


Oracle 1z0-060



Upgrade to Oracle Database 12c

Version: 6.6

QUESTION NO: 1

Your multitenant container (CDB) contains two pluggable databases (PDB), HR_PDB and ACCOUNTS_PDB, both of which use the CDB tablespace. The temp file is called temp01.tmp.

A user issues a query on a table on one of the PDBs and receives the following error:

ERROR at line 1:

ORA-01565: error in identifying file '/u01/app/oracle/oradata/CDB1/temp01.tmp'

ORA-27037: unable to obtain file status

Identify two ways to rectify the error.

- A. Add a new temp file to the temporary tablespace and drop the temp file that that produced the error.
- B. Shut down the database instance, restore the temp01.tmp file from the backup, and then restart the database.
- C. Take the temporary tablespace offline, recover the missing temp file by applying redo logs, and then bring the temporary tablespace online.
- D. Shutdown the database instance, restore and recover the temp file from the backup, and then open the database with RESETLOGS.
- E. Shut down the database instance and then restart the CDB and PDBs.

Answer: A,E

Explanation: * Because temp files cannot be backed up and because no redo is ever generated for them, RMAN never restores or recovers temp files. RMAN does track the names of temp files, but only so that it can automatically re-create them when needed.

* If you use RMAN in a Data Guard environment, then RMAN transparently converts primary control files to standby control files and vice versa. RMAN automatically updates file names for data files, online redo logs, standby redo logs, and temp files when you issue RESTORE and RECOVER.

QUESTION NO: 2

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.CONSTRAINT_USE_PK,
    orderby_cols  => NULL,
    part_name      => NULL,
    copy_vpd_opt  => DBMS_REDEFINITION.CONSTRAINT_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 the objects they were cloned from.

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

QUESTION NO: 3

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: 4

Which statement is true about Oracle Net Listener?

- A. It acts as the listening endpoint for the Oracle database instance for all local and non-local user connections.
- B. A single listener can service only one database instance and multiple remote client connections.
- C. Service registration with the listener is performed by the process monitor (LREG) process of each database instance.
- D. The listener.ora configuration file must be configured with one or more listening protocol addresses to allow remote users to connect to a database instance.
- E. The listener.ora configuration file must be located in the ORACLE_HOME/network/admin directly.

Answer: C

Explanation:

QUESTION NO: 5

You are administering a database stored in Automatic Storage Management (ASM). You use RMAN to back up the database and the MD_BACKUP command to back up the ASM metadata regularly. You lost an ASM disk group DG1 due to hardware failure.

In which three ways can you re-create the lost disk group and restore the data?

- A. Use the MD_RESTORE command to restore metadata for an existing disk group by passing the existing disk group name as an input parameter and use RMAN to restore the data.
- B. Use the MKDGM command to restore the disk group with the same configuration as the backed-up disk group and data on the disk group.
- C. Use the MD_RESTORE command to restore the disk group with the changed disk group specification, failure group specification, name, and other attributes and use RMAN to restore the

data.

D. Use the MKDG command to restore the disk group with the same configuration as the backed-up disk group name and same set of disks and failure group configuration, and use RMAN to restore the data.

E. Use the MD_RESTORE command to restore both the metadata and data for the failed disk group.

F. Use the MKDG command to add a new disk group DG1 with the same or different specifications for failure group and other attributes and use RMAN to restore the data.

Answer: C,D,F

Explanation:

QUESTION NO: 6

Your multitenant container database, CDB1, is running in ARCHIVELOG mode and has two pluggable databases, HR_PDB and ACCOUNTS_PDB. An RMAN backup exists for the database.

You issue the command to open ACCOUNTS_PDB and find that the USERDATA.DBF data file for the default permanent tablespace USERDATA belonging to ACCOUNTS_PDB is corrupted.

What should you do before executing the commands to restore and recover the data file in ACCOUNTS_PDB?

A. Place CDB1 in the mount stage and then the USERDATA tablespace offline in ACCOUNTS_PDB.

B. Place CDB1 in the mount stage and issue the ALTER PLUGGABLE DATABASE accounts_pdb CLOSE IMMEDIATE command.

C. Issue the ALTER PLUGGABLE DATABASE accounts_pdb RESTRICTED command.

D. Take the USERDATA tablespace offline in ACCOUNTS_PDB.

Answer: D

Explanation: * You can take an online tablespace offline so that it is temporarily unavailable for general use. The rest of the database remains open and available for users to access data.

Conversely, you can bring an offline tablespace online to make the schema objects within the tablespace available to database users. The database must be open to alter the availability of a tablespace.

QUESTION NO: 7

Which Oracle Database component is audited by default if the unified Auditing option is enabled?

- A. Oracle Data Pump
- B. Oracle Recovery Manager (RMAN)
- C. Oracle Label Security
- D. Oracle Database Vault
- E. Oracle Real Application Security

Answer: B

Explanation:

QUESTION NO: 8

Your multitenant container (CDB) containing three pluggable databases (PDBs) is running in ARCHIVELOG mode. You find that the SYSAUX tablespace is corrupted in the root container.

The steps to recover the tablespace are as follows:

1. Mount the CDB.
2. Close all the PDBs.
3. Open the database.
4. Apply the archive redo logs.
5. Restore the data file.
6. Take the SYSAUX tablespace offline.
7. Place the SYSAUX tablespace online.
8. Open all the PDBs with RESETLOGS.
9. Open the database with RESETLOGS.
10. Execute the command SHUTDOWN ABORT.

Which option identifies the correct sequence to recover the SYSAUX tablespace?

- A. 6, 5, 4, 7
- B. 10, 1, 2, 5, 8
- C. 10, 1, 2, 5, 4, 9, 8
- D. 10, 1, 5, 8, 10

Answer: A

Explanation: RMAN> ALTER TABLESPACE sysaux OFFLINE IMMEDIATE;

```
RMAN> RESTORE TABLESPACE sysaux;  
RMAN> RECOVER TABLESPACE sysaux;  
RMAN> ALTER TABLESPACE sysaux ONLINE;
```

* Example:

While evaluating the 12c beta3 I was not able to do the recover while testing "all pdb files lost".
Cannot close the pdb as the system datafile was missing...

So only option to recover was:

Shutdown cdb (10)

startup mount; (1)

restore pluggable database

recover pluggable database

alter database open;

alter pluggable database name open;

Oracle support says: You should be able to close the pdb and restore/recover the system tablespace of PDB.

* Inconsistent backups are usually created by taking online database backups. You can also make an inconsistent backup by backing up data files while a database is closed, either:

/ Immediately after the crash of an Oracle instance (or, in an Oracle RAC configuration, all instances)

/ After shutting down the database using SHUTDOWN ABORT

Inconsistent backups are only useful if the database is in ARCHIVELOG mode and all archived redo logs created since the backup are available.

* Open the database with the RESETLOGS option after finishing recovery:

```
SQL> ALTER DATABASE OPEN RESETLOGS;
```

QUESTION NO: 9

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: 10

In order to exploit some new storage tiers that have been provisioned by a storage administrator, the partitions of a large heap table must be moved to other tablespaces in your Oracle 12c database?

Both local and global partitioned B-tree Indexes are defined on the table.

A high volume of transactions access the table during the day and a medium volume of transactions access it at night and during weekends.

Minimal disruption to availability is required.

Which three statements are true about this requirement?

- A. The partitions can be moved online to new tablespaces.
- B. Global indexes must be rebuilt manually after moving the partitions.
- C. The partitions can be compressed in the same tablespaces.
- D. The partitions can be compressed in the new tablespaces.
- E. Local indexes must be rebuilt manually after moving the partitions.

Answer: A,C,D

Explanation: A: You can create and rebuild indexes online. Therefore, you can update base tables at

the same time you are building or rebuilding indexes on that table. You can perform DML operations while the index build is taking place, but DDL operations are not allowed. Parallel execution is not supported when creating or rebuilding an index online.

D: Moving (Rebuilding) Index-Organized Tables

Because index-organized tables are primarily stored in a B-tree index, you can encounter fragmentation as a consequence of incremental updates. However, you can use the ALTER TABLE...MOVE statement to rebuild the index and reduce this fragmentation.

C: If a table can be compressed in the new tablespace, also it can be compressed in the same tablespace.

<http://www.oracle.com/technetwork/issue-archive/2014/14-may/o34dba-2193424.html>

Incorrect:

Not B, not E: Local and Global indexes can be automatically rebuild with UPDATE INDEXES when you move the table.

QUESTION NO: 11

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: 12

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: 13

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

Explanation:

QUESTION NO: 14

You plan to use the In Database Archiving feature of Oracle Database 12c, and store rows that are inactive for over three months, in Hybrid Columnar Compressed (HCC) format.

Which three storage options support the use of HCC?

- A.** ASM disk groups with ASM disks consisting of Exadata Grid Disks.
- B.** ASM disk groups with ASM disks consisting of LUNS on any Storage Area Network array
- C.** ASM disk groups with ASM disks consisting of any zero padded NFS-mounted files
- D.** Database files stored in ZFS and accessed using conventional NFS mounts.

- E. Database files stored in ZFS and accessed using the Oracle Direct NFS feature
- F. Database files stored in any file system and accessed using the Oracle Direct NFS feature
- G. ASM disk groups with ASM disks consisting of LUNs on Pillar Axiom Storage arrays

Answer: A,E,G

Explanation: HCC requires the use of Oracle Storage – Exadata (A), Pillar Axiom (G) or Sun ZFS Storage Appliance (ZFSSA).

Note:

* Hybrid Columnar Compression, initially only available on Exadata, has been extended to support Pillar Axiom and Sun ZFS Storage Appliance (ZFSSA) storage when used with Oracle Database Enterprise Edition 11.2.0.3 and above

* Oracle offers the ability to manage NFS using a feature called Oracle Direct NFS (dNFS). Oracle Direct NFS implements NFS V3 protocol within the Oracle database kernel itself. Oracle Direct NFS client overcomes many of the challenges associated with using NFS with the Oracle Database with simple configuration, better performance than traditional NFS clients, and offers consistent configuration across platforms.

QUESTION NO: 15

In your multitenant container database (CDB) containing pluggable databases (PDB), users complain about performance degradation.

How does real-time Automatic database Diagnostic Monitor (ADDM) check performance degradation and provide solutions?

- A. It collects data from SGA and compares it with a preserved snapshot.
- B. It collects data from SGA, analyzes it, and provides a report.
- C. It collects data from SGA and compares it with the latest snapshot.
- D. It collects data from both SGA and PGA, analyzes it, and provides a report.

Answer: B

Explanation:

Note:

* The multitenant architecture enables an Oracle database to function as a multitenant container database (CDB) that includes zero, one, or many customer-created pluggable databases (PDBs). A PDB is a portable collection of schemas, schema objects, and nonschema objects that appears to an Oracle Net client as a non-CDB. All Oracle databases before Oracle Database 12c were non-CDBs.

* The System Global Area (SGA) is a group of shared memory areas that are dedicated to an

Oracle "instance" (an instance is your database programs and RAM).

* The PGA (Program or Process Global Area) is a memory area (RAM) that stores data and control information for a single process.

QUESTION NO: 16

The tnsnames.ora file has an entry for the service alias ORCL as follows:

```
ORCL =  
  ( DESCRIPTION =  
    ( ADDRESS = ( PROTOCOL = TCP ) ( HOST = 10.156.24.216 ) ( PORT = 1521 ) )  
    ( CONNECT_DATA =  
      ( SERVER = DEDICATED )  
      ( SERVICE_NAME = orcl.oracle.com )  
    )  
  )  
)
```

The TNS ping command executes successfully when tested with ORCL; however, from the same OS user session, you are not able to connect to the database instance with the following command:

```
SQL > CONNECT scott/tiger@orcl
```

What could be the reason for this?

- A. The listener is not running on the database node.
- B. The TNS_ADMIN environment variable is set to the wrong value.
- C. The orcl.oracle.com database service is not registered with the listener.
- D. The DEFAULT_DOMAIN parameter is set to the wrong value in the sqlnet.ora file.
- E. The listener is running on a different port.

Answer: C

Explanation: Service registration enables the listener to determine whether a database service and its service handlers are available. A service handler is a dedicated server process or dispatcher that acts as a connection point to a database. During registration, the LREG process provides the listener with the instance name, database service names, and the type and addresses of service handlers. This information enables the listener to start a service handler when a client request arrives.

QUESTION NO: 17

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: 18

You database is running an ARCHIVELOG mode.

The following parameter are set in your database instance:

LOG_ARCHIVE_FORMAT = arch+%t_%r.arc

LOG_ARCHIVE_DEST_1 = 'LOCATION = /disk1/archive'

DB_RECOVERY_FILE_DEST_SIZE = 50G

DB_RECOVERY_FILE = '/u01/oradata'

Which statement is true about the archived redo log files?

- A. They are created only in the location specified by the LOG_ARCHIVE_DEST_1 parameter.
- B. They are created only in the Fast Recovery Area.
- C. They are created in the location specified by the LOG_ARCHIVE_DEST_1 parameter and in the default location \$ORACLE_HOME/dbs/arch.
- D. They are created in the location specified by the LOG_ARCHIVE_DEST_1 parameter and the location specified by the DB_RECOVERY_FILE_DEST parameter.

Answer: A

Explanation: You can choose to archive redo logs to a single destination or to multiple destinations.

Destinations can be local—within the local file system or an Oracle Automatic Storage Management (Oracle ASM) disk group—or remote (on a standby database). When you archive to multiple destinations, a copy of each filled redo log file is written to each destination. These redundant copies help ensure that archived logs are always available in the event of a failure at one of the destinations.

To archive to only a single destination, specify that destination using the LOG_ARCHIVE_DEST and LOG_ARCHIVE_DUPLEX_DEST initialization parameters.

ARCHIVE_DEST initialization parameter. To archive to multiple destinations, you can choose to archive to two or more locations using the LOG_ARCHIVE_DEST_n initialization parameters, or to archive only to a primary and secondary destination using the LOG_ARCHIVE_DEST and LOG_ARCHIVE_DUPLEX_DEST initialization parameters.

QUESTION NO: 19

Your multitenant container database (CDB) is running in ARCHIVELOG mode. You connect to the CDB RMAN.

