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.certbus.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

What data can you compress using Advanced Compression in Oracle Database 11g?

A. Read only data

B. Data that can be updated, inserted and/or deleted (DML)

C. Only data being archived

D. Data warehousing data

Correct Answer: B

Explanation:

Oracle Database 11g has new option named as Oracle Advanced Table Compression option which aims

at reducing space occupied by data for both OLTP and warehouse databases. This option provides the

following types of compression:

*

Compression of data tables even for OLTP environment. (Previous versions had compression option for tables that are mostly read only).

*

Compression of unstructured data in SecureFiles.

*

Compression of RMAN backups.

*

Compression in Data Pump Export files.

*

Compression of redo data transmitted to a standby database during redo gap resolution (when data guard is configured).

QUESTION 3

You can perform what-if analysis of potential changes with Oracle Warehouse Builder.

A. TRUE

B. FALSE

Correct Answer: A

Explanation: The Metadata Dependency Manager (MDM) enables you to plan your project by previewing the impact of the changes or future changes for "what-if" analysis. When you plan to introduce changes to your source systems, you can gauge the impact of that change on your warehouse design. If changes have already been introduced, then you can plan the time required to update your ETL design and rebuild your data warehouse.

References:

QUESTION 4

Which is NOT among Oracle SQL Analytic functions included in Oracle Database 11g?

A. Ranking functions

B. Substring functions

C. Window aggregate functions

D. LAG/LEAD functions

Correct Answer: B

Explanation:

Substring functions are not analytic.

QUESTION 5

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 6

Indentify the true statement about REF partitions.

A. REF partitions have no impact on partition-wise joins.

B. Changes to partitioning in the parent table are automatically reflected in the child table.

C. Changes in the data in a parent table are reflected in a child table.

D. REF partitions can save storage space in the parent table.

Correct Answer: B

Explanation:

Reference partitioning is a partitioning method introduced in Oracle 11g. Using reference partitioning, a

child table can inherit the partitioning characteristics from a parent table.

QUESTION 7

How many Exadata Storage Server cells can be used in a grid?

A. 7

B. 14

C. 128

D. No practical limit

Correct Answer: D

Explanation:

There is no practical limit to number of cells that can be in the grid.

References:

QUESTION 8

Your customer wants to implement an ILM strategy. The customer must have which option when deploying Oracle\\'s ILM Assistant to implement this strategy?

A. RAC

B. Partitioning

C. OLAP

D. Oracle Clusterware

Correct Answer: B

Explanation: Information Lifecycle Management (ILM) is a set of policies and procedures for managing data during its lifetime. The ILM Assistant manages information by recommending the correct placement of data on logical storage tiers as specified by a lifecycle definition, where a lifecycle definition describes the stages and storage tiers that data resides on during its lifetime. Each stage specifies a retention period during which the data resides on a logical storage tier. A logical storage tier is a collection of Oracle tablespaces in which partitions may reside.

Note: Information today comes in a wide variety of types, for example an E-mail message, a photograph, or an order in an Online Transaction Processing System. Therefore, once you know the type of data and how it will be used, you already have an understanding of what its evolution and final destiny is likely to be.

One of the challenges facing each organization is to understand how its data evolves and grows, monitor how its usage changes over time, and decide how long it should survive, while adhering to all the rules and regulations that now apply

to that data. Information Lifecycle Management (ILM) is designed to address these issues, with a combination of processes, policies, software, and hardware so that the appropriate technology can be used for each stage in the lifecycle of the data.

References:

QUESTION 9

You customer wants to segment their customers1 demographic data into those that use and do not use loyalty card. What would you recommend?

A. Use Oracle OLAP Option.

B. Use Oracle SQL Analytic Functions.

C. Use classification algorithm in Oracle Data Mining.

D. Use non-negative matrix factorization in Oracle Data Mining.

Correct Answer: C

Explanation:

Classification is a data mining function that assigns items in a collection to target categories or classes. The goal of classification is to accurately predict the target class for each case in the data. For example, a

classification model could be used to identify loan applicants as low, medium, or high credit risks.

The simplest type of classification problem is binary classification. In binary classification, the target

attribute has only two possible values: for example, high credit rating or low credit rating

Note:

Oracle Data Mining provides the following algorithms for classification:

*

Decision Tree

Decision trees automatically generate rules, which are conditional statements that reveal the logic used to

build the tree.

*

Naive Bayes

Naive Bayes uses Bayes\\' Theorem, a formula that calculates a probability by counting the frequency of

values and combinations of values in the historical data.

*

Generalized Linear Models (GLM)

GLM is a popular statistical technique for linear modeling. Oracle Data Mining implements GLM for binary classification and for regression. GLM provides extensive coefficient statistics and model statistics, as well as row diagnostics. GLM also supports confidence bounds.

