

Exam Questions 1Z0-053

Oracle Database 11g: Administration II

<https://www.2passeasy.com/dumps/1Z0-053/>



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 are managing an Oracle 11g database with ASM storage, for which the `COMPATIBLE` initialization parameter is set to 11.1.0. In the ASM instance, the `COMPATIBLE.RDBMS` attribute for the disk group is set to 10.2 and the `COMPATIBLE.ASM` attribute is set to 11.1.

Which two statements are true in this scenario for the features enabled for ASM? (Choose two.)

- A. The ASM-preferred mirror read feature is enabled.
- B. The ASM supports variable sizes for extents of 1, 8, and 64 allocation units.
- C. The ASM disk is dropped immediately from a disk group when it becomes unavailable.
- D. The RDBMS always reads the primary copy of a mirrored extent of the ASM disk group.

Answer: AB

NEW QUESTION 3

- (Topic 1)

When starting up your ASM instance, you receive the following error: `SQL> startup pfile=$ORACLE_HOME/dbs/init+ASM.ora`

ASM instance started

Total System Global Area 104611840 bytes Fixed Size 1298220 bytes

Variable Size 78147796 bytes ASM Cache 25165824 bytes

ORA-15032: not all alternations performed

ORA-15063: ASM discovered an insufficient number of disks for diskgroup "DGROU3" ORA-15063: ASM discovered an insufficient number of disks for

diskgroup "DGROU2" ORA-15063: ASM discovered an insufficient number of disks for diskgroup "DGROU1"

In trying to determine the cause of the problem, you issue this query: `SQL> show parameter asm`

NAME	TYPE	VALUE
asm_allow_only_raw_disks	boolean	FALSE
asm_diskgroups	string	DGROU1, DGROU2, DGROU3
asm_diskstring	string	
asm_power_limit	integer	1
asm_preferred_read_failure_groups	string	

What is the cause of the error?

- A. The `ASM_DISKGROUPS` parameter is configured for three disk groups: `DGROU1`, `DGROU2`, and `DGROU3`. The underlying disks for these disk groups have apparently been lost.
- B. The format of the `ASM_DISKGROUPS` parameter is incorrect.
- C. It should reference the disk group numbers, not the names of the disk groups.
- D. The `ASM_POWER_LIMIT` parameter is incorrectly set to 1. It should be set to the number of disk groups being attached to the ASM instance.
- E. The `ASM_DISKSTRING` parameter is not set; therefore disk discovery is not possible.
- F. There is insufficient information to solve this problem.

Answer: D

Explanation:

`ASM_DISKSTRING` specifies an operating system-dependent value used by Automatic Storage Management to limit the set of disks considered for discovery. When a new disk is added to a disk group, each Automatic Storage Management instance that has the disk group mounted must be able to discover the new disk using the value of `ASM_DISKSTRING`.

In most cases, the default value will be sufficient. Using a more restrictive value may reduce the time required for Automatic Storage Management to perform discovery, and thus improve disk group mount time or the time for adding a disk to a disk group. A "?" at the beginning of the string gets expanded to the Oracle home directory. Depending on the operating system, wildcard characters can be used. It may be necessary to dynamically change `ASM_DISKSTRING` before adding a disk so that the new disk will be discovered. An attempt to dynamically modify `ASM_DISKSTRING` will be rejected and the old value retained if the new value cannot be used to discover a disk that is in a disk group that is already mounted.

Refer to here

NEW QUESTION 4

- (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 5

- (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 6

- (Topic 1)

As DBA for the Rebalance, you have decided that you need to facilitate some redundancy in your database. Using ASM, you want to create a disk group that will provide for the greatest amount of redundancy for your ASM data (you do not have advanced SAN mirroring technology available to you, unfortunately).

Which of the following commands would create a disk group that would offer the maximum in data redundancy?

- A.

```
CREATE DISKGROUP dg_alliance1 NORMAL REDUNDANCY
  FAILGROUP diskcontrol1 DISK 'c:\oracle\asm_disk\file_disk3' NAME file_disk1
  FAILGROUP diskcontrol2 DISK 'c:\oracle\asm_disk\file_disk4' NAME
  file_disk1;
```
- B.

```
CREATE DISKGROUP dg_alliance1 EXTERNAL REDUNDANCY
  FAILGROUP diskcontrol1 DISK 'c:\oracle\asm_disk\file_disk3' NAME
  file_disk1;
```
- C.

```
CREATE DISKGROUP dg_alliance1 HIGH REDUNDANCY
  FAILGROUP diskcontrol1 DISK 'c:\oracle\asm_disk\file_disk1' NAME file_disk1
  FAILGROUP diskcontrol2 DISK 'c:\oracle\asm_disk\file_disk2' NAME file_disk2
  FAILGROUP diskcontrol3 DISK 'c:\oracle\asm_disk\file_disk3' NAME file_disk3;
```
- D.

```
CREATE DISKGROUP dg_alliance1 MAXIMUM REDUNDANCY
  FAILGROUP diskcontrol1 DISK 'c:\oracle\asm_disk\file_disk1' NAME file_disk1
  FAILGROUP diskcontrol2 DISK 'c:\oracle\asm_disk\file_disk2' NAME file_disk2
  FAILGROUP diskcontrol2 DISK 'c:\oracle\asm_disk\file_disk3' NAME file_disk3
  FAILGROUP diskcontrol2 DISK 'c:\oracle\asm_disk\file_disk4' NAME file_disk4;
```
- E. None of the above

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

Answer: C

Explanation:

No SAN mirroring available means no external redundancy available.

The highest redundancy of ASM is the HIGH redundancy with 3 mirror copies.

NEW QUESTION 7

- (Topic 1)

You are managing an Oracle Database 11g ASM instance having three disks in a disk group with ASM compatibility attribute set to 11.1.0 and redundancy set to high. One of the disks in the disk group becomes unavailable because of power failure.

Which statements will be true in this scenario? (Choose all that apply.)

- A. The disk is immediately dropped from the disk group.
- B. The ASM tracks the extents that are modified during the outage.
- C. The ASM migrates the extents from the unavailable disk to the remaining disks.
- D. The disk automatically goes offline.

Answer: BD

NEW QUESTION 8

- (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 9

- (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 10

- (Topic 1)

How can you reverse the effects of an ALTER DISKGROUP ... DROP DISK command if it has already completed?

- A. Issue the ALTER DISKGROUP ... ADD DISK command.
- B. Issue the ALTER DISKGROUP ... UNDROP DISKS command.
- C. Issue the ALTER DISKGROUP ... DROP DISK CANCEL command.
- D. Retrieve the disk from the Recycle Bin after the operation completes.

Answer: A

Explanation:

You cannot UNDROP DISKS if the DROP DISK command has already completed.

NEW QUESTION 10

- (Topic 1)

A database instance is using an Automatic Storage Management (ASM) instance, which has a disk group, DGROUP1, created as follows:

```
SQL> CREATE DISKGROUP dgroup1 NORMAL REDUNDANCY  
FAILGROUP controller1 DISK '/devices/diska1', '/devices/diska2' FAILGROUP controller2 DISK '/devices/diskb1', '/devices/diskb2';  
What happens when the whole CONTROLLER1 Failure group is damaged?
```

- A. The transactions that use the disk group will halt.
- B. The mirroring of allocation units occurs within the CONTROLLER2 failure group.
- C. The data in the CONTROLLER1 failure group is shifted to the CONTROLLER2 failure group and implicit rebalancing is triggered.
- D. The ASM does not mirror any data and newly allocated primary allocation units (AU) are stored in the CONTROLLER2 failure group.

Answer: C

NEW QUESTION 13

- (Topic 1)

You are an Oracle DBA responsible for an ASM instance. The disk controller on your system fails. You suspect that the disk itself is okay. You know it will take 24 hours to replace the controller and you don't want to have to rebuild the disks from scratch. What do you do?

- A. Take the whole disk group offline and wait for the controller card to be installed
- B. Once it's installed, bring the disk group online again.
- C. Change the ASM parameter ASM_PREFERRED_READ_FAILURE_GROUPS to indicate that you want to read from the non-failed disks
- D. Once the disk controller is replaced, reset the parameter to its original value.
- E. You have no choice but to rebuild the disks
- F. Drop the disk from the disk group and wait for the controller to be replaced
- G. Once the controller is replaced, add the disk back into the disk group and allow ASM to rebuild it.
- H. If you are using any setting other than REDUNDANCY EXTERNAL for your disk group, you will have to recover any data on that disk from a backup
- I. The database will be unavailable until you can correct the problem and perform recovery.
- J. Change the attribute DISK_REPAIR_TIME on the disk group to a time greater than 24 hours.

Answer: E

NEW QUESTION 16

- (Topic 1)

View Exhibit1 to examine the CRS11G disk group when all the disks are online. View Exhibit2 to examine the CRS11G disk group when one disk is offline.

Why is the rebalancing not performed and the content of the disk group not empty in Exhibit2?

ASM-disk-group-1 (exhibit):

General
 Name: CRS11G
 State: MOUNTED
 Redundancy: NORMAL
 Total Size (GB): 4
 Pending Operations: 0
 Allocation Unit (MB): 1

Advanced Attributes
 Database Compatibility: 11.1.0.0.0
 ASM Compatibility: 11.2.0.0.0
 ASM Volume Compatibility:
 Disk Repair Time (Hours): 3.6

Member Disks

Select Disk	Failure Group	Path	Library	Read/Write Errors	State	Mode	Size (GB)	Used (GB)	Us
<input type="checkbox"/> CRS11G_0000	CRS11G_0000	/dev/oracle/vot_sh1_1122	SYSTEM	0	NORMAL	✓	1.00	0.24	
<input type="checkbox"/> CRS11G_0001	CRS11G_0001	/dev/oracle/vot_ve1_1123	SYSTEM	0	NORMAL	✓	1.00	0.24	
<input type="checkbox"/> CRS11G_0002	CRS11G_0002	/dev/oracle/vot_ve1_1191	SYSTEM	0	NORMAL	✓	1.00	0.24	
<input type="checkbox"/> CRS11G_0003	CRS11G_0003	/dev/oracle/vot_sh2_0287	SYSTEM	0	NORMAL	✓	1.00	0.20	

ASM-disk-group-2 (exhibit):

General
 Name: CRS11G
 State: MOUNTED
 Redundancy: NORMAL
 Total Size (GB): 4
 Pending Operations: 0
 Allocation Unit (MB): 1

Advanced Attributes
 Database Compatibility: 11.1.0.0.0
 ASM Compatibility: 11.2.0.0.0
 ASM Volume Compatibility:
 Disk Repair Time (Hours): 3.6

Member Disks

Select Disk	Failure Group	Path	Library	Read/Write Errors	State	Mode	Size (GB)	Used (GB)	Us
<input type="checkbox"/> CRS11G_0000	CRS11G_0000	/dev/oracle/vot_sh1_1122	SYSTEM	0	NORMAL	✓	1.00	0.24	
<input type="checkbox"/> CRS11G_0001	CRS11G_0001		SYSTEM	0	NORMAL	✗	1.00	0.24	
<input type="checkbox"/> CRS11G_0002	CRS11G_0002	/dev/oracle/vot_ve1_1191	SYSTEM	0	NORMAL	✓	1.00	0.24	
<input type="checkbox"/> CRS11G_0003	CRS11G_0003	/dev/oracle/vot_sh2_0287	SYSTEM	0	NORMAL	✓	1.00	0.20	

- A. Because the disk group is created with NORMAL redundancy
- B. Because the disk repair time attribute is set to a nonzero value
- C. Because the mirrored extents cannot be rebalanced across the other three disks
- D. Because the other three disks have 60% free space, the disk rebalancing is delayed

Answer: B

Explanation:

Key points:

1. The COMPATIBILITY.ASM >= 11.1 and COMPATIBILITY.RDBMS >= 11.1
2. The DISK_REPAIR_TIME is set to non-zero. by default it is 3.6hr. Refer to here.

Note: To use this feature, the disk group compatibility attributes must be set to 11.1 or higher. For more information, refer to "Disk Group Compatibility".

Oracle ASM fast resync keeps track of pending changes to extents on an OFFLINE disk during an outage. The extents are resynced when the disk is brought back online.

By default, Oracle ASM drops a disk in 3.6 hours after it is taken offline. You can set the DISK_REPAIR_TIME disk group attribute to delay the drop operation by specifying a time interval to repair the disk and bring it back online.

