



# 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

What would you use to evenly distribute data across the disk in your Oracle data warehouse?

- A. Range Partitioning
- B. Automatic Storage Management (ASM)
- C. List Partitioning
- D. RAC

Correct Answer: B

Explanation: Automatic Storage Management (ASM) is a feature provided by Oracle Corporation within the Oracle Database from release Oracle 10g (revision 1) onwards. ASM aims to simplify the management of database files. To do so, it provides tools to manage file systems and volumes directly inside the database, allowing database administrators (DBAs) to control volumes and disks with familiar SQL statements in standard Oracle environments. Thus DBAs do not need extra skills in specific file systems or volume managers (which usually operate at the level of the operating system).

With ASM:

\*

IO channels can take advantage of data striping and software mirroring

\*

DBAs can automate online redistribution of data, along with the addition and removal of disks/storage

\*

the system maintains redundant copies and provides 3rd-party[citation needed] RAID functionality

\*

Oracle supports third-party multipathing IO technologies (such as failover or load balancing to SAN access) the need for hot spares diminishes

References:

---

### QUESTION 2

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 3

You can use Oracle Data Mining unstructured data.

A. TRUE

B. FALSE

Correct Answer: A

Explanation: Data that cannot be meaningfully interpreted as numerical or categorical is considered unstructured for purposes of data mining. It has been estimated that as much as 85% of enterprise data falls into this category. Extracting meaningful information from this unstructured data can be critical to the success of a business.

Unstructured data may be binary objects, such as image or audio files, or text objects, which are language-based. Oracle Data Mining supports text objects. Text must undergo a transformation process before it can be mined. Once the data has been properly transformed, the case table can be used for building, testing, or scoring data mining models. Most Oracle Data Mining algorithms support text

References:

---

### QUESTION 4

Identify the type of refresh that is NOT supported by materialized views.

A. Deferred

B. Incremental

C. Full

D. Heuristic

Correct Answer: D

Explanation:

Use the CREATE MATERIALIZED VIEW statement to create a materialized view. A materialized view is a database object that contains the results of a query.



Incorrect answer:

A: Specify DEFERRED to indicate that the materialized view is to be populated by the next REFRESH operation.

B: Oracle Database uses the default index to speed up incremental (FAST) refresh of the materialized view.

C: By default, Oracle Database creates a primary key materialized view with refresh on demand only. If a materialized view log exists on the table, then the column can be altered to be capable of fast refresh. If no such log exists, then only full refresh of the column is possible.

References:

---

#### QUESTION 5

How many Exadata Storage Server cells are there in a Full Rack Exadata database machine configuration that has 8 Database Server nodes?

- A. 2
- B. 14
- C. 16
- D. 24

Correct Answer: B

Explanation:

[Latest 1Z0-515 Dumps](#)

[1Z0-515 Study Guide](#)

[1Z0-515 Braindumps](#)