

[2017-May-NewExam 1Z0-062 VCE Dumps 314q from Braindump2go[41-50]

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 41Which 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, better 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. 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 the plan.* Adaptive plans, introduced in Oracle Database 12c, enable the optimizer to defer the final plan decision for a statement until execution time. The optimizer instruments its chosen plan (the default plan) with statistics collectors so that it can detect at runtime, if its cardinality estimates differ greatly from the actual number of rows seen by the operations in the plan. If there is a significant difference, then the plan or a portion of it will be automatically adapted to avoid suboptimal performance on the first execution of a SQL statement. QUESTION 42You create a new pluggable database, HR_PDB, from the seed database. Which three tablespaces are created by default in HR_PDB? A. SYSTEMB. 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 43Which 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: ACEExplanation: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. QUESTION 44You executed a DROP USER CASCADE on an Oracle 11g release 1 database and immediately realized that you forgot to copy the OCA.EXAM_RESULTS table to the OCP schema.The RECYCLE_BIN enabled before the DROP USER was executed and the OCP user has been granted the FLASHBACK ANY TABLE system privilege.What is the quickest way to recover the contents of the OCA.EXAM_RESULTS table to the OCP schema? A. Execute FLASHBACK TABLE OCA.EXAM_RESULTS TO BEFORE DROP RENAME TO OCP.EXAM_RESULTS; connected as SYSTEM.B. Recover the table using traditional Tablespace Point In Time Recovery.C. Recover the table using Automated Tablespace Point In Time Recovery.D. Recovery the table using Database Point In Time Recovery.E. Execute FLASHBACK TABLE OCA.EXAM_RESULTS TO BEFORE DROP RENAME TO EXAM_RESULTS; connected as the OCP user. Answer: CEExplanation:RMAN tablespace point-in-time recovery (TSPITR). Recovery Manager (RMAN) TSPITR enables quick recovery of one or more tablespaces in a database to an earlier time without affecting the rest of the tablespaces and objects in

the database. Fully Automated (the default) In this mode, RMAN manages the entire TSPITR process including the auxiliary instance. You specify the tablespaces of the recovery set, an auxiliary destination, the target time, and you allow RMAN to manage all other aspects of TSPITR. The default mode is recommended unless you specifically need more control over the location of recovery set files after TSPITR, auxiliary set files during TSPITR, channel settings and parameters or some other aspect of your auxiliary instance.

