

[Home](#)

[Membership](#)

[Subscribe](#)

[Home](#)

[CheckPoint](#)

[Cisco](#)

[CIW](#)

[CompTIA](#)

[CWNP](#)

[EC-Council](#)

[EMC](#)

[EXIN](#)

[HP](#)

[IBM](#)

[ISACA](#)

[ISC2](#)

[Isilon](#)

[Juniper](#)

[Linux](#)

[Microsoft](#)

[Oracle](#)

[PMI](#)

[Sun](#)

[VMware](#)

## EXAM INFORMATION

[RSS feed for this section](#)

This category contains 139 posts

*Exam 1Z0-058 : Oracle Real Application Clusters 11g Release 2 and Grid Infrastructure Administration (update 11 March, 2013)*

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

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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 ocssd log for "Begin Dump" or "End Dump" messages
- E.  
inspecting the database alert log for reboot messages

#### Explanation:

Determining Which Process Caused Reboot

Log File Locations for Processes Causing Reboots.

- oclskd
  - <Grid\_Home>/log/<hostname>/client/oclskd.log
- ocssd
  - /var/log/messages
  - <Grid\_Home>/log/<hostname>/cssd/ocssd.log
- cssdagent
  - <Grid\_Home>/log/<hostname>/agent/ohasd/oracssdagent\_root
- cssdmonitor
  - <Grid\_Home>/log/<hostname>/agent/ohasd/oracssdmonitor\_root
- hangcheck-timer
  - /var/log/messages

First, determine the time of the node reboot by using the uptime command and subtracting the up time from the current system time. The reboot time will be used when examining log files.

When the OCSSD daemon is responsible for rebooting a node, a message similar to Oracle CSSD failure.

Rebooting for cluster integrity is written into the system messages log at /var/log/messages.

The cssd daemon log file that is located at <Grid\_Home>/log/<hostname>/cssd/ocssd.log may also contain messages similar to “Begin Dump” or “End Dump” just before the reboot.

If hangcheck-timer is being used, it will provide message logging to the system messages log when a node restart is initiated by the module. To verify whether this process was responsible for the node reboot, examine the /var/log/messages file and look for an error message similar to: “Hangcheck: hangcheck is restarting the machine.”

Other useful log files include the Clusterware alert log in <Grid\_home>/log/<hostname> and the lastgasp log in /etc/oracle/lastgasp or /var/opt/oracle/lastgasp.

If no indication of which process caused the reboot can be determined from these files, additional debugging and tracing may need to be enabled.

Note: The oclsonmon and the oproc background processes have been eliminated in Oracle Database 11g Release 2.

D60488GC11

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 6 – 21

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

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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 -v verbose command to verify the entire Grid Infrastructure installation.

#### Explanation:

Verifying the Grid Infrastructure Installation

Execute the crsctl command as shown in the slide to confirm that all cluster resources are up and running.

```
root@racnode01 ~]# /u01/app/11.2.0/grid/bin/crsctl stat res -t
```

In addition, you should confirm that your DNS is properly forwarding address requests for your application and SCAN VIPs to your GNS and that they are resolved properly. You can do this with dig. Execute the dig command with DNS and VIP addresses as shown:

```
# dig @myDNS.example.com racnode01-cluster01.example.com
```

```
:: QUESTION SECTION:
```

```
;racnode01-vip.cluster01.example.com. IN A
```

```
:: ANSWER SECTION:
```

```
racnode01-vip.cluster01.example.com. 120 IN A
```

```
192.0.2.103
```

```
# dig @myDNS.example.com cluster01-scan.cluster01.example.com
```

```
:: QUESTION SECTION:
```

```
;cluster01-scan.cluster01.example.com. IN A
```

```
:: ANSWER SECTION:
```

```
cluster01-scan.cluster01.example.com. 120 IN A 192.0.2.248
```

```
cluster01-scan.cluster01.example.com. 120 IN A 192.0.2.253
```

```
cluster01-scan.cluster01.example.com. 120 IN A 192.0.2.254
```

D60488GC11

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 2 44

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

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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

Explanation:

2.6.2 IP Address Requirements

Before starting the installation, you must have at least two network adapters configured on each node: One for the private IP address and one for the public IP address.

You can configure IP addresses with one of the following options:

Dynamic IP address assignment using Oracle Grid Naming Service (GNS). If you select this option, then network administrators assign static IP address for the physical host name and dynamically allocated IPs for the Oracle Clusterware managed VIP addresses.

Oracle Grid Infrastructure Installation Guide

Implementing GNS

To implement GNS, you must collaborate with your network administrator to obtain an IP address on the public network for the GNS VIP. DNS uses the GNS VIP to forward requests for access to

the cluster to GNS. You must also collaborate with your DNS administrator to delegate a domain to the cluster. This can be a separate domain or a subdomain of an existing domain. The DNS server must be configured to forward all requests for this new domain to the GNS VIP. Since each cluster has its own GNS, it must be allocated a unique domain of which to be in control.

Oracle Clusterware Administration and Deployment Guide

**Which statement is correct about this configuration?**

---

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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.

Explanation:

With Oracle RAC 11g Release 2, there is only a database resource define to Oracle Clusterware. This resource will contain the Oracle Home, the spfile, the server pool(s) and the ASM diskgroup(s) required for the database. The database resource will have a weak start dependency on the VIP which means it will try to start the VIP for the node when the instance starts however if the VIP does not start successfully, the instance will still be started. When reviewing the database resource for an administrator Managed database, you will see a server pool defined with the same name as the Oracle database. This pool will be part of a special Oracle defined server pool called Generic. The Generic server pool is managed by Oracle to support Administrator Managed databases. The server pools that are part of Generic will be automatically created or removed when you add or remove an Administrator Managed database.

Oracle Database 11gR2 Upgrade Companion (Version 2.70)

**The Instance Initialization parameters are set to:**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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

Explanation:

The following table lists the initialization parameters that enable the use of Oracle Managed Files.

Initialization Parameter

DB\_CREATE\_FILE\_DEST

Defines the location of the default file system directory or Oracle ASM disk group where the database creates datafiles or tempfiles when no file specification is given in the create operation.

Also used as the default location for redo log and control files if

DB\_CREATE\_ONLINE\_LOG\_DEST\_n are not specified.

DB\_CREATE\_ONLINE\_LOG\_DEST\_n

Defines the location of the default file system directory or Oracle ASM disk group for redo log files and control file creation when no file specification is given in the create operation. By changing n, you can use this initialization parameter multiple times, where n specifies a multiplexed copy of the redo log or control file. You can specify up to five multiplexed copies.

DB\_RECOVERY\_FILE\_DEST

Defines the location of the Fast Recovery Area, which is the default file system directory or Oracle ASM disk group where the database creates RMAN backups when no format option is used, archived logs when no other local destination is configured, and flashback logs. Also used as the default location for redo log and control files or multiplexed copies of redo log and control files if DB\_CREATE\_ONLINE\_LOG\_DEST\_n are not specified.

Oracle Database Administrator's Guide

**Which three statements define a cluster?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Which three statements define a cluster?

- A.  
**is a group of independent, but interconnected computers that act as a single system**
- B.  
**can be deployed to increase availability and performance**
- C.  
**can be deployed to balance a dynamically changing workload**

D.  
should appear to an application as multiple servers

Explanation:

Oracle Real Application Clusters  
Computer cluster

**Which four steps must be performed to achieve this?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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

1. The dynamic volume file must use space in the VOLFILE disk group with a size of 500 M and be called prodvol.
2. 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.

Explanation:

Creating an ACFS Volume

Create the volume:

```
$ asmcmd volcreate -G DATA -s 100M testvol
```

View the volume information:

```
$ asmcmd volinfo -G DATA testvol
```

Make a mount point directory:

```
$ mkdir /u01/app/oracle/acfsdata/testvol
```

Make the file system (as root):

```
# mkfs -t acfs /dev/asm/testvol-403
```

Mount the file system to the mount point:

```
# mount -t acfs /dev/asm/testvol-403 \
```

```
/u01/app/oracle/acfsdata/testvol
```

```
D60488GC11
```

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 10 – 13

Create a file system with the Oracle ACFS `mkfs` command.

Create a file system using an existing volume device.

For example:

```
$/sbin/mkfs -t acfs /dev/asm/volume1-123
```

```
mkfs.acfs: version = 11.2.0.1.0.0
```

```
mkfs.acfs: on-disk version = 39.0
```

```
mkfs.acfs: volume = /dev/asm/volume1-123
```

```
mkfs.acfs: volume size = 10737418240
```

```
mkfs.acfs: Format complete.
```

See “`mkfs`” (Linux or UNIX) or “`acfsformat`” (Windows). The root privilege is not required. The ownership of the volume device file dictates who can run this command.

Oracle Automatic Storage Management Administrator’s Guide

11g Release 2 (11.2)

**How can this be done?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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>`

Explanation:

Adding, Deleting, or Migrating Voting Disks

Modifying voting disks that are stored in Oracle ASM

To migrate voting disks from Oracle ASM to an alternative storage device, specify the path to the non-Oracle

ASM storage device with which you want to replace the Oracle ASM disk group using the following command:

```
$ crsctl replace votedisk path_to_voting_disk
```

You can run this command on any node in the cluster.

To replace all voting disks not stored in Oracle ASM with voting disks managed by Oracle ASM in an Oracle

ASM disk group, run the following command:

```
$ crsctl replace votedisk +asm_disk_group
```

Modifying voting disks that are not stored on Oracle ASM:

To add one or more voting disks, run the following command, replacing the `path_to_voting_disk` variable with one or more space-delimited, complete paths to the voting disks you want to add:

```
$ crsctl add css votedisk path_to_voting_disk [...]
```

To replace voting disk A with voting disk B, you must add voting disk B, and then delete voting disk A. To add a new disk and remove the existing disk, run the following command, replacing the `path_to_voting_diskB` variable with the fully qualified path name of voting disk B:

```
$ crsctl add css votedisk path_to_voting_diskB -purge
```

The `-purge` option deletes existing voting disks.

To remove a voting disk, run the following command, specifying one or more space-delimited, voting disk FUIDs or comma-delimited directory paths to the voting disks you want to remove:

```
$ crsctl delete css votedisk {FUID | path_to_voting_disk[...]}
```

Oracle Clusterware Administration and Deployment Guide

11g Release 2 (11.2)

Which three statements are true about using RMAN with ASM?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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

Explanation:

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. Oracle Database Backup and Recovery Advanced User’s Guide

**Which two statements are true about ACFS snapshots?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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.

Explanation:

About Oracle ACFS Snapshots  
Oracle ACFS snapshot storage is maintained within the file system, eliminating the management of separate storage pools for file systems and snapshots. Oracle ACFS file systems can be dynamically resized to accommodate additional file and snapshot storage requirements.  
Oracle Automatic Storage Management Administrator’s Guide  
11g Release 2 (11.2)

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

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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

Explanation:

Applications remain available while software is upgraded in the new ORACLE\_HOME. The configuration inside the ORACLE\_HOME is retained because the cloning procedure involves physically copying the software

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

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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

Explanation:

Oracle Automatic Storage Management Cluster File System (Oracle ACFS) is a multi-platform, scalable file system, and storage management technology that extends Oracle Automatic Storage Management (Oracle ASM) functionality to support customer files maintained outside of Oracle Database. Oracle ACFS supports many database and application files, including executables, database trace files, database alert logs, application reports, BFILEs, and configuration files. Other supported files are video, audio, text, images, engineering drawings, and other general-purpose application file data.

Notes:

Not supported means Oracle Support Services does not take calls and development does not fix bugs associated with storing unsupported file types in Oracle ACFS.

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

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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/sdc1 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/sdal 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.

Explanation:

Oracle Automatic Storage Management



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

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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

Explanation:

Controlling Oracle Clusterware

The crsctl utility can be used to control Oracle Clusterware.

To start or stop Oracle Clusterware on a specific node:

```
# crsctl start crs
```

```
# crsctl stop crs
```

To enable or disable Oracle Clusterware on a specific node:

```
# crsctl enable crs
```

```
# crsctl disable crs
```

```
D60488GC11
```

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 3 – 5



**Which four statements are true about ADVM interoperability?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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. **ADVM volumes cannot be used as a boot device or a root file system.**

Explanation:

Oracle Automatic Storage Management Cluster File System (Oracle ACFS) and Oracle ASM Dynamic Volume Manager (Oracle ADVM) extend Oracle ASM support to include database and application executables, database trace files, database alert logs, application reports, BFILEs, and configuration files. Other supported files are video, audio, text, images, engineering drawings, and other general-purpose application file data.

Because of the fact that Oracle ADVM Volumes are technically spoken ASM files located on ASM Disk groups, and the fact that the Dynamic Volumes do not use the traditional device partitioning, it enables Oracle to extend some of the ASM features to the ASM Clustered File Systems, which are created inside these ADVM Volumes, such as dynamic resizing or dynamically adding volumes. This makes ADVM and ACFS a far more flexible solution than traditional physical

devices.

Important Notes:

### Which two are true regarding the procedure for removing one or more cluster nodes?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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.

#### Explanation:

```
# crsctl delete node -n node_to_be_deleted
```

```
$/runInstaller -updateNodeList ORACLE_HOME=Grid_home "CLUSTER_NODES={node_to_be_deleted}" CRS=TRUE -silent -local
```

On any node other than the node you are deleting, run the following command from the Grid\_home /oui/bin directory where remaining\_nodes\_list is a comma-delimited list of the nodes that are going to remain part of your cluster:

```
$/runInstaller -updateNodeList ORACLE_HOME=Grid_home "CLUSTER_NODES={remaining_nodes_list}" CRS=TRUE -silent
```

```
$ cluvfy stage -post nodedel -n node_list [-verbose]
```

Oracle Clusterware Administration and Deployment Guide

### Which three methods are available to add the new nodes?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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

### How do you diagnose this error?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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 vasm\_operation for operations halted by the error.

B.

Check vasm\_attributes to determine the asm.compatible setting.

C.

Check vasm\_usergroup to determine the owning user group for this diskgroup.

D.

Check vxs\_session\_role to determine the role privileges of the current user.

E.

**Check the connect string for the privileged role.**

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]

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

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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.

Explanation:

Resource Attributes

NAME

A case-sensitive alphanumeric string that names the resource. Oracle recommends a naming convention that starts with an alphanumeric prefix, such as myApache, and complete the name with an identifier to describe it.

A resource name can contain any platform-supported characters except the exclamation point (!) and the tilde

(~). A resource name cannot begin with a period (.) nor with the string ora.

TYPE

The type of resource indicated when you create a resource. This attribute is required when creating a resource.

Local resource: Instances of local resource type name is local\_resourcerun on each server of

the cluster. When a server joins the cluster, Oracle Clusterware automatically extends local resources to have instances tied to the new server. When a server leaves the cluster, Oracle Clusterware automatically sheds the instances of local resources that ran on the departing server. Instances of local resources are pinned to their servers; they do not fail over from one server to another.

Cluster resource: Cluster-aware resource type name is cluster\_resource aware of the cluster environment and are subject to cardinality and cross-server switchover and failover.

TARGET

An internal, read-only attribute that describes the desired state of a resource. Using the crsctl start resource\_name or crsctl stop resource\_name commands, however, can affect the value of this attribute.

STATE

An internally-managed attribute that reflects the current state of the resource as reported by Oracle Clusterware. The state of a resource can be one of the following:

ONLINE: The resource is online and resource monitoring is enabled (see CHECK\_INTERVAL).

OFFLINE: The resource is offline and only offline resource monitoring is enabled, if configured (see OFFLINE\_CHECK\_INTERVAL).

INTERMEDIATE: The resource is either partially online or was known to be online before and subsequent attempts to determine its state have failed; resource monitoring is enabled (see CHECK\_INTERVAL).

UNKNOWN: The resource is unmanageable and its current state is unknown; manual intervention is required to resume its operation. A resource in this state is not monitored.

Oracle Clusterware Administration and Deployment Guide

11g Release 2 (11.2)

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

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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 administrators.

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

#### Explanation:

```
$ cluvfy stage -post nodeadd -n node3 [-verbose]
```

Use the cluvfy comp peer component verification command to check the compatibility and properties of the specified nodes against a reference node. You can check compatibility for non-default user group names and for different releases of the Oracle software. This command

compares physical attributes, such as memory and swap space, as well as user and group values, kernel settings, and installed operating system packages.

Syntax

```
cluvfy comp peer -n node_list [-refnode node] [-r {10gR1 | 10gR2 | 11gR1 | 11gR2}] [-orainv orainventory_group] [-osdba osdba_group] [-verbose]
```

Usage Notes

Peer comparison with the -refnode option compares the system properties of other nodes against

the reference node. If the value does not match (the value is not equal to reference node value), then CVU flags that comparison as a deviation from the reference node. If a group or user does not exist on reference node as well as on the other node, CVU reports this comparison as 'passed' because there is no deviation from the reference node. Similarly, CVU reports as 'failed' a comparison with a node that has more total memory than the reference node.

```
$ cluvfy stage -pre nodeadd -n node3 [-fixup [-fixupdir fixup_dir]] [-verbose]
```

You can specify the -fixup option and a directory into which CVU prints instructions to fix the cluster or node if the verification fails.

Oracle Clusterware Administration and Deployment Guide

### What are the default connect strings used by SQL\*Plus and ASMCMD when connecting to ASM instances?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

What are the default connect strings used by SQL\*Plus and ASMCMD when connecting to ASM instances?

A.  
operating system authentication for ASMCMD;  
none for SQL\*Plus

B.  
“/ AS SYSASM” for ASMCMD;  
“/” for SQL\*Plus

C.  
“SYS/ AS SYSASM” for ASMCMD;  
“/ AS SYSDBA” for SQL\*Plus

D.  
operating system authentication for ASMCMD;  
operating system authentication for SQL\*Plus

E.  
“/ AS SYSASM” for ASMCMD;  
operating system authentication for SQL\*Plus

#### Explanation:

On Linux and UNIX systems, dba is the default operating system group designated as OSASM, OSOPER, and OSDBA for Oracle ASM. On Windows systems, ora\_dba is the default name designated as OSASM, OSOPER, and OSDBA.

SQL\*Plus commands, ASMCMD commands, and ASMCA use operating system authentication. To connect locally as SYSASM to an Oracle ASM instance using operating system authentication with SQL\*Plus, use the following statement:

```
sqlplus / AS SYSASM
```

OSASM group

This group is granted the SYSASM privilege, which provides full administrative privileges for the Oracle ASM instance.

You can specify the -privilege option to choose the type of connection, either SYSASM or SYSDBA. The default value is SYSASM and is used when administering an Oracle ASM instance. Connect as SYSDBA when connecting to the database instance.

For example:

```
$ asmcmd -privilege sysasm
```

```
usage: asmcmd [-V] [-v <errors|warnings|normal|info|debug>] [--privilege <sysasm|sysdba>] [-p] [command]
```

### Which two statements are true about the procedure for minimizing down time?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

You want to install a database patch on your eight-node cluster by using Opatch with the minimum amount of down time to your cluster nodes.

You have already been prompted for the first set of nodes and you replied with node names RACNODE1 RACNODE2, and RACNODE3.

Which two statements are true about the procedure for minimizing down time?

A.  
The patch must be propagated to, applied, and the inventory updated on all the remaining nodes before restarting the instances on the first set of nodes.

B.

When the first set of nodes has been patched, the instances on the remaining nodes are shut down automatically by opatch.

C.

The instances must be started on the first set of nodes and then stopped on the remaining set to make certain that some nodes are always available.

D.

When the first set of nodes has been patched, the administrator is prompted to shut down the instances on the remaining nodes.

E.

The patch must be propagated to and the inventory updated on all the remaining nodes before restarting the instances on the first set of nodes.