Examine the following command and its output:

```

RMAN> SELECT con_id, name, open_mode FROM V$PDBS;
  CON_ID NAME                OPEN_MODE
-----
  2 PDB$SEED                 READ ONLY
  3 PDB2_1                   MOUNTED
  4 PDB2_2                   MOUNTED

```

You execute the following command:

```
RMAN > BACKUP DATABASE PLUS ARCHIVELOG;
```

Which data files will be backed up?

- A. Data files that belong to only the root container
- B. Data files that belong to the root container and all the pluggable databases (PDBs)
- C. Data files that belong to only the root container and PDB\$SEED
- D. Data files that belong to the root container and all the PDBs excluding PDB\$SEED

Answer: B

Explanation: Backing Up a Whole CDB

Backing up a whole CDB is similar to backing up a non-CDB. When you back up a whole CDB, RMAN backs up the root, all the PDBs, and the archived redo logs. You can then recover either the whole CDB, the root only, or one or more PDBs from the CDB backup.

Note:

* You can back up and recover a whole CDB, the root only, or one or more PDBs.

* Backing Up Archived Redo Logs with RMAN

Archived redo logs are the key to successful media recovery. Back them up regularly. You can back up logs with BACKUP ARCHIVELOG, or back up logs while backing up datafiles and control files by specifying BACKUP ... PLUS ARCHIVELOG.

QUESTION NO: 20

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: 21

Which three functions are performed by the SQL Tuning Advisor?

- A. Building and implementing SQL profiles
- B. Recommending the optimization of materialized views
- C. Checking query objects for missing and stale statistics
- D. Recommending bitmap, function-based, and B-tree indexes
- E. Recommending the restructuring of SQL queries that are using bad plans

Answer: A,C,E

Explanation: The SQL Tuning Advisor takes one or more SQL statements as an input and invokes the Automatic Tuning Optimizer to perform SQL tuning on the statements. The output of the SQL Tuning Advisor is in the form of an advice or recommendations, along with a rationale for each recommendation and its expected benefit. The recommendation relates to collection of statistics on objects (C), creation of new indexes, restructuring of the SQL statement (E), or creation of a SQL profile (A). You can choose to accept the recommendation to complete the tuning of the SQL statements.

QUESTION NO: 22

Examine the following command:

```
ALTER SYSTEM SET enable_ddl_logging=FALSE;
```

Which statement is true?

- A. None of the data definition language (DDL) statements are logged in the trace file.
- B. Only DDL commands that resulted in errors are logged in the alert log file.
- C. A new log.xml file that contains the DDL statements is created, and the DDL command details are removed from the alert log file.
- D. Only DDL commands that resulted in the creation of new database files are logged.

Answer: A

Explanation: ENABLE_DDL_LOGGING enables or disables the writing of a subset of data definition language (DDL) statements to a DDL alert log.

The DDL log is a file that has the same format and basic behavior as the alert log, but it only contains the DDL statements issued by the database. The DDL log is created only for the RDBMS component and only if the ENABLE_DDL_LOGGING initialization parameter is set to true. When this parameter is set to false, DDL statements are not included in any log.

QUESTION NO: 23

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 RECOVER command, 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: 24

A new report process containing a complex query is written, with high impact on the database. You want to collect basic statistics about query, such as the level of parallelism, total database time, and the number of I/O requests.

For the database instance STATISTICS_LEVEL, the initialization parameter is set to TYPICAL and the CONTROL_MANAGEMENT_PACK_ACCESS parameter is set to DIAGNOSTIC+TUNING.

What should you do to accomplish this task?

- A. Execute the query and view Active Session History (ASH) for information about the query.
- B. Enable SQL trace for the query.
- C. Create a database operation, execute the query, and use the DBMS_SQL_MONITOR.REPORT_SQL_MONITOR function to view the report.
- D. Use the DBMS_APPLICATION_INFO.SET_SESSION_LONGOPS procedure to monitor query execution and view the information from the V\$SESSION_LONGOPS view.

Answer: C

Explanation: The REPORT_SQL_MONITOR function is used to return a SQL monitoring report for a specific SQL statement.

Incorrect:

Not A: Not interested in session statistics, only in statistics for the particular SQL query.

Not B: We are interested in statistics, not tracing.

Not D: SET_SESSION_LONGOPS Procedure

This procedure sets a row in the V\$SESSION_LONGOPS view. This is a view that is used to indicate the on-going progress of a long running operation. Some Oracle functions, such as parallel execution and Server Managed Recovery, use rows in this view to indicate the status of, for example, a database backup.

Applications may use the SET_SESSION_LONGOPS procedure to advertise information on the progress of application specific long running tasks so that the progress can be monitored by way of the V\$SESSION_LONGOPS view.

QUESTION NO: 25

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

Explanation:

QUESTION NO: 26

Your database supports a DSS workload that involves the execution of complex queries: Currently, the library cache contains the ideal workload for analysis. You want to analyze some of the queries for an application that are cached in the library cache.

What must you do to receive recommendations about the efficient use of indexes and materialized views to improve query performance?

- A. Create a SQL Tuning Set (STS) that contains the queries cached in the library cache and run the SQL Tuning Advisor (STA) on the workload captured in the STS.
- B. Run the Automatic Workload Repository Monitor (AWR).
- C. Create an STS that contains the queries cached in the library cache and run the SQL Performance Analyzer (SPA) on the workload captured in the STS.
- D. Create an STS that contains the queries cached in the library cache and run the SQL Access Advisor on the workload captured in the STS.

Answer: D

Explanation: * SQL Access Advisor is primarily responsible for making schema modification recommendations, such as adding or dropping indexes and materialized views. SQL Tuning Advisor makes other types of recommendations, such as creating SQL profiles and restructuring SQL statements.

* The query optimizer can also help you tune SQL statements. By using SQL Tuning Advisor and SQL Access Advisor, you can invoke the query optimizer in advisory mode to examine a SQL statement or set of statements and determine how to improve their efficiency. SQL Tuning Advisor and SQL Access Advisor can make various recommendations, such as creating SQL profiles, restructuring SQL statements, creating additional indexes or materialized views, and refreshing optimizer statistics.

Note:

* Decision support system (DSS) workload

* The library cache is a shared pool memory structure that stores executable SQL and PL/SQL code. This cache contains the shared SQL and PL/SQL areas and control structures such as locks and library cache handles.

Reference: Tuning SQL Statements

QUESTION NO: 27

The following parameter are set for your Oracle 12c database instance:

OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES=FALSE

OPTIMIZER_USE_SQL_PLAN_BASELINES=TRUE

You want to manage the SQL plan evolution task manually. Examine the following steps:

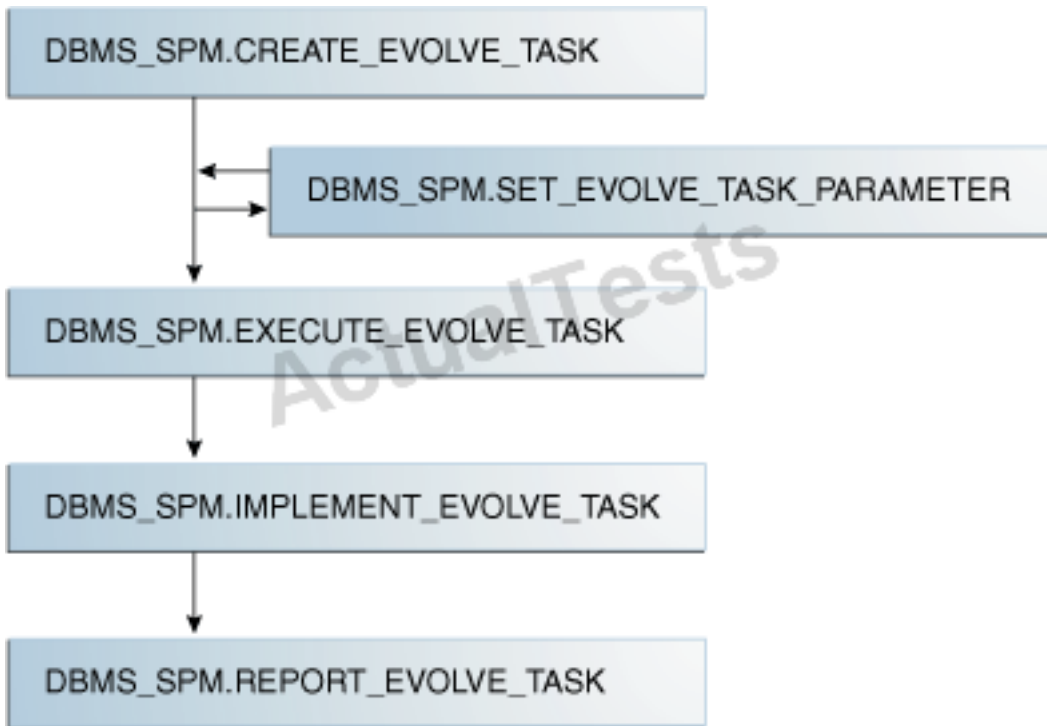
1. Set the evolve task parameters.
2. Create the evolve task by using the DBMS_SPM.CREATE_EVOLVE_TASK function.
3. Implement the recommendations in the task by using the DBMS_SPM.IMPLEMENT_EVOLVE_TASK function.
4. Execute the evolve task by using the DBMS_SPM.EXECUTE_EVOLVE_TASK function.
5. Report the task outcome by using the DBMS_SPM.REPORT_EVOLVE_TASK function.

Identify the correct sequence of steps:

- A. 2, 4, 5
- B. 2, 1, 4, 3, 5
- C. 1, 2, 3, 4, 5
- D. 1, 2, 4, 5

Answer: B

Explanation: * Evolving SQL Plan Baselines



*

2. Create the evolve task by using the `DBMS_SPM.CREATE_EVOLVE_TASK` function. This function creates an advisor task to prepare the plan evolution of one or more plans for a specified SQL statement. The input parameters can be a SQL handle, plan name or a list of plan names, time limit, task name, and description.

1. Set the evolve task parameters.

`SET_EVOLVE_TASK_PARAMETER`

This function updates the value of an evolve task parameter. In this release, the only valid parameter is `TIME_LIMIT`.

4. Execute the evolve task by using the `DBMS_SPM.EXECUTE_EVOLVE_TASK` function.

This function executes an evolution task. The input parameters can be the task name, execution name, and execution description. If not specified, the advisor generates the name, which is returned by the function.

3: `IMPLEMENT_EVOLVE_TASK`

This function implements all recommendations for an evolve task. Essentially, this function is equivalent to using `ACCEPT_SQL_PLAN_BASELINE` for all recommended plans. Input parameters include task name, plan name, owner name, and execution name.

5. Report the task outcome by using the `DBMS_SPM_EVOLVE_TASK` function.

This function displays the results of an evolve task as a CLOB. Input parameters include the task name and section of the report to include.

Reference: Oracle Database SQL Tuning Guide 12c, Managing SQL Plan Baselines

QUESTION NO: 28

In a recent Automatic Workload Repository (AWR) report for your database, you notice a high number of buffer busy waits. The database consists of locally managed tablespaces with free list managed segments.

On further investigation, you find that buffer busy waits is caused by contention on data blocks.

Which option would you consider first to decrease the wait event immediately?

- A. Decreasing PCTUSED
- B. Decreasing PCTFREE
- C. Increasing the number of DBWN process
- D. Using Automatic Segment Space Management (ASSM)
- E. Increasing db_buffer_cache based on the V\$DB_CACHE_ADVICE recommendation

Answer: D

Explanation: * Automatic segment space management (ASSM) is a simpler and more efficient way of managing space within a segment. It completely eliminates any need to specify and tune the *pctused*, *freelists*, and *freelist groups* storage parameters for schema objects created in the tablespace. If any of these attributes are specified, they are ignored.

* Oracle introduced Automatic Segment Storage Management (ASSM) as a replacement for traditional freelists management which used one-way linked-lists to manage free blocks with tables and indexes. ASSM is commonly called "bitmap freelists" because that is how Oracle implement the internal data structures for free block management.

Note:

* Buffer busy waits are most commonly associated with segment header contention outside the data buffer pool (db_cache_size, etc.).

* The most common remedies for high buffer busy waits include database writer (DBWR) contention tuning, adding freelists (or ASSM), and adding missing indexes.

QUESTION NO: 29

Examine this command:

```
SQL > exec DBMS_STATS.SET_TABLE_PREFS ('SH', 'CUSTOMERS', 'PUBLISH', 'false');
```

Which three statements are true about the effect of this command?

- A. Statistics collection is not done for the CUSTOMERS table when schema stats are gathered.
- B. Statistics collection is not done for the CUSTOMERS table when database stats are gathered.
- C. Any existing statistics for the CUSTOMERS table are still available to the optimizer at parse time.
- D. Statistics gathered on the CUSTOMERS table when schema stats are gathered are stored as pending statistics.
- E. Statistics gathered on the CUSTOMERS table when database stats are gathered are stored as pending statistics.

Answer: C,D,E

Explanation: * SET_TABLE_PREFS Procedure

This procedure is used to set the statistics preferences of the specified table in the specified schema.

* Example:

Using Pending Statistics

Assume many modifications have been made to the employees table since the last time statistics were gathered. To ensure that the cost-based optimizer is still picking the best plan, statistics should be gathered once again; however, the user is concerned that new statistics will cause the optimizer to choose bad plans when the current ones are acceptable. The user can do the following:

```
EXEC DBMS_STATS.SET_TABLE_PREFS('hr', 'employees', 'PUBLISH', 'false');
```

By setting the employees tables publish preference to FALSE, any statistics gather from now on will not be automatically published. The newly gathered statistics will be marked as pending.

QUESTION NO: 30

Examine the following impdp command to import a database over the network from a pre-12c Oracle database (source):

```
$> impdp <user_name> full=Y network link=hrdb_test transportable=always
transport_datafiles=
  '/u01/app/oracle/oradata/hrdb/sales01.dbf',
  '/u01/app/oracle/oradata/hrdb/cust01.dbf',
  '/u01/app/oracle/oradata/hrdb/emp01.dbf'
version=12 logfile=import.log
```


Which three are prerequisites for successful execution of the command?

- A. The import operation must be performed by a user on the target database with the DATAPUMP_IMP_FULL_DATABASE role, and the database link must connect to a user on the source database with the DATAPUMP_EXD_FULL_DATABASE role.
- B. All the user-defined tablespaces must be in read-only mode on the source database.
- C. The export dump file must be created before starting the import on the target database.
- D. The source and target database must be running on the same platform with the same endianness.
- E. The path of data files on the target database must be the same as that on the source database.
- F. The impdp operation must be performed by the same user that performed the expdp operation.

