Oracle

Exam 1z0-060

Upgrade to Oracle Database 12c

Version: 8.6

[Total Questions: 150]

Question No : 1

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_SZIE
F. DB_KEEP_CACHE_SIZE

Answer: A,E,F

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

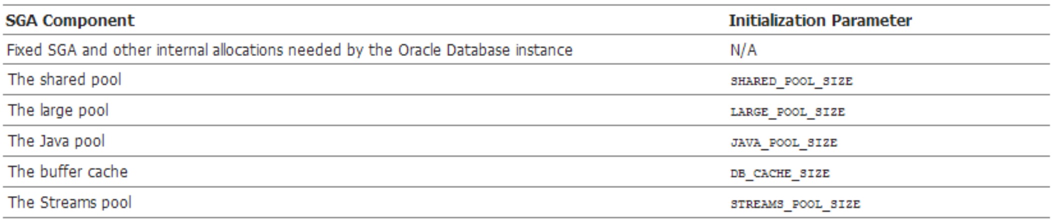

SGA Component	Initialization Parameter
Fixed SGA and other internal allocations needed by the Oracle Database instance	N/A
The shared pool	SHARED_POOL_SIZE
The large pool	LARGE_POOL_SIZE
The Java pool	JAVA_POOL_SIZE
The buffer cache	DB_CACHE_SIZE
The Streams pool	STREAMS_POOL_SIZE

Question No : 2

Which three are direct benefits of the multiprocess, multithreaded architecture of Oracle Database 12c when it is enabled?

A. Reduced logical I/O
B. Reduced virtual memory utilization
C. Improved parallel Execution performance
D. Improved Serial Execution performance
E. Reduced physical I/O
F. Reduced CPU utilization

Answer: B,C,F

Explanation: * Multiprocess and Multithreaded Oracle Database Systems

Multiprocess Oracle Database (also called multiuser Oracle Database) uses several processes to run different parts of the Oracle Database code and additional Oracle processes for the users—either one process for each connected user or one or more processes shared by multiple users. Most databases are multiuser because a primary advantage of a database is managing data needed by multiple users simultaneously.

Each process in a database instance performs a specific job. By dividing the work of the database and applications into several processes, multiple users and applications can connect to an instance simultaneously while the system gives good performance.

* In previous releases, Oracle processes did not run as threads on UNIX and Linux systems. Starting in Oracle Database 12c, the multithreaded Oracle Database model enables Oracle processes to execute as operating system threads in separate address spaces.

Question No : 3

A database is stored in an Automatic Storage Management (ASM) disk group, disk group, DGROUP1 with SQL:

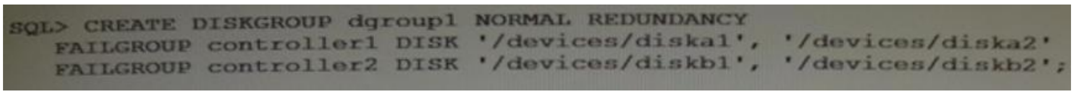

```
SQL> CREATE DISKGROUP dgroup1 NORMAL REDUNDANCY
   FAILGROUP controller1 DISK '/devices/diska1', '/devices/diska2'
   FAILGROUP controller2 DISK '/devices/diskb1', '/devices/diskb2';
```

There is enough free space in the disk group for mirroring to be done.

What happens if the CONTROLLER1 failure group becomes unavailable due to error of for maintenance?

A. Transactions and queries accessing database objects contained in any tablespace stored in DGROUP1 will fall.
B. Mirroring of allocation units will be done to ASM disks in the CONTROLLER2 failure group until the CONTROLLER1 for failure group is brought back online.
C. The data in the CONTROLLER1 failure group is copied to the controller2 failure group and rebalancing is initiated.
D. ASM does not mirror any data until the controller failure group is brought back online, and newly allocated primary allocation units (AU) are stored in the controller2 failure group, without mirroring.
E. Transactions accessing database objects contained in any tablespace stored in DGROUP1 will fail but queries will succeed.

Answer: B
Explanation: CREATE DISKGROUP NORMAL REDUNDANCY

* For Oracle ASM to mirror files, specify the redundancy level as NORMAL REDUNDANCY (2-way mirroring by default for most file types) or HIGH REDUNDANCY (3-way mirroring for all files).

Question No : 4

You performed an incremental level 0 backup of a database:

RMAN > BACKUP INCREMENTAL LEVEL 0 DATABASE;

To enable block change tracking after the incremental level 0 backup, you issued this command:

SQL > ALTER DATABASE ENABLE BLOCK CHANGE TRACKING USING FILE

‘ /mydir/rman_change_track.f’;

To perform an incremental level 1 cumulative backup, you issued this command:

RMAN> BACKUP INCREMENTAL LEVEL 1 CUMULATIVE DATABASE;

Which three statements are true?

A. Backup change tracking will sometimes reduce I/O performed during cumulative incremental backups.
B. The change tracking file must always be backed up when you perform a full database backup.
C. Block change tracking will always reduce I/O performed during cumulative incremental backups.
D. More than one database block may be read by an incremental backup for a change made to a single block.
E. The incremental level 1 backup that immediately follows the enabling of block change tracking will not read the change tracking file to discover changed blocks.

Answer: A,D,E

Explanation: A: In a cumulative level 1 backup, RMAN backs up all the blocks used since the most recent level 0 incremental backup.
E: Oracle Block Change Tracking
Once enabled; this new 10g feature records the modified since last backup and stores the log of it in a block change tracking file using the CTW (Change Tracking Writer) process. During backups RMAN uses the log file to identify the specific blocks that must be backed up. This improves RMAN's performance as it does not have to scan whole datafiles to detect changed blocks. Logging of changed blocks is performed by the CTRW process which is also responsible for writing data to the block change tracking file.

Note:
* An incremental level 0 backup backs up all blocks that have ever been in use in this database.

Question No : 5

On your Oracle Database, you issue the following commands to create indexes:

SQL > CREATE INDEX oe.ord_customer_ix1 ON oe.orders (customer_id, sales_rep_id) INVISIBLE;

SQL> CREATE BITMAP INDEX oe.ord_customer_ix2 ON oe.orders (customer_id, sales_rep_id);

Which two statements are true?

A. Only the ORD_CUSTOMER_IX1 index created.
B. Both the indexes are updated when a row is inserted, updated, or deleted in the ORDERS table.
C. Both the indexes are created: however, only ORD_CUSTOMERS_IX1 is used by the optimizer for queries on the ORDERS table.
D. The ORD_CUSTOMER_IX1 index is not used by the optimizer even when the OPTIMIZER_USE_INVISIBLE_INDEXES parameters is set to true.
E. Both the indexes are created and used by the optimizer for queries on the ORDERS table.
F. Both the indexes are created: however, only ORD_CUSTOMERS_IX2 is used by the optimizer for queries on the ORDERS table.

Answer: B,F

Explanation: Not A: Both indexes are created fine.
B: The invisible index ORD_CUSTOMERS_IX1 and the bitmap index are both updated by DML operations on the Orders table.
F: Since ORD_CUSTOMERS_IX1 is invisible only ORD_CUSTOMERS_IX2 is used by the query optimizer.

