Oracle

Exam 1z0-027

Oracle Exadata Database Machine X3 Administration

Version: 6.13

[Total Questions: 72]

Question No : 1

To guarantee proper cooling, you plan to place perforated floor tiles near your Database Machine.

Where, in relation to the cabinet, should they be placed?

A. On the left side, because the air flow is from left to right
B. At the back, because the air flow is from back to front
C. On the right side, because the air flow is from right to left.
D. At the front, because the air flow is from front to back
E. Underneath the cabinet, because the air flow is from bottom to top

Answer: D

Explanation: Airflow must be front-to-back.

Reference: Oracle White Paper, ORACLE Exadata Database Machine X3-8

Question No : 2

You plan to migrate your Oracle Version 11.1.0.2 database to your Exadata Database Machine.

The database supports an online transaction processing (OLTP) workload and is currently hosted on a Little Endian platform

Which two are the supported and appropriate migration methods to minimize downtime?

A. Upgrade source database to 11.2.0 and migrate using a physical standby database.
B. Migrate using Data Pump.
C. Migrate using GoldenGate.
D. Migrate using cross platform Transportable Database.
E. Migrate using ASM online migration.

Answer: A,B

Reference: Migrating the Oracle E-Business Suite Database to Oracle Exadata Database Machine Using Transportable Tablespaces

Reference: Migrating Oracle E-Business Suite to Oracle Exadata Database Machine Using Oracle Data Pump

Question No : 3

You recently upgraded your Exadata image to the latest release; previously you were using 11.2.0.3.

At the same time, you decide to address some performance problems as follows:

You noticed increased latency for the database log writer, especially during the quarterly battery learn cycle on the cells.

You have complaints of erratic performance from certain write-intensive applications.

Which two actions could improve performance in these areas?

A. Enable write-back flashcache by setting lunWriteCacheMode to Write Back Mode.
B. Use ALTER TABLE in the database to set CELL_FLASH_CACHE = KEEP for the tables belonging to the affected application.
C. Configure Smart Flash Log on the cells to use some of these of the space on the cell flash devices.
D. Configure the table belonging to the affected application using CELLCLI, to the set CELL_FLASH_CACHE = KEEP.
E. Configure Smart Flash Log on the database server to use server flash memory.

Answer: B,C

Explanation: B: The following command could be used to pin the table CUSTOMERS in Exadata
Smart Flash Cache
ALTER TABLE customers STORAGE (CELL_FLASH_CACHE KEEP)

C: Creating Flash Disks Out Of The Flash Cache
When an Exadata cell is installed, by default, all the flash is assigned to be used as flash cache and
user data is automatically cached using the default caching behavior. Optionally, a portion of the
cache can be reserved and used as logical flash disks. These flash disks are treated like

any
Exadata cell disk in the Exadata cell except they actually reside and are stored as non-volatile
disks in the cache.

Note:
* Pinning Objects In The Flash Cache
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

* There are two techniques provided to manually use and manage the cache. The first enables the
pinning of objects in the flash cache. The second supports the creation of logical disks out of the
flash for the permanent placement of objects on flash disks.

Question No : 4

You are about to run the oplan utility to patch the servers on your test Database Machine before patching the production environment.

The following task might be performed:

A) Test the failback procedure

B) Run the exachk utility

C) Read the README file.

D) Automate the patch application process as appropriate.

E) Verify that the patch provides the functionality it is meant to.

F) Apply the patch.

G) Evaluate the system performance.

In which order should the tasks be performed to patch in the recommended fashion?

A. C, B, D, F, B, E, A
B. C, D, F, B, E, G, A, B
C. C, B, D, F, E, G, A
D. C, B, D, F, E, A, G
E. C, B, D, F, B, E, G, A

Answer: E

Question No : 5

Which two are true about Smart Scan?

A. a query rewrite may occur to a container table stored in Exadata but will never benefit From Smart scan.
B. Column projection does not contribute to the performance benefit of Smart Scan
C. It is possible to offload single row functions to the storage servers.
D. Some joins can be offloaded to the storage servers.
E. A query rewrite may occur to a container table stored Exadata, and it will always benefit from Smart Scan.
F. All joins can be offloaded to the storage servers.

