

**100%** Money Back  
**Guarantee**

**Vendor:**Oracle

**Exam Code:**1Z0-027

**Exam Name:**Oracle Exadata X3 and X4  
Administration

**Version:**Demo

## QUESTION 1

You plan to migrate a database supporting an OLTP workload to your Database Machine

This is part of a consolidation project and several other databases already exist on the Database Machine.

Which three Exadata features may help to improve the performance of this OLTP workload?

- A. Hybrid Columnar Compression
- B. I/O Resource Manager
- C. Smart Flash Cache
- D. Smart Flash Log
- E. Smart Scan
- F. Storage Index

Correct Answer: BCD

Explanation: C: OLTP performance benefits with Exadata Smart flash log for low latency commits / (C)Smart flash cache for low latency reads. KEEP in Flash for critical objects

---

## QUESTION 2

Your Database Machine 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 decide to compress these tables without causing unacceptable performance overheads to the OLTP application.

Which three are true regarding this requirement?

- A. Using `compress for oltp` will compress the data less than if using Hybrid Columnar Compression when specified with `compress for query low`.
- B. The compression is performed on the storage servers when using `compress for oltp` in an Exadata environment.
- C. The compression method `compress for archive high` is the worst fit for this requirement.
- D. Using `compress for oltp` will compress the data more than if using Hybrid Columnar Compression when specified with `compress for archive low`.
- E. The compression is performed on the database servers when using `compress for oltp` in an Exadata environment.

Correct Answer: ACE

Note:

(E not B):

\*Types of compression

Basic compression

OLTP compression

Warehouse compression

Online archival compressio

\*

/OLTP compression allows compression during DML operations. /Basic compression works at the data block level.

\*When you enable table compression by specifying COMPRESS FOR OLTP, you enable OLTP table compression. Oracle Database compresses data during all DML operations on the table. This form of compression is recommended for

OLTP environments.

\* When you specify COMPRESS FOR QUERY or COMPRESS FOR ARCHIVE, you enable hybrid columnar compression. With hybrid columnar compression, data can be compressed during bulk load operations. During the load process, data is transformed into a column-oriented format and then compressed. Oracle Database uses a compression algorithm appropriate for the level you specify. In general, the higher the level, the greater the compression ratio.

Hybrid columnar compression can result in higher compression ratios, at a greater CPU cost. Therefore, this form of compression is recommended for data that is not frequently updated.

---

### QUESTION 3

Which two may act as building blocks for the creation of two or more griddisks residing on the same physical device and which have more balanced performance characteristics?

A. griddisk on hard drive-based interleaved celldisks.

B. griddisk on hard drive-based celldisks consumed by ASM when creating diskgroups using ASM-based intelligent Data Placement (IDP)

C. griddisk on hard drive-based non-interleaved celldisks.

D. griddisk on Flash-base celldisks.

Correct Answer: AB

Explanation: A (not C):Interleaved Griddisks mean equally fast Griddisks, whereas with non Interleaved Griddisks (the default) the Griddisks created first are being taken from the outer tracks of the underlying physical disk.

B: Intelligent Data Placement, a feature of ASM that allows placing data in such a way that more frequently accessed data is located close to the periphery of the disk where the access is faster.

---

### QUESTION 4

Which three are true about Enterprise Manager plug- in configuration for the Database Machine?

- A. There are several separate plug-ins for Grid Control.
- B. There is one plug-in for Cloud Control.
- C. There are several separate plug-ins for Cloud Control.
- D. Some plug-ins require SNMP trap forwarders.
- E. All plug-ins require SNMP trap forwarders.
- F. There is one plug-in for Grid Control.

Correct Answer: ABD

Explanation: A:Oracle Enterprise manager Grid control, is hands down the best monitoring and management tool, for the oracle exadata database machine. It comes with plugins to monitor all the hardware components of the database machine, and sensible, preset thresholds for proactive monitoring.

---

#### QUESTION 5

Which three must be true for Smart Scans to be done?

- A. Executing a query in parallel
- B. Setting `_serial_direct_read=true` in the session issuing the SQL statements
- C. Having direct path reads used at run time
- D. Having a 4 megAUsize for the ASM diskgroup containing the tablespace in which tables accessed by a query are stored
- E. `Cell_offload_process = true` for the ASM diskgroup containing the tablespace in which tables accessed by a query are stored.
- F. `cell.smart_scan_capable=true` for the ASM diskgroup containing the tablespace in which tables accessed by a query are stored.

