

Oracle

Exam Questions 1Z0-062

Oracle Database 12c: Installation and Administration



NEW QUESTION 1

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. References:

NEW QUESTION 2

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:

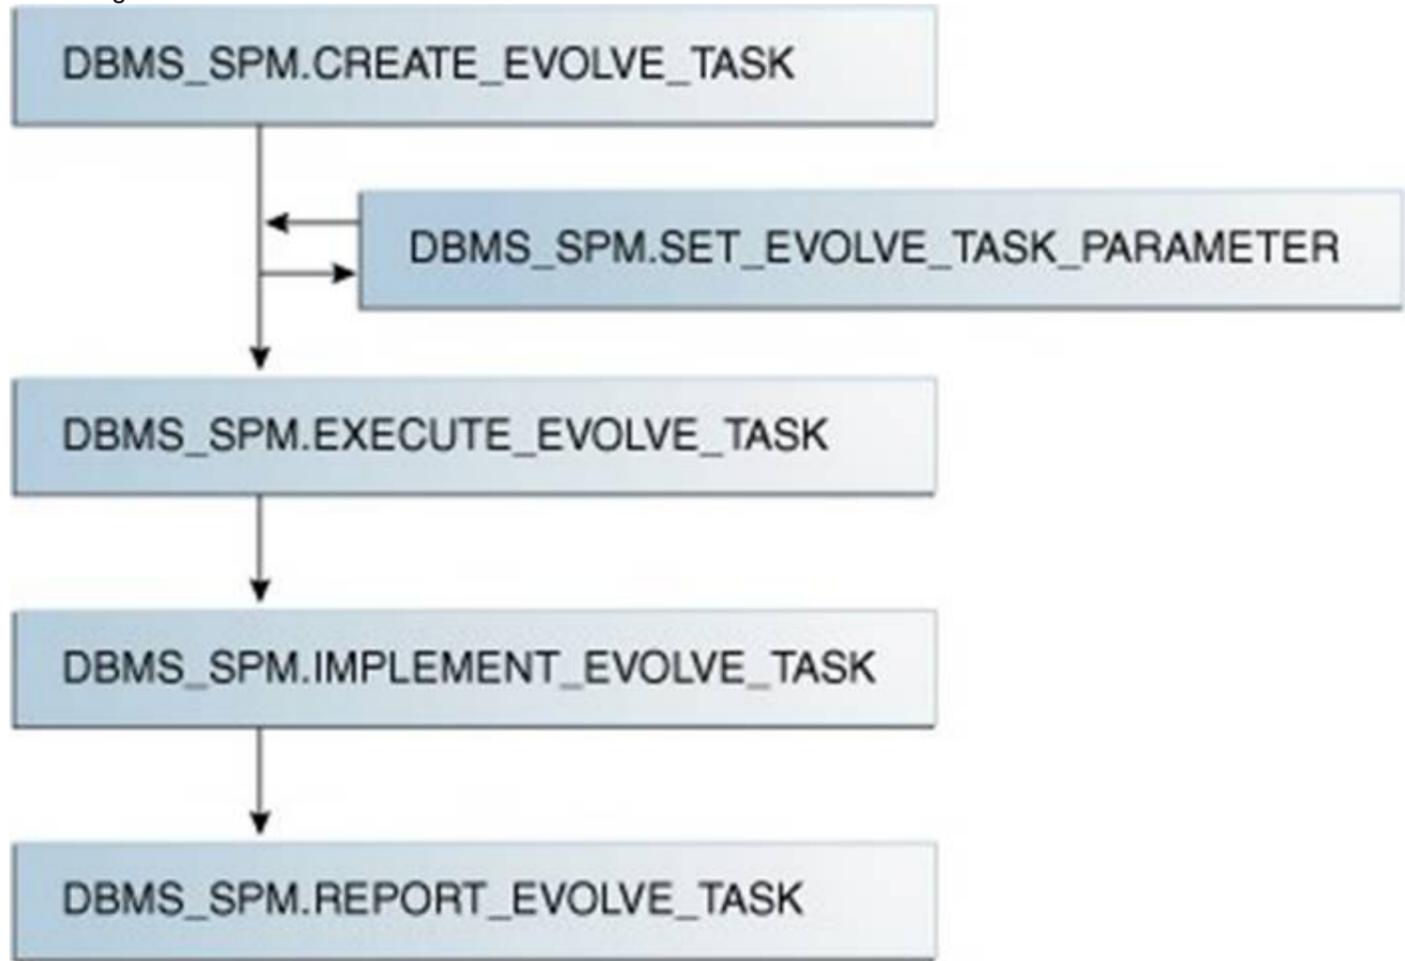
1. Set the evolve task parameters.
2. Create the evolve task by using the DBMS_SPM.CREATE_EVOLVE_TASK function.
3. Implement the recommendations in the task by using the DBMS_SPM.IMPLEMENT_EVOLVE_TASK function.
4. Execute the evolve task by using the DBMS_SPM.EXECUTE_EVOLVE_TASK function.
5. Report the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

References:

NEW QUESTION 3

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? (Choose two.)

- 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: AC

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.

References:

NEW QUESTION 4

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? (Choose two.)

- 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 trial.

Answer: AB

Explanation:

DBMS_FGA.add_policy

* The DBMS_FGA package provides fine-grained security functions.

* ADD_POLICY Procedure

This procedure creates an audit policy using the supplied predicate as the audit condition. Incorrect:

Not C: object_schema

The schema of the object to be audited. (If NULL, the current log-on user schema is assumed.)

NEW QUESTION 5

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.

NEW QUESTION 6

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: AC

Explanation:

C: To use Oracle ASM or Oracle Restart, you must first install Oracle Grid Infrastructure for a standalone server before you install and create the database. Otherwise, you must manually register the database with Oracle Restart.

Desupport of Block and Raw Devices

With the release of Oracle Database 11g release 2 (11.2) and Oracle RAC 11g release 2 (11.2), using Database Configuration Assistant or the installer to store Oracle Clusterware or Oracle Database files directly on block or raw devices is not supported.

If you intend to upgrade an existing Oracle RAC database, or an Oracle RAC database with Oracle ASM instances, then you can use an existing raw or block device partition, and perform a rolling upgrade of your existing installation.

Performing a new installation using block or raw devices is not allowed. References:

NEW QUESTION 7

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? (Choose two.)

- 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: AD

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 DATABASEFORMAT '+DISK' TAG 'ORA_ASM_MIGRATION';

References:

NEW QUESTION 8

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.

NEW QUESTION 9

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 use

Answer: A

NEW QUESTION 10

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.

NEW QUESTION 10

Which three statements are true PFILEs, SPFILEs or both? (Choose three.)

- A. SPFILEs and PFILEs may both be edited with an O/S editing utility
- B. Some SPFILE parameters can be modified successfully with the SCOPE=MEMORY clause
- C. A SPFILE can be created by an idle instance
- D. A PFILE can be created by an idle instance
- E. All SPFILE parameters can be modified successfully with the SCOPE=BOTH clause
- F. All SPFILE parameters can be modified successfully with the SCOPE=MEMORY clause

Answer: BDE

NEW QUESTION 11

Examine the query and its output:

```
SQL> SELECT reason, metric_value FROM dba_outstanding_alerts;
```

REASON	METRIC_VALUE
Tablespace [TEST] is [28 percent] full	28.125
Metrics "Current Logons Count" is at 29	29
Metrics "Database Time Spent Waiting (%)" is at 99.03754 for event class "Application"	99.0375405
db_recovery_file_dest_size of 4294967296 bytes is 97.298 used and has 116228096 remaining bytes available.	97

After 30 minutes, you execute the same query:

```
SQL> SELECT reason,metric_value FROM dba_outstanding_alerets;
```

REASON	METRIC_VALUE
Tablespace [TEST] is [28 percent] full	28.125

What might have caused three of the alerts to disappear?