Not C,Not D,Not E:
* ord_customer_ix1 is an invisible index and is therefore not used by the optimizer.
* VISIBLE | INVISIBLE Use this clause to specify whether the index is visible or invisible to the optimizer. An invisible index is maintained by DML operations, but it is not be used by the optimizer during queries unless you explicitly set the parameter OPTIMIZER_USE_INVISIBLE_INDEXES to TRUE at the session or system level.

Note: Specify BITMAP to indicate that index is to be created with a bitmap for each distinct key, rather than indexing each row separately. Bitmap indexes store the rowids associated with a key value as a bitmap. Each bit in the bitmap corresponds to a possible rowid. If the bit is set, then it means that the row with the corresponding rowid contains the key value. The internal representation of bitmaps is best suited for applications with low levels of concurrent transactions, such as data warehousing.

Question No : 6

Which three statements are true about the working of system privileges in a multitenant control database (CDB) that has pluggable databases (PDBs)?

A. System privileges apply only to the PDB in which they are used.
B. Local users cannot use local system privileges on the schema of a common user.
C. The granter of system privileges must possess the set container privilege.
D. Common users connected to a PDB can exercise privileges across other PDBs.
E. System privileges with the with grant option container all clause must be granted to a common user before the common user can grant privileges to other users.

Answer: A,C,E

Explanation: A, Not D: In a CDB, PUBLIC is a common role. In a PDB, privileges granted locally to PUBLIC enable all local and common users to exercise these privileges in this PDB only.

C: A user can only perform common operations on a common role, for example, granting privileges commonly to the role, when the following criteria are met:

The user is a common user whose current container is root.

The user has the SET CONTAINER privilege granted commonly, which means that the privilege applies in all containers.

The user has privilege controlling the ability to perform the specified operation, and this privilege has been granted commonly

Incorrect:

Note:
* Every privilege and role granted to Oracle-supplied users and roles is granted commonly except for system privileges granted to PUBLIC, which are granted locally.

Question No : 7

Your multitenant container database (CDB) contains three pluggable database (PDBs). You find that the control file is damaged. You plan to use RMAN to recover the control file. There are no startup triggers associated with the PDBs.

Which three steps should you perform to recover the control file and make the database fully operational?

A. Mount the container database (CDB) and restore the control file from the control file auto backup.
B. Recover and open the CDB in NORMAL mode.
C. Mount the CDB and then recover and open the database, with the RESETLOGS option.
D. Open all the pluggable databases.
E. Recover each pluggable database.
F. Start the database instance in the nomount stage and restore the control file from control file auto backup.

Answer: C,D,F

Explanation: Step 1: F

Step 2: D

Step 3: C: If all copies of the current control file are lost or damaged, then you must restore and mount a backup control file. You must then run the RECOVERcommand, even if no data files have been restored, and open the database with the RESETLOGS option.

Note:

* RMAN and Oracle Enterprise Manager Cloud Control (Cloud Control) provide full support for backup and recovery in a multitenant environment. You can back up and recover a whole multitenant container database (CDB), root only, or one or more pluggable databases (PDBs).

Question No : 8

What are two benefits of installing Grid Infrastructure software for a stand-alone server before installing and creating an Oracle database?

A. Effectively implements role separation
B. 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: C,E

Explanation: “Oracle Grid Infrastructure for a standalone server includes Oracle Restart and Oracle Automatic Storage Management. Oracle combined the two infrastructure products into a single set of binaries that is installed into an Oracle Restart home.”

http://docs.oracle.com/cd/E16655_01/install.121/e17735/oraclerestart.htm#NTDBI999

Question No : 9

You configure your database Instance to support shared server connections.

Which two memory areas that are part of PGA are stored in SGA instead, for shared server connection?

A. User session data
B. Stack space
C. Private SQL area
D. Location of the runtime area for DML and DDL Statements
E. Location of a part of the runtime area for SELECT statements

Answer: A,C

Explanation: A: PGA itself is subdivided. The UGA (User Global Area) contains session state information, including stuff like package-level variables, cursor state, etc. Note that, with shared server, the UGA is in the SGA. It has to be, because shared server means that the session state needs to be accessible to all server processes, as any one of them could be assigned a particular session. However, with dedicated server (which likely what you're using), the UGA is allocated in the PGA.

C: The Location of a private SQL area depends on the type of connection established for a session. If a session is connected through a dedicated server, private SQL areas are located in the server process’ PGA. However, if a session is connected through a shared server, part of the private SQL area is kept in the SGA.

Note:
* System global area (SGA)
The SGA is a group of shared memory structures, known as *SGA components*, that contain data and control information for one Oracle Database instance. The SGA is shared by all server and background processes. Examples of data stored in the SGA include cached data blocks and shared SQL areas.

* Program global area (PGA)

A PGA is a memory region that contains data and control information for a server process. It is nonshared memory created by Oracle Database when a server process is started. Access to the PGA is exclusive to the server process. There is one PGA for each server process. Background processes also allocate their own PGAs. The total memory used by all individual PGAs is known as the total instance PGA memory, and the collection of individual PGAs is referred to as the total instance PGA, or just instance PGA. You use database initialization parameters to set the size of the instance PGA, not individual PGAs.

Reference: Oracle Database Concepts 12c

Question No : 10

Examine the following command:

```
CREATE TABLE (prod_id number(4),

Prod_name varchar2 (20),

Category_id number(30),

Quantity_on_hand number (3) INVISIBLE);
```

Which three statements are true about using an invisible column in the PRODUCTS table?

A. The %ROWTYPE attribute declarations in PL/SQL to access a row will not display the invisible column in the output.
B. The DESCRIBE commands in SQL *Plus will not display the invisible column in the

output.
C. Referential integrity constraint cannot be set on the invisible column.
D. The invisible column cannot be made visible and can only be marked as unused.
E. A primary key constraint can be added on the invisible column.

Answer: A,B,E

Explanation: AB: You can make individual table columns invisible. Any generic access of a table does not show the invisible columns in the table. For example, the following operations do not display invisible columns in the output:

* SELECT * FROM statements in SQL
* DESCRIBE commands in SQL*Plus
* %ROWTYPE attribute declarations in PL/SQL
* Describes in Oracle Call Interface (OCI)

Incorrect:
Not D: You can make invisible columns visible.
You can make a column invisible during table creation or when you add a column to a table, and you can later alter the table to make the same column visible.

Reference: Understand Invisible Columns

Question No : 11

Identify two situations in which the alert log file is updated.

A. Running a query on a table returns ORA-600: Internal Error.
B. Inserting a value into a table returns ORA-01722: invalid number.
C. Creating a table returns ORA-00955: name us already in used by an existing objects.
D. Inserting a value into a table returns ORA-00001: unique constraint (SYS.OK_TECHP) violated.
E. Rebuilding an index using ALTER INDEX . . . REBUILD fails with an ORA-01578: ORACLE data block corrupted (file # 14, block # 50) error.

Answer: A,E

Explanation: The alert log is a chronological log of messages and errors, and includes the following items:

*All internal errors (ORA-600), block corruption errors (ORA-1578), and deadlock errors

(ORA-60) that occur

* Administrative operations, such as CREATE, ALTER, and DROP statements and STARTUP, SHUTDOWN, and ARCHIVELOG statements

* Messages and errors relating to the functions of shared server and dispatcher processes

* Errors occurring during the automatic refresh of a materialized view

* The values of all initialization parameters that had nondefault values at the time the database and instance start

Note:
* The alert log file (also referred to as the ALERT.LOG) is a chronological log of messages and errors written out by an Oracle Database. Typical messages found in this file is: database startup, shutdown, log switches, space errors, etc. This file should constantly be monitored to detect unexpected messages and corruptions.

