

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

**Vendor:**Oracle

**Exam Code:**1Z0-070

**Exam Name:**Oracle Exadata X5 Administration

**Version:**Demo

## QUESTION 1

Which statement is true about operating systems in an X5 Database Machine multirack configuration consisting of two full racks and one Exadata storage expansion rack?

- A. All Exadata storage servers used by the same virtual cluster nodes must run the same O/S but Exadata Storage Servers in different clusters may run different operating systems.
- B. All Exadata storage servers must run the Oracle Solaris O/S and all database servers within the same cluster must run Oracle Linux.
- C. All Exadata storage servers may run Oracle Virtual Machine (OVM).
- D. All Exadata storage servers must run Oracle Linux.
- E. All Exadata storage servers must run the Oracle Linux O/S and all database servers within the same cluster must run the same version of Oracle Virtual Machine (OVM).

Correct Answer: D

Explanation:

On both physical and virtual deployments, Exadata systems use minimal Linux distributions to ensure that just the RPMs needed to run Oracle database, are installed and enabled.

References: <http://www.oracle.com/technetwork/database/exadata/exadata-x5-2-ds-2406241.pdf>

---

## QUESTION 2

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

```
alter iormplan objective = low_latency  
alter iormplan active
```

There are no database or category plans defined.

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

DSS and Batch workloads perform well some of the time.

Further investigation shows that at times, the workloads are all OLTP I/Os, at other times all batch I/Os, and sometimes a bit of each.

You wish to have disk I/O managed so that performance will be optimized for all workloads.

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

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

D. alter iormplan objective=auto

E. alter iormplan objective=' '

Correct Answer: D

Explanation:

The supported IORM objectives are auto, low\_latency, balanced, and high\_throughput. 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.

References: [http://docs.oracle.com/cd/E80920\\_01/SAGUG/exadata-storage-server-iorm.htm](http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-server-iorm.htm)

---

### QUESTION 3

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

---

### QUESTION 4

Examine this list of software components:

1.

Oracle Virtual Machine (OVM)

2.

Oracle Enterprise Manager Agent (OMA)

3.

ASM instance

4.

RDBMS instance

5.

Automatic Diagnostic Repository Command Interpreter (ADRCI)

6.

CELLCLI

7.

Cell Server (CELLSRV)

8.

diskmon

Identify the location where these software components can run in the standard Exadata Database Machine deployment.

- A. 1, 2, 3 and 4 run on the database servers; 5, 6, 7, and 8 run on the Exadata storage servers.
- B. 1, 2, 3, 4, 5, and 8 run on the database servers; 5, 6, and 7 run on the Exadata storage servers.
- C. 1, 2, 3, 4 and 8 run on the database servers; 5, 6, and 7 run on the Exadata storage servers.
- D. 3, 4 and 8 run on the database servers; 1, 2, 5, 6 and 7 run on the Exadata storage servers.
- E. 2, 3, 4 and 8 run on the database servers; 1, 5, 6 and 7 run on the Exadata storage servers.

Correct Answer: B

Explanation:

Automatic Diagnostic Repository Command Interpreter (ADRCI) can be used on an Exadata storage server.

Incorrect Answers:

A: The diskmon process is a fundamental component of Oracle Exadata Storage Server Software, and is responsible for implementing I/O fencing. The process is located on the database server host computer, and is part of Oracle Clusterware Cluster Ready Services (Oracle Clusterware CRS).

D, E: Oracle Virtual Machine (OVM) runs on the database servers.

---

## QUESTION 5

Which three are true concerning Storage Indexes?

- A. A maximum of eight table columns for any table are indexed per storage region.
- B. The use of Storage Indexes for particular categories of I/O can be displayed by using an I/O Resource Manager Category Plan.
- C. Storage Indexes persist across any Exadata storage server reboots.
- D. The use of Storage Indexes for a particular database can be disabled by using an I/O Resource Manager Database Plan.
- E. A Storage Index is automatically maintained by CELLSRV based on the filter columns of the offloaded SQL.
- F. Different storage regions may have different columns indexed for the same table.

Correct Answer: ADE

Explanation:

A: Each disk in the Exadata storage cell is divided into equal sized pieces called storage regions (default 1MB). There is an index entry for every storage regions (1MB of data stored on disk). Each entry contains the minimum and maximum value for columns seen in 'where' clause predicates. Information for up to 8 columns can be stored. The index is then used to eliminate disk IO by identifying which storage regions don't match the 'where' clause of a query.

