



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.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

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 2

Identify the statement about Oracle OLAP that is NOT true.

- A. Oracle OLAP cubes are stored in the Oracle relational database
- B. Oracle OLAP uses standard Oracle database security.
- C. Meta data for Oracle OLAP is accessible in an external data dictionary
- D. Oracle OLAP can be deployed using RAC.

Correct Answer: C

Explanation:



All metadata for cubes and dimensions is stored in the Oracle database.

References:

QUESTION 3

You want partitions to be automatically created when data that does not fit into current date range loaded. Which type of partitioning would you implement?

- A. Hash
- B. List
- C. Invisible
- D. Interval

Correct Answer: D

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.

Note: Partitioning is one of the most sought after options for data warehousing. Almost all Oracle data warehouses use partitioning to improve the performance of queries and also to ease the day-to-day maintenance complexities. Starting with 11G, more partitioning options have been provided and these should reduce the burden of the DBA to a great extent.

References:

QUESTION 4

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:

QUESTION 5

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:



VCE & PDF

GeekCert.com

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

2024 Latest geekcert 1Z0-515 PDF and VCE dumps Download

[Latest 1Z0-515 Dumps](#)

[1Z0-515 Exam Questions](#)

[1Z0-515 Braindumps](#)