Explanation:

Installing a Patch with Minimum Down Time with OPatch

In minimum down-time patching, the nodes are divided into two sets. One set of nodes is shut down and the patch is applied to those nodes. After the first set of nodes has been patched, the second set of nodes is shut down. The first set of nodes is then restarted and the patch is applied to the second set of nodes. After the patch has been applied to the second set of nodes, those

nodes are restarted. This method leads to less down time for Oracle RAC, compared to having all the nodes shut down at the same time. When you use the minimum down-time patching method, the following actions occur:

The local node is always patched first.

The local node is used as a base to patch the other nodes.

The user is prompted for the first set of nodes to patch.

For each node in this first set, the user is asked to stop the instance and then the patch is propagated and applied to that node before continuing to the next node.

When the first set of nodes has been patched, the user is asked to shut down Clusterware on the remaining nodes.

The instances are stopped on the last set of remote nodes.

The patch is propagated to the last set of nodes and the inventory is updated.

You can then start up the patched nodes (the first set of nodes) before patching the remaining nodes.

D60488GC11

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 4 44

What is the correct way to silently add the nodes?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

You are ready to add two new nodes called RACNODE5 and RACNODE6 to your existing four-node cluster using addNode.sh

You have run cluvfy -peer to check the new nodes against a reference node.

When you originally created the cluster, the network administrators chose to statically define the scan vip addresses in the corporate DNS server, and you installed the Oracle Grid Infrastructure without using GNS.

What is the correct way to silently add the nodes? r

A.

addNode . sh -silent "CLUSTER\_NEW\_nodes={ RACNODE5, RACNODE6 } > "

B.

addNode . sh -silent "CLUSTER\_NEW\_VIRTUAL\_HOSTNAMES=<(RACNODE5 -VI P, RACNODE6-VIP) < "

C.

addNode. sh -silent "CLUSTER\_NEW\_NODES= { RACNODE5, RACNODE6 } " "CLUSTER\_NEW\_VIRTUAL\_HOSTNAMES= RACNODE5-VIP, RACNODE6-VIP > "

D.

addNode.sh -silent -responseFile mynewnodea . txt With the response file containing only

E.

CLUSTER\_NEW\_NODES= { RACNODE5, RACNODE6 }

F.

addNode.sh -silent -responseFile mynewnodes . txt With the response file containing only CLUSTER NEW VIRTUAL HOSTNAMES= { " RACNODE3-VI P , RACNODE4 -VI P >

Explanation:

If you are using Grid Naming Service (GNS):

```
$ ./addNode.sh -silent "CLUSTER_NEW_NODES={node2}"
```

If you are not using GNS:

```
$ ./addNode.sh -silent "CLUSTER_NEW_NODES={node2}" "CLUSTER_NEW_VIRTUAL_HOSTNAMES={node2-vip}"
```

Alternatively, you can specify the entries in a response file, where file\_name is the name of the file, and run the addNode.sh script, as follows:

```
$ addNode.sh -silent -responseFile file_name
```

When the addNode.sh script completes, a message window displays a list of nodes in the cluster and root scripts that must be run on those nodes.

Oracle Clusterware Administration and Deployment Guide

### Which four statements about mounting ASM cluster file systems are true?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Which four statements about mounting ASM cluster file systems are true?

A.

An ACFS volume can be mounted by using ASMCA.

B.

The standard Linux/UNIX mount command can be used to mount an ACFS volume, provided the ACFS type is specified; (mount -t acfs).

C.

ACFS volumes can be mounted by using the ASMCMD utility.

D.

The acfs mountvol command can be used to mount ACFS volumes on Windows platforms.

E.

Oracle Enterprise Manager can be used to mount ACFS volumes.

#### Explanation:

There are buttons for Create, Show Mount All, and Show Dismount All commands

root privilege is required to run mount.

Windows Administrator privileges are required to mount an Oracle ACFS

On this page, you can choose to mount, dismount, delete, create snapshot, view content, register, and deregister a selected file system. In addition, you can create a file system, mount all file systems, or dismount all file systems.

Oracle Automatic Storage Management Administrator's Guide

### What will be the volume extent and Volume Allocation Unit (VAU) if the stripe column is 4 and AU is 1 MB? Which are the default values?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

You are allocating space from the ASM disk group for an ADVN volume. What will be the volume extent and Volume Allocation Unit (VAU) if the stripe column is 4 and AU is 1 MB? Which are the default values?

A.

The volume extent is 64 MB and the VAU will be 512 MB.

B.

The volume extent is 64 MB and the VAU will be 256 MB.

C.

The volume extent is 32 MB and the VAU will be 256 MB.

D.

It is not possible to calculate these values with the given information.

#### Explanation:

Stripe column and stripe width are two important attributes that can be specified for a volume determining how space is allocated for a volume and how space is allocated within a volume after ACFS or third party file system is created on a volume and a file is created or extended on the file system built on that volume. Both attributes are specified at volume creation time and cannot be changed later. If there is no value a default value is used.

Volume Allocation Unit (VAU) is the smallest allocation for a volume. Whenever a volume is created or extended space is allocated in multiple of VAU. VAU size is determined by the Volume Extent (VE) and the stripe column. VAU is the product of VE and the stripe column.

Volume Extent is based on the Allocation Unit (AU) specified on a disk group and for AU with size of 1MB the VE is 64MB. Whenever a VAU is allocated VE are allocated in a round robin fashion

across the disks in a disk group. Volume size is always multiple of the VAU. For example if a volume is requested with size 200MB and stripe column is 4 and AU is 1MB one VAU with size 256MB will be allocated and the volume size will be 256MB

**Which two start dependencies would you use for the WEBAPP application resource definition?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

You are creating a resource definition called WEBAPP for a web-based application that requires an application called WEBVIP.

Both WEBVIP and WEBAPP are normally active on a single node, but may fail over if required to another cluster node in your eight-node cluster.

A second VIP application resource called appsvip exists, used by a second application resource called secapp. webvip and webapp must start in the correct sequence.

Which two start dependencies would you use for the WEBAPP application resource definition?

- A. hard(SECAPP)
- B. pullup(SECAPP)
- C. pullup(WEBVIP)
- D. hard(WEBVIP)
- E. dispersion(WEBVIP)
- F. weak (WEBVIP)

Explanation:

Start Dependencies

hard

Define a hard start dependency for a resource if another resource must be running before the dependent resource can start. For example, if resource A has a hard start dependency on resource B, then resource B must be running before resource A can start.

Note:

Oracle recommends that resources with hard start dependencies also have pullup start dependencies. weak

If resource A has a weak start dependency on resource B, then an attempt to start resource A attempts to start resource B, if resource B is not running. The result of the attempt to start resource B is, however, of no consequence to the result of starting resource A.

attraction

If resource A has an attraction dependency on resource B, then Oracle Clusterware prefers to place resource A on servers hosting resource B. Dependent resources, such as resource A in this case, are more likely to run on servers on which resources to which they have attraction dependencies are running. Oracle Clusterware places dependent resources on servers with resources to which they are attracted.

pullup

Use the pullup start dependency if resource A must automatically start whenever resource B starts. This dependency only affects resource A if it is not running. As is the case for other dependencies, pullup may cause the dependent resource to start on any server. Use the pullup dependency whenever there is a hard stop dependency, so that if resource A depends on resource B and resource B fails and then recovers, then resource A is restarted.

dispersion

If you specify the dispersion start dependency for a resource, then Oracle Clusterware starts this resource on a server that has the fewest number of resources to which this resource has dispersion. Resources with dispersion may still end up running on the same server if there are not enough servers to disperse them to.

Oracle Clusterware Administration and Deployment Guide

11g Release 2 (11.2)

**how many failure groups are automatically created?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

The database administrator is tasked with creating an ASM disk group. Exadata is not being used.



If failure groups are not specified when creating an ASM disk group containing 10 disks, how many failure groups are automatically created?

- A.  
one
- B.  
two
- C.  
five
- D.  
ten

Explanation:

There are always failure groups even if they are not explicitly created. If you do not specify a failure group for a disk, then Oracle automatically creates a new failure group containing just that disk, except for disk groups containing disks on Oracle Exadata cells.

Oracle Automatic Storage Management Administrator's Guide

What are the two reasons for doing this

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

The Oracle Grid Infrastructure administrator runs ocrconfig -backuploc <new\_location> after completing the installation of the Grid Infrastructure. What are the two reasons for doing this

A.  
The master node may change if the node running the master CRS daemon shuts down causing another node to become the master. Because the default location is a nonshared storage, and only the master node performs backups, the backups could be spread around on many nodes making management or recovery more difficult.

B.  
The cluster node currently acting as the master node for CRS is the only node that backs up the OCR to the default location, and by changing the location, all cluster nodes will then take backups.

C.  
This is done to move the backup location into an ASM Cluster File System directory so that the backups benefit from ASM striping and mirroring.

D.  
If for any reason CRS must be stopped on all nodes for a time spanning a scheduled backup, then on restart, the backup timer will be reset. This could result in longer time duration between automatic backups than the standard four-hour interval.

Explanation:

Changing the Automatic OCR Backup Location

The automatic backup location should be changed to a location shared by all nodes.

```
# ocrconfig backuploc <path to shared CFS or NFS>
```

The backup location will be used for both automatic and manual backups.

It is recommended that these files be included in routine scheduled backups to an offline location.

If CRS has been stopped on all nodes, the schedule of backups is suspended.

On restart, a backup is not immediately taken and the backup timer is reset.

D60488GC11

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 3 – 12

Which two commands will create a 500 MB volume called TESTVOL in the ACFSDBG ACFS diskgroup by storing only one copy of the volume file extents in the diskgroup

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

From a command line, you can create an ASM volume with the same properties by using either SQL or ASMCMD. Which two commands will create a 500 MB volume called TESTVOL in the ACFSDBG ACFS diskgroup by storing only one copy of the volume file extents in the diskgroup

- A.  
`SQL> alter diskgroup ACFSDBG add volume TESTVOL size 500M unprotected;`
- B.  
`ASMCMD> volcreate -G ASMFS -s 500M`

C.  
SQL> alter diskgroup ACFS add volume TESTVOL size 500M;

D.  
ASMCMD> volcreate -G ACFS DG -s 500M -redundancy unprotected TESTVOL

E.  
SQL> create volume TESTVOL diskgroup ACFS DG size 500M unprotected;

Explanation:

The ALTER DISKGROUP VOLUMESQL statements enable you to manage Oracle ADVM volumes, including the functionality to add, modify, resize, disable, enable, and drop volumes.

You can create an Oracle ASM Dynamic Volume Manager (Oracle ADVM) volume in a disk group. The volume device associated with the dynamic volume can then be used to host an Oracle ACFS file system.

If the volume is hosting an Oracle ACFS file system, then you cannot resize that volume with the SQL ALTER DISKGROUP statement. Instead you must use the acfsutil size command.

volcreate  
Creates an Oracle ADVM volume in the specified disk group.

Syntax and Description

```
volcreate -G diskgroup -s size [ --column number ] [ --width stripe_width ]  
[ --redundancy {high|mirror|unprotected} ]
```

Redundancy of the Oracle ADVM volume which can be specified for normal redundancy disk groups. The range of values is as follows: unprotected for non-mirrored redundancy, mirror for double-mirrored redundancy, or high for triple-mirrored redundancy. If redundancy is not specified, the setting defaults to the redundancy level of the disk group.

**Which procedure will do this properly if run as the root user?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

You installed the Oracle Grid Infrastructure on a four-node cluster before discussing the network requirements with the network administrator who was on holiday.

You created a single scan named mydb-scan.myclust.example.com by adding this name to the /etc/hosts file. As a result, the Grid Infrastructure has four node listeners and node VIP but only a single SCAN listener and SCAN VIP.

The network administrator has returned and modified the corporate DNS server to associate three IP addresses with the mydb-scan.myclust.example.com scan name. The SCAN VIPs are on the same network as the node VIPs.

You now must replace the single SCAN VIP and listener with three of each for high availability purposes and make certain that the SCANS and listeners are active. Which procedure will do this properly if run as the root user?

A.  
srvctl stop scan\_listener  
srvctl stop scan  
srvctl start scan  
srvctl start scan\_listener

B.  
srvctl stop scan\_listener  
srvctl stop scan  
srvctl remove scan  
srvctl add scan -n MYDB-SCAN.MYCLUST.EXAMPLE.COM  
srvctl start scan  
srvctl start scan\_listener

C.  
srvctl add scan -n MYDB-SCAN.MYCLUST.EXAMPLE.COM  
srvctl start scan  
srvctl start scan\_listener

D.  
srvctl stop scan\_listener  
srvctl stop scan  
srvctl remove scan  
srvctl add scan  
srvctl start scan  
srvctl start scan\_listener

Explanation:

How to update the IP address of the SCAN VIP resources (ora.scan.vip) [ID 952903.1]

Modified 03-JAN-2012 Type HOWTO Status PUBLISHED

In this Document

Goal

Solution

Applies to:

Oracle Server – Enterprise Edition – Version.2.0.1 to 11.2.0.1 – Release.2 to 11.2

Information in this document applies to any platform.

Goal

The purpose of this document is to explain how to change the IP addresses associated with the SCAN VIPs in a 11gR2 Grid (CRS) environment.

The IP addresses associated with the SCAN VIP resources are initially set when the SCAN resources are created.

Any changes to the DNS entry for the SCAN are not automatically propagated to the clusterware and need to be done manually.

This applies only to installations that are not using GNS.

The information in this note can also be helpful in cases where SCAN was originally configured with just one address and is now being expanded to accommodate three IP addresses.

Solution

Before the SCAN VIPs can be changed, the entry for the SCAN name on the Domain Name Server (DNS) needs to be updated with the new IP addresses. This usually will be done by a network administrator. To check the current setting, the following command can be used:

```
nslookup <scan_name>
```

To check the current IP address(es) of the SCAN VIPs, run the following commands as the root user:

```
$GRID_HOME/bin/srvctl config scan
```

Next refresh the SCAN VIPs with the new IP addresses from the DNS entry:

```
$GRID_HOME/bin/srvctl modify scan -n <scan_name>
```

To check if the SCAN VIPs have been changed, run the following command, it should now show the new IP addresses.

```
$GRID_HOME/bin/srvctl config scan
```

Below is an example using the following configuration:

The name of the SCAN is sales-scan.example.com

subnet of the public network is 10.100.10.0

netmask for the public network is 255.255.255.0

name of the public interface is eth1

old IP addresses.100.10.81, 10.100.10.82 & 10.100.10.83

new IP addresses.100.10.121, 10.100.10.122 & 10.100.10.123

A lookup of the SCAN on the DNS server shows that the entry has already been updated with the new IP addresses:

```
$ nslookup sales-scan.example.com
```

```
Server: dns1.example.com
```

```
Address: 100.10.70#53
```

```
Name: sales-scan.example.com
```

```
Address: 100.10.123
```

```
Name: sales-scan.example.com
```

```
Address: 100.10.122
```

```
Name: sales-scan.example.com
```

```
Address: 100.10.121
```

Stop the SCAN listener and the SCAN VIP resources:

```
# $GRID_HOME/bin/srvctl stop scan_listener
```

```
# $GRID_HOME/bin/srvctl stop scan
```

```
# $GRID_HOME/bin/srvctl status scan
```

```
SCAN VIP scan1 is enabled
```

```
SCAN VIP scan1 is not running
```

```
SCAN VIP scan2 is enabled
```

```
SCAN VIP scan2 is not running
```

```
SCAN VIP scan3 is enabled
```

```
SCAN VIP scan3 is not running
```

```
# $GRID_HOME/bin/srvctl status scan_listener
```

```
SCAN Listener LISTENER_SCAN1 is enabled
```

```
SCAN listener LISTENER_SCAN1 is not running
```

```
SCAN Listener LISTENER_SCAN2 is enabled
```

```
SCAN listener LISTENER_SCAN2 is not running
```

```
SCAN Listener LISTENER_SCAN3 is enabled
```

```
SCAN listener LISTENER_SCAN3 is not running
```

The SCAN VIP resources still show the old IP addresses:

```
# $GRID_HOME/bin/srvctl config scan
SCAN name: sales-scan, Network/10.100.10.0/255.255.255.0/eth1
SCAN VIP name: scan1, IP: /sales-scan.example.com/10.100.10.81
SCAN VIP name: scan2, IP: /sales-scan.example.com/10.100.10.82
SCAN VIP name: scan3, IP: /sales-scan.example.com/10.100.10.83
Now tell CRS to update the SCAN VIP resources:
# $GRID_HOME/bin/srvctl modify scan -n sales-scan.example.com
To verify that the change was successful, check the SCAN configuration again:
# $GRID_HOME/bin/srvctl config scan
SCAN name: sales-scan, Network/10.100.10.0/255.255.255.0/eth1
SCAN VIP name: scan1, IP: /sales-scan.example.com/10.100.10.121
SCAN VIP name: scan2, IP: /sales-scan.example.com/10.100.10.122
SCAN VIP name: scan3, IP: /sales-scan.example.com/10.100.10.123

Start SCAN and the SCAN listener:
# $GRID_HOME/bin/srvctl start scan
# $GRID_HOME/bin/srvctl start scan_listener
Please note that if the SCAN VIPs are being changed because of a change of the subnet of the
public network additional changes may be required , e.g. the node VIPs and the network resource
(ora.net1.network). For more information please refer to Document 276434.1 and the 11.2
documentation.
My Oracle Support
```

### Which three statements are true about the logging or tracing of Java-based Oracle tools?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Which three statements are true about the logging or tracing of Java-based Oracle tools?

- A.  
By default, Oracle logs actions for DBCA and DBUA.
- B.  
For tools such as SRVCTL and CLUVFY, tracing can be enabled by setting `srvm trace = true` your shell environment before executing the commands.
- C.  
The `srvm_trace` environment variable can be set only in the user's shell environment file.
- D.  
To disable tracing for Java-based tools, unset the `srvm_trace` variable: `export srvm_trace=`

#### Explanation:

Enabling Tracing for Java-Based Tools and Utilities in Oracle RAC  
All Java-based tools and utilities that are available in Oracle RAC are called by executing scripts of the same name as the tool or utility. This includes the Cluster Verification Utility (CVU), Database Configuration Assistant (DBCA), the Net Configuration Assistant (NETCA), Server Control Utility (SRVCTL), and the Global Services Daemon (GSD). For example to run DBCA, enter the command `dbca`. By default, Oracle Database enables traces for DBCA and the Database Upgrade Assistant (DBUA). For the CVU, GSDCTL, and SRVCTL, you can set the `SRVM_TRACE` environment variable to `TRUE` to make Oracle Database generate traces. Oracle Database writes traces to log files. For example, Oracle Database writes traces to log files in Oracle `home/cfgtoollogs/dbca` and Oracle `home/cfgtoollogs/dbua` for DBCA and DBUA, respectively. Oracle Real Application Clusters Administration and Deployment Guide

### Which two statements are true about sharing servers among the various server pools on your cluster?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

You have created four top-level server pools called OLTP, BATCH, DSS, and support for a policy-

managed, eight-node cluster. The OLTP pool has two child pools called OLTP\_CHILD1 and OLTP\_CHILD2. Which two statements are true about sharing servers among the various server pools on your cluster?

- A.  
Servers are never shared and are always mutually exclusive between top-level server pools.
- B.  
Use the `exclusive_pools` attribute to prevent servers being shared between top-level server pools.
- C.

Use the `exclusive_pools` attribute to prevent servers being shared between `OLTP_CHILD1` and `OLTP_CHILD2`.

D.  
Use the `exclusive_pools` attribute to prevent servers being shared between batch, dss, support, `oltp_child1`, and `OLTP_CHILD2`.

E.  
Servers are never shared and are always mutually exclusive between top-level server pools and between sibling child pool.

Explanation:

Top-level server pools:  
`EXCLUSIVE_POOLS`

This optional attribute indicates if servers assigned to this server pool are shared with other server pools. A server pool can explicitly state that it is exclusive of any other server pool that has the same value for this attribute. Two or more server pools are mutually exclusive when the sets of servers assigned to them do not have a single server in common. For example, server pools A and B must be exclusive if they both set the value of this attribute to `foo_A_B`.