NEW QUESTION 19

- (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 20

- (Topic 1)

ASM supports all but which of the following file types? (Choose all that apply.)

- A. Database files
- B. SPFILEs
- C. Redo-log files
- D. Archived log files
- E. RMAN backup sets
- F. Password files
- G. init.ora files

Answer: FG

Explanation:

What Types of Files Does Oracle ASM Support?

Table 7-1 File Types Supported by Automatic Storage Management

File Type	Default Templates
Control files	CONTROLFILE
Data files	DATAFILE
Redo log files	ONLINELOG
Archive log files	ARCHIVELOG
Temporary files	TEMPFILE
Data file backup pieces	BACKUPSET
Data file incremental backup pieces	BACKUPSET
Archive log backup piece	BACKUPSET
Data file copy	DATAFILE
Persistent initialization parameter file (SPFILE)	PARAMETERFILE
Flashback logs	FLASHBACK
Change tracking file	CHANGETRACKING
Data Pump dumpset	DUMPSET
Automatically generated control file backup	AUTOBACKUP
Cross-platform transportable data files	XTRANSPORT
Flash file	FLASHFILE
Oracle ASM Persistent initialization parameter file (SPFILE)	ASMPARAMETERFILE
Oracle ASM Persistent initialization parameter file (SPFILE) backup	ASMPARAMETERFILEBACKUP
Oracle Cluster Registry file	OCRFILE
Oracle ASM Dynamic Volume Manager volumes	n/a

NEW QUESTION 23

- (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 28

- (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 33

- (Topic 1)

Which initialization parameter in an ASM instance specifies the disk groups to be automatically mounted at instance startup?

- A. ASM_DISKMOUNT
- B. ASM_DISKGROUP
- C. ASM_DISKSTRING
- D. ASM_MOUNTGROUP

Answer: B

Explanation:

Refer to here

When you run the STARTUP command, this command attempts to mount the disk groups specified by the initialization parameter ASM_DISKGROUPS. If you have not entered a value for ASM_DISKGROUPS, then the ASM instance starts and Oracle displays an error that no disk groups were mounted. You can then mount disk groups with the ALTER DISKGROUP...MOUNT command.

NEW QUESTION 36

- (Topic 2)

Which RMAN backup command is used to create the block-change tracking file?

- A. alter database create block change tracking file
- B. alter database enable block change file
- C. alter database enable block change tracking using file '/ora01/opt/block_change_tracking.fil'
- D. alter system enable block change tracking using file '/ora01/opt/block_change_tracking.fil'
- E. alter system block change tracking on

Answer: C

NEW QUESTION 40

- (Topic 2)

What does it mean if a backup is expired?

- A. The backup set has exceeded the retention criteria set in RMAN and is eligible for removal.
- B. The backup set has one or more invalid blocks in it and is not usable for recovery.
- C. The backup set contains one or more tablespaces no longer in the database.
- D. The backup set contains one or more missing backup set pieces.
- E. The backup set is from a previous version of RMAN and was not upgraded.

Answer: D

NEW QUESTION 41

- (Topic 2)

As soon as you discover that you have lost an online redo log, if the database is still functioning, what should be your first action?

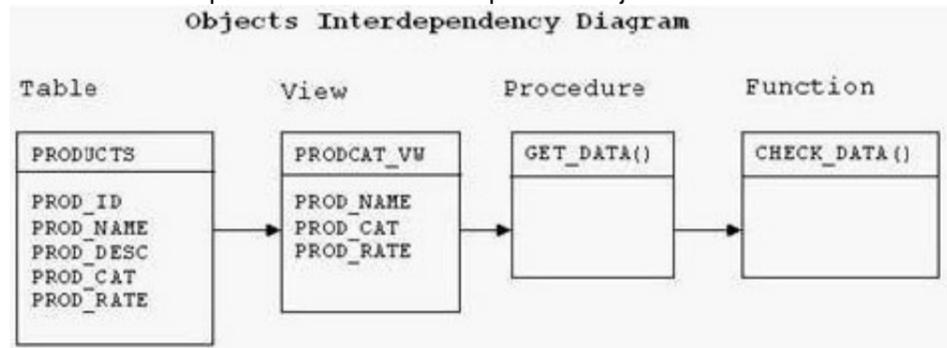
- A. Shut down the database
- B. Clear the online redo log
- C. Back up the database
- D. Checkpoint the database
- E. Call Oracle support

Answer: D

NEW QUESTION 46

- (Topic 2)

View the Exhibit for the object interdependency diagram.
 ? The PRODUCTS table is used to create the PRODCAT_VW view.
 ? PRODCAT_VW is used in the GET_DATA procedure.
 ? GET_DATA is called in the CHECK_DATA function.
 ? A new column PROD_QTY is added to the PRODUCTS table.
 How does this impact the status of the dependent objects? Exhibit:



- A. All dependent objects remain valid.
- B. Only the procedure and function become invalid and must be recompiled.
- C. Only the view becomes invalid and gets automatically revalidated the next time it is used.

D. Only the procedure and function become invalid and get automatically revalidated the next time they are called.

Answer: A

NEW QUESTION 51

- (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 55

- (Topic 2)

You want to make sure that your database backup does not exceed 10 hours in length. What command would you issue that would meet this condition?

- A. backup database plus archivelog;
- B. backup database plus archivelog until time 10:00;
- C. backup database plus archivelog timeout 10:00;
- D. backup database plus archivelog duration 10:00;
- E. backup database plus archivelog timeout 10:00;

Answer: D

NEW QUESTION 56

- (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 58

- (Topic 2)

Which options would you consider while configuring a flash recovery area (fast recovery area in 11g Release 2) for your production database that is running in ARCHIVELOG mode? (Choose all that apply.)

- A. Setting the FAST_START_MTTR_TARGET to set the mean time to recover
- B. Setting the RECOVERY_PARALLELISM parameter to twice the number of CPUs
- C. Using the DB_RECOVERY_FILE_DEST parameter to set the location for flash recovery area
- D. Using the DB_RECOVERY_FILE_DEST_SIZE parameter to define the disk space limit for the recovery files created in the flash recovery area

Answer: CD

NEW QUESTION 60

- (Topic 2)

Which of the following does the recover command not do?

- A. Restore archived redo logs.
- B. Apply archived redo logs.
- C. Restore incremental backups.
- D. Apply incremental backups.
- E. Restore datafile images.

Answer: A

NEW QUESTION 64

- (Topic 2)

You run the following commands:

```
RMAN> list expired backup; RMAN> delete expired backup;
```

What will happen to the backup set pieces associated with the backups that appear in the list expired backup command?

- A. They will be renamed.
- B. Nothing will happen to the
- C. The backup set pieces do not exist.

- D. They will be deleted immediately since they are not in the flash recovery area.
- E. You will need to manually remove the physical files listed in the output of the commands.
- F. They will become hidden files and removed 10 days later.

Answer: B

NEW QUESTION 67

- (Topic 2)

What will be the result of the following configuration?

```
Log_archive_dest_1 = 'location=c:\oracle\arch\mydb' Log_archive_dest_2 = 'location=z:\oracle\arch\mydb'
```

- A. An error will occur during database startup because the second parameter is not valid.
- B. An error will occur during database startup since you are trying to create archived redo logs in two different locations.
- C. Archived redo logs will be created in two different locations by the ARCH process.
- D. Archived redo logs will be created in two different locations by the LGWR process.
- E. Neither parameter setting is valid, so the database will not start up.

Answer: C

NEW QUESTION 71

- (Topic 2)

Which is the correct command to put the database in ARCHIVELOG mode?

- A. alter database archivelog
- B. alter system enable archivelog mode
- C. alter database enable archive
- D. alter database archivelog enable
- E. None of the above

Answer: A

NEW QUESTION 73

- (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 alterlogging.

Answer: A

Explanation:

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

NEW QUESTION 78

- (Topic 2)

Observe the following warning in an RMAN session of your database instance: WARNING: new failures were found since last LIST FAILURE command
Which statement describes the scenario that must have produced this warning?

- A. The CHANGE FAILURE command has detected new failures recorded in the Automatic Diagnostic Repository(ADR).
- B. The VALIDATE DATABASE command has detected new failures recorded in the Automatic Diagnostic Repository (ADR).
- C. The ADVISE FAILURE command has detected new failures recorded in the Automatic Diagnostic Repository (ADR) since the last LIST FAILURE.
- D. The RECOVER command has detected new failures recorded in the Automatic Diagnostic Repository (ADR) since the last LIST FAILURE command was executed.

Answer: C

NEW QUESTION 81

- (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 82

- (Topic 3)

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

- A. Make backups faster
- B. Store RMAN metadata
- C. Store RMAN scripts
- D. Provide the ability to do centralized backup reporting.
- E. Make recovery faster

Answer: BCD

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 87

- (Topic 3)

Given the script

```
create script db_backup_datafile_script
{backup datafile and 1, and2 plus archivelog delete input;}
```

What is the result of running this command?

```
Run {execute script db_backup_datafile_script using 2;}
```

- A. The script will fail since you instructed RMAN to back up only one datafile rather than two.
- B. The script will successfully back up datafile 3 without error.
- C. The script will fail since it uses a substitution variable which is not supported.
- D. The execute script command will prompt for the value of and2 since it's not included in the command.
- E. The script will fail because you cannot use the plus archivelog command when backing up database datafiles.

Answer: D

NEW QUESTION 88

- (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 93

- (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 98

- (Topic 3)

Identify two advantages of using a recovery catalog in-load of the control File of the target database Recovery Manager (RMAN). (Choose two.)

- A. You can use RMAN stored scripts.
- B. Recovery is faster if data is stored in catalog in addition to the control file.
- C. You can store backup Information of all registered databases in one place.
- D. Database backups are automatically deleted when they are older than the specified time period.

Answer: AC

NEW QUESTION 103

- (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 105

- (Topic 4)

You want the ability to recovery any time within the last seven days and therefore you configured the recovery window retention policy using the command:

```
RMAN> CONFIGURE RETENTION POLICY TO RECOVERY WINDOW OF 7 DAYS;
```

After configuring the recovery window, you performed the database backup as follows:

- A. Backup RB1 at log sequence number 12871 on 5th Jan
- B. Backup RB2 at log sequence number 15622 on 12th Jan
- C. Backup RB3 at log sequence 16721 on 15th Jan
- D. On 20th Jan when the log sequence number was 18112 you realize that there is a need to a point in time at the beginning of the recovery window
- E. You have all archived redo log files to date.

Answer: D

NEW QUESTION 106

- (Topic 4)

Which type of backup contains only the blocks that have changed since the last level 0 incremental backup?

- A. a cumulative level 1 backup
- B. a differential level 1 backup
- C. a full backup
- D. a whole backup

Answer: A

NEW QUESTION 108

- (Topic 4)

Which of the following best describes a full backup?

- A. All datafiles of a database
- B. All datafiles, archive logs, and control files
- C. All datafiles and control files
- D. All the used blocks in a datafile

Answer: D

Explanation:

From the training book:

“Full backup: Makes a copy of each data block that contains data and that is within the files being backed up.”

“A full backup contains all used data file blocks.”

“A full backup is different from a whole database backup.”

Conclusion – it means that not all data files are backed up when you do full backup.

NEW QUESTION 109

- (Topic 4)

You want to set the following initialization parameters for your database instance:

LOG_ARCHIVE_DEST_1 = 'LOCATION=/disk1/arch' LOG_ARCHIVE_DEST_2 = 'LOCATION=/disk2/arch' LOG_ARCHIVE_DEST_3 = 'LOCATION=/disk3/arch'

LOG_ARCHIVE_DEST_4 = 'LOCATION=/disk4/arch MANDATORY'

Identify the statement that correctly describes this setting.

- A. The MANDATORY location must be a flash recovery area.
- B. The optional destinations may not use the flash recovery area.
- C. This setting is not allowed because the first destination is not set as MANDATORY.
- D. The online redo log file is not allowed to be overwritten if the archived log cannot be created in the fourth destination.

Answer: D

Explanation:

MANDATORY

Specifies that filled online log files must be successfully archived to the destination before they can be reused.

If MANDATORY is not specified, then, by default, the destination is considered to be optional.

At least one destination must succeed, even if all destinations are optional. If archiving to an optional destination fails, the online redo log file is still available for reuse and may be overwritten eventually. However, if the archival operation of a mandatory destination fails, online redo log files cannot be overwritten.

