



# 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.passapply.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 Cognos) will be used to query an Oracle database that includes the Oracle BI tool generate in submitting queries that might include data stored in cubes?

- A. SQL
- B. PIVSQL
- C. Proprietary API code
- D. SQL for relational and proprietary API code for OLAP

Correct Answer: A

Explanation:

Oracle Business Intelligence Enterprise Edition is most commonly used with the Oracle Database using SQL as the query language. Although the OLAP cube is a multidimensional data type, it is represented in the Oracle database as a collection of relational views and is easily queried by SQL.

Note #1: The wording of the question is strange. SQL can be used and is the first choice. So it seems to be the best answer.

Note #2: Oracle Business Intelligence Enterprise Edition (OBI EE) is a product suite based on the OBI EE Server. The OBI EE Server can map a logical business model to many different physical data sources and present the logical model for query to variety of client applications including Interactive Dashboards, Answers and Oracle Business Intelligence Plug-in for Microsoft Office.

References:

---

### QUESTION 2

You are looking to size a data warehouse configuration. If the I/O throughput for the CPUs is 25 GB/s, the I/O throughput for the HBA is 18 GB/s, and the I/O throughput for the disk subsystem is 6 GB/s, what is the overall throughput of the data warehouse?

- A. 25 GB/s
- B. 18 GB/s
- C. 6 GB/s
- D. It depends on how many processors are in the servers.

Correct Answer: C



Explanation:

In this scenario the disk subsystem is the bottleneck. It determines the throughput.

Note: Each of the components must provide sufficient I/O bandwidth to ensure a well-balanced I/O system.

The end-to-end I/O system consists of more components than just the CPUs and disks. A well-balanced I/O system must provide approximately the same bandwidth across all components in the I/O system.

These components include:

\*

Host bus adapters (HBAs), the connectors between the server and the storage.

\*

Switches, in between the servers and a storage area network (SAN) or network attached storage (NAS).

\*

Ethernet adapters for network connectivity (GigE NIC or Infiniband). In an Oracle Real Application Clusters (Oracle RAC) environment, you need an additional private port for the interconnect between the nodes that you should not include when sizing the system for I/O throughput. The interconnect must be sized separately, taking into account factors such as internode parallel execution.

\*

Wires that connect the individual components.

References:

---

### QUESTION 3

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 4

What are three advantages provided by proper partitioning in a data warehouse?

- A. Partition pruning will occur
- B. Faster sorting
- C. Efficient parallel joins
- D. Efficient data loading
- E. Reduced disk usage

Correct Answer: ACD

Explanation:

There are three major advantages of partitioning.

\*

Partition Pruning - Oracle only accesses a limited set of table partitions if the FROM and WHERE clause



permit it to.

\*

Partition-wise Joins - Where two tables that have compatible partitioning schemes are joined, Oracle improves the efficiency of parallel operations by performing the join between individual partitions of the tables.

\*

Manageability - Partitioning allows DDL operations on a large subset of table rows with some element of commonality defined through the partitioning type.

References:

---

#### QUESTION 5

What are Oracle Data Integrator templates used for?

- A. To model SAP applications
- B. To define how to transform data
- C. As reports to monitor ETL activity
- D. None of these

Correct Answer: B

Explanation: Oracle Data Integrator streamlines the highperformance movement and transformation of data between disparate systems in batch, real-time, synchronous, and asynchronous modes. Knowledge Modules are at the core of the Oracle Data Integrator architecture. They make all Oracle Data Integrator processes modular, flexible, and extensible. Knowledge Modules implement the actual data flows and define the templates for generating code across the multiple systems involved in each process. Knowledge Modules are generic, because they allow data flows to be generated regardless of the transformation rules. And they are highly specific, because the code they generate and the integration strategy they implement are finely tuned for a given technology. Oracle Data Integrator provides a comprehensive library of Knowledge Modules, which can be tailored to implement existing best practices (for example, for highest performance, for adhering to corporate standards, or for specific vertical know-how). By helping companies capture and reuse technical expertise and best practices, Oracle Data Integrator's Knowledge Module framework reduces the cost of ownership. It also enables metadata-driven extensibility of product functionality to meet the most demanding data integration challenges.

References:

[1Z0-515 VCE Dumps](#)

[1Z0-515 Study Guide](#)

[1Z0-515 Braindumps](#)