

Oracle

1Z0-068

Questions & Answers

Oracle Database 12c: RAC and Grid Infrastructure Administration

(Demo Version - Limited Content)



Version: 10.0

Question: 1

Which two statements are true regarding ASM Dynamic Volume manager (ADVM)?

- A. To create an ADVM volume, disk group attributes COMAPTIBLE.ASM and COMPATIBLE.ADVM must be minimally set to 12.1.
- B. An ADVM volume is individually named and can contain only one file system.
- C. Only one ADVM volume can be created in an ASM disk group.
- D. ADVM extends ASM by providing a device driver interface to storage backed by an ASM disk group.
- E. File systems contained in ADVM volumes may only contain database files.

Answer: B,D

Explanation:

- B: Each volume is individually named and may be configured for a single file system.
- D: Oracle ADVM extends Oracle ASM by providing a disk driver interface to Oracle ASM storage allocated as Oracle ADVM volume files.

Question: 2

Which three statements are true about ASM Cloud File System (ACFS) replication?

- A. ACFS auditing information is replicated from the primary file system to the standby file system.
- B. One site of an ACFS replication configuration can be host both primary and standby file systems.
- C. Replication is automatically terminated if the primaries file system has less than 2GB free space.
- D. Standby redo log files are required on the standby site for synchronous redo transport.
- E. The privilege SYSREPL has been introduced for ACFS replication.

Answer: B,C,D

Explanation:

- B: A site can host both primary and standby file systems.
- C: If the primary file system has less than 2 GB available free disk space, Oracle ACFS attempts to automatically terminate replication on the primary file system.
- D: On the standby system, the remote file server (RFS) receives redo data over the network from the LGWR process and writes the redo data to the standby redo log files.

Question: 3

Examine this command to create a volume in the DATA disk group:

```
SQL>ALTER DISKGROUP DATA ADD VOLUME vo11 SIZE 10g HIGH
STRIPE_WIDTH 1M;
STRIPE_COLUMNS 1;
```

The DATA disk group has 50GB free space.

Which two are prerequisites for successful execution of this command?

- A. COMPATIBLE.ASM and COMPATIBLE.ADVM must be set to 11.2 or higher for the DATA disk group.
- B. The DATA disk group must not be created with external redundancy.
- C. The DATA disk group must not contain any other volume.
- D. The DATA disk group must have at least three failure groups.
- E. The DATA disk group must have an AU size of 1MB.

Answer: A,D

Explanation:

The compatibility parameters COMPATIBLE.ASM and COMPATIBLE.ADVM must be set to 11.2 or higher for the disk group. A high redundancy disk group must contain at least three failure groups.

Question: 4

Which two statements are true about vsam views in a clustered environment ?

- A. They exist both in ASM and RDBMS instances, and display the same output.
- B. Their names can be seen in dict when connected to an ASM instance.
- C. Their names can be seen in vfixed_table when connected to an instance of any type that is started.
- D. They exist in both in ASM and RDBSM instances, but they can display different outputs depending on the instance type.
- E. gvasm_attribute can be used to display the ASM_POWER_LIMIT attribute of all ASM instances in the cluster.

Answer: B,D

Explanation:

You can use the views to obtain information about Oracle Automatic Storage Management Cluster File System (Oracle ACFS). These views are accessible from the Oracle ASM instance.

References:

https://docs.oracle.com/cd/E11882_01/server.112/e18951/asmviews.htm#OSTMG94187

Question: 5

Which two statements are true about ASM default templates?

- A. Default data file templates are only created by ASM when a normal redundancy disk group is created.
- B. Default template settings for redundancy depends on the disk group redundancy.
- C. Default templates for a disk group can be modified.

- D. Default template setting for striping depends on disk group redundancy.
- E. Templates may be shared across disk groups.

Answer: B,C

Explanation:

Templates are used to set attributes of files created in an Oracle ASM disk group.

When a file is created, redundancy and striping attributes are set for that file based on an explicitly named template or the system template that is the default template for the file type.

When a disk group is created, Oracle ASM creates a set of default templates for that disk group. The set consists of one template for each file type (data file, control file, redo log file, and so on) that is supported by Oracle ASM. For example, a template named ONLINELOG provides the default file redundancy and striping attributes for all redo log files written to Oracle ASM disks. Default template settings depend on the disk group type. The default template for data files for a normal redundancy disk group sets two-way mirroring, while the corresponding default template in a high redundancy disk group sets three-way mirroring. You can modify these default templates.