NEW QUESTION 112

- (Topic 4)

Examine the following RMAN command:

RMAN> CONFIGURE ENCRYPTION FOR DATABASE ON; RMAN> BACKUP DATABASE PLUS ARCHIVELOG;

Which prerequisite must be met before accomplishing the backup?

- A. Provide a password for the encryption
- B. Set up an Oracle wallet for the encryption
- C. No setup is required as it is a default encryption method
- D. Both Oracle wallet and password must be set up for the encryption

Answer: B

NEW QUESTION 115

- (Topic 4)

You issue the following RMAN command to set a retention policy on a database: RMAN>CONFIGURE RETENTION POLICY TO REDUNDANCY 2;

What will be the outcome of issuing this command?

- A. After two days, a backup will be marked obsolete
- B. After two days, a backup will be deleted from the media
- C. If the RMAN repository has records of two or more recent backups of a file, then older backups will be deleted from the media.
- D. If the RMAN repository has records of two or more recent backups of a file, then older backups will be marked obsolete.

Answer: D

NEW QUESTION 117

- (Topic 5)

Compressed backups work with which of the following commands?

- A. copy as backup
- B. backup as copy
- C. backup
- D. copy

Answer: C

Explanation:

The backup compression only works with backup set, NONE of image copies can work with compression.

NEW QUESTION 121

- (Topic 5)

What is the impact of the following backup if it exceeds the duration allowance? backup as compressed backupset duration 2:00 partial minimize load database ;

- A. The entire backup will fail
- B. It will not be usable for recovery.
- C. The entire backup will fail, but any datafile successfully backed up will be usable for recovery.
- D. If this backup fails, subsequent backups will prioritize datafiles not backed up.
- E. If this backup fails, an error will be raised and any other commands will not be executed.
- F. If this backup fails, no error will be raised and any other commands will be executed.

Answer: B

NEW QUESTION 124

- (Topic 5)

What type of backup is stored in a proprietary RMAN format?

- A. Backup set
- B. Image copy
- C. Backup section
- D. Backup group

Answer: A

Explanation:

A backup set is an RMAN-specific proprietary format, whereas an image copy is a bit-for-bit copy of a file. By default, RMAN creates backup sets. Refer to here

NEW QUESTION 129

- (Topic 5)

You execute the following RMAN command to perform the backup operation:

```
RMAN> RUN
```

```
{  
ALLOCATE CHANNEL c1 DEVICE TYPE disk MAXOPENFILES 8; BACKUP DATABASE FILESPERSET 4;  
}
```

What is the multiplexing level in the preceding backup process?

- A. 4
- B. 8
- C. 7

Answer: A

NEW QUESTION 130

- (Topic 5)

You executed the following command:

```
RMAN> RECOVER COPY OF DATAFILE '/u01/app/oracle/oradata/orcl/users01.dbf';
```

 Which statement regarding the above command is correct?

- A. The '/u01/app/oracle/oradata/orcl/users01.dbf' data file is recovered from the image copy.
- B. The '/u01/app/oracle/oradata/orcl/users01.dbf' data file is recovered from the last incremental backup.
- C. Image copies of the '/u01/app/oracle/oradata/orcl/users01.dbf' data file are updated with all changes up to incremental backup SCN.
- D. Image copies of the '/u01/app/oracle/oradata/orcl/users01.dbf' data file are recovered using the above command if data file recovery fails.

Answer: C

NEW QUESTION 131

- (Topic 5)

Which of the following statements are true about the BACKUP command? (Choose all that apply.)

- A. The BACKUP command can not be used to make image copies of a datafile.
- B. The BACKUP command can improve performance by multiplexing backup files.
- C. The BACKUP can take advantage of the block-change tracking capability.
- D. The BACKUP command cannot store data in incremental backups.
- E. The BACKUP command can store data in cumulative incremental backups only.

Answer: BC

NEW QUESTION 134

- (Topic 5)

If a backup set is expired, what can you do to correct the problem?

- A. Change the retention criteria.
- B. Make the lost backup set pieces available to RMAN again.
- C. Run the crosscheck command to correct the location for the backup set piece contained in the metadata.
- D. Nothin
- E. The backup set piece is lost forever.
- F. Call Oracle support, their assistance is required.

Answer: B

NEW QUESTION 137

- (Topic 5)

What is the result of this command? RMAN> Report need backup days 3;

- A. Lists all datafiles created in the last three days that are not backed up.
- B. Lists all datafiles not recoverable based on the current retention criteria.
- C. Lists all datafiles not backed up in the last three day
- D. The datafile is not recoverable.

- E. Lists all datafiles that need to be backed up due to unrecoverable operations.
- F. Lists all datafiles not backed up in the last three day
- G. It does not imply that the datafile is not recoverable.

Answer: E

NEW QUESTION 140

- (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 141

- (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 143

- (Topic 5)

You are using RMAN to backup your ARCHIVELOG mode database. You have enabled control-file autobackups. Which files are not backed up during the RMAN backup?

- A. Database Datafiles
- B. Database Control Files
- C. Online redo logs
- D. Archived redo logs
- E. The database SPFILE
- F. None of the above, all these files are backed up.

Answer: C

NEW QUESTION 144

- (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 146

- (Topic 6)

You need to restore your database back to 9/30/2008 at 18:00. In what order would you run the following commands to complete this task?

- A. restore controlfile until time '09/30/2008:18:00:00';
- B. restore database until time '09/30/2008:18:00:00';
- C. restore spfile until time '09/30/2008:18:00:00';
- D. recover database until time '09/30/2008:18:00:00';
- E. alter database open resetlogs;
- F. alter database open;
- G. b, d, e
- H. b, d, f
- I. c, a, b, d, e

J. c, a, b, d, f
K. a, b, d, e

Answer: A

NEW QUESTION 150

- (Topic 6)

You have lost all your SYSTEM tablespace datafiles (system_01.dbf and system_02.dbf) and the database has crashed. What would be the appropriate order of operations to correct the situation?

- A. Mount the database with the startup mount command.
- B. Take the SYSTEM data file offline with the alter database command.
- C. Restore the SYSTEM_01.dbf data file from backup media with the required archived redo logs.
- D. Restore all SYSTEM tablespace-related datafiles from backup media.
- E. Issue the recover tablespace SYSTEM command.
- F. Issue the recover data file SYSTEM_01.dbf command.
- G. Open the database with the alter database open command.
- H. Open the database with the alter database open RESETLOGS command.
- I. a, c, f, g
- J. b, d, e, h
- K. a, b, c, f, g
- L. d, a, e, g
- M. b, c, f, e, g

Answer: D

Explanation:

Because there is NO controlfile damaged, and there is NO PITR recover, you don't need to use RESETLOGS option.

NEW QUESTION 152

- (Topic 6)

Which command will result in a trace file being created with the create controlfile command contained in it?

- A. alter database backup controlfile;
- B. alter database backup controlfile to trace;
- C. alter database controlfile backup;
- D. alter database controlfile backup to '/ora01/oracle/ctrl_backup.ctl';
- E. alter database begin controlfile backup;

Answer: B

Explanation:

ALTER DATABASE BACKUP CONTROLFILE TO TRACE

Specify TO TRACE if you want Oracle Database to write SQL statements to a trace file rather than making a physical backup of the control file. You can use SQL statements written to the trace file to start up the database, re-create the control file, and recover and open the database appropriately, based on the created control file. If you issue an ALTER DATABASE BACKUP CONTROLFILE TO TRACE statement while block change tracking is enabled, then the resulting trace file will contain a command to reenable block change tracking.

The trace file will also include ALTER DATABASE REGISTER LOGFILE statements for existing logfiles that reside in the current archivelog destinations. This will implicitly create database incarnation records for the branches of redo to which the logfiles apply.

You can copy the statements from the trace file into a script file, edit the statements as necessary, and use the script if all copies of the control file are lost (or to change the size of the control file).

NEW QUESTION 156

- (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 157

- (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 159

- (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 161

- (Topic 6)

A database has three online redo log groups with one member each. A redo log member with the status ACTIVE is damaged while the database is running. What is the first step you should take to solve this problem?

- A. Attempt to Issue a checkpoint.
- B. Restart the database using the RESETLOGS option.
- C. Drop the redo log number and create it in a different location.
- D. Perform and incomplete recovery up to the most recent available redo log.

Answer: A

Explanation:

Recovering After Losing All Members of an Online Redo Log Group

If a media failure damages all members of an online redo log group, then different scenarios can occur depending on the type of online redo log group affected by the failure and the archiving mode of the database. If the damaged online redo log group is current and active, then it is needed for crash recovery; otherwise, it is not.

Recovering After the Loss of an Online Redo Log Group Inactive

It is not needed for crash recovery Clear the archived or unarchived group. Active

It is needed for crash recovery Attempt to issue a checkpoint and clear the log; if impossible, then you must either use Flashback Database or restore a backup and perform incomplete recovery up to the most recent available redo log.

Current

It is the redo log that the database is currently writing to Attempt to clear the log; if impossible, then you must either use Flashback Database or restore a backup and perform incomplete recovery up to the most recent available redo log.

NEW QUESTION 163

- (Topic 6)

The database is configured in ARCHIVELOG mode and regular complete database backups are taken. The loss of which two types of files may require a recovery with the RESETLOGS option? (Choose two)

- A. Control files
- B. Password files
- C. Inactive online redo log file
- D. Archived log files required to perform recovery
- E. Newly created tablespaces which is not backed up

Answer: AD

Explanation:

The RESETLOGS options is required in:

1. Incomplete Recovery
2. Change of control file

So that, the control file and archived redo logs are required.

NEW QUESTION 165

- (Topic 6)

Another DBA issues a shutdown abort command on a database on which you were running an online backup. What will happen when you try to restart the database?

- A. Oracle will automatically take the datafile out of hot backup mode, generate a warning message, and then open the database.
- B. Oracle will automatically take the datafile out of hot backup mode and then open the database.
- C. Oracle will generate an error when trying to open the database, indicating that a datafile is in hot backup mode
- D. You will need to correct this error before you can open the database.
- E. The database will open with the file in hot backup mode
- F. You can restart the backup at any time.
- G. The datafile in hot backup mode will be corrupted and you will have to recover it.

Answer: C

NEW QUESTION 170

- (Topic 6)

You are peer reviewing a fellow DBA's backup plan for his NOARCHIVELOG mode database, as shown here:

1. Put the tablespaces in backup mode.
2. Back up the datafiles for all tablespaces.

3. Take the tablespaces out of backup mode.
 4. Back up all archived redo logs.
 Your colleague asks for you to comment on his plan. Which response would be correct?

- A. The plan will work as is.
- B. The plan needs to be modified to allow for an archive-log switch after step 3.
- C. The plan needs to be modified so that a backup of the archived redo logs occurs before step 1.
- D. The plan needs to be adjusted to shut down the database after step 1 and to restart the database after step 2.
- E. The plan cannot work as presented.

Answer: B

Explanation:

Without command ALTER SYSTEM SWITCH LOGFILE, the backup of archive redo logs will be useless.

NEW QUESTION 175

- (Topic 6)

Your database is running in ARCHIVELOG mode. One of the data files, USERDATA01.DBF, in the users tablespace is damaged and you need to recover the file until the point of failure. The backup for the data file is available.

Which three files would be used in the user-managed recovery process performed by the database administrator (DBA)? (Choose three.)

- A. Redo logs
- B. Control file
- C. The latest backup of only the damaged data file
- D. The latest backup of all the data file in the USERS tablespace
- E. Temporary files of temporary tablespace
- F. Archive Logs since the latest backup to point of failure

Answer: ACF

NEW QUESTION 178

- (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 181

- (Topic 6)

The database is running in the ARCHIVELOG mode. It has three redo log groups with one member each. One of the redo log groups has become corrupted. You have issued the following command during the recovery of a damaged redo log file:

ALTER DATABASE CLEAR UNARCHIVED LOGFILE GROUP 3;

Which action should you perform immediately after using this command?

- A. You should perform a log switch
- B. You should make a backup of the database
- C. You should switch the database to the NOARCHIVELOG mode
- D. You should shut down the database instance and perform a complete database recovery

Answer: B

NEW QUESTION 182

- (Topic 6)

Which are the correct steps, in order, to deal with the loss of an online redo log if the database has not yet crashed?

- a: Issue a checkpoint.
- b: Shut down the database.
- c: Issue an alter database open command to open the database. d: Startup mount the database.
- e: Issue an alter database clear logfile command. f: Recover all database datafiles.

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

Answer: C

NEW QUESTION 185

- (Topic 6)

Which statement about recovering from the loss of a redo log group is true?

