

1Z0-053 Dumps

Oracle Database 11g: Administration II

<https://www.certleader.com/1Z0-053-dumps.html>



NEW QUESTION 1

- (Topic 1)

After executing the command

```
ALTER DISKGROUP diskgroup2 DROP DISK dg2a;
```

You issue the following command from the ASM instance: `SELECT group_number, COUNT(*) FROM v$asm_operation;`

What is the implication if the query against `V$ASM_OPERATION` returns zero rows?

- A. The drop disk operation is still proceeding and you cannot yet run the undrop disks operation.
- B. The drop disk operation is complete and you can run the undrop disks command if needed.
- C. The drop disk operation is complete and you cannot run the undrop disks command.
- D. The query will fail since there is not a `V$ASM_OPERATION` view available in an ASM instance.
- E. None of the above is true.

Answer: C

Explanation:

Once the `DROP DISK` operation is completed, you **CANNOT** run the `UNDROP DISKS` command any more.

NEW QUESTION 2

- (Topic 1)

You issued the following command:

```
CREATE GLOBAL TEMPORARY TABLE admin_work_area (startdate DATE, enddate DATE, class CHAR(20))  
ON COMMIT DELETE ROWS TABLESPACE tbs_t1;
```

An index is then created on the `ADMIN_WORK_AREA` temporary table.

Which two statements are true regarding the `TBS_T1` tablespace in the above command? (Choose two.)

- A. It stores only the temporary table but not its indexes.
- B. It stores both the temporary table as well as its indexes.
- C. It must be a nondefault temporary tablespace for the database.
- D. It can be a default or nondefault temporary tablespace for the database.
- E. It must be the default temporary tablespace of the user who issues the command.

Answer: BD

NEW QUESTION 3

- (Topic 1)

The following command is executed to shut down an Automatic Storage Management (ASM) instance:

```
SQL>SHUTDOWN ABORT;
```

Which two statements describe the consequences of the above command? (Choose two.)

- A. The disk groups are orderly dismounted.
- B. The ASM instance requires recovery when it is started.
- C. The database instance that are currently clients of the ASM instance are aborted.
- D. The CSS daemon stops and has to be restarted before the ASM instance is restarted.

Answer: BC

NEW QUESTION 4

- (Topic 1)

What is the proper command to shut down the database in a consistent manner?

- A. `shutdown abort`
- B. `shutdown kill`
- C. `shutdown nowait`
- D. `shutdown immediate`
- E. `shutdown halt`

Answer: D

NEW QUESTION 5

- (Topic 1)

You issued the following command to mount the `DATA` disk group in restricted mode: `ALTER DISKGROUP data MOUNT RESTRICT;`

What is the implication of this command?

- A. The client RDBMS instance can access the file as a read-only file.
- B. A new disk cannot be added to a disk group.
- C. A disk in a disk group can be taken offline.
- D. The client RDBMS instance cannot access the files in the disk group.

Answer: D

Explanation:

`MOUNT` ([link](#))

Specify `MOUNT` to mount the disk groups in the local Oracle ASM instance. Specify `ALL MOUNT` to mount all disk groups specified in the `ASM_DISKGROUPS` initialization parameter. File operations can only be performed when a disk group is mounted. If Oracle ASM is running in a cluster or a standalone server managed by Oracle Restart, then the `MOUNT` clause automatically brings the corresponding resource online.

RESTRICTED | NORMAL Use these clauses to determine the manner in which the disk groups are mounted.

In the RESTRICTED mode, the disk group is mounted in single-instance exclusive mode. No other Oracle ASM instance in the same cluster can mount that disk group. In this mode the disk group is not usable by any Oracle ASM client.

In the NORMAL mode, the disk group is mounted in shared mode, so that other Oracle ASM instances and clients can access the disk group. This is the default.

NEW QUESTION 6

- (Topic 1)

Immediately after adding a new disk to or removing an existing disk from an ASM instance, you find that the performance of the database goes down initially until the time the addition or removal process is completed, and then gradually becomes normal.

Which two activities would you perform to maintain a consistent performance of the database while adding or removing disks? (Choose two.)

- A. Define the POWER option while adding or removing the disks.
- B. Increase the number of ARB processes by setting up a higher value for ASM_POWER_LIMIT.
- C. Increase the number of DBWR processes by setting up a higher value for DB_WRITER_PROCESSES.
- D. Increase the number of slave database writer processes by setting up a higher value for DBWR_IO_SLAVES.

Answer: AB

Explanation:

ARBn (ASM Rebalance Process): Rebalances data extents within an ASM disk group, possible processes are ARB0-ARB9 and ARBA.

ALTER DISKGROUP..POWER clause, specify a value from 0 to 11, where 0 stops the rebalance operation and 11 permits Oracle ASM to execute the rebalance as fast as possible. The value you specify in the POWER clause defaults to the value of the ASM_POWER_LIMIT initialization parameter. If you omit the POWER clause, then Oracle ASM executes both automatic and specified rebalance operations at the power determined by the value of the ASM_POWER_LIMIT initialization parameter.

Note:

Beginning with Oracle Database 11g Release 2 (11.2.0.2), if the COMPATIBLE.ASM disk group attribute is set to 11.2.0.2 or higher, then you can specify a value from 0 to 1024 in the POWER clause.

NEW QUESTION 7

- (Topic 1)

Examine the following ALTER command;

```
SQL> ALTER DISKGROUP dgroup1 UNDROP DISKS;
```

What is the purpose of the command?

- A. It cancels all pending disk drops within the disk group.
- B. It adds previously dropped disks back into the disk group.
- C. It restores disks that are being dropped as the result of a DROP DISKGROUP operation.
- D. It mounts disks in the disk group for which the drop-disk operation has already been completed.
- E. It restores all the dropped disks in the disk group for which the drop-disk operation has already been completed.

Answer: A

Explanation:

The key point is PENDING.

NEW QUESTION 8

- (Topic 1)

You are managing an Oracle Database 11g instance and an Oracle Database 10g instance on the same machine. Both instances use the ASM instance as storage. Which statement regarding the ASM disk group compatibility attributes are true in this scenario? (Choose all that apply.)

- A. The database-compatibility version settings for each instance must be greater than or equal to the RDBMS compatibility of all ASM disk groups used by that database instances.
- B. RDBMS compatibility and the database version determines whether a database instance can mount the ASM disk group.
- C. The RDBMS compatibility settings for a disk group control the format of data structures for ASM metadata on the disk.
- D. ASM compatibility controls which features for the ASM will be enabled.

Answer: ABD

NEW QUESTION 9

- (Topic 1)

You are managing an Oracle Database 11g ASM instance with a disk group dg01 having three disks. One of the disks in the disk group becomes unavailable because of power failure. You issued the following command to change the DISK_REPAIR_TIME attribute from 3.6 hours to 5 hours:

```
ALTER DISKGROUP dg01 SET ATTRIBUTE 'disk_repair_time' = '5h';
```

 To which disks in the disk group will the new value be applicable?

- A. all disks in the disk group
- B. all disks that are currently in OFFLINE mode
- C. all disks that are not currently in OFFLINE mode
- D. all disks in the disk group only if all of them are ONLINE

Answer: C

Explanation:

Check out the answer options, it is tricky. The NOT OFFLINE disks equals ONLINE disks. Refer to Set the DISK_REPAIR_TIME Disk Group Attribute Appropriately.

The DISK_REPAIR_TIME disk group attribute specifies how long a disk remains offline before Oracle ASM drops the disk. If a disk is made available before the DISK_REPAIR_TIME parameter has expired, the storage administrator can issue the ONLINE DISK command and Oracle ASM resynchronizes the stale data from the mirror side. In Oracle Database 11g, the online disk operation does not restart if there is a failure of the instance on which the disk is running. You must reissue the command manually to bring the disk online.

NEW QUESTION 10

- (Topic 1)

What is the default AU size of an ASM disk group? What is the maximum AU size in an ASM disk group?

- A. 100KB default, 10TB maximum
- B. 256KB default, 1024MB maximum
- C. 10MB default, 126PB maximum
- D. 64KB default, 1EB maximum
- E. 1MB default, 64MB maximum

Answer: E

Explanation:

The AU size is determined at creation time with the allocation unit size (AU_SIZE) disk group attribute. The values can be 1, 2, 4, 8, 16, 32, and 64 MB.

NEW QUESTION 10

- (Topic 1)

What is the advantage of setting the ASM-preferred mirror read for the stretch cluster configuration?

- A. It improves resync operations.
- B. This feature enables much faster file opens.
- C. It improves performance as fewer extent pointers are needed in the shared pool.
- D. It improves performance by reading from a copy of an extent closest to the node.

Answer: D

Explanation:

Preferred Read Failure Groups

When you configure Oracle ASM failure groups, it might be more efficient for a node to read from an extent that is closest to the node, even if that extent is a secondary extent. In other words, you can configure Oracle ASM to read from a secondary extent if that extent is closer to the node instead of Oracle ASM reading from the primary copy which might be farther from the node. Using the preferred read failure groups feature is most useful in extended clusters.

NEW QUESTION 12

- (Topic 1)

What is the result of increasing the value of the parameter ASM_POWER_LIMIT during a rebalance operation?

- A. The ASM rebalance operation will likely consume fewer resources and complete in a shorter amount of time.
- B. The ASM rebalance operation will consume fewer resources and complete in a longer amount of time.
- C. The ASM rebalance operation will be parallelized and should complete in a shorter amount of time.
- D. There is no ASM_POWER_LIMIT setting used in ASM.
- E. None of the above

Answer: C

NEW QUESTION 13

- (Topic 1)

What are the advantages of variable extent size support for large ASM files? (Choose two.)

- A. It improves resync operations when the disk comes online after being taken offline for maintenance purposes.
- B. It improves performance in the extended cluster configuration by reading from a local copy of an extent.
- C. Fewer extent pointers are needed to describe the file and less memory is required to manage the extent maps in the shared pool.
- D. This feature enables faster file opens because of the reduction in the amount of memory that is required to store file extents.

Answer: CD

NEW QUESTION 18

- (Topic 1)

When an ASM instance receives a SHUTDOWN NORMAL command, what command does it pass on to all database instances that rely on the ASM instances disk groups?

- A. TRANSACTIONAL
- B. IMMEDIATE
- C. ABORT
- D. NORMAL

Answer: A

NEW QUESTION 21

- (Topic 1)

Users are connected to a database instance that is using Automatic Storage Management (ASM). The DBA executes the command as follows to shut down the ASM instance:

```
SQL> SHUTDOWN IMMEDIATE;
```

What happens to the database instance?

- A. It shuts down long with the ASM instance.
- B. It is aborted and the ASM instance shuts down normally.
- C. It stays open and SHUTDOWN command for the ASM instance fails.
- D. It shuts down only after all pending transactions are completed and the ASM instance waits for this before shutting down.

Answer: C

Explanation:

IMMEDIATE or TRANSACTIONAL Clause ([link](#))

Oracle ASM waits for any in-progress SQL to complete before performing an orderly dismount of all of the disk groups and shutting down the Oracle ASM instance. Oracle ASM does not wait for users currently connected to the instance to disconnect. If any database instances are connected to the Oracle ASM instance, then the SHUTDOWN command returns an error and leaves the Oracle ASM instance running. Because the Oracle ASM instance does not contain any transactions, the TRANSACTIONAL mode behaves the same as IMMEDIATE mode.

NEW QUESTION 22

- (Topic 2)

What is the purpose of the recover command? (Choose all that apply.)

- A. Recover database datafiles from physical disk backup sets.
- B. Recover required incremental backups from physical disk backup sets.
- C. Recover required archived redo logs from physical disk backup sets.
- D. Apply incremental backups to recover the database.
- E. Apply archived redo logs to recover the database.

Answer: BCDE

NEW QUESTION 25

- (Topic 2)

During recovery, you need to know if log sequence 11 is in the online redo logs, and if so, you need to know the names of the online redo logs so you can apply them during recovery. Which view or views would you use to determine this information? (Choose all that apply.)

- A. V\$LOGFILE
- B. V\$RECOVER_LOG
- C. V\$RECOVER_DATABASE
- D. V\$LOG_RECOVER
- E. V\$LOG

Answer: AE

NEW QUESTION 30

- (Topic 2)

Archived redo logs can be copied to more than one destination by Oracle.

- A. True
- B. False

Answer: A

NEW QUESTION 32

- (Topic 2)

Your database has a backup that was taken yesterday (Tuesday) between 13:00 and 15:00 hours. This is the only backup you have. You have lost all the archived redo logs generated since the previous Monday, but you have archived redo logs available from the previous Sunday and earlier. You now need to restore your backup due to database loss.

To which point can you restore your database?

- A. 13:00 on Tuesday.
- B. 15:00 on Tuesday.
- C. Up until the last available archived redo log on Sunday.
- D. To any point; all the redo should still be available in the online redo logs.
- E. The database is not recoverable.

Answer: E

NEW QUESTION 35

- (Topic 2)

Which statement is true regarding the VALIDATE DATABASE command?

- A. It checks the database for intrablock corruptions only.
- B. It checks for block corruption in the valid backups of the database.
- C. It checks the database for both intrablock and interblock corruptions.
- D. It checks for only those corrupted blocks that are associated with data files.

Answer: A

Explanation:

interblock corruption

A type of block corruption in which the corruption occurs between blocks rather than within the block itself. This type of corruption can only be logical corruption. intrablock corruption A type of block corruption in which the corruption occurs within the block itself. this type of corruption can be either a physical corruption or logical corruption.

