

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.leads4pass.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 heap-organized table in one of your database schemas contains only date, char, varchar2, and number data type columns.

Which three operations can be offloaded to Exadata Storage Servers when performing a Smart Scan on this table?

- A. Column filtering
- B. MIN/MAX scans
- C. Virtual column filtering
- D. Nested loop join filtering
- E. Sort-merge join filtering
- F. Predicate filtering

Correct Answer: AFC

Reference: <https://www.oracle.com/ocom/groups/public/@otn/documents/webcontent/1367127.pdf>

Reference: <https://www.oracle.com/ocom/groups/public/@otn/documents/webcontent/1367127.pdf>

QUESTION 2

You plan to migrate a database supporting an OLTP workload to your new X5 Database Machine.

The current database instance supports a large number of short duration sessions and a very high volume of short transactions.

Which three X5 Database Machine features can improve performance for this type of workload?

- A. An improved highly efficient undo and redo architecture
- B. Faster optimization due to an Exadata-specific optimizer
- C. Reduced I/O latency for writes due to writeback flashcache on all Exadata X5 and later models
- D. Ultra high I/O performance for reads and writes when using Exadata Extreme Flash in X5 and later models
- E. Reduced I/O latency for reads due to read flashcache on all Exadata X5 and later models

Correct Answer: ACD

Explanation:

A: To further accelerate OLTP workloads, the Exadata Smart Flash Cache also implements a special algorithm to reduce the latency of log write I/Os called Exadata Smart Flash Logging.

C: Use the Write-Back Flash Cache feature to leverage the Exadata Flash hardware and make Exadata Database

Machine a faster system for Oracle Database Deployments.

D: Exadata X5-2 introduces Extreme Flash Storage Servers. Each Extreme Flash storage server contains eight 1.6 TB state-of-the-art PCI Flash drives. PCI flash delivers ultra-high performance by placing flash memory directly on the high speed PCI bus rather than behind slow disk controllers and directors.

References: <http://www.oracle.com/technetwork/database/exadata/exadata-x5-2-ds-2406241.pdf>
<http://www.oracle.com/technetwork/articles/database/exadata-write-back-flash-2179184.html>

QUESTION 3

You must replace the battery on the disk controller in three Exadata storage servers and must power off the affected storage servers.

Which two commands must you execute to safely power off the storage servers in an X5 Database Machine?

- A. 'shutdown -h now'on the affected storage servers
- B. CellCLI> LIST GRIDDISK WHERE STATUS != 'inactive'
- C. CellCLI > LIST GRIDDISK ATTRIBUTES name WHERE asmdeactivationoutcome != 'Yes'
- D. 'crsctl stop cluster -all'on one of the database servers
- E. CellCLI> ALTER GRIDDISK ALL INACTIVE

Correct Answer: CE

Explanation:

C: Step 1:

Run the following command to check if there are other offline disks

```
CellCLI> LIST GRIDDISK ATTRIBUTES name WHERE asmdeactivationoutcome != \"Yes\"
```

If any grid disks are returned, then it is not safe to take the storage server offline because proper Oracle ASM disk group redundancy will not be intact.

E: Step 2:

Inactivate all the grid disks when Oracle Exadata Storage Server is safe to take offline using the following

command:

```
CellCLI> ALTER GRIDDISK ALL INACTIVE
```

Taking the storage server offline when one or more grid disks are in this state will cause Oracle ASM to dismount the affected disk group, causing the databases to shut down abruptly.

QUESTION 4

Which two statements are true about Exadata X5 cell metrics and alerts?

- A. Cell alerts on each storage server are accumulated in memory by the CELLSRV process and stored on a filesystem-based repository.
- B. Cell metrics are written to disk every hour by default.
- C. Cell metrics on each storage server are accumulated in memory by the CELLSRV process and written to a filesystem-based repository.
- D. Cell alerts on each server are accumulated in memory by the MS process and stored on a filesystembased repository.
- E. Cell alerts are written to disk every hour by default.

Correct Answer: BD

Explanation:

Metrics are a series of measurements that are computed and retained in memory for an interval of time, and stored on a disk for a more permanent history.

On the storage servers, the CELLSRV process provides the majority of Oracle Exadata storage services and is the primary storage software component. One of its functions is to process, collect, and store metrics. The Management Server (MS) process receives the metrics data from CELLSRV, keeps a subset of metrics in memory, and writes to an internal disk-based repository hourly.