Answer: A,B,D

Explanation: In this case we have run the impdp without performing any conversion if endian format is different then we have to first perform conversion.

QUESTION NO: 31

Which two are true concerning a multitenant container database with three pluggable database?

- A. All administration tasks must be done to a specific pluggable database.
- B. The pluggable databases increase patching time.
- C. The pluggable databases reduce administration effort.
- D. The pluggable databases are patched together.
- E. Pluggable databases are only used for database consolidation.

Answer: C,E

Explanation: The benefits of Oracle Multitenant are brought by implementing a pure deployment choice. The following list calls out the most compelling examples.

* High consolidation density. (E)

The many pluggable databases in a single multitenant container database share its memory and background processes, letting you operate many more pluggable databases on a particular platform than you can single databases that use the old architecture. This is the same benefit that schema-based consolidation brings.

* Rapid provisioning and cloning using SQL.

* New paradigms for rapid patching and upgrades. (D, not B)

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.

* (C, not A) Manage many databases as one.

By consolidating existing databases as pluggable databases, administrators can manage many databases as one. For example, tasks like backup and disaster recovery are performed at the multitenant container database level.

* Dynamic between pluggable database resource management. In Oracle Database 12c, Resource Manager is extended with specific functionality to control the competition for resources between the pluggable databases within a multitenant container database.

Note:

* 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 multitenant 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.

Reference: 12c Oracle Multitenant

QUESTION NO: 32

Examine the current value for the following parameters in your database instance:

SGA_MAX_SIZE = 1024M

SGA_TARGET = 700M

DB_8K_CACHE_SIZE = 124M

LOG_BUFFER = 200M

You issue the following command to increase the value of DB_8K_CACHE_SIZE:

```
SQL> ALTER SYSTEM SET DB_8K_CACHE_SIZE=140M;
```

Which statement is true?

- A. It fails because the DB_8K_CACHE_SIZE parameter cannot be changed dynamically.
- B. It succeeds only if memory is available from the autotuned components if SGA.
- C. It fails because an increase in DB_8K_CACHE_SIZE cannot be accommodated within SGA_TARGET.
- D. It fails because an increase in DB_8K_CACHE_SIZE cannot be accommodated within SGA_MAX_SIZE.

Answer: D

Explanation: * The SGA_TARGET parameter can be dynamically increased up to the value specified for the SGA_MAX_SIZE parameter, and it can also be reduced.

* Example:

For example, suppose you have an environment with the following configuration:

SGA_MAX_SIZE = 1024M

SGA_TARGET = 512M

DB_8K_CACHE_SIZE = 128M

In this example, the value of SGA_TARGET can be resized up to 1024M and can also be reduced until one or more of the automatically sized components reaches its minimum size. The exact value depends on environmental factors such as the number of CPUs on the system. However, the value of DB_8K_CACHE_SIZE remains fixed at all times at 128M

* DB_8K_CACHE_SIZE

Size of cache for 8K buffers

* For example, consider this configuration:

SGA_TARGET = 512M

DB_8K_CACHE_SIZE = 128M

In this example, increasing DB_8K_CACHE_SIZE by 16 M to 144M means that the 16M is taken away from the automatically sized components. Likewise, reducing DB_8K_CACHE_SIZE by 16M to 112M means that the 16M is given to the automatically sized components.

QUESTION NO: 33

Which three statements are true concerning unplugging a pluggable database (PDB)?

- A. The PDB must be open in read only mode.
- B. The PDB must be closed.
- C. The unplugged PDB becomes a non-CDB.
- D. The unplugged PDB can be plugged into the same multitenant container database (CDB)
- E. The unplugged PDB can be plugged into another CDB.
- F. The PDB data files are automatically removed from disk.

Answer: A,D,E

Explanation:

D: An unplugged PDB contains data dictionary tables, and some of the columns in these encode information in an endianness-sensitive way. There is no supported way to handle the conversion of such columns automatically. This means, quite simply, that an unplugged PDB cannot be moved across an endianness difference.

E (not F): To exploit the new unplug/plug paradigm for patching the Oracle version most effectively, the source and destination CDBs should share a filesystem so that the PDB's datafiles can remain in place.

Reference: Oracle White Paper, Oracle Multitenant

QUESTION NO: 34

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: 35

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: 36

On your Oracle 12c database, you invoked SQL *Loader to load data into the EMPLOYEES table in the HR schema by issuing the following command:

```
$> sqlldr hr/hr@pdb table=employees
```

Which two statements are true regarding the command?

- A. It succeeds with default settings if the EMPLOYEES table belonging to HR is already defined in the database.
- B. It fails because no SQL *Loader data file location is specified.
- C. It fails if the HR user does not have the CREATE ANY DIRECTORY privilege.
- D. It fails because no SQL *Loader control file location is specified.

Answer: A,C

Explanation:

Note:

* SQL *Loader is invoked when you specify the sqlldr command and, optionally, parameters that establish session characteristics.

QUESTION NO: 37

After implementing full Oracle Data Redaction, you change the default value for the NUMBER data type as follows:

```
SQL> SELECT NUMBER_VALUE FROM REDACTION_VALUES_FOR_TYPE_FULL;
NUMBER_VALUE
-----
0

SQL> EXEC DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES(-1)
PL/SQL procedure successfully completed.

SQL> select number_value from redaction_values_for_type_full;
NUMBER_VALUE
-----
-1
```

After changing the value, you notice that FULL redaction continues to redact numeric data with zero.

What must you do to activate the new default value for numeric full redaction?

- A. Re-enable redaction policies that use FULL data redaction.
- B. Re-create redaction policies that use FULL data redaction.
- C. Re-connect the sessions that access objects with redaction policies defined on them.
- D. Flush the shared pool.
- E. Restart the database instance.

Answer: E

Explanation: About Altering the Default Full Data Redaction Value

You can alter the default displayed values for full Data Redaction policies. By default, 0 is the redacted value when Oracle Database performs full redaction (DBMS_REDACT.FULL) on a column of the NUMBER data type. If you want to change it to another value (for example, 7), then you can run the DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES procedure to modify this value. The modification applies to all of the Data Redaction policies in the current database instance. After you modify a value, you must restart the database for it to take effect.

Note:

* 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.

* UPDATE_FULL_REDACTION_VALUES Procedure

This procedure modifies the default displayed values for a Data Redaction policy for full redaction.

* After you create the Data Redaction policy, it is automatically enabled and ready to redact data.

* Oracle Data Redaction enables you to mask (redact) data that is returned from queries issued by

low-privileged users or applications. You can redact column data by using one of the following methods:

- / Full redaction.
- / Partial redaction.
- / Regular expressions.
- / Random redaction.
- / No redaction.

Reference: Oracle Database Advanced Security Guide 12c, About Altering the Default Full Data Redaction Value

QUESTION NO: 38

You must track all transactions that modify certain tables in the sales schema for at least three years.

Automatic undo management is enabled for the database with a retention of one day.

Which two must you do to track the transactions?

- A. Enable supplemental logging for the database.
- B. Specify undo retention guarantee for the database.
- C. Create a Flashback Data Archive in the tablespace where the tables are stored.
- D. Create a Flashback Data Archive in any suitable tablespace.
- E. Enable Flashback Data Archiving for the tables that require tracking.

Answer: D,E

Explanation: E: By default, flashback archiving is disabled for any table. You can enable flashback archiving for a table if you have the FLASHBACK ARCHIVE object privilege on the Flashback Data Archive that you want to use for that table.

D: Creating a Flashback Data Archive

/ Create a Flashback Data Archive with the CREATE FLASHBACK ARCHIVE statement, specifying the following:

Name of the Flashback Data Archive

Name of the first tablespace of the Flashback Data Archive

(Optional) Maximum amount of space that the Flashback Data Archive can use in the first tablespace

/ Create a Flashback Data Archive named fla2 that uses tablespace tbs2, whose data will be retained for two years:

```
CREATE FLASHBACK ARCHIVE fla2 TABLESPACE tbs2 RETENTION 2 YEAR;
```

QUESTION NO: 39

You are the DBA supporting an Oracle 11g Release 2 database and wish to move a table containing several DATE, CHAR, VARCHAR2, and NUMBER data types, and the table's indexes, to another tablespace.

The table does not have a primary key and is used by an OLTP application.

Which technique will move the table and indexes while maintaining the highest level of availability to the application?

- A. Oracle Data Pump.
- B. An ALTER TABLE MOVE to move the table and ALTER INDEX REBUILD to move the indexes.
- C. An ALTER TABLE MOVE to move the table and ALTER INDEX REBUILD ONLINE to move the indexes.
- D. Online Table Redefinition.
- E. Edition-Based Table Redefinition.

Answer: D

Explanation: * Oracle Database provides a mechanism to make table structure modifications without significantly affecting the availability of the table. The mechanism is called online table redefinition. Redefining tables online provides a substantial increase in availability compared to traditional methods of redefining tables.

* To redefine a table online:

Choose the redefinition method: by key or by rowid

* By key—Select a primary key or pseudo-primary key to use for the redefinition. Pseudo-primary keys are unique keys with all component columns having NOT NULL constraints. For this method, the versions of the tables before and after redefinition should have the same primary key columns.

This is the preferred and default method of redefinition.

* By rowid—Use this method if no key is available. In this method, a hidden column named M_ROW\$\$ is added to the post-redefined version of the table. It is recommended that this column be dropped or marked as unused after the redefinition is complete. If COMPATIBLE is set to 10.2.0 or higher, the final phase of redefinition automatically sets this column unused. You can then use the ALTER TABLE ... DROP UNUSED COLUMNS statement to drop it.

You cannot use this method on index-organized tables.

Note:

* When you rebuild an index, you use an existing index as the data source. Creating an index in this manner enables you to change storage characteristics or move to a new tablespace. Rebuilding an index based on an existing data source removes intra-block fragmentation. Compared to dropping the index and using the CREATE INDEX statement, re-creating an existing index offers better performance.

Incorrect:

Not E: Edition-based redefinition enables you to upgrade the database component of an application while it is in use, thereby minimizing or eliminating down time.

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

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: 42

Identify 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: B,E

Explanation: 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 NO: 43

You upgrade your Oracle database in a multiprocessor environment. As a recommended you execute the following script:

```
SQL > @utlrp.sql
```

Which two actions does the script perform?

- A. Parallel compilation of only the stored PL/SQL code
- B. Sequential recompilation of only the stored PL/SQL code
- C. Parallel recompilation of any stored PL/SQL code
- D. Sequential recompilation of any stored PL/SQL code
- E. Parallel recompilation of Java code
- F. Sequential recompilation of Java code

Answer: C,E

Explanation: utlrp.sql and utlprp.sql

The 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.

Reference: Recompiling Invalid Schema Objects

QUESTION NO: 44

Which two 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: C,D

Explanation: http://docs.oracle.com/database/121/ADMIN/cdb_plug.htm#ADMIN13658

When dropping a PDB, you can either keep or delete the PDB's data files by using one of the following clauses:

KEEP DATAFILES, the default, retains the data files.

The PDB's temp file is removed even when KEEP DATAFILES is specified because the temp file is no longer needed.

INCLUDING DATAFILES removes the data files from disk.

If a PDB was created with the SNAPSHOT COPY clause, then you must specify INCLUDING DATAFILES when you drop the PDB.

The following prerequisites must be met:

The PDB must be in mounted mode, or it must be unplugged.

QUESTION NO: 45

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: A,C,D

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 NO: 46

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: 47

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 awrddrpt.sql report is the Automated Workload Repository Compare Period Report. The awrddrpt.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 NO: 48

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.created_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 NO: 49

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

Explanation:

QUESTION NO: 50

You upgraded your database from pre-12c to a multitenant container database (CDB) containing pluggable databases (PDBs).

Examine the query and its output:

```
SQL> SELECT * FROM v$PFILE_users;
-----
USERNAME                               SYSDB SYSOP  SYSAS  SYSBA  SYSDG  SYSKM  CON_ID
-----
SYS                                     TRUE  TRUE   FALSE FALSE  FALSE FALSE   0
```

Which two tasks must you perform to add users with SYSBACKUP, SYSDG, and SYSKM privilege to the password file?

- A. Assign the appropriate operating system groups to SYSBACKUP, SYSDG, SYSKM.
- B. Grant SYSBACKUP, SYSDG, and SYSKM privileges to the intended users.
- C. Re-create the password file with SYSBACKUP, SYSDG, and SYSKM privilege and the FORCE argument set to No.
- D. Re-create the password file with SYSBACKUP, SYSDG, and SYSKM privilege, and FORCE arguments set to Yes.
- E. Re-create the password file in the Oracle Database 12c format.

Answer: B,D

Explanation:

* orapwd

/ You can create a database password file using the password file creation utility, ORAPWD.

The syntax of the ORAPWD command is as follows:

```
orapwd FILE=filename [ENTRIES=numusers] [FORCE={y|n}] [ASM={y|n}]
[DBUNIQUENAME=dbname] [FORMAT={12|legacy}] [SYSBACKUP={y|n}] [SYSDG={y|n}]
[SYSKM={y|n}] [DELETE={y|n}] [INPUT_FILE=input-fname]
```

force - whether to overwrite existing file (optional),

* v\$PWFILe_users

/ 12c: V\$PWFILe_USERS lists all users in the password file, and indicates whether the user has been granted the SYSDBA, SYSOPER, SYSASM, SYSBACKUP, SYSDG, and SYSKM privileges.

/ 10c: sts users who have been granted SYSDBA and SYSOPER privileges as derived from the password file.

ColumnDatatypeDescription

USERNAMEVARCHAR2(30)The name of the user that is contained in the password file

SYSDBAVARCHAR2(5)If TRUE, the user can connect with SYSDBA privileges

SYSOPeR VARCHAR2(5)If TRUE, the user can connect with SYSOPeR privileges

Incorrect:

not E: The format of the v\$PWFILe_users file is already in 12c format.

QUESTION NO: 51

An application accesses a small lookup table frequently. You notice that the required data blocks are getting aged out of the default buffer cache.

How would you guarantee that the blocks for the table never age out?

