



**Vendor: Oracle**

**Exam Code: 1Z0-053**

**Exam Name: Oracle Database 11g: Administration II**

**Version: Demo**

### QUESTION 1

The INV\_HISTORY table is created using the command:

```
SQL>CREATE TABLE INV_HISTORY (inv_no NUMBER(3), inv_date DATE, inv_amt NUMBER
(10,2))
partition by range (inv_date) interval(numtoyminterval(1,'month'))
(partition p0 values less than (to_date('01-01-2005','dd-mm-yyyy')),
partition p1 values less than (to_date('01-01-2006','dd-mm-yyyy')));
```

The following data has been inserted into the INV\_HISTORY table:

INV_NO	INV_DATE	INV_AMT
=====	=====	=====
1	30-dec-2004	1000
2	30-dec-2005	2000
3	1-feb-2006	3000
4	1-mar-2006	4000
5	1-apr-2006	5000

You would like to store the data belonging to the year 2006 in a single partition and issue the command:

```
SQL> ALTER TABLE inv_history
MERGE PARTITIONS
FOR(TO_DATE('15-feb-2006','dd-mon-yyyy')),
FOR(TO_DATE('15-apr-2006'))
INTO PARTITION sys_py;
```

What would be the outcome of this command?

- A. It executes successfully, and the transition point is set to '1-apr-2006'.
- B. It executes successfully, and the transition point is set to '15-apr-2006'.
- C. It produces an error because the partitions specified for merging are not adjacent.
- D. It produces an error because the date values specified in the merge do not match the date values stored in the table.

**Correct Answer: C**

### QUESTION 2

You want to perform the following operations for the DATA ASM disk group:

- Verify the consistency of the disk.
- Cross-check all the file extent maps and allocation tables for consistency.
- Check whether the alias metadata directory and file directory are linked correctly.
- Check that ASM metadata directories do not have unreachable allocated blocks.

Which command accomplishes these tasks?

- A. ALTER DISKGROUP data CHECK;
- B. ALTER DISKGROUP data CHECK DISK;
- C. ALTER DISKGROUP data CHECK FILE;
- D. ALTER DISKGROUP data CHECK DISK IN FAILURE GROUP 1;

**Correct Answer:** A

**Explanation:**

Syntax: ALTER DISKGROUP <disk\_group\_id> CHECK [REPAIR | NOREPAIR];

The check\_diskgroup\_clause lets you verify the internal consistency of Oracle ASM disk group metadata. The disk group must be mounted. Oracle ASM displays summary errors and writes the details of the detected errors in the alert log.

The CHECK keyword performs the following operations:

Checks the consistency of the disk.

Cross checks all the file extent maps and allocation tables for consistency.

Checks that the alias metadata directory and file directory are linked correctly.

Checks that the alias directory tree is linked correctly.

Checks that Oracle ASM metadata directories do not have unreachable allocated blocks.

### QUESTION 3

Which two statements are true regarding the functionality of the remap command in ASMCMD? (Choose two.)

- A. It repairs blocks that have read disk I/O errors.
- B. It checks whether the alias metadata directory and the file directory are linked correctly.
- C. It repairs blocks by always reading them from the mirror copy and writing them to the original location.
- D. It reads the blocks from a good copy of an ASM mirror and rewrites them to an alternate location on disk if the blocks on the original location cannot be read properly.

**Correct Answer:** AD

**Explanation:**

Reference from the Oracle document release v11.1 at here:

Repairs a range of physical blocks on a disk. The remap command only repairs blocks that have read disk I/O errors. It does not repair blocks that contain corrupted contents, whether or not those blocks can be read. The command assumes a physical block size of 512 bytes and supports all allocation unit sizes (1 to 64 MB). Reference from the Oracle document release v11.2 at here:

The remap command marks a range of blocks as unusable on the disk and relocates any data allocated in that range.