Table 16-1 Detection, Repair, and Monitoring of Block Corruption

Response	Intrablock Corruption	Interblock Corruption
Detection	All database utilities detect intrablock corruption, including RMAN (for example, the BACKUP command) and the DBVERIFY utility. If a database process can encounter the ORA-1578 error, then it can detect the corruption and monitor it.	Only DBVERIFY and the interblock corruption.
Tracking	The V\$DATABASE_BLOCK_CORRUPTION view displays blocks marked corrupt by Oracle Database components such as RMAN commands, ANALYZE, dbv, SQL queries, and so on. Any process that encounters an intrablock corruption records the block corruption in this view and in ADR.	The database monitors in ADR.
Repair	Repair techniques include block media recovery , restoring data files, recovering with incremental backups, and block newing. Block media recovery can repair physical corruptions, but not logical corruptions. Any RMAN command that fixes or detects that a block is repaired updates V\$DATABASE_BLOCK_CORRUPTION. For example, RMAN updates the repository at end of successful block media recovery. If a BACKUP, RESTORE, or VALIDATE command detects that a block is no longer corrupted, then it removes the repaired block from the view.	You must fix interblock techniques such as drop index, and so on.

NEW QUESTION 40

- (Topic 2)

To enable faster incremental backups, you enabled block change tracking for the database. Which two statements are true about the block change tracking file? (Choose two.)

- A. Multiple change tracking files can be created for a database.
- B. The change tracking file must be created after the first level 0 backup.
- C. RMAN does not support backup and recovery of the change tracking file.
- D. The database clears the change tracking file and starts tracking changes again, after whole database restore and recovery operations.

Answer: CD

NEW QUESTION 42

- (Topic 2)

You have the following requirements in relation to the detection of block corruption for your database instance:

- ? Check for logical self-consistency of data blocks when modified in memory.
- ? Checksums are calculated before and after the block change.
- ? Checks are performed for the lost writes to the physical standby database.

Which method would help you perform the above checks automatically?

- A. Set the DB_SECUREFILE parameter to PERMITTED.
- B. Set the DB_ULTRA_SAFE parameter to DATA_ONLY.
- C. Set the DB_LOCK_CHECKSUM parameter to TYPICAL.
- D. Set the DB_LOST_WRITE_PROTECT parameter to TYPICAL.

Answer: B

Explanation:

```
Parameter type      String
Syntax              DB_ULTRA_SAFE = { OFF | DATA_ONLY | DATA_AND_INDEX }
Default value       OFF
Modifiable         No
Basic               No
```

DB_ULTRA_SAFE sets the default values for other parameters that control protection levels.

C:\Users\albo\Desktop\1-1.jpg Values:

OFF

When any of DB_BLOCK_CHECKING, DB_BLOCK_CHECKSUM, or DB_LOST_WRITE_PROTECT are explicitly set, no changes are made.

DATA_ONLY

? DB_BLOCK_CHECKING will be set to MEDIUM.

? DB_LOST_WRITE_PROTECT will be set to TYPICAL.

? DB_BLOCK_CHECKSUM will be set to FULL.

DATA_AND_INDEX

DB_BLOCK_CHECKING will be set to FULL. DB_LOST_WRITE_PROTECT will be set to TYPICAL. DB_BLOCK_CHECKSUM will be set to FULL.

```
Parameter type      String
Syntax              DB_BLOCK_CHECKING = { FALSE | OFF | LOW | MEDIUM | TRUE | FULL }
Default value       FALSE
Modifiable         ALTER SYSTEM
Basic               No
```

DB_BLOCK_CHECKING specifies whether or not Oracle performs block checking for database blocks.

C:\Users\albo\Desktop\1-1.jpg Values:

OFF or FALSE

No block checking is performed for blocks in user tablespaces. However, semantic block checking for SYSTEM tablespace blocks is always turned on.

LOW

Basic block header checks are performed after block contents change in memory (for example, after UPDATE or INSERT statements, on-disk reads, or inter-instance block transfers in Oracle RAC).

MEDIUM

All LOW checks and full semantic checks are performed for all objects except indexes (whose contents can be reconstructed by a drop+rebuild on encountering a

corruption).
FULL or TRUE
All LOW and MEDIUM checks and full semantic checks are performed for all objects.

NEW QUESTION 47

- (Topic 2)

If you issue the command shutdown abort prior to trying to put the database in ARCHIVELOG mode, what will be the result when you issue the command alter database archivelog?

- A. The alter database archivelog command will fail.
- B. The alter database archivelog inconsistent command must be used to put the database in ARCHIVELOG mode.
- C. The alter database archivelog command will succeed.
- D. The alter database archivelog command will ask if you want to make the database consistent first.
- E. There is no alter database archivelog command.
- F. The correct command is alter database alter logging.

Answer: A

Explanation:

Before you change database to archivelog mode, you need to have a clean database shutdown.

NEW QUESTION 51

- (Topic 2)

How many individual archive-log destination directories are supported by Oracle Database 11g?

- A. 7
- B. 1
- C. 10
- D. 11
- E. 21

Answer: C

NEW QUESTION 52

- (Topic 2)

Which of the following parameters defines the location where Oracle should create archived redo logs?

- A. LOG_ARCHIVE_1
- B. LOG_DESTINATION_1
- C. LOG_ARCHIVED_DESTINATION_1
- D. LOG_ARCHIVE_DEST_1
- E. LOG_ARCHIVE_SOURCE_1

Answer: D

NEW QUESTION 55

- (Topic 2)

Which command would you use to determine what database backups are currently available for restore?

- A. list database backup;
- B. report database backup;
- C. list backup of database;
- D. list summary backup;
- E. report backup of database;

Answer: C

NEW QUESTION 57

- (Topic 3)

You are working on a CATDB database that contains an Oracle Database version 11.1 catalog schema owned by the user RCO11. The INST1 database contains an Oracle Database version 10.1 catalog schema owned by the user RCAT10.

You want the RMAN to import metadata for database IDs 1423241 and 1423242, registered in RCAT10, into the recovery catalog owned by RCO11. You executed the following commands:

```
RMAN> CONNECT CATALOG rco11/password@catdb  
RMAN> IMPORT CATALOG rcat10/oracle@inst1 NO UNREGISTER;
```

Which two statements are true regarding the tasks accomplished with these commands? (Choose two.)

- A. They import all metadata from the RCAT10 catalog.
- B. They unregister the database from the RCAT10 catalog.
- C. They do not register the databases registered in the RCAT10 catalog.
- D. They register all databases registered in the RCAT10 catalog.

Answer: AD

Explanation:

```
IMPORT CATALOG <connectStringSpec>
[DBID = <dbid> [, <dbid>,...]]
[DB_NAME=<dbname>[, <dbname>,...]]
[ NO_UNREGISTER ];
```

NO_UNREGISTER option forces to remain the database registration kept in the source RC.
By default, after the IMPORT command completed, it will unregister database from the source RC and register the databases to the target RC.

C:\Users\albo\Desktop\1-1.jpg

NEW QUESTION 59

- (Topic 3)

In your production database, you:

? Are using Recovery Manager (RMAN) with a recovery catalog to perform the backup operation at regular intervals

? Set the control file autobackup to "on"

? Are maintaining image copies of the database files

You have lost the server parameter file (SPFILE) and the control file.

Which option must you consider before restoring the SPFILE and the control file by using the control file autobackup?

- A. Setting DBID for the database
- B. Using the RMAN SWITCH command
- C. Using the RMAN SET NEWNAME command
- D. Starting up the database instance in the NOMOUNT state

Answer: D

NEW QUESTION 63

- (Topic 3)

You want to create the Recovery Manager (RMAN) Virtual Private Catalog (VPC) to maintain a separation of responsibilities along with a consolidation of RMAN repository.

Which condition must be met before you create the VPC?

- A. A base catalog exists
- B. The recovery catalog is empty
- C. The base recovery catalog must be dropped
- D. A target database is registered in the recovery catalog

Answer: A

NEW QUESTION 67

- (Topic 3)

RMAN provides more granular catalog security through which feature?

- A. Virtual private database
- B. Virtual private catalog
- C. RMAN virtual database
- D. RMAN secure catalog
- E. Oracle Database Vault

Answer: B

Explanation:

About Virtual Private Catalogs

By default, all of the users of an RMAN recovery catalog have full privileges to insert, update, and delete any metadata in the catalog. For example, if the administrators of two unrelated databases share the same recovery catalog, each administrator could, whether inadvertently or maliciously, destroy catalog data for the other's database. In many enterprises, this situation is tolerated because the same people manage many different databases and also manage the recovery catalog. But in other enterprises where clear separation of duty exists between administrators of various databases, and between the DBA and the administrator of the recovery catalog, you may desire to restrict each database administrator to modify only backup metadata belonging to those databases that they are responsible for, while still keeping the benefits of a single, centrally managed, RMAN recovery catalog. This goal can be achieved by implementing virtual private catalogs.

NEW QUESTION 69

- (Topic 3)

What privileges must be granted to allow an account to create the recovery catalog? (Choose all that apply.)

- A. RECOVERY_CATALOG_OWNER
- B. DBA
- C. RESOURCE
- D. SELECT ANY DICTIONARY
- E. CONNECT

Answer: AC

NEW QUESTION 73

- (Topic 3)

The following databases are registered in the base recovery catalog: PROD1, PROD2, and PROD3. The database user CATOWNER owns the base recovery catalog. You want a new user VPC1 to have access to only the PROD1 database and create a virtual private catalog.

Given below are some of the commands required to achieve this:

1. SQL> GRANT recovery_catalog_owner TO vpc1;
 2. RMAN> CONNECT CATALOG vpc1/password@catdb;
 3. RMAN> GRANT CATALOG FOR DATABASE prod1 TO vpc1;
 4. RMAN> CONNECT CATALOG catowner/password@catdb;
 5. RMAN> CREATE VIRTUAL CATALOG;
- What is the correct sequence in which the commands have to be executed?

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

Answer: B

Explanation:

Refer to here

The basic steps for creating a virtual private catalog are as follows:

1. Create the database user who will own the virtual private catalog (if this user does not exist) and grant this user access privileges.

1.1 Start SQL*Plus and connect to the recovery catalog database with administrator privileges.

1.2 If the user that will own the virtual private catalog does not exist, then create the user. SQL> CREATE USER vpc1 IDENTIFIED BY password DEFAULT TABLESPACE vpcusers QUOTA UNLIMITED ON vpcusers;

1.3 Grant the RECOVERY_CATALOG_OWNER role to the database user that will own the virtual private catalog, and then exit SQL*Plus.

SQL> GRANT recovery_catalog_owner TO vpc1;

1.4 Start RMAN and connect to the recovery catalog database as the base recovery catalog owner

(not the virtual private catalog owner).

RMAN> CONNECT CATALOG catowner@catdb;

1.5 Grant desired privileges to the virtual private catalog owner. RMAN> GRANT CATALOG FOR DATABASE prod1 TO vpc1;

You can also use a DBID rather than a database name. The virtual private catalog user does not have access to the metadata for any other databases registered in the recovery catalog.

You can also grant the user the ability to register new target databases in the recovery catalog. For example:

RMAN> GRANT REGISTER DATABASE TO vpc1;

2. Create the virtual private catalog.

2.1 Start RMAN and connect to the recovery catalog database as the virtual private catalog owner (not the base recovery catalog owner).

RMAN> CONNECT CATALOG vpc1@catdb;

2.2 Create the virtual private catalog. RMAN> CREATE VIRTUAL CATALOG;

If you intend to use a 10.2 or earlier release of RMAN with this virtual private catalog, then execute the following PL/SQL procedure (where base_catalog_owner is the database user who owns the base recovery catalog):

SQL> EXECUTE base_catalog_owner.DBMS_RCVCAT.CREATE_VIRTUAL_CATALOG;

NEW QUESTION 76

- (Topic 3)

What RMAN command must you use before you can back up a database using the recovery catalog?

- A. create catalog
- B. install database
- C. catalog database
- D. merge Catalog with database
- E. register database

Answer: E

NEW QUESTION 79

- (Topic 3)

Which is the correct way to connect to both the target database and the recovery catalog from the RMAN command line? Assume that the target database is called ORCL and that the recovery catalog database is called RCAT. Also assume that the recovery-catalog owner is called RCAT_OWN. Assume the environment is configured for the ORCL database. (Choose all that apply.)

- A. rman target=/ catalog=@rcat
- B. rman target=/ catalog=rcat_own/rcat_own
- C. rman target=/ catalog=rcat_own/rcat_own@RCAT
- D. rman target=sys/robert@orcl catalog=rcat_own/rcat_own@RCAT
- E. You cannot connect to the target database and the recovery catalog at the same time.

Answer: CD

NEW QUESTION 84

- (Topic 3)

What is the purpose of the RMAN recovery catalog? (Choose all that apply.)

- A. It must be used because all RMAN-related backup and recovery metadata information is contained in it.
- B. It provides a convenient, optional, repository of backup- and recovery-related metadata.
- C. It provides the ability to store RMAN scripts for global use by any database that has access to the repository.
- D. It provides a means of storing all RMAN backup sets physically in an Oracle database server.
- E. It provides the ability to store backup records for more than a year.

Answer: BCE

Explanation:

A recovery catalog is a database schema used by RMAN to store metadata about one or more Oracle databases. Typically, you store the catalog in a dedicated database. A recovery catalog provides the following benefits:

? A recovery catalog creates redundancy for the RMAN repository stored in the control file of each target database. The recovery catalog serves as a secondary metadata repository. If the target control file and all backups are lost, then the RMAN metadata still exists in the recovery catalog.

? A recovery catalog centralizes metadata for all your target databases. Storing the metadata in a single place makes reporting and administration tasks easier to perform.

? A recovery catalog can store metadata history much longer than the control file. This capability is useful if you must do a recovery that goes further back in time than the history in the control file. The added complexity of managing a recovery catalog database can be offset by the convenience of having the extended backup history available.

Some RMAN features function only when you use a recovery catalog. For example, you can store RMAN scripts in a recovery catalog. The chief advantage of a stored script is that it is available to any RMAN client that can connect to the target database and recovery catalog. Command files are only available if the RMAN client has access to the file system on which they are stored.

A recovery catalog is required when you use RMAN in a Data Guard environment. By storing backup metadata for all primary and standby databases, the catalog enables you to offload backup tasks to one standby database while enabling you to restore backups on other databases in the environment.

NEW QUESTION 88

- (Topic 3)

While performing a regular check on your recovery catalog you realized that the catalog database is running out of space and you do not have options to increase the space. However, you have another database where more space is available and you want to move your existing recovery catalog to this database.

The options that can be considered while moving the recovery catalog are as follows:

1. Using one of the Oracle expdp utilities to export the catalog data
2. Creating a recovery catalog user and granting the necessary privileges in the other database
3. Creating the recovery catalog using the CREATE CATALOG command
4. Using the corresponding impdp utility to import the catalog data into the other database
5. Registering the target database in the new catalog database using the REGISTER DATABASE command.

Identify the option with the correct sequence for moving the recovery catalog.

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

Answer: B

Explanation:

The exp/imp tools can export and import the complete data structure and data extents to the destination database, so that you don't need to do create catalog and register database.

NEW QUESTION 92

- (Topic 4)

Identify the persistent configuration setting for the target database that can be set for the backup by using RMAN. (Choose all that apply.)

- A. Backup retention policy
- B. Default backup device type
- C. Default destinations for backups
- D. Multiple backup device types for single backup
- E. Default section size for backups

Answer: ABC

Explanation:

SECTION SIZE cannot be configured through CONFIGURE command.

NEW QUESTION 97

- (Topic 4)

Which command is used to configure RMAN to perform a compressed backup for every backup executed?

- A. BACKUP AS COMPRESSED BACKUPSET DATABASE
- B. BACKUP AS COMPRESSED COPY OF DATABASE
- C. CONFIGURE DEVICE TYPE DISK BACKUP TYPE TO COMPRESSED BACKUPSET
- D. CONFIGURE DEVICE TYPE DISK BACKUP TYPE COMPRESS
- E. BACKUP DATABASE COMPRESS

Answer: C

NEW QUESTION 98

- (Topic 4)

Which type of backup backs up only data blocks modified since the most recent backup at the same level or lower?

- A. Differential incremental backup
- B. Different incremental backup
- C. Cumulative backup
- D. Cumulative incremental backup

Answer: A

NEW QUESTION 99

- (Topic 5)

What is the purpose of the catalog command?

- A. To review RMAN control file and recovery catalog metadata and ensure that its correct
- B. To delete RMAN backup-related metadata from the recovery catalog
- C. To create metadata in the control file and the recovery catalog related to backup set pieces
- D. To create a report that lists database backups
- E. To rebuild the recovery catalog

Answer: C

Explanation:

Use the CATALOG command to do the following:

? Add backup pieces and image copies on disk to the RMAN repository

? Record a data file copy as a level 0 incremental backup in the RMAN repository, which enables you to use it as part of an incremental backup strategy

NEW QUESTION 101

- (Topic 5)

You want to take a complete database backup using RMAN. The backup should consist only the used blocks from your database.

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

- A. Backup compression should be enabled
- B. Parallelism for the channel should be set to 2
- C. All the files must be backed up as backup sets
- D. The backup may be stored either on disk or on media with media manager

Answer: CD

NEW QUESTION 103

- (Topic 5)

Which of the following RMAN structures can data from a datafile span?

- A. RMAN backup-set pieces spanning backup sets
- B. RMAN backup-set pieces within a given backup set
- C. RMAN backups
- D. RMAN channels
- E. None of the above

Answer: B

NEW QUESTION 108

- (Topic 5)

Identify the scenarios in which the RMAN CROSSCHECK command is useful. (Choose all that apply.)

- A. To check that the obsolete backups are deleted from the repository records and from the disk.
- B. To update the RMAN repository if you delete archived redo logs with operating system commands.
- C. To update outdated information about backups that disappeared from disk or tape or became corrupted and inaccessible.
- D. To synchronize the actual files on disk or in the media management catalog with the RMAN repository for which the backup was not taken using RMAN.

Answer: BC

Explanation:

Crosscheck are useful because they can do the following: (Link)

? Update outdated information about backups that disappeared from disk or tape or became corrupted

? Update the repository if you delete archived redo logs or other files with operating system commands

Use the CROSSCHECK command to synchronize the physical reality of backups and copies with their logical records in the RMAN repository.

The CROSSCHECK command checks only objects marked AVAILABLE or EXPIRED in the repository by examining the files on disk for DISK channels or by querying the media manager for sbt channels.

Meaning of Crosscheck Status: EXPIRED

Object is not found either in file system (for DISK) or in the media manager (for sbt). A backup set is EXPIRED if any backup piece in the set is EXPIRED.

The CROSSCHECK command does not delete the repository records of the files that it does not find, but updates their repository records to EXPIRED. You can run DELETE EXPIRED to remove the repository records for expired files and any existing physical files whose status is EXPIRED.

If backups are EXPIRED, then you can re-execute the crosscheck later and determine whether expired backups are available. This precaution is especially useful when you use RMAN with a media manager. For example, if some backup pieces or copies were erroneously marked as EXPIRED because the PARMS channel settings were incorrect, then after ensuring that the files really do exist in the media manager, run the CROSSCHECK BACKUP command again to restore those files to AVAILABLE status. AVAILABLE

Object is available for use by RMAN. For a backup set to be AVAILABLE, all backup pieces in the set must have the status AVAILABLE.

NEW QUESTION 109

- (Topic 5)

What is the impact of the results of the output of the following command?

RMAN> report unrecoverable database;

Report of files that need backup due to unrecoverable operations File Type of Backup Required Name

4 full or incremental C:\ORACLE\ORADATA\ORCL\USERS01.DBF

- A. There are no backup sets with any backups of the users01.dbf datafile.
- B. The users01.dbf datafile has had unrecoverable operations occur in i
- C. It will need to be backed up or some data loss is possible during a recovery.
- D. The users01.dbf datafile is corrupted.
- E. The users01.dbf datafile backup exceeds the retention criteria.
- F. The last backup of the users01.dbf datafile failed and must be rerun.

Answer: D

NEW QUESTION 114

- (Topic 5)

Examine the following set of RMAN commands:

```
RMAN> CONFIGURE CHANNEL dc1 DEVICE TYPE DISK FORMAT '/u02/backup/%U'; RMAN> RUN (  
ALLOCATE CHANNEL Chi DEVICE TYPE DISK;  
EXECUTE SCRIPT full_backup;  
)
```

Which statement is true when the RMAN RUN block is executed?

- A. The execution of the script fails because multiple channels cannot coexist.
- B. The script is executed and both the channels are used for the script execution.
- C. The new channel 'CHI' is ignored because a channel has been configured already.
- D. configuration parameter dc1 is overridden because a new channel is allocated in RMAN RUN block.

Answer: D

NEW QUESTION 117

- (Topic 5)

Why would you execute the report obsolete command?

- A. To list all backups that were no longer available for restore operations
- B. To list all backups that had aged beyond the RMAN retention criteria
- C. To list all backup set pieces listed in control-file or recovery-catalog metadata that are not on the backup media
- D. To list all datafiles that are no longer part of the database and thus do not need to be backed up
- E. To list all archived redo logs that are no longer needed for any database recovery

Answer: B

Explanation:

You can report backup sets, backup pieces, and data file copies that are obsolete that is, not needed to meet a specified retention policy by specifying the OBSOLETE keyword.

NEW QUESTION 119

- (Topic 5)

Given the following steps, which would be the correct order to create a backup of an Oracle database in NOARCHIVELOG mode?

- 7. shutdown immediate from RMAN
- 8. Log into RMAN
- 9. startup mount from RMAN 10.backup database
- 11. alter database open
- 12. backup database plus archive log delete input

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

Answer: E

Explanation:

Backing Up a Database in NOARCHIVELOG Mode

If a database runs in NOARCHIVELOG mode, then the only valid database backup is a consistent backup. For the backup to be consistent, the database must be mounted after a consistent shutdown. No recovery is required after restoring the backup.

To make a consistent database backup:

- 1. Start RMAN and connect to a target database.
- 2. Shut down the database consistently and then mount it.

For example, enter the following commands to guarantee that the database is in a consistent state for a backup:

```
RMAN> SHUTDOWN IMMEDIATE; RMAN> STARTUP FORCE DBA; RMAN> SHUTDOWN IMMEDIATE; RMAN> STARTUP MOUNT;
```

- 3. Run the BACKUP DATABASE command.

For example, enter the following command at the RMAN prompt to back up the database to the default backup device:

```
RMAN> BACKUP DATABASE;
```

The following variation of the command creates image copy backups of all data files in the database:

```
RMAN> BACKUP AS COPY DATABASE;
```

- 4. Open the database and resume normal operations. The following command opens the database:

```
RMAN> ALTER DATABASE OPEN;
```

NEW QUESTION 122

- (Topic 6)

If you lost your entire database, including the database spfile, control files, online redo logs, and database datafiles, what kind of recovery would be required with RMAN?

- A. Complete database recovery.
- B. Incomplete database recovery.
- C. Approximate database recovery.
- D. Archived database recovery.
- E. The database could not be recovered with RMAN.

Answer: B

NEW QUESTION 124

- (Topic 6)

Which of the following represents the correct way to perform an online recovery of datafile 4, which is assigned to a tablespace called USERS?

- A. shutdown restore datafile 4; recover datafile 4; alter database open;
- B. Sql alter database datafile 4 offline; restore datafile 4; recover datafile 4; alter database open;
- C. Sql alter database datafile 4 offline; restore datafile 4; Sql alter database datafile 4 online;
- D. Sql alter database datafile 4 offline; restore database datafile 4; recover database datafile 4; Sql alter database datafile 4 online;
- E. Sql alter database datafile 4 offline; restore datafile 4; recover datafile 4; Sql alter database datafile 4 online;

Answer: E

NEW QUESTION 126

- (Topic 6)

You are managing a 24*7 database. The backup strategy for the database is to perform user-managed backups. Identify two prerequisites to perform the backups. (Choose two.)

- A. The database must be opened in restricted mode.
- B. The database must be configured to run in ARCHIVELOG mode.
- C. The tablespaces are required to be in backup mode before taking the backup.
- D. The tablespaces are required to be in read-only mode before taking the backup

Answer: BC

NEW QUESTION 127

- (Topic 6)

Examine the following scenario:

- ? Database is running in ARCHIVELOG mode.
 - ? Complete consistent backup is taken every Sunday.
 - ? On Tuesday the instance terminates abnormally because the disk on which control files are located gets corrupted
 - ? The disk having active online redo log files is also corrupted.
 - ? The hardware is repaired and the paths for online redo log files and control files are still valid.
- Which option would you use to perform the recovery of database till the point of failure?

- A. Restore the latest whole backup, perform complete recovery, and open the database normally
- B. Restore the latest whole backup, perform incomplete recovery, and open the database with the RESETLOGS option.
- C. Restore the latest backups control file, perform complete recovery, and open the database with the RESETLOGS option.
- D. Restore the latest backup control file, perform incomplete recovery using backup control file, and open the database with the RESETLOG option.

Answer: D

NEW QUESTION 132

- (Topic 6)

In your test database:

- ? You are using Recovery Manager (RMAN) to perform incremental backups of your test database
 - ? The test database is running in NOARCHIVELOG mode
 - ? One of the data files is corrupted
 - ? All online redo log files are lost because of a media failure
- Which option must you consider in this scenario?

- A. Configuring the database in ARCHIVELOG mode and then using incremental backup to recover the database
- B. Using incremental backup to recover the damaged data file and then manually creating the online redo log files
- C. Creating a new test database because the database is not recoverable due to the fact that the database is configured in NOARCHIVELOG mode
- D. Using incremental backups to recover the database by using the RECOVER DATABASE NOREDO command and then using the RESETLOGS option to open the database.

Answer: D

Explanation:

Example 3-6 Recovering a NOARCHIVELOG Database

You can perform limited recovery of changes to a database running in NOARCHIVELOG mode by applying incremental backups. The incremental backups must be consistent, like all backups of a database run in NOARCHIVELOG mode, so you cannot back up the database when it is open.

Assume that you run database prod in NOARCHIVELOG mode with a recovery catalog.

You shut down the database consistently and make a level 0 backup of database prod to tape on Sunday afternoon. You shut down the database consistently and make a level 1 differential incremental backup to tape at 3:00 a.m. on Wednesday and Friday.

On Saturday, a media failure destroys half the data files and the online redo logs. Because the online logs are lost, you must specify the NOREDO option in the RECOVER command. Otherwise, RMAN searches for the redo logs after applying the Friday incremental backup and issues an error message when it does not find them.

After connecting RMAN to prod and the catalog database, recover as follows: STARTUP FORCE NOMOUNT;

RESTORE CONTROLFILE; # restore control file from consistent backup

ALTER DATABASE MOUNT;

RESTORE DATABASE; # restore data files from consistent backup

RECOVER DATABASE NOREDO; # specify NOREDO because online redo logs are lost

ALTER DATABASE OPEN RESETLOGS;

The recovered database reflects only changes up through the time of the Friday incremental backup. Because there are no archived redo log files, there is no way to recover changes made after the incremental backup.

NEW QUESTION 133

- (Topic 6)

Your database has experienced a loss of datafile users_01.dbf, which is associated with a tablespace called USERS. The database is still running. Which answer properly describes the order of the steps that you would use to recover from this error?