Question No : 12

You wish to enable an audit policy for all database users, except SYS, SYSTEM, and SCOTT.

You issue the following statements:

```
SQL> AUDIT POLICY ORA_DATABASE_PARAMETER EXCEPT SYS;

SQL> AUDIT POLICY ORA_DATABASE_PARAMETER EXCEPT SYSTEM;

SQL> AUDIT POLICY ORA_DATABASE_PARAMETER EXCEPT SCOTT;
```

For which database users is the audit policy now active?

A. All users except SYS
B. All users except SCOTT
C. All users except sys and SCOTT
D. All users except sys, system, and SCOTT

Answer: B

Explanation: If you run multiple AUDIT statements on the same unified audit policy but specify different EXCEPT users, then Oracle Database uses the last exception user list, not any of the users from the preceding lists. This means the effect of the earlier AUDIT POLICY ... EXCEPT statements are overridden by the latest AUDIT POLICY ... EXCEPT statement.

Note:
* The ORA_DATABASE_PARAMETER policy audits commonly used Oracle Database parameter settings. By default, this policy is not enabled.
* You can use the keyword ALL to audit all actions. The following example shows how to audit all actions on the HR.EMPLOYEES table, except actions by user pmulligan.

Example Auditing All Actions on a Table

```
CREATE AUDIT POLICY all_actions_on_hr_emp_pol
ACTIONS ALL ON HR.EMPLOYEES;

AUDIT POLICY all_actions_on_hr_emp_pol EXCEPT pmulligan;
```

Reference: Oracle Database Security Guide 12c, About Enabling Unified Audit Policies

Question No : 13

Which three resources might be prioritized between competing pluggable databases when creating a multitenant container database plan (CDB plan) using Oracle Database Resource Manager?

A. Maximum Undo per consumer group
B. Maximum Idle time
C. Parallel server limit
D. CPU
E. Exadata I/O
F. Local file system I/O

Answer: C,D,E

Explanation: http://docs.oracle.com/database/121/ADMIN/dbrm.htm#ADMIN11852

Question No : 14

You executed this command to create a password file:

```
$ orapwd file = orapworcl entries = 10 ignorecase = N
```

Which two statements are true about the password file?

A. It will permit the use of uppercase passwords for database users who have been granted the SYSOPER role.
B. It contains username and passwords of database users who are members of the OSOPER operating system group.
C. It contains usernames and passwords of database users who are members of the OSDBA operating system group.
D. It will permit the use of lowercase passwords for database users who have granted the SYSDBA role.
E. It will not permit the use of mixed case passwords for the database users who have been granted the SYSDBA role.

Answer: A,D

Explanation: * You can create a password file using the password file creation utility, ORAPWD.
* Adding Users to a Password File

When you grant SYSDBA or SYSOPER privileges to a user, that user's name and privilege information are added to the password file. If the server does not have an EXCLUSIVE password file (that is, if the initialization parameter REMOTE_LOGIN_PASSWORDFILE is NONE or SHARED, or the password file is missing), Oracle Database issues an error if you attempt to grant these privileges.

A user's name remains in the password file only as long as that user has at least one of these two privileges. If you revoke both of these privileges, Oracle Database removes the user from the password file.
* The syntax of the ORAPWD command is as follows:

```
ORAPWD FILE=filename [ENTRIES=numusers]
[FORCE={Y|N}] [IGNORECASE={Y|N}] [NOSYSDBA={Y|N}]
```

* IGNORECASE

If this argument is set to y, passwords are case-insensitive. That is, case is ignored when comparing the password that the user supplies during login with the password in the password file.

Question No : 15

You created an encrypted tablespace:

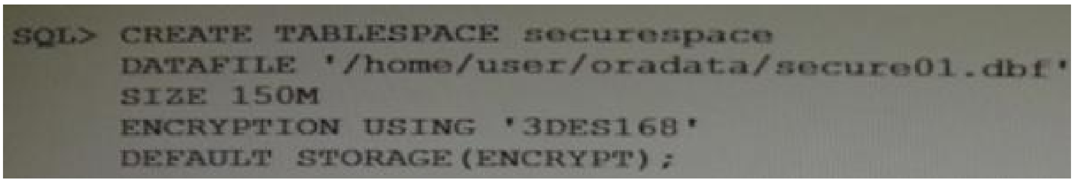

```
SQL> CREATE TABLESPACE securespace
     DATAFILE '/home/user/oradata/secure01.dbf'
     SIZE 150M
     ENCRYPTION USING '3DES168'
     DEFAULT STORAGE(ENCRYPT);
```

You then closed the encryption wallet because you were advised that this is secure.

Later in the day, you attempt to create the EMPLOYEES table in the SECURESPACE tablespace with the SALT option on the EMPLOYEE column.

Which is true about the result?

A. It creates the table successfully but does not encrypt any inserted data in the EMPNAME column because the wallet must be opened to encrypt columns with SALT.
B. It generates an error when creating the table because the wallet is closed.
C. It creates the table successfully, and encrypts any inserted data in the EMPNAME column because the wallet needs to be open only for tablespace creation.
D. It generates error when creating the table, because the salt option cannot be used with encrypted tablespaces.

Answer: B

Question No : 16

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 No : 17

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: A,C,E

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 No : 18

Examine the commands executed to monitor database operations:

```
$> conn sys oracle/oracle@prod as sysdba

SQL > VAR eid NUMBER

SQL > EXEC: eid := DBMS_SQL_MONITOR.BEGIN_OPERATION (‘batch_job’ ,
FORCED_TRACKING => ‘Y’);
```

Which two statements are true?

A. Database operations will be monitored only when they consume a significant amount of resource.
B. Database operations for all sessions will be monitored.
C. Database operations will be monitored only if the STATISTICS_LEVEL parameter is set

to TYPICAL and CONTROL_MANAGEMENT_PACK_ACCESS is set DIAGNISTIC + TUNING.
D. Only DML and DDL statements will be monitored for the session.
E. All subsequent statements in the session will be treated as one database operation and will be monitored.

Answer: C,E

Explanation: C: Setting the CONTROL_MANAGEMENT_PACK_ACCESS initialization parameter to DIAGNOSTIC+TUNING (default) enables monitoring of database operations. Real-Time SQL Monitoring is a feature of the Oracle Database Tuning Pack.

Note:
* The DBMS_SQL_MONITOR package provides information about Real-time SQL Monitoring and Real-time Database Operation Monitoring.
*(not B) BEGIN_OPERATION Function
starts a composite database operation in the current session.

/ (E) FORCE_TRACKING - forces the composite database operation to be tracked when the operation starts. You can also use the string variable 'Y'.

/ (not A) NO_FORCE_TRACKING - the operation will be tracked only when it has consumed at least 5 seconds of CPU or I/O time. You can also use the string variable 'N'.

Question No : 19

Which three tasks can be automatically performed by the Automatic Data Optimization feature of Information lifecycle Management (ILM)?

A. Tracking the most recent read time for a table segment in a user tablespace
B. Tracking the most recent write time for a table segment in a user tablespace
C. Tracking insert time by row for table rows
D. Tracking the most recent write time for a table block
E. Tracking the most recent read time for a table segment in the SYSAUX tablespace
F. Tracking the most recent write time for a table segment in the SYSAUX tablespace

Answer: A,B,D

Explanation:
Incorrect:
Not E, Not F When Heat Map is enabled, all accesses are tracked by the in-memory

activity tracking module. Objects in the SYSTEM and SYSAUX tablespaces are not tracked.

