[2016-6-NEWOracle 1Z0-060 Book PDF Free Offered By Braindump2go for Downloading[NQ41-NQ50

 2016 June Oracle Official - 1Z0-060: Upgrade to Oracle Database 12c Exam Questions New Updated Today in Braindump2go.com. 100% 1Z0-060 Exam Pass Guaranteed!NEW QUESTION 41 - NEW QUESTION 50: QUESTION 41You are planning the creation of a new multitenant container database (CDB) and want to store the ROOT and SEED container data files in separate directories.You plan to create the database using SQL statements.Which three techniques can you use to achieve this? A.   Use Oracle Managed Files (OMF).B.   Specify the SEED FILE_NAME_CONVERT clause.C.   Specify the PDB_FILE_NAME_CONVERT initialization parameter.D.   Specify the DB_FILE_NAMECONVERT initialization parameter.E.   Specify all files in the CREATE DATABASE statement without using Oracle managed Files (OMF). Answer: ABCExplanation: You must specify the names and locations of the seed's files in one of the following ways:* (A) Oracle Managed Files* (B) The SEED FILE_NAME_CONVERT Clause* (C) The PDB_FILE_NAME_CONVERT Initialization Parameter QUESTION 42You are about to plug a multi-terabyte non-CDB into an existing multitenant container database (CDB).The characteristics of the non-CDB are as follows:- Version: Oracle Database 11gRelease 2 (11.2.0.2.0) 64-bit- Character set: AL32UTF8- National character set: AL16UTF16- O/S: Oracle Linux 6 64-bitThe characteristics of the CDB are as follows:- Version: Oracle Database 12c Release 1 64-bit- Character Set: AL32UTF8- National character set: AL16UTF16- O/S: Oracle Linux 6 64-bitWhich technique should you use to minimize down time while plugging this non-CDB into the CDB? A.   Transportable databaseB.   Transportable tablespaceC.   Data Pump full export/importD.   The DBMS_PDB packageE.   RMAN Answer: BExplanation:If the non-CDB is version 11.2.0.3 onward, you can consider using Transport Database, as described here. If the non-CDB is pre-11.2.0.3, then you can still consider using transportable tablespaces.http://www.oracle-base.com/articles/12c/multitenant-migrate-non-cdb-to-pdb-12cr1.php QUESTION 43Your database supports an online transaction processing (OLTP) application. The application is undergoing some major schema changes, such as addition of new indexes and materialized views. You want to check the impact of these changes on workload performance.What should you use to achieve this? A.   Database replayB.   SQL Tuning AdvisorC.   SQL Access AdvisorD.   SQL Performance AnalyzerE.   Automatic Workload Repository compare reports Answer: DExplanation:You can use the SQL Performance Analyzer to analyze the SQL performance impact of any type of system change. Examples of common system changes include:?Database upgrades?Configuration changes to the operating system, hardware, or database?Database initialization parameter changes?Schema changes, such as adding new indexes or materialized views?Gathering optimizer statistics?SQL tuning actions, such as creating SQL profileshttp://docs.oracle.com/cd/B28359_01/server.111/b28318/intro.htm#CNCPT961 QUESTION 44An administrator account is granted the CREATE SESSION and SET CONTAINER system privileges.A multitenant container database (CDB) instant has the following parameter set:THREADED_EXECUTION = FALSEWhich four statements are true about this administrator establishing connections to root in a CDB that has been opened in read only mode? A.   You can conned as a common user by using the connect statement.B.   You can connect as a local user by using the connect statement.C.   You can connect by using easy connect.D.   You can connect by using OS authentication.E.   You can connect by using a Net Service name.F.   You can connect as a local user by using the SET CONTAINER statement. Answer: ACDEExplanation:* The choice of threading model is dictated by the THREADED_EXECUTION initialization parameter.THREADED_EXECUTION=FALSE: The default value causes Oracle to run using the multiprocess model.THREADED_EXECUTION=TRUE: Oracle runs with the multithreaded model. *OS Authentication is not supported with the multithreaded model.* THREADED_EXECUTIONWhen this initialization parameter is set to TRUE, which enables the multithreaded Oracle model, operating system authentication is not supported. Attempts to connect to the database using operating system authentication (for example, CONNECT / AS SYSDBA or CONNECT /) when this initialization parameter is set to TRUE receive an ORA-01031"insufficient privileges" error. QUESTION 45Examine the following query output:You issue the following command to import tables into the hr schema:$ > impdp hr/hr directory = dumpdir dumpfile = hr_new.dmp schemas=hr TRANSFORM=DISABLE_ARCHIVE_LOGGING: YWhich statement is true? A.   