*

Support Vector Machine

Support Vector Machine (SVM) is a powerful, state-of-the-art algorithm based on linear and nonlinear regression. Oracle Data Mining implements SVM for binary and multiclass classification.

References:

QUESTION 10

Your customer wants to determine "market baskets." What do you recommend?

A. Use Oracle OLAP Option.

B. Use Oracle SQL Analytic Functions.

C. Use associations algorithm in Oracle Data Mining.

D. Use regression analysis in Oracle Data Mining

Correct Answer: C

Explanation:

Association is a data mining function that discovers the probability of the co-occurrence of items in a collection. The relationships between co-occurring items are expressed as association rules.

Market-Basket Analysis

Association rules are often used to analyze sales transactions. For example, it might be noted that customers who buy cereal at the grocery store often buy milk at the same time. In fact, association analysis might find that 85% of the checkout sessions that include cereal also include milk. This relationship could be formulated as the following rule.

Cereal implies milk with 85% confidence

This application of association modeling is called market-basket analysis. It is valuable for direct marketing, sales promotions, and for discovering business trends. Market-basket analysis can also be used effectively for store layout, catalog design, and cross-sell.

Association Algorithm

Oracle Data Mining uses the Apriori algorithm to calculate association rules for items in frequent itemsets.

References:

QUESTION 11

Which feature of Oracle Warehouse Builder can be used to help ensure data quality?

A. Data extraction

B. Data profiling

C. Logical mapping

D. Exception reporting

Correct Answer: B

Explanation: After you connect to your data sources in Oracle Warehouse Builder (including Oracle databases, sources accessed through gateways, and flat file sources) you can apply full-featured data profiling to generate statistics about data quality, and to discover complex patterns, foreign key relationships, and functional dependencies. You can then design complex data rules and create data auditors to monitor compliance with those rules in any source or target system in your landscape, regardless of whether those sources are loaded using Oracle Warehouse Builder or other ETL tools.

For customers who have selected solutions other than Oracle Warehouse Builder for data profiling and data quality, these can be applied independently of Oracle Warehouse Builder ETL and design features.

Note: Oracle Warehouse Builder is a full-featured data integration, data warehousing, data quality and metadata management solution designed for the Oracle database. Oracle Warehouse Builder is an integral part of Oracle Database 11g Release 2 (11.2) and is installed as part of every database installation (other than Oracle Database XE).

The major feature areas of Oracle Warehouse Builder include: *Data modeling *Extraction, Transformation, and Load (ETL) *Data profiling and data quality *Metadata management *Business-level integration of ERP application data *Integration with Oracle business intelligence tools for reporting purposes *Advanced data lineage and impact analysis Oracle Warehouse Builder is also an extensible data integration and data quality solutions platform. Oracle Warehouse Builder can be extended to manage metadata specific to any application, and can integrate with new data source and target types, and implement support for new data access mechanisms and platforms, enforce your organization\\'s best practices, and foster the reuse of components across solutions.

References:

QUESTION 12

What areas can SQL Access Advisor give advice on?

A. Partitioning advice, index advice, and materialized views advice

B. Index advice and compression advice

C. Index advice and data masking advice

D. Partitioning advice and compression advice

Correct Answer: A

Explanation:

The SQL Access Advisor was introduced in Oracle 10g to make suggestions about additional indexes and materialized views which might improve system performance. Oracle 11g has made two significant changes to the SQL Access Advisor:

The advisor now includes advice on partitioning schemes that may improve performance. The original workload manipulation has been deprecated and replaced by SQL tuning sets.

References:

QUESTION 13

You are looking for some general design principles that could be used in designing every large scale data warehouse you create. Identify the principle that would have the widest applicability.

A. Partition your tables appropriately to produce partition-wise joins.

B. Always use a star schema or snowflake schema design.

C. Do as much analytics as possible in your BI tools.

D. Always use Oracle OLAP.

Correct Answer: A

Explanation:

Partition-wise joins can be full or partial. Oracle decides which type of join to use. A full partition-wise join divides a large join into smaller joins between a pair of partitions from the two joined tables. To use this feature, you must equipartition both tables on their join keys, or use reference partitioning.

Oracle Database can perform partial partition-wise joins only in parallel. Unlike full partition-wise joins, partial partition-wise joins require you to partition only one table on the join key, not both tables.

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.

References:

QUESTION 14

How does compression affect resource utilization?

A. Reduces the amount of CPU and disk utilization

B. Increases the amount of CPU and disk utilization

C. Reduces the amount of disk but increases CPU utilization for loading

D. Increases the amount of disk but reduces CPU utilization for loading!

Correct Answer: C

Explanation:

Compression is useful because it helps reduce the consumption of resources such as data space or transmission capacity. Because compressed data must be decompressed to be used, this extra processing imposes computational or other costs through decompression.

QUESTION 15

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 Practice Test 1Z0-515 Exam Questions