- A. Configure the KEEP buffer pool and alter the table with the corresponding storage clause.
- B. Increase the database buffer cache size.
- C. Configure the RECYCLE buffer pool and alter the table with the corresponding storage clause.
- D. Configure Automata Shared Memory Management.
- E. Configure Automatic Memory Management-

Answer: A

Explanation: Schema objects are referenced with varying usage patterns; therefore, their cache behavior may be quite different. Multiple buffer pools enable you to address these differences. You can use a KEEP buffer pool to maintain objects in the buffer cache and a RECYCLE buffer pool to prevent objects from consuming unnecessary space in the cache. When an object is allocated to a cache, all blocks from that object are placed in that cache. Oracle maintains a DEFAULT buffer pool for objects that have not been assigned to one of the buffer pools.

QUESTION NO: 52

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 after 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 back.
- 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: 53

You 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_NAME_CONVERT initialization parameter.
- E. Specify all files in the CREATE DATABASE statement without using Oracle managed Files

(OMF).

Answer: A,B,C

Explanation: 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 NO: 54

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

Explanation:

QUESTION NO: 55

Your 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 replay
- B. SQL Tuning Advisor
- C. SQL Access Advisor
- D. SQL Performance Analyzer
- E. Automatic Workload Repository compare reports

Answer: D

Explanation: 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 profiles

http://docs.oracle.com/cd/B28359_01/server.111/b28318/intro.htm#CNCPT961

QUESTION NO: 56

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 connect 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: 57

Examine the following query output:

```
SQL> SELECT name, force_logging FROM v$databases;
NAME          FORCE_LOGGING
-----
PROD          NO
```

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: Y
```

Which 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: C

Explanation:

Note from Oracle Documentation:

With redo logging disabled, the disk space required for redo logs during an Oracle Data Pump import will be smaller. However, to ensure recovery from media failure, the DBA should do an

RMAN backup after the import completes.

Even with this parameter specified, there is still redo logging for other operations of Oracle Data Pump. This includes all CREATE and ALTER statements, except CREATE INDEX, and all operations against the master table used by Oracle Data Pump during the import.

QUESTION NO: 58

You 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: A,B,D

Explanation: Keyword: shows the difference.

* Full ADDM analysis across two AWR snapshot periods

Detects causes, measure effects, then correlates them

Causes: workload changes, configuration changes

Effects: 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.

Reference: Resolving Performance Degradation Over Time

QUESTION NO: 59

Examine the parameter for your database instance:

NAME	TYPE	VALUE
optimizer_adaptive_reporting_only	boolean	FALSE
optimizer_capture_sql_plan_baselines	boolean	FALSE
optimizer_dynamic_sampling	integer	2
optimizer_features_enable	string	12.1.0.1

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:

```
SQL> SELECT product_name
FROM   order_items o, product_information p
WHERE  o.unit_price < 15
AND    quantity > 1
AND    p.product_id = o.product_id;

30 rows selected.
```

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: B

Explanation: * optimizer_dynamic_sampling

OPTIMIZER_DYNAMIC_SAMPLING controls both when the database gathers dynamic statistics, and the size of the sample that the optimizer uses to gather the statistics.

Range of values 0 to 11

QUESTION NO: 60

Which 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.
- 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: A,D,E

Explanation: 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.

Reference: SQL Plan Management with Oracle Database 12c

QUESTION NO: 61

You create a new pluggable database, HR_PDB, from the seed database. Which three tablespaces are created by default in HR_PDB?

- A. SYSTEM
- B. SYSAUX
- C. EXAMPLE
- D. UNDO
- E. TEMP
- F. USERS

Answer: A,B,E

Explanation: * A PDB would have its SYSTEM, SYSAUX, TEMP tablespaces. It can also contain other user-created tablespaces in it.

*

* Oracle Database creates both the SYSTEM and SYSAUX tablespaces as part of every database.

* tablespace_datafile_clauses

Use 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@pdb1
```

```
SELECT tablespace_name FROM dba_tablespaces;
```

```
TABLESPACE_NAME
```

```
-----
```

```
SYSTEM  
SYSAUX  
TEMP
```

USERS

SQL>

QUESTION NO: 62

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: 63

You 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: C

Explanation: 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 NO: 64

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: 65

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

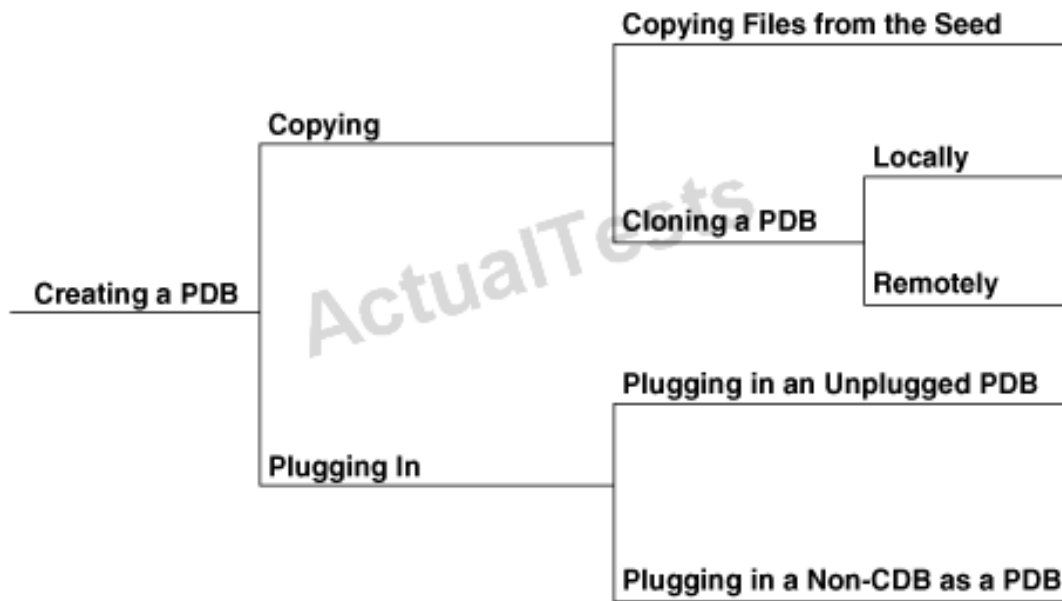
plugged in.

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:



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: 66

What is the effect of specifying the "ENABLE PLUGGABLE DATABASE" clause in a "CREATE DATABASE" statement?

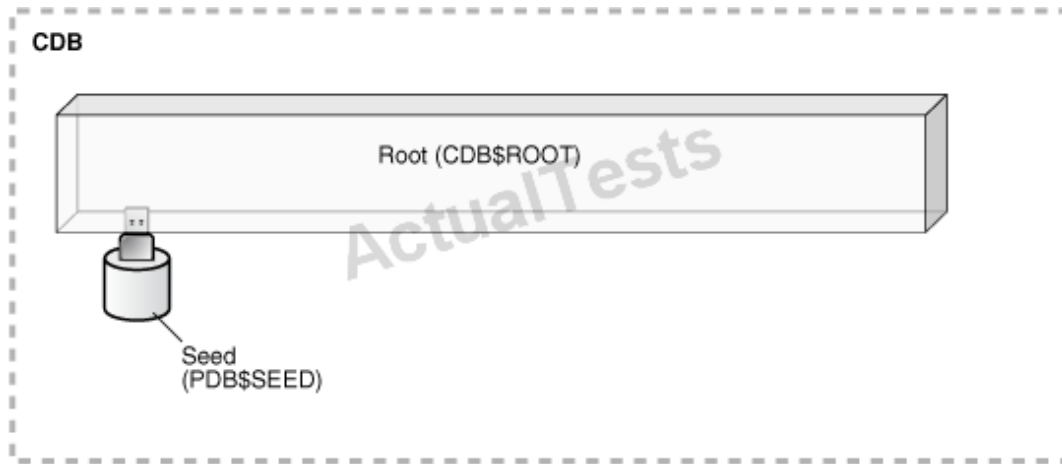
- A. It will create a multitenant container database (CDB) with only the root opened.
- B. It will create a CDB with root opened and seed read only.
- C. It will create a CDB with root and seed opened and one PDB mounted.
- D. It will create a CDB that must be plugged into an existing CDB.
- E. It will 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 NO: 67

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_ZISE = 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 NO: 68

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: 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: 69

Examine the contents of SQL loader control file:

```

LOAD DATA
INFILE myfile1.dat
INFILE myfile2.dat
FIELD NAMES FIRST FILE
APPEND
INTO TABLE EMP
FIELDS CSV WITH EMBEDDED
DATE FORMAT "DD-Month-YYYY"
(empno,
ename,
job,
mgr,
hiredate DATE,
sal,
comm,
deptno,
entrydate DATE)

```

Which three statements are true regarding the SQL* Loader operation performed using the control file?

- A. An EMP table is created if a table does not exist. Otherwise, if the EMP table is appended with the loaded data.
- B. The SQL* Loader data file myfile1.dat has the column names for the EMP table.
- C. The SQL* Loader operation fails because no record terminators are specified.
- D. Field names should be the first line in the both the SQL* Loader data files.
- E. The SQL* Loader operation assumes that the file must be a stream record format file with the normal carriage return string as the record terminator.

Answer: A,B,E

Explanation: A: The APPEND keyword tells SQL*Loader to preserve any preexisting data in the table. Other options allow you to delete preexisting data, or to fail with an error if the table is not empty to begin with.

B (not D):

Note:

* SQL*Loader-00210: first data file is empty, cannot process the FIELD NAMES record

Cause: The data file listed in the next message was empty. Therefore, the FIELD NAMES FIRST FILE directive could not be processed.

Action: Check the listed data file and fix it. Then retry the operation

E:

* A comma-separated values (CSV) (also sometimes called character-separated values, because the separator character does not have to be a comma) file stores tabular data (numbers and text) in plain-text form. Plain text means that the file is a sequence of characters, with no data that has to be interpreted instead, as binary numbers. A CSV file consists of any number of records, separated by line breaks of some kind; each record consists of fields, separated by some other character or string, most commonly a literal comma or tab. Usually, all records have an identical sequence of fields.

* Fields with embedded commas must be quoted.

Example:

```
1997,Ford,E350,"Super, luxurious truck"
```

Note:

* SQL*Loader is a bulk loader utility used for moving data from external files into the Oracle database.

QUESTION NO: 70

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 executes 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: 71

Which two statements are true concerning the Resource Manager plans for individual pluggable databases (PDB plans) in a multitenant container database (CDB)?

- A. If no PDB plan is enabled for a pluggable database, then all sessions for that PDB are treated to an equal degree of the resource share of that PDB.
- B. In a PDB plan, subplans may be used with up to eight consumer groups.
- C. If a PDB plan is enabled for a pluggable database, then resources are allocated to consumer groups across all PDBs in the CDB.
- D. If no PDB plan is enabled for a pluggable database, then the PDB share in the CDB plan is dynamically calculated.
- E. If a PDB plan is enabled for a pluggable database, then resources are allocated to consumer groups based on the shares provided to the PDB in the CDB plan and the shares provided to the consumer groups in the PDB plan.

Answer: A,E

Explanation: A: Setting a PDB resource plan is optional. If not specified, all sessions within the PDB are treated equally.

*

In a non-CDB database, workloads within a database are managed with resource plans. In a PDB, workloads are also managed with resource plans, also called PDB resource plans. The functionality is similar except for the following differences:

/ Non-CDB Database

Multi-level resource plans
Up to 32 consumer groups
Subplans
/ PDB Database
Single-level resource plans only
Up to 8 consumer groups
(not B) No subplans

QUESTION NO: 72

You use a recovery catalog for maintaining your database backups.

You execute the following command:

```
$rman TARGET / CATALOG rman / cat@catdb
```

```
RMAN > BACKUP VALIDATE DATABASE ARCHIVELOG ALL;
```

Which two statements are true?

- A. Corrupted blocks, if any, are repaired.
- B. Checks are performed for physical corruptions.
- C. Checks are performed for logical corruptions.
- D. Checks are performed to confirm whether all database files exist in correct locations
- E. Backup sets containing both data files and archive logs are created.

Answer: B,D

Explanation: B (not C): You can validate that all database files and archived redo logs can be backed up by running a command as follows:

```
RMAN> BACKUP VALIDATE DATABASE ARCHIVELOG ALL;
```

This form of the command would check for physical corruption. To check for logical corruption, RMAN> BACKUP VALIDATE CHECK LOGICAL DATABASE ARCHIVELOG ALL;

D: You can use the VALIDATE keyword of the BACKUP command to do the following:

Check datafiles for physical and logical corruption

Confirm that all database files exist and are in the correct locations.

Note:

You can use the VALIDATE option of the BACKUP command to verify that database files exist and are in the correct locations (D), and have no physical or logical corruptions that would prevent RMAN from creating backups of them. When performing a BACKUP...VALIDATE, RMAN reads the files to be backed up in their entirety, as it would during a real backup. It does not, however, actually produce any backup sets or image copies (Not A, not E).

QUESTION NO: 73

Which three statements are true concerning the multitenant architecture?

- A. Each pluggable database (PDB) has its own set of background processes.
- B. A PDB can have a private temp tablespace.
- C. PDBs can share the sysaux tablespace.
- D. Log switches occur only at the multitenant container database (CDB) level.
- E. Different PDBs can have different default block sizes.
- F. PDBs share a common system tablespace.
- G. Instance recovery is always performed at the CDB level.

Answer: B,D,G

Explanation: B:

* A PDB would have its SYSTEM, SYSAUX, TEMP tablespaces. It can also contain other user created tablespaces in it.

* There is one default temporary tablespace for the entire CDB. However, you can create additional temporary tablespaces in individual PDBs.

D:

* There is a single redo log and a single control file for an entire CDB

* A log switch is the point at which the database stops writing to one redo log file and begins writing to another. Normally, a log switch occurs when the current redo log file is completely filled and writing must continue to the next redo log file.

G: instance recovery

The automatic application of redo log records to uncommitted data blocks when a database instance is restarted after a failure.

Incorrect:

Not A:

- * There is one set of background processes shared by the root and all PDBs. –
- * High consolidation density. The many pluggable databases in a single container database share its memory and background processes, letting you operate many more pluggable databases on a particular platform than you can single databases that use the old architecture.

Not C: There is a separate SYSAUX tablespace for the root and for each PDB.

Not F: There is a separate SYSTEM tablespace for the root and for each PDB. -

QUESTION NO: 74

You notice that the elapsed time for an important database scheduler Job is unacceptably long.

