



# 1Z0-515<sup>Q&As</sup>

Data Warehousing 11g Essentials

**Pass Oracle 1Z0-515 Exam with 100% Guarantee**

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

<https://www.geekcert.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

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.

---

### QUESTION 2

Which statement is true for you to get the benefits of partition-wise joins?

- A. The parent table must be partitioned on the join Key and the child table must be partitioned on] the join key.
- B. The parent table must be partitioned on the primary key and the child table must be partition the join key.
- C. The child table must use a reference partition.
- D. The parent table must be partitioned on the primary key and the child table must use a ref partition.

Correct Answer: A

Explanation:

Note:

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.



### QUESTION 3

What would you do to compress data in partitions that are frequently updated in Oracle Database 11g?

- A. Use Hybrid Columnar Compression.
- B. Use Advanced Compression Option.
- C. Use Hybrid Partitions.
- D. Avoid compressing any data.

Correct Answer: B

Explanation:

Advanced Compression features in Oracle Database 11g include:

\* Online Transaction Processing (OLTP) Table Compression: This breakthrough compression feature compresses table data during all types of data manipulation operations, including conventional INSERT or UPDATE. OLTP Table Compression leverages a sophisticated and intelligent algorithm that minimizes compression overhead during write operations, thereby making it viable for highly transactional workloads

References:

---

### QUESTION 4

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 5

Identify the true statement about a data warehouse

- A. The data warehouse is typically refreshed as often as a transactional system,
- B. Data warehouse queries are simpler than OLTP queries.
- C. A data warehouse typically contains historical data.
- D. Queries against a data warehouse never need summarized information.

Correct Answer: C

Explanation: A data warehouse is a relational database that is designed for query and analysis rather than for transaction processing. It usually contains historical data derived from transaction data, but it can include data from other sources. It separates analysis workload from transaction workload and enables an organization to consolidate data from several sources.

[1Z0-515 Practice Test](#)

[1Z0-515 Study Guide](#)

[1Z0-515 Braindumps](#)