References: <http://www.oracle.com/technetwork/articles/servers-storage-admin/monitor-exadata-em12-2291964.html>

References: <http://www.oracle.com/technetwork/articles/servers-storage-admin/monitor-exadata-em12-2291964.html>

QUESTION 5

Identify three valid reasons for creating multiple griddisks on a single hddisk-based celldisk.

- A. to implement storage realms so that storage that can be reserved for specific resource consumer groups in the same database
- B. to enable the creation of normal or high redundancy ASM diskgroups
- C. to segregate storage into multiple pools with different performance characteristics
- D. to enable disk mirroring for the system area
- E. to implement storage realms so that storage can be reserved for specific databases
- F. to implement storage realms so that storage that can be reserved for specific Grid Infrastructure clusters

Correct Answer: BCF

Explanation:

Creating multiple grid disks per cell disk allows you to create multiple pools of storage on the same Exadata Storage Server. The multiple grid disks can be assigned to separate ASM diskgroups, which can be provisioned to different databases.

Note: Griddisk is a logical disk that can be created on a celldisk. In a standard Exadata deployment we create griddisks on hard disk based celldisks only. While it is possible to create griddisks on flashdisks, this is not a standard practice.

F: After you complete the cell configuration, you can perform the following optional steps on the storage cell:

1.

Add the storage cell to the Exadata Cell realm

2.

Configure security on the Oracle Exadata Storage Server grid disks

References: https://docs.oracle.com/cd/E80920_01/SAGUG/SAGUG.pdf

QUESTION 6

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 7

Which three are true concerning Hybrid Columnar Compression (HCC) deployed on Exadata storage?

- A. HCC can be used only when the Exadata Smart Flash Cache is configured in Write-Through mode.
- B. HCC data is never cached in the Exadata Smart Flash Cache.
- C. By default, decompression is performed by Exadata Storage Servers.
- D. Rows residing in HCC compressed segments, are always self- contained in a single database block.
- E. Rows residing in HCC compressed segments, are always self- contained in a single compression unit.
- F. Row-level locks are supported on HCC compressed data.

Correct Answer: CEF

Explanation:

C: The decompression process typically takes place on the Oracle Exadata Storage Server in order to maximize performance and offload processing from the database server.

E: A logical construct called the compression unit is used to store a set of hybrid columnar compressed rows. When data is loaded, column values for a set of rows are grouped together and compressed. After the column data for a set of rows has been compressed, it is stored in a compression unit.

F: What happens when I update a row on compressed tables? What about locks? BASIC and OLTP: the updated row stays in the compressed block 'usual' Oracle's row-level locks

Hybrid columnar: Updated row is moved, as in a delete + insert How to see that? With dbms_rowid package New row is OLTP compressed if possible Lock affects the entire CU that contains the row

Note: Oracle's Hybrid Columnar Compression technology is a new method for organizing data within a database block. As the name implies, this technology utilizes a combination of both row and columnar methods for storing data. This hybrid approach achieves the compression benefits of columnar storage, while avoiding the performance shortfalls of a pure columnar format.

Incorrect Answers:

B: Data remains compressed not only on disk, but also remains compressed in the Exadata Smart Flash Cache, on Infiniband, in the database server buffer cache, as well as when doing back-ups or log shipping to Data Guard.

References: <http://www.oracle.com/technetwork/database/exadata/ehcc-twp-131254.pdf>

QUESTION 8

Your Exadata X6-2 Database Machine has the exachk utility preinstalled. You decide to use it periodically to validate the installation against Oracle's recommended best practices.

Which three actions could you take to do this?

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

Correct Answer: ADE

Explanation:

AD: To schedule exachk, we can create a cron job or we can create a job in Enterprise Manager..

To add an Enterprise Manager job to execute Exachk on a periodic basis. Execute the following steps.

First go to

Enterprise, Job, Library. Then create a new OS Command job

E: exachk options include AUTORUN_SCHEDULE, which schedules recurring health checks using the AUTORUN_SCHEDULE daemon option.

References:

<http://www.oracle.com/technetwork/database/availability/exadata-health-resource-usage-2021227.pdf>

http://docs.oracle.com/cd/E68491_01/OEXUG/automated-daemon-mode-operation.htm

Best practices: Execute the latest version of the exachk utility before and after planned maintenance activities and at least once every month

QUESTION 9

Which two are true about sparse griddisks and their use in disk groups on an X5 Exadata Database Machine?