Top-level server pools are mutually exclusive, by default.

**Which parameter is mandatory?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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`

Explanation:

Oracle ASM Parameter Setting Recommendations

This section contains information about the following parameters for Oracle ASM:  
Oracle Automatic Storage Management Administrator's Guide

**Which three fragments will complete this statement correctly?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Which three fragments will complete this statement correctly? In a cluster environment, an ACFS volume \_\_\_\_\_.

- A.  
Will be automatically mounted by a node on reboot by default
- B.  
Must be manually mounted after a node reboot
- C.  
Will be automatically mounted by a node if it is defined as cluster stack startup if it is included in the ACFS mount registry.
- D.  
Will be automatically mounted to all node if it is defined as cluster resource when dependent cluster resources requires access
- E.  
Will be automatically mounted to all node in the cluster when the file system is registered

F.  
Must be mounted before it can be registered

Explanation:

By default, an Oracle ACFS file system that is inserted into the cluster mount registry is automatically mounted on all cluster members, including cluster members that are added after the registry addition.

The Oracle ACFS registry resource actions are designed to automatically mount a file system only one time for each Oracle Grid Infrastructure initialization to avoid potential conflicts with administrative actions to dismount a given file system.

Oracle Automatic Storage Management Administrator's Guide

**Which three actions are required to create a general purpose ASM cluster file system (ACFS) to be automatically mounted by Oracle Clusterware?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Which three actions are required to create a general purpose ASM cluster file system (ACFS) to be automatically mounted by Oracle Clusterware?

A.  
**Format an ASM volume with an ASM cluster file system.**

B.  
**Create mount points on all cluster nodes where the ASM cluster file system will be mounted.**

C.  
Manually add an entry to /etc/fstab defining the volume, mount point, and mount options on each node in the cluster.

D.  
**Register the mount point.**

Explanation:

By default, an Oracle ACFS file system that is inserted into the cluster mount registry is automatically mounted on all cluster members, including cluster members that are added after the registry addition. However, the cluster mount registry also accommodates single-

node and multi-node (subset of cluster nodes) file system registrations. The mount registry actions for each cluster member mount only registered file systems that have been designated for mounting on that member.

The Oracle ACFS registry resource actions are designed to automatically mount a file system only one time for each Oracle Grid Infrastructure initialization to avoid potential conflicts with administrative actions to dismount a given file system.

Oracle Automatic Storage Management Administrator's Guide

**Which three statements are true about cluster (OCR) and local registries (OLR)?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Which three statements are true about cluster (OCR) and local registries (OLR)?

A.  
**The ocrdump command can be used to dump the contents of either the cluster registry (OCR) or the local registry (OLR).**

B.  
**Information in the OCR is organized by keys that are associated with privileges. Therefore, the root user will see more information when using ocrdump than the grid user.**

C.  
After the cluster is created, the size of the OCR is static unless the Clusterware software is upgraded.

D.  
**The ocrdump command can be used to dump information for a single OCR key.**

Explanation:

You can use several options for OCRDUMP. For example, you can limit the output to a key and its descendants. You can also write the contents to an XML file that you can view using a browser. OCRDUMP writes the OCR keys as ASCII strings and values in a data type format. OCRDUMP retrieves header information based on a best effort basis.

Displays OCR content in the terminal window in XML format

ocrdump -stdout -backupfile Grid\_home/cdata/cluster\_name/file\_name

Displays the content of the in the Grid\_home/cdata/ cluster\_name/ file\_name directory. You must

**What is the correct interpretation?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

In your two-instance RAC database, you find that FAST\_START\_MTTR\_TARGET is set to 300 on both the nodes. The ESTD\_CLUSTER\_AVAILABLE\_TIME column in the V\$INSTANCE\_RECOVERY view shows the value 100.

What is the correct interpretation? (Choose 2)

- A.  
**In the event of an instance failure, the database will become partially available In 100 seconds.**
- B.  
In the event of an instance failure, the database will become fully available in 100 seconds.
- C.  
In the event of an instance failure, the failed instance will become fully available In 300 seconds.
- D.  
In the event of an instance failure, the database will become partially available in 200 seconds.
- E.  
**In the event of an instance failure, the failed instance will become available In 300 seconds.**

**Explanation:**

FAST\_START\_MTTR\_TARGET

| Property        | Description   |
|-----------------|---|
| Parameter type  | Integer   |
| Default value   | 0   |
| Modifiable      | ALTER SYSTEM  |
| Range of values | 0 to 3600 seconds   |
| Basic           | No  |
| Oracle RAC      | Multiple instances can have different values, and you can change the values at runtime. |

FAST\_START\_MTTR\_TARGET enables you to specify the number of seconds the database takes to perform crash recovery of a single instance. When specified, FAST\_START\_MTTR\_TARGET is overridden by LOG\_CHECKPOINT\_INTERVAL. V\$INSTANCE\_RECOVERY Estimated time (in seconds) that the cluster would become partially available should this instance fail. This column is only meaningful in an Oracle Real Application Clusters (Oracle RAC) environment. In a non-Oracle RAC environment, the value of this column is null. Oracle Database Reference 11g Release 2 (11.2)

**What is true about this situation?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

A policy-managed RAC database that hosts the ACCOUNTS service has only three instances running rather than the usual four on a six-node cluster.

Five of the six cluster nodes are active and the sixth (RACNODE6) is down for maintenance. You decide to check the state of the servers in the server pools to see if there are problems with

the OLTP pool to which the ACCOUNTS service has been assigned. You used the crsctl status server-f command and see that the RACNODE3 node has STATE=VISIBLE. What is true about this situation?

- A.  
**The Cluster Synchronization Services Daemon (cssd) is running and RACNODE3 is considered to be part of the cluster, but the Cluster Ready Services Daemon (crsd) is currently not running.**
- B.  
Both the Cluster Ready Services Daemon (crsd) and the Cluster Synchronization Service Daemon (cssd) are running on RACNODE3, but RACNODE3 is currently being moved to another

server pool.

C.

The Cluster Ready Services Daemon (crsd) is running and RACNODE3 is considered to be a part of the cluster, but the Cluster Synchronization Services Daemon (cssd) is currently not running.

D.

Neither the Cluster Ready Services Daemon (crsd) nor the Cluster Synchronization Services Daemon (cssd) is running on RACNODE3.

Explanation:

Cluster Synchronization Service daemon (CSSD)

An Oracle Clusterware component that discovers and tracks the membership state of each node by providing a common view of membership across the cluster. CSS also monitors process health, specifically the health of the database instance. The Global Enqueue Service Monitor (LMON), a background process that monitors the health of the cluster database environment and registers and de-registers from CSS.

Cluster Ready Services Daemon (CRSD)

The primary Oracle Clusterware process that performs high availability recovery and management operations, such as maintaining OCR. Also manages application resources and runs as root user (or by a user in the admin group on Mac OS X-based systems) and restarts automatically upon failure

STATE=VISIBLE

Servers that have Oracle Clusterware running, but not the Cluster Ready Services daemon (crsd), are put into the VISIBLE state. This usually indicates an intermittent issue or failure and Oracle Clusterware trying to recover (restart) the daemon. Oracle Clusterware cannot manage resources on servers while the servers are in this state.

Oracle Clusterware Administration and Deployment Guide

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

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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.

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



### What is true about this situation?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

A policy-managed RAC database that hosts the ACCOUNTS service has only three instances running rather than the usual four on a six-node cluster.

Five of the six cluster nodes are active and the sixth (RACNODE6) is down for maintenance.

You decide to check the state of the servers in the server pools to see if there are problems with the OLTP pool to which the ACCOUNTS service has been assigned.

You used the crsctl status server-f command and see that the RACNODE3 node has STATE=VISIBLE.

What is true about this situation?

A.

The Cluster Synchronization Services Daemon (cssd) is running and RACNODE3 is considered to be part of the cluster, but the Cluster Ready Services Daemon (crsd) is currently not running.

B.

Both the Cluster Ready Services Daemon (crsd) and the Cluster Synchronization Service Daemon (cssd) are running on RACNODE3, but RACNODE3 is currently being moved to another server pool.

C.

The Cluster Ready Services Daemon (crsd) is running and RACNODE3 is considered to be a part of the cluster, but the Cluster Synchronization Services Daemon (cssd) is currently not running.

D.

Neither the Cluster Ready Services Daemon (crsd) nor the Cluster Synchronization Services Daemon (cssd) is running on RACNODE3.

#### Explanation:

Cluster Synchronization Service daemon (CSSD)

An Oracle Clusterware component that discovers and tracks the membership state of each node by providing a common view of membership across the cluster. CSS also monitors process health, specifically the health of the database instance. The Global Enqueue Service Monitor (LMON), a background process that monitors the health of the cluster database environment and registers and de-registers from CSS.

Cluster Ready Services Daemon (CRSD)

The primary Oracle Clusterware process that performs high availability recovery and management operations, such as maintaining OCR. Also manages application resources and runs as root user (or by a user in the admin group on Mac OS X-based systems) and restarts automatically upon failure

STATE=VISIBLE

Servers that have Oracle Clusterware running, but not the Cluster Ready Services daemon (crsd), are put into the VISIBLE state. This usually indicates an intermittent issue or failure and Oracle Clusterware trying to recover (restart) the daemon. Oracle Clusterware cannot manage resources on servers while the servers are in this state.

Oracle Clusterware Administration and Deployment Guide

### Select three responses that describe the VIPs that will exist in this configuration.

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

You have configured your eight-node cluster to use GNS. The network administrator has established delegated subdomain for the Cluster which is MYCLUSTER.EXAMPLE.COM. DHCP has been configured so that the cluster now manages IP addresses within the cluster.

Select three responses that describe the VIPs that will exist in this configuration.

A.

3 GNS VIPs

B.

8 Node VIPs

C.

3 SCAN VIPs

D.

1 GNS VIP

E.  
3 Node VIPs

Explanation:

Implementing GNS

If you use GNS, then you must specify a static IP address for the GNS VIP address, and delegate a subdomain to be delegated to that static GNS VIP address.

Dynamic IP address assignment using Oracle Grid Naming Service (GNS)

If you select this option, then network administrators assign static IP address for the physical host name and dynamically allocated IPs for the Oracle Clusterware managed VIP addresses. In this case, IP addresses for the VIPs are assigned by a DHCP and resolved using a multicast domain name server configured as part of Oracle Clusterware within the cluster. If you plan to use GNS, then you must have the following:

Oracle Grid Infrastructure Installation Guide

**What will be the SCAN Name field default to if you enter cluster01 in the cluster Name field and cluster01.example.com in the GNS Sub Domain field?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

On the OUI Grid Plug and Play information page, you can configure GRID Naming Service (GNS). What will be the SCAN Name field default to if you enter cluster01 in the cluster Name field and cluster01.example.com in the GNS Sub Domain field?

- A.  
cluster01.example.com
- B.  
cluster01-qns.example.com
- C.  
**cluster01-scan.cluster01.example.com**
- D.  
cluster-vip.example.com

Explanation:

If you specify a GNS domain, then the SCAN name defaults to clustername-scan.

GNS\_domain. Otherwise, it defaults to clustername-scan.current\_domain. For example, if you start Oracle Grid Infrastructure installation from the server node1, the cluster name is mycluster, and the GNS domain is grid.example.com, then the SCAN Name is mycluster-scan.grid.example.com.

Oracle Grid Infrastructure Installation Guide

**Which two statements are true about the vip application resource that must be created?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

You have a four-node cluster with four node VIPs and three scan VIPs already in use after Oracle Grid Infrastructure installation.

You plan to manage a customer-developed, web-based application with the Oracle Grid Infrastructure to provide high availability.

Which two statements are true about the vip application resource that must be created?

- A.  
The vip application must be active on the same subnet and ethernet adaptor as the Interconnect.
- B.  
**The vip application must be active on the same subnet and ethernet adaptor as the public ip Address.**
- C.  
The vip application must not be online on the same node as any of the scan vips.
- D.  
The vip application must be active on a different node than the application depending on Che by using the dispersion start dependency in the resource definition.
- E.  
**The vip application must be enabled to run as the root user.**

Explanation:

If clients of an application access the application through a network, and the placement policy for the application allows it to fail over to another node, then you must register a

virtual internet protocol address (VIP) on which the application depends. An application VIP is a cluster resource that Oracle Clusterware manages (Oracle Clusterware provides a standard VIP agent for application VIPs). You should base any new application VIPs on this VIP type to ensure that your system experiences consistent behavior among all of the VIPs that you deploy in your cluster

A VIP resource is typically owned by root because VIP related operations require root privileges. You must assign a VIP address to each node in the cluster. Each VIP address must be on the same subnet as the public IP address for the node and should be an address that is assigned a name in the DNS

Oracle Clusterware Administration and Deployment Guide

### Which two statements are true about the high availability capabilities of Oracle HA services?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

In the high availability services' provided by Oracle Clusterware are used to protect Oracle resource such as RAC database instances, RAC database services, and other components of the Oracle Infrastructure, and non-Oracle resources as well.

Which two statements are true about the high availability capabilities of Oracle HA services?

- A.  
RAC databases may have their instances failed over in some cases.
- B.  
ASM instances may be failed over if fewer than three nodes remain In the cluster, so that there are always at least three ASM instances available.
- C.  
If a node fails, then all resources that were active on that node will be failed over to a surviving node if any exists.
- D.  
If a node fails, then cluster resources that were active on that node may be failed over to a
- E.  
surviving node if any exists, but local resources are not failed over.
- F.  
HA services will only fail over a resource upon failure of the node where the resource was active.

### which account(s) can the Oracle Clusterware processes run for the Linux platform?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Under which account(s) can the Oracle Clusterware processes run for the Linux platform?

- A.  
only the root account
- B.  
only the Grid Infrastructure owner account
- C.  
only the database software owner account
- D.  
some under the Grid Infrastructure owner account and some under the root account
- E.  
some under the database software owner account and some under the root account

#### Explanation:

Because Oracle Clusterware works closely with the operating system, system administrator privileges are required for some installation tasks. In addition, some Oracle Clusterware processes must run as the special operating system user, root.

Oracle Database 2 Day + Real Application Clusters Guide

### Identify the three valid storage options for Grid Infrastructure voting disk and Install.

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Identify the three valid storage options for Grid Infrastructure voting disk and Install.

- A. a certified Cluster File System (CFS)
- B. a certified Network File System (NFS)
- C. ASM Cluster File System (ACFS)
- D. Automatic Storage Management (ASM)
- E. shared disk slices (block or raw devices)

Explanation:

**Shared Storage Planning for Grid Infrastructure**

| Storage Option                            | Voting/OCR | Oracle Software |
|---|------------|-----------------|
| Automatic Storage Manager (ASM)           | Yes        | No              |
| ASM Cluster File System (ACFS)            | No         | Yes             |
| Oracle Cluster File System (OCFS2)        | Yes        | Yes             |
| Red Hat Global File System (GFS)          | Yes        | Yes             |
| NFS (certified only)                      | Yes        | Yes             |
| Shared disk slices (block or raw devices) | No         | No              |

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 2 – 4

**Which two conditions are required by the ASM fast mirror resynchronization to track block changes for a set period of time before dropping the disk from the disk group?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Which two conditions are required by the ASM fast mirror resynchronization to track block changes for a set period of time before dropping the disk from the disk group?

- A. Redundancy is normal or high.
- B. compatibility.rdbms is set to a value of at least 11.1.
- C. disk\_repair\_time is set to a nondefault value.
- D. block\_change\_tracking IS enabled.
- E. db\_block\_checking is enabled.
- F. resumable\_timeout is set to a nondefault value.

Explanation:

ASM Fast Mirror Resync

Enabled when COMPATIBLE.RDBMS >= 11.1

Whenever ASM is unable to write an extent, ASM takes the associated disk offline. If the corresponding disk group uses ASM mirroring (NORMAL or HIGH redundancy), at least one mirror copy of the same extent exists on another disk in the disk group.

Before Oracle Database 11g, ASM assumed that an offline disk contains only stale data and no longer reads from such disks. Shortly after a disk is put offline, ASM drops it from the disk group by re-creating the extents allocated to the disk on the remaining disks in the disk group using mirrored extent copies. This process is quite resource intensive and can take hours to complete. If the disk is replaced or the failure is repaired, the disk must be added again and another rebalance operation must take place.

D60488GC11

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 8 – 32

**What is the proper setting for the ASM diskgroup attribute on diskgroups being used by all the instances?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

The ASM instance in your environment can support databases from version 10.2.0.4 through

11.2.0.1 You want to use all ASM functionality possible for each database, including OCR and voting files in the ASM. What is the proper setting for the ASM diskgroup attribute on diskgroups being used by all the instances?

- A.  
Set compatible.asm to 11.2 and compatible.rdbms to 10.2
- B.  
Set compatible.asm to 10.2 and compatible.rdbms to 10.2
- C.  
Set compatible.asm to 11.2 and compatible.rdbms to 11.2
- D.  
Set compatible.asm to 10.2 and compatible.rdbms to 11.2

Explanation:

Disk Group Compatibility Attributes  
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. This setting also affects the format of the data structures for the Oracle ASM metadata on the disk. The format of other file contents is determined by Oracle ASM Dynamic Volume Manager (Oracle ADVM) and the database instance.

For Oracle ASM in Oracle Database 11g, 10.1 is the default setting for the COMPATIBLE.ASM attribute when using the SQL CREATE DISKGROUP statement, the ASMCMD mkdg command, and Oracle Enterprise Manager Create Disk Group page. When creating a disk group with ASMCA, the default setting is 11.2.

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. Before advancing the COMPATIBLE.RDBMS attribute, ensure that the values for the COMPATIBLE initialization parameter for all of the databases that access the disk group are set to at least the value of the new setting for COMPATIBLE.RDBMS.

For example, if the COMPATIBLE initialization parameters of the databases are set to either 11.1 or 11.2, then COMPATIBLE.RDBMS can be set to any value between 10.1 and 11.1 inclusively. Oracle Automatic Storage Management Administrator's Guide  
11g Release 2 (11.2)

**Which two statements are true about the behavior of this resource after it is added?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

A third party application has provided files containing the command syntax required to register and manage the application for high availability with Oracle Grid Infrastructure.

The command to define the application is as follows:

```
crsctl add resource LongApp -type cluster_resource -attr  
"ACTION_SCRIPT=/u01/ogi/scripts/longapp.scr , PLACEMENT=restricted ,  
HOSTING_MEMBERS=RACNODE1 RACNODE2 , CHECK_INTERVAL=30  
,START_DEPENDENCIES=hard(LongAppvip) , STOP_DEPENDENCIES=hard (LongAppvip),  
RESTART_ATTEMPTS=2"
```

Which two statements are true about the behavior of this resource after it is added?

- A.  
It will be a policy-managed resource using the server pool containing nodes RACNODE1 and RACNODE2.
- B.  
The LongAppvip vip resource must be activated to successfully activate this resource, but will not be activated automatically if not already active.
- C.  
The LongAppvip vip resource must be stopped before this resource is stopped.
- D.  
It will be an administrator-managed resource hosted by nodes RACNODE1 and RACNODE2.

Explanation:

Configurable Resource Attributes  
HOSTING\_MEMBERS

A space-separated, ordered list of cluster server names that can host a resource. This attribute is required only when using administrator management, and when the value of the PLACEMENT

attribute is set to favored or restricted. When registering applications as Oracle Clusterware resources, use the SERVER\_POOLS attribute, instead.

#### START\_DEPENDENCIES

hard([intermediate:][global:]{resource\_name | type:resource\_type})Specify a hard start dependency for a resource when you want the resource to start only when a particular resource or resource of a particular type starts.

weak([concurrent:][global:][uniform:]{resource\_name | type:resource\_type})Specify a weak start dependency for a resource when you want that resource to start despite whether named resources are running, or not. An attempt to start this resource also attempts to start any resources on which this resource depends if they are not running.

