

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.pass2lead.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

Which can be used in scenario where there are large data loads of a sensitive nature into a data warehouse?

- A. Direct path loading
- B. External tables for loading flat files
- C. Partition exchange loading
- D. Any of these are valid for certain situations.

Correct Answer: A

Explanation: Instead of filling a bind array buffer and passing it to the Oracle database with a SQL INSERT statement, a direct path load uses the direct path API to pass the data to be loaded to the load engine in the server. The load engine builds a column array structure from the data passed to it. The direct path load engine uses the column array structure to format Oracle data blocks and build index keys. The newly formatted database blocks are written directly to the database (multiple blocks per I/O request using asynchronous writes if the host platform supports asynchronous I/O).

Internally, multiple buffers are used for the formatted blocks. While one buffer is being filled, one or more buffers are being written if asynchronous I/O is available on the host platform. Overlapping computation with I/O increases load performance.

http://download.oracle.com/docs/cd/B19306_01/server.102/b14215/ldr_modes.htm#i1008815

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

How can you use Oracle Data Mining with Oracle Warehouse builder?

- A. To identify records to extract
- B. As a standard transform operation
- C. To increase write performance
- D. To eliminate ETL logging

Correct Answer: A

Explanation: Data Mining and Data Warehousing Data can be mined whether it is stored in flat files, spreadsheets, database tables, or some other storage format. The important criteria for the data is not the storage format, but its applicability to the problem to be solved. Proper data cleansing and preparation are very important for data mining, and a data warehouse can facilitate these activities. However, a data warehouse will be of no use if it does not contain the data you need to solve your problem. Oracle Data Mining requires that the data be presented as a case table in single-record case format. All the data for each record (case) must be contained within a row. Most typically, the case table is a view that presents the data in the required format for mining Note: Oracle Warehouse Builder (OWB) enables the design and deployment of enterprise data warehouses, data marts, and e-business intelligence applications.

References:

QUESTION 4

Identify the benefit of using interval partitioning.

- A. Automatic creation of new partitions based on hash values
- B. Automatic creation of new partitions based on the value of data being entered

- C. Improved performance compared to range partitions
- D. Automatic transfer of older partitions lower cost storage

Correct Answer: B

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.

References:

QUESTION 5

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:

[1Z0-515 PDF Dumps](#)

[1Z0-515 VCE Dumps](#)

[1Z0-515 Braindumps](#)