



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**

You can use Oracle Data Mining unstructured data.

- A. TRUE
- B. FALSE

Correct Answer: A

Explanation: Data that cannot be meaningfully interpreted as numerical or categorical is considered unstructured for purposes of data mining. It has been estimated that as much as 85% of enterprise data falls into this category. Extracting meaningful information from this unstructured data can be critical to the success of a business.

Unstructured data may be binary objects, such as image or audio files, or text objects, which are language-based. Oracle Data Mining supports text objects. Text must undergo a transformation process before it can be mined. Once the data has been properly transformed, the case table can be used for building, testing, or scoring data mining models. Most Oracle Data Mining algorithms support text

References:

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

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 4

Which questions CANNOT be addressed by Oracle Data Mining?

- A. Fraud detection
- B. Prediction of customer behavior
- C. Root cause de
- D. Identify factors associated with a business problem

Correct Answer: C

Explanation:

Data Mining can provide valuable results:

*Predict customer behavior (Classification) (not B)

*Predict or estimate a value (Regression)

*Segment a population (Clustering)

*Identify factors more associated with a business problem (Attribute Importance) (not D)

*

Find profiles of targeted people or items (Decision Trees)

*

Determine important relationships and market baskets within the population (Associations)

*

Find fraudulent or rare events (Anomaly Detection) (not A)

References:

QUESTION 5

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 6

What are two ways in which query performance can be improved with partitioning?

- A. Partition pruning
- B. Partition optimization
- C. Partition compression
- D. Partition-wise joins

Correct Answer: AD

Explanation:

A: Even when you don't name a specific partition in a SQL statement, the fact that a table is partitioned might still influence the manner in which the statement accesses the table. When a SQL statement accesses one or more partitioned tables, the Oracle optimizer attempts to use the information in the WHERE clause to eliminate some of the partitions from consideration during statement execution. This process, called partition pruning, speeds statement execution by ignoring any partitions that cannot satisfy the statement's WHERE clause. To do so, the optimizer uses information from the table definition combined with information from the statement's WHERE clause.

D: A partition wise join is a join between (for simplicity) two tables that are partitioned on the same column with the same partitioning scheme. In shared nothing this is effectively hard partitioning locating data on a specific node / storage combo. In Oracle is is logical partitioning.

If you now join the two tables on that partitioned column you can break up the join in smaller joins exactly along the partitions in the data. Since they are partitioned (grouped) into the same buckets, all values required to do the join live in



the equivalent bucket on either sides. No need to talk to anyone else, no need to redistribute data to anyone else... in short, the optimal join method for parallel processing of two large data sets.

QUESTION 7

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 8

What is the difference between an ETL (Extraction Transformation Load) approach and an ELT (Extraction Load Transformation) approach to data integration? Select one.

- A. ETL can operate between heterogeneous data sources.



- B. ELT requires a separate transformation server.
- C. ELT transforms data on the target server.
- D. ELT cannot be used for incremental data loading.

Correct Answer: C

Explanation:

There are two approaches to consider for data integration: ELT and ETL. The difference between ETL and ELT lies in the environment in which the data transformations are applied. In traditional ETL, the transformation takes place when the data is en route from the source to the target system. In ELT, the data is loaded into the target system, and then transformed within the target system environment.

References:

QUESTION 9

You think that result set caching might provide some benefits for your current data warehouse scenario. You perform some analysis on the composition of the queries used in the scenario. Identify the result of the analysis that would indicate the most potential for improvement with result set caching.

- A. The scenario consists mainly of queries that are used infrequently.
- B. The scenario consists mainly of queries that work on data which changes frequently.
- C. The scenario consists mainly of queries with long run times and small result sets.
- D. All data warehouse scenarios will benefit from result set caching.

Correct Answer: C

Explanation: As its name suggests, the query result cache is used to store the results of SQL queries for re-use in subsequent executions. By caching the results of queries, Oracle can avoid having to repeat the potentially time-consuming and intensive operations that generated the resultset in the first place (for example, sorting/aggregation, physical I/O, joins etc). The cache results themselves are available across the instance (i.e. for use by sessions other than the one that first executed the query) and are maintained by Oracle in a dedicated area of memory. Unlike our homegrown solutions using associative arrays or global temporary tables, the query result cache is completely transparent to our applications. It is also maintained for consistency automatically, unlike our own caching programs.

References:

QUESTION 10

You want to create an optimally performing data warehouse hardware configuration for your customer. Which way of creating a hardware configuration will reduce the implementation time the most?

- A. Use reference configurations or an appliance-like configuration.
- B. Use the existing system and add on relevant components.
- C. Customize a configuration from a vendor.



D. Build the system from scratch.

Correct Answer: A

Explanation: Oracle Optimized Warehouse Reference Configurations are best practice guides to choosing the right server, storage and networking components to build an Oracle data warehouse. These best practice guides encapsulate years of configuration expertise from Oracle and its partners, helping customers take the risk out of implementing a data warehouse.

References:

QUESTION 11

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 12

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 13

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 14

What would you use to evenly distribute data across the disk in your Oracle data warehouse?

- A. Range Partitioning
- B. Automatic Storage Management (ASM)
- C. List Partitioning
- D. RAC

Correct Answer: B

Explanation: Automatic Storage Management (ASM) is a feature provided by Oracle Corporation within the Oracle Database from release Oracle 10g (revision 1) onwards. ASM aims to simplify the management of database files. To do so, it provides tools to manage file systems and volumes directly inside the database, allowing database administrators (DBAs) to control volumes and disks with familiar SQL statements in standard Oracle environments. Thus DBAs do not need extra skills in specific file systems or volume managers (which usually operate at the level of the operating system).

With ASM:

*



IO channels can take advantage of data striping and software mirroring

*

DBAs can automate online redistribution of data, along with the addition and removal of disks/storage

*

the system maintains redundant copies and provides 3rd-party[citation needed] RAID functionality

*

Oracle supports third-party multipathing IO technologies (such as failover or load balancing to SAN access) the need for hot spares diminishes

References:

QUESTION 15

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:



VCE & PDF

Pass4itSure.com

<https://www.pass4itsure.com/1z0-515.html>

2024 Latest pass4itsure 1Z0-515 PDF and VCE dumps Download

[Latest 1Z0-515 Dumps](#)

[1Z0-515 PDF Dumps](#)

[1Z0-515 Braindumps](#)