



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

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 two are true about the allocation of I/O resources by IORM within the CELLSRV process?

- A. Database Writer I/O is always prioritized over all user I/O.
- B. If two consumer groups P and Q in the PROD database get 20% and 10% respectively of resource allocation, then the percentage of I/O resource would be the same if they got 10% and 5% respectively, and the interdatabase plan has changed, provided that the category plan is unchanged, and consumer groups P and Q are still mapped to the same categories.
- C. If two consumer groups C and D in the PROD database get 10% and 15% respectively, of resource allocation, then the percentage of I/O resource would be the same if they still got 10% and 15% respectively, and were remapped to a different category by the DBA, provided that the category plans and the interdatabase plans are unchanged.
- D. If two consumer groups A and B in the PROD database get 10% and 15% respectively, of resource allocation, then the percentage of I/O resource would be the same if they got 20% and 30% respectively, provided that the category plans and interdatabase plans are unchanged, and consumer groups A and B are still mapped to the same category.
- E. Log Writer I/O and Control File I/O are always prioritized over all user I/O.

Correct Answer: BE

Explanation:

B: Rules in an interdatabase resource plan specify allocations to databases, not consumer groups.

E: Redo and control file writes always take precedence.

Reference: Using IORM with Exadata

QUESTION 3

Which two Exadata X5 security features would you configure to control which databases can access which griddisks, when multiple databases share space on your storage servers in an unpartitioned storage grid?

- A. Using EXADCLI instead of CELLCLI to create the griddisks
- B. Exadata storage realms using database-scoped security mode
- C. file permissions on the griddisks in each storage server
- D. file permissions on the griddisks in each database server
- E. Exadata storage realms using ASM-scoped security mode

Correct Answer: BE

Reference:



<https://docs.oracle.com/en/engineered-systems/exadata-database-machine/dbmsq/exadata-security-features.html#GUID-F53D9493-5927-4106-8D86-65D759419E46>

<https://docs.oracle.com/en/engineered-systems/exadata-database-machine/dbmsq/security-guide-exadata-database-machine.pdf>

<https://docs.oracle.com/en/engineered-systems/exadata-database-machine/dbmsq/exadata-security-features.html#GUID-F53D9493-5927-4106-8D86-65D759419E46> <https://docs.oracle.com/en/engineered-systems/exadata-database-machine/dbmsq/security-guide-exadata-database-machine.pdf>

QUESTION 4

Which two are benefits of an active/inactive configured InfiniBand network on Exadata Database Machine X5?

- A. Improved performance for Oracle Network traffic
- B. Improved reliability for Cache Fusion RAC network traffic
- C. Improved reliability when executing Distributed Command Line Interface (DCLI) to run CELLCLI commands
- D. Improved performance for ASM rebalance network traffic
- E. Improved performance when executing Distributed Command Line Interface (DCLI) to run CELLCLI commands

Correct Answer: BC

Explanation:

Active-passive bonding provides reliability through failover.

QUESTION 5

Your X6 Exadata Database Machine is running Oracle Database 12c, and has a large database with some very large tables supporting OLTP workloads.

High-volume insert applications and high-volume update workloads access the same tables.

You wish to compress these tables without causing unacceptable performance overheads to the OLTP workload.

Which three are true regarding this requirement?

- A. Compression is performed on database servers when using row store compress advanced in an Exadata environment.
- B. Using row store compress advanced will compress the data more than if using column store compress for archive low.
- C. The compression method column store compress for archive high is the worst fit for this requirement.
- D. Compression is performed on Exadata Storage Servers when using row store compress advanced in an Exadata environment.



E. Using row store compress advanced will compress the data less than if using column store compress for query low.

Correct Answer: ABD

Explanation:

A: Creating a Table with Advanced Row Compression

The following example enables advanced row compression on the table orders:

```
CREATE TABLE orders ... ROW STORE COMPRESS ADVANCED;
```

B: ARCHIVE LOW compression (Exadata only), recommended for Archival Data with Load Time as a critical factor

Incorrect Answers:

E: QUERY LOW compression (Exadata only), recommended for Data Warehouse with Load Time as a critical factor.

References: <https://docs.oracle.com/database/121/ADMIN/tables.htm#ADMIN015>

<https://uhesse.com/2011/01/21/exadata-part-iii-compression/>

[1Z0-070 PDF Dumps](#)

[1Z0-070 VCE Dumps](#)

[1Z0-070 Braindumps](#)