Oracle

Exam 1z0-515

Data Warehousing 11g Essentials

Verson: Demo

[Total Questions: 10]

Question No : 1

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

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 No : 2

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.

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.

Reference:

http://msdn.microsoft.com/en-us/library/aa480064.aspx

Question No : 3

You have analyzed your client's workload and the SQL Access Advisor in Enterprise Manager recommends that you create some materialized views to improve performance. What should you do in order to most simply implement this change?

A. Rewrite all the queries in the application to identify materialized view
B. Rewrite existing queries. New queries will automatically use the views.
C. Respond positively to the Advisor to create the materialized views.
D. Build virtual views on a third normal form schema.

Answer: C

Explanation: Enterprise Manager provides a very simple interface for the SQL Access Advisor (Advisor Central > SQL Advisor >SQL Access Advisor). The first page allows you to create tasks to test existing indexes, materialized view and partitions, or create tasks to suggest new structures.
The "Workload Source" page allows you to define the workload to associate with the task. The basic options allow the workload to be gathered from the cursor cache, an existing SQL tuning set, or a hypothetical workload based on specific schema objects.
The "Recommendation Options" page allows you to define which type of recommendations you are interested in (Indexes, Materialized Views and Partitioning).
After reviewing the result of the analysis you can decide if you should accept or ignore the suggested recommendations.

Note: The SQL Access Advisor was introduced in Oracle 10g to make suggestions about additional indexes and materialized views which might improve system performance.

Reference: QL Access Advisor in Oracle Database 11g Release 1

http://www.oracle-base.com/articles/11g/SQLAccessAdvisor_11gR1.php

Question No : 4

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

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.

Reference: Oracle Data Integrator, Oracle Data Sheet

Question No : 5

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!

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 No : 6

Identify the statement about ASM that is NOT true.

A. ASM is easier to manage than file systems.
B. ASM delivers the performance of raw partitions.
C. ASM is an extra cost option for Oracle databases.
D. ASM delivers automatic striping and mirroring.

Answer: B

Explanation: ASM is a management tool, not a RAW performance tool.

Note:
Automatic Storage Management (ASM) is a new feature that has be introduced in Oracle 10g to simplify the storage of Oracle datafiles, controlfiles and logfiles.

Automatic Storage Management (ASM) simplifies administration of Oracle related files by allowing the administrator to reference disk groups rather than individual disks and files, which are managed by ASM. The ASM functionality is an extention of the Oracle Managed Files (OMF) functionality that also includes striping and mirroring to provide balanced and secure storage. The new ASM functionality can be used in combination with existing raw and cooked file systems, along with OMF and manually managed files.

The ASM functionality is controlled by an ASM instance. This is not a full database instance, just the memory structures and as such is very small and lightweight.

The main components of ASM are disk groups, each of which comprise of several physical disks that are controlled as a single unit. The physical disks are known as ASM disks, while the files that reside on the disks are known as ASM files. The locations and names for the

files are controlled by ASM, but user-friendly aliases and directory structures can be defined for ease of reference.
The level of redundancy and the granularity of the striping can be controlled using templates. Default templates are provided for each file type stored by ASM, but additional templates can be defined as needed.
Failure groups are defined within a disk group to support the required level of redundancy. For two-way mirroring you would expect a disk group to contain two failure groups so individual files are written to two locations.
In summary ASM provides the following functionality:
*Manages groups of disks, called disk groups.
*Manages disk redundancy within a disk group.
*Provides near-optimal I/O balancing without any manual tuning.
*Enables management of database objects without specifying mount points and filenames.
*Supports large files.
Reference: Automatic Storage Management (ASM) in Oracle Database 10g

http://www.oracle-base.com/articles/10g/AutomaticStorageManagement10g.php

Question No : 7

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

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 No : 8

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.

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.

Reference:

http://www.dba-oracle.com/art_builder_bitmap_join_idx.htm

Question No : 9

Which statement is true for you to get the benefits of partition-wise joins?

A. The parent table must be partitioned on the join Key and the child table must be partitioned on] the join key.
B. The parent table must be partitioned on the primary key and the child table must be partition the join key.
C. The child table must use a reference partition.
D. The parent table must be partitioned on the primary key and the child table must use a ref partition.

Answer: A

Explanation:

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

Question No : 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

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.

Reference: Oracle Data Mining Concepts 11g Release 2

http://docs.oracle.com/cd/E11882_01/datamine.112/e16808/market_basket.htm