Correct Answer: ACF

F:ASM Diskgroup has an attribute: `cell.smart_scan_capable` ?Must be set to TRUE for Smart Scans to work

---

#### QUESTION 6

Which two statements are true about the use of direct path loads when selecting from external tables in a database on a Database Machine?

- A. INSERT INTO . . . SELECT FROM statements, executed serially, which select from external tables, require the APPEND hint to use direct path loading.
- B. CREATE TABLE . . . AS SELECT statements, which select from external tables, attempt to use in direct path loading automatically.
- C. CREATE TABLE . . . AS SELECT statements, which select from external tables, require the APPEND hint to use

direct path loading.

D. INSERT INTO . . . SELECT FROM statements, executed serially, which select from external tables, are unable to use direct path loading.

Correct Answer: AB

Explanation: A CTAS(Create table as select)will always use direct path(B, not C)load but IAS(Insert as select)statement will not. In order to achieve direct path load with an IAS statement you must add the APPEND hint to the command(A, not

D).

Direct path loads can also run in parallel. You can set the parallel degree for a direct path load

either by adding the PARALLEL hint to the CTAS or IAS statement or by setting the PARALLEL clause on both the external table and the table into which the data will be loaded.

Once the parallel degree has been set at CTAS will automatically do direct path load in parallel

but an IAS will not. In order to enable an IAS to do direct path load in parallel you must alter the session to enable parallel DML.

Note:

\*Parallel Direct Path Load

The key to good load performance is to use direct path loads wherever possible. A direct path

load parses the input data according to the description given in the external table definition, converts the data for each input field to its corresponding Oracle data type, then builds a column

array structure for the data. These column array structures are used to format Oracle data blocks

and build index keys. The newly formatted database blocks are then written directly to the database, bypassing the standard SQL processing engine and the database buffer cache.

Reference:Best Practices for Implementing a Data Warehouse on the Oracle Exadata Database Machine;Using CTAS and Exchange Partition Replace IAS for Copying Partition on Exadata

---

## QUESTION 7

Your customer wants you to partition the database and storage grids in his X3-2 full rack, creating database clusters and two storage grids.

One cluster will be used for production and should consist of 6 database servers and 11 cells from the first storage grid.

The other cluster will be used for test and development, and should consist of 2 database servers and 3 cells from the second storage grid.

The storage must be partitioned so that the cells are visible only to the appropriate database servers based on the description above.

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 cells in the storage grid associated with cluster to which that database server belongs.
- D. Edit the CELLINIT.ORA file on each database server to contain IP addresses of cells 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 cells in the same storage grid.
- F. Edit the CELLIP.ORA file on each cell to contain IP addresses of database servers in the database server grid that are associated with the storage grid to which that cell belongs.

Correct Answer: C

Explanation: cellip.ora

The cellip.ora is the configuration file, on every compute node, that tells ASM instances which cells are available to this cluster.

Here is a content of a typical cellip.ora file for a quarter rack system:

```
$ cat /etc/oracle/cell/network-config/cellip.ora
```

```
cell="192.168.10.3"
```

```
cell="192.168.10.4"
```

```
cell="192.168.10.5"
```

Now that we see what is in the cellip.ora, the grid disk path, in the examples above, should make more sense.

Note:

\*cellinit.ora decides which network takes storage traffic. \*cellip.ora - list of cells, new cells can be added dynamically without shutdown

## QUESTION 8

You are examining the existing IORM configuration on the cells of Database Machine, to see if they require any modifications based on recent changes to various workloads. All seven cells in your X3-2 half-rack show the following:

```
CellCLI> list iormplan detail

name:          dmorlcel04_IORMPLAN
catPlan:       name=interactive, level=1, allocation=90
               name=batch, level=2, allocation=80
               name=maintenance, level=3, allocation=50
               name=other, level=3, allocation=50
dbPlan:        name=sales, level=1, allocation=45, flashcache=on, flashlog=on
               name=finance, level=1, allocation=45, flashcache=on, flashlog=off
               name=other, level=1, allocation=10, flashcache=off, flashlog=off
objective:     off
status:        active
```

Which two are true about I/O to the cells using this plan?

