

1Z0-515^{Q&As}

Data Warehousing 11g Essentials

Pass Oracle 1Z0-515 Exam with 100% Guarantee

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

<https://www.pass2lead.com/1z0-515.html>

100% Passing Guarantee
100% Money Back Assurance

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

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers



QUESTION 1

Identify the benefit of using interval partitioning.

- A. Automatic creation of new partitions based on hash values
- B. Automatic creation of new partitions based on the value of data being entered
- C. Improved performance compared to range partitions
- D. Automatic transfer of older partitions lower cost storage

Correct Answer: B

Explanation:

Interval Partitioning was introduced in 11g, interval partitions are extensions to range partitioning. These provide automation for equi-sized range partitions. Partitions are created as metadata and only the start partition is made persistent. The additional segments are allocated as the data arrives. The additional partitions and local indexes are automatically created.

References:

QUESTION 2

What two types of results can be cached in the Result Set Cache?

- A. Results of an SQL query
- B. Results from a PL/SQL function
- C. Sequence object results
- D. Result sets derived from data dictionary tables

Correct Answer: AB

Explanation:

Your applications sometime send repetitive queries to the database. To improve the response time of repetitive queries, results of queries, query fragments, and PL/SQL functions can be cached in memory. A result cache stores the results of queries shared across all sessions. When these queries are executed repeatedly, the results are retrieved directly from the cache memory.

You must annotate a query or query fragment with a result cache hint to indicate that results are to be stored in the query result cache.

The query result set can be cached in the following ways:

*

Server-side Cache

*

Client Result Cache

Oracle Database 11g Release 1 (11.1) provides support for server-side Result Set caching for both JDBC types. The server-side result cache is used to cache the results of the current queries, query fragments, and PL/SQL functions in memory and then to use the cached results in future executions of the query, query fragment, or PL/SQL function. The cached results reside in the result cache memory portion of the SGA. A cached result is automatically invalidated whenever a database object used in its creation is successfully modified. The server-side caching can be of the following two types:

*

SQL Query Result Cache (A)

*

PL/SQL Function Result Cache (B)

QUESTION 3

Identify the true statement about REF partitions.

- A. REF partitions have no impact on partition-wise joins.
- B. Changes to partitioning in the parent table are automatically reflected in the child table.
- C. Changes in the data in a parent table are reflected in a child table.
- D. REF partitions can save storage space in the parent table.

Correct Answer: B

Explanation:

Reference partitioning is a partitioning method introduced in Oracle 11g. Using reference partitioning, a child table can inherit the partitioning characteristics from a parent table.

QUESTION 4

Which is NOT among Oracle SQL Analytic functions included in Oracle Database 11g?

- A. Ranking functions
- B. Substring functions

C. Window aggregate functions

D. LAG/LEAD functions

Correct Answer: B

Explanation:

Substring functions are not analytic.

QUESTION 5

Exadata uses smart scans, which are executed in_____.

A. Exadata Storage Server cells

B. Database Server node memory

C. Database Server node CPUs

D. Exadata does not use smart scans.

Correct Answer: A

Explanation: The Oracle Exadata Database Machine brings database performance to a whole new level, but have you ever wondered what exactly makes it so fast? Several components of the Oracle Exadata Database Machine, such as Oracle Database 11g Release 2; Oracle Exadata's Smart Flash Cache, Hybrid Columnar Compression, and SmartScan features; and InfiniBand interconnect, help deliver high performance. One of the key technologies that supports this performance is the storage index, which is not a regular database index. 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.

Oracle Exadata I/O and Smart Scan Storage in Oracle Exadata changes query processing so that not all blocks have to go to the database server for that server to determine which rows might satisfy a query. Oracle Exadata's Smart Scan feature enables certain types of query processing to be done in the storage cell. With Smart Scan technology, the database nodes send query details to the storage cells via a protocol known as iDB (Intelligent Database). With this information, the storage cells can take over a large portion of the data-intensive query processing. Oracle Exadata storage cells can search storage disks with added intelligence about the query and send only the relevant bytes, not all the database blocks, to the database nodes--hence the term smart scan.

References:

[1Z0-515 PDF Dumps](#)

[1Z0-515 Practice Test](#)

[1Z0-515 Braindumps](#)