

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.leads4pass.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

What are two ways in which query performance can be improved with partitioning?

- A. Partition pruning
- B. Partition optimization
- C. Partition compression
- D. Partition-wise joins

Correct Answer: AD

Explanation:

A: Even when you don't name a specific partition in a SQL statement, the fact that a table is partitioned might still influence the manner in which the statement accesses the table. When a SQL statement accesses one or more partitioned tables, the Oracle optimizer attempts to use the information in the WHERE clause to eliminate some of the partitions from consideration during statement execution. This process, called partition pruning, speeds statement execution by ignoring any partitions that cannot satisfy the statement's WHERE clause. To do so, the optimizer uses information from the table definition combined with information from the statement's WHERE clause.

D: A partition wise join is a join between (for simplicity) two tables that are partitioned on the same column with the same partitioning scheme. In shared nothing this is effectively hard partitioning locating data on a specific node / storage combo. In Oracle is is logical partitioning.

If you now join the two tables on that partitioned column you can break up the join in smaller joins exactly along the partitions in the data. Since they are partitioned (grouped) into the same buckets, all values required to do the join live in the equivalent bucket on either sides. No need to talk to anyone else, no need to redistribute data to anyone else... in short, the optimal join method for parallel processing of two large data sets.

QUESTION 2

Which unique method of improving performance is NOT used by the Oracle Exadata Database Machine?

- A. Flash to improve query performance
- B. Reduces the amount of data required to flow through I/O
- C. Increases the I/O using InfiniBand
- D. Performs analysis in a special in-memory database

Correct Answer: D

References:

QUESTION 3

What is the estimated maximum speed of data loads for a Quarter Rack with the Exadata Storage Server?

- A. 1 TB/hr
- B. 2 TB/hr
- C. 4 TB/hr
- D. 5 TB/hr
- E. It depends on the number of CPUs in the server.

Correct Answer: A

Explanation: References:

QUESTION 4

Why does partitioning help parallelism with RAC?

- A. The ability to do partition-wise joins reduces interconnect traffic.
- B. Partitioning allows you to split data storage across nodes.
- C. Partitioning reduces storage requirements.
- D. RAC will spawn additional parallel servers to meet the needs of requesting applications.

Correct Answer: A

Explanation:

Partition-wise joins reduce query response time by minimizing the amount of data exchanged among parallel execution servers when joins execute in parallel. This significantly reduces response time and improves the use of both CPU and memory resources. In Oracle Real Application Clusters (RAC) environments, partition-wise joins also avoid or at least limit the data traffic over the interconnect, which is the key to achieving good scalability for massive join operations.

Partition-wise joins can be full or partial. Oracle decides which type of join to use.

References:

QUESTION 5

Your BI tool (for example, Oracle Business Intelligence Enterprise Edition Plus, Business Objects and Cognos) will be used to query an Oracle database that includes the Oracle OLAP Option. What does tool generate in submitting queries that might include data stored in relational tables and OLAP cubes?

- A. SQL
- B. PL/SQL

C. Proprietary API code

D. SQL for relational and proprietary API code for OLAP

Correct Answer: B

Explanation:

DBMS_CUBE PL/SQL Package. In Database 11gR2, a new feature was added that allows cubes and dimensions to be entirely defined via PL/SQL calls, thus making it a much simpler job to automate the creation and refresh of cubes within the context of an application.

[Latest 1Z0-515 Dumps](#)

[1Z0-515 Practice Test](#)

[1Z0-515 Exam Questions](#)