1Z0-068

Oracle Database 12c: RAC and Grid
Infrastructure Administration

DEMO

https://www.islever.com/1z0-068.html
https://www.islever.com/oracle.html

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: BD

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.

Reference:

https://docs.oracle.com/database/121/OSTMG/GUID-B7A61F3B-C22A-4021-A846-F5EA749E79FF.htm

#OSTMG32000 (See Note and the third para from the top.)

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: BCD

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.

Reference:

https://docs.oracle.com/database/121/OSTMG/GUID-8522E0B6-BE6E-4BE3-B7C7-5EA2ED4A602C.htm#OSTMG34600 https://docs.oracle.com/cd/B19306_01/server.102/b14239/log_transport.htm#i1280979

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. COMAPTIBLE.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: AD

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.

Reference: https://docs.oracle.com/cd/E11882_01/server.112/e18951/asmdiskgrps.htm#OSTMG10072

https://docs.oracle.com/cd/E18283_01/server.112/e16102/asmdiskgrps.htm

QUESTION 4

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: CD

Explanation:

C: Oracle ASM automatically rebalances when the configuration of a disk group changes. By default, the

ALTER DISKGROUP statement does not wait until the operation is complete before returning. Query the

V$ASM_OPERATION view to monitor the status of this operation.

D: Verify that the rebalance is waiting:

SQL> select * from v$asm_operation

where group_number = (select GROUP_NUMBER from v$asm_diskgroup where NAME='BDT');

References: https://docs.oracle.com/cd/E18283_01/server.112/e16102/asmdiskgrps.htm

QUESTION 5

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