



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**

Your database Machine has the exachk utility pre-installed and you decide to use it periodically, to validate the installation against Oracle's recommended best practices.

Which two actions could you take to do this?

- A. Use a cron job on a database node to run it at regular intervals.
- B. Run it once from a database node and it will then perform periodic monitoring automatically.
- C. Use a cron job on each cell to run it at regular intervals.
- D. Run it once on each cell and it will then perform periodic monitoring automatically.
- E. Create a Job in Enterprise Manager to run the exachk utility at regular intervals.

Correct Answer: BE

Note:

*Oracle Exadata Database Machine Cells are added as targets during the database machine discovery workflow and are grouped automatically under the group Exadata Grid.

Based on the Exadata Grid group, a group of Exadata cells can be managed together through Enterprise Manager Cloud Control 12c.

QUESTION 2

What is the benefit of bonding the client access network configuration?

- A. Improved performance
- B. Improved reliability
- C. Both improved performance and reliability
- D. A Single Client Access Name (SCAN)
- E. Improved monitoring

Correct Answer: C

Explanation: The network ports can be connected directly between the components or the ports can be bonded together to form a logical interface. The bonded ports can be configured for load balancing, fault tolerance, and better network utilization

Reference: Oracle Exadata Database Machine -Backup and Recovery Sizing:Tape Backups

QUESTION 3



Which three statements are true about the initial storage configuration after the standard deployment of a new database Machine?

- A. The Data_ and RECO_ ASM diskgroups are built on part of the interleaved griddisks.
- B. The Data_ and RECO_ ASM diskgroups are built on top of the non-interleaved griddisks.
- C. There is a free space available on the hard disks inside the database servers for possible extension of the /u01 file system.
- D. There is no free disk space available on the hard disk inside the database servers for possible extensions of the /u01 file system.
- E. There is free space available on the flashdisk inside the Exadata storage servers for possible use as flash-based griddisks.
- F. There is no free disk available on the flashdisk inside the Exadata storage servers for possible use as flash-based griddisks.

Correct Answer: BCF

Explanation: B(not A): 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. If we follow the recommendations, we will create 3 Diskgroups upon our Griddisks: DATA, RECO and SYSTEMDG.

Note:

*non-rootpartition (/u01).

*Oracle Database files: DATA disk group.

*Flashback log files, archived redo files, and backup files: RECO disk group.

QUESTION 4

Which two are true about Smart Scan?

- A. a query rewrite may occur to a container table stored in Exadata but will never benefit From Smart scan.
- B. Column projection does not contribute to the performance benefit of Smart Scan
- C. It is possible to offload single row functions to the storage servers.
- D. Some joins can be offloaded to the storage servers.
- E. A query rewrite may occur to a container table stored Exadata, and it will always benefit from Smart Scan.
- F. All joins can be offloaded to the storage servers.

Correct Answer: CD

Explanation: C:With Exadata storage, database operations are handled much more efficiently. Queries that perform



table scans can be processed within Exadata storage with only the required subset of data returned to the database server.

Row filtering, column filtering and some join processing (among other functions) are performed within the Exadata storage cells. When this takes place only the relevant and required data is returned to the database server.

D (not F):

*Exadata performs joins between large tables and small lookup tables, a very common scenario for data warehouses with star schemas. Joining large tables and small lookup tables is implemented using Bloom Filters, which are a very

efficient probabilistic method to determine whether a row is a member of the desired result set. *If storage indexes are so great, why doesn't Oracle Exadata use them all the time? The short answer is that they are created and used only when

they will be beneficial. *To use storage indexes, Oracle Exadata queries must use smart scans, so not all types of applications can benefit from storage indexes. Applications with queries that include predicates and perform a lot of full table

scans or fast full scans of indexes--typically those used in data warehousing environments--will benefit greatly from storage indexes. Online transaction processing (OLTP) applications, on the other hand, typically access a small number of

rows through standard indexes and do not perform full table scans, so they may not benefit from storage indexes.

Note:

*Storage indexes reside in the memory of the storage servers--also called storage cells--and significantly reduce unnecessary I/O by excluding irrelevant database blocks in the storage cells.

*To use storage indexes, Oracle Exadata queries must use smart scans, so not all types of applications can benefit from storage indexes.

Incorrect:

Not B: Exadata provides column filtering, also called column projection, for table scans. Only the columns requested are returned to the database server rather than all columns in a table. For example, when the following SQL is issued, only

the employee_name and employee_number columns are returned from Exadata to the database kernel.

SELECT employee_name, employee_number FROM employee_table. For tables with many columns, or columns containing LOBs (Large Objects), the I/O bandwidth saved can be very large. Using both predicate and column filtering

dramatically improves performance and reduces I/O bandwidth consumption. In addition, column filtering also applies to indexes, allowing for even faster query performance.

Reference: Oracle Communications Data Model Implementation and Operations Guide, Exadata Smart Scan Processing and Storage Index

QUESTION 5

Identify the three components that serve a purpose only in the Database Machine.

A. ASM intelligent Data Placement (IDP)



- B. Intelligent Database Protocol (IDB)
- C. Database Resource Manager (DBRM)
- D. I/O Resource Manager (IORM)
- E. Database Filesystem (DBFS)
- F. The DISKMON process

Correct Answer: ABD

Explanation: A: 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.

B: The Exadata software is optimally divided between the database servers and Exadata cells. The database servers and Exadata Storage Server Software communicate using the iDB ? the Intelligent Database protocol. iDB is implemented in the database kernel and transparently maps database operations to Exadata-enhanced operations. iDB implements a function shipping architecture in addition to the traditional data block shipping provided by the database. iDB is used to ship SQL operations down to the Exadata cells for execution and to return query result sets to the database kernel. Instead of returning database blocks, Exadata cells return only the

D: The inter-database I/O allocations are defined within the software in the Exadata cell and managed by the I/O Resource Manager (IORM). The Exadata cell software ensures that inter-database I/O resources are managed and properly allocated within, and between, databases.

[Latest 1Z0-027 Dumps](#)

[1Z0-027 Practice Test](#)

[1Z0-027 Exam Questions](#)



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.