1. Shut down the database.
2. Take the users_01.dbf datafile offline with the alter database command.
3. Restore the users_01.dbf datafile from backup media with the required archived redo logs.
4. Restore all users tablespace-related datafiles from backup media.
5. Issue the recover tablespace users command.
6. Issue the recover datafile users_01.dbf command.
7. Start up the database.
8. Bring the users_01.dbf datafile online with the alter database command.

- A. 1, 3, 6, 7
- B. 2, 3, 6, 8
- C. 1, 2,3,6,7
- D. 1, 2, 3, 6, 7, 8
- E. 2, 3,6,5,7

Answer: B

NEW QUESTION 135

- (Topic 6)

What is the correct order of the following commands if you wanted to restore datafile 4, which was accidentally removed from the file system?

- a: sql 'alter database datafile 4 online'; b: restore datafile 4;
- c: recover datafile 4;
- d: sql 'alter database datafile 4 offline';
- e: startup
- f: shutdown

- A. a, c, b, d
- B. d, b, c, a
- C. f, d, b, c, a, e
- D. c, a, b, d, f
- E. a, b, d, e

Answer: B

NEW QUESTION 138

- (Topic 6)

Which files are required for a full recovery of the database in ARCHIVELOG mode? (Choose three.)

- A. Database datafiles
- B. Online redo logs
- C. Archived redo logs
- D. Backup control file
- E. Control file from a backup

Answer: ACD

NEW QUESTION 141

- (Topic 6)

What is the correct order of steps to perform an online database backup?

- a: alter database begin backup; b: alter database end backup;
- c: Back up the database datafiles. d: Back up the archive log files.
- e: alter system switch logfile;

- A. a, b, c, d, e
- B. e, d, a, b, c
- C. a, c, b, d, e
- D. d, b, c, a, e
- E. a, c, b, e, d

Answer: E

Explanation:

(link)

To back up online read/write tablespaces in an open database:

1. Use the DBA_DATA_FILES data dictionary view to identify all of the data files in the tablespace.
2. ALTER TABLESPACE users BEGIN BACKUP;
3. Back up the online data files of the online tablespace with operating system commands.
4. ALTER TABLESPACE users END BACKUP;
5. ALTER SYSTEM ARCHIVE LOG CURRENT; Archive the unarchived redo logs so that the redo required to recover the tablespace backup is archived.

NEW QUESTION 145

- (Topic 6)

Your production database is running in archivelog mode and you are using recovery manager (RMAN) with recovery catalog to perform the database backup at regular intervals. When you attempt to restart the database instance after a regular maintenance task on Sunday, the database fails to open displaying the message that the data file belonging to the users tablespace are corrupted.

The steps to recover the damaged data files are follows:

1. Mount the database
2. Open the database
3. Recover the data file
4. Restore the data file
5. Make the data file offline
6. Make the data file online

Which option identifies the correct sequence that you must use to recover the data files?

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

Answer: D

NEW QUESTION 146

- (Topic 6)

Your ARCHIVELOG mode database has lost three datafiles and shut down. One is assigned to the SYSTEM tablespace and two are assigned to the USERS tablespace. You can choose from the following steps to recover your database:

- a: Restore the three database datafiles that were lost.
- b: Issue the Startup Mount command to mount the database.
- c: Issue the alter database open command.
- d: Issue the alter database open resetlogs command.
- e: Recover the database using the recover database command.
- f: Recover the datafiles with the recover datafile command.
- g: Take the datafiles offline.

Which is the correct order of these steps in this case?

- A. a, b, e, c
- B. b, e, d
- C. a, b, d, c
- D. b, g, c, f
- E. a, b, d, f

Answer: A

Explanation:

Because the system critical data files is damaged, so that the RESETLOGS option is not applied. four steps to recover the system critical data (SYSTEM tablespace or the tablespace with UNDO):

1. SHUTDOWN ABORT, if the instance is started.
2. MOUNT the instance
3. restore and recover the damaged data files;
4. OPEN the database

NEW QUESTION 150

- (Topic 7)

When running the tablespace point-in-time command

recover tablespace users

until time '10/06/2008:22:42:00'

auxiliary destination 'c:\oracle\auxiliary'; You receive the following error:

```

RMAN-00571: =====
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====
RMAN-00571: =====
RMAN-03002: failure of recover command at 10/08/2008 16:00:30
RMAN-20202: Tablespace not found in the recovery catalog
RMAN-06019: could not translate tablespace name "USERS"

```

What is the likely cause of the error?

- A. The database is in ARCHIVELOG mode.
- B. There is not a current backup of the database available.
- C. The USERS tablespace has dependent objects in other tablespaces and can not be a part of a TSPITR alone.
- D. The USERS tablespace is not eligible for TSPITR because it has invalid objects.
- E. The recover tablespace command is incorrect and generates the error.

Answer: B

NEW QUESTION 155

- (Topic 7)

You have control-file autobackups enabled. When starting your database from SQL*Plus, you receive the following error message:

SQL> startup

ORA-01078: failure in processing system parameters LRM-00109: could not open parameter file

'C:\ORACLE\PRODUCT\11.1.0\DB_1\DATABASE

\INITORCL.ORA'

Using RMAN, how would you respond to this error?

- A. Issue the startup nomount command and then issue the restore parameter file command from the RMAN prompt.
- B. Issue the startup nomount command and then issue the restore spfile command from the RMAN prompt.
- C. Issue the startup nomount command and then issue the restore spfile from autobackup command from the RMAN prompt.
- D. Issue the startup nomount command and then issue the restore spfile from backup command from the RMAN prompt.

E. Issue the restore spfile from autobackup command from the RMAN prompt.

Answer: C

NEW QUESTION 159

- (Topic 7)

Which commands are used for RMAN database recovery? (Choose all that apply.)

- A. restore
- B. repair
- C. copy
- D. recover
- E. replace

Answer: AD

NEW QUESTION 163

- (Topic 7)

When performing a full database disaster recovery with RMAN, in what order would you execute these steps?

- A. Restore the control file from autobackups.
- B. Run the RMAN restore and recover command.
- C. Restore the database spfile from autobackups.
- D. Make the RMAN backup set pieces available.
- E. Open the database with the alter database open resetlogs command.
- F. Open the database with the alter database open command.
- G. a, b, c, d, e, f
- H. c, d, a, b, f
- I. d, c, a, b, f
- J. d, b, d, c, e
- K. d, c, a, b, e

Answer: E

NEW QUESTION 164

- (Topic 7)

You are using Recovery Manager (RMAN) for backup and recovery operations with a recovery catalog. You have been taken database backups every evening. On November 15, 2007, at 11:30 AM, you were informed that the USER_DATA tablespace was accidentally dropped. On investigation, you found that the tablespace existed until 11:00 AM, and important transactions were done after that.

So you decided to perform incomplete recovery until 11:00 AM. All the archive logs needed to perform recovery are intact. In NOMOUNT state you restored the control file that has information about the USER_DATA tablespace from the latest backup. Then you mounted the database. Identify the next set of commands that are required to accomplish the task?

- A. RMAN> run
{
 SET UNTIL TIME 'Nov 15 2007 11:00:00';
 RESTORE DATABASE;
 RECOVER DATABASE;
}
- B. RMAN> run
{
 SET UNTIL TIME 'Nov 15 2007 11:00:00';
 RESTORE DATABASE;
 RECOVER DATABASE USING BACKUP CONTROLFILE;
}
- C. RMAN> run
{
 RESTORE DATABASE;
 RECOVER DATABASE UNTIL TIME 'Nov 15 2007 11:00:00';
}
- D. RMAN> run
{
 RESTORE TABLESPACE user_data;
 RECOVER TABLESPACE user_data UNTIL TIME 'Nov 15 2007 11:00:00';
}

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: A

NEW QUESTION 166

- (Topic 7)

View the Exhibit to examine the error during the database startup. You open an RMAN session for the database instance. To repair the failure, you executed the

following as the first command in the RMAN session:

RMAN> REPAIR FAILURE;

Which statement describes the consequence of the command? Exhibit:

```
SQL> 3STARTUP
Total System Global Area 426864640 bytes
Fixed Size 1300352 bytes
Variable Size 180357248 bytes
Database Buffers 239075328 bytes
Redo Buffers 6131712 bytes
Database mounted
ORA-01157: cannot identify/lock data file 4 - see DBWR trace file
ORA-01110: data file 4: '/u01/app/oracle/oradata/orcl/users01.dbf'
```

- A. The command performs the recovery and closes the failures.
- B. The command only displays the advice and the RMAN script required for repair.
- C. The command produces an error because the ADVISE FAILURE command has not been executed before the REPAIR FAILURE command.
- D. The command executes the RMAN script to repair the failure and removes the entry from the Automatic Diagnostic Repository (ADR).

Answer: C

NEW QUESTION 169

- (Topic 8)

Which of the following are included in a transportable tablespace set? (Choose all that apply.)

- A. The datafiles that make up a self-contained group of tablespaces required for copy
- B. The system tablespace
- C. An export of the tablespace metadata
- D. The spfile
- E. All of the above

Answer: AC

NEW QUESTION 171

- (Topic 8)

Examine the following scenario: The target database instance is running. The most recent backup available for the target database was taken two days ago. Log files switches have occurred in last two days. The target database is duplicated on the same host, using the Recovery Manager (RMAN) duplicate command as follows:

RMAN> RUN

```
{
ALLOCATE AUXILIARY CHANNEL aux 1 DEVICE TYPE DISK; DUPLICATE TARGET DATABASE TO auxdb;
}
```

Which statement is true about the duplicate database in this scenario?

- A. It contains data till the last backup
- B. It contains all data from target database until the current time
- C. It contains all data from only the committed transactions in the target database
- D. It contains all data except that which is used by the transactions in the current online redo file of target database

Answer: D

NEW QUESTION 173

- (Topic 8)

You are managing the APPPROD database as a DBA. You plan to duplicate this database in the same system with the name DUPDB. You issued the following RMAN commands to create a duplicate database:

```
RMAN> CONNECT target sys/sys@APPPROD
RMAN> DUPLICATE TARGET DATABASE TO dupdb FROM ACTIVE DATABASE
DB_FILE_NAME_CONVERT '/oracle/oradata/prod/', '/scratch/oracle/oradata/
dupdb/'
SPILE
PARAMETER_VALUE_CONVERT '/oracle/oradata/prod/', '/scratch/oracle/oradata/
dupdb/'
SET SGA_MAX_SIZE = '300M'
SET SGA_TARGET = '250M'
SET LOG_FILE_NAME_CONVERT '/oracle/oradata/prod/redo/', '/scratch/oracle/
oradata/dupdb/redo/';
```

Which three are the prerequisites for the successful execution of the above command? (Choose three.)

- A. The source database should be open.
- B. The target database should be in ARCHIVELOG mode if it is open.
- C. RMAN should be connected to both the instances as SYSDBA.
- D. The target database backups should be copied to the source database backup directories.
- E. The password file must exist for the source database and have the same SYS user password as the target.

Answer: BCE

NEW QUESTION 177

- (Topic 9)

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

NEW QUESTION 178

- (Topic 9)

Why should you back up a duplicated tablespace after a TSPITR is complete?

- A. The tablespace cannot be duplicated or restored to any point in time after the duplication.
- B. The tablespace cannot be duplicated or restored to the point in time before the duplication.
- C. The entire database cannot be restored after a TSPITR, so a backup is required.
- D. You cannot bring the tablespace online until its been backed up.
- E. There is no requirement to do so, as RMAN will back up the tablespace after the TSPITR.

Answer: B

NEW QUESTION 179

- (Topic 9)

Which command would correctly start a TSPITR of the USERS tablespace?

- A. `recover tablespace users until time '10/06/2008:22:42:00' auxiliary 'c:\oracle\auxiliary';`
- B. `recover tablespace users time '10/06/2008:22:42:00' auxiliary destination 'c:\oracle\auxiliary';`
- C. `recover tablespace users to point-in-time '10/06/2008:22:42:00' auxiliary destination 'c:\oracle\auxiliary';`
- D. `recover tablespace users except time '10/06/2008:22:42:00' auxiliary destination 'c:\oracle\auxiliary';`
- E. `recover tablespace users until time '10/06/2008:22:42:00' auxiliary destination 'c:\oracle\auxiliary';`

- A. Option A
- B. Option B
- C. Option C
- D. Option D
- E. Option E

Answer: E

NEW QUESTION 180

- (Topic 10)

You performed the RMAN database backup having a backupset key number 231 with the KEEP FOREVER option.

After some days, you want to change the status of the database backup and you issued the following command:

```
RMAN>CHANGE BACKUPSET 231 NOKEEP;
```

What is the implication of this command?

- A. The backup is deleted.
- B. The backup is marked unavailable.
- C. The backup overrides the backup retention policy.
- D. the backup becomes eligible for deletion according to the existing retention policy

Answer: D

NEW QUESTION 185

- (Topic 10)

You want to back up your 100-GB database on a remote tape device. You are required to ensure that minimum network bandwidth is consumed while transferring the backups to the tape device. The current consumption of your CPU is approximately 40 percent.

Which type of backup should you perform?

- A. standard backup set
- B. image copy
- C. compressed backup
- D. user-managed backup

Answer: C

NEW QUESTION 189

- (Topic 10)

You have enabled backup optimization in RMAN. You issue the following RMAN command to configure a redundancy-based retention policy:

CONFIGURE RETENTION POLICY TO REDUNDANCY 3;
Which statement is true?

