



# 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.passapply.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 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 documentation<sup>1</sup>, 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 2

You want to monitor how a large production table is accessed. Especially, you are interested to see how the access on that particular table leverages the benefits of the Exadata Platform.

Which two actions are NOT appropriate for that purpose?

- A. YOU query V\$SYSTEM\_EVENTS and filter for the event `cell physical IO interconnect bytes returned by smart scan`, associated to your table.
- B. You query `v$segment_statistics` and filter for the Object ID of your table from `dba_objects` and the the column `STATISTIC_NAME='optimized physical reads'`.
- C. You query `v$SYSTAT` and filter for the statistic `cell smart table scan`, associated to your table.
- D. You run the `CellCli-command list activerequest`, filtering for the attributes `ioReason` and `objectNumber`, that you specify as `Smart Scan` and the Object ID of your table from `DBA_OBJECTS`.

Correct Answer: AC

- A. YOU query V\$SYSTEM\_EVENTS and filter for the event `cell physical IO interconnect bytes returned by smart scan`, associated to your table. This does not show how much data was filtered by Smart Scan, but only how much data was returned after Smart Scan<sup>12</sup>.
- C. You query `v$SYSTAT` and filter for the statistic `cell smart table scan`, associated to your table. This does not show how much data was filtered by Smart Scan for a specific table, but only for all tables in a session<sup>12</sup>. <https://>



[www.databasejournal.com/oracle/monitoring-smart-scans-in-oracle-exadata/](http://www.databasejournal.com/oracle/monitoring-smart-scans-in-oracle-exadata/)

---

### QUESTION 3

Examine this list of software components:

1.  
Oracle KVM Guest
2.  
Oracle Enterprise Manager Agent (OMA)
3.  
ASM instance
4.  
RDBMS instance
5.  
Automatic Diagnostic Repository Command Interpreter (ADRCI)
6.  
CELLCLI
7.  
Cell Server(CELLSRV)
8.  
diskmon
9.  
Restart Server (RS)
10.  
Management Server (MS)

What is the correct location where these software components can run in the standard Exadata Database Machine deployment?

- A. 2, 3, 4, 8, and 10 run on the database servers; 1, 5, 6, 7 and 9 run on the Exadata storage servers.
- B. 1, 2, 3, 4, 9 and 10 run on the database servers; 5, 6, 7, 8, 9, and 10 run on the Exadata storage servers.
- C. 1, 2, 3, 4, 5, 8, 9 and 10 run on the database servers; 5, 6, 7, 9 and 10 run on the Exadata storage servers.



D. 3, 4, 8, and 10 run on the database servers; 1, 2, 5, 6, 7 and 9 run on the Exadata storage servers.

E. 1, 2, 3, 4, 8 and 9 run on the database servers; 5, 6, 7, 9 and 10 run on the Exadata storage servers.

Correct Answer: C

Oracle KVM Guest: This is a virtual machine that runs on top of Oracle Linux KVM hypervisor. It can be used to run Oracle Database or other applications on Exadata Database Machine<sup>2</sup>. Therefore, it runs on the Database Servers.

Oracle Enterprise Manager Agent (OMA): This is a software agent that communicates with Oracle Enterprise Manager Cloud Control and provides monitoring and management capabilities for Exadata Database Machine<sup>2</sup>. Therefore, it runs on both Database Servers and Storage Servers.

ASM instance: This is an instance of Oracle Automatic Storage Management (ASM), which is a volume manager and a file system for Oracle Database files. It manages diskgroups that span across multiple Storage Servers<sup>2</sup>. Therefore, it runs on the Database Servers.

RDBMS instance: This is an instance of Oracle Database that processes SQL statements and executes transactions. It uses ASM disk groups to store data files, control files, redo log files, etc<sup>2</sup>. Therefore, it runs on the Database Servers.

