

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.leads4pass.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 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.
 - 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. <https://www.databasejournal.com/oracle/monitoring-smart-scans-in-oracle-exadata/>
-

QUESTION 2

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 3

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 4

I/O performance of the prod database on your Exadata Database Machine has degraded slightly over the past month. The database has been allocated to the OLTP I/O Resource Management (IORM) category. Which two monitoring tools might be useful in examining I/O performance for the prod database?

- A. OS I/O metrics using Enterprise Manager host pages for the storage servers
- B. OS I/O metrics using OS tools such as iostat on the database servers
- C. I/O-specific dynamic performance views such as v\$iostat_fiie, v\$iostat_function, and v\$iostat_consumer_group from the prod database instances using SQL *p1us
- D. cellcli (or exaclcli/exadcli) to examine storage server metrics such as database, category, ceiidisk, and griddisk
- E. OS I/O metrics using OS tools such as iostat on the storage servers

Correct Answer: CD

Explanation: According to the Oracle documentation¹, two monitoring tools that might be useful in examining I/O performance for the prod database are:

cellcli (or exaclcli/exadcli) to examine storage server metrics such as database, category, ceidisk, and griddisk (D). This tool can help you monitor the I/O Resource Management (IORM) metrics and identify any bottlenecks or imbalances in the

storage layer.

I/O-specific dynamic performance views such as v\$iostat_fiie, v\$iostat_function, and v\$iostat_consumer_group from the prod database instances using SQL *p1us ? These views can help you monitor the I/O activity and latency at the file,

function, and consumer group level.

QUESTION 5

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>

[Latest 1Z0-902 Dumps](#)

[1Z0-902 VCE Dumps](#)

[1Z0-902 Study Guide](#)