Answer: C,D

Explanation: C: With Exadata storage, database operations are handled much more efficiently. Queries that perform table scans can be processed within Exadata storage with only the required subset of data returned to the database server. Row filtering, column filtering and some join processing (among other functions) are performed within the Exadata storage cells. When this takes place only the relevant and required data is returned to the database server.

D (not F):
* Exadata performs joins between large tables and small lookup tables, a very common scenario for data warehouses with star schemas. Joining large tables and small lookup tables is implemented using Bloom Filters, which are a very efficient probabilistic method to determine whether a row is a member of the desired result set.

* If storage indexes are so great, why doesn't Oracle Exadata use them all the time? The short answer is that they are created and used only when they will be beneficial.
* To use storage indexes, Oracle Exadata queries must use smart scans, so not all types of applications can benefit from storage indexes. Applications with queries that include predicates and perform a lot of full table scans or fast full scans of indexes—typically those used in data warehousing environments—will benefit greatly from storage indexes. Online transaction processing (OLTP) applications, on the other hand, typically access a small number of rows through standard indexes and do not perform full table scans, so they may not benefit from storage indexes.

Note:
* Storage indexes reside in the memory of the storage servers—also called storage cells—and significantly reduce unnecessary I/O by excluding irrelevant database blocks in the storage cells.
* To use storage indexes, Oracle Exadata queries must use smart scans, so not all types of applications can benefit from storage indexes.

Incorrect:
Not B: Exadata provides column filtering, also called column projection, for table scans. Only the columns requested are returned to the database server rather than all columns in a table. For example, when the following SQL is issued, only the employee_name and employee_number columns are returned from Exadata to the database kernel.

SELECT employee_name, employee_number FROM employee_table.
For tables with many columns, or columns containing LOBs (Large Objects), the I/O bandwidth saved can be very large. Using both predicate and column filtering dramatically improves performance and reduces I/O bandwidth consumption. In addition, column filtering also applies to indexes, allowing for even faster query performance.

Reference: Oracle Communications Data Model Implementation and Operations Guide, Exadata Smart Scan Processing and Storage Index

Question No : 6

You are about to replace one memory DIMM in an Exadata storage server and need to power off the affected cell.

Which four commands must you execute to safely power off the storage server in your standard deployed quarter rack Database Machine assuming that redundancy is not compromised?

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

Answer: B,C,D,E

Explanation:

B: 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.

C: 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.

D: Step 3:
-- Verify all grid disks are INACTIVE to allow safe storage server shut down by running the following command:CellCLI> LIST GRIDDISK ATTRIBUTES name, asmmodestatusCellCLI> LIST GRIDDISKIf all grid disks are INACTIVE, then the storage server can be shutdown without affecting database availability

E: Step 4:
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 : 7

Which two statements are true about the IPTables firewall configuration on a Database Machine- Machine after the default Initial deployment?

A. IPTables is configured with Oracle supplied rules on the cells.
B. IPTables is configured with Oracle supplied rules on the database servers.
C. IPTables is installed and available but not configured on any servers.
D. IPTables is installed and available but not configured on the database servers.
E. IPTables is installed and available but not configured on the cells.

Answer: A,D

Question No : 8

Which three are true about Smart Flash log?

A. I/O Resource Manager database plans can be used to enable or disable Smart Flash Log for different databases.
B. LGWR will not wait for writes to Smart Flash log if the write to a disk based log file completes first.
C. Smart Flash Log is enabled by default, using 1024 MB of Rash storage on each storage server.
D. You can remove Smart Flash Log from a single storage server with the drop flashing command.
E. I/O Manager category plans can be used to enable or disable Smart Flash Log For different I/O categories.

Answer: A,B,D

Explanation: A: the Exadata I/O Resource Manager (IORM) has been enhanced to enable or disable Smart Flash Logging for the different databases running on the Database Machine.

B: Smart Flash Logging works as follows. When receiving a redo log write request, Exadata will do
parallel writes to the on-disk redo logs as well as a small amount of space reserved in the flash

hardware. When either of these writes has successfully completed the database will be immediately notified of completion. If the disk drives hosting the logs experience slow response
times, then the Exadata Smart Flash Cache will provide a faster log write response time. Conversely, if the Exadata Smart Flash Cache is temporarily experiencing slow response times
(e.g., due to wear leveling algorithms), then the disk drive will provide a faster response time.
This algorithm will significantly smooth out redo write response times and provide overall better
database performance.

