**Vendor:**Oracle

**Exam Code:**1Z0-515

**Exam Name:**Data Warehousing 11g Essentials

**Version:**Demo

**QUESTION 1**

The most performant way to load data from an external table that will also guarantee direct path loading is:

A. Using Create Table as Select (CTAS)

B. Using Data Pump

C. Using Insert as Select (IAS)

D. Using transparent gateways

Correct Answer: A

Explanation: CTAS refers to a CREATE TABLE AS statement - a new table is created and populated with row from a specified query.

The most common uses of CTAS are in these scenarios:

*

Creating a table identical to another table in structure, but with a filter criteria applied to its data.

*

Creating a table with small structural differences from an existing table.

For best performance use Direct Path Load. The conventional path uses standard insert statements whereas the direct path loader loads directly into the Oracle data files and creates blocks in Oracle database block format. During direct-path INSERT operations, the database appends the inserted data after existing data in the table. Data is written directly into datafiles, bypassing the buffer cache. Free space in the existing data is not reused, and referential integrity constraints are ignored. These procedures combined can enhance performance.

References:

---

**QUESTION 2**

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 3**

Your customer is looking to implement ad-hoc analysis in a data warehouse. Which approach is least likely to be used assuming that the customer does not want the expense of managing view?

A. Star schema

B. Snowflake schema

C. Third normal form schema

D. OLAP

Correct Answer: C

Explanation: Data warehouses often use denormalized or partially denormalized schemas (such as a star schema) to optimize query performance. On the other hand OLTP (Online Transaction Processing) systems often use fully normalized schemas to optimize update/insert/delete performance, and to guarantee data consistency.

References:

---

**QUESTION 4**

Which best describes Oracle\\'s OLAP Option for Oracle Database 11g Release 2?

A. Is stored as relational tables and is considered a ROLAP solution

B. Uses bitmap indexes

C. Physically stores OLAP cubes as objects within the relational database

D. Is available both within the Oracle Database and as a stand-alone solution

Correct Answer: A

Explanation: Oracle OLAP is a world class multidimensional analytic engine embedded in Oracle Database 11g. Oracle OLAP cubes deliver sophisticated calculations using simple SQL queries - producing results with speed of thought response times. This outstanding query performance may be leveraged transparently when deploying OLAP cubes as materialized views enhancing the performance of summary queries against detail relational tables. Because Oracle OLAP is embedded in Oracle Database 11g, it allows centralized management of data and business rules in a secure, scalable and enterprise-ready platform.

---

**QUESTION 5**

Knowledge Modules are:

A. Reusable code templates for Oracle Data Integrator

B. Prebuilt applications for Oracle Business Intelligence

C. Options for Oracle Enterprise Manager

D. Algorithms for data mining

Correct Answer: A

Explanation:

Knowledge modules (KMs) in Oracle Data Integrator are components that implement reusable

transformation and ELT (extract, load, and transform) strategies across different technologies.

References:

---

**QUESTION 6**

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 7**

You are looking to create a RAC cluster to deliver high performance for your client\\'s data warehouse. Which statement is true about a configuration with a few large nodes versus a configuration with many smaller nodes?

A. A few large nodes always perform better than many small nodes.

B. A few large nodes always perform worse than many small nodes.

C. It depends on the workload specifics and the effect of a node failure.

D. Performance should be the same with either option.

Correct Answer: D

Explanation:

---

**QUESTION 8**

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:

---

**QUESTION 9**

For data warehousing, identify the benefits that would NOT be provided by the use of RAC.

A. Distribute workload across all the nodes.

B. Distribute workload to some of the nodes.

C. Provide parallel query servers.

D. Provide high availability for all the operations.

Correct Answer: B

Explanation:

With Oracle RAC the workload can be distributed access all cluster nodes, parallel query servers can be provided through the Parallel Query tool, and high availability can be obtained through, for example, Oracle Clusterware.

Note: Oracle RAC (Real Application Clusters) is a cluster database with a shared cache architecture that overcomes the limitations of traditional shared-nothing and shared-disk approaches to provide highly scalable and available database solutions for all your business applications. Oracle RAC is a key component of Oracle\\'s private cloud architecture. Oracle RAC support is included in the Oracle Database Standard Edition for higher levels of system uptime.
References:

---

**QUESTION 10**

Which two statements are true about the advantages of using a data warehouse?

A. A data warehouse uses fewer database structures, so access to answers is faster and easier

B. A data warehouse is typically implemented with a different design, making access faster.

C. A data warehouse is optimized for ongoing write activity, making response faster.

D. A data warehouse uses specialized features of the Oracle database, like materialized views and star transformations, making response faster.

Correct Answer: BD

Explanation:

Data warehouses often use denormalized or partially denormalized schemas (such as a star schema) to

optimize query performance.

Note: A materialized view is a pre-computed table comprising aggregated or joined data from fact and

possibly dimension tables. Also known as a summary or aggregate table.

References:

---

**QUESTION 11**

What is the estimated maximum speed of data loads for a Quarter Rack with the Exadata Storage Server?

A. 1 TB/hr

B. 2 TB/hr

C. 4 TB/hr

D. 5 TB/hr

E. It depends on the number of CPUs in the server.

Correct Answer: A

Explanation: References:

---

**QUESTION 12**

Why does partitioning help parallelism with RAC?

A. The ability to do partition-wise joins reduces interconnect traffic.

B. Partitioning allows you to split data storage across nodes.

C. Partitioning reduces storage requirements.

D. RAC will spawn additional parallel servers to meet the needs of requesting applications.

Correct Answer: A

Explanation:

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.

Partition-wise joins can be full or partial. Oracle decides which type of join to use.

References: