

Oracle

Exam 1z0-058

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

Verson: Demo

[Total Questions: 10]

Question No : 1

Choose three correct statements about ASM disk group attributes set from the ASMCA Create Disk Group

- A. The field labeled ASM Compatibility sets the compatible.asm attribute.
- B. When creating a disk group with ASMCA, the default setting for the ASM Compatibility field (compatible.asm) is 10.1.
- C. The Database Compatibility field (compatible.rdbms) sets the minimum version level for any database instance that is allowed to mount the disk group.
- D. To use ADVM volumes, ADVM Compatibility (compatible.advm) must be set to 11.2.0 or later and ASM Compatibility must be 11.2.0 or later.
- E. Advancing the values for disk group compatibility attributes can be reversed if desired.

Answer: A,C,D

Explanation: COMPATIBLE.ASM

The value for the disk group COMPATIBLE.ASM attribute determines the minimum software version for an Oracle ASM instance that can use the disk group.

COMPATIBLE.RDBMS

The value for the disk group COMPATIBLE.RDBMS attribute determines the minimum COMPATIBLE database initialization parameter setting for any database instance that is allowed to use the disk group.

COMPATIBLE.ADVM

The value for the disk group COMPATIBLE.ADVM attribute determines whether the disk group can contain Oracle ASM volumes. The value must be set to 11.2 or higher.

Before setting this attribute, the COMPATIBLE.ASM value must be 11.2 or higher

The default disk group compatibility settings are 11.2 for Oracle ASM compatibility, 10.1 for database compatibility, and no value for Oracle ADVM compatibility. The disk group compatibility settings can only be advanced; you cannot revert to a lower compatibility setting.

Oracle Database Storage Administrator's Guide

Question No : 2

Your two-instance RAC database is running in ARCHIVELOG. The instance running on node1., fails and this failure is ejected by the instance running on node2, which indicates the recovery process for node1.

Which, process will read the redo log of the failed instance to identify the database blocks that need to be recovered during the recovery process, after the configuration of enqueue and cache part of Global resource Directory (GRD)?

- A. PMON
- B. MMON
- C. SMON
- D. GES
- E. GCS

Answer: C

Explanation:

Instance Recovery and Database Availability

The graphic illustrates the degree of database availability during each step of Oracle instance recovery:

1. Real Application Clusters is running on multiple nodes.
2. Node failure is detected.
3. The enqueue part of the GRD is reconfigured; resource management is redistributed to the surviving nodes.

This operation occurs relatively quickly.

4. The cache part of the GRD is reconfigured and SMON reads the redo log of the failed instance to identify the database blocks that it needs to recover.
5. SMON issues the GRD requests to obtain all the database blocks it needs for recovery. After the requests are complete, all other blocks are accessible.
6. The Oracle server performs roll forward recovery. Redo logs of the failed threads are applied to the database, and blocks are available right after their recovery is completed.
7. The Oracle server performs rollback recovery. Undo blocks are applied to the database for all uncommitted transactions.
8. Instance recovery is complete and all data is accessible.

Note: The dashed line represents the blocks identified in step 2 in the previous slide. Also, the dotted steps represent the ones identified in the previous slide.

D60488GC11

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 14 - 7

Identify the three forms of link aggregation that are supported by Oracle Clusterware for the interconnect.

- A. single switch active/standby configuration to increase redundancy for high availability
- B. single switch active/active configuration to increase bandwidth for performance
- C. multiswitch active/standby configuration to increase redundancy for high availability
- D. multiswitch active/active configuration to increase bandwidth for performance

Answer: A,B,C

Explanation:

Interconnect Link Aggregation: Single Switch

- ✍ Link aggregation can be used to increase redundancy for higher availability with an Active/Standby configuration.
- ✍ Link aggregation can be used to increase bandwidth for performance with an Active/Active configuration.

Interconnect Link Aggregation: Multiswitch

- ✍ Redundant switches connected with an Inter-Switch Trunk may be used for an enhanced highly available design.
- ✍ This is the best practice configuration for the interconnect.