D: Category plans are configured and enabled using the CellCLI utility on the cell. Only one category plan can be enabled at a time

Incorrect:
Not C: By default, 512 MB of the Exadata flash is allocated to Smart Flash Logging

Question No : 9

Which two are Oracle recommendations for media based backups performed for a database running on a Database Machine?

A. Allocate equivalent number of channels and instances per tape drive.
B. Perform periodic level 0 backups and daily cumulative level-1 backups.
C. Use the InfiniBand network between the database server and media server.
D. Configure Recovery Manager (RMAN) channels to connect to the least loaded instances.
E. Use InfiniBand network between the media server and the storage servers.

Answer: A,C

Explanation:
A: Configure one RMAN channel per tape drive and add tape drives to scale backup rates.

C: Configure the Preferred Network Interface (PNI) to direct the Oracle Secure Backup traffic over the InfiniBand network interface.

Example:
ob> lspni (List Preferred Network Interface)
mediaserver1:
PNI 1:
interface: mediaserver1-ib
clients: dbnode1, dbnode2, dbnode3, dbnode4,
dbnode5, dbnode6, dbnode7, dbnode8
PNI 2:
interface: mediaserver1
clients: adminserver
dbnode1:
PNI 1:
interface: dbnode1-ib
clients: mediaserver1

Note:
* Using the Sun ZFS Backup Appliance as an Oracle RMAN backup target for an Oracle Exadata system delivers much faster backup and recovery, enabling organizations to achieve shorter recovery time objectives and shrink backup windows. The appliance is designed for high sustained read and write I/O performance, and it is connected to the Oracle Exadata system via a high-throughput InfiniBand network fabric.

* As the only unified storage vendor to support InfiniBand as a storage network for backup and restore operations, Oracle is leading the way with native high-bandwidth interconnects.

* The InfiniBand network provides 40 Gb of bandwidth per port between the database servers, storage cells, and the Sun ZFS Backup Appliance. Backup and restore operations can be automatically parallelized across all database nodes, Oracle Exadata storage cells, Sun ZFS Backup Appliance channels, and controllers.

Question No : 10

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

This is part of a consolidation project and several other databases already exist on the Database Machine.

Which three Exadata features may help to improve the performance of this OLTP workload?

A. Hybrid Columnar Compression
B. I/O Resource Manager
C. Smart Flash Cache
D. Smart Flash Log
E. Smart Scan
F. Storage Index

Answer: B,C,D

Explanation: C:
OLTP performance benefits with Exadata
/ (D) Smart flash log for low latency commits
/ (C) Smart flash cache for low latency reads. KEEP in Flash for
critical objects

Question No : 11

Identify the three components that serve a purpose only in the Database Machine.

A. ASM intelligent Data Placement (IDP)
B. Intelligent Database Protocol (IDB)
C. Database Resource Manager (DBRM)
D. I/O Resource Manager (IORM)
E. Database Filesystem (DBFS)
F. The DISKMON process

Answer: A,B,D

Explanation:
Intelligent Data Placement, a feature of ASM that allows placing data in such a way that more frequently accessed data is located close to the periphery of the disk where the access is faster.

The Exadata software is optimally divided between the database servers and Exadata cells. The database servers and Exadata Storage Server Software communicate using the iDB –
the Intelligent Database protocol. iDB is implemented in the database kernel and

transparently
maps database operations to Exadata-enhanced operations. iDB implements a function shipping
architecture in addition to the traditional data block shipping provided by the database. iDB is
used to ship SQL operations down to the Exadata cells for execution and to return query result
sets to the database kernel. Instead of returning database blocks, Exadata cells return only the

The inter-database I/O allocations are defined within the software in the Exadata cell and managed by the I/O Resource Manager (IORM). The Exadata cell software ensures that inter-database I/O resources are managed and properly allocated within, and between, databases.

Question No : 12

Which two statements are true about the use of direct path loads when selecting from external tables in a database on a Database Machine?

A. INSERT INTO . . . SELECT FROM statements, executed serially, which select from external tables, require the APPEND hint to use direct path loading.
B. CREATE TABLE . . . AS SELECT statements, which select from external tables, attempt to use in direct path loading automatically.
C. CREATE TABLE . . . AS SELECT statements, which select from external tables, require the APPEND hint to use direct path loading.
D. INSERT INTO . . . SELECT FROM statements, executed serially, which select from external tables, are unable to use direct path loading.