All database operations performed by the impdp command are logged.B.   Only CREATE INDEX and CREATE TABLE statements generated by the import are logged.C.   Only CREATE TABLE and ALTER TABLE statements generated by the import are logged.D.   None of the operations against the master table used by Oracle Data Pump to coordinate its activities are logged. Answer: CExplanation:Oracle Data Pump disable redo logging when loading data into tables and when creating indexes.The new TRANSFORM option introduced in data pumps import provides the flexibility to turn off the redo generation for the objects during the course of import. The Master Table is used to track the detailed progress information of a Data Pump job.The Master Table is created in the schema of the current user running the Pump Dump export or import, and it keeps tracks of lots of detailed information. QUESTION 46You notice a performance change in your production Oracle database and you want to know which

change has made this performance difference.You generate the Compare Period Automatic Database Diagnostic Monitor (ADDM) report to further investigation.Which three findings would you get from the report? A.   It detects any configuration change that caused a performance difference in both time periods.B.   It identifies any workload change that caused a performance difference in both time periods.C.   It detects the top wait events causing performance degradation.D.   It shows the resource usage for CPU, memory, and I/O in both time periods.E.   It shows the difference in the size of memory pools in both time periods.F.   It gives information about statistics collection in both time periods. Answer: ABDExplanation:Keyword: shows the difference.*Full ADDM analysis across two AWR snapshot periodsDetects causes, measure effects, then correlates themCauses: workload changes, configuration changesEffects: regressed SQL, reach resource limits (CPU, I/O, memory, interconnect) Makes actionable recommendations along with quantified impact *Identify what changed/Configuration changes, workload changes*Performance degradation of the database occurs when your database was performing optimally in the past, such as 6 months ago, but has gradually degraded to a point where it becomes noticeable to the users. The Automatic Workload Repository (AWR) Compare Periods report enables you to compare database performance between two periods of time.While an AWR report shows AWR data between two snapshots (or two points in time), the AWR Compare Periods report shows the difference (ABE) between two periods (or two AWR reports with a total of four snapshots). Using the AWR Compare Periods report helps you to identify detailed performance attributes and configuration settings that differ between two time periods. QUESTION 47Examine the parameter for your database instance:You generated the execution plan for the following query in the plan table and noticed that the nested loop join was done. After actual execution of the query, you notice that the hash join was done in the execution plan:Identify the reason why the optimizer chose different execution plans. A.   The optimizer used a dynamic plan for the query.B.   The optimizer chose different plans because automatic dynamic sampling was enabled.C.   The optimizer used re-optimization cardinality feedback for the query.D.   The optimizer chose different plan because extended statistics were created for the columns used. Answer: C Explanation:https://blogs.oracle.com/optimizer/entry/cardinality_feedback QUESTION 48Which three statements are true about adaptive SQL plan management? A.   It automatically performs verification or evolves non-accepted plans, in COMPREHENSIVE mode when they perform better than existing accepted plans.B.   The optimizer always uses the fixed plan, if the fixed plan exists in the plan baseline.C.   It adds new, bettor plans automatically as fixed plans to the baseline.D.   The non-accepted plans are automatically accepted and become usable by the optimizer if they perform better than the existing accepted plans.E.   