The job belongs to a scheduler job class and window.

Which two actions would reduce the job's elapsed time?

- A.** Increasing the priority of the job class to which the job belongs
- B.** Increasing the job's relative priority within the Job class to which it belongs
- C.** Increasing the resource allocation for the consumer group mapped to the scheduler job's job class within the plan mapped to the scheduler window
- D.** Moving the job to an existing higher priority scheduler window with the same schedule and duration
- E.** Increasing the value of the JOB_QUEUE_PROCESSES parameter
- F.** Increasing the priority of the scheduler window to which the job belongs

Answer: C,F

Explanation: http://www.dba-oracle.com/job_scheduling/job_classes.htm<http://docs.oracle.com/database/121/ADMIN/schedover.htm#i1106396>

QUESTION NO: 75

You plan to migrate your database from a File system to Automata Storage Management (ASM) on same platform.

Which two methods or commands would you use to accomplish this task?

- A. RMAN CONVERT command
- B. Data Pump Export and import
- C. Conventional Export and Import
- D. The BACKUP AS COPY DATABASE . . . command of RMAN
- E. DBMS_FILE_TRANSFER with transportable tablespace

Answer: A,D

Explanation: A:

1. Get the list of all datafiles.

Note: RMAN Backup of ASM Storage

There is often a need to move the files from the file system to the ASM storage and vice versa. This may come in handy when one of the file systems is corrupted by some means and then the file may need to be moved to the other file system.

D: Migrating a Database into ASM

* To take advantage of Automatic Storage Management with an existing database you must migrate that database into ASM. This migration is performed using Recovery Manager (RMAN) even if you are not using RMAN for your primary backup and recovery strategy.

* Example:

Back up your database files as copies to the ASM disk group.

```
BACKUP AS COPY INCREMENTAL LEVEL 0 DATABASE  
FORMAT '+DISK' TAG 'ORA_ASM_MIGRATION';
```

Reference: Migrating Databases To and From ASM with Recovery Manager

QUESTION NO: 76

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: 77

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: 78

Your database is open and the LISTENER listener running. You stopped the wrong listener LISTENER by issuing the following command:

```
1snrctl > STOP
```

What happens to the sessions that are presently connected to the database Instance?

- A. They are able to perform only queries.
- B. They are not affected and continue to function normally.
- C. They are terminated and the active transactions are rolled back.
- D. They are not allowed to perform any operations until the listener LISTENER is started.

Answer: B

Explanation: The listener is used when the connection is established. The immediate impact of stopping the listener will be that no new session can be established from a remote host. Existing sessions are not compromised.

QUESTION NO: 79

Which three statements are true about using flashback database in a multitenant container database (CDB)?

- A. The root container can be flashed back without flashing back the pluggable databases (PDBs).
- B. To enable flashback database, the CDB must be mounted.
- C. Individual PDBs can be flashed back without flashing back the entire CDB.
- D. The DB_FLASHBACK_RETENTION_TARGET parameter must be set to enable flashback of the CDB.
- E. A CDB can be flashed back specifying the desired target point in time or an SCN, but not a restore point.

Answer: B,D,E

Explanation:

QUESTION NO: 80

You execute the following PL/SQL:

```
BEGIN
DBMS_FGA.add_policy(
object_schema => 'JIM',
object_name => 'PRODUCTS',
policy_name => 'PROD_AUDIT',
audit_condition => 'PRICE > 10000',
audit_column => 'PRICE');
END;
/
```

Which two statements are true?

- A. Fine-Grained Auditing (FGA) is enabled for the PRICE column in the PRODUCTS table for SELECT statements only when a row with PRICE > 10000 is accessed.

- B.** FGA is enabled for the PRODUCTS.PRICE column and an audit record is written whenever a row with PRICE > 10000 is accessed.
- C.** FGA is enabled for all DML operations by JIM on the PRODUCTS.PRICE column.
- D.** FGA is enabled for the PRICE column of the PRODUCTS table and the SQL statements is captured in the FGA audit trail.

Answer: A,D

Explanation: DBMS_FGA.ADD_POLICY(object_schema => 'hr',object_name => 'emp',policy_name => 'chk_hr_emp',audit_condition => 'dept = "SALES" ',audit_column => 'salary',statement_types => 'insert,update,delete,select');

Default value for statement_types is SELECT

Setting audit_trail to DBMS_FGA.DB sends the audit trail to the SYS.FGA_LOG\$ table in the database and omits SQL Text and SQL Bind.

Setting audit_trail to DBMS_FGA.DB+EXTENDED sends the audit trail to the SYS.FGA_LOG\$ table in the database and includes SQL Text and SQL Bind.

Setting audit_trail to DBMS_FGA.XML writes the audit trail in XML files sent to the operating system and omits SQL Text and SQL Bind.

Setting audit_trail to DBMS_FGA.XML+EXTENDED writes the audit trail in XML files sent to the operating system and includes SQL Text and SQL Bind.

Default value fo audit_trail parameter in DBMS_FGA.ADD_POLICY is DB+EXTENDED

QUESTION NO: 81

You execute the following commands to audit database activities:

```
SQL > ALTER SYSTEM SET AUDIT_TRIAL=DB, EXTENDED SCOPE=SPFILE;
```

```
SQL > AUDIT SELECT TABLE, INSERT TABLE, DELETE TABLE BY JOHN By SESSION
WHENEVER SUCCESSFUL;
```

Which statement is true about the audit record that generated when auditing after instance restarts?

- A.** One audit record is created for every successful execution of a SELECT, INSERT OR DELETE command on a table, and contains the SQL text for the SQL Statements.
- B.** One audit record is created for every successful execution of a SELECT, INSERT OR DELETE command, and contains the execution plan for the SQL statements.
- C.** One audit record is created for the whole session if john successfully executes a SELECT, INSERT, or DELETE command, and contains the execution plan for the SQL statements.

D. One audit record is created for the whole session if JOHN successfully executes a select command, and contains the SQL text and bind variables used.

E. One audit record is created for the whole session if john successfully executes a SELECT, INSERT, or DELETE command on a table, and contains the execution plan, SQL text, and bind variables used.

Answer: D

Explanation:

BY SESSION means:For any type of audit (schema object, statement, or privilege), BY SESSION inserts only one audit record in the audit trail, for each user and schema object, during the session that includes an audited action.

AUDIT_TRAIL=db, extended meansPerforms all actions of AUDIT_TRAIL=db, and also populates the SQL bind and SQL text CLOB-type columns of the SYS.AUD\$ table, when available. These two columns are populated only when this parameter is specified.

QUESTION NO: 82

You support Oracle Database 12c Oracle Database 11g, and Oracle Database log on the same server.

All databases of all versions use Automatic Storage Management (ASM).

Which three statements are true about the ASM disk group compatibility attributes that are set for a disk group?

A. The ASM compatibility attribute controls the format of the disk group metadata.

B. RDBMS compatibility together with the database version determines whether a database Instance can mount the ASM disk group.

C. The RDBMS compatibility setting allows only databases set to the same version as the compatibility value, to mount the ASM disk group.

D. The ASM compatibility attribute determines some of the ASM features that may be used by the Oracle disk group.

E. The ADVN compatibility attribute determines the ACFS features that may be used by the Oracle 10 g database.

Answer: A,B,D

Explanation: AD: The value for the disk group COMPATIBLE.ASM attribute determines the minimum software version for an Oracle ASM instance that can use the disk group. This setting also affects the format of the data structures for the Oracle ASM metadata on the disk.

B: The value for the disk group COMPATIBLE.RDBMS attribute determines the minimum COMPATIBLE database initialization parameter setting for any database instance that is allowed to use the disk group. Before advancing the COMPATIBLE.RDBMS attribute, ensure that the values for the COMPATIBLE initialization parameter for all of the databases that access the disk group are set to at least the value of the new setting for COMPATIBLE.RDBMS.

For example, if the COMPATIBLE initialization parameters of the databases are set to either 11.1 or 11.2, then COMPATIBLE.RDBMS can be set to any value between 10.1 and 11.1 inclusively.

Not E:

/The value for the disk group COMPATIBLE.ADVM attribute determines whether the disk group can contain Oracle ASM volumes. The value must be set to 11.2 or higher. Before setting this attribute, the COMPATIBLE.ASM value must be 11.2 or higher. Also, the Oracle ADVM volume drivers must be loaded in the supported environment.

/ You can create an Oracle ASM Dynamic Volume Manager (Oracle ADVM) volume in a disk group. The volume device associated with the dynamic volume can then be used to host an Oracle ACFS file system.

The compatibility parameters COMPATIBLE.ASM and COMPATIBLE.ADVM must be set to 11.2 or higher for the disk group.

Note:

* The disk group attributes that determine compatibility are COMPATIBLE.ASM, COMPATIBLE.RDBMS. and COMPATIBLE.ADVM. The COMPATIBLE.ASM and COMPATIBLE.RDBMS attribute settings determine the minimum Oracle Database software version numbers that a system can use for Oracle ASM and the database instance types respectively. For example, if the Oracle ASM compatibility setting is 11.2, and RDBMS compatibility is set to 11.1, then the Oracle ASM software version must be at least 11.2, and the Oracle Database client software version must be at least 11.1. The COMPATIBLE.ADVM attribute determines whether the Oracle ASM Dynamic Volume Manager feature can create an volume in a disk group.

QUESTION NO: 83

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

Explanation:

QUESTION NO: 84

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: 85

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: 86

You administer an online transaction processing (OLTP) system whose database is stored in Automatic Storage Management (ASM) and whose disk group use normal redundancy.

One of the ASM disks goes offline, and is then dropped because it was not brought online before DISK_REPAIR_TIME elapsed.

When the disk is replaced and added back to the disk group, the ensuing rebalance operation is too slow.

Which two recommendations should you make to speed up the rebalance operation if this type of failure happens again?

- A. Increase the value of the ASM_POWER_LIMIT parameter.
- B. Set the DISK_REPAIR_TIME disk attribute to a lower value.
- C. Specify the statement that adds the disk back to the disk group.
- D. Increase the number of ASMB processes.
- E. Increase the number of DBWR_IO_SLAVES in the ASM instance.

Answer: A,C

Explanation: ASM_POWER_LIMIT specifies the maximum power on an Automatic Storage Management instance for disk rebalancing. The higher the limit, the faster rebalancing will complete. Lower values will take longer, but consume fewer processing and I/O resources.

Grouping operations in a single ALTER DISKGROUP statement can reduce rebalancing operations.

http://docs.oracle.com/cd/E11882_01/server.112/e18951/asmdiskgrps.htm#OSTMG10070

QUESTION NO: 87

You are administering a database and you receive a requirement to apply the following restrictions:

1. A connection must be terminated after four unsuccessful login attempts by user.
2. A user should not be able to create more than four simultaneous sessions.

3. User session must be terminated after 15 minutes of inactivity.
4. Users must be prompted to change their passwords every 15 days.

How would you accomplish these requirements?

- A. by granting a secure application role to the users
- B. by creating and assigning a profile to the users and setting the REMOTE_OS_AUTHENT parameter to FALSE
- C. By creating and assigning a profile to the users and setting the SEC_MAX_FAILED_LOGIN_ATTEMPTS parameter to 4
- D. By Implementing Fine-Grained Auditing (FGA) and setting the REMOTE_LOGIN_PASSWORD_FILE parameter to NONE.
- E. By implementing the database resource Manager plan and setting the SEC_MAX_FAILED_LOGIN_ATTEMPTS parameters to 4.

Answer: C

Explanation: SEC_MAX_FAILED_LOGIN_ATTEMPTS specifies the number of authentication attempts that can be made by a client on a connection to the server process. These login attempts can be for multiple user accounts in the same connection. After the specified number of failure attempts, the connection will be automatically dropped by the server process.

<http://docs.oracle.com/database/121/REFRN/refrn10274.htm#REFRN10274>

QUESTION NO: 88

A senior DBA asked you to execute the following command to improve performance:

```
SQL> ALTER TABLE subscribe log STORAGE (BUFFER_POOL recycle);
```

You checked the data in the SUBSCRIBE_LOG table and found that it is a large table containing one million rows.

What could be a reason for this recommendation?

- A. The keep pool is not configured.
- B. Automatic Workarea Management is not configured.
- C. Automatic Shared Memory Management is not enabled.
- D. The data blocks in the SUBSCRIBE_LOG table are rarely accessed.
- E. All the queries on the SUBSCRIBE_LOG table are rewritten to a materialized view.

Answer: D

Explanation: The most of the rows in SUBSCRIBE_LOG table are accessed once a week.

QUESTION NO: 89

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: 90

Which two partitioned table maintenance operations support asynchronous Global Index Maintenance in Oracle database 12c?

- A. ALTER TABLE SPLIT PARTITION
- B. ALTER TABLE MERGE PARTITION
- C. ALTER TABLE TRUNCATE PARTITION
- D. ALTER TABLE ADD PARTITION
- E. ALTER TABLE DROP PARTITION
- F. ALTER TABLE MOVE PARTITION

Answer: C,E

Explanation: Asynchronous Global Index Maintenance for DROP and TRUNCATE PARTITION
This feature enables global index maintenance to be delayed and decoupled from a DROP and TRUNCATE partition without making a global index unusable. Enhancements include faster DROP and TRUNCATE partition operations and the ability to delay index maintenance to off-peak time.

Reference: Oracle Database VLDB and Partitioning Guide 12c

QUESTION NO: 91

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: 92

Which two statements are true about Oracle Managed Files (OMF)?

- A. OMF cannot be used in a database that already has data files created with user-specified directions.
- B. The file system directions that are specified by OMF parameters are created automatically.
- C. OMF can be used with ASM disk groups, as well as with raw devices, for better file management.
- D. OMF automatically creates unique file names for table spaces and control files.
- E. OMF may affect the location of the redo log files and archived log files.

Answer: D,E

Explanation:

D: The database internally uses standard file system interfaces to create and delete files as needed for the following database structures:

Tablespaces
Redo log files
Control files

Archived logs
Block change tracking files
Flashback logs
RMAN backups

Note:

* Using Oracle-managed files simplifies the administration of an Oracle Database. Oracle-managed files eliminate the need for you, the DBA, to directly manage the operating system files that make up an Oracle Database. With Oracle-managed files, you specify file system directories in which the database automatically creates, names, and manages files at the database object level. For example, you need only specify that you want to create a tablespace; you do not need to specify the name and path of the tablespace's datafile with the DATAFILE clause.

