



Oracle

Exam Questions 1Z0-062

Oracle Database 12c: Installation and Administration

About ExamBible

[Your Partner of IT Exam](#)

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

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

Answer: A

Explanation:

REVOKE ..FROM

If the current container is the root:

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

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

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

NEW QUESTION 2

Which two are true concerning a multitenant container database with three pluggable database? (Choose two.)

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

Answer: CD

NEW QUESTION 3

Examine the following command: CREATE TABLE (prod_id number(4), Prod_name varchar2 (20), Category_id number(30), Quantity_on_hand number (3) INVISIBLE);

Which three statements are true about using an invisible column in the PRODUCTS table? (Choose three.)

- A. The %ROWTYPE attribute declarations in PL/SQL to access a row will not display the invisible column in the output.
- B. The DESCRIBE commands in SQL *Plus will not display the invisible column in the output.
- C. Referential integrity constraint cannot be set on the invisible column.
- D. The invisible column cannot be made visible and can only be marked as unused.
- E. A primary key constraint can be added on the invisible column.

Answer: ABE

Explanation:

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

* SELECT * FROM statements in SQL

* DESCRIBE commands in SQL*Plus

* %ROWTYPE attribute declarations in PL/SQL

* Describes in Oracle Call Interface (OCI) Incorrect: Not D: You can make invisible columns visible.

You can make a column invisible during table creation or when you add a column to a table, and you can later alter the table to make the same column visible.

NEW QUESTION 4

Which action takes place when a file checkpoint occurs?

- A. The checkpoint position is advanced in the checkpoint queue.
- B. All buffers for a checkpointed file that were modified before a specific SCN are written to disk by DBWn and the SCN is stored in the control file.
- C. The Database Writer process (DBWn) writes all dirty buffers in the buffer cache to data files.
- D. The Log Writer process (LGWR) writes all redo entries in the log buffer to online redo log file

Answer: B

NEW QUESTION 5

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 6

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 7

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

lsnrctl > 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 8

Examine the contents of SQL loader control file:

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

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

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

Answer: ABE

Explanation:

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

B (not D): Note:

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

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

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

- * A comma-separated values (CSV) (also sometimes called character-separated values, because the separator character does not have to be a comma) file stores tabular data (numbers and text) in plain-text form. Plain text means that the file is a sequence of characters, with no data that has to be interpreted instead, as binary numbers. A CSV file consists of any number of records, separated by line breaks of some kind; each record consists of fields, separated by some other character or string, most commonly a literal comma or tab. Usually, all records have an identical sequence of fields.
- * Fields with embedded commas must be quoted. Example:
1997,Ford,E350,"Super, luxurious truck" Note:
- * SQL*Loader is a bulk loader utility used for moving data from external files into the Oracle database.

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

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 10

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
OTHER__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 14

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 18

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 21

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 22

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 23

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

The following are the possible steps to accomplish this task:

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

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

Answer: C

Explanation:

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

Oracle White Paper - Upgrading to Oracle Database 12c -August 2013

NEW QUESTION 26

You plan to create a database by using the Database Configuration Assistant (DBCA), with the following specifications:

- Applications will connect to the database via a middle tier.
- The number of concurrent user connections will be high.
- The database will have mixed workload, with the execution of complex BI queries scheduled at night. Which DBCA option must you choose to create the database?

- A. a General Purpose database template with default memory allocation
- B. a Data Warehouse database template, with the dedicated server mode option and AMM enabled
- C. a General Purpose database template, with the shared server mode option and Automatic Memory Management (AMM) enabled
- D. a default database configuration

Answer: C

Explanation:

References:

<http://www.oracledistilled.com/oracle-database/administration/creating-a-database-using-database-configuration>

NEW QUESTION 27

You notice a performance change in your production Oracle database and you want to know which change has made this performance difference.

You generate the Compare Period Automatic Database Diagnostic Monitor (ADDM) report to further investigation. Which three findings would you get from the report? (Choose three.)

- A. It detects any configuration change that caused a performance difference in both time periods.
- B. It identifies any workload change that caused a performance difference in both time periods.
- C. It detects the top wait events causing performance degradation.
- D. It shows the resource usage for CPU, memory, and I/O in both time periods.
- E. It shows the difference in the size of memory pools in both time periods.
- F. It gives information about statistics collection in both time periods.

Answer: ABD

Explanation:

Keyword: shows the difference.

* Full ADDM analysis across two AWR snapshot periods Detects causes, measure effects, then correlates them Causes: workload changes, configuration changes Effects: regressed SQL, reach resource limits (CPU, I/O, memory, interconnect) Makes actionable recommendations along with quantified impact

* Identify what changed

/ Configuration changes, workload changes

* Performance degradation of the database occurs when your database was performing optimally in the past, such as 6 months ago, but has gradually degraded to a point where it becomes noticeable to the users. The Automatic Workload Repository (AWR) Compare Periods report enables you to compare database performance between two periods of time. While an AWR report shows AWR data between two snapshots (or two points in time), the AWR Compare Periods report shows the difference (ABE) between two periods (or two AWR reports with a total of four snapshots). Using the AWR Compare Periods report helps you to identify detailed performance attributes and configuration settings that differ between two time periods.

NEW QUESTION 31

You notice a performance change in your production Oracle 12c database. You want to know which change caused this performance difference. Which method or feature should you use?

- A. Compare Period ADDM report
- B. AWR Compare Period report
- C. Active Session History (ASH) report
- D. Taking a new snapshot and comparing it with a preserved snapshot

Answer: A

NEW QUESTION 35

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 38

The ORCL database is configured to support shared server mode. You want to ensure that a user connecting remotely to the database instance has a one-to-one ratio between client and server processes.

Which connection method guarantees that this requirement is met?