- A. If the lost redo log group is ACTIVE, you should first attempt to clear the log file.
- B. If the lost redo log group is CURRENT, you must clear the log file.
- C. If the lost redo log group is ACTIVE, you must restore, perform cancel-based incomplete recovery, and open the database using the RESETLOGS option.
- D. If the lost redo log group is CURRENT, you must restore, perform cancel-based incomplete recovery, and open the database using the RESETLOGS option.

Answer: D

NEW QUESTION 188

- (Topic 6)

You have lost all your online redo logs. As a result, your database has crashed. You have tried to restart the database and clear the online redo log files, but when you try to open the database you get the following error.

```
SQL> startup
```

```
ORACLE instance started.
```

```
Total System Global Area 167395328 bytes Fixed Size 1298612 bytes
```

```
Variable Size 142610252 bytes Database Buffers 20971520 bytes Redo Buffers 2514944 bytes Database mounted.
```

```
ORA-00313: open failed for members of log group 2 of thread 1
```

```
ORA-00312: online log 2 thread 1: '/oracle01/oradata/orcl/redo02a.log'
```

```
ORA-27037: unable to obtain file status Linux Error: 2: No such file or directory Additional information: 3
```

```
ORA-00312: online log 2 thread 1: '/oracle01/oradata/orcl/redo02.log'
```

```
ORA-27037: unable to obtain file status Linux Error: 2: No such file or directory Additional information: 3
```

```
SQL> alter database clear logfile group 2;
```

```
alter database clear logfile group 2 * ERROR at line 1:
```

```
ORA-01624: log 2 needed for crash recovery of instance orcl (thread 1) ORA-00312: online log 2 thread 1: '/oracle01/oradata/orcl/redo02.log' ORA-00312: online
```

```
log 2 thread 1: '/oracle01/oradata/orcl/redo02a.log'
```

What steps must you take to resolve the error?

- a: Issue the recover database redo logs command.
- b: Issue the Startup Mount command to mount the database. c: Restore the last full database backup.
- d: Perform a point-in-time recovery, applying all archived redo logs that are available.
- e: Restore all archived redo logs generated during and after the last full database backup. f: Open the database using the alter database open resetlogs command.
- g: Issue the alter database open command.

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

Answer: E

Explanation:

If the online redo log is in ACTIVE or CURRENT status, you cannot issue CLEAR LOGFILE GROUP n command, it occurs ORA-01624 error.

The option (a) is invalid, there is NO such recover database redo log command, so that the answer must be (c, e, b, d, f).

It applies an incomplete recovery, then open database with RESETLOGS option.

NEW QUESTION 193

- (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 194

- (Topic 6)

Your database is up and running and one of your three control files is accidentally erased. You start RMAN and run the following command:

```
RESTORE CONTROLFILE FROM AUTOBACKUP;
```

Which of the following statements is true? (Choose all that apply.)

- A. The command restores only the missing control file.
- B. The command restores all the control files.
- C. The command fails because the database is running.
- D. This is the correct way to address this problem.
- E. This is not the correct way to address this problem.

Answer: CE

Explanation:

During the database running, the control files are locked by the database instance, you must shutdown the database and startup at NOMOUNT status to restore a missing control file.

And you have to open database with RESETLOGS option, due to control file change.

NEW QUESTION 196

- (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 201

- (Topic 6)

Your database is in NOARCHIVELOG mode. You start to do a backup, but your users complain that they don't want you to shut down the database to perform the backup. What options are available to you?

- A. Put the database in hot backup mode and perform an online backup, including backing up the archived redo logs.
- B. Just back up the database datafiles without shutting down the database.
- C. You will have to wait until you can shut down the database to perform the backup.
- D. Mark each datafile as backup in progress, back them up individually, and then mark them as backup not in progress.
- E. No archived redo logs will need to be backed up.
- F. Only back up the datafiles that the user will not be touching.
- G. Once the user has finished what they were doing, you can shut down the database and back up the datafiles the user changed during the course of the remaining backup.

Answer: C

NEW QUESTION 206

- (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 210

- (Topic 6)

Your database is running in ARCHIVELOG mode. One of the data files, USERDATA01.dbf, in the USERS tablespace is damaged and you need to recover the file until the point of failure. The backup for the datafile is available.

Which three files would be used in the user-managed recovery process performed by the database administrator (DBA)? (Choose Three)

- A. Redo logs
- B. Control file

- C. The latest backup of only the damaged data file
- D. The latest backup of all the data file in the USERS tablespace
- E. Temporary files of temporary tablespace
- F. Archive Logs since the latest backup to point of failure

Answer: ACF

NEW QUESTION 211

- (Topic 6)

A database is running In ARCHIVELOG mode. It has two online redo log groups and each group has one member. A LGWR Input/output (I/O) falls due to permanent media failure that has resulted In the loss of redo log file and the LWGR terminates causing the instance to crash. The steps to recover from the loss of a current redo log group member in the random order are as follow.

- 1) Restore the corrupted redo log group.
- 2) Restore from a whole database backup.
- 3) Perform incomplete recovery.
- 4) Relocate by renaming the member of the damaged online redo log group to a new location.
- 5) Open the database with the RESETLOGS option.
- 6) Restart the database instance.
- 7) Issue a checkpoint and clear the log.

Identify the option with the correct sequential steps to accomplish the task efficiently.

- A. 1, 3, 4, and 5
- B. 7, 3, 4. and 5
- C. 2, 3, 4, and 5
- D. 7, 4, 3. and 5
- E. Only 6 is required

Answer: C

Explanation:

Recovering After Losing All Members of an Online Redo Log Group

If a media failure damages all members of an online redo log group, then different scenarios can occur depending on the type of online redo log group affected by the failure and the archiving mode of the database.

If the damaged online redo log group is current and active, then it is needed for crash recovery; otherwise, it is not. Table 30-4 outlines the various recovery scenarios.

If the Group Is...	Then...	And You Should...
Inactive	It is not needed for crash recovery	Clear the archived or unarchived group.
Active	It is needed for crash recovery	Attempt to issue a checkpoint and clear the log; if impossible, then you must either use Flashback Database or restore a backup and perform incomplete recovery up to the most recent available redo log.
Current	It is the redo log that the database is currently writing to	Attempt to clear the log; if impossible, then you must either use Flashback Database or restore a backup and perform incomplete recovery up to the most recent available redo log.

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

NEW QUESTION 216

- (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 219

- (Topic 6)

Which mode of database shutdown requires an instance recovery at the time of the next database startup?

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

Answer: A

NEW QUESTION 222

- (Topic 6)

You are trying to recover your database. During the recovery process, you receive the following error:

ORA-00279: change 5033391 generated at 08/17/2008 06:37:40 needed for thread 1 ORA-00289: suggestion:

/oracle01/flash_recovery_area/ORCL/archivelog/2008_08_17 / o1_mf_1_11_%u_.arc

ORA-00280: change 5033391 for thread 1 is in sequence #11

ORA-00278: log file '/oracle01/flash_recovery_area/ORCL/archivelog/2008_08_17 / o1_mf_1_10_4bj6wnqm_.arc' no longer needed for this recovery Specify log:

{<RET>=suggested | filename | AUTO | CANCEL}

ORA-00308: cannot open archived log '/oracle01/flash_recovery_area/ORCL/ archivelog/2008_08_17 /o1_mf_1_11_%u_.arc'

ORA-27037: unable to obtain file status Linux Error: 2: No such file or directory Additional information: 3

How do you respond to this error? (Choose two.)

- A. Restore the archived redo log that is missing and attempt recovery again.
- B. Recovery is complete and you can open the database.
- C. Recovery needs redo that is not available in any archived redo lo
- D. Attempt to apply an online redo log if available.
- E. Recover the entire database and apply all archived redo logs again.
- F. Recovery is not possible because an archived redo log has been lost.

Answer: AC

NEW QUESTION 223

- (Topic 7)

You are using Recovery Manager (RMAN) with a recovery catalog to back up your production database. The backups and the archived redo log files are copied to a tape drive on a daily basis. Because of media failure, you lost your production database completely along with the recovery catalog database. You want to recover the target database and make it functional. You consider performing the following steps to accomplish the task:

1. Restore an autobackup of the server parameter file.
2. Restore the control file
3. Start the target database instance
4. Mount the database
5. Restore the data files
6. Open the database with RESETLOGS option
7. Recover the data files
8. Set DBID for the target database

Which option illustrates the correct sequence that you must use?

- A. 8, 1, 3, 2, 4, 5, 7, 6
- B. 1, 8, 3, 4, 2, 5, 7, 6
- C. 1, 3, 4, 2, 8, 5, 6, 7
- D. 1, 3, 2, 4, 6, 5, 7, 8

Answer: A

Explanation:

Recovering the Database After a Disaster

The procedure for disaster recovery is similar to the procedure for recovering the database with a backup control file in NOCATALOG mode. If you are restoring the database to a new host, then you should also review the considerations described in "Restoring a Database on a New Host".

This scenario assumes that the Linux server on which your database was running has been damaged beyond repair. Fortunately, you backed up the database to Oracle Secure Backup and have the tapes available. The scenario assumes the following:

? Oracle Database is already installed on the new host.

? You are restoring the database to a new Linux host with the same directory structure as the old host.

? You have one tape drive containing backups of all the data files and archived redo logs through log 1124, and autobackups of the control file and server parameter file.

? You do not use a recovery catalog with the database.

To recover the database on the new host:

1. If possible, restore or re-create all relevant network files such as tnsnames.ora and listener.ora and a password file.
2. Start RMAN and connect to the target database instance.

At this stage, no initialization parameter file exists. If you have set ORACLE_SID and ORACLE_HOME, then you can use operating system authentication to connect as SYSDBA. For example, start RMAN as follows:

```
% rman
RMAN> CONNECT TARGET
/
```

3. Specify the DBID for the target database with the SET DBID command, as described in "Restoring the Server Parameter File".

For example, enter the following command: SET DBID 676549873;

4. Run the STARTUP NOMOUNT command.

When the server parameter file is not available, RMAN attempts to start the instance with a dummy server parameter file.

5. Allocate a channel to the media manager and then restore the server parameter file from autobackup.

For example, enter the following command to restore the server parameter file from Oracle Secure Backup:

```
RUN
{
ALLOCATE CHANNEL c1 DEVICE TYPE sbt; RESTORE SPFILE FROM AUTOBACKUP;
}
```

6. Restart the instance with the restored server parameter file. STARTUP FORCE NOMOUNT;
7. Write a command file to perform the restore and recovery operation, and then execute the command file.
 The command file should do the following:
 - a. Allocate a channel to the media manager.
 - b. Restore a control file autobackup (see "Performing Recovery with a Backup Control File and No Recovery Catalog").
 - c. Mount the restored control file.
 - d. Catalog any backups not recorded in the repository with the CATALOG command.
 - e. Restore the data files to their original locations. If volume names have changed, then run SET NEWNAME commands before the restore operation and perform a switch after the restore operation to update the control file with the new locations for the data files, as shown in the following example.
 - f. Recover the data files. RMAN stops recovery when it reaches the log sequence number specified.

```
RMAN> RUN
{
# Manually allocate a channel to the media manager ALLOCATE CHANNEL t1 DEVICE TYPE sbt;
# Restore autobackup of the control file. This example assumes that you have
# accepted the default format for the autobackup name. RESTORE CONTROLFILE FROM AUTOBACKUP;
# The set until command is used in case the database
# structure has changed in the most recent backups, and you want to
# recover to that point in time. In this way RMAN restores the database
# to the same structure that the database had at the specified time. ALTER DATABASE MOUNT;
SET UNTIL SEQUENCE 1124 THREAD 1; RESTORE DATABASE;
RECOVER DATABASE;
}
```

The following example of the RUN command shows the same scenario except with new file names for the restored data files:

```
RMAN> RUN
{
# If you must restore the files to new locations,
# use SET NEWNAME commands:
SET NEWNAME FOR DATAFILE 1 TO '/dev/vgd_1_0/r1vt5_500M_1'; SET NEWNAME FOR DATAFILE 2 TO '/dev/vgd_1_0/r1vt5_500M_2'; SET NEWNAME FOR
DATAFILE 3 TO '/dev/vgd_1_0/r1vt5_500M_3'; ALLOCATE CHANNEL t1 DEVICE TYPE sbt;
RESTORE CONTROLFILE FROM AUTOBACKUP; ALTER DATABASE MOUNT;
SET UNTIL SEQUENCE 124 THREAD 1; RESTORE DATABASE;
SWITCH DATAFILE ALL; # Update control file with new location of data files. RECOVER DATABASE;
}
```

8. If recovery was successful, then open the database and reset the online logs: ALTER DATABASE OPEN RESETLOGS;

NEW QUESTION 227

- (Topic 7)

David managed to accidentally delete the datafiles for database called DSL. He called Heber and Heber tried to help but he managed to delete the control files of the database. Heber called Bill and Bill saved the day.