### QUESTION 4

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

- A. It improves resync operations.
- B. This feature enables much faster file opens.
- C. It improves performance as fewer extent pointers are needed in the shared pool.

D. It improves performance by reading from a copy of an extent closest to the node.

**Correct Answer:** D

**Explanation:**

Preferred Read Failure Groups

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

#### QUESTION 5

Examine the following command:

```
ALTER DISKGROUP data MOUNT FORCE;
```

In which scenario can you use the above command to mount the disk group?

- A. when ASM disk goes offline
- B. when one or more ASM files are dropped
- C. when some disks in a disk group are offline
- D. when some disks in a failure group for a disk group are rebalancing

**Correct Answer:** C

**Explanation:**

In the FORCE mode, Oracle ASM attempts to mount the disk group even if it cannot discover all of the devices that belong to the disk group. This setting is useful if some of the disks in a normal or high redundancy disk group became unavailable while the disk group was dismounted. When MOUNT FORCE succeeds, Oracle ASM takes the missing disks offline.

If Oracle ASM discovers all of the disks in the disk group, then MOUNT FORCE fails. Therefore, use the MOUNT FORCE setting only if some disks are unavailable. Otherwise, use NOFORCE. In normal- and high-redundancy disk groups, disks from one failure group can be unavailable and MOUNT FORCE will succeed. Also in high-redundancy disk groups, two disks in two different failure groups can be unavailable and MOUNT FORCE will succeed. Any other combination of unavailable disks causes the operation to fail, because Oracle ASM cannot guarantee that a valid copy of all user data or metadata exists on the available disks.

#### QUESTION 6

Which background process of a database instance, using Automatic Storage Management (ASM), connects as a foreground process into the ASM instance?

- A. ASMB
- B. PMON
- C. RBAL
- D. SMON

**Correct Answer:** A

**Explanation:**

ASMB (ASM Background Process): Communicates with the ASM instance, managing storage and providing statistics, runs in ASM instances when the ASMCMD cp command runs or when the database instance first starts if the server parameter file is stored in ASM. ASMB also runs with Oracle Cluster Registry on ASM.

RBAL (ASM Rebalance Master Process): In an ASM instance, it coordinates rebalance activity for disk groups. In a database instances, it manages ASM disk groups.

PMON (Process Monitor): Monitors the other background processes and performs process recovery when a server or dispatcher process terminates abnormally.

SMON (System Monitor Process): Performs critical tasks such as instance recovery and dead transaction recovery, and maintenance tasks such as temporary space reclamation, data dictionary cleanup, and undo tablespace management.

### QUESTION 7

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.

**Correct 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.

### QUESTION 8

Identify three key features of ASM. (Choose three.)

- A. file striping
- B. allocation unit mirroring
- C. automatic disk rebalancing
- D. automatic file size increment
- E. automatic undo management

**Correct Answer:** ABC

### QUESTION 9

You have three production databases, HRDB, FINDB, and ORGDB, that use the same ASM instance. At the end of the day, while all three production database instances are running, you execute the following command on the ASM instance:

```
SQL> shutdown immediate;
```

What is the result of executing this command?

- A. The ASM instance is shut down, but the other instances are still running.
- B. It results in an error because other database instances are connected to it.
- C. All the instances, including the ASM instance, are shut down in the IMMEDIATE mode.
- D. HRDB, FINDB, and ORGDB instances are shut down in the ABORT mode and the ASM instance is shut down in the IMMEDIATE mode.

**Correct Answer: B**

### QUESTION 10

You are managing an ASM instance. You previously issued the following statements:

```
ALTER DISKGROUP dg1 DROP DISK disk2;  
ALTER DISKGROUP dg1 DROP DISK disk3;  
ALTER DISKGROUP dg1 DROP DISK disk5;
```

You want to cancel the disk drops that are pending for the DG1 disk group.

Which statement should you issue?

