

Oracle 1z0-536

Oracle Exadata 11g Essentials
Practice Test
Version: 17.20

QUESTION NO: 1

What model of the Exadata Database Machine come with 3 Infiniband switches?

- A. Quarter Rack of X2-2
- B. Half Rack of X2-2
- C. Full Rack of X2-2
- D. All models

Answer: C

Explanation:

| | Database Machine Full Rack | Database Machine Half Rack | Database Machine Quarter Rack |
|-------------------------|--|----------------------------|--|
| Database Servers | 8 | 4 | 2 |
| Exadata Storage Servers | 14 | 7 | 3 |
| InfiniBand Switches | 3 | 2 | 2 |
| Upgradability | Connect multiple Full Racks via the included InfiniBand fabric | Field upgrade to Full Rack | Field upgrade from Quarter Rack to Half Rack |

http://en.wikipedia.org/wiki/Oracle_Exadata

http://jde.fors.ru/portal/page?_pageid=105,260779&_dad=portal&_schema=PORTAL

QUESTION NO: 2

Which types of data are most likely to be cached In the Exadata Smart Flash Cache?

- A. Results of random reads.
- B. Results of table scans.
- C. Write to a mirror
- D. Redo data
- E. All data is cached In the Flash Cache

Answer: A

Explanation: The Oracle Database and Exadata Storage Server Software work closely together to cache frequently accessed data. When the database sends a read or write request to Sun Oracle Exadata Storage Server, it includes additional information in the request about whether the data is likely to be read again and therefore whether it should be cached. Based on the information the database sends the Exadata Storage Server Software intelligently decides which data will be re-read, and is worth caching, and those operations that would just waste cache. **Random reads Exadata Smart Flash Cache and the Sun Oracle Database Machine against tables and indexes are likely to have subsequent reads and normally will be cached and have their data delivered from the flash cache.** Scans, or sequentially reading tables,

generally would not be cached since sequentially accessed data is unlikely to be subsequently followed by reads of the same data. Write operations are written through to the disk and staged back to cache if the software determines they are likely to be subsequently re-read. Knowing what not to cache is of great importance to realize the performance potential of the cache. **For example, when writing redo, backups or to a mirrored copy of a block, the software avoids caching these blocks. Since these blocks will not be re-read in the near term there is no reason to devote valuable cache space to these objects or blocks.** Only the Oracle Database and Exadata Storage Server software has this visibility and understands the nature of all the I/O operations taking place on the system. Having the visibility through the complete I/O stack allows optimized use of the Exadata Smart Flash Cache to store only the most frequently accessed data. All of this functionality occurs automatically without customer configuration or tuning and in most cases is the best use of the Exadata Smart Flash Cache.

<http://www.oracle.com/technetwork/middleware/bi-foundation/exadata-smart-flash-cache-twp-v5-1-128560.pdf>

QUESTION NO: 3

Your customer wants to use Hybrid Columnar Compression to get maximum compression on their data. Which option represents a best practice for achieving this goal?

- A. Load compressed files with Direct Path loading
- B. Sort incoming data on a column with a low cardinality
- C. Use DBFS and external files to load the data
- D. Use Data Pump to load data

Answer: A

Explanation: Hybrid Columnar Compressed Tables

- New approach to compressed table storage
- Useful for data that is bulk loaded and queried
- Update activity is light

Data loaded using Direct Load uses Hybrid Columnar Compression

- Parallel DML, INSERT /*+ APPEND */, Direct Path SQL*LDR

Reference:<http://www.oracle.com/technetwork/database/features/availability/311358-132337.pdf>

QUESTION NO: 4

Which statement would you make about sharing I/O resources in an Exadata environment?

- A. You can manage workloads within a database with Database Resource Manager.
- B. You can manage workloads across multiple databases with Database Resource Manager.
- C. You can manage workloads within a database with I/O Resource Manager.
- D. You cannot manage workloads across multiple databases.

Answer: C

Explanation: <http://www.slideshare.net/mrabne/exadata-v2-for-oltp-and-dwh-2236244>
slide 60

QUESTION NO: 5

How do you back up the software used on an Exadata storage cell?

