Oracle 1Z0-902 Practice Questions

Shared by Carpenter on 10-11-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Which two statements are true about the initial storage configuration after the standard (non-virtualized) deployment of a new Exadata Database Machine with High Capacity storage servers?

Options:

A- The sparse_<DBM_NAME> diskgroup is created automatically.
B- There is free space available on the hard disks inside the database servers for possible extension of the /uoi file system.
C- The DATA_<DBM_Name> and RECO_<DBM_NAME> ASM diskgroups are built on with DATA on the outer-most tracks and RECO on the inner-most tracks of the physical disk.
D- There is free space available on flashdisks inside the Exadata storage servers for possible use for storage indexes.
E- There is free space available on flashdisks inside the Exadata storage servers to configure Exadata Smart Flash Logs.

Answer:

D, E

Explanation:

According to the Oracle Exadata Database Machine Technical Architecture1, the initial storage configuration after the standard (non-virtualized) deployment of a new Exadata Database Machine with High Capacity storage servers includes two ASM disk groups: DATA_<DBM_Name> and RECO_<DBM_NAME>.These disk groups are built on hard disks and flash disks inside the Exadata Storage Servers1.

The correct statements about this configuration are:

Option D: There is free space available on flashdisks inside the Exadata storage servers for possible use for storage indexes. Storage indexes are a feature of Exadata Storage Software that can improve query performance by avoiding unnecessary I/O operations.Storage indexes use a small amount of flash memory to store metadata about data blocks stored on disk1.

Option E: There is free space available on flashdisks inside the Exadata storage servers to configure Exadata Smart Flash Logs. Exadata Smart Flash Logs are another feature of Exadata Storage Software that can improve database performance by using flash memory as an extension of the database redo log buffer.This can reduce latency and increase throughput for redo log writes1.

Question 2

Question Type: MultipleChoice

You are concerned about the operating temperature of the database servers in your Exadata Full Rack and want to be alerted if the server exceeds 30C.

What command should use use to generate alerts for such an event?

Options:

A- dbmcli -e 'set alert ds_temp> 30' on each database server
B- dbmcli -e 'set threshold ds_temp comparison=' >', critical=30' on each database server
C- dbmcli -e 'alert metriccurrent ds_temp where metricValue > 30' on each database server
D- dbmcli -e 'create threshold ds_temp comparison='>', critical=30' on each database server

Answer:

D

Explanation:

According to Oracle's documentation1, to set a threshold for a metric on a database server, you need to use thecreate thresholdcommand with the appropriate parameters.Theset thresholdcommand is used to modify an existing threshold2.

Therefore, the command that you should use to generate alerts for such an event is:

dbmcli -e ''create threshold ds_temp comparison='>', critical=30'' on each database server1

Question 3

Question Type: MultipleChoice

Which two statements are true for the Oracle Exadata Configuration Assistant (OECA)?

Options:

A- OECA reconfigures the size of disk groups and recreates grid disks.
B- OECA extends the hardware for the elastic configuration only.

C- OECA allows one XT storage server in the configuration.
D- OECA's 'Add Equipment' input option allows allocation of RU slots for customer equipment.
E- OECA facilitates PDU power selection initially and after equipment addition.

Answer:

B, E

Explanation:

B and E are true statements for the Oracle Exadata Configuration Assistant (OECA). OECA facilitates PDU power selection initially and after equipment addition, and it allows for the elastic configuration of the hardware. It does not reconfigure the size of disk groups or recreate grid disks, nor does it allow for a single XT storage server in the configuration. The 'Add Equipment' input option does allow for the allocation of RU slots for customer equipment.

References for this information can be found in the Oracle Exadata Database Machine X9M Implementation Essentials Official Text Book and Resources.

Search results: [1] Oracle Exadata Database Machine X8-2[1][2]: Configuration Assistant Guide ... The Oracle Exadata Configuration Assistant (OECA) is a web-based tool that can ... [2] The Oracle Exadata Configuration Assistant (OECA) is a web-based tool that can be used to ... Add Equipment: This input option allows you to allocate RU slots for ... [3] Oracle Exadata Database Machine X7-2 and X7-2L Installation Guide ... The Oracle Exadata Configuration Assistant (OECA) is a web-based tool that can ... [4] The Oracle Exadata Configuration Assistant (OECA) is a web-based tool that can be used to ... Power Selection[1][2]: This input option allows you to select PDU power initially ... [5] Oracle Exadata Database Machine X7-2 and X7-2L Installation Guide ... The Oracle Exadata Configuration Assistant (OECA) is a web-based tool that can ...

Question 4

Question Type: MultipleChoice

Which are two valid reasons for executing an X9M-2 Exadata storage server rescue procedure?

Options:

A- the failure of physical disk 1
B- corruption in the / (root) filesystem
C- the failure of both physical M.2 disks
D- the failure of physical disk 0 and 11
E- moving all disks from one cell to another as part of a chassis-level component failure
F- accidental loss of all data from all griddisks in a storage server

G- corruption in a normal or high redundancy ASM diskgroup

Answer:

B, C

Explanation:

The two valid reasons are:

corruption in the / (root) filesystem.According to1, this is one of the scenarios that requires a storage server rescue procedure , because it prevents the storage server from booting up normally .The rescue procedure will re-image the root filesystem with a fresh copy of the operating system and restore the configuration files from a backup1.

the failure of both physical M.2 disks.According to2, each Exadata Storage Server X9M-2 has two physical M.2 disks that store the operating system and configuration files for the storage server . If both M.2 disks fail , then the storage server cannot boot up or function properly .The rescue procedure will replace the failed M.2 disks with new ones and re-image them with the operating system and configuration files2.

Question 5

Question Type: MultipleChoice

Which two Exadata security features would you configure to control which databases can access which griddisks, when multiple databases share space on your storage servers in an unpartitioned storage grid?

Options:

A- Exadata storage realms using ASM-scoped security mode
B- File permissions on the griddisks in each database server
C- Using EXADCLI instead of CELLCLI to create the griddisks
D- File permissions on the griddisks in each storage server
E- Exadata storage realms using database-scoped security mode

Answer:

A, E

Explanation:

Exadata storage realms using ASM-scoped security mode allows you to control which database can access which griddisks by creating realms at the ASM level, thereby allowing you to limit the access of each database to the griddisks that it needs to access. This is covered in section 3.15.2 of the Oracle Exadata Database Machine X9M Implementation Essentials Official Text Book and Resources [1].

Exadata storage realms using database-scoped security mode allows you to control which database can access which griddisks by creating realms at the database level[1], thereby allowing you to limit the access of each database to the griddisks that it needs to access. This is covered in section 3.15.3 of the Oracle Exadata Database Machine X9M Implementation Essentials Official Text Book and Resources [2].

[1]https://docs.oracle.com/en/engineered-systems/exadata-database-machine/x9m/exad-implementing-database-machine-x9m.pdf[2]https://docs.oracle.com/en/engineered-systems/exadata-database-machine/x9m/exad-implementing-database-machine-x9m.pdf

Question 6

Question Type: MultipleChoice

Which statement is true about the disks that are used in ASM diskgroups on the Exadata platform?

Options:

A- In an Exadata storage server, each disk is placed into a separate failure group.
B- The grid disks in ASM diskgroups must have similar performance characteristics in addition lo having the same size.
C- Sparse diskgroups are virtual objects allocated from the DATA diskgroup.
D- If High redundancy is selected, grid disks from the primary and secondary Exadata Database Machine must be specified.
E- Extreme Flash (EF) and High Capacity (HC) disks can be mixed to any capacity within the same diskgroup.

Answer:

B

Explanation:

On Exadata systems, ASM disk group configurations require that:

Grid Disks Must Have Similar Performance Characteristics: Within a disk group, grid disks must not only be of the same size but also need to exhibit similar performance traits. This consistency ensures balanced I/O operations and optimizes the performance for ASM mirroring and striping across disks.

Sparse Diskgroups: These are specialized virtual objects that reside within the DATA disk group. Sparse diskgroups are used to support snapshots and other space-efficient storage solutions by creating virtual representations of grid disks within an existing disk group without consuming additional physical space.

These characteristics are fundamental to ensuring that ASM maintains optimal performance and storage efficiency on Exadata platforms.

Question 7

Question Type: MultipleChoice

Examine these commands:

1. Execute "crscti stop cluster -all" as the grid user from one database server.

2. Execute "crscti stop cluster -all" as root from one database server.

3. Power off all network switches.

4. Execute "crscti stop cluster" as root from one database server.

5. Execute "crscti stop cluster" as the grid user from one database server.

6. Power off the rack using the power switches on the PDUs.

7. Execute "shutdown -h now" on all database servers.

8. Execute "shutdown -h now" on all Exadata storage servers.

Which is the correct order or the required commands to completely power off an Exadata Database Machine in an orderly fashion?

Options:

A- 5, 8, 7, and 6
B- 4, 7, 8, 3, and 6
C- 2, 8, 7, 3, and 6
D- 2, 7, 8, and 6
E- 1, 8, 7, 3 and 6

Answer:

A

Explanation:

Option 5 stops the clusterware on all database servers as the grid user3.

Option 8 shuts down all Exadata storage servers using their ILOM interfaces1.

Option 7 shuts down all database servers using their ILOM interfaces1.

Option 6 powers off the rack using the power switches on the PDUs2.

To Get Premium Files for 1Z0-902 Visit

https://www.p2pexams.com/products/1z0-902

For More Free Questions Visit

https://www.p2pexams.com/oracle/pdf/1z0-902