They are using a recovery catalog for this database.

What steps did Bill perform to recover the database and in what order?

- A. Restored the control file with the RMAN restore controlfile command.
- B. Mounted the DSL instance with the alter database mount command.
- C. Restored the datafiles for the DSL database with the RMAN restore command.
- D. Opened the DSL database with the alter database open resetlogs command.
- E. Recovered the datafiles for the DSL database with the RMAN recover command.
- F. Started the DSL instance.
- G. Connected to the recovery catalog with RMAN.
- H. a, b, c, d, e, f, g
- I. b, c, d, g, f, e, a
- J. g, f, a, b, c, e, d
- K. c, a, d, b, f, e, g
- L. g, f, a, b, e, c, d

Answer: C

Explanation:

About Recovery with a Backup Control File

If all copies of the current control file are lost or damaged, then you must restore and mount a backup control file. You must then run the RECOVER command, even if no data files have been restored, and open the database with the RESETLOGS option. If some copies of the current control file are usable, however, then you can follow the procedure in "Responding to the Loss of a Subset of the Current Control Files" and avoid the recovery and RESETLOGS operation.

When RMAN is connected to a recovery catalog, the recovery procedure with a backup control file is identical to recovery with a current control file. The RMAN metadata missing from the backup control file is available from the recovery catalog. The only exception is if the database name is not unique in the catalog, in which case you must use SET DBID command before restoring the control file.

1. Start RMAN and connect to a target database.
2. Start the target instance without mounting the database. RMAN>STARTUP NOMOUNT;
3. Restore the control file
 RMAN> SET DBID 320066378; # (Optional) If the database name is not unique, you need to specify the DBID
 RMAN> RUN

```
{
SET CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO
'autobackup_format';
RESTORE CONTROLFILE FROM AUTOBACKUP;
}
```
4. Start the target instance with mounting the database. RMAN>STARTUP MOUNT;
5. Restore the data files; RMAN>RESTORE DATABASE;
6. Recover the database; RMAN>RECOVER DATABASE;
7. Open the database with RESETLOGS option; RMAN> ALTER DATABASE OPEN RESETLOGS;

NEW QUESTION 229

- (Topic 7)

You are performing incomplete recovery using RMAN. You execute the following RUN block:

```
RUN
{
SET UNTIL SCN 1107600; RESTORE DATABASE; RECOVER DATABASE;
}
```

Which statement is true about the result?

- A. RMAN restores all datafiles from the most recent backup available since the failure and applies the redo logs necessary to recover the database to SCN 1107600
- B. RMAN restores all datafiles needed to restore the database through SCN 1107599 and applies the redo logs necessary to recover the database through SCN 1107599.
- C. RMAN restores all datafiles and control files from the most recent backup
- D. The RUN block fails because you did not specify an UNTIL clause in your RECOVER DATABASE command

Answer: B

NEW QUESTION 231

- (Topic 7)

You executed the following commands in an RMAN session for your database instance that has failures:

```
RMAN> LIST FAILURE;
```

After some time, you executed the following command in the same session: RMAN> ADVISE FAILURE;

But there are new failures recorded in the Automatic Diagnostic Repository (ADR) after the execution of the last LIST FAILURE command.

Which statement is true for the above ADVISE FAILURE command in this scenario?

- A. It produces a warning for new failures before advising for CRITICAL and HIGH failures.
- B. It ignores new failures and considers the failures listed in the last LIST FAILURE command only.
- C. It produces an error with recommendation to run the LIST FAILURE command before the ADVISE FAILURE command.
- D. It produces advice only for new failures and the failures listed in the last LIST FAILURE command are ignored.

Answer: A

NEW QUESTION 232

- (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 236

- (Topic 8)

What are the two different types of database duplication? (Choose two.)

- A. Active
- B. Passive
- C. Online
- D. Backup-based
- E. Failure driven

Answer: AD

NEW QUESTION 237

- (Topic 8)

How many database instances are used during a database-duplication process?

- A. One
- B. Two
- C. Three

- D. Four
- E. Five

Answer: B

NEW QUESTION 241

- (Topic 8)

When performing a database duplication, which duplicate database parameter would you set to ensure that the online redo logs are created in the correct location?

- A. log_file_name_convert
- B. convert_log_file_name
- C. file_name_convert_log
- D. redo_log_file_name_convert
- E. logfile_convert_directory

Answer: A

NEW QUESTION 246

- (Topic 8)

As part of archiving the historical data, you want to transfer data from one database to another database, which is on another server. All tablespaces in the source database are read/write and online. The source and target databases use the same compatibility level and character sets. View the Exhibit and examine the features in the source and target database.

Which of the following steps are required to transport a tablespace from the database to the target database:

1. Make the tablespace read-only at the source database.
2. Export metadata from the source database.
3. Convert data files by using Recovery Manager (RMAN).
4. Transfer the dump file and data files to the target machine.
5. Import metadata at the target database.
6. Make the tablespace read/write at the target database. Exhibit:

Source:

```
SQL> SELECT tp.endian_format, d.platform_name
 2 FROM v$tablespace tp,
 3 v$tablespace d
 4 WHERE tp.platform_name = d.platform_name;
```

ENDIAN_FORMAT	PLATFORM_NAME
Little	Microsoft Windows IA (32-bit)

Target:

```
SQL> SELECT tp.endian_format, d.platform_name
 2 FROM v$tablespace tp,
 3 v$tablespace d
 4 WHERE tp.platform_name = d.platform_name;
```

ENDIAN_FORMAT	PLATFORM_NAME
Little	Linux IA (32-bit)

- A. 2, 4, and 5
- B. All the steps
- C. 2, 3, 4 and 5
- D. 1, 2, 4, 5 and 6

Answer: D

Explanation:

Refer to here. Generate a Transportable Tablespace Set read_only--expdp--4-import-read_wrtie

1, 2, 4, 5, 6 (optional)

If both platforms have the same endianness, no conversion is necessary. Otherwise you must do a conversion of the tablespace set either at the source or destination database. Transport the dump file to the directory pointed to by the DATA_PUMP_DIR directory object, or to any other directory of your choosing.

Run the following query to determine the location of DATA_PUMP_DIR: SELECT * FROM DBA_DIRECTORIES WHERE DIRECTORY_NAME = 'DATA_PUMP_DIR';

OWNER	DIRECTORY_NAME	DIRECTORY_PATH
SYS	DATA_PUMP_DIR	C:\app\orauser\admin\orawin\dpdump\

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

Transport the data files to the location of the existing data files of the destination database. On the UNIX and Linux platforms, this location is typically

/u01/app/oracle/oradata/SID/ or

+DISKGROUP/SID/ datafile/.

NEW QUESTION 251

- (Topic 8)

Which of the following are prerequisite steps to transport a database? (Choose all that apply.)

- A. Query the V\$TRANSPORTABLE_PLATFORMS view in the source database to determine if the intended destination is listed.
- B. Verify that there are no restrictions or limitations that the source or destination database may encounter.
- C. Verify that the source and destination have the same Oracle version, critical updates, patch-set version, and patch-set exceptions.
- D. Determine if you will perform the conversion on the source or destination platform.
- E. None of the above.

Answer: ABD

NEW QUESTION 252

- (Topic 8)

Which three statements must be true before transporting a tablespace from a database on one platform to a database on another platform? (Choose three.)

- A. Both source and target database must be the same character set
- B. Both source and target database must have the same endian format
- C. The COMPATIBLE parameter must be the same in the source and target databases.
- D. The minimum compatibility level for both the source and target database must be 10.0.0.
- E. All read-only and offline data files that belong to the tablespace to be transported must be platform aware.

Answer: ABD

NEW QUESTION 257

- (Topic 8)

Which two statements are true about a duplicate database that is created by using the DUPLICATE command in RMAN? (Choose two.)

- A. It is a copy or a subset of the target database.
- B. It is opened in RESTRICT mode after a duplicating operation.
- C. It is created by using backups and archived redo log files from the target database.
- D. It is created with the same database identifier (DBID) as that of target database.

Answer: AC

Explanation:

Duplicating a Database

Overview of RMAN Database Duplication

Database duplication is the use of the DUPLICATE command to copy all or a subset of the data in a source database. The duplicate database (the copied database) functions entirely independently from the source database (the database being copied).

Purpose of Database Duplication

If you copy a database with operating system utilities instead of the DUPLICATE command, then the DBID of the copied database remains the same as the original database. To register the copy database in the same recovery catalog with the original, you must change the DBID with the DBNEWID utility (see Oracle Database Utilities). In contrast, the DUPLICATE command automatically assigns the duplicate database a different DBID so that it can be registered in the same recovery catalog as the source database.

Backup-Based Duplication

In backup-based duplication, RMAN creates the duplicate database by using pre-existing RMAN backups and copies. This technique of duplication uses one of the following mutually exclusive subtechniques: Duplication without a target database connection, RMAN obtains metadata about backups from a recovery catalog.

Duplication without a target database connection and without a recovery catalog. RMAN obtains metadata about where backups and copies reside from BACKUP LOCATION. Duplication with a target database connection. RMAN obtains metadata about backups

from the target database control file or from the recovery catalog. How RMAN Duplicates a Database

For backup-based duplication, the principal work of the duplication is performed by the auxiliary channels. These channels correspond to a server session on the auxiliary instance on the destination host. For active database duplication the primary work is performed by target channels. RMAN must perform database point-in-time recovery, even when no explicit point in time is provided for duplication. Point-in-time recovery is required because the online redo log files in the source database are not backed up and cannot be applied to the duplicate database. The farthest point of recovery of the duplicate database is the most recent redo log file archived by the source database.

As part of the duplicating operation, RMAN automates the following steps:

1. Creates a default server parameter file for the auxiliary instance if the following conditions are true:
 - ? Duplication does not involve a standby database.
 - ? Server parameter files are not being duplicated.
 - ? The auxiliary instance was not started with a server parameter file.
2. Restores from backup or copies from active database the latest control file that satisfies the UNTIL clause requirements.
3. Mounts the restored or copied backup control file from the active database.
4. Uses the RMAN repository to select the backups for restoring the data files to the auxiliary instance. This step applies to backup-based duplication.
5. Restores and copies the duplicate data files and recovers them with incremental backups and archived redo log files to a noncurrent point in time.
6. Shuts down and restarts the database instance in NOMOUNT mode.
7. Creates a new control file, which then creates and stores the new DBID in the data files.
8. Opens the duplicate database with the RESETLOGS option and creates the online redo log for the new database.

NEW QUESTION 258

- (Topic 8)

The following query will provide what information about transportable tablespaces for the current database? (Choose all that apply.)

```
SELECT d.platform_name "Source", t.platform_name "Compatible Targets",
       endian_format
FROM v$transportable_platform t, v$database d
WHERE t.endian_format =
       (select endian_format
        from v$transportable_platform t, v$database d
        where d.platform_name = platform_name);
```

- A. The list of target platforms having the same endian format as the source database
- B. The list of target platforms requiring endian conversion
- C. The list of target platforms that will not require endian conversion
- D. The list of all target platforms that can receive transportable tablespaces from the source database
- E. None of the above

Answer: AC

NEW QUESTION 259

- (Topic 8)

Which two operations are NOT performed by the DUPLICATE command in Recovery Manager (RMAN) while duplicating a running database? (Choose Two)

- A. Creating the control file for the duplicate database
- B. Restoring the target data files to the duplicate database
- C. Performing complete recovery using all available backups
- D. Generating a new, unique DBID for the duplicate database
- E. Copying the online redo log files from the target database to the duplicate database

Answer: CE

NEW QUESTION 262

- (Topic 8)

Which operation requires that you create an auxiliary instance manually before executing the operation? (Choose all that apply.)

- A. Backup-based database duplication.
- B. Active database duplication.
- C. Tablespace point-in-time recovery.
- D. No operation requires the creation of an auxiliary instance.

Answer: AB

NEW QUESTION 266

- (Topic 8)

When exporting metadata for the transportable tablespaces, what is the correct next step after confirming endian format?