- A. The command fails because you cannot configure a redundancy-based retention policy when backup optimization is enabled
- B. Backup optimization is performed, but RMAN considers the redundancy-based retention policy when it determines which datafiles should be backed up
- C. Backup optimization is permanently disabled
- D. Backup optimization is temporarily disabled because a redundancy-based retention policy is specified

Answer: B

NEW QUESTION 190

- (Topic 11)

Which of the following can be used in conjunction with a Flashback Versions Query to filter the results? (Choose all that apply.)

- A. A range of SCN values
- B. A list of SCN values
- C. A starting and ending timestamp
- D. Minimum and maximum sequence values
- E. A list of sequence values

Answer: AC

NEW QUESTION 195

- (Topic 11)

You plan to use Flashback Drop feature to recover a dropped table SALES_EMP. No other table with the same name exists in the schema. You query RECYCLEBIN and find multiple entries for the SALES_EMP table as follows: You then issue the following statement to recover the table:

SQL> SELECT object_name, original_name, droptime FROM recyclebin;

OBJECT_NAME	ORIGINAL_NAME	DROPTIME
BIN\$/m0DrBV9RFGOAA53dC-FPW==\$0	SALES_EMP	2007-12-07:11:08:55
BIN\$2DeIssLeQTqgH/n80Rm2JQ==\$0	SALES_EMP	2007-12-07:11:11:38
BIN\$UuqroNodQy6ouDtaA+XOVw==\$0	SALES_EMP	2007-12-07:11:08:18

SQL> FLASHBACK TABLE sales_emp TO BEFORE DROP;

What would be the outcome of the precedent statement?

- A. It retrieves the latest version of the table from the recycle bin
- B. It retrieves the oldest version of the table from the recycle bin
- C. It retrieves the version of the table for which undo information is available
- D. It returns an error because the table name is not specified as per the names in the OBJECT_NAME column

Answer: A

NEW QUESTION 200

- (Topic 11)

Which pseudocolumn could you use to identify a unique row in a Flashback Versions Query?

- A. XID
- B. VERSIONS_PK
- C. VERSIONS_XID
- D. VERSIONS_UNIQUE

Answer: C

NEW QUESTION 203

- (Topic 11)

Before a Flashback Table operation, you execute the following command: ALTER TABLE employees ENABLE ROW MOVEMENT; Why would you need this to be executed?

- A. Because row IDs may change during the flashback operation
- B. Because the object number changes after the flashback operation
- C. Because the rows are retrieved from the recycle bin during the flashback operation
- D. Because the table is moved forward and back to a temporary during the flashback operation

Answer: A

NEW QUESTION 205

- (Topic 11)

You executed the following commands in a database session:

```
SQL> SELECT object_name, original_name FROM user_recyclebin;
```

```
OBJECT_NAME                                ORIGINAL_NAME
-----
BIN$QJwAldMyB1LgQJYK+xUptw= = s0 MYSpace
```

```
SQL> CREATE TABLE mvspace AS SELECT * FROM my region;
Table created.
```

```
SQL> SELECT object_name, original_name FROM user_recyclebin;
No rows selected
```

What could have caused the recycle bin to clean up?

- A. There is demand for space from the new table
- B. The undo tablespace does not have sufficient free space
- C. The new table name is the same as the table name in the recycle bin
- D. The temporary tablespace that is assigned to you does not have sufficient free space

Answer: A

NEW QUESTION 208

- (Topic 11)

You executed the following command to drop a user: DROP USER scott CASCADE;
Which two statements regarding the above command are correct? (Choose two.)

- A. All the objects of scott are moved to the Recycle Bin.
- B. Any objects in the Recycle Bin belonging to scott are purged.
- C. All the objects owned by scott are permanently dropped from the database.
- D. All the objects of scott in the Recycle Bin must be purged before executing the DROP command.
- E. Any objects in the Recycle Bin belonging to scott will not be affected by the above DROP command.

Answer: BC

NEW QUESTION 213

- (Topic 11)

View the following SQL statements: Transaction T1

```
SQL> INSERT INTO hr.regions 2 VALUES (5,'Pole');
3 COMMIT;
```

Transaction T2

```
SQL> UPDATE hr.regions 2 SET region_name='Poles'
3 WHERE region_id = 5;
4 COMMIT;
```

Transaction T3

```
SQL> UPDATE hr.regions
2 SET region_name='North and South Poles'
3 WHERE region_id = 5;
```

You want to back out transaction T2. Which option would you use?

- A. It is possible, but transaction T3 also backs out.
- B. It is possible with the NOCASCADE_FORCE option.
- C. It is possible with the NONCONFLICT_ONLY option.
- D. It is not possible because it has conflicts with transaction T3.

Answer: B

Explanation:

Table 12-2 Flashback TRANSACTION_BACKOUT Options

CASCADE

Backs out specified transactions and all dependent transactions in a post-order fashion (that is, children are backed out before parents are backed out).

Without CASCADE, if any dependent transaction is not specified, an error occurs.

NOCASCADE

Default. Backs out specified transactions, which are expected to have no dependent transactions. First dependent transactions causes an error and appears in *_FLASHBACK_TXN_REPORT.

NOCASCADE_FORCE

Backs out specified transactions, ignoring dependent transactions. Server runs undo SQL statements for specified transactions in reverse order of commit times. If no constraints break and you are satisfied with the result, you can commit the changes; otherwise, you can roll them back.

NONCONFLICT_ONLY

Backs out changes to nonconflicting rows of the specified transactions. Database remains consistent, but transaction atomicity is lost.

NEW QUESTION 215

- (Topic 11)

Examine the output of the query that you executed to list the objects in the recycle bin:

```
SQL> SELECT original_name, droptime, dropscn FROM user_recyclebin;
```

ORIGINAL_NAME	DROPTIME	DROPSCN
SALES_TAB	2007-12-11:13:37:11	4472036
SALES_TAB	2007-12-11:13:49:30	4472988
SALES_TAB	2007-12-11:13:55:39	4473100

You verified that no table named SALES_TAB exists in the schema. Then you executed the following command to purge the objects in the recycle bin:

```
SQL> PURGE TABLE sales_tab;
```

What would be the outcome of this command?

- A. All three tables in the recycle bin are purged
- B. Only the table with the oldest DROPSCN is purged
- C. The command returns an error because multiple entries with the same name exist in the recycle bin
- D. Only the table with the latest DROPSCN is purged

Answer: B

NEW QUESTION 218

- (Topic 11)

The EMP table exists in your schema. You want to execute the following query:

```
SELECT ename, sal FROM emp
```

```
AS OF TIMESTAMP (SYSTIMESTAMP - INTERVAL '6' MINUTE) WHERE ename = 'ALLEN';
```

What are the minimum requirements for the statement to execute successfully? (Choose all that apply)

- A. ARCHIVELOG mode must be enabled
- B. Row Movement must be enabled for the table
- C. FLASHBACK must be set to ON for the database
- D. The UNDO_MANAGEMENT parameter must be set to AUTO
- E. The UNDO_RETENTION parameter must be set appropriately

Answer: DE

NEW QUESTION 223

- (Topic 12)

You want to track and store all transactional changes to a table over its lifetime.

To accomplish this task, you enabled Flashback Data Archive with the retention of 5 years. After some time, the business requirement changed and you were asked to change the retention from 5 years to 3 years.

To accomplish this, you issued the following command:

```
ALTER FLASHBACK ARCHIVE fla1 MODIFY RETENTION 3 YEAR;
```

What is the implication of this command?

- A. The command produces an error because the retention time cannot be reduced.
- B. All historical data is retained but the subsequent flashback data archives are maintained for only three years.
- C. All historical data is purged and the new flashback data archives are maintained for three years.
- D. All historical data older than three years is purged from the flashback archive FLA1.

Answer: D

NEW QUESTION 228

- (Topic 12)

What does the DB_FLASHBACK_RETENTION_TARGET parameter configure?

- A. An upper limit on how far you can flash back the database, depending on the information in the redo logs
- B. An upper limit on how far you can flash back the database, depending on the information in the undo tablespace
- C. The amount of time for which the flashback data is to be kept in the flash recovery area, provided that there is enough space
- D. The amount of time for which the flashback data is guaranteed to be kept in the undo tablespace, provided that there is enough space

Answer: C

NEW QUESTION 232

- (Topic 12)

You need to maintain a record of all transactions on some tables for at least three years. Automatic undo management is enabled for the database.

What must you do accomplish this task?

- A. Enable supplemental logging for the database.
- B. Specify undo retention guarantee for the database
- C. Create Flashback Data Archive in the tablespace where the tables are stored.
- D. Create Flashback Data Archive and enable Flashback Data Archive for specific tables

Answer: D

NEW QUESTION 237

- (Topic 13)

Which statement is true for enabling Enterprise Manager Support Workbench in Oracle Database 11g to upload the physical files generated by Incident Packaging

Service (IPS) to MetaLink?

- A. The database must be running in ARCHIVELOG mode.
- B. No special setup is required, and the feature is enabled by default.
- C. The path for the Automatic Diagnostic Repository (ADR) must be configured with the DIAGNOSTIC_DEST initialization parameter.
- D. The Enterprise Manager Support Workbench can be enabled only if the background process manageability monitor (MMON) is configured.
- E. Select the Enable option in the Oracle Configuration Manager Registration window during the installation of the Oracle Database 11g software, provide valid MetaLink credentials and select license agreement.

Answer: E

NEW QUESTION 240

- (Topic 13)

You plan to collect the Automatic Workload Repository (AWR) data every Monday morning for a month. You want Oracle Database to automatically create a baseline every Monday and remove the old baseline. What is the correct action to achieve this?

- A. Create and populate a SQL Tuning Set from the AWR on every Monday.
- B. Change the RETENTION setting for the AWR snapshots to 7 days on Monday.
- C. Create a repeating baseline template.
- D. Insert a finding directive for future ADDM tasks.

Answer: C

NEW QUESTION 242

- (Topic 13)

Identify the three predefined server-generated alerts. (Choose three.)

- A. Drop User
- B. Tablespace Space Usage
- C. Resumable Session Suspended
- D. Recovery Area Low On Free Space
- E. SYSTEM Tablespace Size Increment

Answer: BCD

NEW QUESTION 243

- (Topic 13)

Which tasks can be accomplished using the Enterprise Manager Support Workbench in Oracle Database 11g? (Choose all that apply.)

- A. Generate reports on data failure such as data file failures.
- B. You can track the Service Request (SR) and implement repairs.
- C. You can package and upload diagnostic data to Oracle Support.
- D. You can manually run health checks to gather diagnostic data for a problem.

Answer: BCD

NEW QUESTION 247

- (Topic 13)

Which of the following initialization parameters have been deprecated in Oracle 11g because of the introduction of the Automatic Workload Repository? (Choose all that apply.)

- A. BACKGROUND_DUMP_DEST
- B. FOREGROUND_DUMP_DEST
- C. CORE_DUMP_DEST
- D. USER_DUMP_DEST
- E. DIAGNOSTIC_DEST
- F. All of the above

Answer: ACD

NEW QUESTION 252

- (Topic 13)

Which statements are true regarding the concept of problems and incidents in the fault diagnosability infrastructure for Oracle Database 11g? (Choose all that apply.)

- A. Only the incident metadata and dumps are stored in the Automatic Diagnostic Repository (ADR).
- B. The problem key is the same as the incident number.
- C. The database sends an incident alert to the Oracle Enterprise Manager Database Home page.
- D. Every problem has a problem key, which is a text string that describes the problem.
- E. The database makes an entry into the alert log file when problems and incidents occur.

Answer: CDE

Explanation:

Reference at here

Fault Diagnosability Infrastructure Overview

The fault diagnosability infrastructure aids in preventing, detecting, diagnosing, and resolving problems. The problems that are targeted in particular are critical errors such as those caused by code bugs, metadata corruption, and customer data corruption.

When a critical error occurs, it is assigned an incident number, and diagnostic data for the error (such as trace files) are immediately captured and tagged with this number. The data is then stored in the Automatic Diagnostic Repository (ADR)—a file-based repository outside the database—where it can later be retrieved by incident number and analyzed. About Incidents and Problems

A problem is a critical error in a database instance, Oracle Automatic Storage Management (Oracle ASM) instance, or other Oracle product or component. Critical errors manifest as internal errors, such as ORA-00600, or other severe errors, such as ORA-07445 (operating system exception) or ORA-04031 (out of memory in the shared pool). Problems are tracked in the ADR. Each problem has a problem key, which is a text string that describes the problem. It includes an error code (such as ORA 600) and in some cases, one or more error parameters.

An incident is a single occurrence of a problem. When a problem (critical error) occurs multiple times, an incident is created for each occurrence. Incidents are timestamped and tracked in the Automatic Diagnostic Repository (ADR). Each incident is identified by a numeric incident ID, which is unique within the ADR. When an incident occurs, the database:

- ? Makes an entry in the alert log.
- ? Sends an incident alert to Oracle Enterprise Manager (Enterprise Manager).
- ? Gathers first-failure diagnostic data about the incident in the form of dump files (incident dumps).
- ? Tags the incident dumps with the incident ID.
- ? Stores the incident dumps in an ADR subdirectory created for that incident.

NEW QUESTION 254

- (Topic 14)

Examine the exhibit to view the parameters set in your parameter file. You restart the instance. To what value will the MEMORY_MAX_TARGET parameter be set by default? Exhibit:

ORACLE Enterprise Manager 11g
Database Control

Database Instance: iocp >

Initialization Parameters

Current | **SPFile**

The parameter values listed here are from the SPFILE /home/oracle/app/oracle/p

Name: target Basic: All ▼ Dynamic: All ▼ Category: Memory

Filter on a name or partial name

Apply changes in SPFile mode to the current running instance(s). For static para

Reset

Select	Name ▲	Help	Revisions	Value
<input checked="" type="radio"/>	memory_max_target			
<input type="radio"/>	memory_target			600M
<input type="radio"/>	pga_aggregate_target			120M
<input type="radio"/>	sga_target			320M

Current | **SPFile**

- A. 120M
- B. 320M
- C. 440M
- D. 600M

Answer: D

NEW QUESTION 257

- (Topic 14)

Which two initialization parameters would you set to enable Automatic Shared Memory Management? (Choose two.)