- A. connecting by using an external naming method
- B. connecting by using the easy connect method
- C. creating a service in the database by using the DBMS_SERVICE.CREATE_SERVICE procedure and using this service for creating a local naming service
- D. connecting by using the local naming method with the SERVER = DEDICATED parameter set in the tnsnames.ora file for the net service
- E. connecting by using a directory naming method

Answer: D

NEW QUESTION 40

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 ADVM 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 45

You upgrade your Oracle database in a multiprocessor environment. As a recommended you execute the following script: SQL > @utlpr.sql
Which two actions does the script perform? (Choose two.)

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

Answer: CE

Explanation:

utlpr.sql and utlprp.sql

The utlpr.sql and utlprp.sql scripts are provided by Oracle to recompile all invalid objects in the database. They are typically run after major database changes such as upgrades or patches. They are located in the

\$ORACLE_HOME/rdbms/admin directory and provide a wrapper on the UTL_RECOMP package. The utlpr.sql script simply calls the utlprp.sql script with a command line parameter of "0". The utlprp.sql accepts a single integer parameter that indicates the level of parallelism as follows.

0 - The level of parallelism is derived based on the CPU_COUNT parameter. 1 - The recompilation is run serially, one object at a time.

N - The recompilation is run in parallel with "N" number of threads.

Both scripts must be run as the SYS user, or another user with SYSDBA, to work correctly. References:

NEW QUESTION 50

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 51

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 55

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 58

Which two statements are true about variable extent size support for large ASM files? (Choose two.)

- A. The metadata used to track extents in SGA is reduced.
- B. Rebalance operations are completed faster than with a fixed extent size
- C. An ASM Instance automatically allocates an appropriate extent size.
- D. Resync operations are completed faster when a disk comes online after being taken offline.
- E. Performance improves in a stretch cluster configuration by reading from a local copy of an extent.

Answer: AC

Explanation:

A: Variable size extents enable support for larger ASM datafiles, reduce SGA memory requirements for very large databases (A), and improve performance for file create and open operations.

C: You don't have to worry about the sizes; the ASM instance automatically allocates the appropriate extent size. Note:

- * The contents of ASM files are stored in a disk group as a set, or collection, of data extents that are stored on individual disks within disk groups. Each extent resides on an individual disk. Extents consist of one or more allocation units (AU). To accommodate increasingly larger files, ASM uses variable size extents.
- * The size of the extent map that defines a file can be smaller by a factor of 64 depending on the file size. The initial extent size is equal to the allocation unit size and it increases by a factor of 64 at predefined thresholds. This feature is automatic for newly created and resized datafiles when the disk group compatibility attributes are set to Oracle Release 11 or higher.

NEW QUESTION 60

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

On further investigation, you find that buffer busy waits is caused by contention on data blocks. Which option would you consider first to decrease the wait event immediately?

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

Answer: D

Explanation:

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

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

Note:

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

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

NEW QUESTION 61

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 63

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

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

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

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

Answer: D

Explanation:

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

* To redefine a table online:

Choose the redefinition method: by key or by rowid

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

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

You cannot use this method on index-organized tables. Note:

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

Incorrect:

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

NEW QUESTION 64

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 67

Examine the following query output:

```
SQL> SELECT name, force_logging FROM v$database;
```

NAME	FORCE_LOGGING
PROD	NO

You issue the following command to import tables into the hr schema:

```
$ > impdp hr/hr directory = dumpdir dumpfile = hr_new.dmp schemas=hr TRANSFORM=DISABLE_ARCHIVE_LOGGING: Y
```

Which statement is true?

- A. All database operations performed by the impdp command are logged.
- B. Only CREATE INDEX and CREATE TABLE statements generated by the import are logged.
- C. Only CREATE TABLE and ALTER TABLE statements generated by the import are logged.
- D. None of the operations against the master table used by Oracle Data Pump to coordinate its activities are logged.

Answer: C

Explanation:

Oracle Data Pump disable redo logging when loading data into tables and when creating indexes.

The new TRANSFORM option introduced in data pumps import provides the flexibility to turn off the redo generation for the objects during the course of import.

The Master Table is used to track the detailed progress information of a Data Pump job.

The Master Table is created in the schema of the current user running the Pump Dump export or import, and it keeps tracks of lots of detailed information.

NEW QUESTION 69

You execute the following commands to audit database activities:

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

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

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

- A. One audit record is created for every successful execution of a SELECT, INSERT OR DELETE command on a table, and contains the SQL text for the SQL Statements.

- B. One audit record is created for every successful execution of a SELECT, INSERT OR DELETE command, and contains the execution plan for the SQL statements.
- C. One audit record is created for the whole session if john successfully executes a SELECT, INSERT, or DELETE command, and contains the execution plan for the SQL statements.
- D. One audit record is created for the whole session if JOHN successfully executes a select command, and contains the SQL text and bind variables used.
- E. One audit record is created for the whole session if john successfully executes a SELECT, INSERT, or DELETE command on a table, and contains the execution plan, SQL text, and bind variables used.

Answer: A

Explanation:

Note:

* BY SESSION

In earlier releases, BY SESSION caused the database to write a single record for all SQL statements or operations of the same type executed on the same schema objects in the same session. Beginning with this release (11g) of Oracle Database, both BY SESSION and BY ACCESS cause Oracle Database to write one audit record for each audited statement and operation.

* BY ACCESS

Specify BY ACCESS if you want Oracle Database to write one record for each audited statement and operation. Note:

If you specify either a SQL statement shortcut or a system privilege that audits a data definition language (DDL) statement, then the database always audits by access. In all other cases, the database honors the BY SESSION or BY ACCESS specification.

* For each audited operation, Oracle Database produces an audit record containing this information:

/ The user performing the operation

/ The type of operation

/ The object involved in the operation

/ The date and time of the operation

References:

NEW QUESTION 70

Which four statements are true about database instance behavior? (Choose four.)