- A. ALTER DISKGROUP dg1 UNDROP disk2, disk3, disk5;
- B. ALTER DISKGROUP dg1 UNDROP;
- C. ALTER DISKGROUP dg1 UNDROP DISKS;
- D. You cannot cancel the pending disk drops.

**Correct Answer: C**

**Explanation:**



Use this clause to cancel the drop of disks from the disk group. You can cancel the pending drop of all the disks in one or more disk groups (by specifying `diskgroup_name`) or of all the disks in all disk groups (by specifying `ALL`).

This clause is not relevant for disks that have already been completely dropped from the disk group or for disk groups that have been completely dropped. This clause results in a long-running operation. You can see the status of the operation by querying the `V$ASM_OPERATION` dynamic performance view.

**QUESTION 11**

What is the effect of increasing the value of the ASM\_POWER\_LIMIT parameter?

- A. The number of DBWR processes increases
- B. The number of ASMB processes increases
- C. The number of DBWR\_TO\_SLAVES increases
- D. The rebalancing operation in an ASM instance completes more quickly, but can result in higher I/O overhead

**Correct Answer:** D

**QUESTION 12**

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

**Correct 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

**QUESTION 13**

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.

**Correct Answer: C**

**Explanation:**

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

**QUESTION 14**

What is the net effect of the following command?

```
alter diskgroup dgroup1 drop disk abc;
```

- A. The disk ABC will be dropped from the disk group. Since you did not issue a rebalance command, the data on that disk will be lost.
- B. The command will raise an error indicating that you need to rebalance the disk group to remove the data from that disk prior to dropping the disk.
- C. The disk group will be automatically rebalanced during the drop operation. Once the rebalancing is complete, the disk will be dropped.
- D. This command will fail because you cannot drop a specific disk in an ASM disk group.
- E. The disk drop command will be suspended for a predetermined amount of time, waiting for you to also issue an alter diskgroup rebalance command. Once you have issued the rebalance command, ASM will proceed to rebalance the disk group and then drop the disk.

**Correct Answer: C**

**QUESTION 15**

Which of the following is not a configurable attribute for an individual disk group?

- A. AU\_SIZE
- B. COMPATIBLE.RDBMS
- C. COMPATIBLE.ASM
- D. DISK\_REPAIR\_TIME
- E. DG\_DROP\_TIME

**Correct Answer: E**

**Explanation:**

DG\_DROP\_TIME is an invalid DG attribute.



#### Disk Group Attributes

The DISK\_REPAIR\_TIME disk group attribute specifies how long a disk remains offline before ASM drops the disk.

The COMPATIBLE.ASM attribute determines the minimum software version for an ASM instance that uses the disk group.

The COMPATIBLE.RDBMS attribute determines the minimum COMPATIBLE database initialization parameter setting for any database instance that uses the disk group. The AU\_SIZE attribute determines the allocation unit size of the disk group. The values can be 1, 2, 4, 8, 16, 32, and 64 MB.

#### QUESTION 16

Your organization decided to upgrade the existing Oracle 10g database to Oracle 11g database in a multiprocessor environment.

At the end of the upgrade, you observe that the DBA executes the following script:

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

What is the significance of executing this script?

- A. It performs parallel recompilation of only the stored PL/SQL code.
- B. It performs sequential recompilation of only the stored PL/SQL code.
- C. It performs parallel recompilation of any stored PL/SQL as well as Java code.
- D. It performs sequential recompilation of any stored PL/SQL as well as Java code.

**Correct Answer:** C

**Explanation:**

Recompile invalid objects with utlrp.sql

#### QUESTION 17

You are maintaining the SALES database. You have added a new disk to a disk group. Automatic Storage Management performs the rebalancing activity. You want to speed up the rebalancing activity. Which parameter should you specify to control the speed of the rebalancing activity?

