

QUESTIONS & ANSWERS

Kill your exam at first Attempt



1Z0-499 Dumps
1Z0-499 Braindumps
1Z0-499 Real Questions
1Z0-499 Practice Test
1Z0-499 dumps free



Oracle

1Z0-499

Oracle ZFS Storage Appliance 2017 Implementation Essentials

<http://killexams.com/pass4sure/exam-detail/1Z0-499>



QUESTION: 62

Which type of I/O workload benefits the most from read SSDs?

- A. sequential write
- B. random write
- C. random read
- D. sequential read
- E. mixed random and sequential write

Answer: C

QUESTION: 63

You are performing a cluster failover of ZFS Storage Appliances and encounter an error that a cluster node fails to rejoin the cluster. Which two are possible causes of this problem?

- A. The cluster links cabling is not correctly configured.
- B. There are cluster-wide locking issues on the working node.
- C. The working node is configured more than one network interface card.
- D. The working node is assigned other resources available to the cluster.

Answer: B, C

Explanation:

Note: rejoin: to retrieve and resynchronize the resource map from the peer

QUESTION: 64

A storage administrator wants to maintain access to a controller regardless of its state by

locking its management interface to make it private. Identify two steps that you must perform to lock the management IP.

- A. In the BUI of the active controller, navigate to Configuration > Cluster. Choose the management interfaces for both controllers from the Resource list. Click the padlock icon to lock the management interface to this controller.
- B. In the BUI of the second controller, navigate to Configuration > Cluster. Choose the management interface from the Resource list. Click the padlock icon to lock the management interface to this controller.
- C. In the BUI of the second controller, navigate to Configuration > Cluster. Choose the management interface for first controller from the Resource list. Click the padlock icon to lock the management interface to this controller.
- D. In the BUI of the first controller, navigate to Configuration > Cluster. Choose the management interface from the Resource list. Click the padlock icon to lock the management interface to this controller.
- E. In the BUI of the first controller, navigate to Configuration > Cluster. Choose the management interfaces for the second controller from the Resource list. Click the padlock icon to lock the management interface to this controller.

Answer: C, D

Explanation:

Locking Cluster Management Interfaces (BUI)

After initial configuration, clustered controllers are in an active-active state. When a failover occurs, an active controller takes over all non-private interfaces, and the peer controller becomes passive and inaccessible by its BUI and CLI. To maintain access to a controller regardless of its state, lock its management interface to make it private. The following procedure locks the management interface on each clustered controller.

References:

https://docs.oracle.com/cd/E76475_01/html/E76483/gpkay.html

QUESTION: 65

Why would you purge Analytics data that is older than a certain chosen age?

- A. to delete the Analytics worksheet
- B. to specify the global dataset data retention policy
- C. to back up the entire dataset
- D. to delete the entire dataset
- E. to prune an individual dataset

Answer: B

Explanation:

Analytics Data Retention Policies, Default Retention Policy

By default, the appliance retains per-second data for 7 days, per-minute data for 14 days, and per-hour data for 90 days. However, it is strongly recommended that you specify a data retention policy that suits your business needs. Retention policies are especially important if you plan to retain large amounts of historical data for an extended period.

The maximum retention period is two years.

Enabling a Retention Policy

A retention policy limits the minimum amount of data collected by a granularity of per-second, per-minute, or per-hour over a period of time, or retention period. You could set one retention policy per granularity. For example, you could define a retention policy to save a minimum of one day of data at the per-second interval, a second policy to save a minimum of one week of data at the per-minute interval, and a third policy to save a minimum of one month of data at the per-hour interval. It is recommended that you only keep the minimum amount of data according to your business requirements, including compliance needs.

References:

https://docs.oracle.com/cd/E71909_01/html/E71920/goyit.html

QUESTION: 66

Identify two classes of resources within the Oracle ZFS Appliance cluster.

- A. network
- B. storage
- C. cache
- D. memory

Answer: A, B

QUESTION: 67

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?

- 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:

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_w_migration_.html
shadow_migration shado

QUESTION: 68

Which configuration provides maximum availability and performance?

- A. quad parity
- B. triple mirrored
- C. triple parity RAID
- D. triple striped
- E. quad parity with mirroring

Answer: B

Explanation:

Triple mirrored

Data is triply mirrored, yielding a very highly reliable and high-performing system (for example, storage for a critical database). This configuration is intended for situations in which maximum performance and availability are required. Compared with a two-way mirror, a three-way mirror adds additional IOPS per stored block and higher level protection against failures.

References:

https://docs.oracle.com/cd/E51475_01/html/E52872/goden.html

QUESTION: 69

A storage administrator wants to allow the dfsadmin SMB local user to manipulate Distributed File System (DFS) namespaces. Identify three steps required to achieve this.

- A. Create a local user account on the server for the dfsuser user.
- B. Log in as dfsadmin on the Windows machine from which the DFS namespace will be modified.
- C. Create a local user account on the server for the dfsadmin user. Be sure to use the same password as when the local user was first created on the Windows machine.
- D. Add dfsadmin to the “Administrators” local SMB group.
- E. Log in as dfuser on the Windows domain and modify the DFS namespace.

Answer: C, D, E

Explanation:

Manipulating DFS Namespaces

References:

Oracle ZFS Storage Appliance Administration Guide (June 2014), page 206

QUESTION: 70

Identify three factors that can improve performance when deduplication occurs inline on an Oracle ZFS Storage Appliance.

- A. more PCIe slots available in the ZFS storage system
- B. more network interfaces available in the ZFS storage system
- C. more processors available in the ZFS storage system
- D. Oracle Solaris multithreading operating system
- E. ZFS file system with 256-bit checksum

Answer: C, D, E

Explanation:

Enabling Data Deduplication will add overhead. A more CPU-intensive checksum and on-disk deduplication table entries. Deduplication forces the use of the cryptographically

strong SHA-256 checksum.

References:

Oracle ZFS Storage Appliance Administration Guide (June 2014), page 306

For More exams visit <https://killexams.com/vendors-exam-list>



Kill your exam at First Attempt....Guaranteed!