- A. The threshold alerts were cleared and transferred to DBA_ALERT_HISTORY.
- B. An Automatic Workload Repository (AWR) snapshot was taken before the execution of the second query.
- C. An Automatic Database Diagnostic Monitor (ADOM) report was generated before the execution of the second query.
- D. The database instance was restarted before the execution of the second query.

Answer: D

NEW QUESTION 14

Examine this command:

```
SQL> ALTER SYSTEM SET ENABLE_DDL_LOGGING=TRUE;
```

Which two statements are true? (Choose two.)

- A. All data definition language (DDL) statements are written to the control file
- B. Some DDL statements are written to an XML file in the ADR home
- C. All DDL statements are logged in to a text file in Automatic Diagnostic Repository (ADR) home
- D. Some data definition language (DDL) statements are written to the control file
- E. Some DDL statements are written to a text file in the ADR home
- F. The Alert Log still contains some DDL statements

Answer: DE

NEW QUESTION 17

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

- 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: AE

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-CDBDatabase Multi-level resource plans Up to 32 consumer groups Subplans

/ PDBDatabase

Single-level resource plans only Up to 8 consumer groups

(not B) No subplans

NEW QUESTION 21

Which two statements are true about the RMAN validate database command? (Choose two.) A. It checks the database for intrablock corruptions.

- A. It can detect corrupt pfiles.
- B. It can detect corrupt spfiles.
- C. It checks the database for interblock corruptions.

D. It can detect corrupt block change tracking files.

Answer: AC

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

NEW QUESTION 25

Examine the resources consumed by a database instance whose current Resource Manager plan is displayed.

```
SQL> SELECT name, active_sessions, queue_length,
           consumed_cpu_time, cpu_waits, cpu_wait_time
           FROM v$rsrc_consumer_group;
```

NAME	ACTIVE_SESSIONS	QUEUE_LENGTH	CONSUMED_CPU_WAITS	CPU_WAITS	CPU_WAIT_TIME
OLTP__ORDER__ENTRY	1	0	29690	467	6709
OTHES__GROUPS	0	0	5982366	4089	60425
SYS_GROUP	1	0	2420704	914	19540
DSS_QUERIES	4	2	4594660	3004	55700

Which two statements are true? (Choose two.)

- A. An attempt to start a new session by a user belonging to DSS_QUERIES fails with an error.
- B. An attempt to start a new session by a user belonging to OTHER_GROUPS fails with an error.
- C. The CPU_WAIT_TIME column indicates the total time that sessions in the consumer group waited for the CPU due to resource management.
- D. The CPU_WAIT_TIME column indicates the total time that sessions in the consumer group waited for the CPU due to I/O waits and latch or enqueue contention.
- E. A user belonging to the DSS QUERIES resource consumer group can create a new session but the session will be queued.

Answer: CE

NEW QUESTION 27

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? (Choose two.)

- 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: AD

Explanation:

A: 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.

D:
 * Normally a separate process is fired up to do that rebalance. This will take a certain amount of time. If you want it to happen faster, fire up more processes. You tell ASM it can add more processes by increasing the rebalance power.

* ASMB
 ASM Background Process
 Communicates with the ASM instance, managing storage and providing statistics Incorrect:
 Not B: A higher, not a lower, value of DISK_REPAIR_TIME would be helpful here.
 Not E: If you implement database writer I/O slaves by setting the DBWR_IO_SLAVES parameter, you configure a single (master) DBWR process that has slave processes that are subservient to it. In addition, I/O slaves can be used to "simulate" asynchronous I/O on platforms that do not support asynchronous I/O or implement it inefficiently. Database I/O slaves provide non-blocking, asynchronous requests to simulate asynchronous I/O.

NEW QUESTION 29

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.

NEW QUESTION 30

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? (Choose two.)

- 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: AC

Explanation:

Note:

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

NEW QUESTION 35

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)? (Choose three.)

- A. LOG_BUFFER
- B. SORT_AREA_SIZE
- C. JAVA_POOL_SIZE
- D. STREAMS_POOL_SIZE
- E. DB_16K_CACHE_SIZE
- F. DB_KEEP_CACHE_SIZE

Answer: AEF

Explanation:

Manually Sized SGAComponents that Use SGA_TARGET Space SGAComponent, 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 SGAComponents 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

NEW QUESTION 37

DAILY_ORDS_LST is created in locally managed tablespace ORDERS_TBS which uses automatic segment space management.

```
CREATE TABLE daily_ords_list
  (ordno NUMBER,
   ord_date DATE)
PCTFREE 20;
```

Which two are true? (Choose two.)

- A. 80% of every data block in daily_ords_list is reserved for row inserts
- B. 20% of each data block in the table is reserved for row updates

- C. PCTFREE can help to minimize row chaining during inserts
- D. PCTFREE can help reduce row migration during updates
- E. PCTFREE eliminates row chaining during inserts

Answer: BD

NEW QUESTION 38

Which statement is true regarding the startup of a database instance?

- A. The instance does not start up normally and requires manual media recovery after a shutdown using the ABORT option.
- B. Uncommitted transactions are rolled back during the startup of the database instance after a shutdown using the immediate option.
- C. There is no difference in the underlying mechanics of the startup whether the database is shut down by using the IMMEDIATE option or the ABORT option.
- D. Media recovery is required when the database is shut down by using either the IMMEDIATE option or the ABORT option.
- E. Instance recovery is not required if the database instance was shut down by using SHUTDOWN IMMEDIATE.

Answer: E

Explanation:

References:

http://docs.oracle.com/cd/A87860_01/doc/server.817/a76956/start.htm

NEW QUESTION 43

A database is open READ WRITE and the instance has multiple sessions some of which have active transactions.

You execute this command:

```
SQL> ALTER SYSTEM ENABLE RESTRICTED SESSION;
```

Which three are true about the active transactions? (Choose three.)

- A. They may issue COMMIT OR ROLLBACK statements
- B. They are suspended and unable to issue any statements
- C. They may continue to issue DML statements
- D. They are rolled back automatically
- E. They may continue to issue queries
- F. They are terminated immediately

Answer: BDF

NEW QUESTION 46

A user establishes a connection to a database instance by using an Oracle Net connection. You want to ensure the following:

1. The user account must be locked after five unsuccessful login attempts.
2. Data read per session must be limited for the user.
3. The user cannot have more than three simultaneous sessions.
4. The user must have a maximum minutes session idle time before being logged off automatically. How would you accomplish this?

- A. by granting a secure application role to the user
- B. by implementing Database Resource Manager
- C. by using Oracle Label Security options
- D. by assigning a profile to the user

Answer: D

NEW QUESTION 48

Which two statements are true concerning dropping a pluggable database (PDB)? (Choose two.)

- 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: BC

Explanation:

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

NEW QUESTION 51

Which three statements are true about adaptive SQL plan management? (Choose three.)

- 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: ADE

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.

NEW QUESTION 52

Examine the structure of the SALES table, which is stored in a locally managed tablespace with Automatic Segment Space Management (ASSM) enabled.