- A. An idle instance is created when a STARTUP NOMOUNT is successful
- B. All dynamic performance views (v\$ views) return data when queried from a session connected to an instance in NOMOUNT state
- C. The consistency of redo logs and data files is checked when mounting the database
- D. Redo log files can be renamed in MOUNT state
- E. An SPFILE can be updated when connected to an idle instance
- F. Datafiles can be renamed in MOUNT state

Answer: CDEF

NEW QUESTION 72

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 76

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

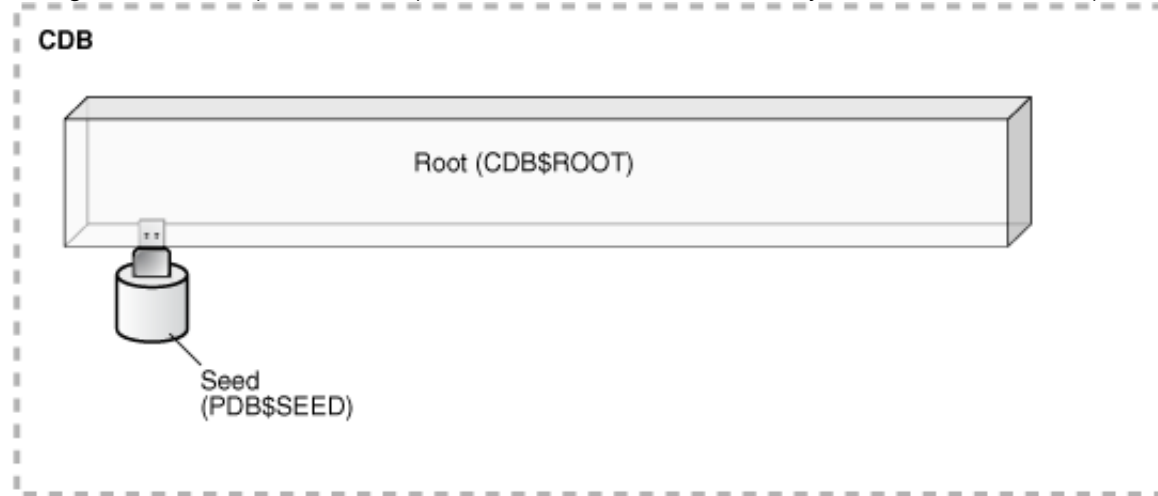
- A. It will create a multitenant container database (CDB) with only the root opened.
- B. It will create a CDB with root opened and seed read only.
- C. It will create a CDB with root and seed opened and one PDB mounted.
- D. It will create a CDB that must be plugged into an existing CDB.
- E. It will create a CDB with root opened and seed mounted.

Answer: B

Explanation:

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

Along with the root (CDB\$ROOT), Oracle Database automatically creates a seed PDB (PDB\$SEED). The following graphic shows a newly created CDB:



* Creating a PDB

Rather than constructing the data dictionary tables that define an empty PDB from scratch, and then populating its Obj\$ and Dependency\$ tables, the empty PDB is created when the CDB is created. (Here, we use empty to mean containing no customer-created artifacts.) It is referred to as the seed PDB and has the name PDB\$Seed. Every CDB non-negotiably contains a seed PDB; it is non-negotiably always open in read-only mode. This has no conceptual significance; rather, it is just an optimization device. The create PDB operation is implemented as a special case of the clone PDB operation.

NEW QUESTION 80

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 83

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 84

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 85

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 88

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 93

Which two must be installed or configured either manually or by DBCA in order to use Enterprise Manager Database Express (EM Express)? (Choose two.)

- A. A port number for Oracle HTTP Server must be configured
- B. The APEX_PUBLIC_USER role must be granted to SYSMAN
- C. A SYSMAN user with SYSDBA privilege must be created
- D. At least one TCP/IP dispatcher must be configured
- E. The Oracle HTTP Server must be installed

Answer: BD

NEW QUESTION 95

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 100

Which four actions are possible during an Online Data file Move operation? (Choose four.)

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

Answer: ACEF

Explanation:

- You can now move On line Datafile without have to stop Monoged Recovery and manually copy and rename Files. This can even be used to move Datafiles from or to ASM.
- New in Oracle Database 12c: FROM METAUNK. Physical Standby Database is in Active Data Guard Mode (opened READ ONLY and Managed Recovery is running):
It is now possible to online move a Datafile while Managed Recovery is running, ie. the Physical Standby Database is in Active Data Guard Mode. You con use this Command to move the Datafile
- A flashback operation does not relocate a moved data file to its previous location. If you move a data file online from one location to another and later flash back the database to a point in time before the move, then the Data file remains in the new location, but the contents of the Data file ore changed to the contents at the time specified in the flashback. Oracle0 Database Administrator's Guide 12c Release 1 (12.1)

NEW QUESTION 103

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

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

Answer: C

NEW QUESTION 104

You want to schedule a job to rebuild a table's indexes after a bulk insert, which must be scheduled as soon as a file containing data arrives on the system. What would you do to accomplish this?

- A. Create a file watcher and an event-based job for bulk insert and then create another job to rebuild indexes on the table.
- B. Create a file watcher for the bulk inserts and then create a job to rebuild indexes.
- C. Create a job array and add a job for bulk insert and a job to rebuild indexes to the job array.
- D. Create an event-based job for the file arrival event, then create a job for bulk insert, and then create a job to rebuild indexes.

Answer: A

NEW QUESTION 109

You plan to implement the distributed database system in your company. You invoke Database Configuration Assistant (DBCA) to create a database on the server. During the installation, DBCA prompts you to specify the Global Database Name. What must this name be made up of?

- A. It must be made up of a database name and a domain name.
- B. It must be made up of the value in ORACLE_SID and HOSTNAME.

- C. It must be made up of the value that you plan to assign for INSTANCE_NAME and HOSTNAME.
D. It must be made up of the value that you plan to assign for ORACLE_SID and SERVICE_NAMES.