- A. ASM\_POWER\_LIMIT
- B. ASM\_DISKSTRING
- C. ASM\_DISKGROUPS
- D. INSTANCE\_TYPE

**Correct Answer:** A

#### QUESTION 18

What are the recommendations for Oracle Database 11g installation to make it Optimal Flexible Architecture (OFA)-compliant? (Choose all that apply.)

- A. ORACLE\_BASE should be set explicitly.
- B. An Oracle base should have only one Oracle home created in it.
- C. Flash recovery area and data file location should be on separate disks.
- D. Flash recovery area and data file location should be created under Oracle base in a non-Automatic Storage Management (ASM) setup.

**Correct Answer:** ACD

**QUESTION 19**

In your database, the LDAP\_DIRECTORY\_SYSAUTH initialization parameter has been set to YES and the users who need to access the database as DBAs have been granted SYSDBA enterprise role in Oracle Internet Directory (OID). SSL and the password file have been configured. A user SCOTT with the SYSDBA privilege tries to connect to the database instance from a remote machine using the command:

```
$ SQLPLUS scott/tiger@DB01 AS SYSDBA
```

Which DB01 is the net service name.

Which authentication method would be used first?

- A. authentication by password file
- B. authentication by using certificates over SSL
- C. authentication by using the Oracle Internet Directory
- D. authentication by using the local OS of the database server

**Correct Answer:** A

**QUESTION 20**

You are managing an Oracle Database 11g database with the ASM storage. The database is having big file tablespaces. You want files to open faster and less memory to be used in the shared pool to manage the extent maps. What configuration would you effect to achieve your objective? (Choose all that apply.)

- A. Set the ASM compatibility attribute for the ASM disk group to 11.1.0.
- B. Set the RDBMS compatibility attribute for the ASM disk group to 11.1.0.
- C. Set the COMPATIBLE initialization parameter for the ASM instance to 11.1.0.
- D. Set the COMPATIBLE initialization parameter for the database instance to 11.1.0.

**Correct Answer:** AD

**QUESTION 21**

Which two statements are true regarding an Automatic Storage Management (ASM) instance? (Choose two.)

- A. An ASM instance mounts an ASM control file
- B. An ASM instance uses the ASMB process for rebalancing of disks within a disk group
- C. Automatic Memory Management is enabled in an ASM instance even when the MEMORY\_TARGET parameter is not set explicitly
- D. An RDBMS instance gets connected to an ASM instance using ASMB as a foreground process when the database instance is started

**Correct Answer:** CD

### QUESTION 22

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.

**Correct 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.

### QUESTION 23

Examine the following ALTER command;

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

What is the purpose of the command?

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

**Correct Answer: A**

#### Explanation:

The key point is PENDING.

### QUESTION 24

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.

**Correct Answer: C**

**QUESTION 25**

Your database instance is running. You are not able to access Oracle Enterprise Manager Database Control because the listener is not started. Which tool or utility would you use to start the listener?

- A. Oracle Net Manager
- B. Listener Control utility
- C. Database Configuration Assistant
- D. Oracle Net Configuration Assistant

**Correct Answer: B**

**QUESTION 26**

View the Exhibit and examine the disk groups created at the time of migrating the database storage to Automatic Storage Management (ASM). Why does the FRA disk group initially have more free space even though both DATA and FRA disk groups are provided with the same size?

**Create Disk Group**

Show SQL Cancel OK

\* Name

Redundancy  HIGH  NORMAL  EXTERNAL

Allocation Unit (MB)

An allocation unit (AU) is the fundamental unit in which contiguous disk space is allocated to ASM files. ASM file extent size is a multiple of AUs. The AU size cannot be modified later.