With the single switch solutions presented in the previous slide, a failure at the switch level would bring down the entire interconnect. A better highly available (HA) design would be to implement a redundant switch strategy as illustrated in the slide, with an Inter-Switch Trunk connecting the switches. This is the best practice design for the Oracle Clusterware interconnect. Only Active/Standby mode is supported in this configuration.

D60488GC11

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 1 - 12,13,14

Question No : 4

: 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

Explanation:

✍ RMAN is also critical to Automatic Storage Management (ASM). Since ASM is critical to GRID operations, RMAN is then critical for GRID based systems. In ASM, RMAN is responsible for tracking the ASM filenames and for performing the deletion of obsolete ASM files. ASM files cannot, however, be accessed through normal operating system interfaces; therefore, RMAN is the preferred means of copying ASM files. It is also possible to use FTP through XDB, but generally, RMAN will be less complex. It is important to learn to use RMAN first because RMAN is the only method for performing backups of a database containing ASM files.

✍ Setting Initialization Parameters for ASM

The procedure for creating a duplicate database to an ASM location is similar to the procedure described in "Settings and Restrictions for OMF Initialization Parameters" on page 25-4. The difference is that you must identify the initialization parameters that control the location where files are created and set these parameters to an ASM disk group. For example, set DB_CREATE_FILE_DEST, DB_CREATE_ONLINE_DEST_n, and CONTROL_FILES to +DISK1.

✍ To take advantage of Automatic Storage Management (ASM) with an existing database you must migrate that database into ASM. This migration is performed using Recovery Manager (RMAN) even if you are not using RMAN for your primary backup and recovery strategy.

Oracle® Database Backup and Recovery Advanced User's Guide

Question No : 5

Which interconnect design is NOT supported for Oracle Clusterware?

- A. crossover cable design
- B. single switch design
- C. multiple switch design
- D. redundant interconnect design

Answer: A

Explanation:

2.7.1 Network Hardware Requirements

For the private network, the interface must support the user datagram protocol (UDP) using high-speed network adapters and switches that support TCP/IP (minimum requirement 1

Gigabit Ethernet).

Note:

UDP is the default interface protocol for Oracle RAC, and TCP is the interconnect protocol for Oracle Clusterware. You must use a switch for the interconnect. Oracle recommends that you use a dedicated switch.

Oracle does not support token-rings or crossover cables for the interconnect.

Oracle® Grid Infrastructure Installation Guide
11g Release 2 (11.2)

Question No : 6

Your network administrator informs you that the Internet service provider is being changed in a month's time in conjunction with a data center move.

You are asked to plan for the changes required in the Oracle Grid Infrastructure, which is set up to use GNS.

The IP addresses and subnets of the public network are to change.

Which two must be done in the Oracle Grid Infrastructure network setup to accommodate this change using the command-line Interfaces available?

- A. The SCAN VIPs and node VIPs must be reconfigured using `srvctl`.
- B. The SCAN VIPs and SCAN listener resources must be removed and added to obtain the new SCAN IP addresses from DHCP.
- C. The interconnect must be reconfigured by using `oifcfg`, `crsctl`, and `ifconfig`.
- D. The SCAN VIPs and node VIPs must be reconfigured by using `oifcfg`.
- E. The Interconnect must be reconfigured by using `srvctl`.

Answer: C,D

Explanation:

How to Modify Public or Private Network Information in Oracle Clusterware [ID 283684.1]

Modified 14-MAR-2012 Type HOWTO Status PUBLISHED

Applies to:

Oracle Server - Enterprise Edition - Version: 10.1.0.2 to 11.2.0.3 - Release: 10.1 to 11.2
Information in this document applies to any platform.

Goal

The purpose of this note is to describe how to change or update the cluster_interconnect and/or public interface information that is stored in OCR.

It may be necessary to change or update interface names, or subnet associated with an interface if there is a network change affecting the servers, or if the original information that was input during the installation was incorrect. It may also be the case that for some reason, the Oracle Interface Configuration Assistant ('oifcfg') did not succeed during the installation.

This note is not intended as a means to change the Public or Private Hostname themselves. Public hostname or Private hostname can only be changed by removing/adding nodes, or reinstalling Oracle Clusterware.

However, node VIP name/IP can be changed, refer to Note 276434.1 for details.