- A. I/O requests in the batch category may use flashcache if the I/O is from the sales finance database, and these I/O requests are guaranteed to get 80% of the I/O if the interactive category I/Os use no more than 20%.
- B. I/O requests made by sessions in the marketing database may use flashing and flashcache if no other categories or database or database are using flashing and flashcache at the same time.
- C. I/O requested in the interactive category may use flashdns if the I/O is from the sales or finance databases, and these I/O requests are guaranteed to get 90% of the I/O if the enough I/Os are issued in this category.
- D. I/O requests from the sales database may use flashing regardless of the I/O category.
- E. No I/Os in any category or from any database may use flashing or flashcache because the objective is off.

Correct Answer: DE

Explanation: list iormplan detail

Note:

\*The I/O Resource Manager (IORM) extends the concept of resource groups with a new attribute known as a category. While resource groups allow DBRM to manage resources within a a database, categories provide I/O resource management among multiple databases.

```
*cellcli> alter iormplan objective=\\'balanced\\'
```

```
-- {balanced | off | low_latency | high_throughput | auto
```

"off" simply turns off the IORM plan\\'s IO metering.

---

## QUESTION 9

You must drop all celldisks on all the cells in a quarter rack as part of a reconfiguration project, to support normal redundancy interleaving.

Select two actions that describe the operating system (O/S) account on the cells to which you should log in, and the tool that may be used to drop the celldisks.

- A. To the CELLMONITOR account using CELLCLI interactively on each cell.
- B. To CELLADMIN account calling CELLCLI on all cells using DCLI
- C. To either the CELIMONITOR or the CELLADMIN account calling CELLCLI on all cells using DCLI
- D. To the CELLMONITOR account calling CELLCLI on all cells using DCLI
- E. To the CELLADMIN account using CELLCLI interactively on each cell
- F. To either the CELLMONITOR or CELLADMIN account calling interactively on each cell

Correct Answer: BE

Note:

\*Each Exadata Storage cell can be monitored, configured, and maintained using the cellcli command line interface. To invoke the CLI, simply login to the Exadata cell as cellmonitor, celladmin, or root, and type "cellcli". Within cellcli, a wide range of commands can be invoked to monitor the cell.

\*While connected to cellcli as the Linux account celladmin, you can perform all cellcli commands except calibrate.

While logged into the cell server as cellmonitor, you can do any "read-only" commands; you cannot change any configuration or attributes.

---

## QUESTION 10

You recently upgraded your Exadata image to the latest release; previously you were using 11.2.0.3.

At the same time, you decide to address some performance problems as follows:

You noticed increased latency for the database log writer, especially during the quarterly battery learn cycle on the cells.

You have complaints of erratic performance from certain write-intensive applications.

Which two actions could improve performance in these areas?

- A. Enable write-back flashcache by setting lunWriteCacheMode to Write Back Mode.
- B. Use ALTER TABLE in the database to set CELL\_FLASH\_CACHE = KEEP for the tables belonging to the affected application.
- C. Configure Smart Flash Log on the cells to use some of these of the space on the cell flash devices.
- D. Configure the table belonging to the affected application using CELLCLI, to the set CELL\_FLASH\_CACHE = KEEP.
- E. Configure Smart Flash Log on the database server to use server flash memory.

Correct Answer: BC

Explanation: B: The following command could be used to pin the table CUSTOMERS in Exadata Smart Flash Cache  
ALTER TABLE customers STORAGE (CELL\_FLASH\_CACHE KEEP)

C:Creating Flash Disks Out Of The Flash Cache When an Exadata cell is installed, by default, all the flash is assigned to be used as flash cache and user data is automatically cached using the default caching behavior. Optionally, a portion of the cache can be reserved and used as logical flash disks. These flash disks are treated like any Exadata cell disk in the Exadata cell except they actually reside and are stored as non- volatile disks in the cache.

Note: \*Pinning Objects In The Flash Cache Preferential treatment over which database objects are cached is also provided with the Exadata Smart Flash Cache. For example, objects can be pinned in the cache and always be cached, or an object can be identified as one which should never be cached. This control is provided by the new storage clause attribute, CELL\_FLASH\_CACHE, which can be assigned to a database table, index, partition and LOB column

\*There are two techniques provided to manually use and manage the cache. The first enables the pinning of objects in the flash cache. The second supports the creation of logical disks out of the flash for the permanent placement of objects on flash disks.