**Candidate Member Disks**

Select	Path	Header Status	Library	Label	ASM Disk Name	Size	Unit	Force Reuse	Failure Group
<input type="checkbox"/>	/devices/diske1	CANDIDATE	SYSTEM			1024	MB	<input type="checkbox"/>	
<input type="checkbox"/>	/devices/diske2	CANDIDATE	SYSTEM			1024	MB	<input type="checkbox"/>	
<input type="checkbox"/>	/devices/diske3	CANDIDATE	SYSTEM			1024	MB	<input type="checkbox"/>	
<input type="checkbox"/>	/devices/diskk1	CANDIDATE	SYSTEM			1024	MB	<input type="checkbox"/>	
<input type="checkbox"/>	/devices/diskk2	CANDIDATE	SYSTEM			1024	MB	<input type="checkbox"/>	

- A. Because the FRA disk group will not support dynamic rebalancing
- B. Because the FRA disk group is not configured to support mirroring
- C. Because disks in the FRA disk group are not formatted at this stage
- D. Because the FRA disk group will support only a single size of allocation unit

**Correct Answer:** B

**QUESTION 27**

What are three benefits of using ASM? (Choose three.)

- A. Ease of disk administration and maintenance
- B. Load balancing across physical disks
- C. Software RAID-1 data redundancy with double or triple mirrors
- D. Automatic recovery of failed disks

**Correct Answer:** ABC

**QUESTION 28**

What components are present in an ASM instance? (Choose three.)

- A. SGA
- B. Database processes
- C. Database datafiles
- D. Control files
- E. Database parameter file or SPFILE

**Correct Answer:** ABE

**QUESTION 29**

Which of the following is a benefit of ASM fast disk resync?

- A. Failed disks are taken offline immediately but are not dropped.
- B. Disk data is never lost.
- C. By default, the failed disk is not dropped from the disk group ever, protecting you from loss of that disk.
- D. The failed disk is automatically reformatted and then resynchronized to speed up the recovery process.
- E. Hot spare disks are automatically configured and added to the disk group.

**Correct Answer:** A

**Explanation:**

ASM Fast Mirror Resync

**QUESTION 30**

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

**Correct Answer:** C

**QUESTION 31**

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

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

**Correct Answer:** E

**Explanation:**

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

**QUESTION 32**

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

**Correct Answer:** B

**Explanation:**

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.

**QUESTION 33**

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

**Correct Answer:** A

**QUESTION 34**

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 "DGROUP3"
```

```
ORA-15063: ASM discovered an insufficient number of disks for diskgroup "DGROUP2"
```

```
ORA-15063: ASM discovered an insufficient number of disks for diskgroup "DGROUP1"
```

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	DGROUP1, DGROUP2, DGROUP3
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: DGROUP1, DGROUP2, and DGROUP3.  
The underlying disks for these disk groups have apparently been lost.
- B. The format of the ASM\_DISKGROUPS parameter is incorrect. It should reference the disk group numbers, not the names of the disk groups
- C. 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.
- D. The ASM\_DISKSTRING parameter is not set; therefore disk discovery is not possible.
- E. There is insufficient information to solve this problem.

**Correct 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.

### QUESTION 35

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

**Correct 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.

### QUESTION 36

How can you reverse the effects of an ALTER DISKGROUP ... DROP DISK command if it has NOT yet 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.

**Correct Answer: B**



### QUESTION 37

To reference existing ASM files, you need to use a fully qualified ASM filename. Your development database has a disk group named DG2A, the database name is DEV19, and the ASM file that you want to reference is a datafile for the USERS02 tablespace. Which of the following is a valid ASM filename for this ASM file?

- A. dev19/+DG2A/datafile/users02.701.2
- B. +DG2A/dev19/datafile/users02.701.2
- C. +DG2A/dev19/users02/datafile.701.2
- D. +DG2A.701.2
- E. +DG2A/datafile/dev19.users.02.701.2

**Correct Answer: B**

#### Explanation:

Fully Qualified File Name Form ([link](#))

A fully qualified file name has the following form:

+diskgroup/dbname/filetype/filetypetag.file.incarnation

Where:

+diskgroup is the disk group name preceded by a plus sign. You can think of the plus sign (+) as the root directory of the Oracle ASM file system, similar to the slash (/) on UNIX or Linux computers.

dbname is the DB\_UNIQUE\_NAME of the database to which the file belongs. filetype is the Oracle file type and can be one of the file types shown in Table 7-3. filetypetag is type-specific information about the file and can be one of the tags shown in Table 7-3.

file.incarnation is the file/incarnation pair, used to ensure uniqueness.

For example,

filetype = 'DATAFILE', Data files and data file copies  
filetypetag = 'tblspname', Name of the tablespace into which the file is added

### QUESTION 38

Which background process coordinates the rebalance activity for disk groups?

- A. ORBn
- B. OSMB
- C. RBAL
- D. ASMn

**Correct Answer: C**

#### Explanation:

RBAL

ASM Rebalance Master Process

Coordinates rebalance activity

In an ASM instance, it coordinates rebalance activity for disk groups. In a database instances, it manages ASM disk groups.

**QUESTION 39**

On the development database rac0, there are six raw devices: /dev/raw/raw1 through /dev/raw/raw6. /dev/raw/raw1 and /dev/raw/raw2 are 8GB each, and the rest are 6GB each. An existing disk group +DATA1, of NORMAL REDUNDANCY, uses /dev/raw/raw1 and /dev/raw/raw2. Which series of the following commands will drop one of the failure groups for +DATA1, create a new disk group +DATA2 using two of the remaining four raw devices, and then cancel the drop operation from +DATA1?

A. ALTER DISKGROUP DATA1 DROP DISK DATA1\_0001;