Answer: A,B

Explanation: A CTAS (Create table as select) will always use direct path (B, not C) load but IAS (Insert as select) statement will not. In order to achieve direct path load with an IAS statement you must add the APPEND hint to the command (A, not D).

Direct path loads can also run in parallel. You can set the parallel degree for a direct path load
either by adding the PARALLEL hint to the CTAS or IAS statement or by setting the PARALLEL clause on both the external table and the table into which the data will be

loaded.
Once the parallel degree has been set at CTAS will automatically do direct path load in parallel
but an IAS will not. In order to enable an IAS to do direct path load in parallel you must alter

the session to enable parallel DML.

Note:
* Parallel Direct Path Load
The key to good load performance is to use direct path loads wherever possible. A direct path
load parses the input data according to the description given in the external table definition,

converts the data for each input field to its corresponding Oracle data type, then builds a column
array structure for the data. These column array structures are used to format Oracle data blocks
and build index keys. The newly formatted database blocks are then written directly to the database, bypassing the standard SQL processing engine and the database buffer cache.

Reference: Best Practices for Implementing a Data Warehouse on the Oracle Exadata Database Machine; Using CTAS & Exchange Partition Replace IAS for Copying Partition on Exadata

Question No : 13

You are evaluating the performance of a SQL statement that accesses a very large table, and have run the following query producing the output shown:

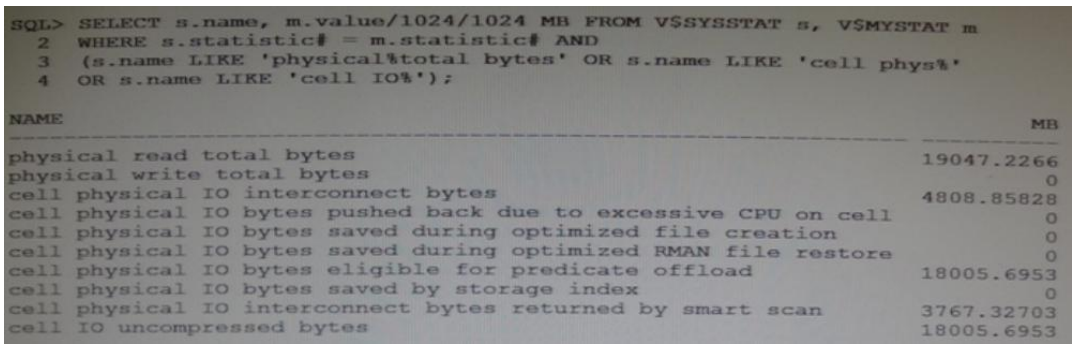

```
SQL> SELECT s.name, m.value/1024/1024 MB FROM V$SYSSTAT s, V$MYSTAT m
  2  WHERE s.statistic# = m.statistic# AND
  3  (s.name LIKE 'physical%total bytes' OR s.name LIKE 'cell phys%'
  4  OR s.name LIKE 'cell IO%');

NAME                                                                  MB
---------------------------------------------------------------- ----------
physical read total bytes                                      19047.2266
physical write total bytes                                              0
cell physical IO interconnect bytes                            4808.85828
cell physical IO bytes pushed back due to excessive CPU on cell         0
cell physical IO bytes saved during optimized file creation             0
cell physical IO bytes saved during optimized RMAN file restore         0
cell physical IO bytes eligible for predicate offload          18005.6953
cell physical IO bytes saved by storage index                           0
cell physical IO interconnect bytes returned by smart scan     3767.32703
cell IO uncompressed bytes                                     18005.6953
```

For which two reasons would the; “physical read total bytes” statistic be greater than the “cell physical IO bytes eligible for predicate offload” statistic?

A. There is an index on the column used in the where clause, causing “cell multiblock physical reads” to be requested by the database instance, resulting in additional I/O for blocks in the cells.
B. The table is an IOT and has an overflow segment, causing “cell multiblock physical reads” to be requested by the database instance, resulting in additional I/O for block in the cells.
C. There is an uncommitted transaction that has modified some of the table blocks, causing some “cell single block physical reads” to be requested by the database instance, resulting in additional I/O for block in the cells.
D. The table is an indexed clustered table, causing “cell single block physical reads” to be requested by the database instance, resulting in additional I/O for blocks in the cells.
E. There are migrated rows in the table, causing some “cell single block physical reads” to be requested by the database instance, resulting in additional I/O for blocks in the cells.