* To implement your ILM strategy, you can use Heat Map in Oracle Database to track data access and modification.
Heat Map provides data access tracking at the segment-level and data modification tracking at the segment and row level.

* To implement your ILM strategy, you can use Heat Map in Oracle Database to track data access and modification. You can also use Automatic Data Optimization (ADO) to automate the compression and movement of data between different tiers of storage within the database.

Reference: Automatic Data Optimization with Oracle Database 12c

with Oracle Database 12c

Question No : 20

Examine the following commands for redefining a table with Virtual Private Database (VPD) policies:

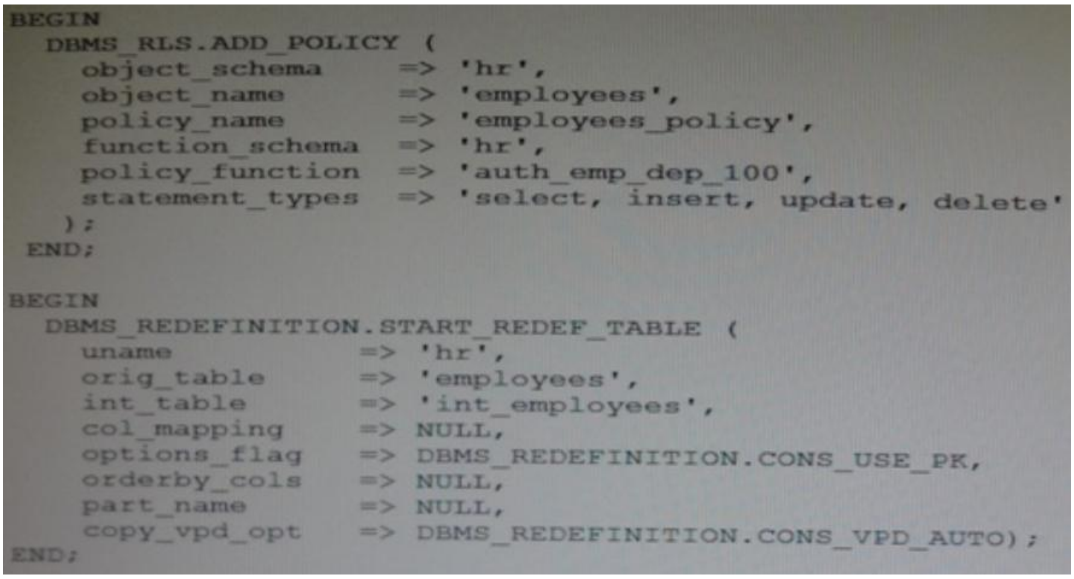

```
BEGIN
  DBMS_RLS.ADD_POLICY (
    object_schema       => 'hr',
    object_name         => 'employees',
    policy_name         => 'employees_policy',
    function_schema     => 'hr',
    policy_function     => 'auth_emp_dep_100',
    statement_types     => 'select, insert, update, delete'
   );
 END;

BEGIN
  DBMS_REDEFINITION.START_REDEF_TABLE (
    uname             => 'hr',
    orig_table        => 'employees',
    int_table         => 'int_employees',
    col_mapping       => NULL,
    options_flag      => DBMS_REDEFINITION.CONS_USE_PK,
    orderby_cols      => NULL,
    part_name         => NULL,
    copy_vpd_opt      => DBMS_REDEFINITION.CONS_VPD_AUTO);
END;
```

Which two statements are true about redefining the table?

A. All the triggers for the table are disabled without changing any of the column names or column types in the table.
B. The primary key constraint on the EMPLOYEES table is disabled during redefinition.
C. VPD policies are copied from the original table to the new table during online redefinition.
D. You must copy the VPD policies manually from the original table to the new table during online redefinition.

Answer: A,C

Explanation: The triggers cloned to the interim table are disabled until the redefinition is completed. Once the redefinition is complete, all cloned objects are renamed to the original names used by they objects they were cloned from.

Ref: http://www.oracle-base.com/articles/10g/online-table-redefinition-enhancements-10gr1.php

Question No : 21

Which two statements are true about the Oracle Direct Network File system (DNFS)?

A. It utilizes the OS file system cache.
B. A traditional NFS mount is not required when using Direct NFS.
C. Oracle Disk Manager can manage NFS on its own, without using the operating kernel NFS driver.
D. Direct NFS is available only in UNIX platforms.
E. Direct NFS can load-balance I/O traffic across multiple network adapters.

Answer: C,E

Explanation: E: Performance is improved by load balancing across multiple network interfaces (if available).

Note:
* To enable Direct NFS Client, you must replace the standard Oracle Disk Manager (ODM) library with one that supports Direct NFS Client.

Incorrect:

Not A: Direct NFS Client is capable of performing concurrent
direct I/O, which bypasses any operating system level caches and eliminates any
operating system write-ordering locks
Not B:
* To use Direct NFS Client, the NFS file systems must first be mounted and available
over regular NFS mounts.
* Oracle Direct NFS (dNFS) is an optimized NFS (Network File System) client that provides
faster and more scalable access to NFS storage located on NAS storage devices
(accessible over TCP/IP).
Not D: Direct NFS is provided as part of the database kernel, and is thus available on all
supported database platforms - even those that don't support NFS natively, like Windows.

Note:
* Oracle Direct NFS (dNFS) is an optimized NFS (Network File System) client that provides
faster and more scalable access to NFS storage located on NAS storage devices
(accessible over TCP/IP). Direct NFS is built directly into the database kernel - just like
ASM which is mainly used when using DAS or SAN storage.

* Oracle Direct NFS (dNFS) is an internal I/O layer that provides faster access to large NFS
files than traditional NFS clients.

Question No : 22

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 size
C. 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: A,C

Explanation: 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 No : 23

To enable the Database Smart Flash Cache, you configure the following parameters:

DB_FLASH_CACHE_FILE = ‘/dev/flash_device_1’ , ‘/dev/flash_device_2’

DB_FLASH_CACHE_SIZE=64G

What is the result when you start up the database instance?

A. It results in an error because these parameter settings are invalid.
B. One 64G flash cache file will be used.
C. Two 64G flash cache files will be used.
D. Two 32G flash cache files will be used.

Answer: A

Question No : 24

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: C,D,E

Question No : 25

Your multitenant container database (CDB) contains some pluggable databases (PDBs), you execute the following command in the root container:

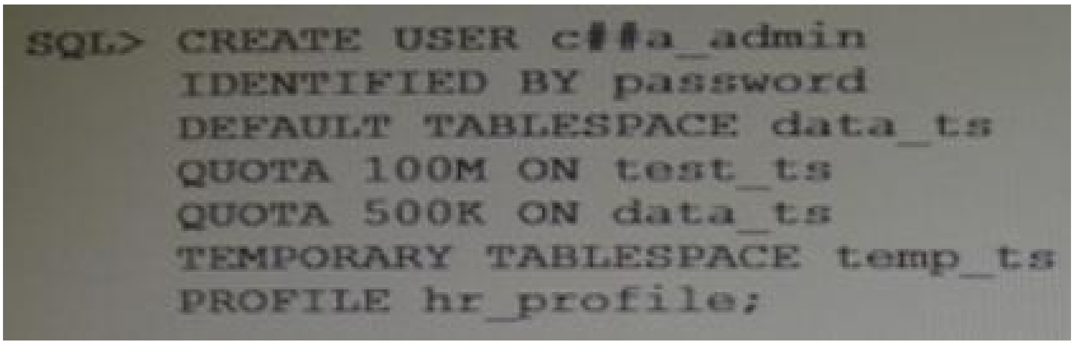

```
SQL> CREATE USER c##a_admin
     IDENTIFIED BY password
     DEFAULT TABLESPACE data_ts
     QUOTA 100M ON test_ts
     QUOTA 500K ON data_ts
     TEMPORARY TABLESPACE temp_ts
     PROFILE hr_profile;
```

Which two statements are true?