- A. Sparse diskgroups must be created using sparse griddisks.
- B. Sparse diskgroups may be created using a combination of sparse and non-sparse griddisks.
- C. Sparse diskgroups may not be used for database snapshots.
- D. Additional space for a sparse griddisk is allocated as soon as newly written data is stored in the flashcache on a cell.
- E. The virtual size of a sparse griddisk may exceed the physical size of the space occupied by the griddisk.

Correct Answer: AE

Explanation:

A: A sparse ASM disk group is composed of sparse grid disks.

E: Sparse grid disks allocate space as new data is written to the disk, and therefore have a virtual size that can be much larger than the actual physical size. Sparse grid disks can be used to create a sparse disk group to store database files that will use a small portion of their allocated space. Sparse disk groups are especially useful for quickly and efficiently creating database snapshots on Oracle Exadata. Traditional databases can also be created using a sparse disk group.

References:

http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-server-snapshots.htm#SAGUG-GUID42945059-13FD-4F6A-B7FA-A1201D16238F http://docs.oracle.com/cd/E80920_01/DBMSO/exadata-whatsnew.htm#DBMSO22120

QUESTION 10

Examine this IORM plan: Which two are true concerning this plan?

```
CellCLI> list iormplan detail
```

```
name:          dmorlcel_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, limit=60,
               flashcache=on, flashlog=on
               name= finance, level=1, allocation=45, limit=60,
               flashcache=on, flashlog=off
               name=other, level=1, allocation=10,
               flashcache=off, flashlog=cn
objective:     balanced
status:        active
```

- A. The Finance database can use at least 45%, but never more than 60%, of the total Flash Cache capacity.
- B. I/Os from the finance database are guaranteed to get a minimum of 45% of the I/O bandwidth if the demand exists, and a maximum of 60% of the I/O bandwidth even if no other databases are doing I/O to the cell, and the demand from the finance database exceeds 60% of the maximum I/O rate of the cell.
- C. If I/Os come from the HR database only, then they may get up to 100% of the I/O bandwidth on the cell.
- D. I/Os from the finance database are guaranteed to get 45% of the I/O bandwidth if the demand is at least 60% of the maximum I/O rate of the cell, but may get 100% of the I/O bandwidth if no other databases are performing I/O to the cell.
- E. If I/Os come from the HR database only, then they may get up to 10% of the I/O bandwidth on the cell.

Correct Answer: BE

Explanation: The IORM plan can be configured using the ALTER IORMPLAN command on command-line interface (CellCLI) utility on each Exadata storage cell. It consists of two parameters - dbplan and catplan. While the "dbplan" is used to create the I/O resource directives for the databases, the "catplan" is used to allocate resources by workload category consolidated on the target system.

allocation/share - Specify the resource allocation to a specific database in terms of percentage or shares. limit - Specify maximum limit of disk utilization for a database.

Incorrect Answers:

A: IORM plans configures % I/O resources, not % of Flash Cache.

References: <https://community.oracle.com/docs/DOC-998939>

QUESTION 11

You issued these commands to all Exadata Storage Servers in an X6 Exadata Database Machine using dcli:

```
alter iormplan objective = offalter iormplan active
```

There are no database or category plans defined.

You are suffering I/O performance problems at certain times, which vary by day and week.

DSS and batch workloads perform well, but your OLTP workloads suffer poor response times when running at the same time as these other workloads.

You must ensure that DSS and batch workloads retain relatively high throughput without causing excessive degradation of OLTP performance.

Which statement would you issue to all Exadata Storage Servers to achieve this?

- A. alter iormplan objective = low_latency
- B. alter iormplan objective = high_throughput
- C. alter iormplan objective = balanced
- D. alter iormplan objective = auto

Correct Answer: A

Explanation:

When the objective option is set to basic, the database resource plan maximum utilization limits are not enforced. For stricter plan conformance, and enforcement of maximum utilization limits, the objective option must be set to something other than basic. The supported IORM objectives are auto, low_latency, balanced, and high_throughput.

The objective may also be changed to low_latency in which case the latency for critical I/Os is very good but there is a significant degradation in scan throughput when both workloads are running concurrently.

Incorrect Answers:

B: The objective can be changed to high_throughput which will increase the scan throughput at the cost of critical I/O latency.

