



IT認證考試題庫 專業平臺

考證寶提供最新考古題與模擬試題
協助您高效通過認證考試

www.kaozhengpro.com

Exam : **1Z0-058**

Title : Oracle Real Application
Clusters 11g Release 2 and
Grid Infrastructure
Administration

Version : Demo

1.Which three actions would be helpful in determining the cause of a node reboot?

- A. determining the time of the node reboot by using the uptime command and subtracting the up time from the current system time
- B. looking for messages such as "Oracle CSSD failure. Rebooting for cluster integrity" in /var/log/messages
- C. using the crsctl command to view tracing information
- D. inspecting the ocspd log for "Begin Dump" or "End Dump" messages
- E. inspecting the database alert log for reboot messages

Answer: A,B,D

2.After Oracle Grid Infrastructure has been installed, you should take a few moments to verify the installation.

Which two actions would be useful in verifying the installation.?

- A. Run the crsctl status resource -t command to confirm that all necessary cluster resources are online.
- B. Use the operating system utilities to verify that your SCAN addresses are being properly resolved.
- C. Start Oracle Enterprise Manager and check all monitored targets.
- D. Run the cluvfy comp nodecon -n all -verbose command to verify the entire Grid Infrastructure installation.

Answer: A,B

3.Which two network addresses are required to be static, non-dhcp addresses when using the Grid Naming?

- A. GNS VIP Address
- B. SCAN VIP Address
- C. Node VIP Address
- D. Node Public Address
- E. Node Private Address

Answer: A,D

4.You are in the planning stages for upgrading your Oracle RAC database from Oracle Database 10g Release 2 to Oracle Database 11g Release 2 to run under the Oracle Grid Infrastructure.

You decide to use an administrator-managed configuration because the cluster is fairly small.

Which statement is correct about this configuration?

- A. A parent pool of the GENERIC server pool will be used.
- B. You must define a new server pool called MANUAL.
- C. A subpool of the GENERIC server pool will be used.
- D. A subpool of the FREE server pool will be used.

Answer: C

5.The Instance Initialization parameters are set to:

DB_CREATE_FILE_DEST = +DATA

DB_CREATE_ONLINE_LOG_DEST_1 = +LOGS

DB_CREATE_ONLINE_LOG_DEST_2 = + FRA

The SQL* Plus command "ALTER DATABASE ADD LOGFILE;" will create:

- A. a new log file in the +DATA disk group, or a log file in the + FRA disk group, if +DATA is not available
- B. a new log file in the +DATA disk group and a log file in the + FRA disk group
- C. a new log file in the +LOGS disk group and a log file in the + FRA disk group
- D. a new log file in the +LOGS disk group, or a log file in the +FRA dls* available
- E. a new log file in the +DATA disk group, a log file in the +LOGS disk group, and a log file in the +FRA disk group
- F. a new log file in the +LOGS disk group, or a log file in the +FRA disk group, if +LOGS is not available

Answer: C

6.Which two actions in a warehousing RAC database may cause concurrent cross-instance calls leading to I/O contention?

- A. truncate table statements
- B. select statements referring to non-partitioned tables
- C. drop table statements
- D. insert statements where each instance inserts into different partitions of a partitioned table

Answer: A,C

7.You want to create an ACFS on an ADVN volume using a shell script and the appropriate command-line utilities. These are the requirements:

- The dynamic volume file must use space in the VOLFILE disk group with a size of 500 M and be called prodvol.
- The mount point called /acfs already exists.

Which four steps must be performed to achieve this?

- A. As the Grid Infrastructure owner, run `mount -t acfs /dev/asm/prodvol-417 /acfs` to mount the file system.
- B. As the Grid Infrastructure owner, run `asmcmd volinfo -d volfile prodvol` to determine the volume information.
- C. As the Grid Infrastructure owner, run `asmcmd volcreate -d volfile -s 500M prodvol` to create the volume file.
- D. As the Grid Infrastructure owner, run `mkfs -t acfs /dev/asm/prodvol-417` to create the file system.
- E. As root, run `mount -t acfs /dev/asm/prodvol-417 /acfs` to mount the file system.
- F. As root, run `mkfs -t acfs /dev/asm/prodvol-417` to create the file system.

Answer: B,C,D,E

8.Some new non-ASM shared storage has been made available by the storage administrator, and the Oracle Grid Infrastructure administrator decides to move the voting disks, which do not reside in ASM, to this new non-ASM location. How can this be done?

