1Z0-027 Q&As

Oracle Exadata X3 and X4 Administration

Pass Oracle 1Z0-027 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.pass4lead.com/1Z0-027.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Oracle Official Exam Center

QUESTION 1

You are monitoringand evaluatinga create index statement on yourDatabaseMachine and haverun the followingqueryafterexecutingstatement, providing the output. Shown:

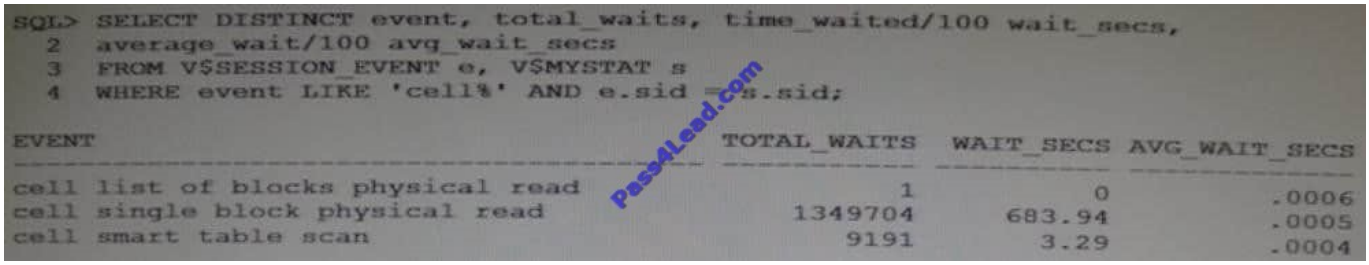

```
SQL> SELECT DISTINCT event, total_waits, time_waited/100 wait_secs,
  2  average_wait/100 avg_wait_secs
  3  FROM V$SESSION_EVENT e, V$MYSTAT s
  4  WHERE event LIKE 'cell%' AND e.sid = s.sid;

EVENT                                    TOTAL_WAITS  WAIT_SECS AVG_WAIT_SECS
---------------------------------------- ----------- ---------- -------------
cell list of blocks physical read                  1          0         .0006
cell single block physical read                1349704     683.94         .0005
cell smart table scan                             9191       3.29         .0004
```

Select two reasons why the statement would have produced so many "cell single block physical read" waits compared to "cell smart table scan"waits.

A. There are huge numbers of migrated rows in the table on which the index is being built.

B. There is an uncommitted transaction that has modified one block of the table on which the index is being built, in each cell.

C. Thereis a transaction that has modified one block of the table on which the index is being built in each cell, which committed after the create index began.

D. There are huge numbers of chained rows in the table on which the index is being built.

E. There is a ROWID column in the table on which the index is being built.

Correct Answer: AD

Explanation: A:It could be that row migration.

D:It could be that row migration or chained rows could cause it.

Note:

*Some facts about scans:

Scans exists in "OLTP" systems

Exadata smart scan requires a direct path read.

A direct path read is chosen at runtime based on internal heuristics The STORAGE clause in an explain plan doesnt necessarily mean you will perform a smart scan.

*The buffer caching in certain "OLTP" environments can occasionally induce conventional reads when smart scan is faster.

*Typically see cell multiblock physical read instead of cell smart table scan waits

*No one-size-fits-all solution can be given here but it is very fixable.

QUESTION 2

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

Correct Answer: ABD

Explanation: A: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.

B: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

D: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 3

You are using Hybrid Columnar Compression for a table stored in a tablespace that is contained in an Exadata-based ASM diskgroup. Identify three statements that correctly explain where the compression and decompression can be done.

A. Decompression can be done on the database servers.

B. Compression can be done on the Exadata storage servers.

C. Compression can be done on the database servers.

D. Decompression can be done on the Exadata storage servers.

Correct Answer: ABD

Explanation: B:

*Exadata storage provides an advanced compression technology, Hybrid Columnar Compression, that typically provides 10x, and higher, levels of data compression.