#### STOP\_DEPENDENCIES

hard([intermediate:][global:][shutdown:]{resource\_name | type:resource\_type})Specify a hard stop dependency for a resource that you want to stop when named resources or resources of a particular resource type stop.

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

---

### Which command will set full debugging for the communications and resource management components; of the CRS module on nodes host01 and host03?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Which command will set full debugging for the communications and resource management components; of the CRS module on nodes host01 and host03?

- A.  
srvctl set log -m CRS -C CRSCTI 5, CRSCOMM 5 -n host01,host03
- B.  
ocrconfig set log CRS "CRSRTI,CRSCOMM:5" -nodelist host01,host03
- C.  
**crsctl set log CRS "CRSRTI:5,CRSCOMM:5" -nodelist host01,host03**
- D.  
crsctl lsmodules crs set log CRSRTI:5,CRSCOMM:5 -nodelist host01, host03

#### Explanation:

crsctl set log

Use the crsctl set log command to set log levels for Oracle Clusterware.

You can enable debugging for Oracle Clusterware modules and their components, and for resources, by setting environment variables or by running crsctl set log commands, using the following syntax:

Syntax

```
crsctl set {log | trace} module_name "component:debugging_level  
[,component:debugging_level][,...]"
```

Usage Notes

You can set log levels for various components of the three modules, CRS, CSS, and EVM. If you choose the all option, then you can set log levels for all components of one module with one command. Use the crsctl lsmodules command to obtain a list of components for each module.

Enter a comma-delimited list of component name-log level pairs enclosed in double quotation marks ("").

The log\_level is a number from 1 to 5 that sets the log level for the component or resource, where 1 is the least amount of log output and 5 provides the most detailed log output. The default log level is 2.

To set log levels for resources, specify the name of a particular resource, or a comma-delimited list of resource name-log level pairs enclosed in double quotation marks ("").

To set debugging levels on specific cluster nodes, include the -nodelist keyword and the names of

the nodes, as follows:

```
crsctl set log crs "CRSRTI:1,CRSCOMM:2" -nodelist node1,node2
```

Oracle Clusterware Administration and Deployment Guide

---

### What must you know to select the appropriate command for this operation?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

You have an ASM cluster and you want to terminate all the ASM instances. Assume that the database Instances supported by the ASM are already shut down. What must you know to select the appropriate command for this operation?

- A.

**Are the OCR and voting disk files stored in ASM diskgroups?**

- B. Are the SYSTEM and SYSAUS tablespaces stored in ASM diskgroups?
- C. Is the undo tablespace stored in ASM diskgroups?
- D. Are the redo log files stored in ASM diskgroups?

Explanation:

Oracle strongly recommends that you shut down all database instances that use the Oracle ASM instance and dismount all file systems mounted on Oracle ASM Dynamic Volume Manager (Oracle ADVM) volumes before attempting to shut down the Oracle ASM instance. If Oracle Cluster Registry (OCR) or voting files are stored in a disk group, the disk group can only be dismounted by shutting down the Oracle ASM instance as part of shutting down the clusterware on a node. To shut down the clusterware, run `crsctl stop crs`.  
Oracle Automatic Storage Management Administrator's Guide

**Which are the key factors that you should consider before converting a single-Instance database Oracle Real Application Cluster...?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Which are the key factors that you should consider before converting a single-Instance database Oracle Real Application Cluster (RAC) database to guarantee a successful media recovery?

- A. **If the database is in archive log mode, the archive file format requires a thread number.**
- B. **The archive logs from all nodes must be accessible to all nodes in the cluster database.**
- C. The storage option must be Automatic Storage Management (ASM).
- D. All database files must be migrated to Oracle Managed Files (OMF).

Explanation:

Issues for Converting Single Instance Databases to Oracle RAC

Backup procedures should be available before conversion takes place.  
Archiving in Oracle RAC environments requires a thread number in the archive file format.  
The archived logs from all instances of an Oracle RAC database are required for media recovery.  
By default, all database files are migrated to Oracle Managed Files (OMF).  
D60488GC11  
Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 11 – 24

**Which three steps are required to enable ODP.NET clients that connect to the GL services to receive FAN High Availability Events?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

You are managing a three-instance policy-managed RAC database PROD. You created a service called GL for the PROD database by using the following command:

```
oracle@gr7597~]$srvctl add service d PROD -s GL g SP1 -c singleton -y manual
```

Examine the following output:

```
(oracle@gr7597~]$srvctl config database -d PROD -a
```

Database unique name: PROD

base name: PROD

Oracle home: /u01/app/oracle/product/11.2.0/dbhome\_1 cle user: oracle Spfile:

+DATA/PROD/spfile6RACDB.ora

Domain:

Start options: open

Stop options: immediate

Database role: PRIMARY

Management policy: AUTOMATIC

Server pools: SP1 Database instances:

Disk Groups: DATA, FRA

Services:

Database is enabled Database is policy managed

```
(oracle@gr7597~]$ crsctl stat res ora.PROD.db
```

NAME=ora.PROD.db  
TYPE=cluster\_resource  
TARGET=ONLINE  
STATE^ONLINE on gr7597 ONLINE on gr7602 ONLINE on gr7633  
[oracle@gr7597~]S crsctl stat serverpool ora.SPI  
NAME=ora.SPI  
ACTIVE\_SERVERS=gr7S97 gr7 602 gr7633  
[oracle@gr7597-JS crsctl NAME=ora.prod.gl.svc  
TYPE=ora.service - type  
TARGET=ONLINE  
STATE=ONLINE on gr7597  
Which three steps are required to enable ODP.NET clients that connect to the GL services to receive FAN High Availability Events?

- A.  
Enable Advanced Queuing notifications by using SRVCTL as shown in the following command:  
srvctl modify service -d prod -s gl -q TRUE -j LONG
- B.  
Execute the following statement for the users that will be connecting by way of the -Net Application, where user\_name is the user name:  
EXECUTE  
DBMS\_AQADM.GRANT\_QUEUE\_PRIVILEGE('DEQUEUE','SYS.SYSSERVICE\_METRICS',  
user\_name);
- C.  
Enable Transparent Application Failover (TAF), either on the client or for the service.
- D.  
Enable Fast Connection Failover for ODP.NET connection pools by subscribing to FAN High Availability events. Do this by setting the HA events connection string attribute to true at connection time.
- E.  
Link client applications with the client thread or operating system library.

Explanation:

Perform the following steps to enable FAN for ODP.NET clients:

```
srvctl modify service -d crm -s odpnet.example.com -q TRUE
```

```
EXECUTE
```

```
DBMS_AQADM.GRANT_QUEUE_PRIVILEGE('DEQUEUE','SYS.SYSSERVICE_METR ICS',  
user_name);
```

Oracle Real Application Clusters Administration and Deployment Guide

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

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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.

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



### Which two statements are true about instance recovery in a RAC environment?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Which two statements are true about instance recovery in a RAC environment?

- A.  
Parallel instance recovery will work even if the recovery\_parallelism initialization parameter set to 0 or 1.
- B.  
**Increasing the size of the default buffer cache can speed up instance recovery because instance recovery may use as much as 50 percent of the default buffer cache for recovery buffers.**
- C.  
The fast\_start\_mttr\_target initialization parameter includes both instance startup and recovery time.
- D.  
**The fast\_start\_mttr\_target initialization parameter specifies only the instance recovery time.**

#### Explanation:

Many sites run with too few redo logs that are too small. Small redo logs cause system checkpoints to continuously put a high load on the buffer cache and I/O system. If there are too few redo logs, then the archive cannot keep up, and the database will wait for the archive

process to catch up.

With the Fast-Start Fault Recovery feature, the FAST\_START\_MTTR\_TARGET initialization parameter simplifies the configuration of recovery time from instance or system failure. FAST\_START\_MTTR\_TARGET specifies a target for the expected mean time to recover (MTTR), that is, the time (in seconds) that it should take to start up the instance and perform cache recovery.

Oracle Database Performance Tuning Guide

### Which two statements are true regarding undo management in the RAC environment?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Which two statements are true regarding undo management in the RAC environment?

- A.  
You can use Automatic Undo Management (AUM) in some of the instances and manual undo management in the rest of the instances in a RAC database.
- B.  
In a policy-managed RAC database, Oracle automatically allocates the undo tablespace even the Oracle Managed Files (OMF) is disabled in a database.
- C.  
**In a policy-managed RAC database, Oracle automatically allocates the undo tablespace if the database is OMF enabled.**
- D.  
**You can dynamically switch undo tablespace assignments by executing the ALTER SYSTEM SET UNDO\_TABLESPACE statement from any instance in a administrator managed database.**

#### Explanation:

You assign undo tablespaces in your Oracle RAC administrator-managed database by specifying a different value for the UNDO\_TABLESPACE parameter for each instance in your SPFILE or individual PFILES. For policy-managed databases, Oracle automatically allocates the undo tablespace when the instance starts if you have Oracle Managed Files enabled.

You can switch from using one undo tablespace to another. Because the UNDO\_TABLESPACE initialization parameter is a dynamic parameter, the ALTER SYSTEM SET statement can be used to assign a new undo tablespace.

### Which three statements are true about services and the Resource Manager?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Which three statements are true about services and the Resource Manager?

- A.  
**The Resource Manager can manage the relative priority of services within an instance by binding services directly to consumer groups if services are mapped to consumer groups by the DBA.**

B.

When a client connects using a service, the service can be mapped to a consumer group, enabling the Resource Manager to manage work requests by service in the order of their importance.

C.

The srvctl utility is used to map services to consumer groups.

D.

The Resource Manager offers benefits in managing workloads because priority is given to business functions rather than the sessions that support those business functions.

Explanation:

A resource consumer group (consumer group) is a collection of user sessions that are grouped together based on their processing needs. When a session is created, it is automatically mapped to a consumer group based on mapping rules that you set up. As a database administrator (DBA), you can manually switch a session to a different consumer group. Before you enable the Resource Manager, you must specify how user sessions are assigned to resource consumer groups. You do this by creating mapping rules that enable the Resource Manager to automatically assign each session to a consumer group upon session startup, based upon session attributes.

Oracle Database Resource Manager (the Resource Manager) enables you to manage multiple workloads within a database that are contending for system and database resources.

In addition, the Database Resource Manager can map services to consumer groups. Therefore, you can automatically manage the priority of one service relative to others. You can use consumer groups to define relative priority in terms of either ratios or resource consumption.

Oracle Database Administrator's Guide

Which two recommendations may reduce this problem?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

You notice that there is a very high percentage of wait time for the enq:HW-contention event in your RAC database that has frequent insert operations.

Which two recommendations may reduce this problem?

A.

shorter transactions

B.

increasing sequence cache sizes

C.

using reverse key indexes

D.

uniform and large extent sizes

E.

automatic segment space management

F.

smaller extent sizes

Explanation:

Segments have High Water Mark (HWM) indicating that blocks below that HWM have been formatted. New tables or truncated tables [that is truncated without reuse storage clause], have HWM value set to segment header block. Meaning, there are zero blocks below

HWM. As new rows inserted or existing rows updated (increasing row length), more blocks are added to the free lists and HWM bumped up to reflect these new blocks. HW enqueues are acquired in Exclusive mode before updating HWM and essentially HW enqueues operate as a serializing mechanism for HWM updates. Allocating additional extent with instance keyword seems to help in non-ASSM tablespace

Serialization of data blocks in the buffer cache due to lack of free lists, free list groups, transaction slots (INTRANS), or shortage of rollback segments.

This is particularly common on INSERT-heavy applications, in applications that have raised the block size above 8K, or in applications with large numbers of active users and few rollback segments. Use automatic segment-space management (ASSM) and automatic undo management to solve this problem.

HW enqueue The HW enqueue is used to serialize the allocation of space beyond the high water mark of a segment.

If this is a point of contention for an object, then manual allocation of extents solves the problem.

### Which two statements are true regarding initialization parameters In a RAC environment?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

You are managing RAC database with policy managed services. The database is started by using an SPFILE.

Which two statements are true regarding initialization parameters In a RAC environment?

- A.  
All initialization parameters must have identical settings on all instances.
- B.  
**All instances in the cluster database use the same SPFILE.**
- C.  
To change values for initialization parameters for an instance, you must log in to that instance.
- D.  
**All initialization parameters for all instances can be changed from any instance In a RAC database.**

#### Explanation:

RAC Initialization Parameter Files

An SPFILE is created if you use the DBCA.

The SPFILE must be created in an ASM disk group or a cluster file system file.

All instances use the same SPFILE.

If the database is created manually, create an SPFILE from a PFILE.

SPFILE Parameter Values and RAC

You can change parameter settings using the ALTER SYSTEM SET command from any instance  
D60488GC11

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 12 – 22, 23

### Which are the two best location options for archive logs so that they can be accessed during recovery without DBA intervention?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

You need to set up a three-Instance RAC database. The data files and fast recovery area will be stored in ASM diskgroups called +data and +fra, respectively. The ASM disk groups will be mounted on all ASM Instances.

Which are the two best location options for archive logs so that they can be accessed during recovery without DBA intervention?

- A.  
**Cluster File System with each instance writing to a shared location**
- B.  
Cluster File System with each instance writing to a separate location as long as all the locations are in directories under the same mount point
- C.  
**the ASM diskgroup +fra with the db\_recovery\_file\_dest parameter set to the same value on all instances**
- D.  
a raw or block device

#### Explanation:

The primary consideration when configuring archiving is to ensure that all archived redo logs can be read from every node during recovery, and if possible during backups. During recovery, because the archived log destinations are visible from the node that performs the recovery, Oracle RAC can successfully recover the archived redo log data.

The fast recovery area for an Oracle RAC database must be placed on an Oracle ASM disk group, a cluster file system, or on a shared directory that is configured through a network file system file for each Oracle RAC instance. In other words, the fast recovery area must be shared among all of the instances of an Oracle RAC database. The preferred configuration for Oracle RAC is to use Oracle Automatic Storage Management (Oracle ASM) for storing the fast recovery area, using a different disk group for your recovery set than for your data files.

The location and disk quota must be the same on all instances. Oracle recommends that you place the fast recovery area on the shared Oracle ASM disks. In addition, you must set the DB\_RECOVERY\_FILE\_DEST and DB\_RECOVERY\_FILE\_DEST\_SIZE parameters to the same values on all instances.

Oracle Database 2 Day + Real Application Clusters Guide

**Which two statements are true regarding the Automatic Workload Repository (AWR) In a RAC environment?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Which two statements are true regarding the Automatic Workload Repository (AWR) In a RAC environment?

- A.  
The AWR includes time model statistics based on time usage for activities displayed in the V\$SYS\_TIME\_MODEL and V\$SESS\_TIME\_MODEL views.
- B.  
The AWR is controlled by the statistics\_level initialization parameter and it must be set to TYPICAL or BASIC
- C.  
The Manageability Monitor Processes (MMON) process gathers statistics every hour from every instance and creates an AWR snapshot and stores it in the SGA.
- D.  
The MMON process gathers statistics from its own instance and kicks off statistics collection from other instances every hour from other instances and creates an AWR snapshot which is written to the sysaux tablespace.

Explanation:

The Automatic Workload Repository (AWR) collects, processes, and maintains performance statistics for problem detection and self-tuning purposes. This data is both in memory and stored in the database. The gathered data can be displayed in both reports and views.

The statistics collected and processed by AWR include:

V\$SYS\_TIME\_MODEL and V\$SESS\_TIME\_MODEL views

V\$SESSTAT views

Activity

MMON (Memory Monitor) is a background process that gathers memory statistics (snapshots) stores this information in the AWR (automatic workload repository). MMON is also responsible for issuing alerts for metrics that exceed their thresholds.

Oracle Database Performance Tuning Guide

**Which two connect strings are correctly configured to connect to the ERP service?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Assume that you want to make a connection to the ERP service associated with a newly created RAC database called PROD on a cluster called cluster01 that consists of three nodes: node1, node2, and node3. Which two connect strings are correctly configured to connect to the ERP service?

- A.  
ERP=(DESCRIPTION=(LOAD\_BALANCE=on)(ADDRESS=(PROTOCOL=TCP)(HOST=cluster01-scan)(PORT=1521))(CONNECT\_DATA=(SERVICE\_NAME=ERP)))
- B.  
ERP=(DESCRIPTION=(LOAD\_BALANCE=on)(ADDRESS=(PROTOCOL=TCP)(HOST=node1-vip)(PORT=1521))(ADDRESS=(PROTOCOL=TCP)(HOST=node2-vip)(PORT=1521))(ADDRESS=(PROTOCOL=TCP)(HOST=node3-vip)(PORT=1521))(CONNECT\_DATA=(SERVICE\_NAME=ORCL)))
- C.  
ERP=(DESCRIPTION=(LOAD\_BALANCE=on)(ADDRESS=(PROTOCOL=TCP)(HOST=cluster01-scan1)(PORT=1521))(ADDRESS=(PROTOCOL=TCP)(HOST=cluster01-scan2)(PORT=1521))(ADDRESS=(PROTOCOL=TCP)(HOST=cluster01-scan3)(PORT=1521))(CONNECT\_DATA=(SERVICE\_NAME=ERP)))
- D.  
url='jdbc:oracle:thin:@(DESCRIPTION=(LOAD\_BALANCE=on)(ADDRESS=(PROTOCOL=TCP)(HOST=cluster01-scan)(PORT=1521))(CONNECT\_DATA=(SERVICE\_NAME=ERP)))'

Explanation:

Use Services with Client Applications

```
ERP=(DESCRIPTION=                ## Using the SCAN ##
(Load_Balance=on)
(ADDRESS=(PROTOCOL=TCP)(HOST=cluster01-scan)(PORT=1521))
(CONNECT_DATA=(SERVICE_NAME=ERP)))
```

```
ERP=(DESCRIPTION=                ## Using VIPs ##
(Load_Balance=on)
(ADDRESS=(PROTOCOL=TCP)(HOST=node1-vip)(PORT=1521))
(ADDRESS=(PROTOCOL=TCP)(HOST=node2-vip)(PORT=1521))
(ADDRESS=(PROTOCOL=TCP)(HOST=node3-vip)(PORT=1521))
(CONNECT_DATA=(SERVICE_NAME=ERP)))
```

```
url="jdbc:oracle:oci:@ERP"        ## Thick JDBC ##
```

```
url="jdbc:oracle:thin:@(DESCRIPTION= ## Thin JDBC ##
(Load_Balance=on)
(ADDRESS=(PROTOCOL=TCP)(HOST=cluster01-scan)(PORT=1521))
(CONNECT_DATA=(SERVICE_NAME=ERP)))"
```

D60488GC11

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 15 - 11

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 15 – 11

**Identify the default block Size used for the database.**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

When creating an Oracle Cluster database using DBCA the “Memory size (SGA and PGA) field is supplied on value of 2000 MB. Identify the default block Size used for the database.

- A.  
2 KB
- B.  
4 KB
- C.  
**8 KB**
- D.  
16 KB
- E.  
32 KB

Explanation:

Type of LimitLimit Value

Minimum2k.Must be a multiple of operating system physical block size

MaximumOperating system dependent, but never more than 32 KB

DBCA tab sizing

In this tab, you specify the smallest block size and the maximum number of operating system user processes that can simultaneously connect to the database.

In the Block Size list, enter the size in bytes or accept the default. Oracle Database data is stored in these blocks. One data block corresponds to a specific number of bytes of physical space on disk. While using pre-defined templates, this field is not enabled since the database will be created with the default block size of 8 KB. But while using the custom option, you can change block size. Selecting a block size other than the default 8 KB value requires advanced knowledge and should only be done when absolutely required.

