



1Z0-027^{Q&As}

Oracle Exadata X3 and X4 Administration

Pass Oracle 1Z0-027 Exam with 100% Guarantee

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

<https://www.pass4lead.com/1Z0-027.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 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 2

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 3

Which two activities are supported on the storage servers in the Database Machine?

- A. Installing an alternative package manager
- B. configuring secure shell user equivalency for the callmonitor user
- C. changing root password
- D. upgrading the Storage Server software package using RPN



E. upgrading a device driver for hard disks when inserting a replacement disk after a hard disk failure

Correct Answer: AD

Explanation: D:What Oracle Solaris Brings to Oracle Exadata Database Machine Simplified, improved updates--Oracle Solaris 11 uses the new Image Package System (IPS), which is a network-based package management that provides a framework for complete software lifecycle management such as installation, upgrade and removal of software packages

QUESTION 4

Which three are true about Smart Flash log?

- A. I/O Resource Manager database plans can be used to enable or disable Smart Flash Log for different databases.
- B. LGWR will not wait for writes to Smart Flash log if the write to a disk based logfile completes first.
- C. Smart Flash Log is enabled by default, using 1024 MB of Flash storage on each storage server.
- D. You can remove Smart Flash Log from a single storage server with the drop flashing command.
- E. I/O Manager category plans can be used to enable or disable Smart Flash Log for different I/O categories.

Correct Answer: ABD

Explanation: A:the Exadata I/O Resource Manager (IORM) has been enhanced to enable or disable Smart Flash Logging for the different databases running on the Database Machine.

B: Smart Flash Logging works as follows. When receiving a redo log write request, Exadata will do parallel writes to the on-disk redo logs as well as a small amount of space reserved in the flash hardware. When either of these writes has successfully completed the database will be immediately notified of completion. If the disk drives hosting the logs experience slow response times, then the Exadata Smart Flash Cache will provide a faster log write response time. Conversely, if the Exadata Smart Flash Cache is temporarily experiencing slow response times (e.g., due to wear leveling algorithms), then the disk drive will provide a faster response time. This algorithm will significantly smooth out redo write response times and provide overall better database performance.

D: Category plans are configured and enabled using the CellCLI utility on the cell. Only one category plan can be enabled at a time

Incorrect:

Not C:By default, 512 MB of the Exadata flash is allocated to Smart Flash Logging

QUESTION 5

You have altered an index supporting a constraint to be invisible on a large read only data warehouse table, to determine if Smart Scan operations will be fast enough to satisfy your performance requirements.

Given the results of your testing, you consider dropping the index.

Which two statements are true?

- A. You must retain the index and set the constraint to DISABLE NOVALIDATE RELY to enforce the constraint



- B. You may drop the index and use a constraint with the DISABLE NOVALIDATE RELY flags
- C. You must retain the index and make it visible again for the constraint to be enforced.
- D. You may drop the index and make the constraint invisible, because this is enough for the constraint to be enforced.
- E. You may retain the index, and leave it as invisible, because this is enough for the constraint to be recognized.

Correct Answer: AC

Note:

*You may have noticed that we introduced Invisible Indexes as an 11g New Feature. Their main benefit is that we can test whether performance differs if we would drop an index without actually dropping it. This is particular useful after an

Exadata Migration because we expect that some conventional indexes migrated are now obsolete and may be substituted by Storage Indexes.

*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.

*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 6

To troubleshoot a possible hardware problem, you consider moving all disk drives from one Exadata storage server to a replacement chassis.

You must contain storage availability while performing task.

The Exadata storage server is an X3-8 Database Machine and storage grid is not partitioned.

Which two factors would prevent you from moving the disks from one Exadata storage server to another one?

- A. The existence of an external redundancy ASM diskgroup
- B. The existence of a normal redundancy ASM diskgroup
- C. The existence of an ASM diskgroup with the repair_time attribute set to 0.
- D. The existence of an ASM diskgroup with its compatible.asm attribute set to 10.2.0.0
- E. Offline or inactive celldisks in another Exadata server



Correct Answer: AD

Explanation: A:If you want Oracle ASM to mirror files, specify the redundancy level as NORMAL REDUNDANCY (2-way mirroring by default for most file types) or HIGH REDUNDANCY (3-way mirroring for all files). You specify EXTERNAL

REDUNDANCY if you do not want mirroring by Oracle ASM. For example, you might choose EXTERNAL REDUNDANCY if you want to use storage array protection features.

