

[2017-May-NewFree Braindump2go 1Z0-062 Exam Dumps 314Q[21-30

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 21To implement Automatic Management (AMM), you set the following parameters: When you try to start the database instance with these parameter settings, you receive the following error message:SQL > startupORA-00824: cannot set SGA_TARGET or MEMORY_TARGET due to existing internal settings, see alert log for more information.Identify the reason the instance failed to start. A. The PGA_AGGREGATE_TARGET parameter is set to zero.B. The STATISTICS_LEVEL parameter is set to BASIC.C. Both the SGA_TARGET and MEMORY_TARGET parameters are set.D. The SGA_MAX_SIZE and SGA_TARGET parameter values are not equal. Answer: BExplanation:Example:SQL> startup forceORA-00824: cannot set SGA_TARGET or MEMORY_TARGET due to existing internal settings ORA-00848: STATISTICS_LEVEL cannot be set to BASIC with SGA_TARGET or MEMORY_TARGET QUESTION 22What are two benefits of installing Grid Infrastructure software for a stand-alone server before installing and creating an Oracle database? A. Effectively implements role separationB. Enables you to take advantage of Oracle Managed Files.C. Automatically registers the database with Oracle Restart.D. Helps you to easily upgrade the database from a prior release.E. Enables the Installation of Grid Infrastructure files on block or raw devices. Answer: ACExplanation:C: To use Oracle ASM or Oracle Restart, you must first install Oracle Grid Infrastructurefor a standalone server before you install and create the database. Otherwise, you must manually register the database with Oracle Restart.Desupport of Block and Raw DevicesWith the release of Oracle Database 11g release 2 (11.2) and Oracle RAC 11g release 2 (11.2), using Database Configuration Assistant or the installer to store Oracle Clusterware or Oracle Database files directly on block or raw devices is not supported. If you intend to upgrade an existing Oracle RAC database, or an Oracle RAC database with Oracle ASM instances, then you can use an existing raw or block device partition, and perform a rolling upgrade of your existing installation. Performing a new installation using block or raw devices is not allowed. QUESTION 23Identify two correct statements about multitenant architectures. A. Multitenant architecture can be deployed only in a Real Application Clusters (RAC) configuration.B. Multiple pluggable databases (PDBs) share certain multitenant container database (CDB) resources.C. Multiple CDBs share certain PDB resources.D. Multiple non-RAC CDB instances can mount the same PDB as long as they are on the same server.E. Patches are always applied at the CDB level.F. A PDB can have a private undo tablespace. Answer: BEExplanation:B: Using 12c Resource manager you will be able control CPU, Exadata I/O, sessions and parallel servers. A new 12c CDB Resource Manager Plan will use so-called "Shares" (resource allocations) to specify how CPU is distributed between PDBs. A CDB Resource Manager Plan also can use "utilization limits" to limit the CPU usage for a PDB. With a default directive, you do not need to modify the resource plan for each PDB plug and unplug.E: New paradigms for rapid patching and upgrades.The investment of time and effort to patch one multitenant container database results in patching all of its many pluggable databases. To patch a single pluggable database, you simply unplug/plug to a multitenant container database at a different Oracle Database software version.Incorrect:Not A:* The Oracle RAC documentation describes special considerations for a CDB in an Oracle RAC environment.* Oracle Multitenant is a new option for Oracle Database 12c Enterprise Edition that helps customers reduce IT costs by simplifying consolidation, provisioning, upgrades, and more. It is supported by a new architecture that allows a container database to hold many pluggable databases. And it fully complements other options, including Oracle Real Application Clusters and Oracle Active Data Guard. An existing database can be simply adopted, with no change, as a pluggable database; and no changes are needed in the other tiers of the application.Not D: You can unplug a PDB from one CDB and plug it into a different CDB without altering your schemas or applications. A PDB can be plugged into only one CDB at a time.not F:* UNDO tablespace can NOT be local and stays on the CDB level.* Redo and undo go hand in hand, and so the CDB as a whole has a single undo tablespace per RAC instance. QUESTION 24You upgrade your Oracle database in a multiprocessor environment. As a recommended you execute the following script:SQL > @utlrp.sqlWhich two actions does the script perform? A. Parallel compilation of only the stored PL/SQL codeB. Sequential recompilation of only the stored PL/SQL codeC. Parallel recompilation of any stored PL/SQL codeD. Sequential recompilation of any stored PL/SQL codeE. Parallel recompilation of Java codeF. Sequential recompilation of Java code Answer: CEExplanation:utlrp.sql and utlprp.sqlThe utlrp.sql and utlprp.sql scripts are provided by Oracle to recompile all invalid objects in the database. They are typically run after major database changes such as upgrades or patches. They are located in the \$ORACLE_HOME/rdbms/admin directory and provide a wrapper on the UTL_RECOMP package. The utlrp.sql script simply calls the utlprp.sql script with a command line parameter of "0". The utlprp.sql accepts a single integer parameter that indicates the level of parallelism as follows.0 - The level of parallelism is

