

Oracle

Exam 1z0-485

Oracle Exadata Database Machine 2014 Implementation Essentials

Version: 6.1

[Total Questions: 71]

Question No : 1

Exadata Database Machine offers an Intelligent Platform Management Interface for the various components in the Exadata product. Which option is true?

- A. IPMI can be used to remotely start and stop servers.
- B. IPMI can be used to remotely log in to the operating system.
- C. IPMI commands can be executed by using SNMP traps.
- D. IPMI settings can be secured by backing up the Oracle Linux installation on the storage system.

Answer: A

Explanation: IPMI – short for Intelligent Platform Management Interface - is an interface standard that allows remote management of a server from another using standardized interface. The servers in the Exadata Database Machine follow that. It's not an Exadata command but rather a general Linux one.

To power on a cell or database server, issue this from another server:

```
# ipmitool -H prolcel01-ilom -U root chassis power on
```

To stop a server, use the shutdown command. To stop immediately and keep it down, i.e. not reboot, execute: **# shutdown -h -y now**

Question No : 2

Which is the best location to point your customer to, for finding the latest Exadata patches?

- A. owner's guide
- B. Patch database in MyOracle Support
- C. MyOracle Support note 888828.1
- D. MyOracle Support for database patches, ULN for OS patches, and Sun Support for Server and InfiniBand patches.

Answer: B

Explanation: Before starting, we would like to share and note here two documents from My Oracle Support, aka metalink. These notes must be the first place that you need to go to review before patching the Exadata environment.

* (B) Database Machine and Exadata Storage Server 11g Release 2 (11.2) Supported

Versions (Doc ID. 888828.1)

- This is for the second and third generation (V2 and X2) for Oracle Exadata, using Sun hardware.

* Database Machine and Exadata Storage Server 11g Release 1 (11.1) Supported

Versions (Doc ID. 835032.1)

* - This is for the first generation (V1) for Oracle Exadata, using HP hardware.

Question No : 3

A customer has three databases named CC, FIN, and DW. The CC database is for their CallCenter. Even a slight decrease in the response time of the database would mean more people "on hold" in their data center. The orders received through the CallCenter are stored in the Finance(FIN) database. Both databases, CC and FIN, serve as sources for the Data Warehouse (DW) database. All databases use the same Automatic Storage Management (ASM) diskgroup and therefore, the same physical disks on Exadata storage. The customer wants to dynamically control the resources that are available for the CallCenter database because this has direct impact on their operations?

Which option should be implemented?

- A. DBRM on the CC database
- B. DBRM on all the databases
- C. IORM, because DBRM cannot be changed for an existing session
- D. IORM and DBRM

Answer: A

Question No : 4

Consider the following software changes that are performed manually on a Linux server:

1. Changes for Linux kernel firewall configuration
2. Changes for custom performance monitoring tools
3. Changes for security scan tools
4. Changes for Linux system performance optimization

Which of the software changes listed are permitted on Exadata Storage Servers?

- A. 1, 2, 3, and 4
- B. only 3
- C. none
- D. only 2
- E. only 1 and 2
- F. only 2 and 3

Answer: E

Explanation: 1: The Storage Server Patch is responsible for keeping our cell nodes always up-to-date, fixing possible problems, and this patch includes different component patches, like kernel patches, firmware, operation system, etc... for the Storage Server.

Incorrect:

3,4: security scan tools changes and Linux system performance optimization changes would be on the database server.

Question No : 5

Which two statements are true about enabling write-back flash cache?

- A. When enabling write-back flash cache in a non rolling manner, it is important to ensure that `asmdeactivationoutcome` is set to YES and `asmModestatus` is set to ONLINE for all grid disks.
- B. Before using write-back flash cache, you need to verify the minimum required versions.
- C. Before write back-flash cache is enabled, you need to drop the Flash Cache first.
- D. The setting `flashCacheMode` should be set to `writeback` by updating `cellinit.ora` and restarting `cellsvr`.
- E. When enabling write-back flash cache in a rolling manner, `dcli` should be used to inactivate the grid disks on all cells first.

Answer: B,C

Explanation: B: Exadata storage version 11.2.3.2.1 is the minimum version required to use this write back flash cache option.

C: Steps for Enabling Write back flash cache:

First of all, you don't need the stop CRS or database (This is ROLLING method) ,you can do it cell by cell.

* drop flashcache

* Be sure `asmdeactivationoutcome` is YES is before disabling grid disk

Question No : 6

Consider the following setup:

User A1 belongs to resource group High on Database A.

User B2 belongs to resource group Low on Database B.

User C3 is a user on Database C without any DBRM setup.

DBRM setup:

Database A: Resource group High gets 80% and Low gets 20%.

Database B: Resource group High gets 60% and Low gets 40%.

IORM setup:

Database A: Share=20, limit=5

Database B: Share=30, limit=10

Database C: 5 shares

Total number of shares in the IORM setup = 100

What percent of I/O will each database user theoretically be using when the Exadata storage unit I/O throughout is used 100% and no other databases but A, B, and C are running?

- A. A1 = 36%, B2=18%, and C3=9%
- B. A1 = 33%, B2=33%, and C3=33%
- C. A1 = 10%, B2=5%, and C3=20%
- D. A1 = 8%, B2=12%, and C3=5%
- E. A1 = 5%, B2=10%, and C3=85%

Answer: E

Explanation: IORM setup limits Database A to 5%, and Database B is limited to 10%,

while Database C has not IORM limit.

