

# 1Z0-902<sup>Q&As</sup>

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.pass2lead.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

Which three statements are true about the CELLCLI command?

- A. It can execute commands on multiple storage servers in parallel.
- B. It has command-line history.
- C. It requires root privileges to create CELLDISKS and GRIDDISKS.
- D. It can be executed on storage servers.
- E. It can be executed using the DCLI utility.

Correct Answer: BDE

The statements that are true about the CellCLI command are:

It has command-line history . According to 1, you can use the up and down arrow keys to scroll through previous commands that you entered in CellCLI . You can also use Ctrl+R to search for a previous command by entering part of it . It can

be executed on storage servers . According to 1, you can invoke CellCLI by logging in to the Exadata cell as cellmonitor , celladmin , or root user and typing "cellcli" . Within CellCLI , you can perform various management and monitoring tasks

for the cell134.

It can be executed using the DCLI utility . According to 5, DCLI (distributed CLI) is another utility that can be used to execute scripts and commands across multiple storagecells from a single interface. You can use DCLI to run CellCLI

commands on multiple cells at once54.

---

### QUESTION 2

Which two statements are true about enabling write-back flash cache?

- A. Write-back flash cache is on by default for High Capacity Storage Servers in High Redundancy configurations.
- B. Write-back Flash Cache cannot be used on Extreme Flash Storage Servers.
- C. When enabling write-back flash cache in a rolling manner, DBCLI should be used to inactivate the grid disks on all cells first.
- D. Before write-back flash cache is enabled, you need to drop the Flash Cache first.
- E. When enabling write-back flash cache in a non-rolling manner, you need to drop the Persistent Memory Commit Accelerator (PMEMLog) first.

Correct Answer: AD

Explanation: According to the Oracle documentation1, two statements that are true about enabling write-back flash cache are:

Write-back flash cache is on by default for High Capacity Storage Servers in High Redundancy configurations (A). This is because write-back flash cache provides better performance and reliability for data file writes than write-through flash cache. Before write-back flash cache is enabled, you need to drop the Flash Cache first (D). This is because changing the flash cache mode requires recreating the flash cache with a different configuration.

---

### QUESTION 3

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 4

Which statement is true about the Persistent Memory Commit Accelerator?

- A. Persistent Memory Commit Accelerator tracks changes to Persistent Memory Data Accelerator to ensure duplicate blocks are not written to Flash.
- B. Persistent Memory Commit Accelerator helps to further reduce redo log write latency by using Persistent Memory and Remote Direct Memory Access (RDMA).
- C. Persistent Memory Commit Accelerator copies redo log data from disk for faster redo apply on Data Guard Standby Databases.
- D. Persistent Memory Commit Accelerator contains logging information from all tiers of the software stack for rapid triage and diagnostics.
- E. Persistent Memory Commit Accelerator reduces redo log write latency by using Persistent Memory and RDMA before flushing to Flash then disk.

Correct Answer: B

Explanation: <https://docs.oracle.com/en/engineered-systems/exadata-database-machine/dbmso/new-features-exadata-system-software-release-23.html>

---

#### QUESTION 5

You have configured a multirack Database Machine with two X9M-8 full racks all in a single cluster and storage grid comprising a total of 4 X9M-8 Database servers and 28 X9M-8 Storage servers. Which two options are true regarding the servers on which Enterprise Manager agents must be deployed in order to monitor all components of this multirack configuration?

- A. on at least two storage servers in both racks
- B. on only one database server in both racks
- C. on all database servers in the first rack
- D. on all storage servers in both racks
- E. on all database servers and at least two storage servers in both racks
- F. on all database servers in the second rack
- G. on all database servers in both racks and one storage server in each rack

Correct Answer: AE

Explanation: In order to monitor all components of this multirack configuration, Enterprise Manager agents must be deployed on at least two storage servers in both racks and all database servers in both racks. This is according to the Oracle Exadata Database Machine X9M Implementation Essentials official text book [1], in which it states "To monitor all components of the multirack configuration, you must deploy agents on all database servers and at least two storage servers in each rack" (page 6-15).