```
CREATE DISKGROUP DATA2 NORMAL REDUNDANCY
  FAILGROUP DATA1A DISK '/dev/raw/raw3'
  FAILGROUP DATA1B DISK '/dev/raw/raw4';
```

```
ALTER DISKGROUP DATA1 UNDROP DISKS;
```

B. ALTER DISKGROUP DATA1 DROP DISK DATA1\_0001;

```
CREATE DISKGROUP DATA2 HIGH REDUNDANCY
  FAILGROUP DATA1A DISK '/dev/raw/raw3'
  FAILGROUP DATA1B DISK '/dev/raw/raw4';
```

```
ALTER DISKGROUP DATA1 UNDROP DISKS;
```

C. ALTER DISKGROUP DATA1 DROP DISK DATA1\_0001;

```
CREATE DISKGROUP DATA2 NORMAL REDUNDANCY
  FAILGROUP DATA1A DISK '/dev/raw/raw3'
  FAILGROUP DATA1B DISK '/dev/raw/raw4';
```

```
ALTER DISKGROUP DATA1 UNDROP DATA1_0001;
```

D. ALTER DISKGROUP DATA1 DROP DISK DATA1\_0001

```
ADD DISK GROUP DATA2 NORMAL REDUNDANCY
  FAILGROUP DATA1A DISK '/dev/raw/raw3'
  FAILGROUP DATA1B DISK '/dev/raw/raw4';
```

```
ALTER DISKGROUP DATA1 UNDROP DISKS;
```

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

**Correct Answer: A**

**QUESTION 40**

Which type of database file is spread across all disks in a disk group?

- A. All types of files are spread across all disks in the disk group.
- B. Datafiles
- C. Redo log files
- D. Archived redo log files
- E. Control files

**Correct Answer:** A

**QUESTION 41**

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.

**Correct Answer:** A

**Explanation:**

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

**QUESTION 42**

Which of the following ALTER DISKGROUP commands does not use V\$ASM\_OPERATION to record the status of the operation?

- A. ADD DIRECTORY
- B. DROP DISK
- C. RESIZE DISK
- D. REBALANCE
- E. ADD FAILGROUP

**Correct Answer:** A

**QUESTION 43**

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.

**Correct Answer:** BD

#### QUESTION 44

You are managing an Oracle Database 11g instance with ASM storage. You lost an ASM disk group DATA. You have RMAN backup of data as well as ASM metadata backup. You want to re-create the missing disk group by using the ASMCMD md\_restore command. Which of these methods would you use to achieve this? (Choose all that apply.)

- A. Restore the disk group with the exact configuration as the backed-up disk group, using the same disk group name, same set of disks, failure group configurations, and data on the disk group.
- B. Restore the disk group with the exact configuration as the backed-up disk group, using the same disk group name, same set of disks, and failure group configurations.
- C. Restore the disk group with changed disk group specification, failure group specification, disk group name, and other disk attributes.
- D. Restore metadata in an existing disk group by passing the existing disk group name as an input parameter.

**Correct Answer:** BCD

#### Explanation:

Purpose

The md\_restore command restores disk groups from a metadata backup file.

Syntax and Description