---

## QUESTION 6

Which three statements are true about the EXADCLI utility?

- A. It may be run interactively.
- B. It can be used to execute EXACLI scripts on multiple storage servers in parallel.
- C. It can be used to execute O/S commands on multiple storage servers in parallel.
- D. It can be used to execute EXACLI commands on multiple storage servers in parallel.
- E. It uses the same security mechanism as the EXACLI command.
- F. It may be used to execute DBMCLI commands on multiple database nodes in parallel.

Correct Answer: DEF

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.

Incorrect Answers:

A: Note that exadcli runs ExaCLI in a "no-prompt" mode. This means that if user interaction is needed (for example, if you need to enter a password or if you need to verify that the certificate from a remote node is valid), then exadcli will exit with an error.

References:

[http://docs.oracle.com/cd/E80920\\_01/DBMMN/exadcli.htm#DBMMN-GUID-4AE469A6-F291-4737-B975-F1B4B91D0BA0](http://docs.oracle.com/cd/E80920_01/DBMMN/exadcli.htm#DBMMN-GUID-4AE469A6-F291-4737-B975-F1B4B91D0BA0)  
[https://docs.oracle.com/cd/E62172\\_01/html/E63692/z400007d1478481.html](https://docs.oracle.com/cd/E62172_01/html/E63692/z400007d1478481.html)

[http://docs.oracle.com/cd/E80920\\_01/DBMMN/exadcli.htm#DBMMN-GUID-4AE469A6-F291-4737-B975-F1B4B91D0BA0](http://docs.oracle.com/cd/E80920_01/DBMMN/exadcli.htm#DBMMN-GUID-4AE469A6-F291-4737-B975-F1B4B91D0BA0)  
[https://docs.oracle.com/cd/E62172\\_01/html/E63692/z400007d1478481.html](https://docs.oracle.com/cd/E62172_01/html/E63692/z400007d1478481.html)

---

**QUESTION 7**

Which two statements are true about Exadata storage server and database server alerts on an X5 Database Machine?

- A. Metric alerts are never stateful.
- B. Metric thresholds, if set, will persist across storage and database server reboots.
- C. SMTP alert notifications must be sent to both ASR manager and Enterprise Manager agents.
- D. SNMP alert notifications can be sent to only one destination.
- E. Metrics have no thresholds set on them by default.

Correct Answer: BC

Reference:

<https://docs.oracle.com/en/engineered-systems/exadata-database-machine/sagug/exadata-storage-server-software-introduction.html#GUID-3E48425A-AB8A-4E62-80C4-BACA65A1F8D3>

[https://docs.oracle.com/cd/E91266\\_01/EMXIG/GUID-FB58204F-2D97-41BC-9AA7-10BFF920B5B4.htm#EMXIG304](https://docs.oracle.com/cd/E91266_01/EMXIG/GUID-FB58204F-2D97-41BC-9AA7-10BFF920B5B4.htm#EMXIG304)

<https://docs.oracle.com/en/engineered-systems/exadata-database-machine/sagug/exadata-storage-server-software-introduction.html#GUID-3E48425A-AB8A-4E62-80C4-BACA65A1F8D3>  
[https://docs.oracle.com/cd/E91266\\_01/EMXIG/GUID-FB58204F-2D97-41BC-9AA7-10BFF920B5B4.htm#EMXIG304](https://docs.oracle.com/cd/E91266_01/EMXIG/GUID-FB58204F-2D97-41BC-9AA7-10BFF920B5B4.htm#EMXIG304)

---

**QUESTION 8**

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 9

An Exadata storage server physical disk on an X5 high-capacity full rack entered the predictive failure state.

Identify the two steps that you must perform to replace this failed physical disk.

- A. Add the griddisks back into the ASM diskgroup they used to be a member of.
- B. Create a new celldisk and new griddisk on the replaced physical disk.
- C. Verify that the griddisks located on the physical disk have been successfully dropped from the associated ASM diskgroups.
- D. Identify the griddisks located on the failed physical disk and drop them from the associated ASM diskgroups.
- E. Replace the failed physical disk.