---

#### QUESTION 6

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

- 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

Correct Answer: BC

Explanation: The two valid reasons are:

corruption in the / (root) filesystem . According to 1, 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 to 2, 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 7

An Exadata X9M-2 Elastic Rack with 4 Database Servers and 8 HC Storage Servers and 3-phase 24kVA PDUs is being installed in a Data Center. However, the Data Center is only providing enough power for a single cable from each PDU. Which statement is correct?

- A. A splitter cable can be used to provide power to all PDU cables.
- B. The power cables from the servers to the PDUs can be rearranged inside the rack following OECA guidance to utilize a single PDU power cable.
- C. The installation cannot proceed until two power feeds are available per PDU.
- D. The installation can go ahead, no change is required.

Correct Answer: C

Explanation: In order for an Exadata X9M-2 Elastic Rack to be installed, two power feeds are required for each of the three-phase 24kVA PDUs. A single cable from each PDU will not be enough to power the rack, and neither a splitter cable nor rearrangement of the power cables from the servers to the PDUs will be able to provide sufficient power for the system. Therefore, the installation cannot proceed until two power feeds are available per PDU. For more information about power requirements for Exadata X9M-2 Elastic Racks, refer to the Oracle Exadata Database Machine X9M Implementation Essentials official text book and resources. Search results: [1] Oracle Exadata Database Machine X9M-2 Full Rack Installation - Oracle Docs [2] Exadata Database Machine X8M-2 Mid-Size Rack Installation - Oracle Docs <https://docs.oracle.com/en/engineered-systems/exadata-database-machine/dbmin/exadata-site-requirements.html>

---

#### QUESTION 8

An Exadata storage server physical disk on an X9M-2 high-capacity full rack entered the predictive failure state. Which two steps must you perform to replace this failed physical disk?

- A. Replace the failed physical disk.
- B. Add the griddisks back into the ASM diskgroup they used to be a member of.
- C. Identify the griddisks located on the failed physical disk and drop them from the associated ASM diskgroups.
- D. Rebalance the data on the effected griddisks before performing a manual drop command.
- E. Create a new celldisk and new griddisks on the replaced physical disk.
- F. Verify that the griddisks located on the physical disk have been successfully dropped from the associated ASM diskgroups.

Correct Answer: AC

Explanation: Once the physical disk has been replaced, you will need to create a new celldisk and griddisks on the replaced physical disk. After the griddisks have been dropped from the associated ASM diskgroups, you can then add the griddisks back into the ASM diskgroup they used to be a member of and rebalance the data on the effected griddisks. Finally, you should verify that the griddisks located on the physical disk have been successfully dropped from the associated ASM diskgroups.

---

#### QUESTION 9

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 GRDDISK where status != 'inactive' on the affected storage server
- C. \crsctl stop cluster -all on one of the database servers
- D. CellCLI LIST GRDDISK ATTRIBUTES name WHERE asmdeactivationoutcome != 'Yes' on the affected Storage server
- E. CellCLI alter GRDDISK 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 GRDDISK all inactive, will deactivate all of the GRDDISKS 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](https://support.oracle.com/knowledge/Oracle%20Database%20Products/1188080_1.html)

---

#### QUESTION 10

Which two of the following network connection types can be deployed on the client network ports of an Exadata X9M-2 Database Server?

- A. 2x 25Gb ports on eth1 and eth2 using SFP28 transceivers and OM4 fiber cable for backup network, and 2x 25Gb ports on eth5 and eth6 using SFP28 transceivers and OM4 fiber cable for client network

B. 2x 25Gb ports on eth1 and eth2 using SFP28 transceivers and OM4 fibre cable for client network, and 2x 10Gb ports on eth9 and eth10 using RJ45 cat6 cable for backup network

C. 2x 10Gb ports on eth1 and eth2 using FU45 cat6 cable for client network, and 2x 25Gb ports on eth11 and eth12 using SFP28 transceivers and OM4 fiber cable for backup network

D. 2x 25Gb ports on eth1 and eth2 using SFP28 transceivers and OM4 fiber cable for backup network, and 2x 10Gb ports on eth3 and eth4 using RJ45 cat6 cable for client network

Correct Answer: AB

Explanation: <https://www.oracle.com/a/ocom/docs/engineered-systems/exadata/exadata-x9m-2-ds.pdf>

[1Z0-902 PDF Dumps](#)

[1Z0-902 Study Guide](#)

[1Z0-902 Braindumps](#)