Automatic Diagnostic Repository Command Interpreter (ADRCI): This is a command-line tool that enables you to view diagnostic data stored in the Automatic Diagnostic Repository (ADR). ADR is a file-based repository for database diagnostic data such as trace files, alert logs, etc<sup>2</sup>. Therefore, ADRCI runs on both Database Servers and Storage Servers, depending on where the ADR is located.

CELLCLI: This is a command-line interface that enables you to configure and manage Exadata Storage Server Software. It allows you to perform tasks such as creating disk groups, monitoring cell health, applying patches, etc<sup>2</sup>. Therefore, it runs on the Storage Servers.

Cell Server(CELLSRV): This is a process that runs on each Storage Server and handles I/O requests from the Database Servers. It implements Exadata Smart Scan , which offloads data-intensive SQL operations from the Database Servers to the Storage Servers<sup>2</sup>. Therefore, it runs on the Storage Servers.

diskmon: This is a process that monitors the status of disks and flash devices on each Storage Server. It reports disk failures and performs automatic disk reclamation<sup>2</sup>. Therefore, it runs on the Storage Servers.

Restart Server (RS): This is a process that manages automatic restarts of critical processes such as CELLSRV , MS , or OMA in case of failures. It also handles graceful shutdowns and startups of all processes on each server<sup>2</sup>. Therefore, it runs on both Database Servers and Storage Servers.

Management Server (MS): This is a process that provides management services for each server such as collecting metrics , logging events , executing commands from CELLCLI , etc<sup>2</sup>. Therefore, it runs on both Database Servers and Storage Servers

#### QUESTION 4

Which three steps are required to expand an Exadata X9M-2 Database Server with the lowest memory configuration available to the highest memory configuration?

- A. Add 12x 32GB DIMMs.
- B. Add 12x 64GB DIMMs.
- C. Shutdown the Database Server if running.
- D. Add 16x 32GB DIMMs.
- E. Add 16x 64GB DIMMs.
- F. Memory cannot be expanded on Exadata X9M-2 Database Servers.
- G. Remove existing memory modules.
- H. Add 32x 64GB DIMMs.
- I. Add 24x 32GB DIMMs.

Correct Answer: CGH



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

<https://chriscraftoracle.wordpress.com/2022/11/09/what-makes-exadata-faster/>

---

#### QUESTION 5

You have a script with several CELLCLI commands that must be executed on each cell in your Exadata full rack. The script must run on each cell simultaneously. How must you achieve this?

- A. Copy and execute the script on all storage servers using the EXACLI command.
- B. Copy the script to all storage servers using the dcli command and manually execute it on all storage servers using the EXACLI command.
- C. Copy and execute the script on all storage servers using the dcli command.
- D. Copy the script to all storage servers using the dcli command and manually execute it on each storage server using the DCLI command.
- E. Copy the script to all storage servers using the CELLCLI command and execute it on all storage servers in parallel using the CELLCLI command.

Correct Answer: C

Explanation: According to Oracle's documentation<sup>1</sup>, the Cell Control Command-Line Interface (CellCLI) utility is the command-line administration tool for Oracle Exadata System Software. It can be used to perform various tasks on Exadata Storage Servers, such as creating and managing objects, monitoring performance, and configuring alerts. To run CELLCLI commands on multiple cells simultaneously, you can use the Distributed Command-Line Interface (dcli) utility<sup>2</sup>. The dcli utility allows you to execute commands or scripts on a group of cells or hosts in parallel. You can use the dcli utility to copy files from one cell or host to multiple cells or hosts using the -f option<sup>3</sup>. You can also use the dcli utility to execute commands or scripts on multiple cells or hosts using the -l option to specify a user name and the -g option to specify a file containing a list of cells or hosts<sup>3</sup>. Therefore, to achieve your requirement of running a script with several CELLCLI commands on each cell in your Exadata full rack simultaneously, you must: Copy and execute the script on all storage servers using the dcli command.

[1Z0-902 VCE Dumps](#)

[1Z0-902 Practice Test](#)

[1Z0-902 Braindumps](#)