



1Z0-068^{Q&As}

Oracle Database 12c: RAC and Grid Infrastructure Administration

Pass Oracle 1Z0-068 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.geekcert.com/1z0-068.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Oracle
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





QUESTION 1

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.

Correct Answer: CE

Section: (none)

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.

Incorrect Answers:

B: 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.

D: The default value of INSTANT_TYPE is RDBMS.

Reference: https://docs.oracle.com/cd/E11882_01/server.112/e18951/asminst.htm#BHCEHJGA

QUESTION 2

Examine this command:

```
SQL> CREATE DISKROUP data NORMAL REDANDUNCY:
```

```
FAILGROUP fgrp1 DISK
```

```
`/dev/disk1` NAME disk1,
```

```
`/dev/disk2` NAME disk2,
```

```
`/dev/disk3` NAME disk3,
```

```
FAILGROUP fgrp2 DISK
```

```
`/dev/disk4` NAME disk4,
```

```
`/dev/disk5` NAME disk5,
```



```
`/dev/disk6` NAME disk6,  
FAILGROUP fgrp3 DISK  
`/dev/disk7` NAME disk7,  
`/dev/disk8` NAME disk8,  
`/dev/disk9` NAME disk9,  
ATTRIBUTE `au_size`='4M',  
`compatible.asm` = '11.12',  
`compatible.rdbms` = '11.12',  
`compatible.advm` = '11.12',  
`content_type` = `recovery`;
```

Which two statements are true about this disk group, created using Grid Infrastructure 12.1?

- A. Each disk in each of the three fail groups is a mirror of the corresponding disk of the other two failure groups.
- B. File created in this disk group always has two copies for each allocation unit: one primary and one secondary.
- C. It can contain ASM Dynamic Volume Manager (ADVM) volumes.
- D. Any database instance with the COMPATIBLE parameter set to 11.2 or higher can access it.
- E. The ASM compatible attributes can be changed to 11.1 for this disk group.

Correct Answer: BC

Section: (none)

B: NORMAL REDUNDANCY requires the existence of at least two failure groups (see the FAILGROUP clause that follows). Automatic Storage Management provides redundancy for all files in the disk group according to the attributes specified in the disk group templates. NORMAL REDUNDANCY disk groups can tolerate the loss of one group.

C: The COMPATIBLE.ADVM attribute determines whether the disk group can contain Oracle ADVM volumes.

Incorrect Answers:

E: The COMPATIBLE.ADVM attribute value must be set to 11.2 or higher.

Reference: https://docs.oracle.com/database/121/SQLRF/statements_5009.htm#SQLRF01114

QUESTION 3

Which three statements are true about using ADVM volumes and mounting ACFS file systems from leaf nodes in a Flex Cluster? (Choose three.)

- A. Partitioning ADVM-managed volumes using fdisk or other utilities on a leaf node is not supported.



- B. An ADVM volume cannot be used as a boot device on a leaf node.
- C. To use ADVM volumes, an ASM proxy instance is not required on a leaf node.
- D. ADVM volumes can be used to store the OCR and voting files to be used by a leaf node.
- E. An ADVM volume cannot be used as the root file system on a leaf node.
- F. To use ADVM volumes on a leaf node, the oraclesacfs, oracleoks, and oracleadv drivers must be loaded on all hub nodes.

Correct Answer: BEF

Section: (none)

Reference: <https://docs.oracle.com/database/121/CWAIX/storage.htm#CWAIX561>

QUESTION 4

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.

Correct Answer: ACF

Section: (none)

Architecturally Oracle Flex Cluster comprises of a Hub and Leaf architecture where in only the Hub nodes will only have direct access to Oracle Cluster Registry (OCR) and Voting Disk (VD). However application can access the database via

Leaf nodes without ASM instance NOT running on Leaf nodes. The connection to the database is through Hub making it transparent for the application.

Note:

Hub Nodes are connected among them via private network and have direct access to the shared storage just like previous versions. These nodes are the ones that access the Oracle Cluster Registry (OCR) and Voting Disk (VD) directly.

Leaf Nodes are lighter and are not connected among them, neither accesses the shared storage like the Hub Nodes. Each Leaf Node communicates with the Hub Node that is attached to, and it is connected to the cluster via the Hub Node

that is linked to.



References: <http://www.oracle.com/technetwork/articles/database/flexasm-flexcluster-benefits-odb12c-2177371.html>

QUESTION 5

Which two types of IP address details are administered using the `srvctl` command for Oracle Clusterware 12c?

- A. the Grid Naming Service (GNS) VIP
- B. public IP addresses
- C. High Available VIP (HAIP) addresses
- D. Single Client Access Name (SCAN) VIPs
- E. Storage Network IP addresses

Correct Answer: AD

Section: (none)

A: To convert a cluster that is not running GNS to a GNS server cluster, run the following command as root, providing a valid IP address and a domain:

```
# srvctl add gns -vip IP_address -domain domain
```

D: Example:

1.

Modify the VIP using a VIP name that resolves to IPv6 by running the following command as root:

```
# srvctl modify vip -node node_name -address vip_name -netnum network_number Do this once for each node.
```

2.

2. Modify the SCAN using a SCAN name that resolves to IPv6 by running the following command:

```
$ srvctl modify scan -scanname scan_name
```

Incorrect Answers:

C: The `oifcfg` command is used to configure HAIP addresses. For example: `oifcfg setif -global eth3/172.16.2.0:cluster_interconnect`

References: <https://docs.oracle.com/database/121/CWADD/admin.htm#CWADD92641>

[Latest 1Z0-068 Dumps](#)

[1Z0-068 VCE Dumps](#)

[1Z0-068 Braindumps](#)