```
md_restore backup_file [--silent]
[--full|--nodg|--newdg -o 'old_diskgroup:new_diskgroup [,...]]
[-S sql_script_file] [-G 'diskgroup [,diskgroup...]]
backup_file
```

Reads the metadata information from backup\_file. -silent

Ignore errors. Typically, if md\_restore encounters an error, it stops. Specifying this flag ignores any errors. -full

Specifies to create a disk group and restore metadata. -nodg

Specifies to restore metadata only. -newdg -o old\_diskgroup:new\_diskgroup]

Specifies to create a disk group with a different name when restoring metadata. The -o option is required with -newdg. S sql\_script\_file

Write SQL commands to the specified SQL script file instead of executing the commands. G diskgroup

Select the disk groups to be restored. If no disk groups are defined, then all disk groups are restored.

#### QUESTION 45

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

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

**Correct Answer:** CD

**QUESTION 46**

Which two are the uses of the ASM metadata backup and restore (AMBR) feature? (Choose two.)

- A. It can be used to back up all data on ASM disks.
- B. It can be used to recover the damaged ASM disk group along with the data.
- C. It can be used to gather information about a pre-existing ASM disk group with disk paths, disk name, failure groups, attributes, templates, and alias directory structure.
- D. It can be used to re-create the ASM disk group with its attributes.

**Correct Answer:** CD

**QUESTION 47**

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.

**Correct Answer:** ABD

**QUESTION 48**

You are managing Oracle Database 11g with an ASM storage with high redundancy. The following command was issued to drop the disks from the dga disk group after five hours:

```
ALTER DISKGROUP dga OFFLINE DISKS IN FAILGROUP f2 DROP AFTER 5H;
```


Which statement is true in this scenario?

- A. It starts the ASM fast mirror resync.
- B. All the disks in the dga disk group would be OFFLINE and the DISK\_REPAIR\_TIME disk attribute would be set to 5 hours.
- C. It drops all disk paths from the dga disk group.
- D. All the disks in the dga disk group in failure group f2 would be OFFLINE and the DISK\_REPAIR\_TIME disk attribute would be set to 5 hours.

**Correct Answer:** D


**Explanation:**

*disk offline clause::=*



Description of the illustration [disk\\_offline\\_clause.gif](#)

*timeout\_clause::=*



Description of the illustration [timeout\\_clause.gif](#)

<http://blog.csdn.net/rlhua>

### QUESTION 49

You are managing an Oracle Database 11g database with ASM storage. The ASM disk group has the COMPATIBLE.ASM attribute set to 11.1. Which statements are true regarding extent management and allocation units in the ASM disk group? (Choose all that apply.)

- A. The `au_size` disk group attribute determines the size of allocation units in the disk group.
- B. The allocation unit size may vary but the extent size is fixed.
- C. The allocation unit size and extent size are fixed for all the disks in a disk group and cannot be changed.
- D. Extent management is completely automated.

**Correct Answer:** AD

### QUESTION 50

You issued the following command to mount the DATA disk group in restricted mode:

```
ALTER DISKGROUP data MOUNT RESTRICT;
```

What is the implication of this command?

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

**Correct Answer:** D

#### Explanation:

MOUNT (link)

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

In the RESTRICTED mode, the disk group is mounted in single-instance exclusive mode. No other Oracle ASM instance in the same cluster can mount that disk group. In this mode the disk

group is not usable by any Oracle ASM client. In the NORMAL mode, the disk group is mounted in shared mode, so that other Oracle ASM instances and clients can access the disk group. This is the default.

## EnsurePass.com Members Features:

1. Verified Answers researched by industry experts.
2. Q&As are downloadable in PDF and VCE format.
3. 98% success Guarantee and **Money Back** Guarantee.
4. Free updates for **180** Days.
5. **Instant Access to download the Items**

View list of All Exam provided:

<http://www.ensurepass.com/certifications?index=A>

To purchase Lifetime Full Access Membership click here: <http://www.ensurepass.com/user/register>

**Valid Discount Code for 2015: JREH-G1A8-XHC6**

To purchase the HOT Exams:

<u>Cisco</u>		<u>CompTIA</u>		<u>Oracle</u>	<u>VMWare</u>	<u>IBM</u>
<a href="#"><u>100-101</u></a>	<a href="#"><u>640-554</u></a>	<a href="#"><u>220-801</u></a>	<a href="#"><u>LX0-101</u></a>	<a href="#"><u>1Z0-051</u></a>	<a href="#"><u>VCAD510</u></a>	<a href="#"><u>C2170-011</u></a>
<a href="#"><u>200-120</u></a>	<a href="#"><u>200-101</u></a>	<a href="#"><u>220-802</u></a>	<a href="#"><u>N10-005</u></a>	<a href="#"><u>1Z0-052</u></a>	<a href="#"><u>VCP510</u></a>	<a href="#"><u>C2180-319</u></a>
<a href="#"><u>300-206</u></a>	<a href="#"><u>640-911</u></a>	<a href="#"><u>BR0-002</u></a>	<a href="#"><u>SG0-001</u></a>	<a href="#"><u>1Z0-053</u></a>	<a href="#"><u>VCP550</u></a>	<a href="#"><u>C4030-670</u></a>
<a href="#"><u>300-207</u></a>	<a href="#"><u>640-916</u></a>	<a href="#"><u>CAS-001</u></a>	<a href="#"><u>SG1-001</u></a>	<a href="#"><u>1Z0-060</u></a>	<a href="#"><u>VCAC510</u></a>	<a href="#"><u>C4040-221</u></a>
<a href="#"><u>300-208</u></a>	<a href="#"><u>640-864</u></a>	<a href="#"><u>CLO-001</u></a>	<a href="#"><u>SK0-003</u></a>	<a href="#"><u>1Z0-474</u></a>	<a href="#"><u>VCP5-DCV</u></a>	<a href="#"><u>RedHat</u></a>
<a href="#"><u>350-018</u></a>	<a href="#"><u>642-467</u></a>	<a href="#"><u>ISS-001</u></a>	<a href="#"><u>SY0-301</u></a>	<a href="#"><u>1Z0-482</u></a>	<a href="#"><u>VCP510PSE</u></a>	<a href="#"><u>EX200</u></a>
<a href="#"><u>352-001</u></a>	<a href="#"><u>642-813</u></a>	<a href="#"><u>JK0-010</u></a>	<a href="#"><u>SY0-401</u></a>	<a href="#"><u>1Z0-485</u></a>		<a href="#"><u>EX300</u></a>
<a href="#"><u>400-101</u></a>	<a href="#"><u>642-832</u></a>	<a href="#"><u>JK0-801</u></a>	<a href="#"><u>PK0-003</u></a>	<a href="#"><u>1Z0-580</u></a>		
<a href="#"><u>640-461</u></a>	<a href="#"><u>642-902</u></a>			<a href="#"><u>1Z0-820</u></a>		

