



Oracle

Exam 1z0-515

Data Warehousing 11g Essentials

Version: 8.3

[Total Questions: 70]

Question No : 1

Identify the true statement about REF partitions.

- A. REF partitions have no impact on partition-wise joins.
- B. Changes to partitioning in the parent table are automatically reflected in the child table.
- C. Changes in the data in a parent table are reflected in a child table.
- D. REF partitions can save storage space in the parent table.

Answer: B

Explanation: Reference partitioning is a partitioning method introduced in Oracle 11g.

Using reference partitioning, a child table can inherit the partitioning characteristics from a parent table.

Question No : 2

Identify the control structure that would NOT be defined as part of a data flow with Oracle Data Integrator.

- A. Loops
- B. Conditions
- C. Error handling
- D. GOTOs

Answer: D

Explanation: GOTOs cannot be used within the Oracle Data Integrator.

Reference: DIJQR.pdf, Page 7 (Oracle Data Integrator)

Question No : 3

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

Answer: C

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

Question No : 4

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

Answer: B

Explanation:

Substring functions are not analytic.

Question No : 5

How many Exadata Storage Server cells can be used in a grid?

- A. 7
- B. 14
- C. 128
- D. No practical limit

Answer: D

Explanation: There is no practical limit to number of cells that can be in the grid.

Reference: Sun Oracle Exadata and Database Machine Overview

Question No : 6

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.

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 No : 7

You will be implementing a data warehouse for one of your customers. In your design process, which index type is most likely to be used to improve the performance of some queries where the data is of low cardinality?

- A. Bitmap indexes
- B. B*-tree indexes
- C. Reverse indexes
- D. Invisible indexes

Answer: A

Explanation: Bitmap indexes are a highly compressed index type that tends to be used primarily for data warehouses.

Characteristic of Bitmap Indexes

- * For columns with very few unique values (low cardinality)
- * Columns that have low cardinality are good candidates (if the cardinality of a column is $\leq 0.1\%$ that the column is ideal candidate, consider also $0.2\% - 1\%$)
- * Tables that have no or little insert/update are good candidates (static data in warehouse)
- * Stream of bits: each bit relates to a column value in a single row of table

Reference: The Secrets of Oracle Bitmap Indexes,
http://www.akadia.com/services/ora_bitmapped_index.html

Question No : 8

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.

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.

Reference: query result cache in oracle 11g, <http://www.oracle-developer.net/display.php?id=503>

Question No : 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.

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.

Reference: Data Warehousing on Oracle RAC Best Practices

Question No : 10

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

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.

Reference: Partitioning in Oracle 11g, Oracle FAQs

Question No : 11

Identify the benefit of using interval partitioning.

- A. Automatic creation of new partitions based on hash values
- B. Automatic creation of new partitions based on the value of data being entered
- C. Improved performance compared to range partitions
- D. Automatic transfer of older partitions lower cost storage

Answer: B

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.

Reference: Partitioning in Oracle 11g, Oracle FAQs

Question No : 12

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

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.

Reference: Implementing Information Lifecycle Management Using the ILM Assistant

Question No : 13

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.

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.

Reference: Oracle® Database VLDB and Partitioning Guide, 11g Release 1 (11.1), 4 Partitioning for Availability, Manageability, and Performance

Question No : 14

You are looking for some general design principles that could be used in designing every large scale data warehouse you create. Identify the principle that would have the widest applicability.

- A. Partition your tables appropriately to produce partition-wise joins.
- B. Always use a star schema or snowflake schema design.
- C. Do as much analytics as possible in your BI tools.
- D. Always use Oracle OLAP.

Answer: A

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

A full partition-wise join divides a large join into smaller joins between a pair of partitions from the two joined tables. To use this feature, you must equipartition both tables on their join keys, or use reference partitioning.

Oracle Database can perform partial partition-wise joins only in parallel. Unlike full partition-wise joins, partial partition-wise joins require you to partition only one table on the join key, not both tables.

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.

Reference: Oracle® Database VLDB and Partitioning Guide, 11g Release 1 (11.1), 4 Partitioning for Availability, Manageability, and Performance

Question No : 15

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 : 16

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 : 17

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

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 No : 18

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 : 19

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

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.

Reference: Oracle® Database Administrator's Guide 11g Release 1 (11.1), Loading Tables

http://www.filibeto.org/sun/lib/nonsun/oracle/11.1.0.6.0/B28359_01/server.111/b28310/tables004.htm

Question No : 20

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.

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.

Reference: ORACLE OPTIMIZED WAREHOUSE REFERENCE CONFIGURATIONS
FREQUENTLY ASKED QUESTIONS

Question No : 21

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 : 22

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