Refer to note 1386709.1 for basics of IPv4 subnet and Oracle Clusterware

Instructions for Changing Interfaces/Subnet

1. Public Network Change

If the change is only public IP address and the new ones are still in the same subnet, nothing needs to be done on clusterware level (all changes needs to be done on OS level to reflect the change).

If the change involves different subnet or interface, as there is not a 'modify' option - you will need to delete the interface and add it back with the correct information. So, in the example here, the subnet is being changed from 10.2.156.0 to 10.2.166.0 via two separate commands - first a 'delif' followed by a 'setif':

```
% $ORA_CRS_HOME/bin/oifcfg delif -global eth0
```

```
% $ORA_CRS_HOME/bin/oifcfg setif -global eth0/10.2.166.0:public syntax: oifcfg setif  
<interface-name>/<subnet>:<cluster_interconnect|public>
```

Note: If public network is changed, it may be necessary to change VIP as well, refer to Note 276434.1 for details; for 11gR2, it may be necessary to change SCAN as well, refer to note 972500.1 for details (This procedure does not apply when GNS is being used).

2. Private Network Change

2A. For pre-11gR2, if you wish to change the cluster_interconnect information and/or private IP address, hosts file needs to be modified on each node to reflect the change while the Oracle Clusterware Stack is down on all nodes. After the stack has restarted, to change the cluster_interconnect used by RDBMS and ASM instances, run oifcfg. In this example:

```
% $ORA_CRS_HOME/bin/oifcfg delif -global eth1
```

```
% $ORA_CRS_HOME/bin/oifcfg setif -global eth1/192.168.1.0:cluster_interconnect
```

2B. For 11gR2 and higher, refer to note 1073502.1

Note: For 11gR2, as clusterware also uses cluster_interconnect, intended private network must be added by "oifcfg setif" before stopping clusterware for any change.

Note: If you are running OCFS2 on Linux and are changing the private IP address for your cluster, you may also need to change the private IP address that OCFS2 is using to communicate with other nodes. For more information on this, please refer to <Note 604958.1>

3. Verify the correct interface subnet is in use by re-running oifcfg with the 'getif' option:

```
% $ORA_CRS_HOME/bin/oifcfg getif
```

```
eth0 10.2.166.0 global public
```

```
eth1 192.168.1.0 global cluster_interconnect
```

How to Modify Private Network Interface in 11.2 Grid Infrastructure [ID 1073502.1]

Modified 08-FEB-2012 Type HOWTO Status PUBLISHED

Applies to:

Oracle Server - Enterprise Edition - Version: 11.2.0.1.0 and later [Release: 11.2 and later]
Information in this document applies to any platform.

Goal

The purpose of this document is to demonstrate how to change the private network interface configuration stored in the OCR. This may be required if the name of the interface for the private network (cluster interconnect) needs to be changed at the OS level, for example, the private network is configured on a single network interface eth0, now you want to replace it with a bond interface bond0 and eth0 will be part of the bond0 interface. It also includes command for adding/deleting a private network interface.

Solution

As of 11.2 Grid Infrastructure, the CRS daemon (crsd.bin) now has a dependency on the private network configuration stored in the gnpn profile and OCR. If the private network is not available or its definition is incorrect, the CRSD process will not start and any subsequent changes to the OCR will be impossible.

Therefore care needs to be taken when making modifications to the configuration of the private network. It is important to perform the changes in the correct order.

Note: If only private network IP is going to be changed, the subnet and network interface remain same (for examples changing private IP from 192.168.0.1 to 192.168.0.10), simply shutdown GI stack, make IP modification at OS level (like /etc/hosts, network config etc) for private network, then restart GI stack will complete the task.

The following procedures apply when subnet or network interface name also requires

change.

Please take a backup of profile.xml on all cluster nodes before proceeding, as grid user:

```
$ cd $GRID_HOME/gpnp/<hostname>/profiles/peer/  
$ cp -p profile.xml profile.xml.bk
```

To modify the private network (cluster_interconnect):

1. Ensure CRS is running on ALL cluster nodes in the cluster
2. As grid user, add new interface:

Find the interface which needs to be removed. For example:

```
$ oifcfg getif  
eth1 100.17.10.0 global public  
eth0 192.168.0.0 global cluster_interconnect
```

Here the eth0 interface will be replaced by bond0 interface.