A. Schema objects owned by the C# # A_ADMIN common user can be shared across all PDBs.
B. The C # # A_ADMIN user will be able to use the TEMP_TS temporary tablespace only in root.
C. The command will, create a common user whose description is contained in the root and each PDB.
D. The schema for the common user C # # A_ADMIN can be different in each container.
E. The command will create a user in the root container only because the container clause is not used.

Answer: C,D

Question No : 26

Your multitenant container database (CDB) contains pluggable databases (PDBs), you are connected to the HR_PDB. You execute the following command:

SQL > CREATE UNDO TABLESPACE undotb01

DATAFILE ‘u01/oracle/rddb1/undotbs01.dbf’ SIZE 60M AUTOEXTEND ON;

What is the result?

A. It executes successfully and creates an UNDO tablespace in HR_PDB.
B. It falls and reports an error because there can be only one undo tablespace in a CDB.
C. It fails and reports an error because the CONTAINER=ALL clause is not specified in the command.
D. It fails and reports an error because the CONTAINER=CURRENT clause is not specified in the command.
E. It executes successfully but neither tablespace nor the data file is created.

Answer: E

Explanation: Interesting behavior in 12.1.0.1 DB of creating an undo tablespace in a PDB. With the new Multitenant architecture the undo tablespace resides at the CDB level and PDBs all share the same UNDO tablespace.

When the current container is a PDB, an attempt to create an undo tablespace fails without returning an error.

Question No : 27

You 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 11g Release 2 (11.2.0.2.0) 64-bit
- Character set: AL32UTF8
- National character set: AL16UTF16

- O/S: Oracle Linux 6 64-bit

The 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-bit

Which technique should you use to minimize down time while plugging this non-CDB into the CDB?

A. Transportable database
B. Transportable tablespace
C. Data Pump full export/import
D. The DBMS_PDB package
E. RMAN

Answer: B

Question No : 28

Identify two 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: A,D

Question No : 29

Identify three valid methods of opening, pluggable databases (PDBs).

A. ALTER PLUGGABLE DATABASE OPEN ALL ISSUED from the root
B. ALTER PLUGGABLE DATABASE OPEN ALL ISSUED from a PDB
C. ALTER PLUGGABLE DATABASE PDB OPEN issued from the seed
D. ALTER DATABASE PDB OPEN issued from the root
E. ALTER DATABASE OPEN issued from that PDB
F. ALTER PLUGGABLE DATABASE PDB OPEN issued from another PDB
G. ALTER PLUGGABLE DATABASE OPEN issued from that PDB

Answer: A,E,G

Explanation: E: You can perform all ALTER PLUGGABLE DATABASE tasks by connecting to a PDB and running the corresponding ALTER DATABASE statement. This functionality is provided to maintain backward compatibility for applications that have been migrated to a CDB environment.

AG: When you issue an ALTER PLUGGABLE DATABASE OPEN statement, READ WRITE is the default unless a PDB being opened belongs to a CDB that is used as a physical standby database, in which case READ ONLY is the default.

You can specify which PDBs to modify in the following ways:

List one or more PDBs.

Specify ALL to modify all of the PDBs.

Specify ALL EXCEPT to modify all of the PDBs, except for the PDBs listed.

Question No : 30

You conned using SQL Plus to the root container of a multitenant container database (CDB) with SYSDBA privilege.

The CDB has several pluggable databases (PDBs) open in the read/write mode.

There are ongoing transactions in both the CDB and PDBs.

What happens alter issuing the SHUTDOWN TRANSACTIONAL statement?

A. The shutdown proceeds immediately.
The shutdown proceeds as soon as all transactions in the PDBs are either committed or rolled hack.
B. The shutdown proceeds as soon as all transactions in the CDB are either committed or rolled back.
C. The shutdown proceeds as soon as all transactions in both the CDB and PDBs are either committed or rolled back.
D. The statement results in an error because there are open PDBs.

Answer: B

Explanation: * SHUTDOWN [ABORT | IMMEDIATE | NORMAL | TRANSACTIONAL [LOCAL]]

Shuts down a currently running Oracle Database instance, optionally closing and dismounting a database. If the current database is a pluggable database, only the pluggable database is closed. The consolidated instance continues to run.

Shutdown commands that wait for current calls to complete or users to disconnect such as SHUTDOWN NORMAL and SHUTDOWN TRANSACTIONAL have a time limit that the SHUTDOWN command will wait. If all events blocking the shutdown have not occurred within the time limit, the shutdown command cancels with the following message:

ORA-01013: user requested cancel of current operation

* If logged into a CDB, shutdown closes the CDB instance.

To shutdown a CDB or non CDB, you must be connected to the CDB or non CDB instance that you want to close, and then enter

SHUTDOWN
Database closed.
Database dismounted.
Oracle instance shut down.

To shutdown a PDB, you must log into the PDB to issue the SHUTDOWN command.

SHUTDOWN
Pluggable Database closed.

Note:
* Prerequisites for PDB Shutdown

When the current container is a pluggable database (PDB), the SHUTDOWN command can only be used if:

The current user has SYSDBA, SYSOPER, SYSBACKUP, or SYSDG system privilege.

The privilege is either commonly granted or locally granted in the PDB.

The current user exercises the privilege using AS SYSDBA, AS SYSOPER, AS SYSBACKUP, or AS SYSDG at connect time.

To close a PDB, the PDB must be open.

Question No : 31

An 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 = FALSE

Which 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: A,C,D,E

Explanation:

http://docs.oracle.com/database/121/ADMIN/cdb_admin.htm

Question No : 32

What are three purposes of the RMAN “FROM” clause?

A. to support PUSH-based active database duplication
B. to support synchronization of a standby database with the primary database in a Data environment
C. To support PULL-based active database duplication
D. To support file restores over the network in a Data Guard environment
E. To support file recovery over the network in a Data Guard environment

Answer: B,D,E

Explanation:

DE:

* With a control file autobackup, **RMAN** can recover the database even if the current control **file, recovery** catalog, and server parameter file are inaccessible.

* RMAN uses a recovery catalog to track filenames for all database files in a Data Guard environment. A recovery catalog is a database schema used by RMAN to store metadata about one or more Oracle databases. The catalog also records where the online redo logs, standby redo logs, tempfiles, archived redo logs, backup sets, and image copies are created.

Question No : 33

You run a script that completes successfully using SQL*Plus that performs these actions:

1. Creates a multitenant container database (CDB)

2. Plugs in three pluggable databases (PDBs)

3. Shuts down the CDB instance

4. Starts up the CDB instance using STARTUP OPEN READ WRITE

Which two statements are true about the outcome after running the script?