Oracle Database 2 Day DBA

**Which two statements are true regarding the srvctl stop command?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Examine the following output:

```
[oracle@gr5153~]$srvctl add service -d RACDB -s erp -g pool1 -c uniform -y manual
[oracle@gr5153~]$srvctl start service -d RACDB -s ERP
[oracle@gr5153~]$scsvctl stat res ora.racdb.erp.svc NAME=ora.racdb.erp.svc
TYPE=ora.service.type
```

TARGET=ONLINE, ONLINE, ONLINE STATE=ONLINE on gr5118, ONLINE on gr5152, ONLINE on gr5153

```
[oracle@gr5153~]$ srvctl config database -d RACDB
```

Database unique name: RACDB

Database name: RACDB

Oracle home:/u01/app/oracle/product/11.2.0/dbhome\_1

Oracle user:oracle

Spfile:+DATA/RACDB/spfileRACDB.ora

Domain:

Start options:open

Stop options:immediate

Database role:PRIMARY

Management policy:AUTOMATIC

Server pools:POOL 1

Database instances:

Disk Groups:DATA, FRA

Services:ERP

Database is policy managed

```
$srvctl stop database -d RACDB -o immediate
```

Which two statements are true regarding the srvctl stop command?

A.

It will shut down all the instances of the RACDB database.

B.

It will shut down only the RACDB instance and the ERP service on the node on which the Command is executed.

C.

It will shut down only the database instance on the node on which the command is executed.

D.

It will stop the ERP service related to the RACDB database on all the nodes.

Explanation:

Shut down all Oracle RAC instances on all nodes. To shut down all Oracle RAC instances for a database, enter the following command, where db\_name is the name of the database:

```
srvctl stop database -d db_name
```

Stops a database, its instances, and its services. When the database later restarts, services with AUTOMATIC management start automatically but services with MANUAL management policy must be started manually.

Oracle Real Application Clusters Administration and Deployment Guide

Which statement is true regarding the srvctl stop instance command?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

You are managing a policy-managed database called RACDB.

Examine the output given below:

```
$ echo $ORACLE_SID
RACDB_1
```

```
$ sqlplus / as sysdba
SQL> SELECT * FROM V$ACTIVE INSTANCES;
```

```
INST NUMBER INST NAME
-----
1          gr7597:RACDB_1
2          gr7602:RACDB_2
3          gr7633:RACDB_3
```

You used these commands to add the service erp:

```
$ srvctl add service -d RACDB -s ERP -g POOL1 -c UNIFORM -y manual
```

```
$ srvctl start service -d RACDB -s ERP
```

```
$ crsctl stat res ora.racdb.erp.svc
NAME=ora.racdb.erp.svc
TYPE=ora.services.type
TARGET= ONLINE , ONLINE , ONLINE
STATE=ONLINE on gr7597, ONLINE on gr7602, ONLINE on gr7633
```

```
$ srvctl stop instance -d RACDB -i RACDB_1
```

Which statement is true regarding the srvctl stop instance command?

- A.  
It will shut down the racdb\_1 instance and the ERP service running on node gr7597.
- B.  
It will shut down the racdb\_1 instance, and the ERP service will fail over to an available Instance.
- C.  
It will shut down the racdb\_1 instance, but the ERP service will keep running on node gr7597
- D.  
It will shut down the racdb\_1 instance and the ERP service running on ih.il Instance However Oracle Clusterware can restart the ERP service running on node gr7597.
- E.  
It will shut down the racdb\_1 instance and stop the ERP service on all the nodes

Explanation:

Stops instances and stops any services running on specified instances, unless you specify the -f option. If you specify -f, then the services fail over to an available instance when the instance stops.

Syntax and Options

Use the srvctl stop instance command with the following syntax:

```
srvctl stop instance -d db_unique_name {[ -n node_name ]} [ -i "instance_name_list" ] [ -o stop_options ] [ -f ]
```

Oracle Real Application Clusters Administration and Deployment Guide

**Which three statements are true about the management of Oracle Services?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Which three statements are true about the management of Oracle Services?

- A.  
To manage workloads or a group of applications, you can define services for a particular application or a subset of an application's operations.
- B.  
Similar work can be categorized under services to ease workload management.
- C.  
Users who share a service should have different service-level requirements.
- D.  
Oracle Enterprise Manager or SRVCTL should be used to control cluster-managed services, not DBMS\_SERVICE.

Explanation:

Using Oracle Services

To manage workloads or a group of applications, you can define services that you assign to a particular application or to a subset of an application's operations. You can also group work by type under services. For example, online users can use one service, while batch processing can use another and reporting can use yet another service to connect to the database. Oracle recommends that all users who share a service have the same service level requirements. You can define specific characteristics for services and each service can represent a separate unit of work. There are many options that you can take advantage of when using services. Although you do not have to implement these options, using them helps optimize application performance.

Oracle does not recommend using the DBMS\_SERVICE package for use with services used by an Oracle RAC database. Use SRVCTL or Oracle Enterprise Manager to create database services for Oracle RAC.

Oracle Real Application Clusters Administration and Deployment Guide

**Which two recommendations may reduce this problem?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

You notice that there is a very high percentage of wait time for RAC database that has frequent insert operations.

Which two recommendations may reduce this problem?

- A.  
shorter transactions

- B. increasing sequence cache sizes
- C. using reverse key indexes
- D. **uniform and large extent sizes**
- E. **automatic segment space management**
- F. smaller extent sizes

Explanation:

Segments have High Water Mark (HWM) indicating that blocks below that HWM have been formatted. New tables or truncated tables [that is truncated without reuse storage clause], have HWM value set to segment header block. Meaning, there are zero blocks below HWM. As new rows inserted or existing rows updated (increasing row length), more blocks are added to the free lists and HWM bumped up to reflect these new blocks. HW enqueues are acquired in Exclusive mode before updating HWM and essentially HW enqueues operate as a serializing mechanism for HWM updates. Allocating additional extent with instance keyword seems to help in non-ASSM tablespace serialization of data blocks in the buffer cache due to lack of free lists, free list groups, transaction slots (INTRANS), or shortage of rollback segments. This is particularly common on INSERT-heavy applications, in applications that have raised the block size above 8K, or in applications with large numbers of active users and few rollback segments. Use automatic segment-space management (ASSM) and automatic undo management to solve this problem.

HW enqueue The HW enqueue is used to serialize the allocation of space beyond the high water mark of a segment.

If this is a point of contention for an object, then manual allocation of extents solves the problem.

**What is the most prominent problem indicated by the above output?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Examine the following details for a six-Instance RAC database:

| Top 5 Timed Events |         |          |               |                  |            |
|--------------------|---------|----------|---------------|------------------|------------|
| Event              | Waits   | Time (s) | Avg wait (ms) | %Total Call Time | Wait Class |
| gc buffer busy     | 177,315 | 29,021   | 164           | 24.3             | Cluster    |
| gc cr block busy   | 110,348 | 5,703    | 52            | 4.8              | Cluster    |
| gc cr block lost   | 4,272   | 4,953    | 1159          | 4.1              | Cluster    |
| cr request retry   | 6,316   | 4,668    | 739           | 3.9              | Other      |
| log file sync      | 6,038   | 4,572    | 174           | 2.7              | Commit     |

What is the most prominent problem indicated by the above output?

- A. high input/output (I/O) delays
- B. CPU saturation and memory depletion
- C. large number of untuned queries from one of the nodes
- D. **misconfigured or faulty interconnect**

Explanation:

In Oracle RAC environments, RDBMS gathers global cache work load statistics which are reported in STATSPACK, AWRs and GRID CONTROL. Global cache lost blocks statistics (“gc cr block lost” and/or “gc current block lost”) for each node in the cluster as well as aggregate statistics for the cluster represent a problem or inefficiencies in packet processing for the interconnect traffic. These statistics should be monitored and evaluated regularly to guarantee



efficient interconnect Global Cache and Enqueue Service (GCS/GES) and cluster processing. Any block loss indicates a problem in network packet processing and should be investigated.

Symptoms:

Primary:

Secondary:

Changes

As explained above, Lost blocks are generally caused by unreliable Private network. This can be caused by a bad patch or faulty network configuration or hardware issue.

Cause

In most cases, gc block lost has been attributed to (a) A missing OS patch (b) Bad network card (c) Bad cable (d) Bad switch (e) One of the network settings.

Oracle Metalinkgc block lost diagnostics [ID 563566.1]

### What does a high gc current block busy event value indicate?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

What does a high gc current block busy event value indicate?

A.

Access to cached data blocks was delayed because they were busy either in the remote or the local cache.

B.

A large number of requested blocks were not cached In any instance.

C.

Asynchronous input/output (I/O) is disabled.

D.

Delay in processing has occurred In the GCS, caused by CPU saturation and would have to be solved by additional CPUs and load-balancing.

#### Explanation:

The gc current block busy wait event indicates that the access to cached data blocks was delayed because they were busy either in the remote or the local cache. This could be caused by any of the following:

Oracle Real Application Clusters Administration and Deployment Guide

### Which inferences is correct?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Examine the following details from the AWR report for your three-instance RAC database:

| To 5 Timed Events       |         |          |               |                  |            |
|-------------------------|---------|----------|---------------|------------------|------------|
| Event                   | Waits   | Time (s) | Avg wait (ms) | %Total Call Time | Wait Class |
| CPU Time                | 4,580   | 65.4     |               |                  |            |
| Log file sync           | 276,281 | 1,501    | 5             | 21.4             | Commit     |
| Log file parallel write | 298,045 | 923      | 3             | 13.2             | System I/O |
| gc current block 3-way  | 605,628 | 631      | 1             | 9.0              | Cluster    |
| gc cr block 3-way       | 514,218 | 533      | 1             | 7.6              | Cluster    |

Which inferences is correct?

A.

There are a large number of requests for cr blocks or current blocks currently in progress.

B.

Global cache access is optimal without any significant delays.

C.

The log file sync waits are clue to cluster interconnect latency.

D.

To determine the frequency of two-way block requests you must examine other events In the report.

Explanation:

Analyzing Cache Fusion Transfer Impact Using GCS Statistics

This section describes how to monitor GCS performance by identifying objects read and modified frequently and the service times imposed by the remote access. Waiting for blocks to arrive may

constitute a significant portion of the response time, in the same way that reading from disk could increase the block access delays, only that cache fusion transfers in most cases are faster than disk access latencies.

The following wait events indicate that the remotely cached blocks were shipped to the local instance without having been busy, pinned or requiring a log flush:

gc current block 2-way

gc current block 3-way

gc cr block 2-way

gc cr block 3-way

The object statistics for gc current blocks received and gc cr blocks received enable quick identification of the indexes and tables which are shared by the active instances. As mentioned earlier, creating an ADDM analysis will, in most cases, point you to the SQL statements and database objects that could be impacted by interinstance contention.

Any increases in the average wait times for the events mentioned in the preceding list could be caused by the following occurrences:

High load: CPU shortages, long run queues, scheduling delays

Misconfiguration: using public instead of private interconnect for message and block traffic

If the average wait times are acceptable and no interconnect or load issues can be diagnosed, then the accumulated time waited can usually be attributed to a few SQL statements which need to be tuned to minimize the number of blocks accessed.

Oracle Real Application Clusters Administration and Deployment Guide  
11g Release 2 (11.2)

Which statement is true regarding the redo log files for the instance?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

You are managing a policy-managed Real Application Cluster (RAC) database, RACDB with Oracle Managed Files (OMF) enabled. Examine the following outputs:

```
[oracle@gr5153 $] srvctl config database -d RACDB
Database unique name: RACDB
Database name: RACDB
Oracle home: /u01/app/oracle/product/11.2.0/dbhome_1
Oracle user: oracle
Spfile +DATA/RACDB/spfileRACDB.ora
Domain:
Start options: open
Stop options: immediate
Database role: PRIMARY
Management policy: AUTOMATIC
Server pools: POOL1
Database instances:
Disk Groups: DATA,FRA
Services:
Database is policy managed.
```

```
SQL> SELECT * FROM v$active_instances;

INST NUMBER INST_NAME
-----
1          gr5118:RACDB_1
2          gr5152:RACDB_2
3          gr5153:RACDB_3
```

```
[oracle@gr5153 ~]$ crsctl stat serverpool ora.POOL1
NAME=ora.POOL1
ACTIVE_SERVERS=gr5118 gr5152 gr5153
```

```
[oracle@gr5153 ~]$ srvctl config srvpool -g POOL1
Server pool name: POOL1
Importance: 0, Min: 0, Max: 3
Candidate server names:
```

You added a new server to the server pool using the following command:

```
[oracle@gr5153 ~]$ srvctl modify srvpool -g POOL1 -i 0 -l 2 -u 4 -n gr5119
```

The instance on the newly added node is started by the Oracle Clusterware. Automatic Storage Management (ASM) is used as the storage option.

Which statement is true regarding the redo log files for the instance?

- A.  
Oracle Clusterware automatically creates and enables a new thread of redo.
- B.  
You must create redo log groups for the newly added service.
- C.  
Redo log members will not be created automatically because storage option used is ASM.
- D.  
The newly added instance must save full online log groups tracked in the control file.

Explanation:

For a policy-managed database, when you add a new node to the cluster, it is placed in the Free pool by default. If you increase the cardinality of the database server pool, then an Oracle RAC instance is added to the new node, racnode3, and it is moved to the database server pool. No further action is necessary.

Add shared storage for the undo tablespace and redo log files.

If OMF is not enabled for your database, then you must manually add an undo tablespace and redo logs.

You should create redo log groups only if you are using administrator-managed databases. For policy-managed databases, if an instance starts due to a change in server pool cardinality, then Oracle Database automatically creates redo log files, enables a redo thread for the instance if there is not a redo thread allocated to that instance, and creates the undo tablespace if there is not an undo tablespace allocated to that instance. The database must be using Oracle Managed Files and Oracle ASM in this situation.

Oracle Database 2 Day + Real Application Clusters Guide

**Which three statements are true about services and Transparent Application Failover (TAF)?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Which three statements are true about services and Transparent Application Failover (TAF)?

- A.  
If TAF has been configured for a service, sessions using that service fail over to a surviving instance when an outage occurs.
- B.  
The TAF setting on a service can be NONE, BASIC, PRECONNECT, or POSTCONNECT, and overrides any TAF setting in the client connection definition.
- C.  
TAF can restart a query after failover has completed but for other statements, such as INSERT, UPDATE, or DELETE, the application must resubmit the transaction.
- D.  
The TAF setting for a client connection overrides any TAF setting in the service definition.
- E.  
Services simplify the deployment of TAF because by defining a TAF policy for a service, all connections using this service will automatically have TAF enabled.

Explanation:

Services and Transparent Application Failover

Services simplify the deployment of Transparent Application Failover (TAF).

You can define a TAF policy for a service and all connections using this service will automatically have TAF enabled.

The TAF setting on a service can be NONE, BASIC, or PRECONNECT and overrides any TAF setting in the client connection definition.

To define a TAF policy for a service, the srvctl utility can be used as shown below:

```
srvctl modify service -s gl.example.com -q TRUE -P
```

```
BASIC -e SELECT -z 180 -w 5 -j LONG
```

Where -z is the number of retries, -w is the delay between retry attempts and -j is the connection load balancing goal.

When Oracle Net Services establishes a connection to an instance, the connection remains open until the client closes the connection, the instance is shut down, or a failure occurs. If you configure TAF for the connection, then Oracle Database moves the session to a surviving instance when an outage occurs.

TAF can restart a query after failover has completed but for other types of transactions, such as

INSERT, UPDATE, or DELETE, the application must roll back the failed transaction and resubmit the transaction. You must re-execute any session customizations, in other words, ALTER SESSION statements, after failover has occurred. However, with TAF, a connection is not moved during normal processing, even if the workload changes over time.

D60488GC11

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 15 13

### What are the prerequisites for converting the single-instance database to an Oracle RAC database?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

You are managing a single-instance database and your company wants to convert the single instance database to an Oracle RAC database. You plan to use the rconfig utility to accomplish this task.

What are the prerequisites for converting the single-instance database to an Oracle RAC database?

- A.  
Oracle Clusterware and Oracle Database software are installed on all target nodes.
- B.  
The Oracle Database binary is enabled for Oracle RAC on all target nodes.
- C.  
**The database being converted has been backed up successfully.**
- D.  
Oracle Clusterware is stopped.
- E.  
Automatic Storage Management (ASM) must be enabled and the ASM disk should be available on all nodes.

#### Explanation:

##### C.3.1 Prerequisites for Converting to Oracle RAC Databases

Before you convert a single-instance database to an Oracle RAC database, ensure that the following conditions are met for each cluster node that you intend to make an Oracle RAC database node:

Oracle Clusterware 11g release 2 (11.2) is installed, configured, and running.

Oracle RAC 11g release 2 (11.2) software is installed.

The Oracle binary is enabled for Oracle RAC.

Shared storage, either Oracle Cluster File System or Oracle ASM, is available and accessible from all nodes.

User equivalence exists for the oracle account.

If you intend to use Oracle Enterprise Manager, then all Oracle Enterprise Manager agents are configured and running, and are configured with cluster and host information.

You have backed up your existing database.

Oracle Real Application Clusters Installation Guide  
11g Release 2 (11.2) for Linux and UNIX

### What is the result of this command?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

The DBA has executed the command:

```
srvctl add service -d PROD -s BATCH -g SP2 -c singleton -y manual
```

What is the result of this command?

- A.  
Creates a singleton service named batch using the server pool sp2 with a manual service management policy.
- B.  
**Creates a singleton service named batch for the prod database by using only server pool SP2 with a manual service management policy.**
- C.  
Creates a singleton service named batch, which runs by default in the prod database by using the server pool SP2 with a manual service management policy, but can fail over to other databases.
- D.  
Creates a singleton service, which is only used for batch workloads for the prod database by

using only server pool sp2 with a manual service management policy.

Explanation:

srvctl add service

Adds services to a database and assigns them to instances. If you have multiple instances of a cluster database on the same node, then always use only one instance on that node for all of the services that node manages.

| Option                   | Description  |
|--------------------------|--|
| -d db_unique_name        | Unique name for the database.  |
| -s service_name          | The service_name.service_domain should be unique within the cluster unless you want to spread connections across multiple databases that offer the same service. If you do specify the service domain as part of the service name (such as sales.foo.com), then the DB_DOMAIN database attribute is appended to the service name.  |
| -c {UNIFORM   SINGLETON} | The cardinality of the service, either UNIFORM (offered on all instances in the server pool) or SINGLETON (runs on only one instance at a time).<br><u>Notes:</u><br>This option is available only with Oracle RAC and only for policy-managed databases.<br>For policy-managed Oracle RAC One Node databases, all services must be SINGLETON.   |
| -g server_pool           | The name of a server pool used when the database is policy managed.<br><u>Note:</u> This option is available only with Oracle RAC and only for policy-managed databases.   |
| -y {AUTOMATIC   MANUAL}  | Service management policy.<br>If AUTOMATIC (the default), then the service is automatically started upon restart of the database, either by a planned restart (with SRVCTL) or after a failure. Automatic restart is also subject to the service role, however (the -l option).<br><br>If MANUAL, then the service is never automatically restarted upon planned restart of the database (with SRVCTL). A MANUAL setting does not prevent Oracle Clusterware from monitoring the service when it is running and restarting it if a failure occurs. |

Oracle Real Application Clusters Administration and Deployment Guide  
11g Release 2 (11.2)

What are two possible reasons for these wait events?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

You are managing a policy-managed three-instance RAC database. You ran database ADDM for the database and noticed gc current block congested and gc cr block congested waits. What are two possible reasons for these wait events?

A.

The wait events indicate a delay in processing has occurred in the Global Cache Services (GCS), which is usually caused by high load.

B.

The wait times indicate that the blocks must wait after initiating a gc block request, for the round trip from the start of the wait until the blocks arrive.

C.

The wait events indicate that there is block contention resulting in multiple requests for access to local blocks.

D.