---



## QUESTION 11

You are about to run the oplan utility to patch the servers on your test Database Machine before patching the production environment.

The following task might be performed:

- A. Test the failback procedure
- B. Run the exachk utility
- C. Read the README file.
- D. Automate the patch application process as appropriate.
- E. Verify that the patch provides the functionality it is meant to.
- F. Apply the patch.
- G. Evaluate the system performance. In which order should the tasks be performed to patch in the recommended fashion?
- H. C, B, D, F, B, E, A
- I. C, D, F, B, E, G, A, B
- J. C, B, D, F, E, G, A
- K. C, B, D, F, E, A, G
- L. C, B, D, F, B, E, G, A

Correct Answer: B

Note: \*Schedule regular health checks with Exachk ?Detects bad disks, faulty hardware, incorrect configuration ?At least every 90 days ?Before and after patching/upgrades ?Follow recommended best practices

\*we are recommending running exachk or Healthcheck (depending on your DBM/Exadata model) \*before\* maintenance and \*after\* maintenance.

---

## QUESTION 12

Yesterday, an Exadata storage server flashdisk entered the poor performance state.

This flashdisk os used by flashcache and has a griddisk which is a member of a normal redundancy diskgroup.

Identify the six steps you must perform to replace flashdisk.

- A. Identify the griddisk located on the poorly performing flashdisk and drop it from the associated ASm diskgroup.
- B. Verify that the griddisk located on the poorly performing flashdisk has been successfully dropped from the associated ASM diskgroup.
- C. Drop the flashcache on all cell.

- D. Safely power off the cell containing the poorly-performing flashdisk.
- E. Replace the poorly performing flashdisk.
- F. Power up the cell containing the replaced flashdisk and activate all griddisks.
- G. Recreate the flashcache on the cell using all flashdisks.
- H. Create a new griddisk on the replaced flashdisk.
- I. Add the griddisk back into the ASM diskgroup to which it belonged.

Correct Answer: ADEFHI

Answer: A,D,E,F,H,I

Note:

\*The default way to deal with the Flash Storage is to use it completely as Flash Cache. You may think of Flash Cache as a prolongation of the Database Buffer Cache. It is populated automatically by the system with objects deemed useful to

cache them.

\*The second possibility to deal with the Flash Storage is to take a part of it for building ASM diskgroups upon. All files on these ASM diskgroups will then reside permanently on Flash Storage:

```
CellCLI> drop flashcache
```

```
Flash cache exa5cel01_FLASHCACHE successfully dropped
```

```
CellCLI> create flashcache all size=100g
```

```
Flash cache exa5cel01_FLASHCACHE successfully created
```

```
CellCLI> create griddisk all flashdisk prefix=flashdrive
```

To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

## Try our product !

100% Guaranteed Success

100% Money Back Guarantee

365 Days Free Update

Instant Download After Purchase

24x7 Customer Support

Average 99.9% Success Rate

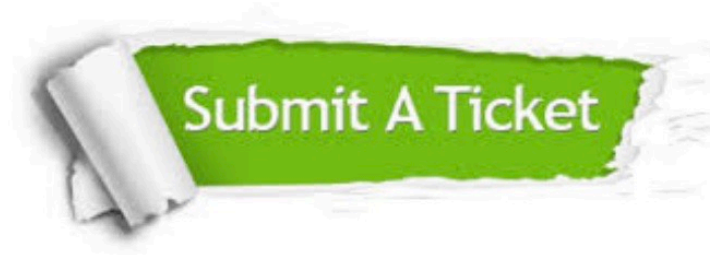
More than 800,000 Satisfied Customers Worldwide

Multi-Platform capabilities - [Windows](#), [Mac](#), [Android](#), [iPhone](#), [iPod](#), [iPad](#), [Kindle](#)

## Need Help

Please provide as much detail as possible so we can best assist you.

To update a previously submitted ticket:



 <p><b>One Year Free Update</b> Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>	 <p><b>Money Back Guarantee</b> To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 <p><b>Security &amp; Privacy</b> We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information &amp; peace of mind.</p>
---	---	--

Any charges made through this site will appear as Global Simulators Limited.

All trademarks are the property of their respective owners.