1Z0-070 Q&As

Oracle Exadata X5 Administration

Pass Oracle 1Z0-070 Exam with 100% Guarantee

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

https://www.pass4itsure.com/1z0-070.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 evaluating the performance of a SQL statement that accesses a very large table. You run this query:

```
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 10%');

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
```

Identify two reasons why the “cell physical IO interconnect bytes” statistic is greater than the “cell physical IO interconnect bytes returned by smart scan” statistic.

A. There is a transaction, which committed after the query began, 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.

B. There are chained rowsin the table, causing some “cell single block physical reads”to be requested by the database instance, resulting in additional I/O.

C. The table is a hash clustered table, causing “cell multiblock physical reads”to be requested by the database instance, resulting in additional I/O.

D. The table is list partitioned, causing “cell list of blocks physical reads”to be requested by the database instance, resulting in additional I/O.

E. There is a local index ona list partitioned tableon the column used in the WHERE clause, causing “cell list of blocks physical reads”to be requested by the database instance, resulting in additional I/O.

Correct Answer: CD

Explanation:

C: Scan on a clustered table can prevent a Smart Scan from occur.

D: Scan on an index-organized table can prevent a Smart Scan from occur.

Note: The Cell physical IO interconnect bytes returned by smart scan metric shows how many bytes of I/O

were returned by a smart scan to the database server.

References: https://uhesse.com/2011/01/19/exadata-part-i-smart-scan/

QUESTION 2

Which three are true concerning Exadata snapshot databases?

A. They are supported on non-container databases.

B. They are based on a read-write copy of an existing database.

C. They are integrated with the Multitenant architecture.

D. They can be created only on sparse ASM disk groups.

E. They don\\'t support all Exadata features.

F. They can be created on any type of ASM disk group.

Correct Answer: ACD

Explanation:

A: An Exadata snapshot database can be either a non-container database (non-CDB) or a container database (CDB). Creating an Exadata snapshot database of a CDB enables access to all of the pluggable databases in that container.

C: You can create two types of Exadata snapshots, depending on the current setup of your environment:

1.

You have a pluggable database (PDB) and want to create a test master from it.

2.

You have a container database (CDB) and want to create test masters from all its PDBs, or you have a simple non-container database and want to create a test master from it.

D: SPARSE disk group based database snapshots is functionality included in Exadata Storage Software. It requires Exadata Storage Software version 12.1.2.1.0 and Oracle Database 12c version 12.1.0.2 with bundle patch 5 or later. This feature is designed to work on native Exadata ASM storage disk groups. It uses ASM SPARSE grid disk based thin provisioning where snapshot databases created on a SPARSE disk group need only the space for changes plus some metadata, thereby enabling storage efficient snapshot databases.

References:

http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-server-snapshots.htm#SAGUG-GUIDE1D6EF45-36EF-40E3-A57E-F80B749E6122
http://www.oracle.com/technetwork/database/exadata/learnmore/exadata-database-copy-twp-2543083.pdf

QUESTION 3

A file contains a script with several EXACLI commands that must be executed on each cell in an X5 full rack.

The script must run on each cell simultaneously.

How would you achieve this?

A. Copy the script to all cells using the DCLI command and manually execute it on each cell using the DCLI command.

B. Copy the script to the cells using the EXACLI command and execute the script on all cells using the EXACLI command.

C. Execute it on all cells using the EXADCLI command specifying the file name containing the EXACLI script.

D. Copy the script to all cells using the CELLCLI command and execute it on all cells in parallel using the CELLCLI command.

E. Execute it on all cells using the DBMCLI command specifying the file name containing the EXACLI script.

Correct Answer: C

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.

References: http://docs.oracle.com/cd/E80920_01/DBMMN/exadcli.htm#DBMMN-GUID-4AE469A6-F291-4737-B975-F1B4B91D0BA0

References: http://docs.oracle.com/cd/E80920_01/DBMMN/exadcli.htm#DBMMN-GUID-4AE469A6-F291-4737-B975-F1B4B91D0BA0

QUESTION 4

You plan to migrate a database supporting both DSS and OLTP workloads to your new X5 Database Machine.