Name	Null?	Type
PROD_ID	NOT NULL	NUMBER
CUST_ID	NOT NULL	NUMBER
TIME_ID	NOT NULL	DATE
CHANNEL_ID	NOT NULL	NUMBER
PROMO_ID	NOT NULL	NUMBER
QUANTITY_SOLD	NOT NULL	NUMBER (10, 2)
AMOUNT_SOLD	NOT NULL	NUMBER (10, 2)

You want to perform online segment shrink to reclaim fragmented free space below the high water mark. What should you ensure before the start of the operation?

- A. Row movement is enabled.
- B. Referential integrity constraints for the table are disabled.
- C. No queries are running on this table.
- D. Extra disk space equivalent to the size of the segment is available in the tablespace.
- E. No pending transaction exists on the table.

Answer: A

NEW QUESTION 55

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.

NEW QUESTION 59

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? (Choose three.)

- 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: ABD

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 a volume in a disk group.

NEW QUESTION 60

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

- 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: CE

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.

References:

NEW QUESTION 62

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; References:
```

NEW QUESTION 67

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? (Choose four.)

- A. You can conned as a common user by using the connect statement.
- B. You can connect as a local user by using the connect statement.
- C. You can connect by using easy connect.
- D. You can connect by using OS authentication.
- E. You can connect by using a Net Service name.
- F. You can connect as a local user by using the SET CONTAINER statemen

Answer: ACDE

NEW QUESTION 71

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

NEW QUESTION 72

Which two statements are true about the logical storage structure of an Oracle database? (Choose two.)

- A. An extent contains data blocks that are always physically contiguous on disk.
- B. An extent can span multiple segments.
- C. Each data block always corresponds to one operating system block.
- D. It is possible to have tablespaces of different block sizes.
- E. A data block is the smallest unit of I/O in data files.

Answer: DE

NEW QUESTION 77

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

- 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: DE

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 References:

NEW QUESTION 78

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: ACD

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.

NEW QUESTION 81

You want to capture column group usage and gather extended statistics for better cardinality estimates for the CUSTOMERS table in the SH schema.

Examine the following steps:

1. Issue the SELECT DBMS_STATS.CREATE_EXTENDED_STATS ('SH', 'CUSTOMERS') FROM dual statement.
2. Execute the DBMS_STATS.SEED_COL_USAGE (null, 'SH', 500) procedure.
3. Execute the required queries on the CUSTOMERS table.
4. Issue the SELECT DBMS_STATS.REPORT_COL_USAGE ('SH', 'CUSTOMERS') FROM dual statement.

Identify the correct sequence of steps.

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

Answer: B

Explanation:

Step 1 (2). Seed column usage

Oracle must observe a representative workload, in order to determine the appropriate column groups. Using the new procedure DBMS_STATS.SEED_COL_USAGE, you tell Oracle how long it should observe the workload.

Step 2: (3) You don't need to execute all of the queries in your work during this window. You can simply run explain plan for some of your longer running queries to

ensure column group information is recorded for these queries.

Step 3. (1) Create the column groups

At this point you can get Oracle to automatically create the column groups for each of the tables based on the usage information captured during the monitoring window. You simply have to call the DBMS_STATS.CREATE_EXTENDED_STATS function for each table. This function requires just two arguments, the schema name and the table name. From then on, statistics will be maintained for each column group whenever statistics are gathered on the table.

Note:

* DBMS_STATS.REPORT_COL_USAGE reports column usage information and records all the SQL operations the database has processed for a given object.

* The Oracle SQL optimizer has always been ignorant of the implied relationships between data columns within the same table. While the optimizer has traditionally analyzed the distribution of values within a column, he does not collect value-based relationships between columns.

* Creating extended statistics Here are the steps to create extended statistics for related table columns with dbms_stats.create_extended_stats:

1 - The first step is to create column histograms for the related columns. 2 - Next, we run dbms_stats.create_extended_stats to relate the columns together.

Unlike a traditional procedure that is invoked via an execute ("exec") statement, Oracle extended statistics are created via a select statement.

NEW QUESTION 86

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:

* Overview, example:

- Log into ncdb12c as sys
- Get the database in a consistent state by shutting it down cleanly.
- Open the database in read only mode
- Run DBMS_PDB.DESCRIBE to create an XML file describing the database.
- Shut down ncdb12c
- Connect to target CDB (CDB2)
- Check whether non-cdb (NCDB12c) can be plugged into CDB(CDB2)
- Plug-in Non-CDB (NCDB12c) as PDB(NCDB12c) into target CDB(CDB2).
- Access the PDB and run the noncdb_to_pdb.sql script.
- Open the new PDB in read/write mode.

* You can easily plug an Oracle Database 12c non-CDB into a CDB. Just create a PDB manifest file for the non-CDB, and then use the manifest file to create a cloned PDB in the CDB.

* Note that to plug in a non-CDB database into a CDB, the non-CDB database needs to be of version 12c as well. So existing 11g databases will need to be upgraded to 12c before they can be part of a 12c CDB.

NEW QUESTION 88

Your database is configured in ARCHIVELOG mode. Examine the RMAN configuration parameters:

```
CONFIGURE RETENTION POLICY TO REDUNDANCY 1; # default
CONFIGURE BACKUP OPTIMIZATION OFF; # default
CONFIGURE CONTROLFILE AUTOBACKUP OFF; # default
CONFIGURE DEVICE TYPE DISK PARALLELISM 1 BACKUP TYPE TO COPY;
CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE DISK TO 1; # default
```

Examine the command:

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

Which two are true? (Choose two.)

- A. It fails because the DELETE INPUT option can be used only with the BACKUP AS COPY command
- B. It creates image copies of the archivelogs
- C. It creates a backupset of archive log files
- D. It creates image copies of the database files
- E. It fails because the DELETE INPUT option can be used only with the BACKUP AS BACKUPSET command

Answer: BE

NEW QUESTION 90

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? (Choose two.)

- 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: BD

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).

NEW QUESTION 92

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: ACE

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.

NEW QUESTION 96

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? (Choose two.)

- 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: AC

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.

NEW QUESTION 99

The HR user receives the following error while inserting data into the sales table:

```
ERROR at line 1:  
ORA-01653: unable to extend table HR.SALES by 128 in tablespace USERS
```

On investigation, you find that the users tablespace uses Automatic Segment Space Management (ASSM). It is the default tablespace for the HR user with an unlimited quota on it.

Which two methods would you use to resolve this error? (Choose two.)

- A. Altering the data file associated with the USERS tablespace to extend automatically
- B. Adding a data file to the USERS tablespace
- C. Changing segment space management for the USERS tablespace to manual
- D. Creating a new tablespace with autoextend enabled and changing the default tablespace of the HR user to the new tablespace
- E. Enabling resumable space allocation by setting the RESUMABLE_TIMEOUT parameter to a nonzero value

Answer: AB

NEW QUESTION 101

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

- 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: CE

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.

NEW QUESTION 104

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.

NEW QUESTION 109

As a user of the ORCL database, you establish a database link to the remote HQ database such that all users in the ORCL database may access tables only from the SCOTT schema in the HQ database. SCOTT's password is TIGER. The service name "HQ" is used to connect to the remote HQ database.

Which command would you execute to create the database link?

- A. CREATE DATABASE LINK HQ USING 'HQ';
- B. CREATE DATABASE LINK HQ CONNECT TO CURRENT_USER USING 'HQ';
- C. CREATE PUBLICDATABASE LINK HQ CONNECT TO scott IDENTIFIED BY tiger USING 'HQ';
- D. CREATE DATABASE LINK HQ CONNECT TO scott IDENTIFIED BY tiger USING 'HQ';

Answer: C

NEW QUESTION 112

Identify three 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: ACD

Explanation:

Use the CREATE PLUGGABLE DATABASE statement to create a pluggable database (PDB). This statement enables you to perform the following tasks:

* (A) Create a PDB by using the seed as a template

Use the create_pdb_from_seed clause to create a PDB by using the seed in the multitenant container database (CDB) as a template. The files associated with the seed are copied to a new location and the copied files are then associated with the new PDB.

* (C) Create a PDB by cloning an existing PDB

Use the create_pdb_clone clause to create a PDB by copying an existing PDB (the source PDB) and then plugging the copy into the CDB. The files associated with the source PDB are copied to a new location and the copied files are associated with the new PDB. This operation is called cloning a PDB.

The source PDB can be plugged in or unplugged. If plugged in, then the source PDB can be in the same CDB or in a remote CDB. If the source PDB is in a remote CDB, then a database link is used to connect to the remote CDB and copy the files.

* Create a PDB by plugging an unplugged PDB or a non-CDB into a CDB

Use the create_pdb_from_xml clause to plug an unplugged PDB or a non-CDB into a CDB, using an XML metadata file.

NEW QUESTION 113

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 (AWRM).
- 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.

NEW QUESTION 115

Which three features work together, to allow a SQL statement to have different cursors for the same statement based on different selectivity ranges? (Choose three.)

- 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: ACD

NEW QUESTION 117

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? (Choose three.)

- A. Use Oracle Managed Files (OMF).
- B. Specify the SEEDFILE_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: ABC

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 SEEDFILE_NAME_CONVERT Clause
- * (C) The PDB_FILE_NAME_CONVERT Initialization Parameter

NEW QUESTION 119

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$PWFIL_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? (Choose two.)

- 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: BD

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\$PWFIL_Users

/ 12c: V\$PWFIL_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 SYSOPERVARCHAR2(5)If TRUE, the user can connect with SYSOPER privileges

Incorrect:

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

NEW QUESTION 120

You created a new database using the "create database" statement without specifying the "ENABLE PLUGGABLE" clause.

What are two effects of not using the "ENABLE PLUGGABLE database" clause?

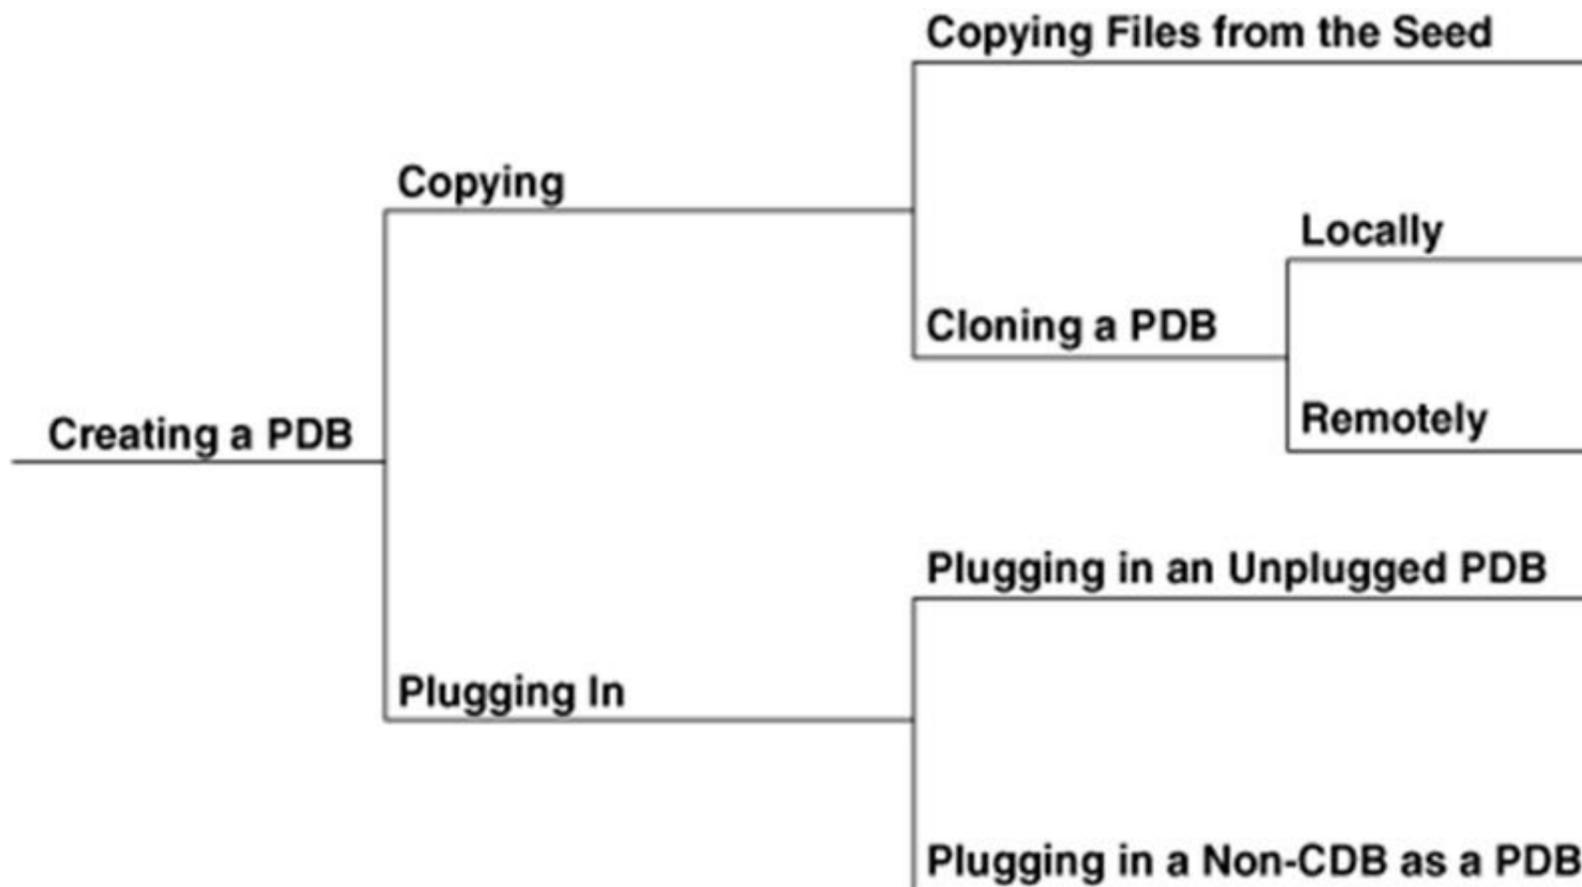
- A. The database is created as a non-CDB and can never contain a PDB.
- B. The database is treated as a PDB and must be plugged into an existing multitenant container database (CDB).
- C. The database is created as a non-CDB and can never be plugged into a CDB.
- D. The database is created as a non-CDB but can be plugged into an existing CDB.
- E. The database is created as a non-CDB but will become a CDB whenever the first PDB is plugged in.

Answer: AD

Explanation:

A (not B,not E): The CREATE DATABASE ... ENABLE PLUGGABLE DATABASE SQL statement creates a new CDB. If you do not specify the ENABLE PLUGGABLE DATABASE clause, then the newly created database is a non- CDB and can never contain PDBs.

D: You can create a PDB by plugging in a Non-CDB as a PDB. The following graphic depicts the options for creating a PDB:



Incorrect:

Not E: For the duration of its existence, a database is either a CDB or a non-CDB. You cannot transform a non-CDB into a CDB or vice versa. You must define a database as a CDB at creation, and then create PDBs within this CDB.

NEW QUESTION 125

Which three statements are true about Automatic Workload Repository (AWR)? (Choose three.)

- A. All AWR tables belong to the SYSTEM schema.
- B. The AWR data is stored in memory and in the database.
- C. The snapshots collected by AWR are used by the self-tuning components in the database
- D. AWR computes time model statistics based on time usage for activities, which are displayed in the v\$SYS time model and V\$SESS_TIME_MODEL views.
- E. AWR contains system wide tracing and logging information.

Answer: BCD

NEW QUESTION 128

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? (Choose two.)

- 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: CE

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'.

NEW QUESTION 130

You set the following parameters in the parameter file and restart the database instance:

```
MEMORY_TARGET=500M
PGA_AGGREGATE_TARGET=90M
SGA_TARGET=270M
```

Which two statements are true? (Choose two.)

- A. The MEMORY_MAX_TARGET parameter is automatically set to 500 MB.
- B. The PGA_AGGREGATE_TARGET and SGA_TARGET parameters are automatically set to zero.
- C. The value of the MEMORY_MAX_TARGET parameter remains zero for the database instance.
- D. The lower limits of the PGA_AGGREGATE_TARGET and SGA_TARGET parameters are set to 90 MB and 270 MB respectively.
- E. The instance does not start up because Automatic Memory Management (AMM) is enabled but PGA_AGGREGATE_TARGET and SGA_TARGET parameters are set to nonzero values.

Answer: AD

NEW QUESTION 134

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.

NEW QUESTION 136

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

- 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: CDE

NEW QUESTION 141

Your production database uses file system storage. You want to move storage to Oracle Automatic Storage Management (ASM). How would you achieve this?

- A. by using a transportable database
- B. by using the Database Upgrade Assistant (DBUA)
- C. by using Data Pump
- D. by using RMAN

Answer: D

Explanation:

References:

http://docs.oracle.com/cd/E11882_01/server.112/e18951.pdf (p.184)

NEW QUESTION 142

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

NEW QUESTION 144

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:

References: <http://www.oracle.com/technetwork/database/manageability/emx-intro-1965965.html>

NEW QUESTION 146

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 (PMON) 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:

<https://docs.oracle.com/database/121/CNCPT/process.htm>

NEW QUESTION 148

The HR user executes the following query on the EMPLOYEES table but does not issue COMMIT, ROLLBACK, or any data definition language (DDL) command after that:

```
SQL> SELECT job
      FROM employees
      WHERE job='CLERK' FOR UPDATE OF empno;
```

HR then opens a second session.

Which two operations wait when executed in HR's second session? (Choose two.)

- A. LOCK TABLE employees IN EXCLUSIVE MODE;
- B. INSERT INTO employees(empno,ename) VALUES (1289, 'Dick');
- C. SELECT job FROM employees WHERE job='CLERK' FOR UPDATE OF empno;
- D. SELECT empno,ename FROM employees WHERE job='CLERK';
- E. INSERT INTO employees(empno,ename,job) VALUES (2001,'Harry','CLERK');

Answer: AC

NEW QUESTION 150

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? (Choose three.)

- 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 CK 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: ACD

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.

NEW QUESTION 151

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;

NEW QUESTION 153

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

NEW QUESTION 155

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? (Choose three.)

- 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 there are 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: ACD

NEW QUESTION 160

Which three statements are true about SQL plan directives? (Choose three.)

- 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: BDF

NEW QUESTION 165

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? (Choose three.)

- 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: ABD

NEW QUESTION 170

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? (Choose three.)

- 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: ADE

NEW QUESTION 171

You want to prevent a group of users in your database from performing long-running transactions that consume huge amounts of space in the undo tablespace. If the quota for these users is exceeded during execution of a data manipulation language (DML) statement, the operation should abort and return an error.

However, queries should still be allowed, even if users have exceeded the undo space limitation.

How would you achieve this?

- A. Specify the maximum amount of quota a user can be allocated in the undo tablespace.
- B. Decrease the number of Interested Transaction List (ITL) slots for the segments on which these users perform transactions.
- C. Implement a profile for these users.
- D. Implement a Database Resource Manager plan

Answer: D

NEW QUESTION 174

Which two statements are true about Oracle Data Pump export and import operations? (Choose two.)

- 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: AB

Explanation:

B: Data Pump can employ multiple worker processes, running in parallel, to increase job performance.
 D: For export jobs, the master table records the location of database objects within a dump file set. / Export builds and maintains the master table for the duration of the job. At the end of an export job, the content of the master table is written to a file in the dump file set.
 / For import jobs, the master table is loaded from the dump file set and is used to control the sequence of operations for locating objects that need to be imported into the target database.

NEW QUESTION 177

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? (Choose three.)

- 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: ACE

Explanation:

Note:
 * Oracle 12c now provides the ability to index a subset of partitions and to exclude the others. Local and global indexes can now be created on a subset of the partitions of a table. Partial Global indexes provide more flexibility in index creation for partitioned tables. For example, index segments can be omitted for the most recent partitions to ensure maximum data ingest rates without impacting the overall data model and access for the partitioned object. Partial Global Indexes save space and improve performance during loads and queries. This feature supports global indexes that include or index a certain subset of table partitions or subpartitions, and exclude the others. This operation is supported using a default table indexing property. When a table is created or altered, a default indexing property can be specified for the table or its partitions.

NEW QUESTION 182

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? (Choose three.)

- 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: BDE

Explanation:

B: SYSDG administrative privilege has ability to perform Data Guard operations (including startup and shutdown) using Data Guard Broker or dgmgrl.
 D: SYSASM
 The new (introduced in 11g) SYSASM role to manage the ASM instance, variable extent sizes to reduce shared pool usage, and the ability of an instance to read from a specific disk of a diskgroup
 E (Not A): SYSDBA is like a role in the sense that it is granted, but SYSDBA is a special built-in privilege to allow the DBA full control over the database
 Incorrect:
 Not C: SYSKM. SYSKM administrative privilege has ability to perform transparent data encryption wallet operations. Note: Use the V\$PWFILE_USERS view to see the users who have been granted administrative privileges.

NEW QUESTION 183

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