C, D: The recommended objective option is auto which allows IORM to continuously monitor the workloads, and select the best mode based on the active workloads currently on the cells. The objective values of auto and balanced have the same behavior.

QUESTION 12

Which two communication methods are used by which components in the Enterprise Manager architecture for a Database Machine?

- A. SNMP traps for alerts are sent by the storage server ILOM directly to the Enterprise Manager agent.
- B. SNMP traps for alerts are sent by the storage server ILOM to the storage server MS process.
- C. SNMP traps for alerts are sent by the storage server MS process to the Enterprise Manager agent.
- D. SNMP traps for alerts are sent by the storage server MS process to the storage server ILOM.
- E. SNMP traps for alerts are sent by the storage server ILOM to the storage server RS process.

Correct Answer: BC

Reference: https://docs.oracle.com/cd/E91266_01/EMXIG/GUID-FB58204F-2D97-41BC-9AA7-10BFF920B5B4.htm#EMXIG145

Reference: https://docs.oracle.com/cd/E91266_01/EMXIG/GUID-FB58204F-2D97-41BC-9AA7-10BFF920B5B4.htm#EMXIG145

QUESTION 13

Which two options should be used to identify a damaged or failing flash card on an X6 Database Machine high-capacity storage server?

- A. Using the CELLCLI CALIBRATE command on the storage server after logging in as the root user
- B. Using the CELLCLI LIST LUN DETAIL command as the celladmin user
- C. Using the CELLCLI CALIBRATE command on the storage server after logging in as the celladmin user
- D. Hardware monitoring using the storage server ILOM

Correct Answer: AD

Explanation:

The CALIBRATE command runs raw performance tests on cell disks, enabling you to verify the disk performance before the cell is put online. You must be logged on to the cell as the root user to run CALIBRATE.

Oracle ILOM automatically detects system hardware faults and environmental conditions on the server. If a problem occurs on the server, Oracle ILOM will automatically do the following:

Illuminate the Service Required status indicator (LED) on the server front and back panels.

Identify the faulted component in the Open Problems table.

Record system information about the faulted component or condition in the event log.

References: http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-server-cellcli.htm#SAGUG20617

References: http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-server-cellcli.htm#SAGUG20617

QUESTION 14

Which three are true concerning Smart Scans?

- A. The decision to perform a Smart Scan is made by the Optimizer.
- B. Smart Scans are only possible during full table scans or fast full index scans.
- C. Smart Scans are possible on Index Organized Tables (IOTs).
- D. Smart Scans always outperform scans that are not offloaded.
- E. Smart Scans are possible on heap organized tables.
- F. Smart Scans are possible on uncompressed B*Tree indexes.

Correct Answer: ABF

Explanation:

B: In order for queries to take advantage of Exadata's Offloading capabilities, the optimizer must decide to execute a statement with a Full Table Scan or a Fast Full Index Scan.

F: Smart Scans can be used for full scanning through B*Tree index segments.

References: Expert Oracle Exadata - K. Osborne, et al., (Apress, 2011) BBS, page 346

QUESTION 15

Which three statements are true regarding the configuration of Auto Service Request (ASR) on an X5 Database Machine?

- A. ASR Manager may be installed on any type of server running Oracle Solaris.
- B. ASR Manager may be installed on any type of server running Oracle Linux.
- C. Oracle Linux must be used on the server where ASR Manager is installed.
- D. SMTP must be enabled on at least one database server.
- E. Oracle Solaris must be used on the server where ASR Manager is installed.
- F. ASR Manager must be installed on one of the database servers.
- G. HTTPS connectivity must be enabled from the ASR Manager host to the Internet.

Correct Answer: ABG

Explanation:

AB: Oracle recommends that you install Oracle ASR Manager on an external, standalone server. This server receives fault telemetry information from Oracle Exadata Database Machine servers. This server must run an Oracle Solaris or Linux operating system.

G: Before installing Oracle Auto Service Request (Oracle ASR) you must ensure connectivity to the Internet using HTTPS.

References: Oracle Auto Service Request Exadata Database Machine Quick Installation Guide

http://docs.oracle.com/cd/E80920_01/ASXQI/toc.htm#GUID-89192315-1B1A-4B5A-89BB94BB5B6ABB05

[Latest 1Z0-070 Dumps](#)

[1Z0-070 Practice Test](#)

[1Z0-070 Exam Questions](#)