QUESTION 45 In your multitenant container database (CDB) containing pluggable database (PDBs), the HR user executes the following commands to create and grant privileges on a procedure: `CREATE OR REPLACE PROCEDURE create_test_v (v_emp_id NUMBER, v_ename VARCHAR2, v_salary NUMBER, v_dept_id NUMBER) BEGIN INSERT INTO hr.test VALUES (V_emp_id, V_ename, V_salary, V_dept_id); END;` / `GRANT EXECUTE ON CREATE_TEST TO john, jim, smith, king;` How can you prevent users having the EXECUTE privilege on the CREATE_TEST procedure from inserting values into tables on which they do not have any privileges? A. Create the CREATE_TEST procedure with definer's rights. B. Grant the EXECUTE privilege to users with GRANT OPTION on the CREATE_TEST procedure. C. Create the CREATE_TEST procedure with invoker's rights. D. Create the CREATE_TEST procedure as part of a package and grant users the EXECUTE privilege the package. Answer: C Explanation: If a program unit does not need to be executed with the escalated privileges of the definer, you should specify that the program unit executes with the privileges of the caller, also known as the invoker. Invoker's rights can mitigate the risk of SQL injection. Incorrect: Not A: By default, stored procedures and SQL methods execute with the privileges of their owner, not their current user. Such definer-rights subprograms are bound to the schema in which they reside. not B: Using the GRANT option, a user can grant an Object privilege to another user or to PUBLIC. **QUESTION 46** You created a new database using the "create database" statement without specifying the "ENABLE PLUGGABLE" clause. What are two effects of not using the "ENABLE PLUGGABLE database" clause? A. The database is created as a non-CDB and can never contain a PDB. B. The database is treated as a PDB and must be plugged into an existing multitenant container database (CDB). C. The database is created as a non-CDB and can never be plugged into a CDB. D. The database is created as a non-CDB but can be plugged into an existing CDB. E. The database is created as a non-CDB but will become a CDB whenever the first PDB is plugged in. Answer: AD Explanation: A (not B, not E): The `CREATE DATABASE ... ENABLE PLUGGABLE DATABASE SQL` statement creates a new CDB. If you do not specify the `ENABLE PLUGGABLE DATABASE` clause, then the newly created database is a non-CDB and can never contain PDBs. D: You can create a PDB by plugging in a Non-CDB as a PDB. The following graphic depicts the options for creating a PDB: Incorrect: Not E: For the duration of its existence, a database is either a CDB or a non-CDB. You cannot transform a non-CDB into a CDB or vice versa. You must define a database as a CDB at creation, and then create PDBs within this CDB. **QUESTION 47** You have installed two 64G flash devices to support the Database Smart Flash Cache feature on your database server that is running on Oracle Linux. You have set the `DB_SMART_FLASH_FILE` parameter: `DB_FLASH_CACHE_FILE = /dev/flash_device_1 `,' /dev/flash_device_2` How should the `DB_FLASH_CACHE_SIZE` be configured to use both devices? A. Set `DB_FLASH_CACHE_SIZE = 64G`. B. Set `DB_FLASH_CACHE_SIZE = 64G, 64G`. C. Set `DB_FLASH_CACHE_SIZE = 128G`. D. `DB_FLASH_CACHE_SIZE` is automatically configured by the instance at startup. Answer: B Explanation: * Smart Flash Cache concept is not new in Oracle 12C - DB Smart Flash Cache in Oracle 11g. In this release Oracle has made changes related to both initialization parameters used by DB Smart Flash cache. Now you can define many files|devices and its sizes for "Database Smart Flash Cache" area. In previous releases only one file|device could be defined. `DB_FLASH_CACHE_FILE = /dev/sda, /dev/sdb, /dev/sdc` `DB_FLASH_CACHE_SIZE = 32G, 32G, 64G` So above settings defines 3 devices which will be in use by "DB Smart Flash Cache" `/dev/sda size 32G /dev/sdb size 32G /dev/sdc size 64G` New view `V$FLASHFILESTAT` it's used to determine the cumulative latency and read counts of each file|device and compute the average latency **QUESTION 48** What is the effect of specifying the "ENABLE PLUGGABLE DATABASE" clause in a "CREATE DATABASE" statement? A. It will create a multitenant container database (CDB) with only the root opened. B. It will create a CDB with root opened and seed read only. C. It will create a CDB with root and seed opened and one PDB mounted. D. It will create a CDB that must be plugged into an existing CDB. E. It will create a CDB with root opened and seed mounted. Answer: B Explanation: * The `CREATE DATABASE ... ENABLE PLUGGABLE DATABASE SQL` statement creates a new CDB. If you do not specify the `ENABLE PLUGGABLE DATABASE` clause, then the newly created database is a non-CDB and can never contain PDBs. Along with the root (CDB\$ROOT), Oracle Database automatically creates a seed PDB (PDB\$SEED). The following graphic shows a newly created CDB: * Creating a PDB Rather than constructing the data dictionary tables that define an empty PDB from scratch, and then populating its Obj\$ and Dependency\$ tables, the empty PDB is created when the CDB is created. (Here, we use empty to mean containing no customer-created artifacts.) It is referred to as the seed PDB and has the name PDB\$Seed. Every CDB non-negotiably contains a seed PDB; it is non-negotiably always open in read-only mode. This has no conceptual significance; rather, it is just an optimization device. The create PDB operation is implemented as a special case of the clone PDB operation. **QUESTION 49** Examine the

following parameters for a database instance: MEMORY_MAX_TARGET=0 MEMORY_TARGET=0 SGA_TARGET=0
PGA_AGGREGATE_TARGET=500m Which three initialization parameters are not controlled by Automatic Shared Memory
Management (ASMM)? A. LOG_BUFFER B. SORT_AREA_SIZE C. JAVA_POOL_SIZE D. STREAMS_POOL_SIZE E.
DB_16K_CACHE_SIZE F. DB_KEEP_CACHE_SIZE Answer: AEF Explanation: Manually Sized SGA Components that Use
SGA_TARGET Space SGA Component, Initialization Parameter/ The log buffer LOG_BUFFER/ The keep and recycle buffer
caches DB_KEEP_CACHE_SIZE DB_RECYCLE_CACHE_SIZE/ Nonstandard block size buffer caches DB_nK_CACHE_SIZE
Note: * In addition to setting SGA_TARGET to a nonzero value, you must set to zero all initialization parameters listed in the table
below to enable full automatic tuning of the automatically sized SGA components. * Table, Automatically Sized SGA Components
and Corresponding Parameters QUESTION 50 Which background process does Automatic Shared Memory Management use to
coordinate the sizing of memory components? A. PMON B. SMON C. MMNL D. MMAN E. MMON Answer: D
!!!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)