

[2017-May-NewFull Version 1Z0-062 Exam Dump (PDF and VCE) 314Q for Free Download[1-10

2017 May New Oracle 1Z0-062 Exam Dumps with VCE and PDF Free Updated in www.Braindump2go.com Today!100% Real Exam Questions! 100% Exam Pass Guaranteed!1.[2017 Version New 1Z0-062 Exam Dumps (PDF & VCE) 314Q&As Download: <http://www.braindump2go.com/1z0-062.html> 2.[2017 Version New 1Z0-062 Exam Questions & Answers Download: <https://drive.google.com/drive/folders/0B75b5xYLjSSNdGhleTVzODBwSnM?usp=sharing> QUESTION 1Which three statements are true about a job chain? A. It can contain a nested chain of jobs.B. It can be used to implement dependency-based scheduling. C. It cannot invoke the same program or nested chain in multiple steps in the chain.D. It cannot have more than one dependency. E. It can be executed using event-based or time-based schedules. Answer: ABExplanation: http://docs.oracle.com/cd/B28359_01/server.111/b28310/scheduse009.htm#ADMIN12459 QUESTION 2The hr user receiver, the following error while inserting data into the sales table: On investigation, you find that the users tablespace uses Automnrf Segment Space Management (ASSM). It is the default tablespace for the HR user with an unlimited quota on it.Which two methods would you use to resolve this error? A. Altering the data life associated with the USERS tablespace to ex automaticallyB. Adding a data life to the USERS tablespaceC. Changing segment space management for the USERS tablespace to manualD. Creating a new tablespace with autoextend enabled and changing the default tablespace of the HR user to the new tablespaceE. Enabling resumable space allocation by setting the RESUMABLE_TIMEOUT parameter to a nonzero value Answer: AB QUESTION 3You execute the commands: Which two statements are true? A. The create user command fails if any role with the name Sidney exists in the database.B. The user sidney can connect to the database instance but cannot perform sort operations because no space quota is specified for the temp tablespace.C. The user sidney is created but cannot connect to the database instance because no profile is D. The user sidney can connect to the database instance but requires relevant privileges to create objects in the users tablespace.E. The user sidney is created and authenticated by the operating system. Answer: AD QUESTION 4Examine the query and its output: What might have caused three of the alerts to disappear? A. The threshold alerts were cleared and transferred to d0A_alert_history. B. An Automatic Workload Repository (AWR) snapshot was taken before the execution of the secondC. An Automatic Database Diagnostic Monitor (ADOM) report was generated before the execution of the second query.D. The database instance was restarted before the execution of the second query. Answer: AExplanation:DBA_OUTSTANDING_ALERTS: The DBA_OUTSTANDING_ALERTS means describes the outstanding alerts in the database and contains current database alerts. DBA_ALERT_HISTORY: The DBA_ALERT_HISTORY means lists a history of alerts that have been cleared and provides a history of nonthreshold alerts. QUESTION 5Which two statements are true? A. A role cannot be assigned external authentication. B. A role can be granted to other roles.C. A role can contain both system and object privileges.D. The predefined resource role includes the unlimited_tablespace privilege.E. All roles are owned by the sys user.F. The predefined connect role is always automatically granted to all new users at the time of their creation. Answer: BCEExplanation: http://docs.oracle.com/cd/E11882_01/network.112/e36292/authorization.htm#DBSEG99878 (the functionality of roles) QUESTION 6Identify three valid options for adding a pluggable database (PDB) to an existing multitenant container database (CDB). A. Use the CREATE PLUGGABLE DATABASE statement to create a PDB using the files from the SEED.B. Use the CREATE DATABASE . . . ENABLE PLUGGABLE DATABASE statement to provision a PDB by copying file from the SEED.C. Use the DBMS_PDB package to clone an existing PDB.D. Use the DBMS_PDB package to plug an Oracle 12c non-CDB database into an existing CDB.E. Use the DBMS_PDB package to plug an Oracle 11 g Release 2 (11.2.0.3.0) non-CDB database into an existing CDB. Answer: ACDEExplanation:Use the CREATE PLUGGABLE DATABASE statement to create a pluggable database (PDB). This statement enables you to perform the following tasks:* (A) Create a PDB by using the seed as a templateUse the create_pdb_from_seed clause to create a PDB by using the seed in the multitenant container database (CDB) as a template. The files associated with the seed are copied to a new location and the copied files are then associated with the new PDB.* (C) Create a PDB by cloning an existing PDBUse the create_pdb_clone clause to create a PDB by copying an existing PDB (the source PDB) and then plugging the copy into the CDB. The files associated with the source PDB are copied to a new location and the copied files are associated with the new PDB. This operation is called cloning a PDB.The source PDB can be plugged in or unplugged. If plugged in, then the source PDB can be in the same CDB or in a remote CDB. If the source PDB is in a remote CDB, then a database link is used to connect to the remote CDB and copy the files.* Create a PDB by plugging an unplugged PDB or a non-CDB into a CDBUse the create_pdb_from_xml clause to plug an unplugged PDB or a non-CDB into a CDB, using an XML metadata file. QUESTION 7 Your database supports a DSS workload that involves the execution of complex queries:Currently, the library cache contains the ideal workload for analysis. You want to analyze some of the queries for an application that are cached in the library cache.What