Answer: B,D

Question No : 14

You plan to monitor storage servers after configuring an I/O resource manager plan with directives for inter-database plans and intra-database plans.

Which two types if metrics would help assess the impact of the intra-database plans on I/O to the storage servers?

A. Category I/O
B. Database I/O
C. Resource Consumer Group I/O
D. Smart Flash Log I/O
E. Smart Flash Cache I/O

Answer: B,C

Explanation: B: Database metrics provide information about the size of the I/O load from each database specified in the interdatabase plan.
C: Consumer group metrics provide information about the size of the I/O load from each consumer group specified in a database resource plan. Each database in the interdatabase plan has metrics for each of its consumer groups.

Note:

* I/O Resource Manager (IROM) Settings

Incorrect:
Not A: Category metrics provide information about the size of the I/O load from each category specified in the current IORM category plan.

Question No : 15

A table in one of your database schemas contains only varchar, number, and date data types for the columns.

Which three operations can be offloaded to the Exadata storage servers when doing a smart scan against this table, if no other situations arise that prevent Smart Scan from occurring?

A. Column filtering
B. Sort merge join filtering
C. Predicate filtering
D. Nested loop Join filtering
E. Hash join filtering
F. Virtual column filtering

Answer: A,C,E
Explanation:

A: Smart Scan Column Filtering
Exadata provides column filtering, also called column projection, for table scans. Only the columns requested are returned to the database server rather than all columns in a table. For example, when the following SQL is issued, only
the employee_name and employee_number columns are returned from Exadata to the database kernel.
SELECT employee_name, employee_number FROM employee_table.
For tables with many columns, or columns containing LOBs (Large Objects), the I/O bandwidth saved can be very large. Using both predicate and column filtering dramatically improves performance and reduces I/O bandwidth consumption. In addition, column

filtering also applies to indexes, allowing for even faster query performance.

C: Smart Scan Predicate Filtering

Exadata enables predicate filtering for table scans. Only the rows requested are returned to the database server rather than all rows in a table. For example, when the following SQL is issued only rows where the employees' hire date is after the specified date are sent from Exadata to the database instance.

SELECT * FROM employee_table WHERE hire_date > '1-Jan-2003'.
This ability to return only relevant rows to the server greatly improves database performance. This performance enhancement also applies as queries become more complicated, so the same benefits also apply to complex queries, including those with subqueries.

Question No : 16

In which two locations should files be staged, to be loaded using external tables into a database on a Database Machine?

A. On a dbfs file system stored in a staging database on the Database Machine
B. On an Exadata-based ACFS file system on the Database Machine
C. On an nfs file system mounted on a database server where the external table will be accessed.
D. On local storage on one or more cells that are accessible to the database server where the load will be performed.

Answer: A,C
Reference: Installing Oracle E-Business Suite Release 12 with the Oracle Exadata Database Machine

Question No : 17

Which two are regarding the case of storage indexes?

A. To increase the chance of using the a storage index, you can make table indexes invisible.
B. To maximize the benefit of storage Indexes, load your data stored on the filtered columns.
C. The cell physical 10 bytes saved by storage index statistic returns multiple rows, one for each storage server.
D. Storage indexes are retained after a cell is rebooted.
E. Avoid the use of bind variables because Storage Indexes do not work with bind variables.

Answer: A,B

Explanation: To use storage indexes, Oracle Exadata queries must use smart scans, so not all types of applications can benefit from storage indexes.
* With Exadata storage, database operations are handled much more efficiently. Queries that perform table scans can be processed within Exadata storage with only the required subset of data returned to the database server. Row filtering, column filtering and some join processing (among other functions) are performed within the Exadata storage cells. When this takes place only the relevant and required data is returned to the database server.

Question No : 18

Which three factors should you consider when choosing a method for migrating a database to the Database Machine?

A. Type of database workload
B. Number of tablespaces in the source database
C. Size of the source database
D. Endian format of the source database
E. ASM attributes and allocation unit (AU) size of the diskgroups used by the source database

Answer: A,C,D

Explanation: A: Use Real Production Workload
– Real Application Testing (RAT)
• Database Workload Replay