- A. Export the tablespaces using data pump.
- B. Determine if the transportable set is self-contained.
- C. Convert the datafiles using RMAN.
- D. Copy the datafiles from source to destination.

Answer: B

NEW QUESTION 269

- (Topic 9)

In your production database, users report that they are unable to generate reports on an important table because it does not contain any data. While investigating the reason, you realize that another user executed the TRUNCATE TABLE command, which accidentally caused the data to be lost. Now you want to recover the lost data of the table without affecting objects in other schemas. Which method must you use to recover the lost data?

- A. Complete Recovery with online redo log
- B. Complete Recovery with archived redo log
- C. Tablespace Point-in-Time Recovery (TSPITR)
- D. Incomplete Recovery with system change number (SCN)

Answer: C

NEW QUESTION 274

- (Topic 9)

Which options must you configure while performing an automated Tablespace Point-in-Time Recovery (TSPITR) by using Recovery Manager (RMAN)?

- A. New channels for restore and recovery tasks
- B. New name for the data files of the tablespace
- C. Auxiliary name for the data files of the tablespace
- D. Auxiliary destinations for an auxiliary set of data files

Answer: D

NEW QUESTION 276

- (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 278

- (Topic 10)

You issued the following RMAN command to back up the database:

```
RMAN> RUN{
ALLOCATE CHANNEL c1 DEVICE TYPE sbt BACKUP DATABASE
TAG quarterly KEEP FOREVER
RESTORE POINT FY06Q4;
}
```

Which two statements are true regarding the backup performed? (Choose two.)

- A. Archived redo log files are backed up along with data files.
- B. Only data files are backed up and a restore point named FY06Q4 is created.
- C. Archived log files are backed up along with data files, and the archived log files are deleted.
- D. The command creates a restore point named FY06Q4 to match the SCN at which this backup is consistent.

Answer: AD

Explanation:

Section: Monitoring and Tuning RMAN Explanation

Refer to here. keepOption

Overrides any configured retention policy for this backup so that the backup is not considered obsolete, as shown in Example 2-26.

You can use the KEEP syntax to generate archival database backups that satisfy business or legal requirements. The KEEP setting is an attribute of the backup set (not individual backup piece) or image copy.

Note: You cannot use KEEP with BACKUP BACKUPSET.

With the KEEP syntax, you can keep the backups so that they are considered obsolete after a specified time (KEEP UNTIL), or make them never obsolete (KEEP FOREVER). As shown in Example 2-27, you must be connected to a recovery catalog when you specify KEEP FOREVER.

Note: You can use CHANGE to alter the status of a backup generated with KEEP. See Also: keepOption for more information about backups made with the KEEP option Creating a Consistent Database Backup for Archival Purposes

This example uses a keepOption to create an archival backup set that cannot be considered obsolete for one year. The example backs up the database, archives the redo in the current online logs to ensure that this new backup is consistent, and backs up only those archived redo log files needed to restore the data file backup to a consistent state. The BACKUP command also creates a restore point to match the SCN at which this backup is consistent. The FORMAT parameter must be capable of creating multiple backup pieces in multiple backup sets.

```
BACKUP DATABASE
```

```
FORMAT '/disk1/archival_backups/db_%U.bck' TAG quarterly
```

```
KEEP UNTIL TIME 'SYSDATE + 365' RESTORE POINT Q1FY06;
```

NEW QUESTION 279

- (Topic 10)

You performed the RMAN database backup with the KEEP option. Which two statements are true regarding this backup? (Choose two.)

- A. The backup contains data files, the server parameter file, and the control file even if the control file autobackup is disabled.
- B. The KEEP option overrides the configured retention policy.
- C. The backup contains only data files and archived redo log files.
- D. The KEEP option is an attribute of an individual backup piece.

Answer: AB

NEW QUESTION 280

- (Topic 10)

Which statement describes the significance of the CHANGE FAILURE command in RMAN? (Choose all that apply.)

- A. It is used to change failure priority only for HIGH or LOW priorities.
- B. It is used to execute the advised repair script.
- C. It is used to change failure priority only for the CRITICAL priority.
- D. It is used to explicitly close the open failures.
- E. It is used to inform the database about the repair after the repair script executes.

Answer: AD

NEW QUESTION 281

- (Topic 10)

What does the minimize load database parameter mean when backing up a database?

- A. RMAN will attempt to make the backup run as fast as possible without any IO limitations.
- B. RMAN will automatically restrict the number of channels in use to one.
- C. RMAN will spread the backup IO over the total duration stated in the backup command.
- D. RMAN will skip any datafile that currently is involved in an IO operatio
- E. RMAN will retry backing up the datafile later and an error will be raised at the end of the backup if the datafile cannot be backed up.
- F. Datafiles will be backed up; those having the lowest current number of IO operations will be backed up first.

Answer: C

NEW QUESTION 285

- (Topic 10)

You have enabled backup optimization for the RMAN environment.

Identify two criteria on which RMAN will skip the file, if it has already been backed up. (Choose two.)

- A. The data file backup is done with multiple channels
- B. The data files is in the read-write mode after being backed up in the read only mode
- C. The backup was taken after the data files was taken offline-normal or is in the read only mode
- D. The data file backup complies with the back retention policy and the backup duplexingfeature

Answer: CD

NEW QUESTION 287

- (Topic 10)

Multiple RMAN sessions are connected to the database instance. Examine the following output when backup commands are running in server sessions: What could have helped you to correlate server sessions with channels?

```
SQL> SELECT s.sid, p.spid, s.client _info FROM v$process p, v$session s
WHERE p.addr = s.paddr
AND CLIENT _INFO LIKE 'rman%';
```

```
SID SPID CLIENT_INFO
-----
103 25280 rman channel 1=ORA_DISK_1
151 25292 rman channel 1=ORA_DISK_2
```

- A. Implement RMAN multiplexing
- B. Set the DEBUG ON in the RMAN script
- C. Specify the command ID in the RMAN script
- D. Use a tag with the RMAN BACKUP command

Answer: C

NEW QUESTION 291

- (Topic 11)

Over the course of a day, a department performed multiple DML statements (inserts, updates, deletes) on multiple rows of data in multiple tables. The manager would like a report showing the time, table name, and DML type for all changes that were made. Which Flashback technology would be the best choice to produce the list?

- A. Flashback Drop
- B. Flashback Query
- C. Flashback Transaction Query
- D. Flashback Versions Query
- E. Flashback Table

Answer: C

NEW QUESTION 293

- (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 297

- (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 302

- (Topic 11)

On which two database objects can the VERSIONS clause of the Flashback Versions Query be used? (Choose two.)

- A. fixed tables
- B. heap tables
- C. external tables
- D. temporary tables
- E. Index-Organized Tables (IOTs)
- F. views

Answer: BE

NEW QUESTION 303

- (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 308

- (Topic 11)

At the request of a user, you issue the following command to restore a dropped table: SQL> flashback table "BIN\$F2JFfMq8Q5unbC0ceE9eJg==\$0" to before drop;

Later, the user notifies you that the data in the table seems to be very old and out of date. What might be the problem?

- A. Because a proper range of SCNs was not specified, the wrong data was restored.
- B. A proper range of timestamps was not specified, so the wrong data was restored.
- C. A previous Flashback Drop operation had been performed, resulting in multiple versions of the table being stored in the Recycle Bin.
- D. Either option A or B could be correct.
- E. Not enough information was provided to determine which.
- F. None of the above.

Answer: C

NEW QUESTION 313

- (Topic 12)

Which are the two prerequisites before setting up Flashback Data Archive? (Choose two.)

- A. Flash recovery area must be defined
- B. Undo retention guarantee must be enabled.
- C. Database must be running in archivelog mode.
- D. Automatic undo management must be enabled.
- E. The tablespace in which the Flashback Data Archive is created must have automatic segment space Management (ASSM).

Answer: DE

NEW QUESTION 314

- (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 319

- (Topic 12)

View the Exhibit and examine the output.

You executed the following command to enable Flashback Data Archive on the EXCHANGE_RATE table:

ALTER TABLE exchange_rate FLASHBACK ARCHIVE; What is the outcome of this command?

Exhibit:

```
SQL> SELECT flashback_archive_name,status FROM
2 dba_flashback_archive
3 /
```

FLASHBACK_ARCHIVE_NAME	STATUS
FLA1	

- A. The Flashback Archive is created on the same tablespace where the tables are stored.
- B. The Flashback Archive is created on the SYSAUX tablespace.
- C. The command generates an error because no Flashback Archive name is specified and there is no default Flashback Archive.
- D. The table uses the default Flashback Archive.

Answer: C

NEW QUESTION 320

- (Topic 13)

View the Exhibit and note the contents of V\$DIAG_INFO. Which statement is true about the ADR?

Exhibit:

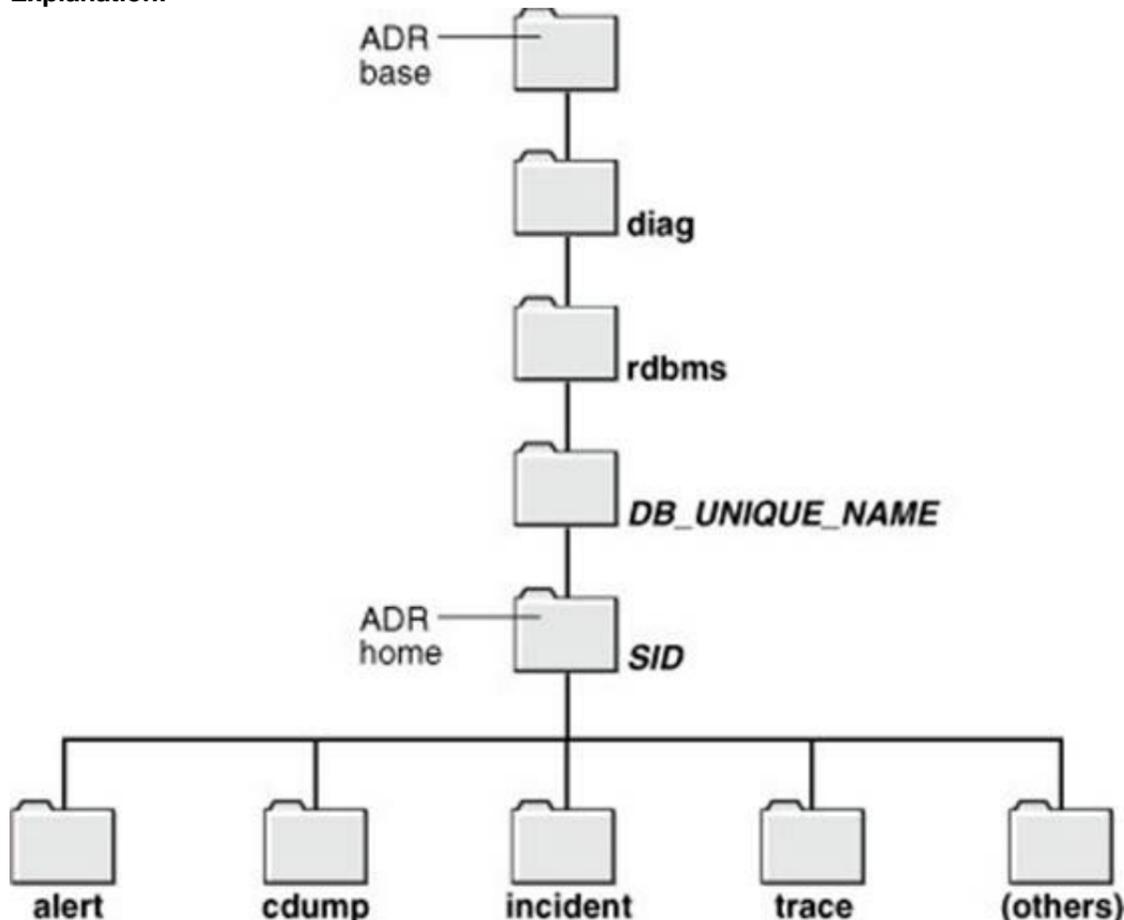
```
SELECT * FROM V$DIAG_INFO;
```

INST_ID	NAME	VALUE
1	Diag	TRUE
1	ADR Base	/u01/oracle
1	ADR Home	/u01/oracle/diag/rdbms/orclbi/orclbi
1	Diag Trace	/u01/oracle/diag/rdbms/orclbi/orclbi/trace
1	Diag Alert	/u01/oracle/diag/rdbms/orclbi/orclbi/alert
1	Diag Incident	/u01/oracle/diag/rdbms/orclbi/orclbi/incident
1	Diag Cdump	/u01/oracle/diag/rdbms/orclbi/orclbi/cdump
1	Health Monitor	/u01/oracle/diag/rdbms/orclbi/orclbi/hm
1	Default Trace File	/u01/oracle/diag/rdbms/orclbi/orclbi/trace/orcl_ora_22769.trc
1	Active Problem Count	8

- A. The text alert log file will be available in Diag Trace
- B. A copy alert log file will be kept in Diag Incident for every incident.
- C. The XML version of the alert log file will be available in Diag Trace.
- D. An Automatic Database Diagnostic Management (ADDM) report is generated and stored in the Health Monitor whenever an incident occurs.