Answer: A

Explanation:

Using the DBCA to Create a Database (continued)

3. Database Identification: Enter the Global Database Name in The form database_name.domain_name, and the system identifier (SID). The SID defaults to the database name and uniquely identifies the instance associated with the database.

4. Management Options: Use this page to set up your database so that it can be managed with Oracle Enterprise Manager. Select the default: "Configure the Database with Enterprise Manager." Optionally, this page allows you to configure alert notifications and daily disk backup area settings.

Note: You must configure the listener before you can configure Enterprise Manager (as shown earlier).

NEW QUESTION 113

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 115

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

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

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

Which two statements are true about redefining the table? (Choose two.)

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

Answer: BC

Explanation:

C (not D): CONS_VPD_AUTO

Used to indicate to copy VPD policies automatically

* DBMS_RLS.ADD_POLICY

/ The DBMS_RLS package contains the fine-grained access control administrative interface, which is used to implement Virtual Private Database (VPD). DBMS_RLS is available with the Enterprise Edition only.

Note:

* CONS_USE_PK and CONS_USE_ROWID are constants used as input to the "options_flag" parameter in both the START_REDEF_TABLE Procedure and CAN_REDEF_TABLE Procedure. CONS_USE_ROWID is used to indicate that the redefinition should be done using rowids while CONS_USE_PK implies that the redefinition should be done using primary keys or pseudo-primary keys (which are unique keys with all component columns having NOT NULL constraints).

* DBMS_REDEFINITION.START_REDEF_TABLE

To achieve online redefinition, incrementally maintainable local materialized views are used. These logs keep track of the changes to the master tables and are used by the materialized views during refresh synchronization.

* START_REDEF_TABLE Procedure

Prior to calling this procedure, you must manually create an empty interim table (in the same schema as the table to be redefined) with the desired attributes of the post-redefinition table, and then call this procedure to initiate the redefinition.

NEW QUESTION 116

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 120

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

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

Answer: ADE

NEW QUESTION 125

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 126

Examine the parameters for a database instance:

NAME	TYPE	VALUE
temp_undo_enabled	boolean	TRUE
undo_management	string	AUTO
undo_retention	integer	900
undo_tablespace	string	UNDOTBS1

Which two statements are true? (Choose two.)

- A. Undo records for temporary tables are stored in a temporary tablespace.
- B. Undo records for temporary tables are stored in the undo tablespace and logged in the redo.
- C. Undo records for temporary tables are stored in the undo tablespace and logged in the redo only for those sessions where temporary undo is enabled.
- D. No redo is generated for the undo records belonging to temporary tables.
- E. No redo and undo records are generated for temporary table

Answer: AD

NEW QUESTION 128

You execute this command:

```
SQL> CREATE TABLESPACE lmtbsb DATAFILE '/u02/oracle/data/lmtbsb01.dbf' SIZE 50M
EXTENT MANAGEMENT LOCAL;
```

Which two statements are true about segment space management for segments in this tablespace? (Choose two.)

- A. Space utilization inside segments is mapped by bitmaps.
- B. Segments are automatically shrunk and compressed when rows are deleted.
- C. The PCTFREE storage parameter has no effect on segments created in this tablespace.

D. The PCTUSED storage parameter has no effect on segments created in this tablespace

Answer: AD

NEW QUESTION 129

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 134

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 pla

Answer: D

NEW QUESTION 136

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

How would you accomplish this?

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

Answer: B

NEW QUESTION 141

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

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

USERNAME	SYSDB	SYSOP	SYSAS	SYSBA	SYSDBG	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 142

In your database, USERS is the default permanent tablespace. Examine the commands and their outcome:

```
SQL> CREATE USER user02 identified by us123 QUOTA 10M ON users;
User created.

SQL> GRANT create session, sysdba TO user02;
Grant succeeded.
```

You plan to execute the commands:

```
SQL> CONN user02/us123 AS SYSDBA
SQL> CREATE TABLE mytab (id number, lname varchar2(20));
```

Which two statements are true? (Choose two.)

- A. The MYTAB table is created in the SYSTEM tablespace but no rows can be inserted into the table by USER02.
- B. The MYTAB table is created in the SYSTEM tablespace and rows can be inserted into the table by USER02.
- C. The MYTAB table is created in the USERS tablespace but no rows can be inserted into the table by USER02.
- D. The CREATE TABLE statement generates an error because the SYSDBA privilege does not provide any space quota on the SYSTEM tablespace by default.
- E. The MYTAB table is owned by the SYS use

Answer: BE

NEW QUESTION 146

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 149

In your Oracle 12c database, you plan to execute the command:

```
SQL> CREATE TABLESPACE tbs1 DATAFILE '/u02/oracle/data/tbs01.dbf' SIZE 50M; The u02 file system has 1 GB of free space available.
```

What is the outcome?

- A. It creates a locally managed tablespace with manual segment space management enabled.
- B. It raises an error because extent management is not specified.
- C. It creates a locally managed tablespace with automatic segment space management enabled.
- D. It creates a dictionary-managed tablespace with manual segment space management enabled.

Answer: C

Explanation:

References: https://docs.oracle.com/cd/B28359_01/server.111/b28310/tspaces002.htm#ADMIN11359

NEW QUESTION 150

Which three statements are true about the purpose of checkpoints? (Choose three.)

- A. They ensure that uncommitted transactions are rolled back in case of an instance failure.
- B. They ensure that all the dirty buffers are written to disk during a normal shutdown.
- C. They ensure that instance recovery starts in the event of an instance failure.
- D. They ensure that dirty buffers in the buffer cache are written to disk regularly.
- E. They reduce the time required for recovery in case of an instance failur

Answer: BDE

NEW QUESTION 152

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 utluiobj.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 vie

Answer:

D

NEW QUESTION 157

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 161

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

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

Answer: BCF

Explanation:

* Multiprocess and Multithreaded Oracle Database Systems

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

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

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

NEW QUESTION 162

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 164

Examine the following command;

ALTER SYSTEM SET enable_ddl_logging = TRUE; Which statement is true?