<http://www.oracle-base.com/articles/9i/oracle-managed-files.php>

http://docs.oracle.com/cd/B10500_01/server.920/a96521/omf.htm

Reference: What Are Oracle-Managed Files?

QUESTION NO: 93

Which three actions are possible during an Online Data file Move operation?

- A. Creating and dropping tables in the data file being moved
- B. Performing file shrink of the data file being moved
- C. Querying tables in the data file being moved
- D. Performing Block Media Recovery for a data block in the data file being moved
- E. Flashing back the database
- F. Executing DML statements on objects stored in the data file being moved

Answer: A,C,F

Explanation: An Online Move data file operation is not compatible when:
• The data file is an OFFLINE data file
• A concurrent flashback database operation is executing
• A media recovery is completing
• A file shrink operation or tablespace offline/drop operation involving the same file is performing

QUESTION NO: 94

Your multitenant container database (CDB) contains a pluggable database, HR_PDB. The default permanent tablespace in HR_PDB is USERDATA. The container database (CDB) is open and you

connect RMAN.

You want to issue the following RMAN command:

```
RMAN > BACKUP TABLESPACE hr_pdb:userdata;
```

Which task should you perform before issuing the command?

- A. Place the root container in ARCHIVELOG mode.
- B. Take the user data tablespace offline.
- C. Place the root container in the nomount stage.
- D. Ensure that HR_PDB is open.

Answer: A

Explanation: RMAN> select name,open_mode from v\$pdb;

using target database control file instead of recovery catalog

```
NAME
OPEN_MODE-----PDB$SEED READ ONLYORA12P1 READ
WRITEORA12P2 MOUNTED
```

```
RMAN> backup tablespace ora12p2:users;
```

```
Starting backup at 31-MAR-14allocated channel: ORA_DISK_1channel ORA_DISK_1: SID=137
```

```
device type=DISKchannel ORA_DISK_1: starting full datafile backup setchannel ORA_DISK_1:
```

```
specifying datafile(s) in backup setinput datafile file number=00013
```

```
name=/appl/oradata/cdbroot/ORA12C1/F5D05369C4B23E83E0430100007F6D99/datafile/o1_mf_
```

```
users_9mhr0o5l_.dbfchannel ORA_DISK_1: starting piece 1 at 31-MAR-14channel ORA_DISK_1:
```

```
finished piece 1 at 31-MAR-14piece
```

```
handle=/appl/oradata/flash_recovery/ORA12C1/F5D05369C4B23E83E0430100007F6D99/backup
```

```
set/2014_03_31/o1_mf_nnndf_TAG20140331T001832_9mhzdb6w_.bkp
```

```
tag=TAG20140331T001832 comment=NONEchannel ORA_DISK_1: backup set complete,
```

```
elapsed time: 00:00:01Finished backup at 31-MAR-14
```

```
Starting Control File and SPFILE Autobackup at 31-MAR-14piece
```

```
handle=/appl/oradata/flash_recovery/ORA12C1/autobackup/2014_03_31/o1_mf_s_843610715_9
```

```
mhzdcv8_.bkp comment=NONEFinished Control File and SPFILE Autobackup at 31-MAR-14
```

```
RMAN>
```

The above example illustrates that if root container is open and in archive log it can do the backup of tablespace whether the pdb is mounted or open.

QUESTION NO: 95

Identify three scenarios in which you would recommend the use of SQL Performance Analyzer to analyze impact on the performance of SQL statements.

- A. Change in the Oracle Database version
- B. Change in your network infrastructure
- C. Change in the hardware configuration of the database server
- D. Migration of database storage from non-ASM to ASM storage
- E. Database and operating system upgrade

Answer: A,C,E

Explanation: Oracle 11g/12c makes further use of SQL tuning sets with the SQL Performance Analyzer, which compares the performance of the statements in a tuning set before and after a database change. The database change can be as major or minor as you like, such as:

- * (E) Database, operating system, or hardware upgrades.
- * (A,C) Database, operating system, or hardware configuration changes.
- * Database initialization parameter changes.
- * Schema changes, such as adding indexes or materialized views.
- * Refreshing optimizer statistics.
- * Creating or changing SQL profiles.

QUESTION NO: 96

Which two statements are true about the RMAN validate database command?

- A. It checks the database for intrablock corruptions.
- B. It can detect corrupt pfiles.
- C. It can detect corrupt spfiles.
- D. It checks the database for interblock corruptions.
- E. It can detect corrupt block change tracking files.

Answer: A,C

Explanation:

Block corruptions can be divided into interblock corruption and intrablock corruption. In intrablock corruption, the corruption occurs within the block itself and can be either physical or logical corruption. In interblock corruption, the corruption occurs between blocks and can only be logical corruption.

(key word) * The VALIDATE command checks for intrablock corruptions only. Only DBVERIFY and the ANALYZE statement detect interblock corruption.

VALIDATE Command Output ••> List of Control File and SPFILE.

File TYPE >>> SPFILE or Control File.

Status >»» OK if no corruption, or FAILED If block corruption is found.

Blocks Failing »»» The number of blocks that fail the corruption check. These blocks are newly corrupt.

Blocks Examined »»» Total number of blocks in the file.

Oracle' Database Backup and Recovery User's Guide

12c Release 1 (12.1) - 16 Validating Database Files and Backups

QUESTION NO: 97

You install a non-RAC Oracle Database. During Installation, the Oracle Universal Installer (OUI) prompts you to enter the path of the Inventory directory and also to specify an operating system group name.

Which statement is true?

- A. The ORACLE_BASE base parameter is not set.
- B. The installation is being performed by the root user.
- C. The operating system group that is specified should have the root user as its member.
- D. The operating system group that is specified must have permission to write to the inventory directory.

Answer: D

Explanation:

Note:

Providing a UNIX Group Name

If you are installing a product on a UNIX system, the Installer will also prompt you to provide the name of the group which should own the base directory.

You must choose a UNIX group name which will have permissions to update, install, and deinstall Oracle software. Members of this group must have write permissions to the base directory chosen.

Only users who belong to this group are able to install or deinstall software on this machine.

QUESTION NO: 98

You are required to migrate your 11.2.0.3 database as a pluggable database (PDB) to a multitenant container database (CDB).

The following are the possible steps to accomplish this task:

1. Place all the user-defined tablespaces in read-only mode on the source database.
2. Upgrade the source database to a 12c version.
3. Create a new PDB in the target container database.
4. Perform a full transportable export on the source database with the VERSION parameter set to 12 using the expdp utility.
5. Copy the associated data files and export the dump file to the desired location in the target database.
6. Invoke the Data Pump import utility on the new PDB database as a user with the DATAPUMP_IMP_FULL_DATABASE role and specify the full transportable import options.
7. Synchronize the PDB on the target container database by using the DBMS_PDS.SYNC_ODB function.

Identify the correct order of the required steps.

- A. 2, 1, 3, 4, 5, 6
- B. 1, 3, 4, 5, 6, 7
- C. 1, 4, 3, 5, 6, 7
- D. 2, 1, 3, 4, 5, 6, 7
- E. 1, 5, 6, 4, 3, 2

Answer: C

Explanation:

1. Set user tablespaces in the source database to READ ONLY.
 2. From the Oracle Database 11g Release 2 {11.2.0.3} environment, export the metadata and any data residing in administrative tablespaces from the source database using the FULL=Y and TRANSPORTABLE=ALWAYS parameters. Note that the VERSION=12 parameter is required only when exporting from an Oracle Database 11g Release 2 database:
 3. Copy the tablespace data files from the source system to the destination system. Note that the log file from the export operation will list the data files required to be moved.
 4. Create a COB on the destination system, including a PDB into which you will import the source database.
 5. In the Oracle Database 12c environment, connect to the pre-created PDB and import the dump file. The act of importing the dump file will plug the tablespace data files into the destination PDB
- Oracle White Paper - Upgrading to Oracle Database 12c -August 2013

QUESTION NO: 99

In your multitenant container database (CDB) with two pluggable database (PDBs). You want to create a new PDB by using SQL Developer.

Which statement is true?

- A. The CDB must be open.
- B. The CDB must be in the mount stage.
- C. The CDB must be in the nomount stage.
- D. All existing PDBs must be closed.

Answer: A

Explanation:

* 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. The size of the seed PDB is only about 1 gigabyte and it takes only a few seconds on a typical machine to copy it.

QUESTION NO: 100

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: 101

Examine the parameters for your database instance:

NAME	TYPE	VALUE
optimizer_adaptive_reporting_only	boolean	FALSE
optimizer_capture_sql_plan_baselines	boolean	FALSE
optimizer_dynamic_sampling	integer	2
optimizer_features_enable	string	12.1.0.1

Which three statements are true about the process of automatic optimization by using cardinality feedback?

- A. The optimizer automatically changes a plan during subsequent execution of a SQL statement if there is a huge difference in optimizer estimates and execution statistics.
- B. The optimizer can re optimize a query only once using cardinality feedback.
- C. The optimizer enables monitoring for cardinality feedback after the first execution of a query.
- D. The optimizer does not monitor cardinality feedback if dynamic sampling and multicolumn statistics are enabled.
- E. After the optimizer identifies a query as a re-optimization candidate, statistics collected by the collectors are submitted to the optimizer.

Answer: A,C,D

Explanation: C: During the first execution of a SQL statement, an execution plan is generated as usual.

D: if multi-column statistics are not present for the relevant combination of columns, the optimizer can fall back on cardinality feedback.

(not B)* Cardinality feedback. This feature, enabled by default in 11.2, is intended to improve plans for repeated executions.

optimizer_dynamic_sampling
optimizer_features_enable

* dynamic sampling or multi-column statistics allow the optimizer to more accurately estimate selectivity of conjunctive predicates.

Note:

* OPTIMIZER_DYNAMIC_SAMPLING controls the level of dynamic sampling performed by the optimizer.

Range of values. 0 to 10

* Cardinality feedback was introduced in Oracle Database 11gR2. The purpose of this feature is to automatically improve plans for queries that are executed repeatedly, for which the optimizer does not estimate cardinalities in the plan properly. The optimizer may misestimate cardinalities for a variety of reasons, such as missing or inaccurate statistics, or complex predicates. Whatever the reason for the misestimate, cardinality feedback may be able to help.

QUESTION NO: 102

Which three statements are true when the listener handles connection requests to an Oracle 12c database instance with multithreaded architecture enabled In UNIX?

- A. Thread creation must be routed through a dispatcher process
- B. The local listener may spawn a new process and have that new process create a thread
- C. Each Oracle process runs an SCM thread.
- D. Each multithreaded Oracle process has an SCM thread.
- E. The local listener may pass the request to an existing process which in turn will create a thread.

Answer: A,D,E

Explanation:

QUESTION NO: 103

Which three operations can be performed as multipartition operations in Oracle?

- A. Merge partitions of a list partitioned table
- B. Drop partitions of a list partitioned table
- C. Coalesce partitions of a hash-partitioned global index.
- D. Move partitions of a range-partitioned table
- E. Rename partitions of a range partitioned table
- F. Merge partitions of a reference partitioned index

Answer: A,B,F

Explanation: Multipartition maintenance enables adding, dropping, truncate, merge, split operations on multiple partitions.

A: Merge Multiple Partitions:

The new "ALTER TABLE ... MERGE PARTITIONS " help merge multiple partitions or subpartitions with a single statement. When merging multiple partitions, local and global index operations and semantics for inheritance of unspecified physical attributes are the same for merging two partitions.

B: Drop Multiple Partitions:

The new "ALTER TABLE ... DROP PARTITIONS " help drop multiple partitions or subpartitions with a single statement.

Example:

view plaincopy to clipboardprint?

```
SQL> ALTER TABLE Tab_tst1 DROP PARTITIONS  
Tab_tst1_PART5, Tab_tst1_PART6, Tab_tst1_PART7;
```


Table altered

SQL>

Restrictions :

- You can't drop all partitions of the table.
- If the table has a single partition, you will get the error: ORA-14083: cannot drop the only partition of a partitioned.

QUESTION NO: 104

You are connected using SQL* Plus to a multitenant container database (CDB) with SYSDBA privileges and execute the following sequence statements:

```
SQL> CREATE PLUGGABLE DATABASE NEW_PDB ADMIN USER PDB_ADMIN IDENTIFIED BY SECRET;
Pluggable database created.

SQL> ALTER PLUGGABLE DATABASE NEW_PDB OPEN;
Pluggable database altered.
SQL> ALTER SESSION SET CONTAINER = NEW_PDB;
Session altered.
SQL> GRANT CONNECT TO PDB_ADMIN;
Grant succeeded.
SQL CONNECT PDB_ADMIN/SECRET@LOCALHOST/NEW_PDB
Connected.
SQL> SELECT * FROM SESSION_PRIVS;

PRIVILEGE
-----
CREATE SESSION
SET CONTAINER

SQL> ALTER SESSION SET CONTAINER = PDB$SEED;
```

What is the result of the last SET CONTAINER statement and why is it so?

- A. It succeeds because the PDB_ADMIN user has the required privileges.
- B. It fails because common users are unable to use the SET CONTAINER statement.
- C. It fails because local users are unable to use the SET CONTAINER statement.
- D. It fails because the SET CONTAINER statement cannot be used with PDB\$SEED as the target pluggable database (PDB).

Answer: C

Explanation:

QUESTION NO: 105

Examine the details of the Top 5 Timed Events in the following Automatic Workloads Repository (AWR) report:

Event	Waits	Time(s)	Avg wait (ms)	% DB time	Wait Class
DB CPU		67		98.21	
db file sequential read	8,371	0	0	0.52	User I/O
latch: row cache objects	16	0	8	0.19	Concurrency
latch: shared pool	956	0	0	0.15	Concurrency
log file sync	25	0	2	0.06	Commit

What are three possible causes for the latch-related wait events?

- A. The size of the shared pool is too small.
- B. Cursors are not being shared.
- C. A large number COMMITS are being performed.
- D. There are frequent logons and logoffs.
- E. The buffers are being read into the buffer cache, but some other session is changing the buffers.

Answer: A,B,E

Explanation:

QUESTION NO: 106

You enabled an audit policy by issuing the following statements:

```
SQL> AUDIT POLICY ORA_DATABASE_PARAMETER BY SCOTT;
```

```
SQL> AUDIT POLICY ORA_DATABASE_PARAMETER BY SYS, SYSTEM;
```

For which database users and for which executions is the audit policy now active? Select two.