- A. by running `crsctl add css votedisk <path_to_new_location>` followed by `crsctl delete css -votedisk <path_to_old_location>`
- B. by running `crsctl replace css votedisk <path_to_old_location,path_to_new_location>`
- C. by running `srvctl replace css votedisk <path_to_old_location, path_to_new_location>`
- D. by running `srvctl add css votedisk <path_to_new_location>` followed by `srvctl delete css votedisk <path_to_old_location>`

Answer: A

9.Which three statements are true about using RMAN with ASM?

- A. RMAN is the only supported method to back up database files stored in ASM.
- B. RMAN is the only supported method to back up ACFS files.
- C. RMAN can use ASM storage for backups.
- D. RMAN cannot use ASM storage for backups.
- E. Using RMAN, database files can be migrated to ASM from a file system.
- F. Using RMAN, database files cannot be moved from ASM to a file system

Answer: A,C,E

10.Which two statements are true about ACFS snapshots?

- A. They can be created for ACFS file systems only if the ASM disk group hosting the ADVM volume file used by the file system has free space available.
- B. They can be created for ACFS file systems only if the ADVM volume file used by the file system has free space available.
- C. They can be created only if the ASM disk group hosting the ADVM volume used by the file system has no other ASM files contained in the disk group.
- D. They can be created when ACFS is used both on clusters and on stand-alone servers.
- E. They are accessible only on the cluster node that was used when creating the snapshot.

Answer: B,D

11.Your cluster is subject to a service-level agreement that allows for little scheduled down time You want to use patching and maintenance methods that permit the Oracle Grid Infrastructure and Oracle RAC Databases to be available as much as possible.

Which two techniques will work some or all of the time to provide you with minimum down time?

- A. rolling upgradeable and in place patch sets
- B. idling upgradeable and out of place patch bundles
- C. rolling upgradeable and out of place patch sets
- D. rolling upgradeable and in place patch bundles
- E. rolling upgradeable and out of place one-off patches

Answer: C,D

12.Which two types of files can be stored in an ASM clustered file system?

- A. OCR and Voting Disk files
- B. data files for external tables
- C. Oracle database executable
- D. Grid Infrastructure executables
- E. data files for tablespaces F. archive log files

Answer: B,C

13.You want to reorganize the DATA diskgroup while continuing database operations. The DATA diskgroup was created using normal redundancy having one disk per failure group. The two disks used are /dev/sdal and /dev/sda2.

You plan to drop the existing disks and add the /dev/sdb1 and /dev/sdb2 disks to failure group FG_C and the /dev/sdcl and /dev/sdc2 disks to failure group FG_D.

Which procedure would you use to minimize the effect of the I/Os of this reorganization on ongoing database operations?

A. Set rebalance power to 0 for diskgroup DATA.

Add failure group FG_C with all the /dev/sdb disks.

Add failure group FG_D with all the /dev/sdc disks.

Drop disks/dev/sda1 and /dev/sda2.

Set rebalance power to 1 for diskgroup DATA.

B. Set rebalance power to 0 for diskgroup DATA.

Add failure group FG_C with all the /dev/sdb disks.

Add failure group FG_D with all the /dev/sdb disks.

Drop disks/dev/sda1 and /dev/sda2.

Set rebalance power to 9 for diskgroup DATA.

C. Set rebalance power to 9 for diskgroup DATA. Add failure group FG_C with all the /dev/sdb disks.

Add failure group FG_D with all the /dev/sdc disks.

Drop disks /dev/sda1 and /dev/sda2.

Set rebalance power to 0 for diskgroup DATA.

D. Set rebalance power to 0 for diskgroup DATA

Drop disks /dev/sda1 and /dev/sdb disks.

Add failure group FG_C with all the /dev/sdb disks.

Add failure group FG_D with all the /dev/sdc disks

Set rebalance power to 1 for diskgroup DATA.

Answer: A

14.Which two Cluster ware stack administration actions must be performed as the root user?

A. checking the health of the Clusterware on one node

B. starting the Clusterware manually on one node

C. disabling the Clusterware from automatic start at node reboot

D. checking the health and viability of the Clusterware on all nodes

E. listing the location of the voting disks

Answer: B,C

15.Which four statements are true about ADVM interoperability?

