



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.passapply.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

A file contains a script with several EXACLI commands that must be executed on each cell in an X5 full rack.

The script must run on each cell simultaneously.

How would you achieve this?

- A. Copy the script to all cells using the DCLI command and manually execute it on each cell using the DCLI command.
- B. Copy the script to the cells using the EXACLI command and execute the script on all cells using the EXACLI command.
- C. Execute it on all cells using the EXADCLI command specifying the file name containing the EXACLI script.
- D. Copy the script to all cells using the CELLCLI command and execute it on all cells in parallel using the CELLCLI command.
- E. Execute it on all cells using the DBMCLI command specifying the file name containing the EXACLI script.

Correct Answer: C

Explanation:

The exadcli utility runs commands on multiple remote nodes in parallel threads.

You can issue an ExaCLI command to be run on multiple remote nodes. Remote nodes are referenced by their host name or IP address. Unlike dcli, exadcli can only execute ExaCLI commands. Other commands, for example, shell commands, cannot be executed using exadcli.

References: http://docs.oracle.com/cd/E80920_01/DBMMN/exadcli.htm#DBMMN-GUID-4AE469A6-F291-4737-B975-F1B4B91D0BA0

References: http://docs.oracle.com/cd/E80920_01/DBMMN/exadcli.htm#DBMMN-GUID-4AE469A6-F291-4737-B975-F1B4B91D0BA0

QUESTION 2

Which three statements are true about bulk data loading capabilities in an X5 Database Machine environment?

- A. DBFS must use the DBFS_DG diskgroup for any DBFS store.
- B. DBFS may be used if Exadata-based shared storage is required to stage data before bulk loading into a database.
- C. DBFS must be used to have a POSIX-compliant shared storage solution that is accessible from the database servers on a Database Machine.
- D. DBFS must be used to bulk load data into a production database on the Database Machine.
- E. ACFS may be used if Exadata-based shared storage is required to stage data before bulk loading into a database.
- F. ACFS must be used to have a POSIX-compliant shared storage solution that is accessible from the database servers on a Database Machine.



G. ACFS may use the DBFS_DG diskgroup to contain the ADVN volume file.

Correct Answer: BCD

Explanation:

External tables on DBFS file-systems provide the probably the most high-performance way to bulk load data into your database.

Bulk Data Loading

Describe the various options for staging data to be bulk loading into Database Machine

Configure the Database File System (DBFS) feature for staging input data files

Use external tables based on input data files stored in DBFS to perform high-performance data loads

QUESTION 3

Which two activities are supported on the storage servers in an Exadata X6 Database Machine?

- A. Upgrading the Storage Server software package using RPM
- B. Configuring secure shell user equivalency for the cellmonitor user
- C. Installing an alternative package manager
- D. Changing the root password
- E. Upgrading a device driver for hard disks when inserting a replacement disk after a hard disk failure

Correct Answer: AC

QUESTION 4

In which scenario will griddisks residing on the same celldisk have balanced performance characteristics in an X5 High Capacity Exadata Storage Server?

- A. griddisks on memory based celldisks
- B. griddisks on hard disk-based celldisk consumed by ASM when creating diskgroups using ASM-based Intelligent Data Placement (IDP)
- C. griddisks on hard disk-based celldisks consumed by the Database Smart Flash Cache
- D. griddisks on Flash-based celldisks

Correct Answer: B

Explanation:



Intelligent Data Placement enables you to specify disk regions on Oracle ASM disks for best performance. Using the disk region settings, you can ensure that frequently accessed data is placed on the outermost (hot) tracks which have greater speed and higher bandwidth. In addition, files with similar access patterns are located physically close, reducing latency. Intelligent Data Placement also enables the placement of primary and mirror extents into different hot or cold regions.

References: https://docs.oracle.com/cd/E11882_01/server.112/e18951/asmdiskgrps.htm#OSTMG10072

QUESTION 5

You plan to migrate a database supporting both DSS and OLTP workloads to your new X5 Database Machine.

The workloads contain many complex aggregating functions and expensive joins on large partitioned tables in the DSS workload and indexed access for OLTP workloads.

Which three benefits accrue as a result of this migration?

- A. Superior compression capability designed specifically for OLTP workloads
- B. Columnar storage capability for data in row major data blocks that is held in flash cache
- C. Superior compression capability designed specifically for data warehouse tables
- D. Superior compression capability designed specifically for archival data
- E. Superior flash cache compression technique
- F. Cell offload processing for indexed-organized table access

Correct Answer: ACD

Explanation:

There are 6 different kinds of Table Data Compression methods:

1.
BASIC compression, introduced in Oracle 8 already and only recommended for Data Warehouse
2.
OLTP compression, introduced in Oracle 11 and recommended for OLTP Databases as well
3.
QUERY LOW compression (Exadata only), recommended for Data Warehouse with Load Time as a critical factor
4.
QUERY HIGH compression (Exadata only), recommended for Data Warehouse with focus on Space Saving



5.
ARCHIVE LOW compression (Exadata only), recommended for Archival Data with Load Time as a critical factor

6.
ARCHIVE HIGH compression (Exadata only), recommended for Archival Data with maximum Space Saving

Incorrect Answers:

B:

1.
BASIC compression, introduced in Oracle 8 already and only recommended for Data Warehouse

2.
OLTP compression, introduced in Oracle 11 and recommended for OLTP Databases as well

3.
QUERY LOW compression (Exadata only), recommended for Data Warehouse with Load Time as a critical factor

4.
QUERY HIGH compression (Exadata only), recommended for Data Warehouse with focus on Space Saving

5.
ARCHIVE LOW compression (Exadata only), recommended for Archival Data with Load Time as a critical factor

6.
ARCHIVE HIGH compression (Exadata only), recommended for Archival Data with maximum Space Saving

Method 1 and 2 are working by compressing the rows inside of Oracle Blocks – the more redundant values inside the block, the better the compression ratio. OLTP compression might be used with Exadata Database Machine, because Exadata V2 is designed to support OLTP as well as Data Warehouse. Methods 3 to 6 use Hybrid Columnar Compression (HCC), which is designed to achieve better compression ratios by organizing sets of rows into compression units and then detaching the columns from the rows and collecting them. It is more likely that the same columns store similar (redundant) data, therefore the compression ratio improves with HCC.

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

[1Z0-070 VCE Dumps](#)

[1Z0-070 Study Guide](#)

[1Z0-070 Braindumps](#)