Correct Answer: CE

Explanation: You may need to replace a physical disk because the disk is in warning - predictive failure status. The predictive failure status indicates that the physical disk will soon fail, and should be replaced at the earliest opportunity. The Oracle ASM disks associated with the grid disks on the physical drive are automatically dropped, and an Oracle ASM rebalance relocates the data from the predictively failed disk to other disks.

Note: After the physical disk is replaced, the grid disks and cell disks that existed on the previous disk in that slot are re-created on the new physical disk. If those grid disks were part of an Oracle ASM group, then they are added back to the disk group, and the data is rebalanced on them, based on the disk group redundancy and ASM\_POWER\_LIMIT parameter.

References: [http://docs.oracle.com/cd/E80920\\_01/DBMMN/maintaining-exadata-storage-servers.htm#DBMMN21047](http://docs.oracle.com/cd/E80920_01/DBMMN/maintaining-exadata-storage-servers.htm#DBMMN21047)

---

## QUESTION 10

Which two statements are true regarding the use of Auto Service Request (ASR) with an X6 Database Machine?

- A. The database server ILOMs must use SMTP over the management network for notifications to ASR Manager.
- B. The database server ILOMs must have SNMP traps configured to use the management network for notifications to ASR Manager.
- C. The storage server ILOMs must have SNMP traps configured to use the management network for notifications to ASR Manager.
- D. The database server MS process must have SNMP traps configured to use the management network for notifications to ASR Manager.

Correct Answer: BC

Explanation:

B: Database Server ILOM plug-in

Monitoring databases and their instances, ASM environments, the Grid Infrastructure, and the host software environment are done by Enterprise Manager in the usual way as these are standard targets. But monitoring the hardware for the database servers requires the ILOM plug-in, as there is no Management Server (MS) on the database servers to receive SNMP traps from the ILOM. The plug-in will receive sensor state and availability data from the ILOM including alerts based on pre-set ILOM thresholds.

C: Exadata Storage Server plug-in extends the monitoring of exadata cells in addition to providing a GUI interface. The plug-in uses an SSH connection to the cellmonitor user on the cells and uses list commands only. This is for interactive monitoring. One may also set thresholds using the plug-in which are distinct from any thresholds set using cellcli utility as the celladmin user. For alerts to be sent to the plug-in, SNMP traps are used as follows:

Cell ILOM alerts are sent to the cell Management Server (MS) via an SNMP trap. The MS then send SNMP notifications onward to the plug-in.

Cell alerts flagged by MS itself, such as cell thresholds being exceeded, or ADR software alerts, are sent to the plug-in using SNMP.

References:

<https://dbatrain.wordpress.com/2011/06/>

[http://docs.oracle.com/cd/E21659\\_01/html/E21660/z40015671004046509.html](http://docs.oracle.com/cd/E21659_01/html/E21660/z40015671004046509.html)

---

## QUESTION 11

Which two statements are true about data paths used during RMAN backup and restore operations when media management servers use InfiniBand to connect to an X5 Database Machine?

- A. During backups, data blocks are always read from Smart Flash Cache by cellsrv.



B. During restores, data blocks read from the media are sent to a database server that sends the blocks to cellsvr to be written.

C. During backups, data blocks are read by cellsvr and sent to a database server, which sends the blocks to the media manager to be written to media.

D. During backups, data blocks are read by cellsvr and sent directly to the media manager to be written to media.

E. During restores, data blocks read from the media are sent directly to cellsvr to be written.

Correct Answer: DE

Explanation:

The fastest database backup is achieved via InfiniBand.

1.

Get dedicated Media Servers that connect to tape library

2.

Offers backup rates of up to 12GB/sec

RMAN does not back up directly to tape. However, it will integrate with media management software such as Oracle Secure Backup and utilize their capabilities to manage tape libraries.

Exadata Database Machine: The Database Machine contains the databases that need to be backed up. Oracle RMAN is the only mechanism to back up the databases that utilize Exadata Storage Servers as the storage. RMAN processes run on the database servers and interact with the Oracle Secure Backup (OSB) agent, which further interacts with the media management software and enables RMAN to communicate with the tape library.

References: [http://apprize.info/data/oracle\\_4/6.html](http://apprize.info/data/oracle_4/6.html)

---

## QUESTION 12

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>