



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.pass4itsure.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 plan to partition the database and storage grids in an X5-2 full rack, creating two clusters and two storage grids without using virtualization.

One cluster will be used for production and should consist of six database servers and nine storage servers.

The other cluster will be used for test and development, and should consist of two database servers and five storage servers.

The storage must be partitioned so that the storage servers are visible only to the appropriate database servers that are meant to access them.

What must be done to achieve this?

- A. Configure Exadata realms using Oracle ASM-scoped security mode.
- B. Configure Exadata realms using Database-scoped security mode.
- C. Edit the CELLIP.ORA file on each database server to contain IP addresses of storage servers in the storage grid associated with the cluster to which that database server belongs.
- D. Edit the CELLINIT.ORA file on each database server to contain IP addresses of storage servers in the storage grid associated with the cluster to which that database server belongs.
- E. Edit the CELLIP.ORA file on each database server to contain IP addresses of database servers which are allowed access to specific storage servers in the same storage grid.
- F. Edit the CELLIP.ORA file on each storage server to contain IP addresses of database servers in the database server grid that are associated with the storage grid to which that storage server belongs.

Correct Answer: C

Explanation:

The cellinit.ora file is host-specific, and contains all database IP addresses that connect to the storage network used by Oracle Exadata Storage Servers. This file must exist for each database that connect to Oracle Exadata Storage Servers.

The cellinit.ora file contains the database IP addresses.

The cellip.ora file contains the storage cell IP addresses.

Both files are located on the database server host.

References: http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-serverconfiguring.htm#SAGUG20369

QUESTION 2

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 archiveologmode.
- 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 noarchiveologmode.
- 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

QUESTION 3

Which two are true concerning Columnar Flash Caching?

- A. It can be enabled or disabled for individual databases by using I/O Resource Manager database plans.
- B. It automatically transforms ROW STORE COMPRESS compressed data into a pure columnar format during Exadata Smart Flash Cache population.
- C. Data is either cached in Hybrid Columnar Compressed format or in pure compressed format, but never in both formats at the same time.
- D. It is enabled by default.
- E. It automatically transforms Hybrid Columnar Compressed (HCC) data into a pure columnar format during Exadata Smart Flash Cache population.
- F. It improves single-row lookup performance.



Correct Answer: DE

Explanation:

D: In-Memory Columnar Caching on cells is enabled by default when the INMEMORY_SIZE is configured. You do not need to do anything to get this enhancement.

E: Columnar Flash Caching implements a dual format architecture in Exadata flash by automatically transforming frequently scanned Hybrid Columnar Compressed data into a pure columnar format as it is loaded into the flash cache.

Incorrect Answers:

F: Columnar Flash Caching accelerates reporting and analytic queries while maintaining excellent performance for OLTP style single row lookups.

References:

http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-server-monitoring.htm#SAGUG20883

<http://www.oracle.com/technetwork/database/exadata/exadata-x5-8-ds-2745934.pdf>

QUESTION 4

You plan to consolidate multiple mission-critical databases onto a single rack Exadata X6 database Machine.

You do not plan to use virtualization on the database servers.

You have identified two database categories, based on differing availability and maintenance objectives.

Which two choices will be the best solution?

- A. Create a single database cluster that accesses a single storage grid.
- B. Create a single database cluster that accesses multiple storage grids.
- C. Partition the Exadata storage servers into two separate storage grids.
- D. Create a single storage grid and share space on the cell disks using Exadata storage realms.
- E. Partition the database servers into two separate clusters.

Correct Answer: BE

QUESTION 5

Which three are true concerning Exadata snapshot databases?

- A. They are supported on non-container databases.
- B. They are based on a read-write copy of an existing database.
- C. They are integrated with the Multitenant architecture.



D. They can be created only on sparse ASM disk groups.

E. They don't support all Exadata features.

F. They can be created on any type of ASM disk group.

Correct Answer: ACD

Explanation:

A: An Exadata snapshot database can be either a non-container database (non-CDB) or a container database (CDB). Creating an Exadata snapshot database of a CDB enables access to all of the pluggable databases in that container.

C: You can create two types of Exadata snapshots, depending on the current setup of your environment:

1.

You have a pluggable database (PDB) and want to create a test master from it.

2.

You have a container database (CDB) and want to create test masters from all its PDBs, or you have a simple non-container database and want to create a test master from it.

D: SPARSE disk group based database snapshots is functionality included in Exadata Storage Software. It requires Exadata Storage Software version 12.1.2.1.0 and Oracle Database 12c version 12.1.0.2 with bundle patch 5 or later. This feature is designed to work on native Exadata ASM storage disk groups. It uses ASM SPARSE grid disk based thin provisioning where snapshot databases created on a SPARSE disk group need only the space for changes plus some metadata, thereby enabling storage efficient snapshot databases.

References:

http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-server-snapshots.htm#SAGUG-GUIDE1D6EF45-36EF-40E3-A57E-F80B749E6122

<http://www.oracle.com/technetwork/database/exadata/learnmore/exadata-database-copy-twp-2543083.pdf>

[Latest 1Z0-070 Dumps](#)

[1Z0-070 PDF Dumps](#)

[1Z0-070 Braindumps](#)