



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.pass4itsure.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 tasked with replacing a memory module of an Exadata Storage Server and need to power off the affected storage server. Which two commands must you execute to safely power off the storage server in an Exadata X9M Database Machine?

- A. CellCLI> alter cell shutdown SERVICES all on the affected storage server
- B. CellCLI> list GRIDDISK where status != 'inactive' on the affected storage server
- C. \crsctl stop cluster -all on one of the database servers
- D. CellCLI LIST GRIDDISK ATTRIBUTES name WHERE asmdeactivationoutcome != 'Yes' on the affected Storage server
- E. CellCLI alter GRIDDISK all inactive on the affected storage server
- F. shutdown -h now on the affected storage server

Correct Answer: EF

Explanation: <https://docs.oracle.com/en/engineered-systems/exadata-database-machine/dbmmn/maintaining-exadata-storage-servers.html#GUID-AE16A1DA-53C6-4E80-94E5-963AA65373AB>

The two commands that must be executed to safely power off the storage server in an Exadata X9M Database Machine are E and F.

Command E, CellCLI alter GRIDDISK all inactive, will deactivate all of the GRIDDISKS on the affected storage server. This will ensure that no data is lost during the power-off process.

Command F, shutdown -h now, will shut down the storage server. This will ensure that the storage server is completely powered off and no data is lost. For more information on how to properly power off an Exadata Storage Server, refer to

the Oracle Exadata Database Machine X9M Implementation Essentials official text book and resources [1][2].

https://support.oracle.com/knowledge/Oracle%20Database%20Products/1188080_1.html

QUESTION 2

Which two options can be used to identify a damaged or failing flash card on an X9M-2 Database Machine High Capacity storage server?

- A. using the CELLCLI CALIBRATE command on the storage server after logging in as the celladmin user
- B. using the CELLCLI CALIBRATE command on the storage server after logging in as the root user
- C. hardware monitoring using the storage server ILOM
- D. using the CELLCLI list LUN DETAIL command as the celladmin user

Correct Answer: AC

The CELLCLI CALIBRATE command can be used to check the health of all flash cards in the storage server, and it



should be run after logging in as the celladmin user. The hardware monitoring using the storage server ILOM should also be used to check for errors and other potential issues with the flash cards. The CELLCLI list LUN DETAIL command should not be used to identify a damaged or failing flash card, since it will not provide any information about the health of the flash card. <https://docs.oracle.com/en/engineered-systems/exadata-database-machine/dbmmn/index.html>

QUESTION 3

You are working on a remote Exadata Database Machine with current software and firmware. You now need to start a storage server after hardware maintenance. Which ILOM command starts the storage server?

- A. -> start /System
- B. -> start /SP/console
- C. -> set /SYS power_state=ON
- D. You must be physically present in the data center to start an Exadata Storage Server.
- E. -> start /SYSTEM/power
- F. -> start /SYS/power

Correct Answer: C

Explanation: This command is based on the Oracle Exadata Database Machine X9M Implementation Essentials Official Text Book and other resources. This command sets the power state of the system to ON and starts the storage server. It is the only command that is able to start the storage server remotely. All the other commands are incorrect or require the user to be physically present in the data center.

QUESTION 4

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: `asmcrad 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/acfs_dataload_vol-123`: This command creates a new ACFS filesystem on the volume created in step 1 `sudo /u01/app/19.0.0.0/grid/bin/srvctl start filesystem -device /dev/asm/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-`

QUESTION 5

You are in the process of upgrading your nonvirtualized X9M-2 Database Machine elastic configuration with 4 database servers and 7 HC storage servers with an additional 4 database servers and 7 HC storage servers.

The new storage servers are called DM01CEL08 through dmoicel14.

After creating 96 new griddisks, you issued this SQL statement:

```
SQL> ALTER DISKGROUP DATA ADD DISK
 2> 'O/*/DATA*DM01CEL08*'
 3> 'O/*/DATA*DM01CEL09*'
 4> 'O/*/DATA*DM01CEL10*'
 5> 'O/*/DATA*DM01CEL11*'
 6> 'O/*/DATA*DM01CEL12*'
 7> 'O/*/DATA*DM01CEL13*'
 8> 'O/*/DATA*DM01CEL14*'
 9> REBALANCE POWER 512;
```

How many failgroups if any, will be added to the DATE diskgroup by executing this SQL statement?

- A. 1 consisting of all 96 griddisks
- B. 96 consisting of one griddisk each
- C. 0 because the new griddisks will be added to the existing faiigroups
- D. 12 consisting of seven griddisks each
- E. 7 consisting of 12 griddisks each

Correct Answer: A

Explanation: This SQL statement is adding the new griddisks to the existing diskgroup "DATA" and creating one new



failgroup, consisting of all 96 griddisks. The "REBALANCE POWER 512" option tells the system to perform a rebalance operation with a power of 512. It means the system distributes the data evenly across all the disks in the diskgroup using a power of 512.

[1Z0-902 PDF Dumps](#)

[1Z0-902 Study Guide](#)

[1Z0-902 Braindumps](#)