and to further increase database performance.

References:

Oracle ZFS Storage Appliance Administration Guide (June 2014), pages 432, 463

<http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/privatecloudapp-config-2972083.pdf>

## Question 5

---

**Question Type:** MultipleChoice

---

What happens to a shadow migration job if a pool is failed over in a cluster, or both system disks fail and a new head node is required while using the shadow migration feature?

**Options:**

---

- A-** All data necessary to continue shadow migration without interruption is stored in the write-optimized cache.
- B-** All data necessary to continue shadow migration without interruption is removed and the job starts from the beginning.
- C-** All data necessary to continue shadow migration without interruption is kept within the storage pool.
- D-** All data necessary to continue shadow migration without interruption is stored in NVRAM.

**Answer:**

---

C

**Explanation:**

---

Explaation:

Shadow migration is implemented using on-disk data within the filesystem, so there is no external database and no data stored locally outside the storage pool. If a pool is failed over in a cluster, or both system disks fail and a new head node is required, all data necessary to continue shadow migration without interruption will be kept with the storage pool.

References:

[https://docs.oracle.com/cd/E51475\\_01/html/E52872/shares\\_\\_shadow\\_migration\\_\\_shadow\\_migration\\_.html](https://docs.oracle.com/cd/E51475_01/html/E52872/shares__shadow_migration__shadow_migration_.html)

## Question 6

---



**Question Type: MultipleChoice**

---

Which three actions can you perform to restore data from a snapshot?

**Options:**

---

- A-** Mount the snapshot LUN and copy the data back.
- B-** Clone the snapshot LUN and copy the data back.
- C-** Click the rollback icon of a snapshot.
- D-** Mount the share, change directories to the .zfs directory of the share, find the snapshot (time sequenced), and copy the data.
- E-** Back up and restore the data from tape.

**Answer:**

---

B, C, D

**Explanation:**

---

B: Cloning a Snapshot.

A clone is a writable copy of a snapshot, and is managed like any other share. Like snapshots of filesystems, it initially consumes no additional space. As the data in the clone changes, it will consume more space. The original snapshot cannot be destroyed without also

destroying the clone. Scheduled snapshots can be safely cloned, and scheduled snapshots with clones will be ignored if they otherwise should be destroyed.

C: Use the following procedure to roll back, or restore, a filesystem or LUN to an existing snapshot.

D: Filesystem snapshots can be accessed over data protocols at `.zfs/snapshot` in the root of the filesystem.

References:

[https://docs.oracle.com/cd/E71909\\_01/html/E71919/gprim.html](https://docs.oracle.com/cd/E71909_01/html/E71919/gprim.html)

## Question 7

---

**Question Type: MultipleChoice**

---

A storage administrator wants to bind the ZFS storage ZS3 to the existing LDAP server.

What is required in order for this to happen?

**Options:**

---

**A-** an LDAP proxy server

- B-** a self-signed certificate
- C-** a Base Search DN (distinguished name)
- D-** a password for an LDAP proxy server

**Answer:**

---

C

## Question 8

---

**Question Type:** MultipleChoice

---

A customer orders a new ZFS Storage Racked System in a data center. Identify two factors that should be considered for site planning.

**Options:**

---

- A-** the total airflow requirements of all equipment installed in the rack to ensure that it operates within its specified temperature range
- B-** the total airflow requirements of the storage server installed in the rack to ensure that it operated within its specified temperature range
- C-** whether the temperature in the rack does not exceed the controller's maximum ambient rated temperature

**D-** whether the temperature in the data center does not exceed the controller's maximum ambient rated temperatures

**Answer:**

---

A, C

**Explanation:**

---

Caution - Reduced airflow: Installation of the equipment in a rack should be such that the amount of airflow required for safe operation of the equipment is not compromised.

Caution - Elevated operating ambient temperature: If the controller is installed in a closed or multi-unit rack assembly, the operating ambient temperature of the rack environment might be greater than room ambient temperature. Therefore, consideration should be given to installing the equipment in an environment compatible with the maximum ambient temperature (T<sub>ma</sub>) specified for the controller

References:

[https://docs.oracle.com/cd/E86025\\_01/pdf/E86027.pdf](https://docs.oracle.com/cd/E86025_01/pdf/E86027.pdf), page 31

## Question 9

---

**Question Type:** MultipleChoice

---

Identify three required tasks to configure a ZFS Storage Appliance cluster.

**Options:**

---

- A- Cable the cluster serial interconnect controllers.
- B- Cable each appliance's HBAs to the shared JBOD
- C- Perform the initial tty-based configuration on one node.
- D- Cable the cluster interconnect controllers.
- E- Perform the initial tty-based configuration on both nodes.
- F- Cable the primary appliance's HBAs to the shared JBOD
- G- Cable the cluster Ethernet interconnect controllers.

**Answer:**

---

B, C, D

**Explanation:**

---

Configuring Clustering

Etc.

References:

[https://docs.oracle.com/cd/E51475\\_01/html/E52872/godjv.html](https://docs.oracle.com/cd/E51475_01/html/E52872/godjv.html)

## Question 10

---

**Question Type: MultipleChoice**

---

Consider the following authorizations for a role called davidadmin, assigned to the david user:

```
prd7110:configuration roles davidadmin> authorization show:  
Auths:
```

NAME	OBJECT	
PERMISSIONS		
auth-000	stat.*	read
create		
auth-001	nas.pool-0.hank.*	
changeAccessProps		
changeGeneralProps		
changeProtocolProps		
changeSpaceProps		
clone		
createShare		
destroy		
rename		
rollback		
scheduleSnap		
takeSnap		

Which statement is true?

**Options:**

---

- A- David can define the protocols used to access LUNs or shares within the "david" project.
- B- David can reboot the appliance.
- C- David can define new users.
- D- David can create snaps and clone shares in the "hank" project.
- E- David can create shares in the "david" project.

**Answer:**

---

D



**To Get Premium Files for 1Z0-499 Visit**

**<https://www.p2pexams.com/products/1z0-499>**

**For More Free Questions Visit**

**<https://www.p2pexams.com/oracle/pdf/1z0-499>**

