



1Z0-070^{Q&As}

Oracle Exadata X5 Administration

Pass Oracle 1Z0-070 Exam with 100% Guarantee

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

<https://www.geekcert.com/1z0-070.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 completely prevent a Smart Scan from occurring?

- A. querying a table containing many chained rows
- B. querying a table containing many migrated rows
- C. performing a minimum or maximum function on an indexed column
- D. performing a Fast Full Index scan on a reverse key index
- E. referencing more than 255 columns from an OLTP compressed table in a query
- F. querying a table containing a JSON column

Correct Answer: BE

Explanation:

B: Migrated Rows is a special case of chained rows. Migrated rows still affect performance, as they do in conventional storage situations, but with the additional overhead of reducing the beneficial effects of "Smart Scan" in addition to increasing the number of I/Os

E: Smart Scans - broadly speaking and ignoring edge cases - can only transport a maximum of 254 columns from a single (non-HCC) segment. Requesting more columns will simply disable Smart Scans for that segment.

An interesting limitation to Exadata Smart Scans - if more than 254 columns from a table (not HCC compressed, more on that in moment) need to be projected, Smart Scans for that particular segment will be disabled and Exadata will fall back to conventional I/O. This means that the number of columns in the projection clause can make a significant difference to performance, since only Smart Scans allow taking advantage of offloading and particularly avoiding I/O via Storage Indexes.

Incorrect Answers:

A: Smart scan can cope with some cases of chained rows

References: <http://oracle-randolf.blogspot.se/2013/01/exadata-smart-scan-projection-limitation.html>

QUESTION 2

Which two are true about Smart Scan?

- A. A query rewrite may occur to a materialized view stored in Exadata, and it will always benefit from Smart Scan.
- B. Column projection does not contribute to the performance benefit of Smart Scan.
- C. A query rewrite may occur to a materialized view stored in Exadata but will never benefit from Smart Scan.
- D. It is possible to offload single row functions to the Exadata Storage Servers.
- E. Join processing can be partly offloaded to Exadata Storage Servers.



F. Join processing can be fully offloaded to Exadata Storage Servers.

Correct Answer: DE

Explanation:

D: With Exadata storage, database operations are handled much more efficiently. Queries that perform table scans can be processed within Exadata storage with only the required subset of data returned to the database server. Row filtering, column filtering and some join processing (among other functions) are performed within the Exadata storage cells. When this takes place only the relevant and required data is returned to the database server.

E: With Smart Scan processing, row filtering, column filtering, some join processing, and other functions are performed in the Exadata cells.

References: <http://www.centroid.com/blog/exadata-smart-scan-processing>

QUESTION 3

Which two options should be used to identify a damaged or failing flash card on an X6 Database Machine high-capacity storage server?

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

Correct Answer: AD

Explanation:

The CALIBRATE command runs raw performance tests on cell disks, enabling you to verify the disk performance before the cell is put online. You must be logged on to the cell as the root user to run CALIBRATE.

Oracle ILOM automatically detects system hardware faults and environmental conditions on the server. If a problem occurs on the server, Oracle ILOM will automatically do the following:

Illuminate the Service Required status indicator (LED) on the server front and back panels.

Identify the faulted component in the Open Problems table.

Record system information about the faulted component or condition in the event log.

References: http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-server-cellcli.htm#SAGUG20617

References: http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-server-cellcli.htm#SAGUG20617



QUESTION 4

You have altered an index supporting a constraint to be invisible on a large data warehouse table in an X5 Database Machine.

Which two statements are true?

- A. You might retain the index, and leave it as invisible, and the constraint will still be recognized and enforced.
- B. You must retain the index and make it visible again for the constraint to be recognized and enforced.
- C. You must retain the index and set the constraint to DISABLE NOVALIDATE RELY for the constraint to be recognized.
- D. You might drop the index and use a constraint with the DISABLE NOVALIDATE RELY flags for the constraint to be recognized.
- E. You might drop the index and make the constraint invisible, for the constraint to be recognized and enforced.

Correct Answer: BC

Explanation:

B: With making indexes invisible, we can easily check whether indexes are useful without having to drop (and in case recreate) them actually. While this may be of interest for "ordinary" Oracle Databases already, it is particular a useful feature for Exadata where we expect some conventional indexes to become obsolete after a migration.

C: DISABLE NOVALIDATE RELY means: "I don't want an index and constraint checking to slow down my batch data loading into datawarehouse, but the optimizer can RELY on my data loading routine and assume this constraint is enforced by other mechanism". This information can greatly help optimizer to use correct materialized view when rewriting queries. So if you don't use materialized views for query rewrite then you can put RELY for all your constraints (or NORELY for all your constraints) and forget about it.

QUESTION 5

Which two statements are true about RMAN incremental level-1 backups for a database running on an X5 Database Machine?

- A. Block Change tracking (BCT) is not needed for databases in an X5 Database Machine, because ASM automatically does block change tracking when diskgroups are built from griddisks.
- B. cellsrv filters blocks and returns only those that have changed since the last backup, for databases in archivelogmode.
- C. cellsrv returns all blocks that have been or might have been changed since the last backup, and RMAN filters the blocks that do not require backup.
- D. cellsrv filters blocks and returns only those that have changed since the last backup, for databases in noarchivelogmode.
- E. Block Change Tracking (BCT) is not needed for databases in an X5 Database Machine, because change tracking is done by cellsrv.



Correct Answer: DE

Explanation:

D: Whenever data blocks change, the Change Tracking Writer (CTWR) background process tracks the changed blocks in a private area of memory.

When RMAN performs an incremental backup on the Exadata platform, cellsrv filters out unwanted blocks and sends back only those that have changed since the last level 0 or level 1 backup.

E: For the Exadata platform, you may choose to allow cellserv to do all the blocking filtering for incremental backups.

Note: RMAN block change tracking is used to improve the performance of incremental backup. The level 0 incremental backup scans the entire database but level 1 incremental backups use the block change tracking file to scan only the blocks that have changed since the last backup. This significantly reduces the amount of reads that are required on the database.

References: Expert Oracle Exadata (29 Aug 2015), page 319

[1Z0-070 VCE Dumps](#)

[1Z0-070 Exam Questions](#)

[1Z0-070 Braindumps](#)