Add new interface bond0:

```
$ oifcfg setif -global <interface>/<subnet>:cluster_interconnect
```

For example:

```
$ oifcfg setif -global bond0/192.168.0.0:cluster_interconnect
```

This can be done with -global option even if the interface is not available yet, but this can not be done with -node option if the interface is not available, it will lead to node eviction.

If the interface is available on the server, subnet address can be identified by command:

```
$ oifcfg iflist
```

It lists the network interface and its subnet address. This command can be run even if CRS is not up and running. Please note, subnet address might not be in the format of x.y.z.0.

For example, it can be:

```
$ oifcfg iflist  
lan1 18.1.2.0  
lan2 10.2.3.64 << this is the private network subnet address associated with private network  
IP: 10.2.3.86
```

If the scenario is just to add a 2nd private network, for example: new interface is eth3 with subnet address:

192.168.1.96, then issue:

```
$ oifcfg setif -global eth3/192.168.1.96:cluster_interconnect
```

Verify the change:

```
$ oifcfg getif
```

3. Shutdown CRS on all nodes and disable the CRS as root user:

```
# crsctl stop crs
# crsctl disable crs
```

4. Make the network configuration change at OS level as required, ensure the new interface is available on all nodes after the change.

```
$ ifconfig -a
$ ping <private hostname>
```

5. Enable CRS and restart CRS on all nodes as root user:

```
# crsctl enable crs
# crsctl start crs
```

6. Remove the old interface:

```
$ oifcfg delif -global eth0
```

Note #1. This step is not required for adding 2nd interface scenario.

#2. If the new interface is added without removing the old interface, eg: old interface still available when CRS restart, then after step 6, CRS needs to be stop and start again to ensure the old interface is no longer in use.

The crsd.log will show:

```
2010-01-30 09:22:47.234: [ default][2926461424] CRS Daemon Starting
...
2010-01-30 09:22:47.273: [ GPnP][2926461424]clsgpnp_Init: [at clsgpnp0.c:837] GPnP client pid=7153, ti=3,
f=0
2010-01-30 09:22:47.282: [ OCRAPI][2926461424]clsu_get_private_ip_addresses: no ip addresses found.
2010-01-30 09:22:47.282: [GIPCXCPT][2926461424] gipcShutdownF: skipping shutdown, count 2, from
[ clsinet.c : 1732], ret gipcretSuccess (0)
2010-01-30 09:22:47.283: [GIPCXCPT][2926461424] gipcShutdownF: skipping shutdown, count 1, from
[ clsgpnp0.c : 1021], ret gipcretSuccess (0)
[ OCRAPI][2926461424]a_init_cls: failed to call clsu_get_private_ip_addr (7)
2010-01-30 09:22:47.285: [ OCRAPI][2926461424]a_init:13!: Clusterware init unsuccessful : [44]
2010-01-30 09:22:47.285: [ CRSOCR][2926461424] OCR context init failure. Error: PROC-44: Error in network
address and interface operations Network address and interface operations error [7]
2010-01-30 09:22:47.285: [ CRSD][2926461424][PANIC] CRSD exiting: Could not init OCR, code: 44
2010-01-30 09:22:47.285: [ CRSD][2926461424] Done.
Above errors indicate a mismatch between OS setting (oifcfg iflist) and gpnp profile setting profile.xml.
```

untitled

Workaround: restore the OS network configuration back to the original status, start CRS. Then follow above steps to make the changes again.

Please consult with Oracle Support Service if after restoring OS network configuration, CRS still could not start.

2. If any one node is down in the cluster, oifcfg command will fail with error:

```
$ oifcfg setif -global bond0/192.168.0.0:cluster_interconnect
```

PRIF-26: Error in update the profiles in the cluster

Workaround: start CRS on the node where it is not running. Ensure CRS is up on all cluster nodes.

3. If a user other than Grid Infrastructure owner issues above command, it will fail with same error: \$ oifcfg setif -global bond0/192.168.0.0:cluster_interconnect

PRIF-26: Error in update the profiles in the cluster

Workaround: ensure to login as Grid Infrastructure owner to perform such command.