- A. set SHARED_POOL_SIZE to zero
- B. set STATISTICS_LEVEL to BASIC
- C. set SGA_TARGET to a non-zero value
- D. set DB_CACHE_SIZE to a non-zero value
- E. set STATISTICS_LEVEL to TYPICAL or ALL

Answer: AC

NEW QUESTION 260

- (Topic 14)

When manually configuring the SGA, which of the following parameter changes requires an instance restart to take effect?

- A. DB_CACHE_SIZE
- B. SHARED_POOL_SIZE
- C. LARGE_POOL_SIZE
- D. JAVA_POOL_SIZE
- E. SGA_MAX_SIZE
- F. SORT_AREA_SIZE

Answer: E

NEW QUESTION 261

- (Topic 14)

Your database interface is running. A user SCOTT starts a SQL *Plus session, and issues the following query:

SQL> SELECT * FROM sales;

Which process would retrieve the result from the database and return it to the client program?

- A. User process
- B. Server process
- C. System Monitor (SMON)
- D. Process Monitor (PMON)
- E. Checkpoint process (CKPT)

Answer: B

NEW QUESTION 262

- (Topic 14)

To manually configure the SGA components using Oracle Enterprise Manager Memory Advisor, you can set values for which of the following initialization parameters? (Choose all that apply.)

- A. DB_CACHE_SIZE
- B. SHARED_POOL_SIZE
- C. LARGE_POOL_SIZE
- D. JAVA_POOL_SIZE
- E. SGA_MAX_SIZE
- F. SORT_AREA_SIZE

Answer: ABCD

NEW QUESTION 266

- (Topic 14)

What Oracle process runs when the database is in ARCHIVELOG mode but not when it is in NOARCHIVELOG mode?

- A. MMON
- B. LGWR
- C. ARCH
- D. ARWR
- E. COPY


Answer: C

NEW QUESTION 267

- (Topic 15)

View the Exhibit-1 to observe the maintenance window property. View the Exhibit-2 to examine the output of the query.

Which two statements describe the conclusions? (Choose two.) EM-window (exhibit):



sql-dba_autotask_client (exhibit):

```
SQL> SELECT client_name,
2 mean_job_duration,
3 resource_percentage
4 FROM DBA_AUTOTASK_CLIENT;
```

CLIENT NAME	MEAN_JOB_DURATION	RESOURCE_PERCENTAGE
auto optimized stats collection	+0000000000 00:18:41.000000000	25
auto space advisor	+0000000000 00:02:36.666666666	25
sql tuning advisor	+0000000000 00:04:36.500000000	25

- A. RESOURCE_PERCENTAGE should be increased.
- B. The repeat time for the window should be decreased.
- C. RESOURCE_PERCENTAGE should be decreased.
- D. The window duration should be increased.

Answer: AD

NEW QUESTION 268

- (Topic 15)

While deploying a new application module, the software vendor ships the application software along with appropriate SQL plan baselines for the new SQLs being introduced. Which two statements describe the consequences? (Choose two.)

- A. The plan baselines can be evolved over time to produce better performance.
- B. The newly generated plans are directly placed into the SQL plan baseline without being verified.
- C. The new SQL statements initially run with the plans that are known to produce good performance under standard test configuration.
- D. The optimizer does not generate new plans for the SQL statements for which the SQL plan baseline has been imported.

Answer: AC

Explanation:

Purpose of SQL Plan Baselines

Common scenarios where SQL plan management can improve or preserve SQL performance include:

A database upgrade that installs a new optimizer version usually results in plan changes for a small percentage of SQL statements. Most of these plan changes result in either no performance change or improvement. However, some plan changes may cause performance regressions. SQL plan baselines significantly minimize potential regressions resulting from an upgrade.

Ongoing system and data changes can impact plans for some SQL statements, potentially causing performance regressions. SQL plan baselines help minimize performance regressions and stabilize SQL performance.

Deployment of new application modules means introducing new SQL statements into the database. The application software may use appropriate SQL execution plans developed in a standard test configuration for the new statements. If the system configuration is significantly different from the test configuration, then the database can evolve SQL plan baselines over time to produce better performance.

NEW QUESTION 272

- (Topic 15)

Which two statements about the SQL Management Base (SMB) are true? (Choose two.)

- A. It contains only SQL profiles generated by SQL Tuning Advisor.
- B. It stores plans generated by the optimizer using a stored outline.
- C. It is part of the data dictionary and stored in the SYSAUX tablespace.
- D. It is part of the data dictionary and stored in the SYSTEM tablespace.
- E. It contains the statement log, the plan history, plan baselines, and SQL profiles.

Answer: CE

NEW QUESTION 274

- (Topic 15)

Which of these appropriately describes the results of a manual SQL Tuning Advisor task?

- A. A list of SQL statements and recommendations for tuning
- B. A list of SQL statements that have been tuned by the Advisor, with before and after metrics
- C. Graphs showing the actual performance improvement made by the Advisor after it implemented the recommended changes
- D. All of the above

Answer: A

NEW QUESTION 276

- (Topic 15)

Which of the following represents the correct sequence of events for Database Replay?

- A. Capture, analyze, preprocess, replay
- B. Capture, preprocess, analyze, replay
- C. Capture, preprocess, replay, analyze
- D. Analyze, capture, preprocess, replay
- E. None of the above

Answer: C

Explanation:

Database Replay

Database Replay allows workloads to be captured from production systems and re-executed with high fidelity on test copies of production databases. This enables detailed analysis of how the proposed changes may affect production systems; for instance, patching or upgrading database software.

Task Name	Description	Go to Task
1 Capture Workload	Capture a workload from the production environment. This can be scheduled to accommodate a database restart if desired.	
2 Preprocess Workload	Preprocessing prepares a captured workload for replay. You must do this once for every captured workload. Preprocessing is best performed in the test database. The captured workload must be accessible from the test database.	
3 Replay Workload	Replay the preprocessed workload on a test copy of the production database.	

C:\Users\albo\Desktop\1-1.jpg

NEW QUESTION 281

- (Topic 15)

You run the SQL Tuning Advisor (STA) to tune a SQL statement that is part of a fixed SQL plan baseline. The STA generates a SQL profile for the SQL statement, which recommends that you accept the profile.

Which statement is true when you accept the suggested SQL profile?

- A. The tuned plan is not added to the SQL plan baseline.
- B. The tuned plan is added to the fixed SQL plan baseline as a fixed plan.
- C. The tuned plan is added to the fixed SQL plan baseline as a nonfixed plan.
- D. The tuned plan is added to a new nonfixed SQL plan baseline as a nonfixed plan.

Answer: C

Explanation:

15.4 Using Fixed SQL Plan Baselines (Refer to here)

When you tune a SQL statement with a fixed SQL plan baseline using SQL Tuning Advisor, a SQL profile recommendation has special meaning. When the SQL profile is accepted, the database adds the tuned plan to the fixed SQL plan baseline as a non-fixed plan. However, as described above, the optimizer does not use the tuned plan when a reproducible fixed plan is present. Therefore, the benefit of SQL tuning may not be realized. To enable the use of the tuned plan, manually alter the tuned plan to a fixed plan by setting its FIXED attribute to YES.

NEW QUESTION 285

- (Topic 15)

Identify the two direct sources from where SQL plans can be loaded into the SQL plan baselines. (Choose two.)

- A. Cursor cache
- B. Stored outline
- C. SQL Tuning Set
- D. Automatic Workload Repository (AWR) snapshots

Answer: AC

Explanation:

The AWR snapshots needs to be loaded to STS at first, then load to SQL plan.

----- (Link)

You can perform manual plan loading by:

Loading Plans from SQL Tuning Sets and AWR Snapshots

To load plans from a SQL tuning set, use the LOAD_PLANS_FROM_SQLSET function of the DBMS_SPM package.

The following example loads the plans stored in the SQL tuning set named tset1:

```
DECLARE
my_plans PLS_INTEGER; BEGIN
my_plans := DBMS_SPM.LOAD_PLANS_FROM_SQLSET( sqlset_name => 'tset1'); END;
/
```

To load plans from Automatic Workload Repository (AWR), load the plans stored in AWR snapshots into a SQL tuning set before using the LOAD_PLANS_FROM_SQLSET function as described in this section.

Loading Plans from the Shared SQL Area

To load plans from the shared SQL area, use the LOAD_PLANS_FROM_CURSOR_CACHE function of the DBMS_SPM package. In the following example, Oracle Database loads the plans located in the shared SQL area for the SQL statement identified by its sql_id:

```
DECLARE
my_plans PLS_INTEGER; BEGIN
my_plans := DBMS_SPM.LOAD_PLANS_FROM_CURSOR_CACHE( sql_id => '99twu5t2dn5xd');
END;
/
```

NEW QUESTION 289

- (Topic 15)

Performance divergence indicated in the Workload Replay report is most likely due to what?

- A. DML and SQL statement results that do not match between the capture and replay systems
- B. When errors that occur in the capture system don't occur in the replay system
- C. Top SQL statements
- D. Infrastructure or system-configuration differences

E. Time-of-day differences between capture and replay systems

Answer: D

NEW QUESTION 293

- (Topic 15)

Which two changes and their effect on the system can be tested by using the Database Replay feature? (Choose two.)

- A. multiplexing of the control file
- B. adding the redo log member to the database
- C. database and operating system upgrades
- D. changing the database storage to ASM-managed storage

Answer: CD

NEW QUESTION 296

- (Topic 15)

To generate recommendations to improve the performance of a set of SQL queries in an application, you execute the following blocks of code:

```
BEGIN
  dbms_advisor.create_task(dbms_advisor.sqlaccess_advisor, 'TASK1');
END;
/
BEGIN
  dbms_advisor.set_task_parameter('TASK1', 'ANALYSIS_SCOPE', 'ALL');
  dbms_advisor.set_task_parameter('TASK1', 'MODE', 'COMPREHENSIVE');
END;
/
BEGIN
  dbms_advisor.execute_task('TASK1');
  dbms_output.put_line(dbms_advisor.get_task_script('TASK1'));
END;
/
```

The blocks of code execute successfully; however, you do not get the required outcome. What could be the reason?

- A. A template needs to be associated with the task.
- B. A workload needs to be associated with the task.
- C. The partial or complete workload scope needs to be associated with the task.
- D. The type of structures (indexes, materialized views, or partitions) to be recommended need to be specified for the task.

Answer: B

NEW QUESTION 298

- (Topic 15)

View the Exhibit1 to examine the series of SQL commands. View the Exhibit2 to examine the plans available in the SQL plan baseline. The baseline in the first row of the Exhibit is created when OPTIMIZER_MODE was set to FIRST_ROWS. Which statement is true if the SQL query in exhibit1 is executed again when the value of OPTIMIZER_MODE is set to FIRST_ROWS?

```
SQL> SELECT signature, sql_handle, plan_name, origin, enabled,
accepted, fixed, autopurge
FROM dba_sql_plan_baselines;
```

SIGNATURE	SQL_HANDLE	PLAN_NAME	ORIGIN	ENABLED	ACCEPTED	FIXED
8.062E+18	SYS_SQL_6fa2	SYS_SQL_PLAN_1ea	AUTO-CAPTURE	YES	NO	NO
8.062E+18	SYS_SQL_6fe2	SYS_SQL_PLAN_4be	AUTO-CAPTURE	YES	YES	NO
...						
...						
...						

parameter-optimizer (exhibit):

```
SQL> SHOW PARAMETER OPTIMIZER

NAME                                TYPE                                VALUE
-----                                -
optimizer_capture_sql_plan_baselines boolean                             TRUE
optimizer_dynamic_sampling           integer                             2
optimizer_features_enable            string                              11.1.0.6
optimizer_index_caching               integer                             0
optimizer_index_cost_adj              integer                             100
optimizer_mode                        string                              ALL_ROWS
optimizer_secure_view_merging         boolean                             TRUE
optimizer_use_invisible_indexes       boolean                             FALSE
optimizer_use_pending_statistics      boolean                             FALSE
optimizer_use_sql_plan_baselines      boolean                             TRUE

SQL> SELECT * FROM sh.sales WHERE quantity_sold > 40 ORDER BY prod_id;
SQL> SELECT * FROM sh.sales WHERE quantity_sold > 40 ORDER BY prod_id;
SQL> ALTER SESSION SET OPTIMIZER_MODE=FIRST_ROWS;
SQL> SELECT * FROM sh.sales WHERE quantity_sold > 40 ORDER BY prod_id;
```

- A. The optimizer uses a new plan because none of the plans in the exhibit2 are fixed plans.
- B. The optimizer uses the plan in the second row of the exhibit2 because it is an accepted plan.
- C. The optimizer uses the plan in the first row of the exhibit2 because it is the latest generated plan.
- D. The optimizer uses the plan in the first row of the exhibit2 because OPTIMIZER_MODE was set to FIRST_ROW during its creation.

Answer: B

Explanation:

Setting the OPTIMIZER_MODE Initialization Parameter(Link)

The OPTIMIZER_MODE initialization parameter establishes the default behavior for choosing an optimization approach for the instance.

OPTIMIZER_MODE Initialization Parameter Values

? ALL_ROWS, The optimizer uses a cost-based approach for all SQL statements in the session regardless of the presence of statistics and optimizes with a goal of best throughput (minimum resource use to complete the entire statement). This is the default value.