A. The seed will be in mount state.
B. The seed will be opened read-only.
C. The seed will be opened read/write.
D. The other PDBs will be in mount state.
E. The other PDBs will be opened read-only.
F. The PDBs will be opened read/write.

Answer: B,D

Explanation: B: The seed is always read-only.

D: Pluggable databases can be started and stopped using SQL*Plus commands or the ALTER PLUGGABLE DATABASE command.

Question No : 34

You execute the following piece of code with appropriate privileges:

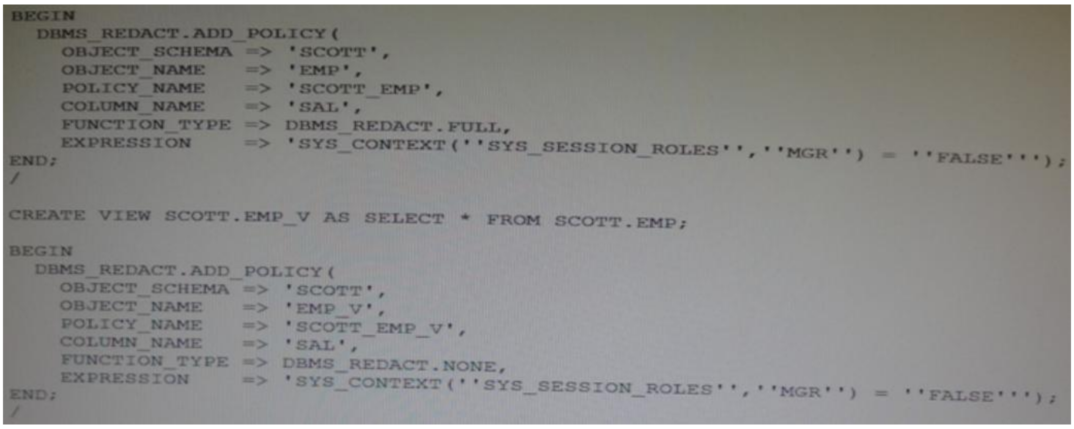

```
BEGIN
  DBMS_REDACT.ADD_POLICY(
    OBJECT_SCHEMA => 'SCOTT',
    OBJECT_NAME   => 'EMP',
    POLICY_NAME   => 'SCOTT_EMP',
    COLUMN_NAME   => 'SAL',
    FUNCTION_TYPE => DBMS_REDACT.FULL,
    EXPRESSION    => 'SYS_CONTEXT(''SYS_SESSION_ROLES'',''MGR'') = ''FALSE''');
END;
/

CREATE VIEW SCOTT.EMP_V AS SELECT * FROM SCOTT.EMP;

BEGIN
  DBMS_REDACT.ADD_POLICY(
    OBJECT_SCHEMA => 'SCOTT',
    OBJECT_NAME   => 'EMP_V',
    POLICY_NAME   => 'SCOTT_EMP_V',
    COLUMN_NAME   => 'SAL',
    FUNCTION_TYPE => DBMS_REDACT.NONE,
    EXPRESSION    => 'SYS_CONTEXT(''SYS_SESSION_ROLES'',''MGR'') = ''FALSE''');
END;
/
```

User SCOTT has been granted the CREATE SESSION privilege and the MGR role.

Which two statements are true when a session logged in as SCOTT queries the SAL column in the view and the table?

A. Data is redacted for the EMP.SAL column only if the SCOTT session does not have the MGR role set.
B. Data is redacted for EMP.SAL column only if the SCOTT session has the MGR role set.
C. Data is never redacted for the EMP_V.SAL column.
D. Data is redacted for the EMP_V.SAL column only if the SCOTT session has the MGR role set.
E. Data is redacted for the EMP_V.SAL column only if the SCOTT session does not have the MGR role set.

Answer: A,C

Explanation:

Note:

* DBMS_REDACT.FULL completely redacts the column data.
* DBMS_REDACT.NONE applies no redaction on the column data. Use this function for development testing purposes. LOB columns are not supported.
* The DBMS_REDACT package provides an interface to Oracle Data Redaction, which enables you to mask (redact) data that is returned from queries issued by low-privileged users or an application.

* If you create a view chain (that is, a view based on another view), then the Data Redaction policy also applies throughout this view chain. The policies remain in effect all of the way up through this view chain, but if another policy is created for one of these views, then for the columns affected in the subsequent views, this new policy takes precedence.

Question No : 35

Identify three benefits of Unified Auditing.

A. Decreased use of storage to store audit trail rows in the database.
B. It improves overall auditing performance.
C. It guarantees zero-loss auditing.
D. The audit trail cannot be easily modified because it is read-only.
E. It automatically audits Recovery Manager (RMAN) events.

Answer: B,D,E

Explanation: https://blogs.oracle.com/imc/entry/oracle_database_12c_new_unified

Question No : 36

A warehouse fact table in your Oracle 12c Database is range-partitioned by month and accessed frequently with queries that span multiple partitions

The table has a local prefixed, range partitioned index.

Some of these queries access very few rows in some partitions and all the rows in other partitions, but these queries still perform a full scan for all accessed partitions.

This commonly occurs when the range of dates begins at the end of a month or ends close to the start of a month.

You want an execution plan to be generated that uses indexed access when only a few rows are accessed from a segment, while still allowing full scans for segments where many rows are returned.

Which three methods could transparently help to achieve this result?

A. Using a partial local Index on the warehouse fact table month column with indexing disabled to the table partitions that return most of their rows to the queries.
B. Using a partial local Index on the warehouse fact table month column with indexing disabled for the table partitions that return a few rows to the queries.
C. Using a partitioned view that does a UNION ALL query on the partitions of the warehouse fact table, which retains the existing local partitioned column.
D. Converting the partitioned table to a partitioned view that does a UNION ALL query on the monthly tables, which retains the existing local partitioned column.
E. Using a partial global index on the warehouse fact table month column with indexing disabling for the table partitions that return most of their rows to the queries.
F. Using a partial global index on the warehouse fact table month column with indexing disabled for the table partitions that return a few rows to the queries.

Answer: A,C,E

Question No : 37

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: A,D

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:

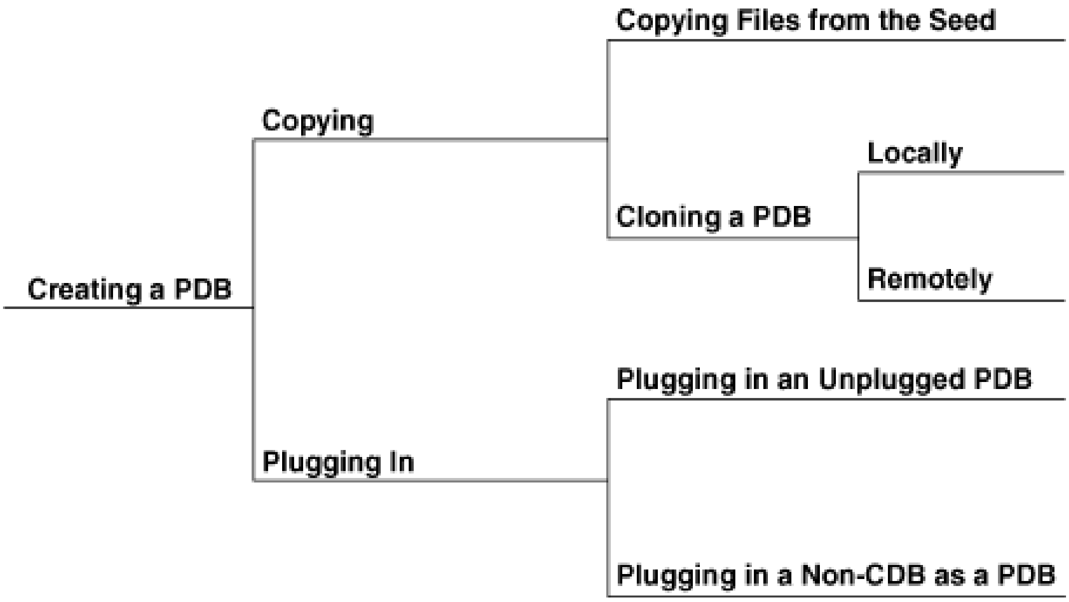
Creating a PDB
Copying
Copying Files from the Seed
Cloning a PDB
Locally
Remotely
Plugging In
Plugging in an Unplugged PDB
Plugging in a Non-CDB as a PDB