D:Restoring the redundancy of an Oracle ASM disk group after a transient disk path failure can be time consuming. This is especially true if the recovery process requires rebuilding an entire Oracle ASM failure group. Oracle ASM fast mirror

resync significantly reduces the time to resynchronize a failed disk in such situations. When you replace the failed disk, Oracle ASM can quickly resynchronize the Oracle ASM disk extents.

To use this feature, the disk group compatibility attributes must be set to 11.1 or higher.

Incorrect:

Not C:You can set the DISK_REPAIR_TIME disk group attribute to delay the drop operation by specifying a time interval to repair the disk and bring it back online.

Note:

*The redundancy levels are:

/External redundancy

Oracle ASM does not provide mirroring redundancy and relies on the storage system to provide RAID functionality. Any write error cause a forced dismount of the disk group. All disks must be located to successfully mount the disk group.

/Normal redundancy

Oracle ASM provides two-way mirroring by default, which means that all files are mirrored so that there are two copies of every extent. A loss of one Oracle ASM disk is tolerated. You can optionally choose three-way or unprotected mirroring.

/High redundancy

Oracle ASM provides triple mirroring by default. A loss of two Oracle ASM disks in different failure groups is tolerated.

Reference:Administering Oracle ASM Disk Groups

QUESTION 7

You plan to migrate an Oracle database that supports an online transaction processing (OLTP) workload to your Database Machine.

Following are details for the source database:

Database version: 10.2.0

Byte order: Big Endian HP-UX (64-bit)

Database size: 24 TB



Storage: ASM with 1 MB allocation unit size

Which two are supported migration methods?

- A. Physical migration using ASM online Migration
- B. Physical migration using Transportable Database
- C. Logical migration using Oracle Streams
- D. Local migration using Oracle Streams
- E. Logical migration using logical standby

Correct Answer: BC

Note:

* physical migration techniques (e.g. Data Guard Physical Standby, Transportable Database, Transportable Tablespaces(B), Rapid Clone) vs. logical migration techniques(C)(e.g. Oracle Data Pump).

*You can use either Data Pump export/import or original export/import for Streams instantiations.

*The Oracle Exadata Database Machine runs the standard Oracle Database. Therefore, any application that runs with the Oracle Database today can be seamlessly migrated to the Exadata Database Machine with no changes to the application.

*There are several techniques for migrating data to a Database Machine. Migration can be done using Oracle Recovery Manager (RMAN) to backup from traditional storage and restore the data onto Exadata. Oracle Data Guard can also be used to facilitate a migration. This is done by first creating a standby database based on Exadata storage. The standby can be using Exadata storage and the production database can be on traditional storage. By executing a fast switchover, taking just seconds, you can transform the standby database into the production database. This provides a built-in safety net as you can undo the migration very gracefully if unforeseen issues arise. Transportable Tablespaces(B)and Data Pump may also be used to migrate to Exadata. Any technique used to move data between Oracle Databases can be used with Exadata.

QUESTION 8

Which is true regarding Enterprise Manager monitoring and configuration?

- A. Agents must be deployed on the- storage servers, in case all the database servers are down.
- B. Monitoring the power distribution units (PDUs) is done by Enterprise Manager agents, but additional ethernet cables are required to connect the PDUs to the internal network switch in the Database Machine
- C. Monitoring the power distribution units (PDUs) is done by Enterprise Manager agents without, any additional network configuration.
- D. Agents may be deployed on the storage servers, in case all the database servers are down.
- E. Monitoring the power distribution units (PDUS) is done by Enterprise Manager agents but additional ethernet cables are required to connect the PDUs to the corporate network switch.

Correct Answer: C



Note:

*The following prerequisites must be met before you can deploy the plug-in:

Set and validate the preferred credentials on all Agents where you want to deploy the plug-in. This is essential in order to deploy the plug-in.

The PDU firmware version must be 1.02 or later.

*This plug-in will be used to monitor PDU's actual current value of Phase1, Phase2, and Phase3 and raises appropriate events if they have crossed the different threshold values. Each PDU will have a plug-in instance added to its own OEM

Grid Control agent.

By deploying the plug-in to your Grid Control environment, you gain the following management features:

Monitor PDU targets.

Raise alerts and violations based on thresholds set on monitoring and configuration data.

