

Oracle

Exam 1z0-053

Oracle Database 11g: Administration II

Version: 14.2

[Total Questions: 698]

Topic break down

Topic	No. of Questions
Topic 1: Database Architecture and ASM	65
Topic 10: Monitoring and Tuning RMAN	24
Topic 11: Using Flashback Technology	36
Topic 12: Additional Flashback Operations	23
Topic 13: Diagnosing the Database	44
Topic 14: Managing Memory	32
Topic 15: Managing Database Performance	72
Topic 16: Space Management	50
Topic 17: Managing Resources	26
Topic 18: Automating Tasks with the Scheduler	19
Topic 19: Administering the Scheduler	17
Topic 2: Configuring for Recoverability	52
Topic 20: Globalization	21
Topic 3: Using the RMAN Recovery Catalog (RC)	31
Topic 4: Configuring Backup Specifications	22
Topic 5: Using RMAN to Create Backups	42
Topic 6: Performing User-Managed Backup and Recovery	58
Topic 7: Using RMAN to Perform Recovery	25
Topic 8: Using RMAN to Duplicate a Database	23
Topic 9: Performing Tablespace Point-in-Time Recovery (TSPITR)	16

Topic 1, Database Architecture and ASM
Question No : 1 - (Topic 1)

View the Exhibit and examine the disk groups created at the time of migrating the database storage to Automatic Storage Management (ASM).

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"/>	

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?

- 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

Answer: B

Question No : 2 - (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. Once it's installed, bring the disk group online again.
- B. Change the ASM parameter `ASM_PREFERRED_READ_FAILURE_GROUPS` to indicate that you want to read from the non-failed disk. Once the disk controller is replaced, reset the parameter to its original value.
- C. You have no choice but to rebuild the disk. Drop the disk from the disk group and wait for the controller to be replaced. Once the controller is replaced, add the disk back into the disk group and allow ASM to rebuild it.
- D. 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. The database will be unavailable until you can correct the problem and perform recovery.
- E. Change the attribute `DISK_REPAIR_TIME` on the disk group to a time greater than 24 hours.

Answer: E

Question No : 3 - (Topic 1)

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;

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.

Refer to here

Question No : 4 - (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

Question No : 5 - (Topic 1)

If you use ALTER DISKGROUP ... ADD DISK and specify a wildcard for the discovery string, what happens to disks that are already a part of the same or another disk group?

- A. The command fails unless you specify the FORCE option.
- B. The command fails unless you specify the REUSE option.
- C. The command must be reissued with a more specific discovery string.
- D. The other disks, already part of the disk group, are ignored.



Answer: C

Explanation:

Refer to an Oracle Example for adding disk to diskgroup ([link](#))



How A Disk is Discovered

While an ASM instance is initialized, ASM discovers and examines the contents of all of the disks that are in the paths that you designated with values in the ASM_DISKSTRING initialization parameter. Disk discovery also occurs when you:

-  Run the ALTER DISKGROUP...ADD DISK and ALTER DISKGROUP...RESIZE DISK commands
-  Query the V\$ASM_DISKGROUP and V\$ASM_DISK views


Disk Discovery Rules([link](#))

The rules for discovering ASM disks are as follows:


-  ASM can discover up to 10,000 disks. That is, if more than 10,000 disks match the ASM_DISKSTRING initialization parameter, then ASM discovers only the first 10,000.
-  ASM only discovers disks that contain a partition table.

Note:

ASM does not discover a disk that contains an operating system partition table, even if the disk is in an ASM disk string search path and ASM has read and write permission for the disk.

-  When adding a disk, the FORCE option must be used if ASM recognizes that the disk was managed by Oracle. Such a disk appears in the V\$ASM_DISK view with a status of FOREIGN. In this case, you can only add the disk to a disk group by using the FORCE keyword.

In addition, ASM identifies the following configuration errors during discovery:

 Multiple paths to the same disk

In this case, if the disk is part of a disk group, then disk group mount fails. If the disk is being added to a disk group with the ADD DISK or CREATE DISKGROUP command, then the command fails. To correct the error, adjust the ASM_DISKSTRING value so that ASM will not discover multiple paths to the same disk. Or if you are using multipathing software, then ensure that you include only the pseudo-device name in the ASM_DISKSTRING value. See "ASM and Multipathing".

 Multiple ASM disks with the same disk header

This can be caused by having copied one disk onto another. In this case, the disk group mount operation fails.

Question No : 6 - (Topic 1)

Examine the following ALTER command;

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

What is the purpose of the command?

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

Answer: A

Explanation:

The key point is PENDING.

Question No : 7 - (Topic 1)

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.

Answer: A,D

Question No : 8 - (Topic 1)

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.

Answer: C

Question No : 9 - (Topic 1)

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.

Answer: B,C,D

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 No : 10 - (Topic 1)

You issued the following command:

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

An index is then created on the ADMIN_WORK_AREA temporary table.

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

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

Answer: B,D

Question No : 11 - (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.