Description of cncpt358.png follows

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 No : 38

Which three are true about the large pool for an Oracle database instance that supports shared server connections?

A. Allocates memory for RMAN backup and restore operations
B. Allocates memory for shared and private SQL areas
C. Contains a cursor area for storing runtime information about cursors
D. Contains stack space
E. Contains a hash area performing hash joins of tables

Answer: A,B,C

Explanation: The large pool can provide large memory allocations for the following:
/ (B)UGA (User Global Area) for the shared server and the Oracle XA interface (used where transactions interact with multiple databases)
/Message buffers used in the parallel execution of statements
/ (A) Buffers for Recovery Manager (RMAN) I/O slaves

Note:
* large pool
Optional area in the SGA that provides large memory allocations for backup and restore operations, I/O server processes, and session memory for the shared server and Oracle XA.

* Oracle XA
An external interface that allows global transactions to be coordinated by a transaction manager other than Oracle Database.

* UGA
User global area. Session memory that stores session variables, such as logon information, and can also contain the OLAP pool.

* Configuring the Large Pool

Unlike the shared pool, the large pool does not have an LRU list (not D). Oracle Database does not attempt to age objects out of the large pool. Consider configuring a large pool if the database instance uses any of the following Oracle Database features:

* Shared server
In a shared server architecture, the session memory for each client process is included in the shared pool.

* Parallel query
Parallel query uses shared pool memory to cache parallel execution message buffers.

* Recovery Manager
Recovery Manager (RMAN) uses the shared pool to cache I/O buffers during backup and restore operations. For I/O server processes, backup, and restore operations, Oracle Database allocates buffers that are a few hundred kilobytes in size.

Question No : 39

You are administering a database stored in Automatic Storage management (ASM). The files are stored in the DATA disk group. You execute the following command:

SQL > ALTER DISKGROUP data ADD ALIAS '+data/prod/myfile.dbf' FOR '
+data/prod/myfile.dbf'

What is the result?

A. The file '+data.231.54769' is physically relocated to '+data/prod' and renamed as 'myfile.dbf'.
B. The file '+data.231.54769' is renamed as 'myfile.dbf', and copied to '+data/prod'.
C. The file '+data.231.54769' remains in the same location and a synonym 'myfile.dbf' is created.
D. The file 'myfile.dbf' is created in '+data/prod' and the reference to '+data.231.54769' in the data dictionary removed.

Answer: C

Explanation: ADD ALIAS
Use this clause to create an alias name for an Oracle ASM filename. The alias_name consists of the full directory path and the alias itself.

Question No : 40

To implement Automatic Management (AMM), you set the following parameters:

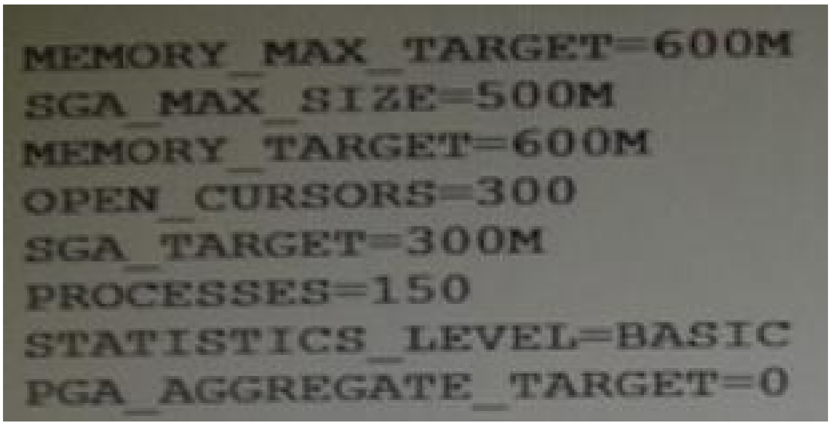

```
MEMORY_MAX_TARGET=600M
SGA_MAX_SIZE=500M
MEMORY_TARGET=600M
OPEN_CURSORS=300
SGA_TARGET=300M
PROCESSES=150
STATISTICS_LEVEL=BASIC
PGA_AGGREGATE_TARGET=0
```

When you try to start the database instance with these parameter settings, you receive the following error message:

SQL > startup

ORA-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: B

Explanation:

Example:

SQL> startup force

ORA-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 No : 41

Examine the query and its output executed In an RDBMS Instance:

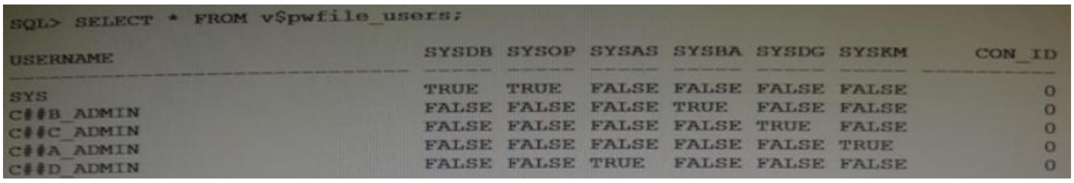

```
SQL> SELECT * FROM v$pwfile_users;

USERNAME                      SYSDB SYSOP SYSAS SYSBA SYSDG SYSKM     CON_ID
----------------------------- ----- ----- ----- ----- ----- ----- ----------
SYS                           TRUE  TRUE  FALSE FALSE FALSE FALSE          0
C##B_ADMIN                    FALSE FALSE FALSE TRUE  FALSE FALSE          0
C##C_ADMIN                    FALSE FALSE FALSE FALSE TRUE  FALSE          0
C##A_ADMIN                    FALSE FALSE FALSE FALSE FALSE TRUE           0
C##D_ADMIN                    FALSE FALSE TRUE  FALSE FALSE FALSE          0
```

Which three statements are true about the users (other than sys) in the output?

A. The C # # B_ADMIN user can perform all backup and recovery operations using RMAN only.
B. The C # # C_ADMIN user can perform the data guard operation with Data Guard Broker.

C. The C # # A_ADMIN user can perform wallet operations.
D. The C # # D_ADMIN user can perform backup and recovery operations for Automatic Storage Management (ASM).
E. The C # # B_ADMIN user can perform all backup and recovery operations using RMAN or SQL* Plus.