A. Using fdisk or similar disk utilities to partition ADVM-managed volumes is not supported

B. On Linux platforms, the raw utility can be used to map ADVM volume block devices to raw volume devices.

C. The creation of multipath devices over ADVM devices is not supported.

D. You may create ASMLIB devices over ADVM devices to simplify volume management.

E. ADVM does not support ASM storage contained in Exadata.

F. F. ADVM volumes cannot be used as a boot device or a rootfile system.

Answer: A,C,E,F

16.Your four-node cluster was originally purchased, installed, and configured three years ago. You recently added another four nodes to the cluster.

Now you want to remove two of the older nodes that are still accessible to be redeployed elsewhere in the

data center. Which two are true regarding the procedure for removing one or more cluster nodes?

- A. The procedure requires that all commands be invoked from one of the surviving cluster nodes.
- B. All commands are run as root regardless of which nodes are used to invoke them.
- C. The procedure requires that some commands be invoked on the node or nodes to be removed and that some be invoked from all surviving cluster nodes.
- D. The procedure requires that some commands be invoked on the node or nodes to be removed and that some be invoked from one surviving cluster node.
- E. Some commands require that the name of the node or nodes to be removed are passed as arguments, and some commands require the name of existing nodes to be passed.

Answer: D,E

17. Your production environment cluster is running Oracle Enterprise Linux and currently has four nodes. You are asked to plan for extending the cluster to six nodes. Which three methods are available to add the new nodes?

- A. silent cloning using `crsctl clone cluster` and `ssh`
- B. a GUI interface from Enterprise Manager
- C. with the Oracle Universal Installer using `runInstaller -clone <nodename>`
- D. silent cloning using `perl clone.pl -silent` either with parameters in a file or in line
- E. using `addNode.sh`

Answer: B,D,E

18. The Oracle 11g Release 2 (version 11.2.0.1) ASM instance has a new diskgroup named DATA that is currently not mounted. You log in to a cluster node as the Grid Infrastructure software owner and set the environment variables to point to the ASM instance on that node. Using SQL*Plus, you issue the commands:

```
CONNECT / AS SYSDBA
```

```
ALTER DISKGROUP DATA MOUNT;
```

The system response is:

```
ERROR at line 1:
```

```
ORA-15032: not all alterations performed
```

```
ORA-15260: permission denied on ASM disk group
```

How do you diagnose this error?

- A. Check `v$asm_operation` for operations halted by the error.
- B. Check `v$asm_attributes` to determine the `asm.compatible` setting.
- C. Check `v$asm_usergroup` to determine the owning user group for this diskgroup.
- D. Check `v$xs_session_role` to determine the role privileges of the current user.
- E. Check the connect string for the privileged role.

Answer: E

19. You enter the following command: `crsctl status resource MyApp`

You get this output: `NAME=MyApp`

```
TYPE=cluster_resource TARGET=ONLINE STATE=ONLINE on RACNODE4
```

MyApp is a policy-managed resource using a server pool with two nodes called RACNODE3 and RACNODE4 and has a cardinality of 1.

What are the meanings of the target and state status values?

- A. MyApp is currently active on RACNODE4 and is meant to be active only on RACNODE4.
- B. MyApp is meant to be active, is currently active on RACNODE4, but the Grid Infrastructure may start MyApp on RACNODE3 due to failovers.
- C. MyApp is active on RACNODE4 and was manually started.
- D. MyApp should also be online on RACNODE3 because it is a cluster_resources type that must be active on at least two nodes in the cluster, thereby overriding the CARDINALITY attribute.

Answer: B

20. After evaluating the various methods for extending a cluster, you decide to use addNode.sh.

The cluster originally consisted of four nodes: RACNODE1, RACNODE2, RACNODE3, and RACNODE4. Now two nodes called RACNODE5 and RACNODE6 have been installed and connected to the cluster by OS administrations.

Which three actions should be performed to check whether the new nodes are ready for running addNode.sh and to help correct any problems?

- A. `cluvfy stage -pre crsinst -n RACNODE5/ RACNODE6 -C + DATA -q +VOTE -orainv`
- B. `<oinstall group> -fixup -verbose`
- C. `cluvfy stage -post hwos -n RACNODE5, RACNODE6 -verbose`
- D. `cluvfy comp peer -refnode RACNODE1 -n RACNODE5, RACNODE6 -orainv <oinstall group> osdba <asmdba group> -verbose`
- E. `cluvfy stage -post hwos -n all -verbose`
- F. `cluvfy stage -pre nodeadd -n RACNODE5, RACNODE6 -fixup`
- G. `cluvfy comp peer -refnode RACNODE5 -n RACNODE6 -orainv <oinstall group> -osdba <asmdba group> -verbose`

Answer: C,D,F