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

Answer: E

NEW QUESTION 166

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 169

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

The characteristics of the non-CDB are as follows:

- Version: Oracle Database 12c Releases 1 64-bit
- Character set: WE8ISO8859P15
- National character set: AL16UTF16
- O/S: Oracle Linux6 64-bit

The characteristics of the CDB are as follows:

- Version: Oracle Database 12c Release 1 64-bit
- Character set: AL32UTF8
- O/S: Oracle Linux 6 64-bit

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

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

Answer: C

NEW QUESTION 170

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

SQL > CREATE INDEX oe.ord_customer_ix1 ON oe.orders (customers_id, sales_rep_id) INVISIBLE; SQL> CREATE BITMAP INDEX oe.ord_customer_ix2 ON oe.orders (customers_id, sales_rep_id); Which two statements are correct? (Choose two.)

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

Answer: AE

Explanation:

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

NEW QUESTION 173

A database instance is started by using an SPFILE. The database is configured in ARCHIVELOG mode and the control file autobackup is configured. Daily full database backups are performed by using RMAN.

You lost all control files due to media failure.

Given the steps to recover from the error in random order:

1. Shut down the instance, if it is not already down.
2. Restore the control file from autobackup to a new location.
3. Start the database instance to NOMOUNT state.
4. Recover the database to the point of failure of the control file.
5. Open the database with the RESETLOGS option.
6. Mount the database.
7. Update the SPFILE with the new location of the control file by using the ALTER SYSTEM command. Identify the correct sequence of the required steps.

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

Answer: A

NEW QUESTION 178

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 180

Examine these two statements:

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

Tablespace created.

SQL> ALTER DATABASE DEFAULT TABLESPACE MRKT;

Database altered.
```

Which three are true about the MRKT tablespace? (Choose three.)

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

Answer: CDE

NEW QUESTION 183

Which three statements are true about the Pre-Upgrade Information Tool? (Choose three.)

- A. It generates a script to recompile invalid objects post-upgrade.
- B. The preupgrade_fixups.sql script is created to list and describe issues in the source database.
- C. A log file, preupgrade.log, is created that contains the output of the Pre-Upgrade Information tool.
- D. It checks for required tablespaces and if they are not available, creates them automatically.
- E. The preupgrade_fixups.sql script is executed automatically to fix issues in the source database.
- F. The postupgrade_fixups.sql script is created to address issues that can be fixed after a database has been upgraded.

Answer: ACE

Explanation:

References <https://docs.oracle.com/database/122/UPGRD/using-preupgrade-information-tool-for-oracle-database.htm#UPG>

NEW QUESTION 187

Which three activities are supported by the Data Recovery Advisor? (Choose three.)

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

Answer: ABE

Explanation:

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

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

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

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

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

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

NEW QUESTION 190

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 193

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

A. Running a query on a table returns ORA-600: Internal Error.

B. Inserting a value into a table returns ORA-01722: invalid number.

C. Creating a table returns ORA-00955: name us already in used by an existing objects.

D. Inserting a value into a table returns ORA-00001: unique constraint (SYS.OK_TECHP) violated.

E. Rebuilding an index using ALTER INDEX . . . REBUILD fails with an ORA-01578: ORACLE data block corrupted (file # 14, block # 50) error.

Answer: AE

Explanation:

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

* All internal errors (ORA-600), block corruption errors (ORA-1578), and deadlock errors (ORA-60) that occur

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

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

* Errors occurring during the automatic refresh of a materialized view

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

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

NEW QUESTION 197

Which statement is true about profiles?

A. Resource limits specified in a profile assigned to a user are always enabled.

B. A user can exist without any profile.

C. A profile can be assigned only to one user.

D. Password management using profiles is always enable

Answer: D

NEW QUESTION 198

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 201

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 205

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 209

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 213

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 215

What is the outcome of the SHUTDOWN ABORT command?

- A. Pending transactions are committed and the database is closed.
- B. Dirty buffers in the buffer cache and unwritten redo are not written to the data files and redo log files respectively.
- C. Uncommitted transactions are rolled back
- D. Instance recovery must be requested by the DBA at the next startup

Answer: B

NEW QUESTION 217

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 222

Which three statements are true about server-generated alerts? (Choose three.)

- A. Server-generated alerts notify administrators of problems that cannot be resolved automatically.
- B. Alerts are not issued for locally managed read-only tablespaces.
- C. Response actions cannot be specified for server-generated alerts.
- D. Stateful alerts can be queried only from the DBA_ALERT_HISTORY view.
- E. When an alert is cleared, it is moved to the DBA_ALERT_HISTORY view.

Answer: ABE

Explanation:

References https://docs.oracle.com/cd/B28359_01/server.111/b28310/schema001.htm#ADMIN10120

NEW QUESTION 224

In your database, the STATISTICS_LEVEL parameter is set to TYPICAL and an Automatic Workload Repository (AWR) snapshot is taken every 30 minutes.

Which two statements are true about the Automatic Database Diagnostic Monitor (ADDM)? (Choose two.)

- A. It measures database performance by analyzing the wait time and CPU time of all non-idle user sessions.
- B. It always compares the latest snapshot with the baseline snapshot for analysis.
- C. It runs after each AWR snapshot is created and it requires at least two snapshots for analysis.
- D. It requires at least four AWR snapshots for analysis.
- E. It calls other advisors if required, but does not provide recommendations about the advisors to be use

Answer: AC

NEW QUESTION 228

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 229

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 233

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 236

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 239

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 244

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 247

The HR schema exists in two databases, BOSTON and DENVER, and has the same password, HR. You have the CREATE DATABASE LINK and CREATE SESSION privileges on both the database. BOSTON is defined as a service name in the tnsnames.ora of both the databases.

You plan to use the command:

```
CREATE DATABASE LINK hr_link CONNECT to hr IDENTIFIED BY hr USING 'denver';
```

What must be done to ensure only the HR user in the BOSTON database can access the HR schema in the DENVER database?

- A. Execute this command as HR user in the BOSTON database and SYS user in the DENVER database.
- B. Execute this command as SYS user in both the databases.
- C. Execute this command as HR user in the DENVER database.
- D. Execute this command as HR user in the BOSTON database

Answer: D

NEW QUESTION 252

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

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

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

Answer: D

NEW QUESTION 253

What must you use to read data from a table in your database and write it to an external table?

- A. Use SQL* Loader conventional path load.
- B. Use SQL* Loader direct path load.
- C. Use CREATE TABLE
- D. . ORGANIZATION EXTERNAL command with ORACLE_LOADER access driver.
- E. Use CREATE TABLE
- F. . ORGANIZATION EXTERNAL command with ORACLE_DATAPUMP access driver

Answer: D

NEW QUESTION 258

Which two services may you see on the My Service Dashboard page? (Choose two.)

- A. Network Cloud Service
- B. User Cloud Service
- C. Compute Cloud Service
- D. Database Cloud Service

Answer: CD

NEW QUESTION 259

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 264

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 266

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 270

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 274

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 277

Which three statements are true about Enterprise Manager Database Express? (Choose three.)

- A. It can be used to perform database backup operations.
- B. It can use the HTTP protocol.
- C. The same port number is used for multiple Database Express configurations on the same host.
- D. It can use the HTTPS protocol.
- E. It is available only when the database is open.

Answer: BDE

NEW QUESTION 278

Which two statements are true about using SQL*Loader? (Choose two.)

- A. It can load data from external files by using the direct path only.
- B. It can load data into multiple tables using the same load statement.
- C. It can load data into only one table at a time.
- D. It can generate unique sequential key values in specified columns.
- E. It can load data from external files by using the conventional path only.

Answer: AC

NEW QUESTION 283

Your database is configured in ARCHIVELOG mode, and a daily full database backup is taken by using RMAN. Control file autobackup is configured. Loss of which three database files can lead to an incomplete recovery? (Choose three.)

- A. inactive online redo log file group
- B. a data file belonging to the default temporary tablespace
- C. a data file belonging to the SYSAUX tablespace
- D. server parameter file (SPFILE)
- E. active online redo log file group

F. all the control flies

Answer: AEF

NEW QUESTION 288

You want to load data from a large file into your database without causing an overhead on the SGA. Which tool would you use.

- A. external table
- B. Oracle data Pump
- C. SQL*Loader with a direct data path
- D. SQL*Loader with a conventional data path
- E. Enterprise Manager Database Express

Answer: C

Explanation:

References: https://docs.oracle.com/cd/B19306_01/server.102/b14215/ldr_modes.htm#i1007501

NEW QUESTION 289

You want to create a locally managed tablespace called NEWTBS to store segments with different extent sizes. Which set of tablespace attributes can be specified for a tablespace that satisfies the requirements?

- A. EXTENT MANAGEMENT LOCAL STORAGE (INITIAL 5M MAXSIZE 10M)
- B. REUSE AUTOEXTEND ON MAXSIZE UNLIMITED
- C. EXTENT MANAGEMENT LOCAL SEGMENT SPACE MANAGEMENT UNIFORM
- D. EXTENT MANAGEMENT LOCAL AUTOALLOCATE

Answer: D

NEW QUESTION 290

Examine the following command:

SQL> DBMS_STATS. SET_TABLE_PREFS ('SH', 'CUSTOMERS', 'PUBLISH', 'false'); What is the effect of executing this command?

- A. Existing statistics for the CUSTOMERS table become unusable for the query optimizer.
- B. Automatic statistics collection is stopped for the CUSTOMERS table.
- C. Statistics for the CUSTOMERS table are locked and cannot be overwritten.
- D. Statistics subsequently gathered on the CUSTOMERS table are stored as pending statistic

Answer: D

NEW QUESTION 291

Which four operations performed after the Oracle Restart installation are automatically added to the Oracle Restart configuration? (Choose four.)

- A. listener configured by using NETCA
- B. database service created by using SRVCTL
- C. database created by using a SQL statement
- D. database created by using DBCA
- E. ASM instance created by using ASMCA
- F. database service created by using DBMS_SERVICE.CREATE_SERVICE
- G. database service created by modifying the SERVICE_NAMES initialization parameter

Answer: ABDE

Explanation:

References https://docs.oracle.com/cd/E18283_01/server.112/e17120/restart002.htm#insertedID3

NEW QUESTION 296

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 300

You want to import the schema objects of the HR user from the development database DEVDB to the production database PRODDB by using Oracle Data Pump. A database link devdb.us.oracle.com is created between PRODDB and DEVDB.

You execute the following command on the PRODDB database server:

```
$ impdp system/manager directory = DB_DATA
  dumpfile = schemas.dat
  schemas = hr
  flashback_time = "TO_TIMESTAMP ('05-01-2012 14:35:00', 'DD-MM-
YYYY HH24:MI:SS') "
```

The command fails, displaying the following error:

