

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

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 across 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 2

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 3

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 4

One goal of your Information Lifecycle Management strategy using Oracle's ILM capabilities is to reduce cost or online

storage. Identify two database options that would help in enabling such a strategy.

- A. RAC and Advanced Compression
- B. RAC and Partitioning
- C. Partitioning and Advanced Compression
- D. RAC One and Advanced Compression

Correct Answer: B

Explanation: Advanced compression: Advanced Compression, an option introduced in Oracle Database 11 g Enterprise Edition, offers a comprehensive set of compression capabilities to help organizations reduce costs, while maintaining or improving performance. It significantly reduces the storage footprint of databases through compression of structured data (numbers, characters) as well as unstructured data (documents, spreadsheets, XML and other files). It provides enhanced compression for database backups and also includes network compression capabilities for faster synchronization of standby databases. Archival Compression:

*

Built on HCC technology

*

Compression algorithm optimized for maximum storage savings

*

Benefits any application with data retention requirements

*

Best approach for ILM and data archival

Partitioning: There are a number of benefits to partitioning data. Partitioning provides an easy way to distribute the data across appropriate storage devices depending on its usage, while still keeping the data online and stored on the most cost-effective device. Since partitioning is completely transparent to anyone accessing the data, no application changes are required, thus partitioning can be implemented at any time.

Note There is a wide variety of information held in an organization today, for example it could be an email message, a picture, or an order in an Online Transaction Processing System. Therefore, once the type of data being retained has been identified, you already have an understanding of what its evolution and final destiny is likely to be. The challenge now before all organizations, is to understand how their data evolves and grows, monitor how its usage changes over time, and decide how long it should survive. In addition, the evolving rules and regulations such as Sarbanes-Oxley, place additional constraints on the data that is being retained and some organizations now require that data is deleted when there is no longer a legal requirement to keep it, to avoid expensive e-discovery when the data is requested for a legal matter.

Implementing ILM using Oracle Database 11g Page 4 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 phase of the lifecycle of the data.1

References:

QUESTION 5

For which task would you NOT use Oracle Data Mining?

- A. Predicting customer behavior
- B. Associating factors with a business issue
- C. Determining associations within a population
- D. Reducing the amount of data used in a data warehouse

Correct Answer: D

Explanation:

Data mining does not reduce the amount of data in the warehouse.

Note:

Data mining (the analysis step of the knowledge discovery in databases process, or KDD), a relatively young and interdisciplinary field of computer science is the process of discovering new patterns from large data sets involving methods at the intersection of artificial intelligence, machine learning, statistics and database systems. The overall goal of the data mining process is to extract knowledge from a data set in a human-understandable structure and besides the raw analysis step involves database and data management aspects, data preprocessing, model and inference considerations, interestingness metrics, complexity considerations, post-processing of found structure, visualization and online updating.

QUESTION 6

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 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

The Analytic Workspace Manager would be used to generate_____.

- A. Materialized views
- B. Oracle OLAP Option cubes
- C. Oracle Data Mining algorithms
- D. Oracle SQL Analytic functions

Correct Answer: B

Explanation: You can use Analytic Workspace Manager for creating measures, dimensions and cubes in the OLAP database if the database was installed with the OLAP option. Workspace Manager, a feature of Oracle Database, enables application developers and DBAs to manage current, proposed and historical versions of data in the same database.

Note: Applications and DBA operations often work with more than one version of the data. Three common reasons to have multiple data versions are concurrency, auditing and scenario creation. Oracle Workspace Manager provides workspaces as a virtual environment to isolate a collection of changes to production data, keep a history of changes to data and create multiple data scenarios for "what if" analysis. It can save money, time and labor over traditional approaches.

QUESTION 9

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 10

Identify the true statement about adaptive parallelism.

- A. It is turned on by default.
- B. It is turned off by default.
- C. You should always leave the default setting
- D. There is no such thing.

Correct Answer: A

Explanation: Adaptive Parallelism: The adaptive multiuser algorithm, which is enabled by default, reduces the degree of

parallelism as the load on the system increases. When using the Oracle Database adaptive parallelism capabilities, the database uses an algorithm at SQL execution time to determine whether a parallel operation should receive the requested DOP or have its DOP lower to ensure the system is not overloaded. In a system that makes aggressive use of parallel execution by using a high DOP, the adaptive algorithm adjusts the DOP down when only a few operations are running in parallel. While the algorithm still ensures optimal resource utilization, users may experience inconsistent response times. Using solely the adaptive parallelism capabilities in an environment that requires deterministic response times is not advised. Adaptive parallelism is controlled through the database initialization parameter `PARALLEL_ADAPTIVE_MULTI_USER`.

References:

QUESTION 11

You want partitions to be automatically created when data that does not fit into current date range loaded. Which type of partitioning would you implement?

- A. Hash
- B. List
- C. Invisible
- D. Interval

Correct Answer: D

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.

Note: Partitioning is one of the most sought after options for data warehousing. Almost all Oracle data warehouses use partitioning to improve the performance of queries and also to ease the day-to-day maintenance complexities. Starting with 11G, more partitioning options have been provided and these should reduce the burden of the DBA to a great extent.

References:

QUESTION 12

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:

QUESTION 13

How can you implement near real time data integration with Oracle Data Integrator?

- A. By accessing Change Data Capture records from logs
- B. By using Exchange Partition
- C. By mining Oracle UNDO segments
- D. By reading operating system logs

Correct Answer: A

Explanation: Conventional "Extract, Transform, Load" (ETL) tools closely intermix data transformation rules with integration process procedures, requiring the development of both data transformations and data flow. Oracle Data Integrator (ODI) takes a different approach to integration by clearly separating the declarative rules (the "what") from the actual implementation (the "how"). With ODI, declarative rules describing mappings and transformations are defined graphically, through a drag-and-drop interface, and stored independently from the implementation. ODI automatically generates the data flow, which can be fine-tuned if required. This innovative approach for declarative design has also been applied to ODI's framework for Changed Data Capture. ODI's CDC moves only changed data to the target systems and can be integrated with Oracle GoldenGate, thereby enabling the kind of real time integration that businesses require.

References:

QUESTION 14

What would you do to compress data in partitions that are frequently updated in Oracle Database 11g?

- A. Use Hybrid Columnar Compression.
- B. Use Advanced Compression Option.
- C. Use Hybrid Partitions.
- D. Avoid compressing any data.

Correct Answer: B

Explanation:

Advanced Compression features in Oracle Database 11g include:

* Online Transaction Processing (OLTP) Table Compression: This breakthrough compression feature compresses table data during all types of data manipulation operations, including conventional INSERT or UPDATE. OLTP Table Compression leverages a sophisticated and intelligent algorithm that minimizes compression overhead during write operations, thereby making it viable for highly transactional workloads

References:

QUESTION 15

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:

[1Z0-515 PDF Dumps](#)

[1Z0-515 VCE Dumps](#)

[1Z0-515 Study Guide](#)