The non-accepted plans in a SQL plan baseline are automatically evolved, in COMPREHENSIVE mode, during the nightly maintenance window and a persistent verification report is generated. Answer: ADEExplanation:With adaptive SQL plan management, DBAs no longer have to manually run the verification or evolve process for non-accepted plans. When automatic SQL tuning is in COMPREHENSIVE mode, it runs a verification or evolve process for all SQL statements that have non-accepted plans during the nightly maintenance window. If the non-accepted plan performs better than the existing accepted plan (or plans) in the SQL plan baseline, then the plan is automatically accepted and becomes usable by the optimizer. After the verification is complete, a persistent report is generated detailing how the non-accepted plan performs compared to the accepted plan performance. Because the evolve process is now an AUTOTASK, DBAs can also schedule their own evolve job at end time.Note:*The optimizer is able to adapt plans on the fly by predetermining multiple subplans for portions of theplan.*Adaptive plans, introduced in Oracle Database 12c, enable the optimizer to defer the final plan decisionfor a statement until execution time. The optimizer instruments its chosen plan (the default plan) withstatistics collectors so that it can detect at runtime, if its cardinality estimates differ greatly from theactual number of rows seen by the operations in the plan. If there is a significant difference, then theplan or a portion of it will be automatically adapted to avoid suboptimal performance on the firstexecution of a SQL statement. QUESTION 49You create a new pluggable database, HR_PDB, from the seed database. Which three tablespaces are created by default in HR_PDB? A.   SYSTEM B.   SYSAUXC.   EXAMPLED.   UNDOE.   TEMPF.   USERS Answer: ABEExplanation:*A PDB would have its SYSTEM, SYSAUX, TEMP tablespaces.It can also contains other user created tablespaces in it.*Oracle Database creates both the SYSTEM and SYSAUX tablespaces as part of every database.*tablespace_datafile_clausesUse these clauses to specify attributes for all data files comprising the SYSTEM and SYSAUX tablespaces in the seed PDB.Incorrect:Not D:a PDB can not have an undo tablespace. Instead, it uses the undo tablespace belonging to the CDB.Note:* Example:CONN pdb_admin@pdb1SELECT tablespace_name FROM dba_tablespaces;TABLESPACE_NAME------------------------------SYSTEMSYSAUXTEMPUSERSSQL> QUESTION 50 Which two statements are true about variable extent size support for large ASM files? A.   The metadata used to track extents in SGA is reduced.B.   Rebalance operations are completed faster than with a fixed extent sizeC.   An ASM Instance automatically allocates an appropriate extent size.D.   Resync operations are completed faster when a disk comes online after being taken offline. E.   Performance improves in a stretch cluster configuration by reading from a local copy of an extent. Answer: ACExplanation: A:Variable size extents enable support for larger ASM datafiles, reduce SGA memory requirements for very large databases(A), and

improve performance for file create and open operations.C:You don't have to worry about the sizes; the ASM instance automatically allocates the appropriate extent size.Note:*The contents of ASM files are stored in a disk group as a set, or collection, of data extents that are stored on individual disks within disk groups. Each extent resides on an individual disk. Extents consist of one or more allocation units (AU). To accommodate increasingly larger files, ASM uses variable size extents.*The size of the extent map that defines a file can be smaller by a factor of 8 and 64 depending on the file size. The initial extent size is equal to the allocation unit size and it increases by a factor of 8 and 64 at predefined thresholds. This feature is automatic for newly created and resized datafiles when the disk group compatibility attributes are set to Oracle Release 11 or higher. 2016 Valid Braindump2go Oracle 1Z0-060 Study Materials: 1.| Braindump2go Latest 1Z0-060 PDF and 1Z0-060 VCE Dumps with New 161q 1Z0-060 Exam Questions: http://www.braindump2go.com/1z0-060.html [100% Exam Pass Guaranteed!]   **2.| Braindump2go New 1Z0-060 Exam Questions and Answers - Google Drive: https://drive.google.com/folderview?id=0B75b5xYLjSSNWm1lN2tlY1ZKVVk&usp=sharing** MORE Practice is the Most Important IF You want to PASS Oracle 1Z0-060 Exam 100%!????? Braindump2go.com????? Pass All IT Exams at the first Try!