? FIRST_ROWS_n, The optimizer uses a cost-based approach, regardless of the presence of statistics, and optimizes with a goal of best response time to return the first n number of rows, where n equals 1, 10, 100, or 1000.

? FIRST_ROWS, The optimizer uses a mix of cost and heuristics to find a best plan for fast delivery of the first few rows.

Note that using heuristics sometimes leads the optimizer to generate a plan with a cost that is significantly larger than the cost of a plan without applying the heuristic. FIRST_ROWS is available for backward compatibility and plan stability; use FIRST_ROWS_n instead.

NEW QUESTION 301

- (Topic 15)

View the Exhibit for some of the parameter settings. You start a session and issue the following command:

```
SQL>CREATE INDEX emp_ename ON emp(ename) TABLESPACE users INVISIBLE;
```

What is the outcome of the above command? Exhibit:

NAME	TYPE	VALUE
optimizer_dynamic_sampling	integer	2
optimizer_features_enable	string	11.1.0.6
optimizer_use_invisible_indexes	boolean	FALSE
optimizer_index_caching	integer	0
optimizer_index_cost_adj	integer	100
skip_unusable_indexes	boolean	TRUE

- A. The index is not used by the optimizer but is maintained during DML operations.
- B. The index is not used by the optimizer and is not maintained during DML operations.
- C. The index is used by the optimizer only if a hint is specified in the query statement and is maintained during DML operations.
- D. The index is used by the optimizer only if a hint is specified in the query statement but is not maintained during DML operations.

Answer: A

NEW QUESTION 304

- (Topic 15)

Which of the following advisors within the Oracle advisory framework will analyze a single SQL statement and make recommendations for performance improvement?

- A. SQL Repair Advisor
- B. SQL Optimizer
- C. SQL Access Advisor
- D. SQL Tuning Advisor

Answer: D

NEW QUESTION 305

- (Topic 15)

View the Exhibit.

Examine the following command that is executed for the TRANSPORT table in the SH schema:

```
SQL> SELECT DBMS_STATS.CREATE_EXTENDED_STATS('sh', 'customers_obe',
'country_id, cust_state_province') FROM dual;
```

Which statement describes the significance of this command? Exhibit:

```
SQL> DESCRIBE CUSTOMERS_OBE
```

Name	Null?	Type
CUST_ID		NUMBER
CUST_FIRST_NAME	NOT NULL	VARCHAR2 (20)
CUST_LAST_NAME	NOT NULL	VARCHAR2 (40)
CUST_GENDER		CHAR (1)
CUST_YEAR_OF_BIRTH		NUMBER (4)
CUST_MARITAL_STATUS		VARCHAR2 (20)
CUST_STREET_ADDRESS	NOT NULL	VARCHAR2 (40)
CUST_POSTAL_CODE	NOT NULL	VARCHAR2 (10)
CUST_CITY	NOT NULL	VARCHAR2 (30)
CUST_STATE_PROVINCE		VARCHAR2 (40)
CUST_RY_ID	NOT NULL	CHAR (2)
CUST_MAIN_PHONE_NUMBER		VARCHAR2 (25)
CUST_INCOME_LEVEL		VARCHAR2 (30)
CUST_CREDIT_LIMIT		NUMBER
CUST_EMAIL		VARCHAR2 (30)

- A. It collects statistics into the pending area in the data dictionary.
- B. It creates a virtual hidden column in the CUSTOMERS_OBE table.
- C. It collects statistics with AUTO_SAMPLE_SIZE for ESTIMATE_PERCENT.
- D. It creates a histogram to hold skewed information about the data in the columns.

Answer: B

Explanation:

DBMS_STATS.CREATE_EXTENDED_STATS Function

Creates a virtual column for a user specified column group or an expression in a table This function creates a column statistics entry in the system for a user specified column

group or an expression in a table. Statistics for this extension will be gathered when user or auto statistics gathering job gathers statistics for the table. We call statistics for such an extension, "extended statistics". This function returns the name of this newly created entry for the extension.

NEW QUESTION 307

- (Topic 15)

Which of the following advisors is run in every maintenance window by the auto-task system?

- A. The Memory Advisor
- B. The SQL Tuning Advisor
- C. The Undo Advisor
- D. The SQL Access Advisor

Answer: B

Explanation:

About Automated Maintenance Tasks ([link](#))

Oracle Database has three predefined automated maintenance tasks:

? Automatic Optimizer Statistics Collection

? Automatic Segment Advisor

? Automatic SQL Tuning Advisor

By default, all three automated maintenance tasks are configured to run in all maintenance windows.

NEW QUESTION 310

- (Topic 15)

Which of the following identifies and creates an index to minimize the DB time for a particular SQL statement?

- A. The SGA Tuning Advisor
- B. The SQL Access Advisor
- C. The SQL Tuning Advisor
- D. The Memory Advisor

Answer: C

Explanation:

The SQL Access Advisor allows the DBA to gather global recommendations for a workload. The SQL Tuning advisor is more granular, tuning a single statement. The main functions of the SQL Access advisor is to recommend missing indexes and materialized views, but a comprehensive task analysis will also create SQL Profiles that can be used within the SQL Tuning advisor. The DBA defines the SQL used in the SQL Access Advisor task, and can choose current SQL, a user-defined set of SQL, a historical workload, or a hypothetical workload.

A hypothetical workload is very useful because the DBA need-only specify the tables that participate in the queries, and the SQL Access Advisor gathers the appropriate SQL statements to create the workload. The SQL Tuning Advisor (STA) is primarily designed to replace the manual tuning of SQL statements and speed up the overall SQL tuning process. The SQL Tuning Advisor studies poorly executing SQL statements and evaluates resource consumption in terms of CPU, I/O, and temporary space.

The SQL Tuning Advisor (STA) works with the Automatic Tuning Optimizer (ATO) to analyze historical SQL workload using data from the AWR, and it generates

recommendations for new indexes and materialized views that will reduce the disk I/O associated with troublesome SQL statements.
http://www.dba-oracle.com/oracle10g_tuning/t_sql_tuning_advisor.htm

NEW QUESTION 312

- (Topic 16)

Which steps are mandatory to enable Direct NFS?

1. Mount all required file systems using the kernel NFS driver.
2. Create an orafstab file containing the attributes for each NFS server to be accessed using Direct NFS.
3. Replace the ODM library libodm11.so_stub with libodm11.so.

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

Answer: B

NEW QUESTION 317

- (Topic 16)

Which two statements are true about the compressed backups in RMAN? (Choose two.)

- A. Compressed backups can only be taken on the tape drives.
- B. The binary compression creates some performance overhead during backup operation.
- C. The ZLIB compression algorithm can be used only if the COMPATIBLE initialization parameter is set to 11.1.0.
- D. The media manager compression for the tape drive should be enabled for taking compressed backups on the tape.

Answer: BD

NEW QUESTION 320

- (Topic 16)

Which statement describes the effect of table redefinition on the triggers attached to the table?

- A. All triggers on the table are invalidated and are automatically revalidated with the next DML execution on the table.
- B. All triggers on the table are invalidated and must be manually recompiled before the next DML execution on the table.
- C. All triggers on the table remain valid.
- D. Only triggers that are affected by the changes to the structure of the table are invalidated and automatically revalidated with the next DML execution on the table.

Answer: A

NEW QUESTION 324

- (Topic 16)

Following is the list of locations in random order where orafstab can be placed.

1. /etc/mtab
2. \$ORACLE_HOME/dbs/orafstab
3. /etc/orafstab

What is the sequence in which Direct NFS will search the locations?

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

Answer: C

NEW QUESTION 327

- (Topic 16)

Which statement about the enabling of table compression in Oracle Database 11g is true?

- A. Compression can be enabled at the table, tablespace, or partition level for direct loads only.
- B. Compression can be enabled only at the table level for both direct loads and conventional DML.
- C. Compression can be enabled at the table, tablespace, or partition level for conventional DML only.
- D. Compression can be enabled at the table, tablespace, or partition level for both direct loads and conventional DML.

Answer: D

NEW QUESTION 328

- (Topic 16)

You have three temporary tablespace groups named G1, G2, and G3 in your database. You are creating a new temporary tablespace as follows:

```
CREATE TEMPORARY TABLESPACE TEMP1 TEMPFILE 'u1/data/temp1.dbf' SIZE 10M TABLESPACE GROUP ";
```

Which statement regarding the above command is correct?

- A. It will create the tablespace TEMP1 in group G1.
- B. It will create the tablespace TEMP1 in group G3.
- C. It will not add the tablespace TEMP1 to any group.
- D. It will create the tablespace TEMP1 in the default group.

Answer: C

NEW QUESTION 332

- (Topic 16)

Sales details are being stored on a daily basis in the SALES_2007 table. A large amount of data is added to the table daily. To save disk space, you issued the following command:

```
ALTER TABLE sales_2007 COMPRESS FOR ALL OPERATIONS;
```

What would be the outcome of this command?

- A. It produces an error because data already exists in the table.
- B. It produces an error because compression can be enabled at table creation only.
- C. It compresses all data added or modified henceforth but the existing data in the table is not compressed immediately.
- D. It immediately compresses all existing data as well as new data, resulting from either fresh additions or modifications to existing data.

Answer: C

NEW QUESTION 334

- (Topic 16)

When shrinking a table segment, you choose to shrink all the indexes for that table using the SHRINK SPACE command. Which clause should you use?

- A. INCLUDING DEPENDENCIES
- B. INCLUDING DEPENDENCIES CASCADE
- C. COMPACT
- D. CASCADE
- E. None of the above

Answer: D

NEW QUESTION 335

- (Topic 16)

Which tasks can be accomplished using the DBMS_LOB.SETOPTIONS procedure?

- A. only encryption and compression settings for all SecureFile LOBs
- B. only encryption and deduplication settings for only SecureFile CLOBs
- C. deduplication, encryption, and compression settings for all SecureFile LOBs
- D. deduplication, encryption, and compression settings only for SecureFile CLOBs

Answer: C

Explanation:

Refer to here.

```
DBMS_LOB.SETOPTIONS()
```

This procedure sets compression, deduplication and encryption features. It enables the features to be set on a per-LOB basis, overriding the default LOB settings. This call incurs a round trip to the server to make the changes persistent.

NEW QUESTION 340

- (Topic 16)

View the Exhibit for some of the current parameter settings. A user logs in to the HR schema and issues the following commands:

```
SQL> CREATE TABLE emp (empno NUMBER(3), ename VARCHAR2(20), sal NUMBER(8,2));
```

```
SQL> INSERT INTO emp(empno,ename) VALUES(1,'JAMES');
```

At this moment, a second user also logs in to the HR schema and issues the following command:

```
SQL> ALTER TABLE emp MODIFY sal NUMBER(10,2);
```

What happens in the above scenario? Exhibit:

NAME	TYPE	VALUE
db_file_multiblock_read_count	integer	107
ddl_lock_timeout	integer	60
distributed_lock_timeout	integer	60
dml_locks	integer	748
lock_sga	boolean	FALSE
enable_ddl_logging	boolean	FALSE
resumable_timeout	integer	0

- A. The second user's session immediately produces the resource busy error.
- B. The second user's command executes successfully.
- C. The second user's session waits for a time period before producing the resource busy error.
- D. A deadlock is created.

Answer: C

NEW QUESTION 342

- (Topic 16)

In which cases is reference partitioning effective in enhancing performance?

- A. It is effective only in partition pruning.
- B. It is effective only in partitionwise joins provided that the query predicates are different from the partitioning key.
- C. It is effective in both partition pruning as well as partitionwise joins provided that the query predicates are identical to the partitioning key.
- D. It is effective in both partition pruning as well as partitionwise joins irrespective of whether the query predicates are different from or identical to the partitioning key.

Answer: D

NEW QUESTION 347

- (Topic 16)

You notice that a long-running transaction is suspended due to a space constraint, and there is no AFTER SUSPEND triggered event addressing the issue. You also note that the critical transaction is just about to reach the RESUMABLE_TIMEOUT value.

Which of these actions is appropriate?

- A. Abort the session, fix the space problem, then resubmit the transaction.
- B. Use the DBMS_RESUMABLE.SET_SESSION_TIMEOUT procedure to extend the time-out for the session while you fix the problem.
- C. Do nothing, let the transaction fail, then fix the problem.
- D. Use Segment Shrink to clean up the table.
- E. Use the DBMS_RESUMABLE.SET_TIMEOUT procedure to extend the time-out for the session while you fix the problem.

Answer: B

NEW QUESTION 350

- (Topic 17)

View the Exhibit to examine the parameters set for your database instance.

You execute the following command to perform I/O calibration after the declaration of bind variables in the session that are used in the command:

```
SQL> EXECUTE dbms_resource_manager.calibrate_io( num_physical_disks=>1,
max_latency=>50, max_iops=>:max_iops, max_mbps=>:max_mbps, actual_latency=>:actual_latency);
```

Which statement describes the consequence? Exhibit:

NAME	TYPE	VALUE
filesystemio_options	string	ASYNCH
backup_tape_io_slaves	boolean	FALSE
dbwr_io_slaves	integer	0
disk_asynch_io	boolean	TRUE
tape_asynch_io	boolean	TRUE
optimizer_use_pending_statistics	boolean	FALSE
statistics_level	string	TYPICAL
timed_os_statistics	integer	0
timed_statistics	boolean	FALSE
aq_tm_processes	integer	0
db_writer_processes	integer	1
gcs_server_processes	integer	0
global_txn_processes	integer	1
job_queue_processes	integer	1000
log_archive_max_processes	integer	4
processes	integer	150

- A. The command produces an error.
- B. The calibration process runs successfully and populates all the bind variables.
- C. The calibration process runs successfully but the latency time is not computed.
- D. The calibration process runs successfully but only the latency time is computed.