- A. You use RMAN.
- B. You back up the software files using an operating system method.
- C. You do a complete copy of their system storage area.
- D. You do not need to backup this software, as there is an automated recovery mechanism built into each cell.

Answer: A

Explanation: Exadata works with Oracle Recovery Manager (RMAN), a command-line and Enterprise Manager-based tool, to allow efficient Oracle database backup and recovery. All existing RMAN scripts work unchanged in the Exadata environment. RMAN is designed to work intimately with the server, providing block-level corruption detection during backup and restore. RMAN optimizes performance and space consumption during backup with file multiplexing and backup set compression, and integrates with Oracle Secure Backup (OSB) and third party media management products for tape backup.

Migration can be done using Oracle Recovery Manager (RMAN) to backup from traditional storage and restore the data onto Exadata.

<http://www.oracle.com/us/solutions/datawarehousing/039572.pdf>

http://docs.google.com/viewer?a=v&q=cache:LRA3-6_fEdsJ:www.oracle.com/technetwork/database/features/availability/maa-tech-wp-sundbm-backup-final-129256.pdf+Exadata+storage+cell+backup+Oracle+Recovery+Manager&hl=en&gl=in&pid=bl&srcid=ADGEEsJdZzFjI9-eUWzRYEY7N6AXoHwkH_Vib2u5xAU1_pK0gvCpcjU86uq5zyDhcgNnkpoz6q7waRnrjJET4mywoyRjZ91u9ehPU1Fafc9pDRDyC1i5h0fJdFs28nsUe3u-ejOSekOZ&sig=AHIEtbR_chRZ2XvxeJ2SJQ2tLIVwQbyZwQ

QUESTION NO: 6

Which statement properly describes potential allocation for a grid disk?

- A. A grid disk must use all of a cell disk.
- B. A grid disk can only use portions of a cell disk across all cells.
- C. A grid disk can use portions of a cell disk on a subset of cells.
- D. A grid disk can span multiple Exadata Storage Server cells.

Answer: C

Explanation:

Space for grid disks can be allocated in an interleaved manner. Grid disks that use this type of space allocation are referred to as interleaved grid disks. This method attempts to equalize the performance of the **grid disks residing on the same cell disk** rather than having the grid disks that occupy the outer tracks getting better performance at the expense of the grid disks on the inner tracks.

<http://webcache.googleusercontent.com/search?q=cache:1jN0wMz6f34J:vishalgupta.com/oracle/Docs/ExadataV2/doc.112/e13861/whatsnew.htm+oracle+exdata+grid+disk&cd=5&hl=en&ct=clnk&gl=in&source=www.google.co.in>

CREATE GRIDDISK ALL PREFIX=prefix—This CELLCLI command automatically creates grid disks on all available Exadata cell disks.

Note that the name you provide for prefix should be the same as the disk group name. For example, to create grid disks for your ASM disk group DATA, use the CREATE GRIDDISK ALL PREFIX='DATA' command. Because no size is specified in this example, the grid disk consumes the whole cell disk.

http://www.oracle86.com/ora_dir/server.111/b28282/configbp001.htm

QUESTION NO: 7

Which resource plan prioritizes workload in the I/O Resource Manager?

- A. Storage plan
- B. Category plan
- C. Inter-database plan
- D. Intra-database plan

Answer: B

Explanation:

http://docs.google.com/viewer?a=v&q=cache:fLI_2oI_sGMJ:www.oracledatabase12g.com/wp-content/uploads/2010/11/041509_65094_ppt.ppt+I/O+Resource+Manager+Category+,Inter-database+,+Intra-database+plan&hl=en&gl=in&pid=bl&srcid=ADGEESj1eixF5XksWdThz5Le4XNNbHJPmaZdulvdA14luDaeH69ptNV7h9-2eATr8UnGkY5D1E_IA2PKGt0cnxFcUU7BToYvC_aQuJAALthnkfKVmmfMh2yVuXHIOIJqli6CnWzD9KjX&sig=AHIEtbRS7Q1lwLxVZqj8bGVopRI9jAZ0ig

slide 29-30

QUESTION NO: 8

What does the role attribute of a DB plan indicate?