Question: 6

Which three statements are true regarding Flex ASM on a four-node cluster consisting of three-hub nodes and a leaf node?

- A. An ASM instance and database instances can coexist on the same hub node.
- B. A database instance on a hub node can be a client of an ASM instance running on any leaf node.
- C. A database instance on a hub node can be a client of an ASM instance running on any hub node.
- D. A database instance on a leaf node can be a client of an ASM instance running on any hub node.
- E. A database instance on a hub node uses an ASM instance as an I/O server when requesting I/O to diskgroups.
- F. ASM Cluster File System (ASFS) can only be deployed on hub nodes.

Answer: A,C,F

Question: 7

Which two statements are true about Flex ASM in a Flex Cluster?

- A. Oracle databases from releases older than 12c cannot use Flex ASM.
- B. ACFS may only be configured on a hub node running an ASM instance.
- C. Instances of ASM-based databases on hosts that have no ASM instance running, require Flex ASM.
- D. Flex ASM requires that I/Os requests be sent through an ASM Proxy Instance, regardless of database and ASM instance placement in a cluster.
- E. Hub nodes with connections from multiple leaf nodes must run an ASM instance.

Answer: A,C

Question: 8

A directory +DATA/ORCL/DATAFILE/USERS exists in ASM.
USERS is a user created directory; ORCL and DATAFILE are system generated directories.
Which two statements are true about user-created directories?

- A. You can create subdirectories under the USERS directory.
- B. You cannot rename a user-created subdirectory.
- C. You cannot drop a directory containing aliases unless you delete the aliases first.
- D. You cannot create a user-defined directory at the root (+) level.
- E. You cannot create a nested subdirectory /USERS/2014 under USERS with a single CREATE DIRECTORY command even if directory USER1 doesn't exist.

Answer: D,E

Question: 9

Which three components are integral parts of High Availability NFS (HANFS)?

- A. Application VIP
- B. DHCP
- C. ExportFS
- D. HAVIP
- E. NFS
- F. HAIP

Answer: C,D,E

Explanation:

CD: In addition to ACFS/ADVM and ASM, HANFS also relies on new Oracle 12.1 Clusterware (CRS) resources, namely the HAVIP and the ExportFS.

DE: High Availability Network File Storage (NFS) for Oracle Grid Infrastructure provides uninterrupted service of NFS V2/V3/V4 exported paths by exposing NFS exports on Highly Available Virtual IPs (HAVIP) and using Oracle Clusterware agents to ensure that the HAVIPs and NFS exports are always online.

References: <https://docs.oracle.com/database/121/OSTMG/GUID-4D7A5911-8FA6-47DC-98BB-593B2C90D43E.htm#OSTMG95489>

<http://www.oracle.com/ocom/groups/public/@otn/documents/webcontent/2011281.pdf>

Question: 10

You just added an ASM disk to the DATA diskgroup.
Which two can be used to monitor the rebalancing?

- A. ams_cmd lsop
- B. v\$asm_disk
- C. v\$asm_operation
- D. v\$asm_diskgroup

- E. v\$session_longops
- F. amscmd lsdg

Answer: A,C

Question: 11

Examine this command:

```
SQL> CREATE DISKGROUP RECO HIGH REDUNDANCY:
```

```
FAILGROUP fgrpl1 DISK
'/dev/disk1' NAME disk1,
'/dev/disk2' NAME disk2,
'/dev/disk3' NAME disk3,
FAILGROUP fgrpl2 DISK
'/dev/disk4' NAME disk4,
'/dev/disk5' NAME disk5,
'/dev/disk6' NAME disk6,
FAILGROUP fgrpl3 DISK
'/dev/disk7' NAME disk7,
'/dev/disk8' NAME disk8,
'/dev/disk9' NAME disk9,
ATTRIBUTE 'au_size'='4M',
'compatible.asm' = '12.1',
'compatible.rdbms' = '12.1',
'compatible.advm' = '12.1',
'content_type' = 'recovery';
```

Which two statements are true about this disk group?