Answer: B,D,E

Explanation:

B: SYSDG administrative privilege has ability to perform Data Guard operations (including startup and shutdown) using Data Guard Broker or dgmgrl.

D: SYSASM

The new (introduced in 11g) SYSASM role to manage the ASM instance, variable extent sizes to reduce shared pool usage, and the ability of an instance to read from a specific disk of a diskgroup

E (Not A): SYSDBA is like a role in the sense that it is granted, but SYSDBA is a special built-in privilege to allow the DBA full control over the database

Incorrect:

Not C: SYSKM. SYSKM administrative privilege has ability to perform transparent data encryption wallet operations.

Note:

Use the V$PWFILE_USERS view to see the users who have been granted administrative privileges.

Question No : 42

You use the segment advisor to help determine objects for which space may be reclaimed.

Which three statements are true about the advisor given by the segment advisor?

A. It may advise the use of online table redefinition for tables in dictionary managed tablespace.
B. It may advise the use of segment shrink for tables in dictionary managed tablespaces it the no chained rows.
C. It may advise the use of online table redefinition for tables in locally managed

tablespaces
D. It will detect and advise about chained rows.
E. It may advise the use of segment shrink for free list managed tables.

Answer: A,D,E

Explanation: (http://docs.oracle.com/database/121/ADMIN/schema.htm#ADMIN11601)
The Segment Advisor generates the following types of advice:

If the Segment Advisor determines that an object has a significant amount of free space, it recommends online segment shrink. If the object is a table that is not eligible for shrinking, as in the case of a table in a tablespace without automatic segment space management, the Segment Advisor recommends online table redefinition.

If the Segment Advisor determines that a table could benefit from compression with the advanced row compression method, it makes a recommendation to that effect. (Automatic Segment Advisor only. See "Automatic Segment Advisor".)

If the Segment Advisor encounters a table with row chaining above a certain threshold, it records that fact that the table has an excess of chained rows.

Question No : 43

Examine these two statements:

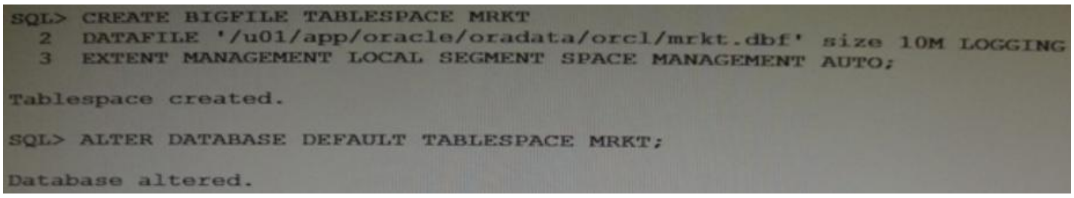

```
SQL> CREATE BIGFILE TABLESPACE MRKT
  2  DATAFILE '/u01/app/oracle/oradata/orcl/mrkt.dbf' size 10M LOGGING
  3  EXTENT MANAGEMENT LOCAL SEGMENT SPACE MANAGEMENT AUTO;

Tablespace created.

SQL> ALTER DATABASE DEFAULT TABLESPACE MRKT;

Database altered.
```

Which three are true about the MRKT tablespace?

A. The MRKT tablespace is created as a small file tablespace, because the file size is less than the minimum required for big file files.
B. The MRKT tablespace may be dropped if it has no contents.
C. Users who were using the old default tablespace will have their default tablespaces changed to the MRKT tablespace.
D. No more data files can be added to the tablespace.
E. The relative file number of the tablespace is not stored in rowids for the table rows that

are stored in the MRKT tablespace.

Answer: C,D,E

Question No : 44

In your multitenant container database (CDB) containing pluggable database (PDBs), you granted the CREATE TABLE privilege to the common user C # # A_ADMIN in root and all PDBs. You execute the following command from the root container:

SQL > REVOKE create table FROM C # # A_ADMIN;

What is the result?

A. It executes successfully and the CREATE TABLE privilege is revoked from C # # A_ADMIN in root only.
B. It fails and reports an error because the CONTAINER=ALL clause is not used.
C. It excludes successfully and the CREATE TABLE privilege is revoked from C # # A_ADMIN in root and all PDBs.
D. It fails and reports an error because the CONTAINER=CURRENT clause is not used.
E. It executes successfully and the CREATE TABLE privilege is revoked from C # # A_ADMIN in all PDBs.

Answer: A
Explanation: REVOKE ..FROM

If the current container is the root:

/ Specify CONTAINER = CURRENT to revoke a locally granted system privilege, object privilege, or role from a common user or common role. The privilege or role is revoked from the user or role only in the root. This clause does not revoke privileges granted with CONTAINER = ALL.

/ Specify CONTAINER = ALL to revoke a commonly granted system privilege, object privilege on a common object, or role from a common user or common role. The privilege or role is revoked from the user or role across the entire CDB. This clause can revoke only a privilege or role granted with CONTAINER = ALL from the specified common user or common role. This clause does not revoke privileges granted locally with CONTAINER = CURRENT. However, any locally granted privileges that depend on the commonly granted privilege being revoked are also revoked.

If you omit this clause, then CONTAINER = CURRENT is the default.

Reference: Oracle Database SQL Language Reference 12c, Revoke

Question No : 45

You notice that the performance of your production 24/7 Oracle database significantly degraded. Sometimes you are not able to connect to the instance because it hangs. You do not want to restart the database instance.

How can you detect the cause of the degraded performance?

A. Enable Memory Access Mode, which reads performance data from SGA.
B. Use emergency monitoring to fetch data directly from SGA analysis.
C. Run Automatic Database Diagnostic Monitor (ADDM) to fetch information from the latest Automatic Workload Repository (AWR) snapshots.
D. Use Active Session History (ASH) data and hang analysis in regular performance monitoring.
E. Run ADDM in diagnostic mode.

Answer: B

Question No : 46

Which two statements are true about the use of the procedures listed in the v$sysaux_occupants.move_procedure column?

A. The procedure may be used for some components to relocate component data to the SYSAUX tablespace from its current tablespace.
B. The procedure may be used for some components to relocate component data from the SYSAUX tablespace to another tablespace.
C. All the components may be moved into SYSAUX tablespace.
D. All the components may be moved from the SYSAUX tablespace.

Answer: A,B

Explanation:

http://www.dba-oracle.com/t_v_sysaux_contents_tips.htm

Question No : 47

Examine the following steps of privilege analysis for checking and revoking excessive, unused privileges granted to users:

1. Create a policy to capture the privilege used by a user for privilege analysis.

2. Generate a report with the data captured for a specified privilege capture.

3. Start analyzing the data captured by the policy.

4. Revoke the unused privileges.

5. Compare the used and unused privileges' lists.

6. Stop analyzing the data.

Identify the correct sequence of steps.

A. 1, 3, 5, 6, 2, 4
B. 1, 3, 6, 2, 5, 4
C. 1, 3, 2, 5, 6, 4
D. 1, 3, 2, 5, 6, 4
E. 1, 3, 5, 2, 6, 4

Answer: B

Explanation: 1. Create a policy to capture the privilege used by a user for privilege analysis.
3. Start analyzing the data captured by the policy.
6. Stop analyzing the data.
2. Generate a report with the data captured for a specified privilege capture.
5. Compare the used and unused privileges' lists.
4. Revoke the unused privileges.

Question No : 48