Provide rich out-of-box metrics and reports based on the gather

Reference: Oracle Enterprise Manager System Monitoring Plug-In Installation Guide for Exadata Power Distribution Unit

QUESTION 9

Which three are among the software components that constitute the QoS Management framework?

- A. Cluster Health Monitor (CHM)
- B. Cluster Verification Utility (CLUVFY)
- C. O/S Resource Manager
- D. Grid Infrastructure for a standalone server
- E. Grid Infrastructure for a cluster
- F. OC4J cluster resource

Correct Answer: AEF

Explanation: A: In conjunction with Cluster Health Monitor, QoS Management's Memory Guard detects nodes that are at risk of failure due to memory over-commitment. It responds by automatically preventing new connections thus preserving

existing workloads and restores connectivity once the sufficient memory is again available.

F: The ora.oc4j is for the QoS (Quality of Service Management), which is only available on Exadata.

Incorrect:

Not B: The Cluster Verification Utility (CVU) performs system checks in preparation for installation, patch updates, or other system changes. Using CVU ensures that you have completed the required system configuration and



preinstallation

steps so that your Oracle grid infrastructure or Oracle Real Application Clusters (Oracle RAC) installation, update, or patch operation completes successfully.

QUESTION 10

You have used setupem.sh to deploy a plug in for Grid Control.

In which two ways do all the targets supported by the plug-in get configured?

- A. The targets must be configured by an Enterprise Manager administrator manually using Grid Control.
- B. Setupem.sh is used again to configure the targets.
- C. The targets must be configured by an O/S administrator manually.
- D. The targets must be configured by a database administrator manually.
- E. There may be more than one target for each plug-in.

Correct Answer: AE

After running the `setupem.sh` script:

- The Cloud Control OMS and Repository will be installed
 - The Exadata plug-ins are deployed on the OMS
-

QUESTION 11

You are in the planning stage of the network configuration for your Database Machine. The requirements are:

- 1.A
fault-tolerant network, providing higher availability for connections to database instances
- 2.Fault tolerance providing higher availability for connections to perform management functions on the database and storage servers.
- 3.Full monitoring of all Database Machine components using Enterprise Manager

Which three components require external Ethernet network cables to connect your existing network infrastructure to your database machine to satisfy this requirement?

- A.
Database servers
- B.



Exadata storage servers

C.

InfiniBand storage servers.

D.

Power distribution units (PDUs)

E.

Cisco Ethernet switch

Correct Answer: ABC

QUESTION 12

Consider the following sequence of CELLCLI commands and SQL statements:

```
CELLCLI> CREATE CELLDISK ALL HARDDISK
CELLCLI> CREATE GRIDDISK ALL HARDDISK PREFIX=DATA01 SIZE=120G
CELLCLI> CREATE GRIDDISK ALL HARDDISK PREFIX=DATA02 SIZE=120G
SQL> CREATE DISKGROUP DATA02 NORMAL REDUNDANCY DISK 'O/*/DATA02*';
SQL> CREATE DISKGROUP DATA01 NORMAL REDUNDANCY DISK 'O/*/DATA01*';
```

All CELLCLI commands and SQL statements are executed successfully without warnings. Which statement is true about the I/O performance of the DATA01 ASM diskgroups?

- A. The DATA01 diskgroups has better performance because the DATA01 griddisks were created first.
- B. The DATA02 ASM diskgroup has better performance because the DATA02 diskgroups was created first.
- C. The DATA01 diskgroup has better performance because the DATA01 griddisks are created on interleaved cell disks.
- D. The performance of both diskgroups is balanced because they are created on interleaved griddisks.

Correct Answer: A

Explanation: Griddisks are the fourth layer of abstraction, and they will be the Candidate Disks to build your ASM diskgroups from. By default (interleaving=none on the Celldisk layer), the first Griddisk that is created upon a Celldisk is placed

on the outer sectors of the underlying Harddisk. It will have the best performance therefore

Incorrect:

Not D: non-interleaved is the default.

[1Z0-027 VCE Dumps](#)

[1Z0-027 Practice Test](#)

[1Z0-027 Braindumps](#)



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](#)


We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications. You can view Vendor list of All Certification Exams offered:

<https://www.pass4lead.com/allproducts>

Need Help

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

To update a previously submitted ticket:



 <p>One Year Free Update 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>Money Back Guarantee 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>Security & Privacy We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & 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.

Copyright © pass4lead, All Rights Reserved.