- A. SYS, SYSTEM
- B. SCOTT
- C. Only for successful executions
- D. Only for failed executions
- E. Both successful and failed executions

Answer: A,E

Explanation: * The ORA_DATABASE_PARAMETER policy audits commonly used Oracle Database parameter settings. By default, this policy is not enabled.

QUESTION NO: 107

A redaction policy was added to the SAL column of the SCOTT.EMP table:

```

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

```

All users have their default set of system privileges.

For which three situations will data not be redacted?

- A. SYS sessions, regardless of the roles that are set in the session
- B. SYSTEM sessions, regardless of the roles that are set in the session
- C. SCOTT sessions, only if the MGR role is set in the session
- D. SCOTT sessions, only if the MGR role is granted to SCOTT
- E. SCOTT sessions, because he is the owner of the table
- F. SYSTEM session, only if the MGR role is set in the session

Answer: A,B,D

Explanation:

Both users SYS and SYSTEM automatically have the EXEMPT REDACTION POLICY system privilege. (SYSTEM has the EXP_FULL_DATABASE role, which includes the EXEMPT REDACTION POLICY system privilege.) This means that the SYS and SYSTEM users can always bypass any existing Oracle Data Redaction policies, and will always be able to view data from tables (or views) that have Data Redaction policies defined on them

QUESTION NO: 108

What is the result of executing a TRUNCATE TABLE command on a table that has Flashback Archiving enabled?

- A. It fails with the ORA-665610 Invalid DDL statement on history-tracked message
- B. The rows in the table are truncated without being archived.
- C. The rows in the table are archived, and then truncated.
- D. The rows in both the table and the archive are truncated.

Answer: C

Explanation: <http://surachartopun.com/2010/06/ddl-on-tables-enabled-for-flashback.html>

QUESTION NO: 109

Which three activities are supported by the Data Recovery Advisor?

- A. Advising on block checksum failures
- B. Advising on inaccessible control files
- C. Advising on inaccessible block change tracking files
- D. Advising on empty password files
- E. Advising on invalid block header field values

Answer: A,B,E

Explanation: * Data Recovery Advisor can diagnose failures such as the following:

/ (B) Components such as datafiles and control files that are not accessible because they do not exist, do not have the correct access permissions, have been taken offline, and so on

/ (A, E) Physical corruptions such as block checksum failures and invalid block header field values

/ Inconsistencies such as a datafile that is older than other database files

/ I/O failures such as hardware errors, operating system driver failures, and exceeding operating system resource limits (for example, the number of open files)

* The Data Recovery Advisor automatically diagnoses corruption or loss of persistent data on disk, determines the appropriate repair options, and executes repairs at the user's request. This reduces the complexity of recovery process, thereby reducing the Mean Time To Recover (MTTR).

QUESTION NO: 110

You create a table with the PERIOD FOR clause to enable the use of the Temporal Validity feature of Oracle Database 12c.

Examine the table definition:

```
create table employees
(empno number, salary number,
deptid number, name varchar2(100),
period for employee_time);
```

Which three statements are true concerning the use of the Valid Time Temporal feature for the EMPLOYEES table?

- A. The valid time columns employee_time_start and employee_time_end are automatically created.
- B. The same statement may filter on both transaction time and valid temporal time by using the AS OF TIMESTAMP and PERIOD FOR clauses.
- C. The valid time columns are not populated by the Oracle Server automatically.
- D. The valid time columns are visible by default when the table is described.
- E. Setting the session valid time using DBMS_FLASHBACK_ARCHIVE.ENABLE_AT_VALID_TIME sets the visibility for data manipulation language (DML), data definition language (DDL), and queries performed by the session.

Answer: A,B,E

Explanation: A: To implement Temporal Validity(TV), 12c offers the option to have two date columns in that table which is having TV enabled using the new clause Period For in the Create Table for the newly created tables or in the Alter Table for the existing ones. The columns that are used can be defined while creating the table itself and will be used in the Period For clause or you can skip having them in the table's definition in the case of which, the Period For clause would be creating them internally.

E: ENABLE_AT_VALID_TIME Procedure

This procedure enables session level valid time flashback.

QUESTION NO: 111

Which three statements are true regarding the use of the Database Migration Assistant for Unicode (DMU)?

- A. A DBA can check specific tables with the DMU
- B. The database to be migrated must be opened read-only.
- C. The release of the database to be converted can be any release since 9.2.0.8.
- D. The DMU can report columns that are too long in the converted character set.
- E. The DMU can report columns that are not represented in the converted character set.

Answer: A,D,E

Explanation: A: In certain situations, you may want to exclude selected columns or tables from scanning or conversion steps of the migration process.

D: Exceed column limit

The cell data will not fit into a column after conversion.

E: Need conversion

The cell data needs to be converted, because its binary representation in the target character set is different than the representation in the current character set, but neither length limit issues nor invalid representation issues have been found.

* Oracle Database Migration Assistant for Unicode (DMU) is a unique next-generation migration tool providing an end-to-end solution for migrating your databases from legacy encodings to Unicode.

Incorrect:

Not C: The release of Oracle Database must be 10.2.0.4, 10.2.0.5, 11.1.0.7, 11.2.0.1, or later.

QUESTION NO: 112

Oracle Grid Infrastructure for a stand-alone server is installed on your production host before installing the Oracle Database server. The database and listener are configured by using Oracle Restart.

Examine the following command and its output:

```
$ crsctl config has
```

```
CRS-4622: Oracle High Availability Services auto start is enabled.
```

What does this imply?

A. When you start an instance on a high with SQL *Plus dependent listeners and ASM disk groups are automatically started.

B. When a database instance is started by using the SRVCTL utility and listener startup fails, the instance is still started.

C. When a database is created by using SQL*Plus, it is automatically added to the Oracle Restart configuration.

D. When you create a database service by modifying the SERVICE_NAMES initialization parameter, it is automatically added to the Oracle Restart configuration.

Answer: B

Explanation: About Startup Dependencies

Oracle Restart ensures that Oracle components are started in the proper order, in accordance with component dependencies. For example, if database files are stored in Oracle ASM disk groups, then before starting the database instance, Oracle Restart ensures that the Oracle ASM instance is started and the required disk groups are mounted. Likewise, if a component must be shut down, Oracle Restart ensures that dependent components are cleanly shut down first.

Oracle Restart also manages the weak dependency between database instances and the Oracle Net listener (the listener): When a database instance is started, Oracle Restart attempts to start the listener. If the listener startup fails, then the database is still started. If the listener later fails, Oracle Restart does not shut down and restart any database instances.

http://docs.oracle.com/cd/E16655_01/server.121/e17636/restart.htm#ADMIN12710

QUESTION NO: 113

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

Explanation:

QUESTION NO: 114

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: 115

You find this query being used in your Oracle 12c database:

```
select employee_id, first_name, salary
from hr.employees
order by employee_id
fetch first 20 percent rows only;
```

Which method is used by the optimizer to limit the rows being returned?

- A. A filter is added to the table query dynamically using ROWNUM to limit the rows to 20 percent of the total rows
- B. All the rows are returned to the client or middle tier but only the first 20 percent are returned to the screen or the application.
- C. A view is created during execution and a filter on the view limits the rows to 20 percent of the total rows.
- D. A TOP-N query is created to limit the rows to 20 percent of the total rows

Answer: C

Explanation:

QUESTION NO: 116

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: 117

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

Explanation:

QUESTION NO: 118

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 parameter 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: 119

Which two statements are true when row archival management is enabled?

- A.** The ORA_ARCHIVE_STATE column visibility is controlled by the ROW ARCHIVAL VISIBILITY session parameter.
- B.** The ORA_ARCHIVE_STATE column is updated manually or by a program that could reference activity tracking columns, to indicate that a row is no longer considered active.
- C.** The ROW ARCHIVAL VISIBILITY session parameter defaults to active rows only.
- D.** The ORA_ARCHIVE_STATE column is visible if referenced in the select list of a query.
- E.** The ORA_ARCHIVE_STATE column is updated automatically by the Oracle Server based on activity tracking columns, to indicate that a row is no longer considered active.

Answer: C,D

Explanation: <http://gavinsoorma.com/2013/08/oracle-12c-new-feature-in-database-archiving/>

QUESTION NO: 120

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

Explanation:

QUESTION NO: 121

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 if 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: 122

You have altered a non-unique index to be invisible to determine if queries execute within an acceptable response time without using this index.

Which two are possible if table updates are performed which affect the invisible index columns?

- A. The index remains invisible.
- B. The index is not updated by the DML statements on the indexed table.
- C. The index automatically becomes visible in order to have it updated by DML on the table.
- D. The index becomes unusable but the table is updated by the DML.
- E. The index is updated by the DML on the table.

Answer: A,E

Explanation: Unlike unusable indexes, an invisible index is maintained during DML statements.

Note:

* Oracle 11g allows indexes to be marked as invisible. Invisible indexes are maintained like any other index, but they are ignored by the optimizer unless the `OPTIMIZER_USE_INVISIBLE_INDEXES` parameter is set to `TRUE` at the instance or session level. Indexes can be created as invisible by using the `INVISIBLE` keyword, and their visibility can be toggled using the `ALTER INDEX` command.

QUESTION NO: 123

In your multitenant container database (CDB) containing same pluggable databases (PDBs), you execute the following commands in the root container:

```
SQL> CREATE ROLE c##role1;
SQL> GRANT create view, create procedure to c##role1;
SQL> GRANT c##role1 to c##a_admin;
```

Which two statements are true?

- A. The C ## ROLE1 role is created in the root database and all the PDBs.
- B. The C ## ROLE1 role is created only in the root database because the container clause is not used.
- C. Privileges are granted to the C##A_ADMIN user only in the root database.
- D. Privileges are granted to the C##A_ADMIN user in the root database and all PDBs.
- E. The statement for granting a role to a user fails because the CONTAINER clause is not used.

Answer: A,C

Explanation: * You can include the CONTAINER clause in several SQL statements, such as the CREATE USER, ALTER USER, CREATE ROLE, GRANT, REVOKE, and ALTER SYSTEM statements.

* * CREATE ROLE with CONTAINER (optional) clause

/ CONTAINER = ALL

Creates a common role.

/ CONTAINER = CURRENT

Creates a local role in the current PDB.

QUESTION NO: 124

The persistent configuration settings for RMAN have default for all parameters.

Identify four RMAN commands that produce a multi-section backup.

- A. BACKUP TABLESPACE SYSTEM SECTION SIZE 100M;
- B. BACKUP AS COPY TABLESPACE SYSTEM SECTION SIZE 100M;
- C. BACKUP ARCHIVELOG ALL SECTION SIZE 25M;
- D. BACKUP TABLESPACE "TEMP" SECTION SIZE 10M;
- E. BACKUP TABLESPACE "UNDO" INCLUDE CURRENT CONTROLFILE SECTION SIZE 100M;
- F. BACKUP SPFILE SECTION SIZE 1M;
- G. BACKUP INCREMENTAL LEVEL 0 TABLESPACE SYSAUX SECTION SIZE 100M;

Answer: A,B,E,G

Explanation:

QUESTION NO: 125

Flashback is enabled for your multitenant container database (CDB), which contains two pluggable database (PDBs). A local user was accidentally dropped from one of the PDBs.

You want to flash back the PDB to the time before the local user was dropped. You connect to the CDB and execute the following commands:

SQL > SHUTDOWN IMMEDIATE

SQL > STARTUP MOUNT

SQL > FLASHBACK DATABASE to TIME "TO_DATE ('08/20/12' , 'MM/DD/YY')";

Examine following commands:

1. ALTER PLUGGABLE DATABASE ALL OPEN;
2. ALTER DATABASE OPEN;
3. ALTER DATABASE OPEN RESETLOGS;

Which command or commands should you execute next to allow updates to the flashback back schema?

- A. Only 1
- B. Only 2
- C. Only 3
- D. 3 and 1
- E. 1 and 2

Answer: D

Explanation: <http://www.oracle-base.com/articles/12c/multitentant-flashback-of-container-database-12cr1.php>

QUESTION NO: 126

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 DIAGNOSTIC + 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: 127

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: 128

You are about to plug a multi-terabyte non-CDB into an existing multitenant container database (CDB) as a pluggable database (PDB).

The characteristics of the non-CDB are as follows:

- Version: Oracle Database 12c Releases 1 64-bit
- Character set: WE8ISO8859P15
- 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
- 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: D

Explanation:

Note:

* Generating a Pluggable Database Manifest File for the Non-CDB

Execute the `dbms_pdb.describe` procedure to generate the manifest file.

```
exec dbms_pdb.describe(pdb_descr_file=>'/u01/app/oracle/oradata/noncdb/noncdb.xml');
```

Shut down the noncdb instance to prepare to copy the data files in the next section.

```
shutdown immediate
```

```
exit
```

QUESTION NO: 129

Your database has the SRV1 service configured for an application that runs on middle-tier application server. The application has multiple modules. You enable tracing at the service level by executing the following command:

```
SQL > exec DBMS_MONITOR.SERV_MOD_ACT_TRACE_ENABLE ('SRV1');
```

The possible outcome and actions to aggregate the trace files are as follows:

1. The command fails because a module name is not specified.
2. A trace file is created for each session that is running the SRV1 service.
3. An aggregated trace file is created for all the sessions that are running the SRV1 service.
4. The trace files may be aggregated by using the `trcess` utility.
5. The trace files be aggregated by using the `tkprof` utility.

Identify the correct outcome and the step to aggregate by using `tkprof` utility?

- A. 1
- B. 2 and 4
- C. 2 and 5
- D. 3 and 4
- E. 3 and 5

Answer: B

Explanation: Tracing information is present in multiple trace files and you must use the `trcess`

tool to collect it into a single file.

Incorrect:

Not 1: Parameter `service_name`

Name of the service for which tracing is enabled.

`module_name`

Name of the MODULE. An optional additional qualifier for the service.

Note:

* The procedure enables a trace for a given combination of Service, MODULE and ACTION name. The specification is strictly hierarchical: Service Name or Service Name/MODULE, or Service Name, MODULE, and ACTION name must be specified. Omitting a qualifier behaves like a wild-card, so that not specifying an ACTION means all ACTIONS. Using the ALL_ACTIONS constant achieves the same purpose.