*The Exadata Storage Server (Exadata storage or Exadata cells) is used as the storage for the Oracle Database in the Database Machine. It runs the Exadata Storage Server

Software that provides the unique and powerful Exadata technology including Smart Scan, Smart

Flash Cache, Smart Flash Logging, IO Resource Manager, Storage Indexes and Hybrid Columnar Compression.

A, D:

*decompression

/Queries run directly on Hybrid Columnar Compressed data does not require the data to be decompressed

/Data that is required to satisfy a query predicate does not need to be decompressed; only the columns and rows being returned to the client are decompressed in memory

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

QUESTION 4

Which two are true about the use of DBFS in a Database Machine environment?

A. DBFS must be used to bulk load data into a database on the Database Machine if the staging area requires Exadata based shared storage.

B. DBFS must be used to have a POSIX compliant shared storage solution that is accessible from the database servers on a Database Machine.

C. DBFS must be used to bulk load data into a production database on the Database Machine.

D. DBFSmust use the DBFS_DG diskgroup for any DBFS store.

E. DBFS must be used to have a POSIX-compliant Exadata-based shared storage solution.

Correct Answer: AB

Note:

*external tables on DBFS file-systems provide the probably the most high-performance way to bulk load data into your database.

QUESTION 5

You are examiningthe existingIORM configurationon the cells of Database Machine, to see if they require mymodifications based on recent changesto variousworkloads. All seven cellsin your X3-2 half-rackshown thefollowing:

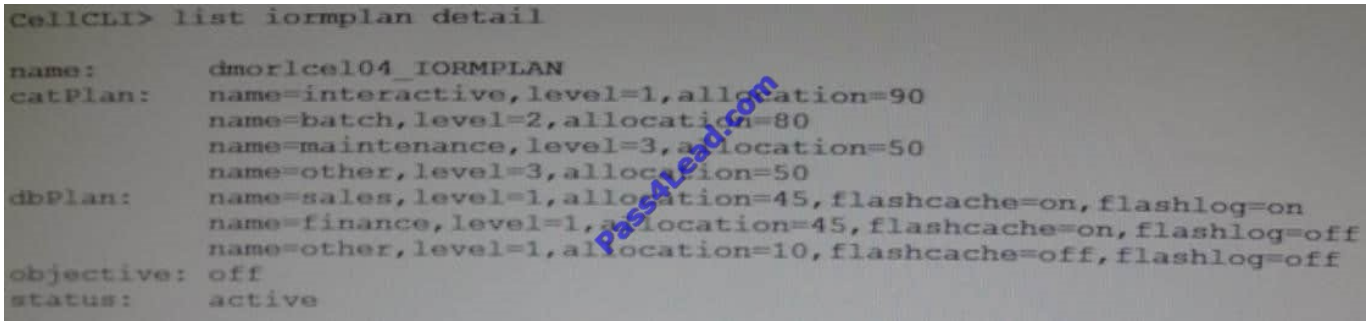

Which two are true about I/O to the cells using this plan?

A. I/O requests in the batch category may use flashcache if the I/O is from the sales finance database, and these I/O requests are guaranteed to get 80% of the I/O if the interactive category I/Os use no more than 20%.

B. I/O requests made by sessions in the marketing database may use flashing and flashcache if no other categories or database or database are using flashing and flashcache at the same time.

C. I/O requested in the interactive category may use flashdns if the I/O is from the sales or finance databases, and these I/O requests are guaranteed to get 90% of the I/O if the enough I/Os are issued in this category.

D. I/O requests from the sales database may use flashing regardless of the I/O category.

E. No I/Os in any category or from any database may use flashing or flashcache because the objective is off.

Correct Answer: DE

Explanation: list iormplan detail

Note:

*The I/O Resource Manager (IORM) extends the concept of resource groups with a new attribute known as a category. While resource groups allow DBRM to manage resources within a a database, categories provide I/O resource

management among multiple databases.

*cellcli> alter iormplan objective=\\'balanced\\'

-- {balanced | off | low_latency | high_throughput | auto

"off" simply turns off the IORM plan\\'s IO metering.