Notethat the resource groups are for CPU allocation.

Question No : 7

Consider this CellCLI command:

```
CellCLI> CREATE GRIDDISK ALL HARDDISK PREFIX=data, size=423G;
```

Which two statements describe what happens when you execute this command?

- A. It creates one 423 GB grid disk on the first available cell hard disk.
- B. It creates one 423 GB grid disk on each available cell hard disk.
- C. It creates grid disks on the outermost 423 GB that is available on each hard disk.
- D. It creates grid disks on the innermost 423 GB that is available on each hard disk.
- E. It creates an Exadata Smart Flash Cache on all flash drives.

Answer: B,C

Explanation: * Example:

```
CellCLI> create griddisk all harddisk prefix=temp_dg, size=570G
```

This command will create 12 Griddisks, each of 570G in size from the outer (fastest) sectors of the underlying Harddisks. It fills up the first 2 Celldisks entirely, because they have just 570G space free – the rest is already consumed by the OS partition.

Question No : 8

You get a Host Unreachable error when you attempt to connect to a server through a network terminal command line. What are two other ways in which you can connect?

- A. Use the ILOM Web GUI.
- B. Use the dcli command at the root prompt on a database node.
- C. Attach a terminal device to the back panel of the server with a serial cable.
- D. Connect by using SQL *Plus.
- E. Log in as root on the database node using the Net1 IP address.

Answer: A,C

Explanation: In addition to gaining shell access via SSH to manage your Exadata servers, you can also access them from the Integrated Lights Out Management (ILOM) console or KVM console.

and should typically not require modifications unless you have changed network information inside your database machine.

Note: A KVM switch (with KVM being an abbreviation for "keyboard, video and mouse") is a hardware device that allows a user to control multiple computers from one or more[1] keyboard, video monitor and mouse. Although multiple computers are connected to the KVM, typically a smaller number of computers can be controlled at any given time

Question No : 9

Identify three best practices for applying asmdeactivationoutcome es on Exadata Database Servers and ExadataStorage Servers?

- A. Backing up database servers and storage cells is not recommended before performing planned maintenance.
- B. Database server updates can be rolled back using the the "yum downgrade" procedure.
- C. Bundle patches do not require testing before being installed on a production system.
- D. It is recommended that Exadata systems with Data Guard configured use the "Standby First" patching approach.
- E. Patching should never be interrupted due to a connection drop. It is therefore recommended that you use VNC or the screen utility.
- F. Before patching cells in a rolling manner, you must check asmdeactivationoutcome amModestatus and make sure that cells on all disks are online and that disks can be deactivated.

Answer: D,E,F

Question No : 10

The mpstat output from OS Watcher shows a database node as being 90% idle on an average. What would you do to get a full picture of CPU utilization on the entire Exadata RAC cluster?

- A. Average the mpstat id1 output from all the nodes.
- B. Ask application users if they have noticed a slowdown in screen response.
- C. Look for an increase in batch job servicing times.
- D. A & B above

Answer: A

Question No : 11

Which two attributes describe key benefits of the InfiniBand network?

- A. All Exadata database servers have a direct path link to each Exadata Storage Server.
- B. Cell-to-cell communication uses Reliable Datagram Sockets (RDS) over InfiniBand to achieve low latency.
- C. Expanding from two Full racks to four only requires adding an external InfiniBand switch to be at the top of the fat-tree topology.
- D. Each InfiniBand link provides 10 Gigabits of bandwidth.
- E. Oracle's interconnect protocol uses direct memory access (DMA) to eliminate buffer copies and reduce CPU use.


Answer: C,E

Explanation: C: Oracle Exadata is architected to scale-out to any level of performance. To achieve higher performance and greater storage capacity, additional database servers and Exadata cells are added to the configuration – e.g., Half Rack to Full Rack upgrade. As more Exadata cells are added to the configuration, storage capacity and I/O performance increases near linearly.

E: Oracle's interconnect protocol uses direct data placement (DMA – direct memory access) to ensure very low CPU overhead by directly moving data from the wire to database buffers with no extra data copies being made. The InfiniBand network has the flexibility of a LAN network, with the efficiency of a SAN.

Incorrect:

Not A: The architecture of the Exadata solution includes components on the database server and in the Exadata cell. The software architecture for a Quarter Rack configuration is shown below.



Not B: No cell-to-cell communication is ever done or required in an Exadata configuration.

Not D: Each InfiniBand link provides 40 Gigabits of bandwidth –

Question No : 12

Identify three Exadata Storage Server software processes and their purpose?

- A. CELLSRV: The Cell Server is responsible for servicing disk I/O and predicate processing offload.
- B. CELLSRV: The Cell Server is responsible for balancing workload to other storage servers.
- C. MS: The Management Server is responsible for storage cell management and configuration.
- D. MS: The Management Server is responsible for starting a local Enterprise Manager agent.
- E. RS: The Restart Server is responsible for Automatic Storage Management (ASM) instance restart.
- F. RS: The Restart Server is responsible for CELLSRV and MS monitoring and restart.

Answer: A,C,F

Reference: <http://www.oracle.com/technetwork/database/exadata/exadata-technical-whitepaper-134575.pdf>(page 22)