* `SERV_MOD_ACT_TRACE_ENABLE` Procedure

This procedure will enable SQL tracing for a given combination of Service Name, MODULE and ACTION globally unless an `instance_name` is specified.

```
* DBMS_MONITOR.SERV_MOD_ACT_TRACE_ENABLE(  
service_name IN VARCHAR2,  
module_name IN VARCHAR2 DEFAULT ANY_MODULE,  
action_name IN VARCHAR2 DEFAULT ANY_ACTION,  
waits IN BOOLEAN DEFAULT TRUE,  
binds IN BOOLEAN DEFAULT FALSE,  
instance_name IN VARCHAR2 DEFAULT NULL);
```

QUESTION NO: 130

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 fails 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: 131

Which three statements are true about SQL plan directives?

- A. They are tied to a specific statement or SQL ID.
- B. They instruct the maintenance job to collect missing statistics or perform dynamic sampling to generate a more optimal plan.
- C. They are used to gather only missing statistics.
- D. They are created for a query expression where statistics are missing or the cardinality estimates by the optimizer are incorrect.
- E. They instruct the optimizer to create only column group statistics.
- F. Improve plan accuracy by persisting both compilation and execution statistics in the SYSAUX tablespace.

Answer: B,D,F

Explanation:

QUESTION NO: 132

You want to flash back a test database by five hours.

You issue this command:

SQL > FLASHBACK DATABASE TO TIMESTAMP (SYSDATE - 5/24);

Which two statements are true about this flashback scenario?

- A. The database must have multiplexed redo logs for the flashback to succeed.
- B. The database must be MOUNTED for the flashback to succeed.
- C. The database must use block change tracking for the flashback to succeed.
- D. The database must be opened in restricted mode for the flashback to succeed.
- E. The database must be opened with the RESETLOGS option after the flashback is complete.
- F. The database must be opened in read-only mode to check if the database has been flashed back to the correct SCN.

Answer: B,E

Explanation:

http://docs.oracle.com/cd/B28359_01/backup.111/b28270/rcmflash.htm#BRADV80055

QUESTION NO: 133

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

Explanation:

QUESTION NO: 134

In your database, you want to ensure that idle sessions that are blocking active are automatically terminated after a specified period of time.

How would you accomplish this?

- A. Setting a metric threshold
- B. Implementing Database Resource Manager
- C. Enabling resumable timeout for user sessions
- D. Decreasing the value of the IDLE_TIME resource limit in the default profile

Answer: D

Explanation: An Oracle session is sniped when you set the idle_time parameter to disconnect inactive sessions. (It's only like sniping on ebay in that a time is set for an action to occur.)

Oracle has several ways to disconnect inactive or idle sessions, both from within SQL*Plus via resources profiles (connect_time, idle_time), and with the SQL*net expire time parameter. Here are two ways to disconnect an idle session:

Set the idle_time parameter in the user profile

Set the sqlnet.ora parameter expire_time

QUESTION NO: 135

You Execute the Following command to create a password file in the database server:

```
$ orapwd file = '+DATA/PROD/orapwprod entries = 5 ignorecase = N format = 12'
```

Which two statements are true about the password file?

- A. It records the usernames and passwords of users when granted the DBA role.
- B. It contains the usernames and passwords of users for whom auditing is enabled.
- C. Is used by Oracle to authenticate users for remote database administration.
- D. It records the usernames and passwords of all users when they are added to the OSDBA or OSOPER operating system groups.
- E. It supports the SYSBACKUP, SYSDG, and SYSKM system privileges.

Answer: C,E

Explanation:

QUESTION NO: 136

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: 137

Which three statements are true about Oracle Data Pump export and import operations?

- A. You can detach from a data pump export job and reattach later.
- B. Data pump uses parallel execution server processes to implement parallel import.
- C. Data pump import requires the import file to be in a directory owned by the oracle owner.
- D. The master table is the last object to be exported by the data pump.
- E. You can detach from a data pump import job and reattach later.

Answer: A,D,E

Explanation: http://docs.oracle.com/cd/E16655_01/server.121/e17639/dp_overview.htm

QUESTION NO: 138

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\$PWFILERS view to see the users who have been granted administrative privileges.

QUESTION NO: 139

In your Database, the TBS PERCENT USED parameter is set to 60 and the TBS PERCENT FREE parameter is set to 20.

Which two storage-tiering actions might be automated when using information Lifecycle Management (ILM) to automate data movement?

- A. The movement of all segments to a target tablespace with a higher degree of compression, on a different storage tier, when the source tablespace exceeds TBS PERCENT USED
- B. Setting the target tablespace to read-only
- C. The movement of some segments to a target tablespace with a higher degree of compression, on a different storage tier, when the source tablespace exceeds TBS PERCENT USED
- D. Setting the target tablespace offline
- E. The movement of some blocks to a target tablespace with a lower degree of compression, on a different storage tier, when the source tablespace exceeds TBS PERCENT USED

Answer: B,C

Explanation:

The value for TBS_PERCENT_USED specifies the percentage of the tablespace quota when a tablespace is considered full. The value for TBS_PERCENT_FREE specifies the targeted free percentage for the tablespace. When the percentage of the tablespace quota reaches the value of TBS_PERCENT_USED, ADO begins to move data so that percent free of the tablespace quota approaches the value of TBS_PERCENT_FREE. This action by ADO is a best effort and not a guarantee.

QUESTION NO: 140

Which three statements are true about Flashback Database?

- A. Flashback logs are written sequentially, and are archived.
- B. Flashback Database uses a restored control file to recover a database.
- C. The Oracle database automatically creates, deletes, and resides flashback logs in the Fast Recovery Area.
- D. Flashback Database can recover a database to the state that it was in before a reset logs operation.
- E. Flashback Database can recover a data file that was dropped during the span of time of the flashback.
- F. Flashback logs are used to restore to the blocks' before images, and then the redo data may be used to roll forward to the desired flashback time.

Answer: C,D,F

Explanation:

<http://docs.oracle.com/database/121/BRADV/flashdb.htm#BRADV529><http://docs.oracle.com/database/121/BRADV/glossary.htm#BRADV90169>

QUESTION NO: 141

Which statement is true about Enterprise Manager (EM) express in Oracle Database 12c?

- A. By default, EM express is available for a database after database creation.
- B. You can use EM express to manage multiple databases running on the same server.
- C. You can perform basic administrative tasks for pluggable databases by using the EM express interface.
- D. You cannot start up or shut down a database Instance by using EM express.
- E. You can create and configure pluggable databases by using EM express.

Answer: D

Explanation: O12c is integrated in database dont have his own agent like in O11G then you cannot start up or shut down a database.

QUESTION NO: 142

Examine the following command;

```
ALTER SYSTEM SET enable_ddl_logging = TRUE;
```

Which statement is true?

- A. Only the data definition language (DDL) commands that resulted in errors are logged in the alert log file.
- B. All DDL commands are logged in the alert log file.
- C. All DDL commands are logged in a different log file that contains DDL statements and their execution dates.
- D. Only DDL commands that resulted in the creation of new segments are logged.
- E. All DDL commands are logged in XML format in the alert directory under the Automatic Diagnostic Repository (ADR) home.

Answer: C

Explanation: The log files are created only in \$ADR_HOME/log

in alert log only alter system is caught. Sun Mar 30 23:30:04 2014 ALTER SYSTEM SET

enable_ddl_logging=TRUE SCOPE=BOTH; Sun Mar 30 23:32:41 2014 ALTER SYSTEM SET

enable_ddl_logging=TRUE SCOPE=BOTH; Sun Mar 30 23:39:35 2014 ALTER SYSTEM SET

enable_ddl_logging=FALSE SCOPE=BOTH; [oracle@ovs trace]\$

only in diag/rdbms/.../log/ two files are capturing this. The xml file is capturing full details and text file having only the commands.

[oracle@ovs log]\$ pwd/oraclebase/diag/rdbms/ora12c1/ora12c1/log [oracle@ovs log]\$ cat ddl/*

create table testing2(a number)

create table testing3(aa number)

[oracle@ovs log]\$ cat ddl_ora12c1.log Sun Mar 30 23:33:26 2014 diag_adl:create table testing2(a

number)diag_adl:create table testing3(aa number)

<http://docs.oracle.com/database/121/REFRN/refrn10302.htm#REFRN10302>

QUESTION NO: 143

In which two scenarios do you use SQL* Loader to load data?

- A. Transform the data while it is being loaded into the database.
- B. Use transparent parallel processing without having to split the external data first.
- C. Load data into multiple tables during the same load statement.
- D. Generate unique sequential key values in specified columns.

Answer: C,D

Explanation: http://docs.oracle.com/cd/B28359_01/server.111/b28319/ldr_concepts.htm

QUESTION NO: 144

You are connected to a pluggable database (PDB) as a common user with DBA privileges.

The STATISTICS_LEVEL parameter is PDB_MODIFIABLE. You execute the following:

```
SQL > ALTER SYSTEM SET STATISTICS_LEVEL = ALL SID = '*' SCOPE = SPFILE;
```

Which is true about the result of this command?

- A. The STATISTICS_LEVEL parameter is set to all whenever this PDB is re-opened.
- B. The STATISTICS_LEVEL parameter is set to ALL whenever any PDB is reopened.
- C. The STATISTICS_LEVEL parameter is set to all whenever the multitenant container database (CDB) is restarted.
- D. Nothing happens; because there is no SPFILE for each PDB, the statement is ignored.

Answer: A

Explanation:

http://docs.oracle.com/database/121/ADMIN/cdb_pdb_admin.htm#ADMIN14017

QUESTION NO: 145

Which two are prerequisites for performing a flashback transaction?

- A. Flashback Database must be enabled.
- B. Undo retention guarantee for the database must be configured.
- C. EXECUTE privilege on the DBMS_FLASHBACK package must be granted to the user flashing back transaction.
- D. Supplemental logging must be enabled.
- E. Recycle bin must be enabled for the database.
- F. Block change tracking must be enabled for the database.

Answer: C,D

Explanation:

http://docs.oracle.com/cd/E11882_01/appdev.112/e41502/adfns_flashback.htm#ADFNS611

To configure your database for the Flashback Transaction feature, you or your database administrator must:

With the database mounted but not open, enable ARCHIVELOG:

```
ALTER DATABASE ARCHIVELOG;
```

Open at least one archive log:

```
ALTER SYSTEM ARCHIVE LOG CURRENT;
```

If not done, enable minimal and primary key supplemental logging:

```
ALTER DATABASE ADD SUPPLEMENTAL LOG DATA;
```

```
ALTER DATABASE ADD SUPPLEMENTAL LOG DATA (PRIMARY KEY) COLUMNS;
```

If you want to track foreign key dependencies, enable foreign key supplemental logging:

```
ALTER DATABASE ADD SUPPLEMENTAL LOG DATA (FOREIGN KEY) COLUMNS;
```

Reference: Oracle Database Advanced Application Developer's Guide 11g, Using Oracle Flashback Technology

QUESTION NO: 146

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 fail.
- 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: 147

On your Oracle 12c database, you Issue the following commands to create indexes

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

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

Which two statement are correct?

- A.** Both the indexes are created; however, only the ORD_COSTOMER index is visible.
- B.** The optimizer evaluates index access from both the Indexes before deciding on which index to use for query execution plan.
- C.** Only the ORD_CUSTOMER_IX1 index is created.
- D.** Only the ORD_CUSTOMER_IX2 index is created.
- E.** Both the indexes are updated when a new row is inserted, updated, or deleted In the orders table.

Answer: A,E

Explanation: 11G has a new feature called Invisible Indexes. An invisible index is invisible to the optimizer as default. Using this feature we can test a new index without effecting the execution plans of the existing sql statements or we can test the effect of dropping an index without dropping it.

QUESTION NO: 148

Your multitenant container database has three pluggable databases (PDBs): PDB1, PDB2, and PDB3.

Which two RMAN commands may be; used to back up only the PDB1 pluggable database?

- A. BACKUP PLUGGABLE DATABASE PDB1 while connected to the root container
- B. BACKUP PLUGGABLE DATABASE PDB1 while connected to the PDB1 container
- C. BACKUP DATABASE while connected to the PDB1 container
- D. BACKUP DATABASE while connected to the boot container
- E. BACKUP PLUGGABLE database PDB1 while connected to PDB2

Answer: A,C

Explanation: To perform operations on a single PDB, you can connect as target either to the root or directly to the PDB.

* (A) If you connect to the root, you must use the PLUGGABLE DATABASE syntax in your RMAN commands. For example, to back up a PDB, you use the BACKUP PLUGGABLE DATABASE command.

* (C) If instead you connect directly to a PDB, you can use the same commands that you would use when connecting to a non-CDB. For example, to back up a PDB, you would use the BACKUP DATABASE command.

Reference: Oracle Database Backup and Recovery User's Guide 12c, About Backup and Recovery of CDBs

QUESTION NO: 149

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: 150

You upgraded from a previous Oracle database version to Oracle Database version to Oracle Database 12c. Your database supports a mixed workload. During the day, lots of insert, update, and delete operations are performed. At night, Extract, Transform, Load (ETL) and batch reporting jobs are run. The ETL jobs perform certain database operations using two or more concurrent sessions.

After the upgrade, you notice that the performance of ETL jobs has degraded. To ascertain the cause of performance degradation, you want to collect basic statistics such as the level of parallelism, total database time, and the number of I/O requests for the ETL jobs.

How do you accomplish this?

- A.** Examine the Active Session History (ASH) reports for the time period of the ETL or batch reporting runs.
- B.** Enable SQL tracing for the queries in the ETL and batch reporting queries and gather diagnostic data from the trace file.
- C.** Enable real-time SQL monitoring for ETL jobs and gather diagnostic data from the V\$SQL_MONITOR view.
- D.** Enable real-time database operation monitoring using the DBMS_SQL_MONITOR.BEGIN_OPERATION function, and then use the DBMS_SQL_MONITOR.REPORT_SQL_MONITOR function to view the required information.

Answer: D

Explanation: * Monitoring database operations

Real-Time Database Operations Monitoring enables you to monitor long running database tasks such as batch jobs, scheduler jobs, and Extraction, Transformation, and Loading (ETL) jobs as a composite business operation. This feature tracks the progress of SQL and PL/SQL queries associated with the business operation being monitored. As a DBA or developer, you can define business operations for monitoring by explicitly specifying the start and end of the operation or implicitly with tags that identify the operation.