```
ORA-39001: invalid argument value
ORA-39000: bad dump file specification
ORA-31640: unable to open dump file "/home/oracle/schema/schemas.
dat" for read
ORA-27037: unable to obtain file status
```

What should you do to resolve the error?

- A. Add network_link = devdb.us.oracle.com.
- B. Add the SYSTEM user to the schemas option.
- C. Change the dumpfile option value to schema.dat@devdb.us.oracle.com.
- D. Replace the schemas option with network_link = devdb.us.oracle.com.
- E. Replace the dumpfile option with network_link = devdb.us.oracle.co

Answer: E

NEW QUESTION 303

Identify three situations in which messages are written to the alert log file. (Choose three.)

- A. Rebuilding an index using ALTER INDEX . . . REBUILD fails with an ORA-01578: ORACLE data block corrupted (file # 14, block # 50)"
- B. Creating a table returns "ORA-00955: name is already in used by an existing object"
- C. Inserting a value into a table returns "ORA-01722: invalid number"
- D. Updating a record in a table returns "ORA-00060: deadlock detected while waiting for resource"
- E. Inserting a value into a table returns "ORA-00001: unique constraint (SYS.PK_XXXX) violated"
- F. Running a query on a table returns "ORA-01578: ORACLE data block corrupted (file # 4, block # 131)"

Answer: ADF

NEW QUESTION 304

For which three requirements would you use the Database Resource Manager? (Choose three.)

- A. specifying an idle time limit that applies to sessions that are idle and blocking other sessions
- B. limiting the degree of parallelism of operations performed by user sessions in a consumer group
- C. specifying the maximum number of concurrent sessions allowed for a user
- D. limiting the CPU used per database call
- E. specifying the amount of private space a session can allocate in the shared pool of the SGA.

Answer: ABC

Explanation:

References:

http://docs.oracle.com/cd/B19306_01/server.102/b14231/dbrm.htm

NEW QUESTION 308

Which three statements are true about checkpointing? (Choose three.)

- A. It prompts the Checkpoint (CKPT) process to write data to the data files and redo information to the online redo log files.
- B. It ensures that all dirty buffers are written to data files during consistent shutdown.
- C. It reduces the time required for recovery in case of an instance failure.
- D. Frequent thread checkpoints can degrade database performance.
- E. It prompts the Database Writer (DBWn) process to write checkpoint information into data file headers and the control file.

Answer: BCD

NEW QUESTION 312

Which component resides in the System Global Area (SGA) of a database instance only in shared server connections?

- A. User Global Area
- B. Program Global Area
- C. SQL Query Result Cache
- D. PL/SQL Function Result Cache

Answer: A

NEW QUESTION 316

Which statement is true about the Oracle central inventory directory (oraInventory)?

- A. oraInventory must not be shared by all Oracle software installations on a single system.
- B. If ORACLE_BASE is set to /u01/app/oracle for the oracle user during an installation, OUI creates the Oracle Inventory directory in the /u01/app/oracle/oraInventory path.
- C. If an OFA-compliant path is not created and the ORACLE_BASE environment variable is not set during an Oracle Database installation, the Oracle Inventory directory is placed in the home directory of the user that is performing the installation.
- D. Oracle software owners must be members of the same central oraInventory group, but they need not have this group as their primary group.

Answer: D

Explanation:

References <https://docs.oracle.com/database/121/CWLIN/usrgtps.htm#CWLIN483>

NEW QUESTION 317

You want to create a database with a block size other than the default 8 kilobytes (KB) by using the Database Configuration Assistant (DBCA). Which option should you use?

- A. Automatic Storage Management (ASM) for storage of data files
- B. a file system for storage of data files
- C. a Data Warehouse database template
- D. a custom database template

Answer: D

NEW QUESTION 318

You executed the following query:

```
SELECT oldest_flashback_scn, oldest_flashback_time FROM V$FLASHBACK_DATABASE_LOG;
```

Considering that all the redo logs are available, what information can you derive from the output of the preceding query?

- A. The time when the last flashback operation in your database was performed
- B. The time when the first flashback operation in our database was performed
- C. The approximate time and the lowest system change number (SCN) to which you can flash back your database
- D. The system change number (SCN) and the time when the Flashback Database was enabled in the database instance

Answer: C

NEW QUESTION 320

Your database is in ARCHIVELOG mode and you want to automate the backup scheduling for your database. Which two tools or utilities would you use to achieve this? (Choose two.)

- A. Oracle Enterprise Manager Database Express (EM Express)
- B. Oracle Enterprise Manager Cloud Control
- C. Database Configuration Assistant (DBCA)
- D. Recovery Manager (RMAN) script invoked by using scheduler

Answer: BD

NEW QUESTION 325

Which two statements are true about availability audit features after migration to unified auditing? (Choose two.)

- A. The ability of users to audit their own schema objects is not available in the post-migrated database.
- B. Operating system audit trail is available in the post-migrated database.
- C. Network auditing is available in the post-migrated database.
- D. Mandatory auditing of audit administrative actions is available in the post-migrated database.

Answer: AD

Explanation:

References: https://docs.oracle.com/database/121/DBSEG/audit_changes.htm#DBSEG341

NEW QUESTION 330

Which background process does Automatic Shared Memory Management use to coordinate the sizing of memory components?

- A. PMON
- B. SMON
- C. MMNL
- D. MMAN
- E. MMON

Answer: D

NEW QUESTION 334

One of your databases has archive logging enabled and RMAN backups are taken at regular intervals. The data file for the USERS tablespace is corrupt. Which command must you execute before starting the recovery of this tablespace?

- A. STARTUP FORCE
- B. ALTER TABLESPACE users OFFLINE IMMEDIATE;
- C. SWITCH DATAFILE ALL;
- D. ALTER TABLESPACE users OFFLINE NORMAL;
- E. ALTER TABLESPACE users OFFLINE TEMPORARY;

Answer: E

NEW QUESTION 337

Because of a logical corruption in the EMPLOYES tables, you want to perform Tablespace Point-in-Time Recovery (TSPITR) to recover the table. Before you started the TSPITR process, you queried the TS_PITR_CHECK view and you realized that the table has a referential constraint with DEPARTMENTS that exists in another tablespace, MASTERTBS. Which two actions will permit the TSPITR to work? (Choose two.)

- A. Taking the MASTERTBS tablespace offline
- B. Dropping the relationship between the tables
- C. Adding the MASTERTBS tablespace to the recovery set
- D. Putting the MASTERTBS tablespace in read-only mode

Answer: BC

Explanation:

http://docs.oracle.com/cd/E11882_01/backup.112/e10642/rcmtspit.htm#BRADV99978

If constraints for the tables in tablespace tbs1 are contained in the tablespace tbs2, then you cannot recover tbs1 without also recovering tbs2.

NEW QUESTION 342

One of your databases supports an OLTP workload. The default UNDO tablespace is fixed size with:

- 1. RETENTION NOGUARANTEE
- 2. UNDO_RETENTION is 12 minutes

User SCOTT gets this error after a query on the SALES table has run for more than 15 minutes: ORA-01555: snapshot too old

Which three factors taken separately or in some combination might be the cause? (Choose three.)

- A. An update was made to the SALES table after the query began
- B. An update to the SALES table was committed after the query began
- C. A committed delete to the SALES table was made more than 12 minutes before the query began
- D. An uncommitted update to the SALES table was made more than 12 minutes before the query began
- E. A committed update to the SALES table was made more than 12 minutes before the query began
- F. An uncommitted delete to the SALES table was made more than 12 minutes before the query began
- G. An update was made to the SALES table before the query began

Answer: ADE

NEW QUESTION 347

You are managing an Oracle Database 12c database. The database is open, and you plan to perform Recovery Manager (RMAN) backups.

Which three statements are true about these backups? (Choose three.)

- A. The backups would be consistent.
- B. The backups would be possible only if the database is running in ARCHIVELOG mode.
- C. The backups need to be restored and the database has to be recovered in case of a media failure.
- D. The backups would be inconsistent.
- E. The backups by default consist of all the data blocks within the chosen files or the full databas

Answer: BCD

NEW QUESTION 351

When does a database checkpoint occur?

- A. When there is an online redo log switch.
- B. When a user session terminates abnormally.
- C. When a server process terminates abnormally.
- D. When the SHUTDOWN ABORT command is issue

Answer: A

NEW QUESTION 356

Which statement is true about the Database as a Service (DBaaS) instances and Database instances in Oracle Public Cloud

- A. An Oracle database instance can support only one DBaaS instance.
- B. ADBaaS instance can support only one Oracle database instance.
- C. An Oracle database instance can support multiple DBaaS instances.
- D. ADBaaS instance can support multiple Oracle database instances.
- E. ADBaaS instance runs in a pluggable database (PDB), which is contained in a multi-tenant container database (CDB).

Answer: D

NEW QUESTION 358

In your Oracle 12c database, you invoke SQL *Loader Express Mode command to load data: \$> sqlldr hr/hr table=employees

Which two statements are true about this command? (Choose two.)

- A. It succeeds and creates the EMPLOYEES table in the HR schema if the table does not exist.
- B. It fails because the SQL *Loader control file location is not specified.
- C. It fails because the SQL *Loader data file location is not specified.
- D. It succeeds with default settings if the EMPLOYEES table belonging to the HR schema is already defined in the database.
- E. It succeeds even if the HR user does not have the CREATE DIRECTORY privileg

Answer: DE

NEW QUESTION 361

Which statement is true about using the Database Upgrade Assistant (DBUA) to upgrade your database from Oracle Database 11g to Oracle Database 12c?

- A. It terminates if the SYSTEM tablespace in the source database is not autoextensible.
- B. It automatically makes necessary changes to Oracle environment variables.
- C. It automatically enables unified auditing in the upgraded database.
- D. It automatically adds new data files if there is not enough disk space to grow.

Answer: D

Explanation:

References: https://docs.oracle.com/cd/E18283_01/server.112/e17222/upgrade.htm#insertedID5

NEW QUESTION 365

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

- A. A resource plan change can be automated by using the Oracle Scheduler.
- B. It can be used to control the consumption of only physical I/Os where excessive physical I/Os can trigger an automatic session termination but excessive logical I/Os cannot.
- C. It can be used to control the usage of the undo tablespace by consumer groups.
- D. A resource plan can have multiple resource plan directives, each of which controls resource allocation for a different consumer group.
- E. It can be used to enable resumable timeout for user sessions.
- F. It can be used to control the usage of the temp tablespace by consumer group

Answer: ACD

NEW QUESTION 368

Which three statements are true about Oracle checkpoint processing? (Choose three.)

- A. Frequent thread checkpoints can degrade database performance
- B. Database Writer (DBWn) processes write checkpoint information to datafile headers and the control file
- C. It reduces the recovery time from instance failures
- D. Incremental checkpoints write some dirty buffers to the datafiles and unwritten redo to the online redo logs.
- E. Thread checkpoints ensure that all dirty buffers are written to data files during a normal shutdown

Answer: BCE

NEW QUESTION 369

Which two tools can be used to configure static service information in the listener.ora file? (Choose two.)

- A. Oracle Net Manager
- B. Oracle Enterprise Manager Cloud Control
- C. Oracle Net Configuration Assistant
- D. Listener Control Utility (LSNRCTL)
- E. Oracle Enterprise Manager Database Express

Answer: AB

NEW QUESTION 374

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

- A. An AWR snapshot shows the SQL statements that are producing the highest load on the system, based on criteria such as elapsed time and CPU time.
- B. AWR data is stored in memory and in a database.
- C. All AWR tables belong to the SYSTEM schema.
- D. The manageability monitor (MMON) process gathers statistics and creates an AWR snapshot that is used by the self- tuning components in a database.
- E. An AWR snapshot contains system-wide tracing and logging informatio

Answer: ABD

NEW QUESTION 378

.....

Relate Links

100% Pass Your 1Z0-062 Exam with Exambible Prep Materials

<https://www.exambible.com/1Z0-062-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>