- A. The role specified for the category
- B. The role specified for the user
- C. The role specified in a Data Guard environment
- D. The role specified for the application

Answer: B

Explanation:

QUESTION NO: 9

If a hard drive is removed from a storage cell, what must you do?

- A. Nothing.
- B. The cell must be rebooted.
- C. You must alter ASM to alert It
- D. You must recreate any grid disks that use the drive.

Answer: A

Explanation:

Automatic addition of replacement disk to the disk group: All the required Exadata operations to re-create the disk groups, and add the grid disks back to the original disk group are **now performed automatically when a replacement disk is inserted after a physical disk failure.**

•Automatic cell restart: Grid disks are automatically changed to online when a cell recovers from a failure, or after a restart.

<http://webcache.googleusercontent.com/search?q=cache:1jN0wMz6f34J:vishalgupta.com/oracle/Docs/ExadataV2/doc.112/e13861/whatsnew.htm+oracle+exadata+when+a+physical+disk+is+replaced+in+a+storage+cell&cd=8&hl=en&ct=clnk&gl=in&source=www.google.co.in>

<http://forums.oracle.com/forums/thread.jspa?threadID=2207219&tstart=0>

QUESTION NO: 10

Your customer has designated a number of database objects to be kept persistently in the Exadata Smart Flash Cache. What happens if the total size of these objects is greater than 80% of the size of the available Exadata Smart Flash Cache?

- A. Nothing
- B. Exadata Smart Flash Cache expands the allocation to hold the objects.
- C. Not all objects will be stored in Exadata Smart Flash Cache.
- D. Overall performance is increased as more objects fit into Exadata Smart Flash Cache.

Answer: C

Explanation: We have 3 possible values for CELL_FLASH_CACHE: DEFAULT (the default), KEEP and NONE. Keep means that the sales table will be stored in the Flash Cache “more aggressively” than the default.

in the Flash Cache, KEEP assigned objects are also subject to an age out policy – if those objects ain’t accessed for a longer time, they will vanish from the flash cache to make room for objects

with the default attribute. And yes, they cannot consume more than 80 % of the Flash Cache. If you assign objects with in summary more than 80 % of the Flash Cache with the CELL_FLASH_CACHE KEEP attribute – they will simply not all be flash cached – at the same time.

<http://uhesse.wordpress.com/2011/02/02/exadata-part-iv-flash-cache/>

QUESTION NO: 11

Which three attributes are likely to result in data being cached in the Exadata Smart Flash Cache?

- A. CELL_FLASH_CACHE attribute on the data object
- B. CACHE hint In the SQL statement
- C. Data from a table scan
- D. Small data less than 128 KBs
- E. Control file I/Os

Answer: A,D,E

Explanation: A: Preferential treatment over which database objects are cached is also provided with the Exadata Smart Flash Cache. For example, objects can be pinned in the cache and always be cached, or an object can be identified as one which should never be cached. This control is provided by the new storage clause attribute, CELL_FLASH_CACHE, which can be assigned to a database table, index, partition and LOB column.

Incorrect:

C: Scans, or sequentially reading tables, generally would not be cached since sequentially accessed data is unlikely to be subsequently followed by reads of the same data.

<http://www.oracle.com/technetwork/middleware/bi-foundation/exadata-smart-flash-cache-twp-v5-1-128560.pdf>

QUESTION NO: 12

You are going to set up grid disks for your customer. Which is the best practice to use when defining grid disks for a Storage Server?

- A. Create a grid disk for every table space.

- B. Create two grid disks, one for data and one for recovery.
- C. Create two grid disks, one for tables and one for Indexes.
- D. Create a grid disk for each table.

Answer: B

Explanation: You have the option to use the disk groups you created using Oracle ASMCA both for database files and recovery files, or you can choose to use different disk groups for each file type. Ideally, you should create separate Automatic Storage Management disk groups for data files and for recovery files.

http://www.filibeto.org/sun/lib/nonsun/oracle/11.2.0.1.0/E11882_01/install.112/e10854/oraclerestart.htm

QUESTION NO: 13

When using IORM, which statement accurately describes when redo log file writes take place?