```
SQL > CREATE UNDO TABLESPACE undotb01  

DATAFILE 'u01/oracle/rddb1/undotbs01.dbf' SIZE 60M AUTOEXTEND ON; What is the result?
```

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

Answer: E

Explanation:

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

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

NEW QUESTION 188

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

- 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: ACE

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.

NEW QUESTION 191

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? (Choose two.)

- 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: AE

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.

NEW QUESTION 196

Which two actions does an incremental checkpoint perform? (Choose two.)

- A. It signals CKPT to write the checkpoint position to the data file headers.
- B. It writes the checkpoint position to the data file headers.
- C. It advances the checkpoint position in the checkpoint queue.
- D. It writes the checkpoint position to the control file.

Answer: CD

Explanation:

References:

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

NEW QUESTION 200

You have just completed a manual upgrade of an Oracle 11g Database to Oracle Database 12c.

The Post-Upgrade Status Tool reports an INVALID status for some of the components after the upgrade. What must you do first in this situation to attempt to fix this problem?

- A. Run catuppst.sql to perform revalidation actions
- B. Run utllobj.sql to filter out objects that were invalidated by the upgrade process.
- C. Run preupgrd.sql and then execute the generated "fix-up" scripts to resolve status issues.
- D. Run utlrp.sql to recompile stored PL/SQL and Java code and check the DBA_REGISTRY view

Answer: D

NEW QUESTION 203

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

- 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 tablespac

Answer: AB

NEW QUESTION 207

You want to create a role that:

- is protected from unauthorized usage
 - does not use a password embedded in the application source code or stored in a table
 - is enabled for a user based on security policies defined in a PL/SQL package
- How would you create this role?

- A. as a secure application role
- B. with definer's rights
- C. with global authentication
- D. with external authentication

Answer: A

Explanation:

References: https://docs.oracle.com/cd/B28359_01/network.111/b28531/authorization.htm#DBSEG97973

NEW QUESTION 208

The user SCOTT owns the CUST table that is placed in the SALES tablespace. The user SCOTT opens a session and executes commands as follows:
SQL> INSERT INTO cust VALUES(101, 'JACK'); 1 row created. SQL> INSERT INTO cust VALUES(102, 'SMITH'); 1 row created.

As a DBA, you execute the following command from another session: ALTER TABLESPACE sales READ ONLY; Which statement is true regarding the effect of this command on the transaction in Scott's session?

- A. The command fails as a transaction is still pending.
- B. The transaction in Scott's session is rolled back and the tablespace becomes readonly.
- C. The command waits and the user SCOTT can execute data manipulation language (DML) statements only as part of the current transaction.
- D. The command hangs until all transactions on the objects in the tablespace commit or rollback, and then the tablespace is placed in readonly mode.

Answer: B

NEW QUESTION 211

Which two statements are true about standard database auditing? (Choose two.)

- A. DDL statements can be audited.
- B. Statements that refer to standalone procedure can be audited.
- C. Operations by the users logged on as SYSDBA cannot be audited.
- D. Only one audit record is ever created for a session per audited statement even though it is executed more than once.

Answer: AB

NEW QUESTION 213

You use multiple temporary tables frequently in your database. Which two are benefits of configuring temporary undo? (Choose two.)

- A. Performance improves because less redo is written to the redo log.
- B. Temporary undo reduces the amount of undo stored in undo tablespaces.
- C. Performance improves because data manipulation language (DML) operations performed on temporary tables do not use the buffer cache.
- D. Performance improves because no redo and undo are generated for the temporary table.

Answer: AB

NEW QUESTION 218

Which three statements are true about Flashback Database? (Choose three.)

- 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: CDF

NEW QUESTION 223

Which two are prerequisites for performing a flashback transaction? (Choose two.)

- 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: BD

Explanation:

References: <http://searchoracle.techtarget.com/tip/How-to-perform-Oracle-Flashback-Transaction-Queries>
https://docs.oracle.com/cd/E11882_01/appdev.112/e41502/adfns_flashback.htm#ADFNS610

NEW QUESTION 228

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: ABEG

NEW QUESTION 231

You have the following entry in the tnsnames.ora of your hq.us.example.com host machine:

```
ORCL =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP) (HOST = hq.us.example.com) (PORT = 1521))
    )
    CONNECTED_DATA =
      (SERVICE_NAME = ORCL.us.example.com)
  )
)
```

You issue the following command at the command prompt: `Sqlplus HR/HR@ORCL`
 Which statement is true about the connection to the ORCL database instance?

- A. The connection succeeds, provided the NAMES.DEFAULT_DOMAIN parameter is set to us.example.com in the sqlnet.ora file on the client side.
- B. The connection fails because the net service name does not have the suffix us.example.com.
- C. The connection succeeds, provided the SERVICE_NAMES initialization parameter is set to ORCL.
- D. The connection succeeds, provided the ORCL.us.example.com database service is registered with a listener, the listener is up, and the database is open.

Answer: D

NEW QUESTION 233

In which two scenarios do you use SQL* Loader to load data? (Choose two.)

- 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 column

Answer: CD

NEW QUESTION 235

Your database instance is started by using a server parameter file (SPFILE). You execute the following command to change the value of the LOG_BUFFER initialization parameter:

```
ALTER SYSTEM SET LOG_BUFFER=32 M;
```

What is the outcome of this command?

- A. The parameter value is changed and it comes into effect as soon as space becomes available in the SGA.
- B. It returns an error because the value of this parameter cannot be changed dynamically.
- C. The parameter value is changed and it comes into effect at the next instance startup.
- D. It returns an error because SCOPE should be set to MEMOR

Answer: B

NEW QUESTION 238

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 or 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: D

NEW QUESTION 239

What can be automatically implemented after the SQL Tuning Advisor is run as part of the Automated Maintenance Task?

- A. statistics recommendations
- B. SQL profile recommendations
- C. SQL statement restructure recommendations
- D. creation of materialized views to improve query performance

Answer: B

NEW QUESTION 244

Which two statements are true when row archival management is enabled? (Choose two.)

- 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: CD

NEW QUESTION 247

You create a table with the PERIODFOR 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? (Choose three.)

- 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 PERIODFOR 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: ABC

NEW QUESTION 251

Which activity is audited by default and recorded in the operating system audit trail irrespective of whether or not database auditing is enabled?

- A. execution of SQL statements by users connected with the SYSDBA privilege
- B. creation of a fine-grained audit policy
- C. configuration of unified auditing mode
- D. usage of the AUDIT statement

Answer: A

Explanation:

References https://docs.oracle.com/cd/B28359_01/network.111/b28531/auditing.htm#DBSEG0622

NEW QUESTION 255

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? (Choose two.)

- 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: AC

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.

NEW QUESTION 258

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? (Choose two.)

- 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: AC

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.

NEW QUESTION 259

You have successfully taken a database backup by using the command: RMAN> BACKUP AS BACKUPSET DATABASE;

Now you execute this command:

RMAN> BACKUP INCREMENTAL LEVEL 1 DATABASE;

What is the outcome?

- A. It fails because an incremental level 1 backup always searches for an image copy as level 0 backup.
- B. It fails because an incremental level 0 backup does not exist.
- C. It takes a backup of blocks that have been formatted since the last full database backup.
- D. It takes an incremental level 0 backup of the database.
- E. It first takes an incremental level 0 backup and then an incremental level 1 backup.

Answer: E

Explanation:

References: https://docs.oracle.com/cd/B19306_01/backup.102/b14192/bkup004.htm (4.4.1.2)

NEW QUESTION 260

Examine this command executed on a client that is remote from the database server. SQL> CONNECT hr/hr@orcl Which two are required for this command to connect the SQLPLUS client to a database instance? (Choose two.)

- A. An orcl TNS entry must be defined in the client-side and server-side tnsnames.ora files
- B. An orcl TNS entry must be defined in the client-side tnsnames.ora file
- C. A service name must be defined to the listener that matches the service name in the orcl TNS entry
- D. An orcl TNS entry must be defined in the server-side tnsnames.ora file
- E. The service name orcl must be defined to the listener

Answer: DE

NEW QUESTION 265

You enabled block change tracking for faster incremental backups in your database. Which background process writes to the change tracking file?

- A. RBAL
- B. CKPT
- C. SMON
- D. PMON
- E. MMON
- F. CTWR
- G. DBWR

Answer: F

NEW QUESTION 269

Which statement is true about a database in ARCHIVELOG mode?

- A. All backups taken prior to switching to ARCHIVELOG mode can be used to perform complete recovery.
- B. Online redo log files have to be multiplexed before putting the database in ARCHIVELOG mode.
- C. A Fast Recovery Area (FRA) must be configured for the database.
- D. Full database backups can be performed when the database is open.

Answer: D

NEW QUESTION 270

You have a production Oracle 12c database running on a host.

You want to install and create databases across multiple new machines that do not have any Oracle database software installed. You also want the new databases to have the same directory structure and components as your existing 12c database.

The steps in random order:

1. Create directory structures similar to the production database on all new machines.
2. Create a response file for Oracle Universal Installer (OUI) with the same configurations as the production database.
3. Create a database clone template for the database.
4. Run the Database Configuration Assistant (DBCA) to create the database.
5. Run OUI in graphical mode on each machine.
6. Run OUI in silent mode using the OUI response file.

Identify the required steps in the correct sequence to achieve the requirement with minimal human intervention.

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

Answer: E

NEW QUESTION 274

The schema SALES exists in two databases, ORCL1 and ORCL2, and has the same password, SALES123. User SALES has CREATE DATABASE LINK and CREATE SESSION privileges on both databases. Examine these commands: Conn SALES/SALES123