derived based on the CPU_COUNT parameter. 1 - The recompilation is run serially, one object at a time. N - The recompilation is run in parallel with "N" number of threads. Both scripts must be run as the SYS user, or another user with SYSDBA, to work correctly. QUESTION 25 Which statement is true concerning dropping a pluggable database (PDB)? A. The PDB must be open in read-only mode. B. The PDB must be in mount state. C. The PDB must be unplugged. D. The PDB data files are always removed from disk. E. A dropped PDB can never be plugged back into a multitenant container database (CDB). Answer: B Explanation: When you unplug a PDB from a CDB, the unplugged PDB is in mounted mode. A prerequisite of dropping PDB is :The PDB must be in mounted mode, or it must be unplugged. QUESTION 26 You notice a high number of waits for the db file scattered read and db file sequential read events in the recent Automatic Database Diagnostic Monitor (ADDM) report. After further investigation, you find that queries are performing too many full table scans and indexes are not being used even though the filter columns are indexed. Identify three possible reasons for this. A. Missing or stale histogram statistics B. Undersized shared pool C. High clustering factor for the indexes D. High value for the DB_FILE_MULTIBLOCK_READ_COUNT parameter E. Oversized buffer cache Answer: ACDE Explanation: D: DB_FILE_MULTIBLOCK_READ_COUNT is one of the parameters you can use to minimize I/O during table scans. It specifies the maximum number of blocks read in one I/O operation during a sequential scan. The total number of I/Os needed to perform a full table scan depends on such factors as the size of the table, the multiblock read count, and whether parallel execution is being utilized for the operation. QUESTION 27 Which three features work together, to allow a SQL statement to have different cursors for the same statement based on different selectivity ranges? A. Bind Variable Peeking B. SQL Plan Baselines C. Adaptive Cursor Sharing D. Bind variable used in a SQL statement E. Literals in a SQL statement Answer: ACE Explanation: * In bind variable peeking (also known as bind peeking), the optimizer looks at the value in a bind variable when the database performs a hard parse of a statement. When a query uses literals, the optimizer can use the literal values to find the best plan. However, when a query uses bind variables, the optimizer must select the best plan without the presence of literals in the SQL text. This task can be extremely difficult. By peeking at bind values the optimizer can determine the selectivity of a WHERE clause condition as if literals had been used, thereby improving the plan. C: Oracle 11g/12g uses Adaptive Cursor Sharing to solve this problem by allowing the server to compare the effectiveness of execution plans between executions with different bind variable values. If it notices suboptimal plans, it allows certain bind variable values, or ranges of values, to use alternate execution plans for the same statement. This functionality requires no additional configuration. QUESTION 28 You notice a performance change in your production Oracle 12c database. You want to know which change caused this performance difference. Which method or feature should you use? A. Compare Period ADDM report B. AWR Compare Period report C. Active Session History (ASH) report D. Taking a new snapshot and comparing it with a preserved snapshot Answer: B Explanation: The awrrdrpt.sql report is the Automated Workload Repository Compare Period Report. The awrrdrpt.sql script is located in the \$ORACLE_HOME/rdbms/admin directory. Incorrect: Not A: Compare Period ADDM Use this report to perform a high-level comparison of one workload replay to its capture or to another replay of the same capture. Only workload replays that contain at least 5 minutes of database time can be compared using this report. QUESTION 29 You want to capture column group usage and gather extended statistics for better cardinality estimates for the CUSTOMERS table in the SH schema. Examine the following steps: 1. Issue the SELECT DBMS_STATS.CREATE_EXTENDED_STATS ('SH', 'CUSTOMERS') FROM dual statement. 2. Execute the DBMS_STATS.SEED_COL_USAGE (null, 'SH', 500) procedure. 3. Execute the required queries on the CUSTOMERS table. 4. Issue the SELECT DBMS_STATS.REPORT_COL_USAGE ('SH', 'CUSTOMERS') FROM dual statement. Identify the correct sequence of steps. A. 3, 2, 1, 4 B. 2, 3, 4, 1 C. 4, 1, 3, 2 D. 3, 2, 4, 1 Answer: B Explanation: Step 1 (2). Seed column usage Oracle must observe a representative workload, in order to determine the appropriate column groups. Using the new procedure DBMS_STATS.SEED_COL_USAGE, you tell Oracle how long it should observe the workload. Step 2: (3) You don't need to execute all of the queries in your work during this window. You can simply run explain plan for some of your longer running queries to ensure column group information is recorded for these queries. Step 3. (1) Create the column groups At this point you can get Oracle to automatically create the column groups for each of the tables based on the usage information captured during the monitoring window. You simply have to call the DBMS_STATS.CREATE_EXTENDED_STATS function for each table. This function requires just two arguments, the schema name and the table name. From then on, statistics will be maintained for each column group whenever statistics are gathered on the table. Note: * DBMS_STATS.REPORT_COL_USAGE reports column usage information and records all the SQL operations the database has processed for a given object. * The Oracle SQL optimizer has always been ignorant of the implied relationships between data columns within the same table. While the optimizer has traditionally analyzed the distribution of values within a column, he does not collect value-based relationships between columns. * Creating extended statistics Here are the steps to create extended statistics for related table columns with dbms_stats.create_extended_stats: 1 - The first step is to create column histograms for the related columns. 2 - Next, we run dbms_stats.create_extended_stats to relate the

columns together. Unlike a traditional procedure that is invoked via an execute ("exec") statement, Oracle extended statistics are created via a select statement. QUESTION 30 Which three statements are true about Automatic Workload Repository (AWR)? A. All AWR tables belong to the SYSTEM schema. B. The AWR data is stored in memory and in the database. C. The snapshots collected by AWR are used by the self-tuning components in the database. D. AWR computes time model statistics based on time usage for activities, which are displayed in the v\$SYS time model and V\$SESS_TIME_MODEL views. E. AWR contains system wide tracing and logging information. Answer: BCD !!!RECOMMEND!!! 1. | 2017 Version New 1Z0-062 Exam Dumps (PDF & VCE) 314 Q&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)