- A. Based on the priority of the user
- B. Immediately
- C. Based on the resource group of the user
- D. Based on the intradatabase resource plan

Answer: B

Explanation: Whenever a transaction is committed, the corresponding redo entries temporarily stored in redo log buffers of the system global area are written to an online redo log file.

Note:I/O Resource Manager =IORM

http://download.oracle.com/docs/cd/A58617_01/server.804/a58397/ch5.htm

QUESTION NO: 14

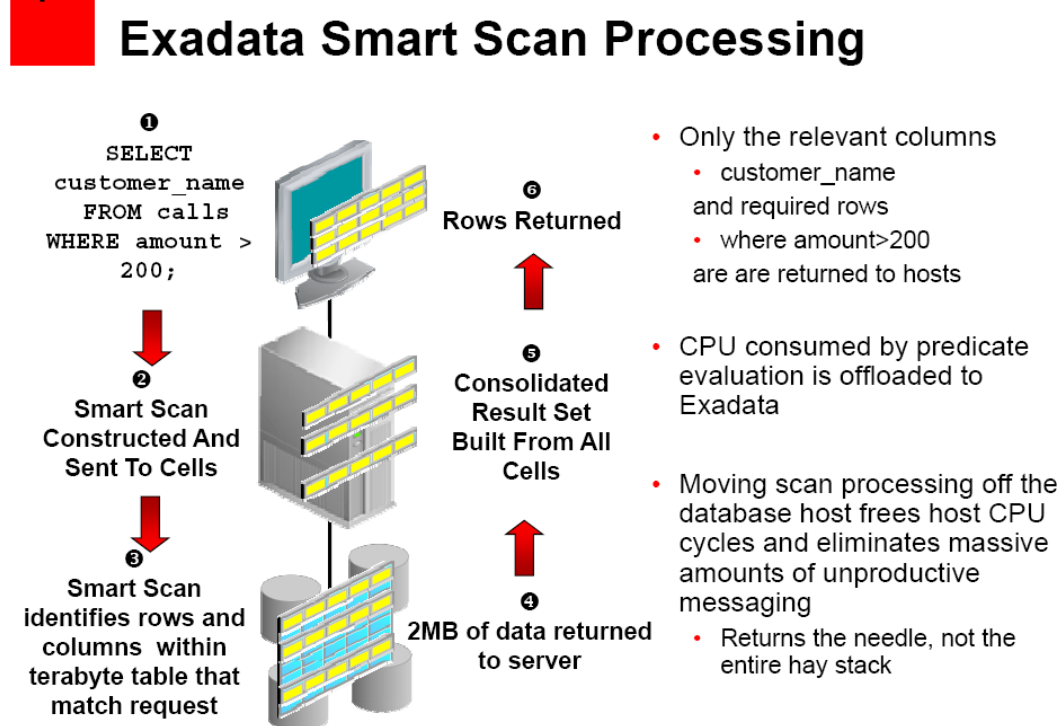
Your customer asks you what advantage Offload processing gives to the Exadata Storage Server. How can you explain it best?

- A. Offload processing allows PL/SQL functions to be executed on the Exadata Storage Server.

- B.** Offload processing aggregates SQL processing into higher level functions, which are shipped to the Exadata Storage Server.
- C.** Offload processing moves some processing closer to the stored data, which allows the Exadata Storage Server to return only the rows and columns requested.
- D.** Offload processing is just marketing - it provides no real advantage.

Answer: C

Explanation:



Reference: [www.sil.mu/newsevents/2010/sun_transforming_the_industry_event_20100215/presentations/Exadata V2 For Sun 2.pdf](http://www.sil.mu/newsevents/2010/sun_transforming_the_industry_event_20100215/presentations/Exadata_V2_For_Sun_2.pdf)

QUESTION NO: 15

How is cell-to-cell communication implemented in an Oracle Exadata Database Machine?

- A.** Through flash cache.
- B.** Through the Infiniband switch.
- C.** Through onboard memory.
- D.** There is never any cell-to-cell communication in an Exadata Storage Server.

Answer: D

Explanation: No cell-to-cell communication is ever done or required in an Exadata configuration.