must you do to receive recommendations about the efficient use of indexes and materialized views to improve query performance?
A. Create a SQL Tuning Set (STS) that contains the queries cached in the library cache and run the SQL Tuning Advisor (STA) on the workload captured in the STS.
B. Run the Automatic Workload Repository Monitor (AWRM).
C. Create an STS that contains the queries cached in the library cache and run the SQL Performance Analyzer (SPA) on the workload captured in the STS.
D. Create an STS that contains the queries cached in the library cache and run the SQL Access Advisor on the workload captured in the STS.

Answer: D
Explanation: * SQL Access Advisor is primarily responsible for making schema modification recommendations, such as adding or dropping indexes and materialized views. SQL Tuning Advisor makes other types of recommendations, such as creating SQL profiles and restructuring SQL statements. * The query optimizer can also help you tune SQL statements. By using SQL Tuning Advisor and SQL Access Advisor, you can invoke the query optimizer in advisory mode to examine a SQL statement or set of statements and determine how to improve their efficiency. SQL Tuning Advisor and SQL Access Advisor can make various recommendations, such as creating SQL profiles, restructuring SQL statements, creating additional indexes or materialized views, and refreshing optimizer statistics.
Note: * Decision support system (DSS) workload * The library cache is a shared pool memory structure that stores executable SQL and PL/SQL code. This cache contains the shared SQL and PL/SQL areas and control structures such as locks and library cache handles.

Tuning SQL Statements QUESTION 8
In a recent Automatic Workload Repository (AWR) report for your database, you notice a high number of buffer busy waits. The database consists of locally managed tablespaces with free list managed segments. On further investigation, you find that buffer busy waits is caused by contention on data blocks. Which option would you consider first to decrease the wait event immediately?
A. Decreasing PCTUSED.
B. Decreasing PCTFREE.
C. Increasing the number of DBWn processes.
D. Using Automatic Segment Space Management (ASSM).
E. Increasing

db_buffer_cache based on the V\$DB_CACHE_ADVICE recommendation.
Answer: D
Explanation: * Automatic segment space management (ASSM) is a simpler and more efficient way of managing space within a segment. It completely eliminates any need to specify and tune the pctused, freelists, and freelist groups storage parameters for schema objects created in the tablespace. If any of these attributes are specified, they are ignored. * Oracle introduced Automatic Segment Storage Management (ASSM) as a replacement for traditional freelists management which used one-way linked-lists to manage free blocks with tables and indexes. ASSM is commonly called "bitmap freelists" because that is how Oracle implement the internal data structures for free block management.
Note: * Buffer busy waits are most commonly associated with segment header contention outside the data buffer pool (db_cache_size, etc.). * The most common remedies for high buffer busy waits include database writer (DBWR) contention tuning, adding freelists (or ASSM), and adding missing indexes.

