



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.passapply.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

You are patching your Exadata X6 Database Machine by applying a new image to the Storage Servers in a rolling fashion.

Your ASM environment on the Database Machine has five diskgroups stored on an unpartitioned Exadata storage grid, with redundancy settings as shown:

1.
DATA_DBM1 – Normal Redundancy
2.
RECO_DBM1 – Normal Redundancy
3.
DBFS_DG – Normal Redundancy
4.
DATA2_DBM1 – High Redundancy
5.
DATA3_DBM1 High Redundancy

Which two diskgroups will not suffer from any data loss throughout the patching process even if there is a single disk failure on one of the cells

- A. DBFS_DG
- B. DATA3_DBM1
- C. DATA2_DBM1
- D. DATA_DBM1
- E. RECO_DBM1

Correct Answer: DE

Explanation:

HIGH redundancy provides protection against 2 simultaneous disk failures from 2 distinct storage servers or 2 entire storage servers. HIGH redundancy provides redundancy during Exadata storage server rolling upgrades.

References: <http://blog.umairmansoob.com/choosing-high-vs-normal-redundancy-with-exadata/>



QUESTION 2

Which type or types of network traffic are transported over the internal InfiniBand interconnect in Exadata Database Machine X5?

- A. IDB protocol traffic, clustered ASM traffic, and clustered database instance traffic
- B. IDB protocol traffic and clustered database instance traffic only
- C. IDB protocol traffic and clustered ASM instance traffic only
- D. Both clustered ASM and clustered database instance traffic only
- E. IDB protocol traffic only

Correct Answer: E

Explanation:

Oracle Exadata uses the Intelligent Database protocol (iDB) to transfer data between Database Node and Storage Cell Node.

iDB is used to ship SQL operations down to the Exadata cells for execution and to return query result sets to the database kernel.

QUESTION 3

Which three factors should you consider when choosing a method for migrating a database to the X5 Database Machine?

- A. The down time allowed for the migration
- B. Endian format of the source database
- C. Number of tablespaces in the source database
- D. The type of database workloads
- E. Size of the source database

Correct Answer: BDE

Explanation:

B: Endian format: Non-Exadata format

D: Use Real Production Workload

– Real Application Testing (RAT)

1.

Database Workload Replay



2.

SQL Performance Analyzer (SPA)

Reference: <http://www.oracle.com/technetwork/database/availability/xmigration-11-133466.pdf>

Reference: <http://www.oracle.com/technetwork/database/availability/xmigration-11-133466.pdf>

Exadata MAA Best Practices, Migrating Oracle Databases

QUESTION 4

Which two network switch failure scenarios on a standalone Exadata Database Machine X5-2 Full Rack will affect database availability?

- A. failure of the Ethernet switch only
- B. failure of both the InfiniBand leaf switches
- C. failure of a single InfiniBand leaf switch and the Ethernet switch
- D. failure of a single InfiniBand leaf switch
- E. failure of both the InfiniBand leaf switches and the Ethernet switch

Correct Answer: BE

Explanation:

Ethernet switch for administrative connectivity to servers in the Database Machine

Connect any combination of up to 18 Exadata Database Machine racks or Exadata Storage Expansion

Racks via the InfiniBand fabric. Larger Configurations can be built with external InfiniBand switches.

Connected racks can be any combination of v2, X2, X3 or X4 generation hardware.

Reference: <https://www.oracle.com/technetwork/database/exadata/exadata-storage-expansion-x5-2-ds-2406252.pdf>

Reference: <https://www.oracle.com/technetwork/database/exadata/exadata-storage-expansion-x5-2-ds-2406252.pdf>

QUESTION 5

You are evaluating the performance of a SQL statement that accesses a very large table. You run this query:



```
SQL> SELECT s.name, m.value/1024/1024 MB FROM V$SYSSTAT s, V$MYSTAT m
2 WHERE s.statistic# = m.statistic# AND
3 (s.name LIKE 'physical % total bytes' OR s.name LIKE 'cell phys%'
4 OR s.name LIKE 'CELL 10%');
```

NAME	MB
physical read total bytes	19047.2266
physical write total bytes	0
cell physical IO interconnect bytes	4808.85828
cell physical IO bytes pushed back due to excessive CPU on cell	0
cell physical IO bytes saved during optimized file creation	0
cell physical IO bytes saved during optimized RMAN file restore	0
cell physical IO bytes eligible for predicate offload	18005.6953
cell physical IO bytes saved by storage index	0
cell physical IO interconnect bytes returned by smart scan	3767.32703
cell IO uncompressed bytes	18005.6953

Identify two reasons why the “cell physical IO interconnect bytes” statistic is greater than the “cell physical IO interconnect bytes returned by smart scan” statistic.

- A. There is a transaction, which committed after the query began, that has modified some of the table blocks, causing some “cell single block physical reads”to be requested by the database instance, resulting in additional I/O.
- B. There are chained rows in the table, causing some “cell single block physical reads”to be requested by the database instance, resulting in additional I/O.
- C. The table is a hash clustered table, causing “cell multiblock physical reads”to be requested by the database instance, resulting in additional I/O.
- D. The table is list partitioned, causing “cell list of blocks physical reads”to be requested by the database instance, resulting in additional I/O.
- E. There is a local index on a list partitioned table on the column used in the WHERE clause, causing “cell list of blocks physical reads”to be requested by the database instance, resulting in additional I/O.

Correct Answer: CD

Explanation:

C: Scan on a clustered table can prevent a Smart Scan from occur.

D: Scan on an index-organized table can prevent a Smart Scan from occur.

Note: The Cell physical IO interconnect bytes returned by smart scan metric shows how many bytes of I/O were returned by a smart scan to the database server.

References: <https://uhesse.com/2011/01/19/exadata-part-i-smart-scan/>