Answer: A

Explanation:



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

alert, The XML-formatted alert log cdump, Core files

incident, Multiple subdirectories, where each subdirectory is named for a particular incident, and where each contains dumps pertaining only to that incident

trace, Background and server process trace files, SQL trace files, and the text-formatted alert log

(others), Other subdirectories of ADR home, which store incident packages, health monitor reports, and other information

NEW QUESTION 324

- (Topic 13)

Which of the following tasks does the tool Incident Packaging Service (IPS) perform?

- A. Cleans up the ADR by deleting files not associated with an incident uploaded to Oracle Support.
- B. Identifies all files associated with a critical error and adds them to a zip file to be sent to Oracle Support.
- C. Automatically opens a Service Request with Oracle Support for each critical error and sends all relevant files.
- D. Displays a high-level view of critical errors on the database home page.

Answer: B

Explanation:

Incident packaging service (IPS) and incident packages

The IPS enables you to automatically and easily gather the diagnostic data—traces, dumps, health check reports, and more—pertaining to a critical error and package the data into a zip file for transmission to Oracle Support. Because all diagnostic data relating to a critical error are tagged with that error's incident number, you do not have to search through trace files and other files to determine the files that are required for analysis; the incident packaging service identifies the required files automatically and adds them to the zip file.

Before creating the zip file, the IPS first collects diagnostic data into an intermediate logical structure called an incident package (package). Packages are stored in the Automatic Diagnostic Repository. If you choose to, you can access this intermediate logical structure, view and modify its contents, add or remove additional diagnostic data at any time, and when you are ready, create the zip file from the package. After these steps are completed, the zip file is ready to be uploaded to Oracle Support.

NEW QUESTION 327

- (Topic 13)

Which of the following statements is true regarding the initialization parameter DIAGNOSTIC_DEST?

- A. The default value is the value of the environment variable \$ORACLE_HOME; if \$ORACLE_HOME isn't set, then the default is set to \$ORACLE_BASE.
- B. The default value is the value of the environment variable \$ORACLE_BASE; if \$ORACLE_BASE isn't set, then it is set to \$ORACLE_HOME.
- C. DIAGNOSTIC_DEST is always equal to \$ORACLE_HOME.
- D. DIAGNOSTIC_DEST is always equal to \$ORACLE_BASE.

Answer: B

Explanation:

The answer is not exact correct.

If the ORACLE_BASE is not set, the default location is \$ORACLE_HOME/log

NEW QUESTION 330

- (Topic 13)

Which setting enables the baselines by default in Oracle Database 11g?

- A. setting the STATISTICS_LEVEL parameter to TYPICAL
- B. adding filters to the Automatic Database Diagnostic Monitor (ADDM)
- C. enabling Automated Maintenance Task using Oracle Enterprise Manager

D. setting the OPTIMIZER_USE_PENDING_STATISTICS parameter to TRUE

Answer: A

NEW QUESTION 333

- (Topic 13)

Which two statements are true regarding Health Monitor checks in Oracle Database 11g? (Choose two.)

- A. Health Monitor checks can be used to scan the contents of the redo log and archive logs for accessibility and corruption.
- B. Health Monitor checks can be used to verify the integrity of database files and report failures if these files are inaccessible, corrupt or inconsistent.
- C. Health Monitor checks can be used to verify the contents of dictionary entries for each dictionary object and fix it automatically.
- D. Health Monitor checks are always initiated manually when there is some critical error.

Answer: AB

NEW QUESTION 338

- (Topic 13)

View the Exhibit to examine the details for an incident. Which statement is true regarding the status of the incident?

Exhibit:

File Name	Size (MB)	Timestamp	Path	View Contents
ora_ora_9007_3937.trc	2.71	August 20, 2007 9:15:14 PM GMT+07:00	/u01/app/oracle/diag/rdbms/orcl/orcl/incident/incdr_3937	
ora_ora_9007.trc	0	August 20, 2007 9:15:14 PM GMT+07:00	/u01/app/oracle/diag/rdbms/orcl/orcl/trace	

- A. The incident has been newly created and is in the process of collecting diagnostic information.
- B. The incident is now in the Done state and the ADR can select the incident to be purged.
- C. The DBA is working on the incident and prefers that the incident be kept in the ADR.
- D. The data collection for the incident is complete and the incident can be packaged and sent to Oracle Support.

Answer: D

NEW QUESTION 341

- (Topic 13)

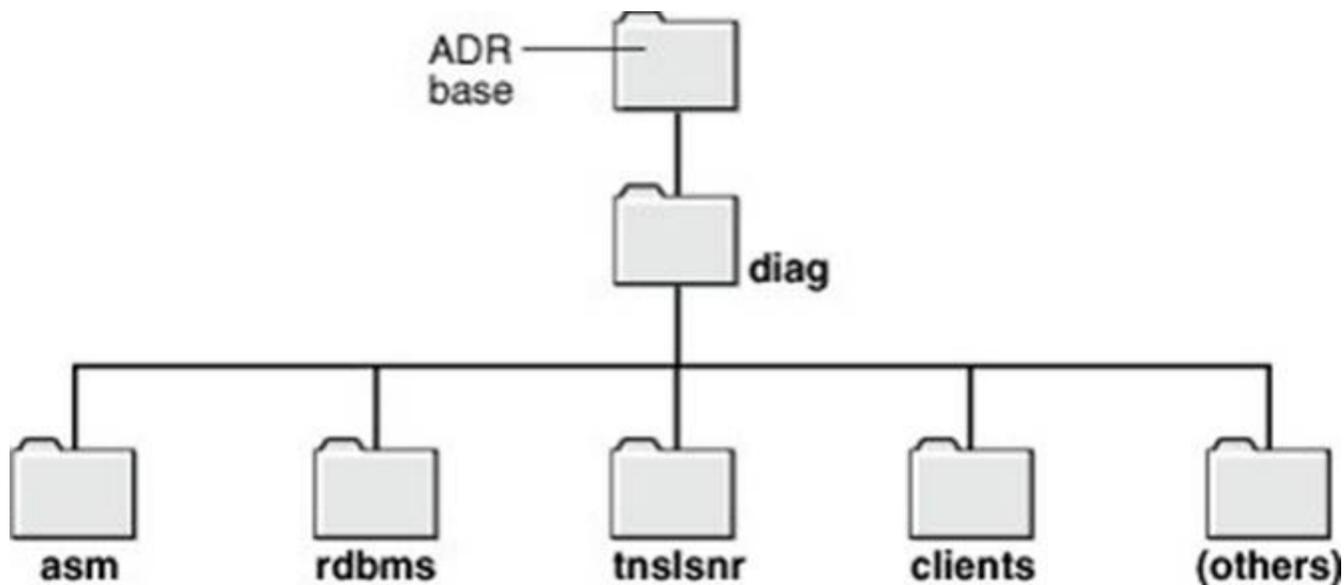
Examine the following values of the initialization parameters in the database having the SID ORCL:
 BACKGROUND_DUMP_DEST=/u01/app/oracle/product/11.1.0/db_1/bdump USER_DUMP_DEST=/u01/app/oracle/product/11.1.0/db_1/udump
 CORE_DUMP_DEST=/u01/app/oracle/product/11.1.0/db_1/cdump DIAGNOSTIC_DEST=
 The environment variables have the following value: ORACLE_BASE=/u01/app/oracle
 ORACLE_HOME=/u01/app/oracle/product/11.1.0/db_1
 What is the location of the Automatic Diagnostic Repository (ADR) home?

- A. /u01/app/oracle/product/11.1.0/db_1
- B. /u01/app/oracle
- C. \$ORACLE_HOME/bdump
- D. \$ORACLE_HOME/log

Answer: B

Explanation:

The Automatic Diagnostic Repository (ADR) is a directory structure that is stored outside of the database. It is therefore available for problem diagnosis when the database is down. The ADR root directory is known as ADR base. Its location is set by the DIAGNOSTIC_DEST initialization parameter. If this parameter is omitted or left null, the database sets DIAGNOSTIC_DEST upon startup as follows:
 ? If environment variable ORACLE_BASE is set, DIAGNOSTIC_DEST is set to the directory designated by ORACLE_BASE.
 ? If environment variable ORACLE_BASE is not set, DIAGNOSTIC_DEST is set to ORACLE_HOME/log.



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

NEW QUESTION 346

- (Topic 13)

Which two statements are true regarding the starting of the database instance using the following command? (Choose two.)
 SQL>STARTUP UPGRADE

- A. It enables all system triggers.
- B. It allows only SYSDBA connections.
- C. It ensures that all job queues remain active during the upgrade process.
- D. It sets system initialization parameters to specific values that are required to enable database upgrade scripts to be run.

Answer: BD

NEW QUESTION 350

- (Topic 13)

Which of the following are true concerning block media recovery? (Choose all that apply.)

- A. Any gap in archive logs ends the recovery.
- B. If a gap in archive logs is encountered, RMAN will search forward for newer versions of the blocks that are not corrupt.
- C. Uncorrupted blocks from the flashback logs may be used to speed recovery.
- D. The database can be in NOARCHIVELOG mode.
- E. None of the above.

Answer: BC

Explanation:

Overview of Block Media Recovery (link) Basic Concepts of Block Media Recovery

Whenever block corruption has been automatically detected, you can perform block media recovery manually with the RECOVER ... BLOCK command. By default, RMAN first searches for good blocks in the real-time query physical standby database, then flashback logs and then blocks in full or level 0 incremental backups.

Prerequisites for Block Media Recovery (link)

The following prerequisites apply to the RECOVER ... BLOCK command:

? The target database must run in ARCHIVELOG mode and be open or mounted with a current control file.

? If the target database is a standby database, then it must be in a consistent state, recovery cannot be in session, and the backup must be older than the corrupted file.

? The backups of the data files containing the corrupt blocks must be full or level 0 backups and not proxy copies.

If only proxy copy backups exist, then you can restore them to a nondefault location on disk, in which case RMAN considers them data file copies and searches them for blocks during block media recovery.

? RMAN can use only archived redo logs for the recovery.

RMAN cannot use level 1 incremental backups. Block media recovery cannot survive a missing or inaccessible archived redo log, although it can sometimes survive missing redo records.

? Flashback Database must be enabled on the target database for RMAN to search the flashback logs for good copies of corrupt blocks.

If flashback logging is enabled and contains older, uncorrupted versions of the corrupt blocks, then RMAN can use these blocks, possibly speeding up the recovery.

? The target database must be associated with a real-time query physical standby database for RMAN to search the database for good copies of corrupt blocks.

Restore Failover (link)

RMAN automatically uses restore failover to skip corrupted or inaccessible backups and look for usable backups. When a backup is not found, or contains corrupt data, RMAN automatically looks for another backup from which to restore the desired files.

NEW QUESTION 355

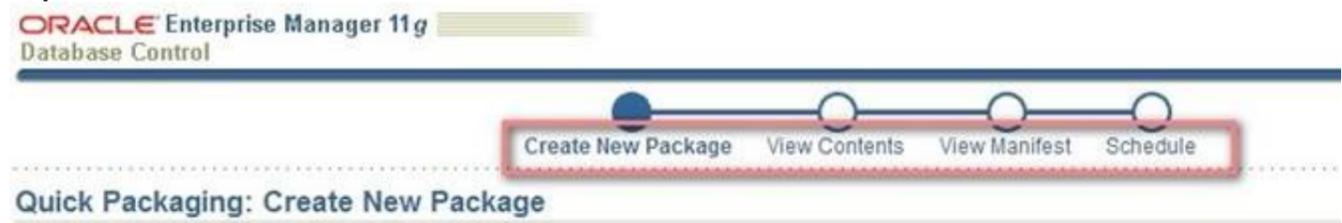
- (Topic 13)

Choose the correct order to package and upload data for an incident to Oracle Support.