Answer: A

Explanation:

Requisition of Calibrate I/O (link)

Before running I/O calibration, ensure that the following requirements are met:

- ? The user must be granted the SYSDBA privilege
- ? timed_statistics must be set to TRUE
- ? Asynchronous I/O must be enabled
- ? When using file systems, asynchronous I/O can be enabled by setting the FILESYSTEMIO_OPTIONS initialization parameter to SETALL.
- ? Ensure that asynchronous I/O is enabled for data files by running the following query:

```
COL NAME FORMAT A50
SELECT NAME,ASYNCH_IO FROM V$DATAFILE F,V$IOSTAT_FILE I WHERE F.FILE#=I.FILE_NO
AND FILETYPE_NAME='Data File';
```

Additionally, only one calibration can be performed on a database instance at a time.

NEW QUESTION 354

- (Topic 17)

Which three statements correctly describe the features of the I/O calibration process? (Choose three.)

- A. Only one I/O calibration process can run at a time.
- B. It automates the resource allocation for the Automated Maintenance Tasks.
- C. It improves the performance of the performance-critical sessions while running.
- D. It can be used to estimate the maximum number of I/Os and maximum latency time for the system.

E. The latency time is computed only when the TIMED_STATISTICS initialization parameter is set to TRUE.

Answer: ADE

NEW QUESTION 355

- (Topic 17)

Which of the following describes how a distributed resumable transaction behaves?

- A. The resumable setting on the initiating session determines the resumable conditions for the entire distributed transaction.
- B. The resumable setting for the initiating instance determines the resumable conditions for the entire distributed transaction.
- C. The resumable setting on the initiating session controls only that part of the transaction that occurs within the local instance; remote resumable settings determine the behavior of the distributed parts of the transaction.
- D. None of the above.

Answer: C

NEW QUESTION 356

- (Topic 17)

The database users regularly complain about the difficulty in performing transactions. On investigation, you find that some users perform long-running transactions that consume huge amounts of space in the undo tablespace, which caused the problem. You want to control the usage of the undo tablespace only for these user sessions and you do not want these sessions to perform long-running operations.

How would you achieve this?

- A. Implement a profile for the users.
- B. Implement external roles for the users.
- C. Set the threshold for the undo tablespace.
- D. Implement a Database Resource Manager plan.

Answer: B

NEW QUESTION 361

- (Topic 17)

For which two situations would you use functionality provided by the Resource Manager? (Choose two.)

- A. setting idle timeout limits on resource plans
- B. saving storage space by using compressed backup sets
- C. creating jobs that will run automatically at a scheduled time
- D. assigning priorities to jobs to manage access to system resources
- E. creating alerts to perform notification when tablespaces are low on available space resources

Answer: AD

NEW QUESTION 364

- (Topic 17)

Within a resource-plan definition, what differentiates a top-level plan from a subplan?

- A. A subplan has the PLAN_SUB parameter value set to SUB.
- B. A top-level plan has the GROUP_OR_PLAN parameter set to the name of the subplan in the resource-plan definition.
- C. There is no difference in the resource-plan definition.
- D. A subplan always has the CPU_MTH parameter value set to RATIO.
- E. The string TOP_LEVEL is appended to the name of top-level resource plans.

Answer: C

NEW QUESTION 365

- (Topic 17)

Which statements describe the capabilities of the DBMS_NETWORK_ACL_ADMIN package? (Choose all that apply.)

- A. It can be used to allow the access privilege settings for users but not roles.
- B. It can be used to allow the access privilege settings for users as well as roles.
- C. It can be used to control the time interval for which the access privilege is available to a user.
- D. It can be used to selectively restrict the access for each user in a database to different host computers.
- E. It can be used to selectively restrict a user's access to different applications in a specific host computer.

Answer: BCD

NEW QUESTION 369

- (Topic 17)

View the exhibit and examine the output. Which statement can be an inference from the output? Exhibit:

Oracle Enterprise Manager 10g Grid Control implementation Guide Cause of I/O Bottlenecks

Request	Web Cache		Web Server		Application Server		Database Server		Server Time Details
	Count	IR Ratio	Count	Average Time (ms)	Count	Average Time (ms)	Count	Average Time (ms)	
http://www.oracle.com/.../.../.../...	1	n/a	2	108,874.13	2	327.05	1	4,443.00	
http://www.oracle.com/.../.../.../...	11	n/a	15	2,549.00	1	5,874.00	1	1,001.00	
http://www.oracle.com/.../.../.../...	34	n/a	38	4,429.17	31	819.71	2	2,213.39	
http://www.oracle.com/.../.../.../...	1	n/a	2	233.45	6	29.16	2	2,693.11	
http://www.oracle.com/.../.../.../...	21	n/a	24	4,554.70	24	2,048.11	1	3,045.00	
http://www.oracle.com/.../.../.../...	51	n/a	54	2,056.62	22	441.00	1	1,763.00	
http://www.oracle.com/.../.../.../...	1	n/a	1	1,036.63	1	171.00	1	1,257.00	
http://www.oracle.com/.../.../.../...	61	n/a	64	2,001.12	40	438.45	1	1,721.00	
http://www.oracle.com/.../.../.../...	4	n/a	5	2,776.21	2	1,328.50	1	1,602.00	
http://www.oracle.com/.../.../.../...	1	n/a	1	2,090.10	1	125.00	1	1,453.00	

- A. The FRA disk group has an asynchronous I/O bottleneck
- B. The least number of I/Os are performed on the last data file in the list
- C. The number of times that the backup or restore process directed the OS to wait until an I/O was complete is the highest for the last data file in the list
- D. The number of times that the backup or restore process made an OS call to poll for I/O completion in Nonblocking mode is the least for the FRA disk group

Answer: A

NEW QUESTION 374

- (Topic 18)

You notice that a job in a chain has not completed on a nonconstrained RAC database. Which of these are valid reasons why that might occur?

- A. The job priority is 1 and the resource consumer group CPU emphasis allocation is a low percentage.
- B. The job affinity is to a service and one node in that service is unavailable.
- C. The job affinity is to an instance and that instance is unavailable.
- D. There is no service affinity.
- E. None of the above.

Answer: C

NEW QUESTION 378

- (Topic 18)

Which of the following is not a valid setting for the PROGRAM_TYPE parameter in a program object or the JOB_TYPE parameter in a job object?

- A. PLSQL_BLOCK
- B. JAVA_STORED_PROCEDURE
- C. STORED_PROCEDURE
- D. EXECUTABLE
- E. None of the above are invalid settings.

Answer: B

Explanation:

job_type
Job action type ('PLSQL_BLOCK', 'STORED_PROCEDURE', 'EXECUTABLE', or 'CHAIN')

NEW QUESTION 380

- (Topic 18)

Which of the following tasks is not performed by the job coordinator?

- A. Update job log when a job completes
- B. Spawn and remove job slaves
- C. Write/read job info to/from memory cache
- D. Query job table
- E. Pass job information to job slaves

Answer: A

NEW QUESTION 383

- (Topic 18)

The user SYS creates a job by using the following command:

```
BEGIN DBMS_SCHEDULER.CREATE_JOB (
job_name => 'update_sales',
job_type => 'STORED-PROCEDURE',
```

```
job_action => 'OPS.SALES_PKG.UPDATE_SALES_SUMMARY',  
start_date => '28-DEC-07 07.00.00 PM Australia/Sydney', repeat_interval => 'FREQ=DAILY; INTERVAL=2', end_date => '20-JAN-08 07.00.00 PM  
Australia/Sydney', comments => 'New sales job');  
END;  
/
```

Which two statements are true about the job that was created by the preceding command?
(Choose two.)

- A. The job is enabled by default after creation
- B. The job is automatically dropped after the end date
- C. The job executes with the privileges of the user SYS
- D. The globalization environment that exists at the time of the job creation prevails at the job runs

Answer: BC

Explanation:

By default, jobs are created with auto_drop set to TRUE. end_date

This attribute specifies the date and time after which the job expires and is no longer run. After the end_date, if auto_drop is TRUE, the job is dropped. If auto_drop is FALSE, the job is disabled and the STATE of the job is set to COMPLETED.

If no value for end_date is specified, the job repeats forever unless max_runs or max_failures is set, in which case the job stops when either value is reached.

The value for end_date must be after the value for start_date. If it is not, an error is generated when the job is enabled.

NEW QUESTION 386

- (Topic 19)

If two windows overlap, which window attribute will determine whether one should be chosen over the other?

- A. WINDOW_PRIORITY
- B. PRIORITY
- C. PRIORITY_LEVEL
- D. WINDOW_PRIORITY_LEVEL
- E. OVERLAP_RULE

Answer: A

NEW QUESTION 389

- (Topic 19)

Which three statements are true regarding persistent lightweight jobs? (Choose three.)

- A. Persistent lightweight jobs modify several tables in the data dictionary to generate a lot of redo.
- B. The user cannot set privileges on persistent lightweight jobs.
- C. Persistent lightweight jobs are useful when users need to create a large number of jobs in a short time.
- D. Persistent lightweight jobs are useful when users need to create a small number of jobs that run infrequently.
- E. The use of a template is mandatory to create persistent lightweight jobs.

Answer: BCE

NEW QUESTION 390

- (Topic 19)

Which of the following is true about job chains?

- A. They consist of one or more Scheduler programs.
- B. They are used to implement dependency scheduling.
- C. They are used to implement time-based scheduling.
- D. They are used to implement event-based scheduling.
- E. None of the above.

Answer: B

Explanation:

Creating and Managing Job Chains

A job chain ("chain") is a named series of tasks that are linked together for a combined objective. Chains are the means by which you can implement dependency based scheduling, in which jobs are started depending on the outcomes of one or more previous jobs.

NEW QUESTION 391

- (Topic 20)

Which NLS parameter can be used to change the default Oracle sort method from binary to linguistic for the SQL SELECT statement?

- A. NLS_LANG
- B. NLS_COMP
- C. NLS_SORT
- D. None of the above

Answer: D

Explanation:

Neither NLS_SORT nor NLS_COMP cannot change sorting from BINARY to LINGUISTIC. Because the NLS_SORT depends on NLS_COMP, if they are different, the sorting method will always be BINARY.

(Refer to NLS_SORT) The exact operators and query clauses that obey the NLS_SORT parameter depend on the value of the NLS_COMP parameter. If an operator or clause does not obey the NLS_SORT value, as determined by NLS_COMP, the collation used is BINARY.

The BINARY comparison is faster and uses less resources than any linguistic comparison but for text in a natural language, it does not provide ordering expected by users.

The value of NLS_SORT affects execution plans of queries. Because a standard index cannot be used as a source of values sorted in a linguistic order, an explicit sort operation must usually be performed instead of an index range scan. A functional index on the NLSSORT function may be defined to provide values sorted in a linguistic order and reintroduce the index range scan to the execution plan.

NLS_COMP specifies the collation behavior of the database session.

NEW QUESTION 393

- (Topic 20)

Which of the following are valid settings for the NLS_COMP parameter? (Choose all that apply.)

- A. ASCII
- B. ANSI
- C. BINARY
- D. MONOLINGUAL
- E. MULTILINGUAL

Answer: BC

Explanation:

Parameter type	String
Syntax	NLS_COMP = { BINARY LINGUISTIC ANSI }
Default value	BINARY
Modifiable	ALTER SESSION
Basic	No

C:\Users\albo\Desktop\1-1.jpg

NLS_COMP specifies the collation behavior of the database session. Values:

BINARY

Normally, comparisons in the WHERE clause and in PL/SQL blocks is binary unless you specify the NLSSORT function.

LINGUISTIC

Comparisons for all SQL operations in the WHERE clause and in PL/SQL blocks should use the linguistic sort specified in the NLS_SORT parameter. To improve the performance, you can also define a linguistic index on the column for which you want linguistic comparisons.

ANSI

A setting of ANSI is for backwards compatibility; in general, you should set NLS_COMP to LINGUISTIC

NEW QUESTION 398

- (Topic 20)

In a database with the database character set of US7ASCII and a national character set of UTF-8, which datatypes would be capable of storing Unicode data by default?

- A. VARCHAR2
- B. CHAR
- C. NVARCHAR2
- D. CLOB
- E. LONG

Answer: C

NEW QUESTION 401

- (Topic 20)

Globalization support is implemented through the text- and character-processing functions provided by which Oracle feature?

- A. RSTLNE
- B. NLSRTL
- C. LISTENER
- D. NLSSORT
- E. Linguistic sorts

Answer: B

Explanation:

NLSRTL: NLS Runtime Type Library.

NEW QUESTION 406

- (Topic 20)

Given two different character sets (A and B), which of the following must be true for A to be considered a strict superset of B? (Choose all that apply.)

- A. A must contain all of the characters defined in B.
- B. A must be Unicode.
- C. The encoded values in A must match the encoded values in B for all characters defined in B.
- D. A must be a multi-byte character set.
- E. The encoded values in A must match the encoded values in B for all numeric and alphabetic characters in B.

Answer: AC

NEW QUESTION 411

- (Topic 20)

What elements of globalization can be explicitly defined using the NLS_LANG environment variable? (Choose all that apply.)

- A. NLS_LANGUAGE
- B. NLS_SORT
- C. NLS_CALENDAR
- D. NLS_CHARACTERSET
- E. NLS_TERRITORY

Answer: ADE

NEW QUESTION 412

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your 1Z0-053 Exam with Our Prep Materials Via below:

<https://www.certleader.com/1Z0-053-dumps.html>