The workloads contain many complex aggregating functions and expensive joins on large partitioned tables in the DSS workload and indexed access for OLTP workloads.

Which three benefits accrue as a result of this migration?

A. Superior compression capability designed specifically for OLTP workloads

B. Columnar storage capability for data in row major data blocks that is held in flash cache

C. Superior compression capability designed specifically for data warehouse tables

D. Superior compression capability designed specifically for archival data

E. Superior flash cache compression technique

F. Cell offload processing for indexed-organized table access

Correct Answer: ACD

Explanation:

There are 6 different kinds of Table Data Compression methods:

1.

BASIC compression, introduced in Oracle 8 already and only recommended for Data Warehouse

2.

OLTP compression, introduced in Oracle 11 and recommended for OLTP Databases as well

3.

QUERY LOW compression (Exadata only), recommended for Data Warehouse with Load Time as a critical factor

4.

QUERY HIGH compression (Exadata only), recommended for Data Warehouse with focus on Space Saving

5.

ARCHIVE LOW compression (Exadata only), recommended for Archival Data with Load Time as a critical factor

6.

ARCHIVE HIGH compression (Exadata only), recommended for Archival Data with maximum Space Saving

Incorrect Answers:

B:

1.

BASIC compression, introduced in Oracle 8 already and only recommended for Data Warehouse

2.

OLTP compression, introduced in Oracle 11 and recommended for OLTP Databases as well

3.

QUERY LOW compression (Exadata only), recommended for Data Warehouse with Load Time as a critical factor

4.

QUERY HIGH compression (Exadata only), recommended for Data Warehouse with focus on Space Saving

5.

ARCHIVE LOW compression (Exadata only), recommended for Archival Data with Load Time as a critical factor

6.

ARCHIVE HIGH compression (Exadata only), recommended for Archival Data with maximum Space Saving

Method 1 and 2 are working by compressing the rows inside of Oracle Blocks – the more redundant values inside the block, the better the compression ratio. OLTP compression might be used with Exadata Database Machine, because Exadata V2 is designed to support OLTP as well as Data Warehouse. Methods 3 to 6 use Hybrid Columnar

Compression (HCC), which is designed to achieve better compression ratios by organizing sets of rows into compression units and then detaching the columns from the rows and collecting them. It is more likely that the same columns store similar (redundant) data, therefore the compression ratio improves with HCC.

References: https://uhesse.com/2011/01/21/exadata-part-iii-compression/

QUESTION 5

Which two completely prevent a Smart Scan from occurring?

A. querying a table containing many chained rows

B. querying a table containing many migrated rows

C. performing a minimum or maximum function on an indexed column

D. performing a Fast Full Index scan on a reverse key index

E. referencing more than 255 columns form an OLTP compressed table in a query

F. querying a table containing a JSON column

Correct Answer: BE

Explanation:

B: Migrated Rows is a special case of chained rows. Migrated rows still affect performance, as they do in conventional storage situations, but with the additional overhead of reducing the beneficial effects of “Smart Scan” in addition to increasing the number of I/Os

E: Smart Scans - broadly speaking and ignoring edge cases - can only transport a maximum of 254 columns from a single (non-HCC) segment. Requesting more columns will simply disable Smart Scans for that segment.

An interesting limitation to Exadata Smart Scans - if more than 254 columns from a table (not HCC compressed, more on that in moment) need to be projected, Smart Scans for that particular segment will be disabled and Exadata will fall back to conventional I/O. This means that the number of columns in the projection clause can make a significant difference to performance, since only Smart Scans allow taking advantage of offloading and particularly avoiding I/O via Storage Indexes.

Incorrect Answers:

A: Smart scan can cope with some cases of chained rows

References: http://oracle-randolf.blogspot.se/2013/01/exadata-smart-scan-projection-limitation.html