- A. Schedule, create new package, view manifest, view contents
- B. Create new package, view manifest, view contents, schedule
- C. Schedule, create new package, view contents, view manifest
- D. Create new package, view contents, view manifest, schedule
- E. None of the above.

Answer: D

Explanation:



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

NEW QUESTION 359

- (Topic 13)

Which of the following is a prerequisite for running DBMS_TDB.CHECK_DB to a successful completion?

- A. The database must be in read-write mode.
- B. The database must have no external files.
- C. The database must open in read-only mode.
- D. The database must be mounted but not opened.

Answer: C

NEW QUESTION 362

- (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 364

- (Topic 13)

Which two statements are true regarding the SQL Repair Advisor? (Choose two.)

- A. The SQL Repair Advisor can be invoked to tune the performance of the regressed SQL statements.
- B. The SQL Repair Advisor can be invoked even when the incident is not active for a SQL statement crash.
- C. The SQL Repair Advisor is invoked by the Health Monitor when it encounters the problematic SQL statement.
- D. The DBA can invoke the SQL Repair Advisor when he or she receives an alert generated when a SQL statement crashes and an incident is created in the ADR.

Answer: BD

Explanation:

You can run SQL Repair Advisor through the EM Support Workbench when an incident raised.

Or You run the SQL Repair Advisor by creating and executing a diagnostic task using the DBMS_SQLDIAG.CREATE_DIAGNOSIS_TASK and DBMS_SQLDIAG. respectively. The SQL Repair Advisor first reproduces the critical error and then tries to produce a workaround in the form of SQL patch.

NEW QUESTION 367

- (Topic 13)

Which two statements are true regarding the Automatic Diagnostic Repository (ADR) in Oracle Database 11g? (Choose two.)

- A. A single ADR can support multiple ADR homes for different database instances.
- B. The alert files are stored in XML file format in the TRACE directory of each ADR home.
- C. If the environmental variable ORACLE_BASE is set, then DIAGNOSTIC_DEST is set to \$ORACLE_BASE.
- D. The BACKGROUND_DUMP_DEST initialization parameter overrides the DIAGNOSTIC_DEST initialization parameter for the location of the alert log file.

Answer: AC

Explanation:

The ADR root directory is known as ADR base. Its location is set by the DIAGNOSTIC_DEST initialization parameter. If this parameter is omitted or left null, the database sets DIAGNOSTIC_DEST upon startup as follows:

? If environment variable ORACLE_BASE is set, DIAGNOSTIC_DEST is set to the directory designated by ORACLE_BASE.

? If environment variable ORACLE_BASE is not set, DIAGNOSTIC_DEST is set to

ORACLE_HOME/log

Within ADR base, there can be multiple ADR homes, where each ADR home is the root directory for all diagnostic data—traces, dumps, the alert log, and so on—for a particular instance of a particular Oracle product or component. For example, in an Oracle Real Application Clusters environment with Oracle ASM, each database instance, Oracle ASM instance, and listener has an ADR home.

NEW QUESTION 369

- (Topic 13)

Examine the section of the Health Check report given below:

DBMS_HM.GET_RUN_REPORT('HM_RUN_1061')

Run Name : HM_RUN_1061 Run Id : 1061

Check Name : Data Block Integrity Check Mode : REACTIVE

Status : COMPLETED

Start Time : 2007-05-12 22:11:02.032292 -07:00

End Time : 2007-05-12 22:11:20.835135 -07:00

Error Encountered : 0 Source Incident Id : 7418

Number of Incidents Created : 0

Which two statements are true regarding the Health Check report? (Choose two.)

- A. Health Check was performed manually.
- B. Health Check was performed to check the disk image block corruptions.
- C. Health Check was performed to check interblock and intersegment corruption.
- D. Health Check was performed to verify the integrity of database files and report failures.
- E. Health Check was performed by the Health Monitor automatically in response to a critical error.

Answer: AB

Explanation:

About Health Monitor Checks

Health Monitor checks (also known as checkers, health checks, or checks) examine various layers and components of the database. Health checks detect file corruptions, physical and logical block corruptions, undo and redo corruptions, data dictionary corruptions, and more. The health checks generate reports of their findings and, in many cases, recommendations for resolving problems. Health checks can be run in two ways:

? Reactive—The fault diagnosability infrastructure can run health checks automatically in response to a critical error.

? Manual—As a DBA, you can manually run health checks using either the

DBMS_HM PL/SQL package or the Enterprise Manager interface. You can run checkers on a regular basis if desired, or Oracle Support may ask you to run a checker while working with you on a service request.

Health Monitor checks store findings, recommendations, and other information in the Automatic Diagnostic Repository (ADR).

Health checks can run in two modes:

? DB-online mode means the check can be run while the database is open (that is, in OPEN mode or MOUNT mode).

? DB-offline mode means the check can be run when the instance is available but

the database itself is closed (that is, in NOMOUNT mode).

All the health checks can be run in DB-online mode. Only the Redo Integrity Check and the DB Structure Integrity Check can be used in DB-offline mode.

Types of Health Checks:

Health monitor runs the following checks:

? DB Structure Integrity Check—This check verifies the integrity of database files and reports failures if these files are inaccessible, corrupt or inconsistent. If the database is in mount or open mode, this check examines the log files and data files listed in the control file. If the database is in NOMOUNT mode, only the control file is checked.

? Data Block Integrity Check—This check detects disk image block corruptions such as checksum failures, head/tail mismatch, and logical inconsistencies within the block. Most corruptions can be repaired using Block Media Recovery. Corrupted block information is also captured in the V\$DATABASE_BLOCK_CORRUPTION view. This check does not detect inter-block or inter-segment corruption.

? Redo Integrity Check—This check scans the contents of the redo log for accessibility and corruption, as well as the archive logs, if available. The Redo Integrity Check reports failures such as archive log or redo corruption.

? Undo Segment Integrity Check—This check finds logical undo corruptions. After locating an undo corruption, this check uses PMON and SMON to try to recover the corrupted transaction. If this recovery fails, then Health Monitor stores information about the corruption in V\$CORRUPT_XID_LIST. Most undo corruptions can be resolved by forcing a commit.

? Transaction Integrity Check—This check is identical to the Undo Segment Integrity Check except that it checks only one specific transaction. Dictionary Integrity Check—This check examines the integrity of core dictionary objects, such as tab\$ and col\$. It performs the following operations:

Verifies the contents of dictionary entries for each dictionary object. Performs a cross-row level check, which verifies that logical constraints on rows in the dictionary are enforced. Performs an object relationship check, which verifies that parent-child relationships between dictionary objects are enforced.

The Dictionary Integrity Check operates on the following dictionary objects:

tab\$, clu\$, fet\$, uet\$, seg\$, undo\$, ts\$, file\$, obj\$, ind\$, icol\$, col\$, user\$, con\$, cdef\$, ccol\$, bootstrap\$, objauth\$, ugroup\$, tsq\$, syn\$, view\$, typed_view\$, superobj\$, seq\$, lob\$, coltype\$, subcoltype\$, ntab\$, refcon\$, opqtype\$, dependency\$, access\$, viewcon\$, icoldep\$, dual\$, sysauth\$, objpriv\$, defrole\$, and ecol\$.

NEW QUESTION 372

- (Topic 13)

Which statement is true when Automatic Workload Repository (AWR) baselines are created using baseline templates?

- A. AWR baselines are always created as repeating baselines.

- B. AWR baselines can be created on the basis of two time values.
- C. AWR baselines are always created with infinite expiration duration.
- D. AWR baselines are always created using the Automatic Workload Repository (AWR) retention period as expiration duration.

Answer: B

NEW QUESTION 377

- (Topic 14)

View the Exhibit to examine the output produced by the following query at three different times since the database instance started and has undergone workloads of different capacities:

SQL> SELECT substr(component, 0, 10) COMP, current_size CS, user_specified_size US FROM v\$memory_dynamic_components WHERE current_size!=0; What do you infer from this? Exhibit:

```
First execution:
=====
```

COMP	CS	US
shared pool	58720256	0
large pool	4194304	0
java pool	4194304	0
SGA Target	176160768	0
DEFAULT bu	100663296	0
Shared IO	8388608	8388608
PGA Target	117440512	0

```
6 rows selected.

Second execution:
=====
```

COMP	CS	US
shared pool	58720256	0
large pool	4194304	0
java pool	4194304	0
SGA Target	192937984	0
DEFAULT bu	117440512	0
Shared IO	8388608	8388608
PGA Target	100663296	0

```
6 rows selected.

Third execution:
=====
```

COMP	CS	US
shared pool	62914560	0
large pool	100663296	0
java pool	4194304	0
SGA Target	192937984	0
DEFAULT bu	8388608	0
Shared IO	8388608	8388608
PGA Target	100663296	0

```
6 rows selected.
```

- A. All sessions are connected to the database instance in dedicated mode, and no RMAN or parallel query operations have been performed.
- B. The database instance is running with manual shared memory management.
- C. The database instance is running with manual PGA management.
- D. The database instance has the MEMORY_TARGET value set to a nonzero value.

Answer: D

NEW QUESTION 379

- (Topic 14)

View the Exhibit to examine the parameter values. You are planning to set the value for the MEMORY_TARGET parameter of your database instance. What value would you assign?

Exhibit:

SQL> SHOW PARAMETER TARGET

NAME	TYPE	VALUE
archive_lag_target	integer	0
db_flashback_retention_target	integer	1440
fast_start_io_target	integer	0
fast_start_mtrr_target	integer	3600
memory_max_target	big integer	0
memory_target	big integer	0
pga_aggregate_target	big integer	90M
sga_target	big integer	272M

- A. 1440 MB
- B. 90 MB
- C. 362 MB
- D. 272 MB

Answer: C

NEW QUESTION 383

- (Topic 14)

View the Exhibit and examine the parameter settings in your server-side parameter file (SPFILE). When you tried to start the database instance, you received the following error:

SQL> startup

ORA-00824: cannot set SGA_TARGET or MEMORY_TARGET due to existing internal settings, see alert log for more information

Why did the instance fail to start? Exhibit:

```
*.audit_file_dest='/u01/app/oracle/admin/orcl/adump'  
*.audit_trail='db'  
*.compatible='11.1.0.0.0'  
*.control_files='/u01/app/oracle/oradata/orcl/control01.ctl', '/  
u01/app/oracle/oradata/orcl/control02.ctl', '/u01/app/oracle/ora  
data/orcl/control03.ctl'  
*.db_block_size=8192  
*.db_domain='us.oracle.com'  
*.db_name='orcl'  
*.db_recovery_file_dest='/u01/app/oracle/flash_recovery_area'  
*.db_recovery_file_dest_size=2147483648  
*.sga_target=436207616  
*.dispatchers='(PROTOCOL=TCP) (SERVICE=orclXDB)'  
*.filesystemio_options='ASYNCH'  
*.job_queue_processes=1000  
*.memory_max_target=629145600  
*.memory_target=629145600  
*.open_cursors=300  
*.processes=150  
*.remote_login_passwordfile='EXCLUSIVE'  
*.statistics_level='BASIC'  
orcl.resource_manager_plan='FORCE:'  
*.undo_tablespace='UNDOTBS1'
```

- A. Because the PGA_AGGREGATE_TARGET parameter is not set
- B. Because the STATISTICS_LEVEL parameter is set to BASIC
- C. Because MEMORY_TARGET and MEMORY_MAX_TARGET cannot be equal
- D. Because both the SGA_TARGET and MEMORY_TARGET parameters are set.

Answer: B

Explanation:

Setting SGA Target Size

You enable the automatic shared memory management feature by setting the SGA_TARGET parameter to a nonzero value. This parameter sets the total size of the SGA. It replaces the parameters that control the memory allocated for a specific set of individual components, which are now automatically and dynamically resized (tuned) as needed.

Note:

The STATISTICS_LEVEL initialization parameter must be set to TYPICAL (the default) or ALL for automatic shared memory management to function.

NEW QUESTION 385

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual 1Z0-053 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the 1Z0-053 Product From:

<https://www.2passeasy.com/dumps/1Z0-053/>

Money Back Guarantee

1Z0-053 Practice Exam Features:

- * 1Z0-053 Questions and Answers Updated Frequently
- * 1Z0-053 Practice Questions Verified by Expert Senior Certified Staff
- * 1Z0-053 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * 1Z0-053 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year