4. From 11.2.0.2 onwards, if attempt to delete the last private interface

(cluster_interconnect) without adding a new one first, following error will occur:

PRIF-31: Failed to delete the specified network interface because it is the last private interface

Workaround: Add new private interface first before deleting the old private interface.

5. If CRS is down on the node, the following error is expected:

\$ oifcfg getif

PRIF-10: failed to initialize the cluster registry

Workaround: Start the CRS on the node

My Oracle Support

Question No : 7

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

Explanation: Cause

You have logged in with SYSDBA privilege that is not allowed for ASM operations

Solution

Login with SYSASM credentials for ASM operations

MOS [ID 889810.1]

Question No : 8

Your cluster has Oracle Grid Infrastructure installed and working to support RAC databases and all the related resources.

There is also a vendor HA software on the cluster providing HA services for some non-Oracle applications. You want to remove the vendor HA software and use Oracle Grid Infrastructure to provide the HA framework for your applications.

The first application to be migrated to Oracle HA is used to display financial currency information at various fixed flat screens at a business park. This application must be owned by the Financial Application owner.

Which three components of the framework are required to support this application for HA?

- A. an application VIP to support sending data to the screens
- B. a resource definition defining the application and various HA attributes
- C. privileges to permit the application to run as the correct user and to have the correct access rights to files
- D. an action program called by the HA framework for starting, stopping, and monitoring the application
- E. a resource dependency to make certain that the application VIP is started before the application is started and to start VIP automatically if it is down

Answer: B,C,D

Explanation:

Overview of Using Oracle Clusterware to Enable High Availability

To manage your applications with Oracle Clusterware:

1. Create an action script or use an existing agent.
2. Register your applications as resources with Oracle Clusterware.

If a single application requires that you register multiple resources, you may be required to define relevant dependencies between the resources.

3. Assign the appropriate privileges to the resource.
4. Start or stop your resources.

Oracle® Clusterware Administration and Deployment Guide
11g Release 2 (11.2)

Question No : 9

In an ASM instance, there are several parameters that have meanings specific to the ASM. Which parameter is mandatory?

- A. ASM_DISKGROUP
- B. ASM_DISKSTRING
- C. INSTANCE_TYPE
- D. CLUSTER_DATABASE
- E. MEMORY_TARGET
- F. ASM_POWER_LIMIT

Answer: C

Explanation: Oracle ASM Parameter Setting Recommendations

This section contains information about the following parameters for Oracle ASM:

- ✍ ASM_DISKGROUPS
- ✍ ASM_DISKSTRING
- ✍ ASM_POWER_LIMIT
- ✍ ASM_PREFERRED_READ_FAILURE_GROUPS
- ✍ DB_CACHE_SIZE
- ✍ DIAGNOSTIC_DEST
- ✍ INSTANCE_TYPE
- ✍ uk.co.certification.simulator.questionpool.PList@347f2a80
- ✍ LARGE_POOL_SIZE
- ✍ PROCESSES

✍ REMOTE_LOGIN_PASSWORDFILE

✍ SHARED_POOL_SIZE

Oracle® Automatic Storage Management Administrator's Guide

Question No : 10

Which two statements are true regarding the Active Session History (ASH) reports for RAC?

- A. They provide details about Oracle databases for all current sessions, and history of past session all RAC nodes.
- B. They provide statistics about Oracle databases for the active sessions on all the RAC nodes.
- C. They report on data captured for active sessions. The volume of data is directly related to the work being performed by sessions.
- D. They report on data captured for active sessions. The volume of data is directly related to the number of sessions on the system.

Answer: A,C

Explanation: ASH report statistics provide details about Oracle Database session activity. Oracle Database records information about active sessions for all active Oracle RAC instances and stores this data in the System Global Area (SGA). Any session that is connected to the database and using CPU is considered an active session. The exception to this is sessions that are waiting for an event that belongs to the idle wait class. ASH reports present a manageable set of data by capturing only information about active sessions. The amount of the data is directly related to the work being performed, rather than the number of sessions allowed on the system.

Oracle Real Application Clusters Administration and Deployment Guide