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.pass4itsure.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

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:

QUESTION 2

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 3

Exadata uses smart scans, which are executed in_____.

A. Exadata Storage Server cells

B. Database Server node memory

C. Database Server node CPUs

D. Exadata does not use smart scans.

Correct Answer: A

Explanation: The Oracle Exadata Database Machine brings database performance to a whole new level, but have you ever wondered what exactly makes it so fast? Several components of the Oracle Exadata Database Machine, such as Oracle Database 11g Release 2; Oracle Exadata\\'s Smart Flash Cache, Hybrid Columnar Compression, and SmartScan features; and InfiniBand interconnect, help deliver high performance. One of the key technologies that supports this performance is the storage index, which is not a regular database index. Storage indexes reside in the memory of the storage servers--also called storage cells--and significantly reduce unnecessary I/O by excluding irrelevant database blocks in the storage cells.

Oracle Exadata I/O and Smart Scan Storage in Oracle Exadata changes query processing so that not all blocks have to go to the database server for that server to determine which rows might satisfy a query. Oracle Exadata\\'s Smart Scan feature enables certain types of query processing to be done in the storage cell. With Smart Scan technology, the database nodes send query details to the storage cells via a protocol known as iDB (Intelligent Database). With this information, the storage cells can take over a large portion of the data- intensive query processing. Oracle Exadata storage cells can search storage disks with added intelligence about the query and send only the relevant bytes, not all the database blocks, to the database nodes--hence the term smart scan.

References:

QUESTION 4

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 5

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 6

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 7

Which is NOT an available composite partition in Oracle Database 11g?

A. range-list

B. list-list

C. list-range

D. interval-hash

Correct Answer: D

Explanation:

Extended Composite Partitioning

In previous releases of Oracle, composite partitioning was limited to Range-Hash and Range-List partitioning. Oracle 11g Release 1 extends this to allow the following composite partitioning schemes:

Range-Hash (available since 8i)

Range-List (available since 9i)

Range-Range

List-Range

List-Hash

List-List

Note: inverval-hash is a valid Interval partitioning.

References:

QUESTION 8

Identify the action that you CANNOT perform using Database Resource Manager.

A. Define Consumer Groups.

B. Create rules to map sessions to Consumer Groups.

C. Define a Resource Plan.

D. Allocate individual CPUs to Consumer Groups.

Correct Answer: D

Explanation:

Oracle Database Resource Management (DRM) provides tools that allow any Oracle DBA to manage a database server\\'s CPU resources effectively for application user groups and during different resource

demand periods.

DRM consists of four basic components:

*Resource Consumer Groups (not A). A resource consumer group is a collection of users with similar requirements for resource consumption. Users can be assigned to more than one resource consumer group, but each user\\'s active session can only be assigned to one resource consumer group at a time.

*Resource Plans (not C). In its simplest form, a resource plan describes the resources allocated to one or more resource consumer group(s).

*Resource Plan Directives (not B). Resource plan directives allocate resources among the resource consumer groups in the resource plan. Essentially, directives connect resource consumer groups or subplans to their resource plans.

* SYSTEM_PLAN. Oracle supplies an initial, default resource plan named SYSTEM_PLAN. This plan implements a CPU utilization resource allocation method to divide and prioritize CPU resources to three resource consumer groups

QUESTION 9

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 10

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 11

Identify the benefit of using bitmap join indexes. Select one.

A. Faster query performance for all queries.

B. Reduced space for indexes.

C. Faster query performance for some queries.

D. Lower memory usage.

Correct Answer: B

Explanation:

Oracle benchmarks claim that bitmap join indexes can run a query more than eight times faster than traditional indexing methods.

However, this speed improvement is dependent upon many factors, and the bitmap join is not a panacea.

Some restrictions on using the bitmap join index include:

The indexed columns must be of low cardinality--usually with less than 300 distinct values. The query must not have any references in the WHERE clause to data columns that are not contained in the index.

The overhead when updating bitmap join indexes is substantial. For practical use, bitmap join indexes are dropped and rebuilt each evening about the daily batch load jobs. This means that bitmap join indexes are useful only for Oracle data warehouses that remain read-only during the processing day.

References:

QUESTION 12

Your BI tool (for example, Oracle Business Intelligence Enterprise Edition Cognos) will be used to query an Oracle database that includes the Oracle BI tool generate in submitting queries that might include data stored in cubes?

A. SQL

B. PIVSQL

C. Proprietary API code

D. SQL for relational and proprietary API code for OLAP

Correct Answer: A

Explanation:

Oracle Business Intelligence Enterprise Edition is most commonly used with the Oracle Database using SQL as the query language. Although the OLAP cube is a multidimensional data type, it is represented in the Oracle database as a collection of relational views and is easily queried by SQL.

Note #1: The wording of the question is strange. SQL can be used and is the first choice. So it seems to be the best answer.

Note #2: Oracle Business Intelligence Enterprise Edition (OBI EE) is a product suite based on the OBI EE Server. The OBI EE Server can map a logical business model to many different physical data sources and present the logical model for query to variety of client applications including Interactive Dashboards,

Answers and Oracle Business Intelligence Plug-in for Microsoft Office.

References:

QUESTION 13

For which type of query is the SQL result cache automatically disabled?

A. Queries that access data which changes frequently

B. Queries that return large amounts of data

C. Queries that use SQL functions such as SYSDATE

D. Queries that are used infrequently

Correct Answer: C

Explanation:

SYSDATE produces a new value every time it is used. Caching such a value would make no sense.

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

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: