



1Z0-902^{Q&As}

Oracle Exadata Database Machine X9M Implementation Essentials

Pass Oracle 1Z0-902 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.geekcert.com/1z0-902.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Oracle
Official Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers



**QUESTION 1**

You have been asked to design a backup solution for an Exadata X9M-2 Quarter Rack with Extreme Flash Storage Servers connected to a new ZFS Storage Appliance ZS7 with 2 Storage Controllers with 100Gb Ethernet cards and 3 Storage Trays. You are using Oracle Exadata Configuration Assistant to validate the rack layout.

1.

Use "Add Equipment" to add the Exadata X9M EF Storage Servers, starting from RU10.

2.

Use drop down to add ZFS Storage Appliance Controllers.

3.

You cannot add ZFS Storage Appliance to an Exadata Rack.

4.

Use "Add Equipment" to add the ZFS Storage Trays, starting from RU1.

5.

Use drop down to add ZFS Storage Trays.

6.

Use "Add Equipment" to add the Exadata X9M Database Servers, starting from RU16.

7.

Use "Add Equipment" to add the Exadata X9M EF Storage Servers, starting from RU1.

8.

Use drop down to add Exadata X9M EF Storage Servers.

9.

Use "Add Equipment" to add the ZFS Storage Controllers, starting from RU27.

10.

Use drop down to add Exadata X9M Database Servers.

11.

Use "Add Equipment" to add the ZFS Storage Trays, starting from RU31. Which of these steps are correct and what is their correct order?

A. 10, 8, 2, 5

B. 4, 1, 6, 9



C. 3

D. 10, 8, 9, 11

E. 7,6,9,11

Correct Answer: D

Explanation: The correct order of steps is 10, 8, 9, 11. The 10th step is to use the drop down to add Exadata X9M Database Servers, the 8th step is to use the drop down to add Exadata X9M EF Storage Servers, the 9th step is to use "Add Equipment" to add the ZFS Storage Controllers, and the 11th step is to use "Add Equipment" to add the ZFS Storage Trays. These steps are referenced in the Oracle Exadata Database Machine X9M Implementation Essentials Official Textbook, which is available online at https://docs.oracle.com/cd/E80437_01/E80437/html/index.html.
<https://docs.oracle.com/en/engineered-systems/exadata-database-machine/dbmin/configuring-exadata.html>

QUESTION 2

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?

A. The sparse_ 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_ and RECO_ 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.

Correct Answer: DE

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_ and RECO_. 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 3

Your customer needs to ensure that their data is available on the Exadata machine during updates. The customer wants to be able to update one server at a time but still be protected against single-node server failure.

What ASM redundancy level should they use?



- A. Normal
- B. Sparse
- C. High
- D. External
- E. Extended

Correct Answer: A

Your customer needs to ensure that their data is available on the Exadata machine during updates. The customer wants to be able to update one server at a time but still be protected against single-node server failure. This means that they need a redundancy level that can tolerate at least one disk failure without losing data availability or performance. Therefore, the best option for them is normal redundancy . Normal redundancy provides two-way mirroring by default, which means that each data extent has a primary copy and a mirror copy on different disks in different failure groups¹. A failure group is a subset of disks within a disk group that are likely to fail at the same time due to common dependencies². On Exadata Database Machine , each Storage Server is considered as a failure group by default². This means that if one Storage Server fails or needs to be updated , the data on its disks can still be accessed from another Storage Server that has the mirror copies of the same data extents². Therefore, normal redundancy ensures data availability and performance during updates or single-node server failures. <https://docs.oracle.com/en/engineered-systems/exadata-databasemachine/sagug/exadata-administering-asm.html>

QUESTION 4

Examine these commands:

1.

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

2.

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

3.

Power off all network switches.

4.

Execute "crsctl stop cluster" as root from one database server.

5.

Execute "crsctl 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?

- 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

Correct Answer: C

QUESTION 5

You are going to create an ACFS Filesystem for use as a staging area for data loads. Examine the following commands:

```
1) /usr/sbin/mkfs -t acfs /dev/asm/acfs_dataload_vol-123
2) asmcmd volinfo -G data acfs_dataload_vol | grep Device
3) acfsutil size +250G -d /dev/asm/acfs_dataload_vol-123 /mnt/acfs_dataload_vol/
4) sudo /u01/app/19.0.0.0/grid/bin/srvctl start filesystem -device /dev/asm/acfs_dataload_vol-123
5) sudo chown -R oracle:dba /mnt/acfs_dataload_vol
6) asmcmd volcreate -G data -s 1024G acfs_dataload_vol
7) /sbin/acfsutil snap create -w data_load /mnt/acfs_dataload_vol
8) /u01/app/19.0.0.0/grid/bin/srvctl add filesystem -device /dev/asm/acfs_dataload_vol-123 -path /mnt/acfs_dataload_vol -
user oracle
9) sudo /u01/app/19.0.0.0/grid/bin/srvctl stop filesystem -device /dev/asm/acfs_dataload_vol-123
```

What steps and in which order should they be executed?

- A. 6, 2, 1, 8, 4, 5
- B. 6, 8, 3, 4, 2, 1
- C. 6, 3, 1, 9, 7, 8
- D. 6, 1, 2, 7, 5, 8

Correct Answer: A

Explanation: The steps and the order in which they should be executed to create an ACFS Filesystem for use as a staging area for data loads are: `asmcmd volcreate -G data -s 1024G acfs_dataload_vol` : This command creates a new volume in the data diskgroup, with a size of 1024 GB, and named `acfs_dataload_vol`. `asmcmd vol info -G data acfs_dataload_vol | grep Device`: This command will show the device name of the newly created volume `/usr/sbin/mkfs`



-t acfs /dev/asm/ac:fs_dataload_vol-123: This command creates a new ACFS filesystem on the volume created in step 1
1 sudo /u01/app/19.0.0.0/grid/bin/srvctl start filesystem -device /dev/asn/acfs_dataload_vol-123: This command starts the filesystem created in step 3
sudo chown -R oracle:dba /mnt/acfs_dataload_vol: This command changes the ownership of the /mnt/acfs_dataload_vol to the oracle:dba user.
./sbin/acfsutil snap create -v data_load /mnt/acfs_dataload_vol: This command creates a snapshot of the filesystem to be used as a staging area.
acfsutil size +2S0G -d /dev/asm/acfs_dataload_vol-123 /mnt/acfs_dataload_vol/ : This command increases the size of the filesystem by 2GB
/u01/app/19-G.O-Q7grid/bin/srvctl add filesystem -device /dev/asm/acfs_dataload_vol-

[Latest 1Z0-902 Dumps](#)[1Z0-902 PDF Dumps](#)[1Z0-902 Braindumps](#)