- A. File created in it have three copies for each allocation unit.
- B. Any database instance with the COMPATIBLE parameter 10.1 and higher can use it.
- C. The Compatible attribute values of the disk group can be changed to lower values.
- D. Files created using a template with the redundancy attributes set to MIRROR will have three-way Mirroring.
- E. The RECO diskgroup may only be used for the FRA.
- F. Files created using a template with the redundancy attribute set to HIGH will have three-way mirroring.

Answer: A,F

Explanation:

For Oracle ASM to mirror files, specify the redundancy level as NORMAL REDUNDANCY (2-way mirroring by default for most file types) or HIGH REDUNDANCY (3-way mirroring for all files).

Note:

When a file is created, redundancy and striping attributes are set for that file based on an explicitly named template or the system template that is the default template for the file type.

Question: 12

Which three statements are true about node listeners in a Grid Infrastructure Cluster using Flex ASM?

- A. They run from the Grid Infrastructure home by default.
- B. They listen only for database instances on the same node.
- C. They listen only for the ASM instance on the same node.
- D. There only can be one node listener per node.
- E. They listen on behalf on any ASM instance and all database instances on the same node.

Answer: A,D,E

Question: 13

Which three statements are true about ASM Cloud File System (ACFS) auditing?

- A. Audit information gathered by each cluster node can be consolidated into the unified ACFS audit trail.
- B. acfsutil audit archive must be used to archive audit files that are larger than 10MB.
- C. acfsutil audit init must be run by a system administrator.
- D. acfsutil audit purge can be run by an audit manager.
- E. acfsutil audit read marks audit files to indicate that it is safe to purge them.

Answer: C,D,E

Explanation:

C: The acfsutil audit init command must be run by the system administrator before enabling auditing for any of the audit sources on a file system.

D: Only an audit manager can run the acfsutil audit purge command.

E: acfsutil audit read marks the audit trail to indicate to the audit manager that the log archive file for the current node has been reviewed, backed up as necessary, and is safe to purge.

Question: 14

Which three statements are true regarding ACFS snapshots in Oracle 12c ASM?

- A. They are contained in a directory on the same volume as their parent file system.
- B. All ACFS snapshot operations for all filesystem are serialized clusterwide.
- C. You can create a snapshot from an existing snapshot.
- D. They can be administered using ASMCMD.
- E. They can only be accessed when the file system is mounted.
- F. They are always point-in-time copies of the parent file system.

Answer: B,C,E

Explanation:

All Oracle ACFS snapshot operations are serialized clusterwide in the kernel.

You chose to create a snapshot image on an existing snapshot group, select a snapshot group from

the existing snapshot group table (make sure you select a snapshot group that has not reached its maximum limit of snapshot images).

Before you can access the files on a file system, you need to mount the file system.

References:

<https://docs.oracle.com/database/121/OSTMG/GUID-5A3EF695-A795-4FEA-8BE2-AF657BD2238C.htm#OSTMG33000>

http://mysupport.netapp.com/NOW/public/eseries/amw/index.html#page/SANtricity_v11.20_Array_Management_Window_Help/create_PIT_image.html (Creating a Snapshot image 3rd point, 1st line).

<https://docs.oracle.com/cd/E19455-01/805-7228/6j6q7ueup/index.html> (1st paragraph, 1st line).

Question: 15

Which two statements are true about initialization parameters for Clustered ASM instances?

- A. A change to ASM_DISKSTRING requires a restart of all ASM instances for the change to take effect.
- B. ASM_POWER_LIMIT controls the number of RDBMS instances that can access a disk group while it is being rebalanced.
- C. The ASM_DISKGROUP parameter is optional.
- D. The default value of INSTANT_TYPE is ASM.
- E. The maximum value of ASM_POWER_LIMIT is 1024.

Answer: C,E

Explanation:

C: The default value of the ASM_DISKGROUPS parameter is a NULL string. If the parameter value is NULL or is not specified, then ASM does not mount any disk groups.

E: The ASM_POWER_LIMIT initialization parameter specifies the default power for disk rebalancing in a disk group. The range of values is 0 to 1024.

EXAMS SUCCESS

Thank you for trying

**Our 1Z0-068 Exam Dumps PDF
Demo**

Try 1Z0-068 practice question

If you want to try 1Z0-068 Exam Practice Test
Questions So go to below link and try it!

<https://www.exams4success.com/1z0-068/practice-questions>