QUESTION 9
The following parameter are set for your Oracle 12c database instance:
OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES=FALSE
OPTIMIZER_USE_SQL_PLAN_BASELINES=TRUE
You want to manage the SQL plan evolution task manually. Examine the following steps:
1. Set the evolve task parameters.
2. Create the evolve task by using the DBMS_SPM.CREATE_EVOLVE_TASK function.
3. Implement the recommendations in the task by using the DBMS_SPM.IMPLEMENT_EVOLVE_TASK function.
4. Execute the evolve task by using the DBMS_SPM.EXECUTE_EVOLVE_TASK function.
5. Report the task outcome by using the DBMS_SPM.REPORT_EVOLVE_TASK function.
Identify the correct sequence of steps:
A. 2, 4, 5
B. 2, 1, 4, 3, 5
C. 1, 2, 3, 4, 5
D. 1, 2, 4, 5
Answer: B

Explanation: * Evolving SQL Plan Baselines
2. Create the evolve task by using the DBMS_SPM.CREATE_EVOLVE_TASK function. This function creates an advisor task to prepare the plan evolution of one or more plans for a specified SQL statement. The input parameters can be a SQL handle, plan name or a list of plan names, time limit, task name, and description.
1. Set the evolve task parameters.
SET_EVOLVE_TASK_PARAMETER This function updates the value of an evolve task parameter. In this release, the only valid parameter is TIME_LIMIT.
4. Execute the evolve task by using the DBMS_SPM.EXECUTE_EVOLVE_TASK function. This function executes an evolution task. The input parameters can be the task name, execution name, and execution description. If not specified, the advisor generates the name, which is returned by the function.

3: IMPLEMENT_EVOLVE_TASK This function implements all recommendations for an evolve task. Essentially, this function is equivalent to using ACCEPT_SQL_PLAN_BASELINE for all recommended plans. Input parameters include task name, plan name, owner name, and execution name.
5. Report the task outcome by using the DBMS_SPM_EVOLVE_TASK function. This function displays the results of an evolve task as a CLOB. Input parameters include the task name and section of the report to include.
Oracle Database SQL Tuning Guide 12c, Managing SQL Plan Baselines

QUESTION 10
Examine this command:
SQL > exec DBMS_STATS.SET_TABLE_PREFS ('SH', 'CUSTOMERS', 'PUBLISH', 'false');
Which three statements are true about the effect of this command?
A. Statistics collection is not done for the CUSTOMERS table when schema stats are gathered.
B. Statistics collection is not done for the CUSTOMERS table when database stats are gathered.
C. Any existing statistics for the CUSTOMERS table are still available to the optimizer at parse time.
D. Statistics gathered on the CUSTOMERS table when schema stats are gathered are stored as pending statistics.
E. Statistics gathered on the CUSTOMERS table when database stats are

gathered are stored as pending statistics. Answer: CDEExplanation: * SET_TABLE_PREFS Procedure This procedure is used to set the statistics preferences of the specified table in the specified schema. * Example: Using Pending Statistics Assume many modifications have been made to the employees table since the last time statistics were gathered. To ensure that the cost-based optimizer is still picking the best plan, statistics should be gathered once again; however, the user is concerned that new statistics will cause the optimizer to choose bad plans when the current ones are acceptable. The user can do the following: EXEC DBMS_STATS.SET_TABLE_PREFS('hr', 'employees', 'PUBLISH', 'false'); By setting the employees tables publish preference to FALSE, any statistics gather from now on will not be automatically published. The newly gathered statistics will be marked as pending. !!!RECOMMEND!!! 1. |2017 Version New 1Z0-062 Exam Dumps (PDF & VCE) 314Q&As Download: <http://www.braindump2go.com/1z0-062.html> 2. |2017 Version New 1Z0-062 Study Guide Video: YouTube Video: [YouTube.com/watch?v=RBY_W-sEQKk](https://www.youtube.com/watch?v=RBY_W-sEQKk)