DUMPSBOSS.COM

Oracle ZFS Storage Appliance 2017 Implementation Essentials

Oracle 1z0-499

Version Demo

Total Demo Questions: 10

Total Premium Questions: 70

Buy Premium PDF

https://dumpsboss.com

support@dumpsboss.com

QUESTION NO: 1

Which configuration should improve ZFS throughput and response time for stable write processing (for example, O_SYMC)?

A. ZIL

B. ARC

C. VDEV

D. ZIO

E. DMU

ANSWER: C

Explanation:

Each pool can contain multiple ZFS virtual devices (vdevs).

QUESTION NO: 2

Which type of devices can be tuned using the synchronous write bias properly by selecting values of either latency or throughput?

A. block devices

B. quantum devices

C. log devices

D. cache devices

ANSWER: C

Explanation:

The Synchronous Write Bias setting controls the behavior when servicing synchronous writes. By default, the system optimizes synchronous writes for latency, which leverages the log devices to provide fast response times. In a system with multiple disjointed filesystems, this can cause contention on the log devices that can increase latency across all consumers. Even with multiple filesystems requesting synchronous semantics, it may be the case that some filesystems are more latency-sensitive than others.

References: https://docs.oracle.com/cd/E51475_01/html/E52872/shares__shares__general__synchronous_write_bias.html

QUESTION NO: 3

Identify two workflows on a ZFS Storage Appliance that are used to prepare a system for Oracle Enterprise Manager Monitoring, or to remove the artifacts created for the monitoring environment.

A. Unconfigure for Oracle Enterprise Manager Monitoring.

B. Add agents for Oracle Enterprise Manager Monitoring.

C. Configure for Oracle Enterprise Manager Monitoring.

D. Assign agents for Oracle Enterprise Manager Monitoring.

ANSWER: A B

Explanation:

A: Unconfiguring Oracle Enterprise Manager Monitoring

This workflow removes artifacts created by Configure for Oracle Enterprise Manager Monitoring. Specifically, it:

Removes the oracle_agent role and user

Removes the Oracle Enterprise Manager worksheet

B: Configuring for Oracle Enterprise Manager Monitoring

This workflow is used to prepare an environment for monitoring, or to reset any of the artifacts that were created by the workflow back to their original state in the event the artifacts were changed during operation by the storage administrator. Executing this workflow makes the following changes to the system:

An oracle_agent Role Properties will be created with limited access to the system, to allow the Oracle Enterprise Manager Grid Controller agent to obtain information required for monitoring, but not to make alterations to the system. An oracle_agent user will be created and assigned this role. Use of this role and user is critical to keeping clean audit records for when and how the agent accesses the appliance.

Advanced Analytics will be enabled, makes an extended set of statistics available to all users of the Oracle ZFS Storage appliance.

The Worksheet Oracle Enterprise Manager will be created, facilitating communication between the grid controller administrator and the storage administrator. All metrics monitored by grid controller are available from this worksheet.

References: https://docs.oracle.com/cd/E56047_01/html/E56080/goleu.html#scrolltoc

QUESTION NO: 4

What happens when the thin provisioning property is set for a LUN?

A. A share reserves exactly enough space to completely fill the share.

B. This property allows volume size to exceed the amount of available space.

C. Thin provisioning allows the storage administrator to divide the pool capacity into evenly sized LUNs.

D. A ZFS file system is created by setting quota and reservation to the same value.

ANSWER: B

Explanation:

Thin Provisioning controls whether space is reserved for the volume. This property is only valid for LUNs. By default, a LUN reserves exactly enough space to completely fill the volume. This ensures that clients will not get out-of-space errors at inopportune times. This property allows the volume size to exceed the amount of available space. When set, the LUN will consume only the space that has been written to the LUN.

References: https://docs.oracle.com/cd/E37831_01/html/E52872/shares__shares__general__space_usage.html

QUESTION NO: 5

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 NO: 6

Which option represents the steps to add L2ARC cache SSDs (Readzillas) to a storage pool by using the command-line interface (CLI)?

A. enter configuration storage, import cache to the pool, verify and done

B. enter configuration storage, add cache to the pool, verify and done

C. enter configuration storage, configure cache to the pool, verify and done

D. enter configuration storage, get cache to the pool, verify and done

ANSWER: B

Explanation:

Configuring Storage Using the CLI, Adding Cache Devices to an Existing Pool

References: Oracle ZFS Storage Appliance Administration Guide (June 2014), page 103

QUESTION NO: 7

Identify three valid thresholds used by the Status Dashboard.

A. Sunny

B. Rainy

C. Monsoon

D. Cat-5 Hurricane

E. Thunderstorm

ANSWER: A B D

Explanation:

Use the Thresholds screen to configure the dashboard activity weather icons. The defaults provided are based on heavy workloads and may not be suitable for your environment.

References: https://docs.oracle.com/cd/E56021_01/html/E55851/goosp.html

QUESTION NO: 8

What is the interface of the ZFS Storage Appliance to access the Oracle Storage Cloud Service?

A. Snap Management tool

B. Upload CLI tool

C. Java Library

D. RESTful management API

ANSWER: D

Explanation:

Your applications can access Oracle Storage Cloud Service programmatically by using either an OpenStack Swift-compatible REST API or Java API.

The Oracle ZFS Storage Appliance (ZFSSA) is a family of enterprise storage products that provides efficient file and block data services over the network. The ZFSSA RESTful Application Programming Interface (API) can be used to manage the ZFS Storage Appliance. This documentation is organized using the same hierarchy as the Browser User Interface (BUI) and Command Line Interface (CLI).

Cloud environments typically use a RESTful architecture, which is based on a layered client-server model. This layered model allows services to be transparently redirected through standard hubs, routers, and other network systems without client configuration.

References:

https://www.cloudberrylab.com/solutions/oracle-cloud

https://docs.oracle.com/cd/E51475_01/html/E52433/makehtml-id-4.html#scrolltoc

QUESTION NO: 9

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

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 p. 30

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

QUESTION NO: 10

Which two actions make a statistic into an archived data set?

A. monitoring the statistic in Analytics > Open Worksheet for more than one hour

B. clicking the archive (tape) icon in Analytics > Open Worksheet

C. using “Add statistic…” in Analytics > Open Worksheet

D. saving the statistic in a Worksheet

E. zooming out to more than one hour, and then pausing the statistic

ANSWER: B D

Explanation:

B: Action: Archive

Sets the statistic to be permanently opened and archived to disk. If the statistic had already been opened, then all cached data in memory is also archived to disk. Archiving statistics creates permanent datasets, visible in the Datasets view (those with a non-zero "on disk" value). This is how statistics may be recorded 24x7, so that activity from days, weeks and months in the past can be viewed after the fact.

Note: The term dataset refers to the in memory cached and on disk saved data for a statistic, and is presented as an entity in Analytics with administration controls.

Datasets are automatically created whenever you view statistics in Open Worksheets. A dataset is not saved to disk for future viewing unless you archive it.

References: https://docs.oracle.com/cd/E27998_01/html/E48490/analytics__concepts.html#scrolltoc