```
CREATE DATABASE LINK orcl2 USING 'orcl2';
```

What is the outcome of executing these commands in the ORCL1 database?

- A. ORCL2 is created as a public database link to connect a single session to the SALES schema in the ORCL2 database.
- B. ORCL2 is created as a shared database link to connect multiple sessions to the SALES schema in the ORCL2 database.
- C. ORCL2 is created as a private database link to connect to only the SALES schema in the ORCL2 database.
- D. ORCL2 database link creation fail

Answer: C

NEW QUESTION 279

Identify two prerequisites for configuring Enterprise Manager Database Express (EM Express).

- A. Grant the APEX_PUBLIC_USER role to the SYSMAN user.
- B. Use the DBMS_XDB_CONFIG.SETHTTPPORT procedure to configure a port number for Oracle HTTP Server.
- C. Install Oracle HTTP Server.
- D. Configure at least one dispatcher for the TCP/IP protocol.
- E. Create a SYSMAN user with the SYSDBA privilege as an administrator for EM Express

Answer: BD

NEW QUESTION 282

You want to create a database and you have the following:

- Oracle Grid Infrastructure is installed and configured.
- Oracle Database Vault is installed in ORACLE_HOME to be used for this database.
- Oracle Enterprise Manager Cloud Control is available and an agent is deployed on the database server. Examine the requirements:

1. configuring the database instance to support shared server mode
2. using Automatic Storage Management (ASM) for storing database files.
3. configuring a naming method to help a remote user connect to a database instance
4. configuring the Fast Recovery Area
5. configuring Database Vault
6. configuring Enterprise Manager (EM) Database Express
7. registering with EM Cloud Control
8. configuring remote log archive destinations
9. enabling daily incremental backups
10. configuring a nondefault block size for nondefault block size tablespaces

Which of these requirements can be met while creating a database by using the Database Configuration Assistant (DBCA)?

- A. 1, 2, 4, 5, 7, 8, 9 and 10
- B. 1, 2, 4, 5, 6 and 7
- C. 1, 2, 3, 8, 9 and 10
- D. 1, 2, 3, 4, 6, 8, 9 and 10
- E. 1, 2, 4, 5, 6, 7 and 8

Answer: D

NEW QUESTION 283

Which three functions can be performed by the SQL Tuning Advisor? (Choose three.)

- A. recommending creation of indexes based on SQL workload
- B. recommending restructuring of SQL statements that have suboptimal plans
- C. checking schema objects for missing and state statistics
- D. recommending optimization of materialized views
- E. generating SQL profiles

Answer: BCE

NEW QUESTION 284

Examine the following steps:

ADBA grants the CREATE TABLE system privilege with ADMIN OPTION to the user SIDNEY. SIDNEY grants the CREATE TABLE system privilege to the HR user.

Which statement is true?

- A. SIDNEY can revoke the CREATE TABLE system privilege only from HR, to whom he granted it.
- B. HR can grant the CREATE TABLE system privilege to other users.
- C. Neither SIDNEY nor HR can create new tables if the DBA revokes the CREATE TABLE privilege from SIDNEY.
- D. HR still retains the CREATE TABLE system privilege if the DBA revokes the CREATE TABLE privilege from SIDNEY.

Answer: D

Explanation:

References:

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

NEW QUESTION 288

The performance of your database degrades between 11:00 AM and 3:00 PM. Automatic Workload Repository (AWR) snapshots are collected on an hourly basis. What is the most efficient way of diagnosing this problem?

- A. Create a custom ADDM task for the period defined by the snapshots taken between 11:00 AM and 3:00 PM.
- B. Analyze the latest Automatic Database Diagnostic Monitor (ADDM) report.
- C. Analyze the hourly ADDM reports generated between 11:00 AM and 3:00 PM.
- D. Create a SQL Tuning Set (STS) for the currently cached SQL statements in the shared pool and run SQL Performance Analyzer (SPA) to generate recommendations.

Answer: A

NEW QUESTION 293

In which situations does the Database Writer process (DBWn) write to data files? (choose two).

- A. when the RMAN recovery process starts
- B. when a user process commits a transaction
- C. when a tablespace is made read-only or taken offline
- D. when PMON cleans up dirty buffers in the database buffer cache
- E. when clean buffers for reading new blocks into the database buffer cache are not found easily

Answer: BD

Explanation:

References https://docs.oracle.com/cd/B19306_01/server.102/b14220/process.htm

NEW QUESTION 297

A database uses Automatic Storage Management (ASM) as database storage, which has a diskgroup, DATA1, which is created as follows:

```
SQL> CREATE DISKGROUP data1 NORMAL REDUNDANCY
      FAILGROUP failgrp1 DISK '/dev/sda1', '/dev/sda2'
      FAILGROUP failgrp2 DISK '/dev/sda3', '/dev/sda4';
```

What happens when the FAILGRP1 failure group is corrupted?

- A. Mirroring of allocation units occurs within the FAILGRP2 failure group.
- B. Transactions that are using the diskgroup fail.
- C. ASM does not mirror any data and newly allocated primary allocation units (AU) are stored in the FAILGRP2 failure group.
- D. Data in the FAILGRP1 failure group is moved to the FAILGRP2 failure group and rebalancing is started.

Answer: D

NEW QUESTION 302

You create a locally managed tablespace ORDERS_TBS with automatic segment management.

You then create the table DAILY_ORDS_LST in the ORDERS_TBS tablespace using the command. CREATE TABLE daily_ords_1st(ordno NUMBER, ord_date DATE) PCTFREE 20;

How does the PCTFREE storage parameter influence data storage for this table?

- A. It allows only 80% of space to be occupied in all data blocks of this table.
- B. It minimizes row chaining during row insertion.
- C. It minimizes row migration during existing row data updation.
- D. It automatically coalesces free space of a data block when it reaches 20% of available space

Answer: A

NEW QUESTION 303

Which statement is true regarding the DEFAULT profile?

- A. The values assigned to the resource limits and password parameters in the default profile can be altered.
- B. A different DEFAULT profile can be created before each user in a database.
- C. It can be dropped and recreated.
- D. it must be explicitly assigned to the user

Answer: A

NEW QUESTION 307

The HR user owns the BONUSSES table. HR grants privileges to the user TOM by using the command: SQL> GRANT SELECT ON bonuses TO tom WITH GRANT OPTION;

TOM then executes this command to grant privileges to the user JIM: SQL> GRANT SELET ON hr.bonuses TO jim; Which statement is true?

- A. TOM cannot revoke the SELECT ON HR.BONUSSES privilege from JIM.
- B. HR can revoke the SELECT ON HR.BONUSSES privilege from JIM.
- C. JIM can grant the SELECT ON HR.BONUSSES privilege to other users, but cannot revoke the privilege from them.
- D. HR can revoke the SELECT ON HR.BONUSSES privilege from TOM, which will automatically revoke the SELECT ON HR.BONUSSES privilege from JIM.

Answer: D

NEW QUESTION 312

In your database, archive logging and control file autobackup are enabled.

The data files and redo log files are intact but control files are impacted due to media failure. In which two recovery scenarios must you use the RESETLOGS option? (Choose two.)

- A. One control file copy is intact so the spfile is changed to refer to only one copy.
- B. One control file copy is intact and damaged control file copies have to be restored to the default location.
- C. All copies of the control file are damaged and the CREATE CONTROLFILE statement is executed manually.
- D. All copies of the control file are damaged and the auto backed up control file is used for recovery.
- E. One control file copy is intact and damaged control file copies have to be restored to a non-default location

Answer: CD

NEW QUESTION 313

Which two statements are true about Oracle network connections? (Choose two.)

- A. A listener may listen on behalf of only one database instance at a time.
- B. A server process checks a user's authentication credentials and creates a session if the credentials are valid.
- C. The listener continuously monitors a connection after the user process connects to a service handler.
- D. The listener always spawns a new server process to deal with each new connection.
- E. A connection request from a client is always first received by a listener running on the port that is used for the connection request for the database server.

Answer: BE

NEW QUESTION 316

Which three statements are true about naming methods? (Choose three.)

- A. Local naming supports multiple protocols, but for any one connection, the client and server must use the same protocol.
- B. In the Easy Connect method, clients can connect to a database server by using the host name of the database with an optional port and service name.
- C. In the Easy Connect method, the listener port and IP address must be provided for the connection to be successful.
- D. The local naming method does not support connect-time failover and load-balancing options.
- E. The directory naming method supports connect-time failover and load-balancing option

Answer: ABE

NEW QUESTION 319

Which three database operations can be performed only at MOUNT state? (Choose three.)

- A. performing Flashback Database
- B. renaming control files
- C. enabling or disabling ARCHIVELOG mode
- D. re-creating control files
- E. performing full database recovery

Answer: ACE

NEW QUESTION 321

Automatic Shared Memory Management is enabled for your database instance. You notice that there are SQL statements performing poorly because of repeated parsing activity.

Which action generates recommendations to overcome the performance issues?

- A. running the Memory Advisor for the buffer cache
- B. running the Memory Advisor for the library cache
- C. running the Memory Advisor for the SGA
- D. running the Memory Advisor for the PGA

Answer: B

NEW QUESTION 323

Which files must you copy from the Oracle home of the database that is being upgraded to the new Oracle home for Oracle Database 12c? (Choose three.)

- A. the tnsnames.ora file
- B. the sqlnet.ora file
- C. the initialization parameter file
- D. the password file
- E. the listener.ora file

Answer: ABE

Explanation:

References https://docs.oracle.com/cd/E11882_01/server.112/e23633/afterup.htm#UPGRD52747

NEW QUESTION 326

Which four statements are true about the components of the Oracle Scheduler? (Choose four.)

- A. A schedule can be specified to a single job only.
- B. A scheduler job can point to a chain instead of pointing to a single program object.
- C. A job may get started automatically when a window opens.
- D. A program and job can be specified as part of a schedule definition.
- E. A job is specified as part of a program specification.
- F. A program can be used in the definition of multiple jobs.
- G. A program and schedule can be specified as part of a job definition.

Answer: BCFG

Explanation:

References: https://docs.oracle.com/cd/B19306_01/server.102/b14231/schedover.htm

NEW QUESTION 330

Which three are activities performed by SMON? (Choose three.)

- A. cleaning up the database buffer cache and freeing resources that a client process was using
- B. applying online redo during instance recovery
- C. cleaning up temporary segments that are no longer needed
- D. performing database services registration with the default listener
- E. restarting a server or a dispatcher process that terminated abnormally
- F. recovering failed transactions that were skipped during instance recovery because of file-read or tablespace offline errors

Answer: BCF

NEW QUESTION 332

The HR.DEPARTMENTS table is the parent of the HR.EMPLOYEES table. The EMPLOYEES.DEPARTMENT_ID column has a foreign key constraint with the ON DELETE CASCADE option that refers to the DEPARTMENTS.DEPARTMENT_ID column. An index exists on the DEPARTMENTS.DEPARTMENT_ID column. A transaction deletes a primary key in the DEPARTMENTS table, which has child rows in the EMPLOYEES table. Which statement is true?

- A. The transaction acquires a table lock only on the DEPARTMENTS table until the transaction is complete.
- B. The transaction acquires a table lock on the DEPARTMENTS tabl
- C. This lock enables other sessions to query but not update the DEPARTMENTS table until the transaction on the DEPARTMENTS table is complete.
- D. The transaction acquires a table lock on the EMPLOYEES tabl
- E. This lock enables other sessions to query but not update the EMPLOYEES table until the transaction on the DEPARTMENTS table is complete.
- F. Only the rows that are deleted in the DEPARTMENTS and EMPLOYEES tables are locked until the transactions on the DEPARTMENTS table is complete.

Answer: C

NEW QUESTION 334

Which two tasks must you perform before you begin the upgrade process to Oracle Database 12c? (Choose two.)

- A. Put all readonly tablespaces in read write mode
- B. Recompile all invalid objects
- C. Set the compatible parameter to 12.1.0.1
- D. Gather dictionary statistics
- E. Empty all user recycle bins

Answer: BE

NEW QUESTION 337

You plan to upgrade your Oracle Database 9i to Oracle Database 12c. Which two methods can you use? (Choose two.)

- A. Perform a rolling upgrade.
- B. Perform a direct upgrade by running the Database Upgrade Assistant (DBUA).
- C. Perform a direct upgrade by manually running the catctl.pl and catupgrd.sql scripts before issuing the STARTUPUPGRADE command.
- D. Install the Oracle Database 12c software, create a new Oracle 12c database, and then use the Oracle Data Pump to import data from the source Oracle 9i database to the target Oracle 12c database.
- E. Upgrade your current database to Oracle Database release 10.2.0.5, and then upgrade to Oracle Database 12c.

Answer: AE

NEW QUESTION 338

Which statement is true about using the Export/Import method for migrating data when upgrading to Oracle Database 12c?

- A. It automatically restarts a Data Pump Export or Import job after a failure is connected and the job continues from the point of failure.
- B. It can be used to migrate a database only if the source and target databases are hosted on the same endian format.
- C. It can be used to migrate a database only if the source database does not have any tablespace in read-only mode.
- D. It allows migration of a database directly over network link

Answer: D

NEW QUESTION 342

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

1Z0-062 Practice Exam Features:

- * 1Z0-062 Questions and Answers Updated Frequently
- * 1Z0-062 Practice Questions Verified by Expert Senior Certified Staff
- * 1Z0-062 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * 1Z0-062 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The 1Z0-062 Practice Test Here](#)