The wait events indicate that the local instance making the request for current or consistent read blocks was waiting for logical I/O from its own buffer cache at the same time.

Explanation:

Load-Related Wait Events

The main wait events for load-related waits are:

The load-related wait events indicate that a delay in processing has occurred in the GCS, which is usually caused by high load, CPU saturation and would have to be solved by additional CPUs, load-balancing, off loading processing to different times or a new cluster node. For the events mentioned, the wait time encompasses the entire round trip from the time a session starts to wait after initiating a block request until the block arrives

Oracle Real Application Clusters Administration and Deployment Guide

Which two statements are true regarding the Average Active Sessions chart on the performance page

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Which two statements are true regarding the Average Active Sessions chart on the performance

page

A.  
It shows active aggregate wait class statistics across all the instances in the cluster database.

B.  
It shows active aggregate wait class per instance for all the instances in the cluster database.

C.  
If the Average Active Sessions chart displays a large number of sessions waiting, indicating internal contention, but the throughput is high, then the situation may be acceptable.

D.  
If the Average Active Sessions chart displays a large number of sessions waiting, indicating internal contention, but the throughput is low, then the CPU needs to be upgraded

Explanation:

Chart for Average Active Sessions

The Average Active Sessions chart in the Cluster Database Performance page shows potential problems inside the database. Verifying the Interconnect Settings for Oracle RAC Categories, called wait classes, show how much of the database is using a resource, such as CPU or disk I/O. Comparing CPU time to wait time helps to determine how much of the response time is consumed with useful work rather than waiting for resources that are potentially held by other processes. Compare the peaks on the Average Active Sessions chart with those on the Database Throughput charts. If the Average Active Sessions chart displays a large number of sessions waiting, indicating internal contention, but throughput is high, then the situation may be acceptable.  
Oracle Database 2 Day + Real Application Clusters Guide

Which two recommendation would you make to reduce this problem?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

You notice that there is a very high percentage of wait time for the gc current split event in your RAC database that has frequent insert operations.  
Which two recommendation would you make to reduce this problem?

- A.  
shorter transactions
- B.  
using hash partitioned global indexes
- C.  
uniform and large extent sizes
- D.  
automatic segment space management
- E.  
smaller extent sizes
- F.  
increasing sequence cache sizes

Which, process will read the redo log of the failed instance to identify the database blocks...?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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

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

**Which two statements describe the reasons for the active versions on the nodes of the cluster?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

You have two administrator-defined server pools on your eight-node cluster called OLTP and DSS.

Hosts RACNODE3, RACNODE4, and RACNODE5 are currently assigned to the DSS Pool.

Hosts RACNODE6, RACNODE7, and RACNODE8 are assigned to the OLTP Pool.

Hosts RACNODE1 and RACNODE2 are assigned to the Generic pool.

You are patching the Oracle Grid Infrastructure in a rolling fashion for your cluster and you have completed patching nodes RACNODE3, RACNODE4, RACNODE5, and RACNODE6, but you have not patched nodes RACNODE1 and RACNODE2.

While examining the status of RACNODE2 software, you get this output:

```
$ crsctl query crs softwareversion
```

```
Oracle Cluster-ware version on node [RACNODE2] is [11.2.0.2.0]
```

```
$ crsctl query crs activeversion
```

```
Oracle Clusterware active version on node [RACNODE2] is [11.2.0.1.0]
```

Which two statements describe the reasons for the active versions on the nodes of the cluster?

A.

The active version is 11.2.0.2.0 on RACNODE3, RACNODE4, and RACNODE5 because all the nodes in the DSS server pool have the same installed version.

B.

**The active version is 11.2.0.1.0 on RACNODE6, RACNODE7, and RACNODE8 because some nodes in the cluster still have version 11.2.0.1.0 installed.**

C.

The active version is 11.2.0.1.0 on RACNODE6, RACNODE7, and RACNODE8 because some nodes in the OLTP Pool still have version 11.2.0.1.0 installed.

D.

**The active version is 11.2.0.1.0 on RACNODE3, RACNODE4, and RACNODE5 because some nodes in the cluster still have version 11.2.0.1.0 installed.**

Explanation:

```
crsctl query crs softwareversion
```

Use the crsctl query crs softwareversion command to display latest version of the software that has been successfully started on the specified node.

```
crsctl query crs activeversion
```

Use the crsctl query crs activeversion command to display the active version of the Oracle

Clusterware software running in the cluster. During a rolling upgrade, however, the active version is not advanced until the upgrade is finished across the cluster, until which time the cluster operates at the pre-upgrade version.

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

**Which statement indicates the proper time to change the compatible.rdbms disk group attribute to 11.2?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

The disk groups on the current ASM instance at version 11.2 were configured to support a version 10.2 database instance. The 10.2 instance has the COMPATIBLE parameter defined as 10.2.0. The compatible.asm attribute is set to 11.2 for each disk group. The database has been upgraded to 11.2. Which statement indicates the proper time to change the compatible.rdbms disk group attribute to 11.2?

- A.  
Change the disk group attribute after the database instance COMPATIBLE parameter is upgraded to 11.2.
- B.  
Change the disk group attribute after the database instance is started with the 11.2 software.
- C.  
Change the disk group attribute after the database instance optimizer\_features\_enabled parameter is set to 11.2.
- D.  
Change each disk group after the 11.2 features are required for use on the disk group.
- E.  
Never, upgrading the attribute is not reversible.

Explanation:

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.

Before advancing the COMPATIBLE.RDBMS attribute, ensure that the values for the COMPATIBLE initialization parameter for all of the databases that access the disk group are set to at least the value of the new setting for COMPATIBLE.RDBMS.

Oracle Automatic Storage Management Administrator's Guide  
11g Release 2 (11.2)

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

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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.

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.1.0.2 to 11.2.0.3 – Release.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.2.0.1.0 and later [Release.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 gnp 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_clsss: 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.

```

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

### Choose three areas that are contained within the ASM SGA.

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

The System Global Area (SGA) for the ASM instance contains distinct memory areas. Choose three areas that are contained within the ASM SGA.

- A. **Shared Pool**
- B. Buffer Cache
- C. Log Buffer
- D. **Large Pool**
- E. **ASM Cache**
- F. Streams Pool

#### Explanation:

The SGA in an ASM instance is different in memory allocation and usage than the SGA in a database instance. The SGA in the ASM instance is divided into four primary areas as follows:

Shared Pool: Used for metadata information

Large Pool: Used for parallel operations

ASM Cache: Used for reading and writing blocks during rebalance operations

Free Memory: Unallocated memory available

D60488GC11

**Which four statements regarding IPMI configuration are true?**Posted by seenagape on [March 12, 2013](#)[No comments](#)

During the installation of Oracle Grid Infrastructure 11.2, it is possible to incorporate failure isolation support by using Intelligent Platform Management Interface (IPMI).

Which four statements regarding IPMI configuration are true?

- A.  
Each cluster member node requires a Baseboard Management Controller (BMC), which runs a firmware compatible with IPMI version 1.5.
- B.  
The ethernet port on each cluster node used by BMC must be connected to the IPMI management network.
- C.  
Each node in the cluster must have an IPMI driver installed.
- D.  
The cluster requires a dedicated network specifically for IPMI.
- E.  
If you intend to use IPMI, you must provide an administration account username and password when prompted during installation.

Explanation:

## 2.13.1 Requirements for Enabling IPMI

You must have the following hardware and software configured to enable cluster nodes to be managed with IPMI:

Have IPMI Configuration completed and have IPMI administrator account information

If you intend to use IPMI, then ensure BMC interfaces are configured, and have an administration account username and password to provide when prompted during installation.

For nonstandard installations, if you must change configuration on one or more nodes after installation (for example, if you have different administrator usernames and passwords for BMC interfaces on cluster nodes), then decide if you want to reconfigure the BMC interface, or modify IPMI administrator account information after installation.

Oracle Grid Infrastructure Installation Guide

11g Release 2 (11.2) for Linux

**What could have caused the DATA2 diskgroup to be mounted based on the information in the V \$ASM\_DISKGROUP view?**Posted by seenagape on [March 12, 2013](#)[No comments](#)

mounted: DATA and DATA2, but the results of the SHOW PARAMETER ASM\_DISKGROUPS

command show only DATA as below:

```
SQL> SELECT GROUP_NUMBER, NAME, STATE, TYPE, OFFLINE_DISKS, VOTING_FILES
2> FROM V$ASM_DISKGROUP;
```

| GROUP<br>NUMBER | NAME  | STATE   | TYPE   | OFFLINE_DISKS | VOTING<br>FILES |
|-----------------|-------|---------|--------|---------------|-----------------|
| 1               | DATA2 | MOUNTED | EXTERN | 0             | Y               |
| 2               | DATA  | MOUNTED | EXTERN | 0             | n               |

```
SQL> SHOW PARAMETER ASM_DISKGROUPS
```

| NAME           | TYPE   | VALUE |
|----------------|--------|-------|
| asm_diskgroups | string | DATA  |

What could have caused the DATA2 diskgroup to be mounted based on the information in the V \$ASM\_DISKGROUP view?

- A.  
The DATA2 diskgroup contains the database files required to start a database instance that depends on this ASM instance.
- B.  
The DATA2 diskgroup contains the voting files required for this cluster.
- C.  
The DATA diskgroup has mirrored objects in the DATA2 diskgroup.

D.  
The DATA value in the parameter implies all diskgroup strings starting with data.

E.  
The DATA2 diskgroup contains the SPFILE needed to start the ASM instance.

Explanation:

Disk Groups Mounted at Startup

At startup, the Oracle ASM instance attempts to mount the following disk groups:

D60488GC11

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 7 – 5

**which interface are the SCAN and node VIPs (Virtual IP addresses) associated?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

ETH0 is the public interface on all your RAC nodes.  
ETH1 is the private interface on all your RAC nodes for use with the cluster interconnect.  
ETH2 is used on all the RAC nodes for accessing storage on a filer using network attached storage (NAS).

With which interface are the SCAN and node VIPs (Virtual IP addresses) associated?

A.  
ETH0 for the node VIPs and ETH1 for the SCAN VIPs

B.  
**ETH0 for both the node VIPs and the SCAN VIPs**

C.  
ETH1 for both the node VIPs and the SCAN VIPs

D.  
ETH1 for the node VIPs and ETH0 for the SCAN VIPs

E.  
with both ETH0 and ETH1 for both the node VIPs and the SCAN VIPs

Explanation:

2.7.2 IP Address Requirements

Before starting the installation, you must have at least two interfaces configured on each node:

One for the private IP address and one for the public IP address.

You can configure IP addresses with one of the following options:

Dynamic IP address assignment using Oracle Grid Naming Service (GNS). If you select this option, then network administrators assign static IP address for the physical host name and dynamically allocated IPs for the Oracle Clusterware managed VIP addresses. In this case, IP addresses for the VIPs are assigned by a DHCP and resolved using a multicast domain name server configured as part of Oracle Clusterware within the cluster. If you plan to use GNS, then you must have the following:

- A DHCP service running on the public network for the cluster
- Enough addresses on the DHCP to provide 1 IP address for each node's virtual IP, and 3 IP addresses for the cluster used by the Single Client Access Name (SCAN) for the cluster Static IP address assignment. If you select this option, then network administrators assign a fixed IP address for each physical host name in the cluster and for IPs for the Oracle Clusterware managed VIPs. In addition, domain name server (DNS) based static name resolution is used for each node. Selecting this option requires that you request network administration updates when you modify the cluster.

Note:

Oracle recommends that you use a static host name for all server node public hostnames.

Public IP addresses and virtual IP addresses must be in the same subnet.

Oracle only supports DHCP-assigned networks for the default network, not for any subsequent networks.

Oracle Grid Infrastructure Installation Guide

11g Release 2 (11.2) for Linux

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

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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

Explanation:

Interconnect Link Aggregation: Single Switch

Interconnect Link Aggregation: Multiswitch

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

**Which three statements are true about ASM performance and scalability?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Which three statements are true about ASM performance and scalability?

A.  
Disks in a diskgroup may be of varying speed and capacity without affecting performance.

B.  
Disks in a diskgroup should match in speed and capacity for best performance.

C.  
A database seldom needs more than two diskgroups: one for data and the other for a fast recovery area.

D.  
A database needs a separate diskgroup for each major tablespace, redo logs, and fast recovery area.

E.  
Multiple databases that use the same diskgroup make the best use of space.

F.  
Multiple databases that use a separate diskgroup for each database make the best use of space.

Explanation:

Performance, Scalability, and Manageability Considerations for Disk Groups

Create separate disk groups for database files and fast recovery area.

Disks in a disk group should have the same size and performance characteristics.

Allows the disk group to deliver consistent performance

Allows ASM to use disk space most effectively

Allows operations with different storage requirements to be matched with different disk groups effectively

Using separate disk groups for each database as opposed to having multiple databases in a disk group has various benefits and drawbacks.

Housing multiple databases in a single disk group affords the most efficient use of space.

However, any faults or maintenance that affects the disk group may affect many databases.

Separate disk groups provide greater isolation from the effects of a fault or maintenance operation.

However, to achieve this may consume more disk space and may require more disk group maintenance to balance disk resources.

D60488GC11

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 8 – 37

**What is the most efficient method to move the application and the VIP?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Your Apache-based application resource called MyApp has a VIP application resource definition and an application resource definition.

The application is administrator managed and is currently active on nodeRACNODE3. You want to move it to host RACNODE4, which is one of the HOSTING\_MEMBERS in the resource definition.

What is the most efficient method to move the application and the VIP?

A.

Run `crsctl relocate resource MyApp -n RACNODE4 -f`.

B.

Run `crsctl stop resource MyApp -n RACNODE3 -f` followed by `crsctl start MyApp -n RACNODE4 -f`.

C.

Run `crsctl stop resource MyApp -n RACNODE3 -f` followed by `crsctl relocate resource MyApp -n RACNODE4`.

D.

Run `crsctl relocate resource MyApp -n RACNODE4`.

Explanation:

Relocating Applications and Application Resources

Use the `crsctl relocate resource` command to relocate applications and application resources. For example, to relocate the Apache Web server application to a server named `rac2`, run the following command:

```
# crsctl relocate resource myApache -n rac2
```

Each time that the action program is called, the `crsctl relocate resource` command waits for the duration specified by the value of the `SCRIPT_TIMEOUT` resource attribute to receive notification of success or failure from the action program. A relocation attempt fails if:

The application has required resources that run on the initial server

Applications that require the specified resource run on the initial server

To relocate an application and its required resources, use the `-f` option with the `crsctl relocate resource` command. Oracle Clusterware relocates or starts all resources that are required by the application regardless of their state.

Oracle Clusterware Administration and Deployment Guide

11g Release 2 (11.2)

How would you start the provisioning daemon?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

You plan to use Enterprise Manager to locate and stage patches to your Oracle Home.

The software library has been configured to be downloaded to `/u01/app/oracle` and your “My Oracle Support” credentials have been entered.

You want to start the provisioning daemon in order to use the deployment procedure manager to view, edit, monitor, and run deployment procedures.

You want to start the provisioning daemon in order to use the deployment procedure manager to view, edit, monitor, and run deployment procedures.

How would you start the provisioning daemon?

How would you start the provisioning daemon?

A.

using `pafctl start`

B.

using `crsctl start paf`

C.

using `srvctl start paf`

D.

using `emctl start paf`

Explanation:

Starting the Provisioning Daemon

The provisioning daemon is started with:

```
$ pafctl start
```

Enter repository user password :

Enter interval [default 3]:

Provisioning Daemon is Up, Interval = 3

D60488GC11

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 4 26

How does the Oracle Grid Infrastructure administrator determine the location of the Oracle Clusterware voting disk?

How does the Oracle Grid Infrastructure administrator determine the location of the Oracle Clusterware voting disk?

- A.  
Run cat /etc/oracle/vote.loc from any node.
- B.  
Run srvctl query css votedisk from any node.
- C.  
**Run crsctl query css votedisk from any node.**
- D.  
Run select name, path from v\$votedisk from any RAC database instance of any database on the cluster.

Explanation:

Determining the Location of Oracle Clusterware Configuration Files  
 You can determine the location of the voting disk by using the crsctl query css votedisk command on any node.  
 D60488GC11  
 Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 3 7

**Which three are components of Oracle HA Framework for protecting third-party applications using Oracle Clusterware?**

Which three are components of Oracle HA Framework for protecting third-party applications using Oracle Clusterware?

- A.  
**resources**
- B.  
**action programs**
- C.  
voting disks
- D.  
**application VIPS**
- E.  
SCAN VIPS

Explanation:

Oracle Clusterware HA Components

| Component           | Definition   |
|---------------------|--|
| Resource            | An entity that Oracle Clusterware manages for HA such as an application  |
| Action Program      | A developed program that provides the logic for starting, stopping, and monitoring a resource                                |
| Privileges          | Access and usage privileges for a resource allowing it to run as a different user than the Cluster Ready Services (CRS) user |
| Resource Dependency | A relationship among resources or applications that implies operational ordering   |
| Application VIP     | A VIP that an application has a dependency with  |
| OCR                 | Storage mechanism for resource profiles, policies, and privileges  |

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 5 – 4

**What is the recommended procedure to shut down an ASM instance with a mounted ACFS file system at ...?**

What is the recommended procedure to shut down an ASM instance with a mounted ACFS file system at /u01/app/oracle/acfsmounts/vol1, where the volume name is VOL1 and the disk group name is DATA?

- A.  
Issue the synchronize command for the file system twice. Then stop the ASM Instance with the SQL\*Plus SHUTDOWN ABORT command.



B.  
Use the `advmutil dismount /all /DATA` command. Then stop the ASM instance with the `SQL*Plus SHUTDOWN IMMEDIATE` command.

C.  
Use the `/bin/umount /u01/app/oracle/acfsmounts/vol1` command. Then shut down the ASM instance with the `SQL*Plus SHUTDOWN IMMEDIATE` command.

D.  
Use the `SQL ALTER DISKGROUP VOLULE VOL1 DISMOUNT` command. Then stop the ASM instance with the `SQL*Plus SHUTDOWN IMMEDIATE` command.

E.  
Stop the ASM instance with the `SQL*Plus SHUTDOWN IMMEDIATE` command.

Explanation:

Oracle ACFS and Dismount or Shutdown Operations

It is important to dismount any active file system configured with an Oracle ADVM volume device file before an Oracle ASM instance is shutdown or a disk group is dismounted. After the file systems are dismounted, all open references to Oracle ASM files are removed and associated disk groups can be dismounted or the instance shut down.

Oracle Automatic Storage Management Administrator's Guide  
11g Release 2 (11.2)

**Which interconnect design is NOT supported for Oracle Clusterware?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

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

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)

**Which three statements are true about Clusterware resource debugging?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Which three statements are true about Clusterware resource debugging?

A.  
**The `crsctl` command can be used to set debugging for Clusterware resources by using the following syntax:  
`crsctl set log resource "resource_name:logging_level".`**

B.  
**Enabling logging for Clusterware resources can negatively affect cluster performance, so it should be used only when required.**

C.  
After collecting debugging data for a specific Clusterware resource, terminate the collection of debugging data by issuing the `crsctl set log resource "resource_name:stop"` command.

D.

Although the crsctl command can be used to dynamically affect logging for Clusterware resources, you can also configure resource debugging automatically by specifying a logging\_level clause in the Grid\_Home/log/hostname/admin/clscfg.ini file.

E.

Debugging cannot be set for user-defined resources.

Explanation:

Enabling Resource Debugging

Change the USR\_ORA\_DEBUG resource attribute to 1 for specific resources:

```
# crsctl set log res "ora.host01.vip:1"
```

After you capture all trace information, change the debug attribute back to 0:

```
# crsctl set log res "ora.host01.vip:0"
```

You can use an initialization file to configure debugging.

The initialization file name includes the name of the process that you are debugging (process\_name.ini).

The file is located in the <Grid\_Home>/log/host\_name/admin/ directory.

Oracle Support may request that you enable tracing to capture additional information for problem resolution with Oracle Clusterware resources. Because the procedures described here may affect performance, perform these activities only with the assistance of Oracle Support.

The initialization file name includes the name of the process that you are debugging (process\_name.ini). The file is located in the <Grid\_Home>/log/host\_name/admin/ directory. For example, the name for the CLSCFG debugging initialization file on node1 would be:

```
<Grid_Home>/log/node1/admin/clscfg.ini
```

```
D60488GC11
```

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 6 – 13

**Which two statements are true?**

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Various clients can access and manipulate ASM files. Which two statements are true?

A.

The DBMS\_FILE\_TRANSFER.COPY\_FILE procedure can move a database file from one ASM to another ASM, but not to an operating system file system.

B.

The ASMCMD cp command can move database files from a file system to ASM, but not from ASM to ASM.

C.

The SQL\*Plus command ALTER DISKGROUP orcl MOVE '+DATA/orcl/example01.dbf' to '+OLDDATA/orcl/example01.dbf' can move the example01 data file to a different diskgroup.

D.

The DBMS\_FILE\_TRANSFER.GET\_FILE procedure reads an ASM file from a remote machine and makes a local copy on an ASM or a file system.

E.

The ASMCMD rm command will delete ASM files and directories, but not database files on an operating system file system.

Explanation:

DBMS\_FILE\_TRANSFER

COPY\_FILE Procedure

This procedure reads a file from a source directory and creates a copy of it in a destination directory. The source and destination directories can both be in a local file system, or both be in an Automatic Storage Management (ASM) disk group, or between local file system and ASM with copying in either direction.

You can copy any type of file to and from a local file system. However, you can copy only database files (such as datafiles, tempfiles, controlfiles, and so on) to and from an ASM disk group.

GET\_FILE Procedure

This procedure contacts a remote database to read a remote file and then creates a copy of the file in the local file system or ASM. The file that is copied is the source file, and the new file that results from the copy is the destination file. The destination file is not closed until the procedure completes successfully.

Examples

```
CREATE OR REPLACE DIRECTORY df AS '+datafile' ;
```

```
GRANT WRITE ON DIRECTORY df TO "user";
```

```

CREATE DIRECTORY DSK_FILES AS '^t_work^';
GRANT WRITE ON DIRECTORY dsk_files TO "user";
– assumes that dbs2 link has been created and we are connected to the instance.
– dbs2 could be a loopback or point to another instance.
BEGIN
– asm file to an os file
– get an asm file from dbs1.asm/a1 to dbs2.^t_work^/oa5.dat
DBMS_FILE_TRANSFER.GET_FILE ( 'df', 'a1', 'dbs1', 'dsk_files', 'oa5.dat' );
– os file to an os file
– get an os file from dbs1.^t_work^/a2.dat to dbs2.^t_work^/a2back.dat
DBMS_FILE_TRANSFER.GET_FILE ( 'dsk_files', 'a2.dat', 'dbs1', 'dsk_files', 'a2back.dat' );
END ;
/

```

Oracle Database PL/SQL Packages and Types Reference

11g Release 2 (11.2)

ASMCMD

cp

Purpose

Enables you to copy files between Oracle ASM disk groups and between a disk group and the operating system.

You can use the cp command to:

Copy files from a disk group to the operating system

Copy files from a disk group to a disk group

Copy files from the operating system to a disk group

rm

Purpose

Deletes the specified Oracle ASM files and directories.

Oracle Automatic Storage Management Administrator's Guide

11g Release 2 (11.2)

### Which three steps can be performed to determine the health of the OCR files and voting disks?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

The Oracle Grid Infrastructure administrator wants to perform daily checks on the integrity of the Clusterware files. Which three steps can be performed to determine the health of the OCR files and voting disks?

A.

Run ocrcheck to verify the health of the OCR file only from the master node, which performs all input/output (I/O) operations on the OCR.

B.

Run `grep voting <grid_home>/log/<hostname>/cssd/ocssd.log` from any cluster node.

C.

Run `cluvfy comp ocr -n all` from any cluster node.

D.

Run `ocrcheck` to verify the health of the OCR file from any cluster node.

E.

Run `votecheck` to verify the health of the voting disk from any cluster node.

#### Explanation:

Checking the Integrity of Oracle Clusterware

Configuration Files

The following techniques are used to validate the integrity of Oracle Cluster configuration files.

Check the `ocssd.log` for voting disks issues.

```
$ grep voting <grid_home>/log/<hostname>/cssd/ocssd.log
```

Use the `cluvfy` utility or the `ocrcheck` command to check the integrity of the OCR.

```
$ cluvfy comp ocr n all -verbose
```

```
$ ocrcheck
```

```
D60488GC11
```

Oracle 11g: RAC and Grid Infrastructure Administration

### Which two tools are available to do this?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

You decide to create a snapshot of your ACFS for backup and recovery purposes. Which two tools

are available to do this?

- A.  
ASMCMD
- B.  
SQL\*Plus
- C.  
ACFSUTIL
- D.  
ADVMUTIL
- E.  
Enterprise Manager

Explanation:

About Oracle ACFS Snapshots

Oracle ACFS snapshots are administered with the acfsutil snap commands. For information about the acfsutil snap commands, refer to “acfsutil snap create”, “acfsutil snap delete”, and “acfsutil snap info”.

You can also manage Oracle ACFS snapshots with Oracle Enterprise Manager. For information about using Oracle Enterprise Manager, see “Managing Oracle ACFS Snapshots with Oracle Enterprise Manager”.

Oracle Automatic Storage Management Administrator’s Guide  
11g Release 2 (11.2)

Which three statements can be verified by from the information provided?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Directory structure in ASM has a default of +<diskgroup>/<database>/<filetype>/<filename>.

The EXAMPLE tablespace was created with:

```
CREATE TABLESPACE EXAMPLE DATAFILE '+DATA/ORCL/oradata/example_01.dbf'
```

The asmcmd ls command shows:

```
ASMCMD>ls -l +DATA/ORCL/oradata/*
```

```
Type Redund Striped Type Sys Name
```

```
N example_01.dbf => +DATA/ORCL/DATAFILE/EXAMPLE.264.695918623
```

Which three statements can be verified by from the information provided?

- A.  
The ASM file +DATA/ORCL/oradata/example\_01.dbf has an alias of +DATA/ORCL/DATAFILE/EXAMPLE.264.695918623.
- B.  
The ASM system file +DATA/ORCL/DATAFILE/EXAMPLE.264.695918623 has an alias of +DATA/ORCL/oradata/example\_01.dbf.
- C.  
The default path indicates that the file is an oradata file belonging to the ORCL database.
- D.  
The default path indicates that the file is a data file belonging to the ORCL database.
- E.  
+DATA/ORCL/oradata/example\_01.dbf is the name of the file in V\$DBFILE in the database instance.
- F.  
+DATA/ORCL/DATAFILE/EXAMPLE.264.695918623 is the name of the file in V\$DBFILE in the database instance.

Explanation:

Fully Qualified ASM File Names

Every file created in ASM gets a system-generated file name, known as the fully qualified file name.

You cannot set the fully qualified file name.

ASM guarantees uniqueness within the ASM environment.

The fully qualified file name is used in database views that display Oracle Database file names.

D60488GC11

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 9 – 11

Choose the two correct commands that will successfully create an ASM cluster file system (ACFS) and register the mount point on a Linux/UNIX system.

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Choose the two correct commands that will successfully create an ASM cluster file system (ACFS) and register the mount point on a Linux/UNIX system.

- A.  
mkfs -t asm /dev/asm/testvol-461 -n "testvol"
- B.  
mkfs -t acfs /dev/asm/testvol-461 "testvol"
- C.  
crsctl -r mountpoint /dev/asm/testvol-461 /u01/app/grid/acfsmounts/acfs\_testvol
- D.  
acfsutil registry -a -f /dev/asm/testvol-461 /u01/app/grid/acfsmounts/acfs\_testvol

Explanation:

mkfs

Purpose

Creates an Oracle ACFS file system.

Syntax and Description

mkfs -t acfs -h

mkfs [-v] [-f]-t acfs [-b blocksize] [-n name ] volume\_device [size]

Table 13-3 Options for the Linux mkfs command

Option Description

- t acfs Specifies the type of file system on Linux. acfs designates the Oracle ACFS type.
- v Specifies verbose mode. The progress is displayed as the operation occurs.
- n name Specifies the name for the file system. A name can be a maximum of 64 characters. acfsutil info fs returns the name if one was specified.
- f Specifies the force option. This action creates the file system even if there is an existing Oracle ACFS on the volume device, although only if the file system is dismounted.
- F Specifies that the file system should be created on the original file system. Use this option with caution.
- h Displays the usage help text and exits
- h blocksize The default block size is 4K and this is the only size supported in 11g Release 2 (11.2)
- size Specifies the size that the file system should consume on the named device. The quantity specified can be in units of K (kilobytes), M (megabytes), G (gigabytes), T (terabytes), or P (petabytes). If a unit is not specified, the default is bytes.
- If the size specified is not a multiple of the block size, then the value is rounded up to the closest multiple. If this option is not specified, the entire device is consumed.
- volume\_device Specifies an Oracle ADVM device file that is to be formatted.

acfsutil registry

Purpose

Registers an Oracle ACFS file system with the Oracle ACFS mount registry.

Syntax and Description

acfsutil registry -h

acfsutil registry

acfsutil registry -a [-f] [-n { nodes|all } ]

[-o moptions] device mount\_point

acfsutil registry -d {device | mount\_point}

acfsutil registry -l [device | mount\_point]

acfsutil registry -m device

acfsutil registry -h displays help text and exits.

Oracle Automatic Storage Management Administrator's Guide

11g Release 2 (11.2)

Which three statements are true about Oracle Clusterware component log files?

Posted by seenagape on [March 12, 2013](#)

[No comments](#)

Which three statements are true about Oracle Clusterware component log files?

- A.  
Oracle RAC uses a common unified log directory structure to store all Oracle Clusterware component log files.
- B.  
The consolidated directory structure simplifies diagnostic information collection and assists

during data retrieval and problem analysis.

C.  
The Clusterware and Database log files are stored under the same unified directory structure.

D.  
The location of the log directory structure is <Grid\_Home>/log/<hostname>.

E.  
The log directory structure is the same on UNIX and Linux systems, but different on Windows platforms.

Explanation:

Oracle Clusterware Main Log Files

Oracle Clusterware uses a unified log directory structure to consolidate the Oracle Clusterware component log files. This consolidated structure simplifies diagnostic information collection and assists during data retrieval and problem analysis.

The slide shows you the main directories used by Oracle Clusterware to store its log files:

CRS logs are in <Grid\_Home>/log/<hostname>/crsd/. The crsd.log file is archived every 10 MB (crsd.i01, crsd.i02, ...).

CSS logs are in Grid\_HOME /log/<hostname>/cssd/. The cssd.log file is archived every 20 MB (cssd.i01, cssd.i02, ...).

EVM logs are in <Grid\_Home>/log/<hostname>/evmd.

SRVM (srvctl) and OCR (ocrdump, ocrconfig, ocrcheck) logs are in <Grid\_Home>/log/<hostname>/client/ and \$ORACLE\_HOME/log/<hostname>/client/.

Important Oracle Clusterware alerts can be found in alert<nodename>.log in the <Grid\_Home>/log/<hostname> directory.

D60488GC11

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 2 – 11

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

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

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

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

**What should be done to verify that the two new nodes called RACNODE5, RACNODE6 are physically connected?**Posted by seenagape on [March 11, 2013](#)[No comments](#)

The original cluster on RACNODE1, RACNODE 2, RACNODE 3, and RACNODE4 had the Oracle Grid Infrastructure installed to support Grid Naming Service (GNS).

What should be done to verify that the two new nodes called RACNODE5, RACNODE6 are physically connected?

- A.  
cluvfy stage -post crsinst -n RACNODE5,RACNODE6 -verbose
- B.  
**cluvfy stage -post nodeadd -n RACNODE5,RACNODE6 -verbose**
- C.  
cluvfy stage -post hacfg -verbose
- D.  
cluvfy stage -post nodeadd -n all  verbose

**Explanation:**

Completing OUI Silent Node Addition

Perform integrity checks on the cluster.

```
[grid@host01]$ cluvfy stage post nodeadd n host03 -verbose
```

```
D60488GC11
```

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 4 – 8

**Which three describe how often OCR backups are taken and how many are retained, according to the default backup schedule?**Posted by seenagape on [March 11, 2013](#)[No comments](#)

Which three describe how often OCR backups are taken and how many are retained, according to the default backup schedule?

- A.  
every 12 hours and the Clusterware keeps the last four copies
- B.  
**every 4 hours and the Clusterware keeps the last three copies**
- C.  
**every day and the Clusterware keeps the last two copies**
- D.  
every day and the Clusterware keeps the last three copies
- E.  
**every week and the Clusterware keeps the last two copies**

**Explanation:**

Locating the OCR Automatic Backups

The OCR is backed up automatically.

Only one node performs the backup.

To determine the node and location of the backup:

```
$ ocrconfig -showbackup auto
```

```
host02 2009/07/28 12:20:42 /u01/app/.../cdata/cluster01/backup00.ocr
```

```
host02 2009/07/28 08:20:41 /u01/app/.../cdata/cluster01/backup01.ocr
```

```
host02 2009/07/28 04:20:40 /u01/app/.../cdata/cluster01/backup02.ocr
```

```
host02 2009/07/27 16:20:37 /u01/app/.../cdata/cluster01/day.ocr
```

```
host02 2009/07/28 00:20:39 /u01/app/.../cdata/cluster01/week.ocr
```

Files could be spread across nodes due to outages.

The backup frequency and retention policies are:

Every four hours: CRS keeps the last three copies.

At the end of every day: CRS keeps the last two copies.

At the end of every week: CRS keeps the last two copies.

```
D60488GC11
```

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 3 – 11

**How can this be done?**

The Oracle Grid Infrastructure administrator decides to make more copies of the voting disks that are currently stored in the ASM disk group +VOTE. How can this be done?

- A.  
by running `crsctl add css votedisk <path_to_new_voting_disk>` to make a copy to a shared location on a shared device or file system
- B.  
by running `crsctl add css votedisk +VOTE`, thereby adding another copy of the voting disk to the +VOTE disk group
- C.  
by running `srvctl replace votedisk +asm_disk_group` on another disk group that has greater redundancy, thereby causing additional copies to be created
- D.  
**by running `crsctl replace votedisk +asm_disk_group` on another disk group that has greater redundancy, thereby causing additional copies to be created**

Explanation:

Storing Voting Disks on Oracle ASM

Using the `crsctl replace votedisk` command, you can move a given set of voting disks from one Oracle ASM disk group into another, or onto a certified file system. If you move voting disks from one Oracle ASM disk group to another, then you can change the number of voting disks by placing them in a disk group of a different redundancy level as the former disk group.

Notes:

You cannot directly influence the number of voting disks in one disk group.

You cannot use the `crsctl add | delete votedisk` commands on voting disks stored in Oracle ASM disk groups because Oracle ASM manages the number of voting disks according to the redundancy level of the disk group.

You cannot add a voting disk to a cluster file system if the voting disks are stored in an Oracle ASM disk group. Oracle does not support having voting disks in Oracle ASM and directly on a cluster file system for the same cluster at the same time.

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

**Which two additional steps are required when deleting a node from a cluster in this situation?**

You plan to remove a node called RACNODE4 from a four-node cluster.

The cluster is running Oracle Grid Infrastructure 11g Release 2 and there is also a database home with cluster databases running Oracle Database version 11g Release 1. This older home supports different applications that have not been certified to run on the latest release.

The cluster databases for Oracle Database 11g Release 1 each had four Instances at one time, but currently only the first two nodes, RACNODE1 and RACNODE2, are running instances of those databases.

Which two additional steps are required when deleting a node from a cluster in this situation?

- A.  
Run `crsctl pin css -n RACNODE1, RACNODE2` as root if the nodes are unpinned.
- B.  
Run `crsctl pin css -n RACNODE4` as the Grid Infrastructure owner if RACNODE4 is unpinned.
- C.  
**Run `olsnodes -t -n` to list all cluster nodes that are pinned.**
- D.  
**Run `crsctl unpin css - RACNODE4` as the Grid Infrastructure owner if RACNODE4 is pinned.**
- E.  
Run `crsctl unpin css - RACNODE4` as root if RACNODE 4 is pinned.
- F.  
Run `crsctl nodelist -t -n` to list all cluster nodes that are pinned.

Explanation:

Deleting a Cluster Node on Linux and UNIX Systems

1. Ensure that `Grid_home` correctly specifies the full directory path for the Oracle Clusterware



home on each node, where Grid\_home is the location of the installed Oracle Clusterware software.

2. Run the following command as either root or the user that installed Oracle Clusterware to determine whether the node you want to delete is active and whether it is pinned:

```
$ olsnodes -s -t
```

If the node is pinned, then run the crsctl unpin css command. Otherwise, proceed to the next step.

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

### What is true about the server pool to which the node will be assigned?

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

You have defined two server pools.

One is called OLTP with MIN\_SIZE=3, MAX\_SIZE=5 and IMPORTANCE=3.

The other is called DSS with MIN\_SIZE=2, MAX\_SIZE=4 and IMPORTANCE=5.

Your cluster currently has four nodes with two nodes in each of the pools. A fifth node is added to the cluster. What is true about the server pool to which the node will be assigned?

A.

The node will have the server assigned to the DSS pool because the current size is less than MAX\_SIZE and the DSS pool has a greater importance than the OLTP pool.

B.

The node will have the server assigned to the OLTP pool because the OLTP pool's MAX\_SIZE is greater than MAX\_SIZE of the DSS pool.

C.

**The node will have the server assigned to the OLTP pool because the OLTP pool's current size is less than its MIN\_SIZE and the DSS pool has enough servers to equal its own MIN\_SIZE.**

D.

The node will have the server assigned to the DSS pool on the basis of having a greater IMPORTANCE.

#### Explanation:

Oracle Clusterware continues to assign servers to server pools until the following conditions are met:

Until all server pools are filled in order of importance to their minimum (MIN\_SIZE).

Until all server pools are filled in order of importance to their maximum (MAX\_SIZE).

By default, any servers not placed in a server pool go into the Free server pool.

Oracle Clusterware Administration and Deployment Guide

11g Release 2 (11.2)

### What steps are required to increase the diskgroup capacity?

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

An ASM diskgroup that contains database tablespaces with six disks and normal redundancy is nearing 90% capacity. You want to add additional capacity to the diskgroup. A set of six additional disks has been made available, and appears as CANDIDATE disks in the V\$ASM\_DISK view.

What steps are required to increase the diskgroup capacity?

A.

Export the database objects.

Drop the diskgroup.

Create a new diskgroup.

Import the database objects.

B.

Shut down the database instances using the diskgroup.

Add the disks to the diskgroup.

Start the database instances using the diskgroup.

C.

**Set the rebalance power limit (or allow it to default).**

**Add the disks to the diskgroup.**

D.

Shut down all database instances and ASM instances.

Back up the ASM diskgroup.

Start the ASM instances.

Add the disks to the diskgroup.

Issue a rebalance command.  
Start the database instances.

Explanation:

You can use the ADD clause of the ALTER DISKGROUP statement to add a disk or a failure group to a disk group. The same syntax that you use to add a disk or failure group with the CREATE DISKGROUP statement can be used with the ALTER DISKGROUP statement. ASM automatically rebalances the disk group when disks are added. By default, the ALTER DISKGROUP statement returns immediately after the disks have been added while the rebalance operation continues to run asynchronously. You can query the V\$ASM\_OPERATION view to monitor the status of the rebalance operation.

D60488GC11

Oracle 11g RAC and Grid Infrastructure Administration Accelerated 8 – 21

What is true about the attributes that may be set to control the application?

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

As part of your data center's high availability strategy, you are creating resource definitions to control the management of a web-based application by the Oracle Grid Infrastructure clusterware stack.

The application and its VIP are normally online on one node of a four-node cluster due to the CARDINALITY of the resource type being set to 1.

You have chosen a policy-managed resource type for the application by using a server pool that uses only RACNODE3 and RACNODE4. The START ATTEMPTS attribute for the resource is set to 2 and FAILURE INTERVAL is set to 60.

What is true about the attributes that may be set to control the application?

A.

The clusterware will attempt to start the application on the same node twice within the server pool as long as that node is up. If the node fails, then the VIP and the application will be failed over to the other node in the server pool immediately.

B.

The clusterware will attempt to start the application on the same node twice within the server pool as long as that node is up. If the node fails, then the VIP and the application will be failed over to the other node in the server pool only after two 60-second intervals have elapsed.

C.

The clusterware will attempt to start the application on the same node twice within the server pool as long as that node is up. If the application fails to start after 60 seconds, but the node is still up, then the VIP and the application will NOT be failed over to the other node in the server pool.

D.

The clusterware will attempt to start the application on the same node twice within the server pool as long as that node is up. If the application fails to start immediately but the node is still up, then the VIP and the application will NOT be failed over to the other node in the server pool.

Explanation:

CARDINALITY

The number of servers on which a resource can run, simultaneously. This is the upper limit for resource cardinality.

RESTART\_ATTEMPTS

The number of times that Oracle Clusterware attempts to restart a resource on the resource's current server before attempting to relocate it. A value of 1 indicates that Oracle Clusterware only attempts to restart the resource once on a server. A second failure causes Oracle Clusterware to attempt to relocate the resource. A value of 0 indicates that there is no attempt to restart but Oracle Clusterware always tries to fail the resource over to another server.

FAILURE\_INTERVAL

The interval, in seconds, before which Oracle Clusterware stops a resource if the resource has exceeded the number of failures specified by the FAILURE\_THRESHOLD attribute. If the value is zero (0), then tracking of failures is disabled.

FAILURE\_THRESHOLD

The number of failures of a resource detected within a specified FAILURE\_INTERVAL for the resource before Oracle Clusterware marks the resource as unavailable and no longer monitors it. If a resource fails the specified number of times, then Oracle Clusterware stops the resource. If the value is zero (0), then tracking of failures is disabled. The maximum value is 20.

Oracle Clusterware Administration and Deployment Guide

11g Release 2 (11.2)

**Choose four tools that can be used to create ASM disk groups.**

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

Choose four tools that can be used to create ASM disk groups.

- A.  
Enterprise Manager
- B.  
SRVCTL
- C.  
ASMCMD
- D.  
CRSCTL
- E.  
ASMCA
- F.  
SQL\*Plus

**Which three statements are true about ASM dynamic volume manager (ADVM)?**

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

Which three statements are true about ASM dynamic volume manager (ADVM)?

- A.  
ADVM provides volume management services and a standard disk device driver interface to file system drivers.
- B.  
The administrator can use ADVM to create volumes that contain bootable vendor operating systems.
- C.  
File systems and other disk-based applications issue I/O requests to ADVM volume devices as they would to other storage devices on a vendor operating system.
- D.  
ADVM extends ASM by providing a disk driver interface to storage backed by an ASM volume.
- E.  
To use the ADVM driver, the oracleacfs, oradeoks, and oracleadvn drivers must be loaded, but an ASM instance is not required.

**Explanation:**

At the operating system (OS) level, the ASM instance provides the disk group, which is a logical container for physical disk space. The disk group can hold ASM database files and ASM dynamic volume files. The ASM Dynamic Volume Manager (ADVM) presents the volume device file to the operating system as a block device.

The mkfs utility can be used to create an ASM file system in the volume device file.

Four OS kernel modules loaded in the OS provide the data service. On Linux, they are:

oracleasm, the ASM module;

oracleadvn, the ASM dynamic volume manager module;

oracleoks, the kernel services module; and oracleacfs, the ASM file system module.

These modules provide the ASM Cluster File System, ACFS snapshots, the ADVM, and cluster services.

The ASM volumes are presented to the OS as a device file at /dev/asm/<volume name>-<number>.

ADVM provides volume management services and a standard disk device driver interface to clients. Clients, such as file systems and other disk-based applications, issue I/O requests to

ADVM volume devices as they would to other storage devices on a vendor operating system.

ADVM extends ASM by providing a disk driver interface to storage backed by an ASM file. The

administrator can use the ADVM to create volumes that contain file systems. These file systems

can be used to support files beyond Oracle database files such as executables, report files, trace

files, alert logs, and other application data files. With the addition of ADVM and ACFS, ASM

becomes a complete storage solution of user data for both database and non-database file needs.

ACFS is intended as a general file system accessible by the standard OS utilities. ACFS can be

used in either a single server or a cluster environment.

Note: Oracle ACFS file systems cannot be used for an Oracle base directory or an Oracle grid

infrastructure home that contains the software for Oracle Clusterware, ASM, Oracle ACFS, and Oracle ADVM components.

Oracle ACFS file systems cannot be used for an OS root directory or boot directory.

ASM volumes serve as containers for storage presented as a block device accessed through ADVM. File systems or user processes can do I/O on this ASM volume device just as they would on any other device. To accomplish this, ADVM is configured into the operating system. A volume device is constructed from an ASM file.

D60488GC11

Oracle 11g RAC and Grid Infrastructure Administration Accelerated 10 – 3,4,5

### which two purposes would you recommend an ASM clustered file system (ACFS)?

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

For which two purposes would you recommend an ASM clustered file system (ACFS)?

- A.  
a shared home directory for Oracle database executables in a single-instance cluster for cold failover
- B.  
a shared home directory for Oracle Grid Infrastructure executables
- C.  
a root file system for the operating system
- D.  
a shared file system for RAC data files
- E.  
a general purpose shared file system for OS files
- F.  
a clustered file system for OCR and voting disk files

#### Explanation:

Overview of Oracle ACFS

Oracle Automatic Storage Management Cluster File System (Oracle ACFS) is a multi-platform, scalable file system, and storage management technology that extends Oracle Automatic Storage Management (Oracle ASM) functionality to support customer files maintained outside of Oracle Database. Oracle ACFS supports many database and application files, including executables, database trace files, database alert logs, application reports, BFILES, and configuration files. Other supported files are video, audio, text, images, engineering drawings, and other general-purpose application file data.

Notes:

Oracle ASM is the preferred storage manager for all database files. It has been specifically designed and optimized to provide the best performance for database file types. For a list of file types supported by Oracle ASM, see Table 7-1, "File types supported by Oracle ASM".

Oracle ACFS is the preferred file manager for non-database files. It is optimized for general purpose files.

Oracle ACFS does not support any file type that can be directly stored in Oracle ASM, except where explicitly noted in the documentation.

Not supported means Oracle Support Services does not take calls and development does not fix bugs associated with storing unsupported file types in Oracle ACFS.

Starting with Oracle Automatic Storage Management 11g Release 2 (11.2.0.3), Oracle ACFS supports RMAN backups (BACKUPSET file type), archive logs (ARCHIVELOG file type), and Data Pump dumpsets

(DUMPSET file type). Note that Oracle ACFS snapshots are not supported with these files.

Oracle ACFS does not support files for the Oracle Grid Infrastructure home.

Oracle ACFS does not support Oracle Cluster Registry (OCR) and voting files.

Oracle ACFS functionality requires that the disk group compatibility attributes for ASM and ADVM be set to

11.2 or greater. For information about disk group compatibility, refer to "Disk Group Compatibility".

Oracle Automatic Storage Management Administrator's Guide

11g Release 2 (11.2)

### Which three items are required?

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

A Grid Plug and Play (GPnP) software image must contain at least three items. Which three items

are required?

- A.  
operating system software
- B.  
Oracle Database software
- C.  
GPnP software
- D.  
security certificate of the provisioning authority
- E.  
application software

Explanation:

GPnP Components

Software image

A software image is a read-only collection of software to be run on nodes of the same type.

At a minimum, the image must contain:

- An operating system
  - The GPnP software
  - A security certificate from the provisioning authority
  
  - Other software required to configure the node when it starts up
- D60488GC11  
Oracle 11g: RAC and Grid Infrastructure Administration

Which two methods fulfill the installation requirements?

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

As part of the preinstallation process for adding two new nodes to your four-node UNIX cluster, you are in discussions with the OS administrators about the operating system installation and setup for the two new nodes called RACNODE5 and RACNODE6.

The nodes have already been connected to the network infrastructure and the administrators are ready for the OS installation. Which two methods fulfill the installation requirements?

- A.  
Install a new image of the OS, then configure SSH for the root user.
- B.  
Install a cloned image of the OS that at least matches the existing node images for drivers, patches, and updates.
- C.  
Install a new image of the OS, that at least matches an existing node for drivers, patches, and updates, and create the necessary OS users and groups with user and group IDs matching those on the existing nodes.
- D.  
Install a new image of the OS that at least matches the existing node images for drivers.
- E.  
Install a new image of the OS, and create the necessary OS users and groups with any user and group IDs.

Explanation:

Prerequisite Steps for Adding Cluster Nodes

1. Make physical connections.

Connect the nodes' hardware to the network infrastructure of your cluster. This includes establishing electrical connections, configuring network interconnects, configuring shared disk subsystem connections, and so on. See your hardware vendor documentation for details about this step.

2. Install the operating system.

Install a cloned image of the operating system that matches the operating system on the other nodes in your cluster. This includes installing required service patches, updates, and drivers. See your operating system vendor documentation for details about this process. Note: Oracle recommends that you use a cloned image. However, if the installation fulfills the installation requirements, then install the operating system according to the vendor documentation.

3. Create Oracle users.

You must create all Oracle users on the new node that exist on the existing nodes. For example, if you are adding a node to a cluster that has two nodes, and those two nodes have different owners for the Grid Infrastructure home and the Oracle home, then you must create those owners on the new node, even if you do not plan to install an Oracle home on the new node.

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

**Which ASM\_DISKSTRING parameter would perform the best by searching the fewest number of devices for ASM discovery?**

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

The storage administrator is adding several disks to a pool of ASM disks. The current ASM\_DISKSTRING parameter is "/dev/sda\*", /dev/sdb\*, /dev/sd\*6". There are several hundred disk partitions currently.

The new disks are /dev/sdf1, /dev/sdg1, /dev/sdh1, and /dev/sdd6.

Which ASM\_DISKSTRING parameter would perform the best by searching the fewest number of devices for ASM discovery?

- A.  
ASM\_DISKSTRING="/dev/sda\*", "/dev/sdb\*", "/dev/sd\*6", "/dev/sd\*1"
- B.  
ASM\_DISKSTRING="/dev/sda\*", "/dev/sdb\*", "/dev/sd\*6", "/dev/sd\*"
- C.  
ASM\_DISKSTRING="/dev/sd\*"
- D.  
ASM\_DISKSTRING=""
- E.  
ASM\_DISKSTRING="/dev/sda\*", "/dev/sdb\*", "/dev/sd\*6", "/dev/sdf1", "/dev/sdg1", "/dev/sdh1"

Explanation:

Improving Disk Discovery Time

The value for the ASM\_DISKSTRING initialization parameter is an operating system-dependent value that Oracle ASM uses to limit the set of paths that the discovery process uses to search for disks. When a new disk is added to a disk group, each Oracle ASM instance that has the disk group mounted must be able to discover the new disk using its ASM\_DISKSTRING.

Oracle Automatic Storage Management Administrator's Guide  
11g Release 2 (11.2)

**Which command-line tool would you use to do this?**

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

You want to check an ACFS for consistency on the Linux platform. Which command-line tool would you use to do this?

- A.  
advmtutil
- B.  
asmcmd
- C.  
fsck
- D.  
acfsutil
- E.  
fdisk

Explanation:

Dismounting an Oracle ACFS File System

You can dismount a file system without deregistering the file system or disabling the volume on which the file system is mounted.

For example, you can dismount a file system and run fsck to check the file system.

```
# /bin/umount /u01/app/acfsmounts/myacfs  
# /sbin/fsck -a -v -y -t acfs /dev/asm/volume1-123
```

After you dismount a file system, you must explicitly mount the file system.

1. Use umount on Linux systems or acfsdismount on Windows systems. For information about the

commands to dismount a file system, see “umount” or “acfsdismount”.

2. Use fsck on Linux systems or acfschkdsk on Windows systems to check a file system. For information about the commands to check a file system, see “fsck” or “acfschkdsk”.

Oracle Automatic Storage Management Administrator’s Guide  
11g Release 2 (11.2)

**Which two statements are true about the high availability capabilities of Oracle HA services?**

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

The high availability services provided by Oracle Clusterware are used to protect Oracle resources, such as RAC database instances, RAC database services, and other components of the Oracle Grid Infrastructure, and non-Oracle resources as well.

Which two statements are true about the high availability capabilities of Oracle HA services?

- A.  
RAC databases may have their instances failed over in some cases.
- B.  
ASM instances may be failed over if fewer than three nodes remain in the cluster, so that there are always at least three ASM instances available.
- C.  
If a node fails, then all resources that were active on that node will be failed over to a surviving node if any exists.
- D.  
**If a node fails, then cluster resources that were active on that node may be failed over to a surviving node if any exists, but local resources are not failed over.**
- E.  
**HA services will only fail over a resource upon failure of the node where the resource was active.**

Explanation:

Resource Types

Local resource: These are server-centric resources; the type name is local\_resource. These run locally on individual servers of the cluster and are not relevant outside of the scope of the server.

Cluster resource: Cluster-aware resource types (type name is cluster\_resource) are aware of the cluster environment and are subject to cardinality and cross-server switchover and failover. The state of cluster resources is relevant in the context of the cluster.

D60488GC11

Oracle 11g: RAC and Grid Infrastructure Administration

**Which three fragments will complete this statement correctly?**

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

Which three fragments will complete this statement correctly? In a cluster environment, an ACFS volume \_\_\_\_\_.

- A.  
will be automatically mounted by a node on reboot by default
- B.  
must be manually mounted after a node reboot
- C.  
**will be automatically mounted by a node on cluster stack startup if it is included in the ACFS mount registry**
- D.  
**will be automatically mounted by a node if it is defined as a cluster resource when a dependent cluster resource requires access**
- E.  
**will be automatically mounted to all nodes in the cluster when the file system is registered**
- F.  
must be mounted before it can be registered

Explanation:

About the Oracle ACFS Mount Registry

The Oracle ACFS mount registry supports both Oracle Restart and Oracle Grid Infrastructure

cluster configurations. File systems that are to be mounted persistently (across reboots) can be registered with the Oracle ACFS mount registry. In cluster configurations, registered Oracle ACFS file systems are automatically mounted by the mount registry, similar to a clusterwide mount table. However, in Oracle Restart configurations the automatic mounting of registered Oracle ACFS file systems is not supported. For more information, see

“Oracle ACFS and Oracle Restart”.

By default, an Oracle ACFS file system that is inserted into the cluster mount registry is automatically mounted on all cluster members, including cluster members that are added after the registry addition. However, the cluster mount registry also accommodates single-node and multi-node (subset of cluster nodes) file system registrations. The mount registry actions for each

cluster member mount only registered file systems that have been designated for mounting on that member.

The Oracle ACFS registry resource actions are designed to automatically mount a file system only one time for each Oracle Grid Infrastructure initialization to avoid potential conflicts with administrative actions to dismount a given file system.

Oracle ACFS Individual File System Resource Management

The Oracle ACFS individual file system resource is supported only for Oracle Grid Infrastructure cluster configurations; it is not supported for Oracle Restart configurations. See “Oracle ACFS and Oracle Restart”.

Oracle ASM Configuration Assistant (ASMCA) facilitates the creation of Oracle ACFS individual file system resources (ora.diskgroup.volume.acfs). During database creation with Database Configuration Assistant (DBCA), the individual file system resource is included in the dependency list of its associated disk group so that stopping the disk group also attempts to stop any dependent Oracle ACFS file systems.

An Oracle ACFS individual file system resource is typically created for use with application resource dependency lists. For example, if an Oracle ACFS file system is configured for use as an Oracle Database home, then a resource created for the file system can be included in the resource dependency list of the Oracle Database application. This dependency causes the file system and stack to be automatically mounted due to the start action of the database application.

An Oracle ACFS file system that is to be mounted from a dependency action should not be included in the Oracle ACFS mount registry.

Oracle Automatic Storage Management Administrator’s Guide

11g Release 2 (11.2)

### Which command will set full debugging for the communications and resource management components of the CRS module on nodes host01 and host03?

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

Which command will set full debugging for the communications and resource management components of the CRS module on nodes host01 and host03?

- A.  
srvctl set log -m CRS -c CRSCTI 5, CRSCOMM 5 -m host01,host03
- B.  
ocrconfig set log CRS “CRSRTI, CRSCOMM” -nodelist host01,host03
- C.  
**crsctl set log CRS “CRSRTI, CRSCOMM” -nodelist host01,host03**
- D.  
crsctl lsmodules crs set log CRSRTX , CRSCOMM -nodelist host01, host03

#### Explanation:

##### Dynamic Debugging

This section includes the following CRSCCTL commands that aid in debugging:

crsctl set log

Use the crsctl set log command to set log levels for Oracle Clusterware.

##### Syntax

```
crsctl set log {[crs | css | evm "component_name=log_level, [...]"] |  
[all=log_level]}
```

You can also set log levels for the agents of specific resources, as follows:

```
crsctl set log res “resource_name=log_level, [...]“
```



Table E-44 Oracle Clusterware (CRS) Modules and Functions

| Module   | Descriptions                      |
|----------|-----------------------------------|
| CRSUI    | User interface module             |
| CRSCOMM  | Communication module              |
| CRSRTI   | Resource management module        |
| CRSMAIN  | Main module/driver                |
| CRSPLACE | CRS placement module              |
| CRSAPP   | CRS application                   |
| CRSRES   | CRS resources                     |
| CRSOCR   | Oracle Cluster Registry interface |
| CRSTIMER | Various timers related to CRS     |
| CRSEVT   | CRS EVM/event interface module    |
| CRSD     | CRS daemon                        |

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

**Which statement describes the requirements for the network interface names, such as eth0, in Oracle Clusterware?**

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

Which statement describes the requirements for the network interface names, such as eth0, in Oracle Clusterware?

- A.  
Only the public interface names must be the same for all nodes.
- B.  
Only the private interface names must be the same for all nodes.
- C.  
**Both the public interface name and the private interface name must be the same for all nodes.**
- D.  
Both the public interface name and the private interface name can vary on different nodes.
- E.  
Only the private interface names can be different on different nodes.
- F.  
Only the public interface names can be different on different nodes.

Explanation:

Checking Network Requirements

Each node must have at least two network interface cards (NICs).

Interface names must be the same on all nodes.

Public NIC must support TCP/IP and Private NIC UDP.

Public IP must be registered in the domain name server (DNS) or the /etc/hosts file.

```
# cat /etc/hosts
```

```
##### Public Interfaces eth0 (odd numbers)#####
```

```
xxx.xxx.100.11
```

```
host01.example.com host01
```

```
xxx.xxx.100.13
```

```
host02.example.com host02
```

IF GNS is used, the Cluster GNS address must be registered in the DNS.

Prevent public network failures when using NAS devices or NFS mounts by starting the Name

Service Cache Daemon.

```
# /sbin/service nscd start
```

```
D60488GC11
```

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 2 18

**What is true about this situation?**

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

A policy-managed RAC database that hosts the ACCOUNTS service has only three instances running rather than the usual four on a six-node cluster.

Five of the six cluster nodes are active and the sixth (RACNODE6) is down for maintenance.

You decide to check the state of the servers in the server pools to see if there are problems with the OLTP pool to which the ACCOUNTS service has been assigned.

You used the `ctctl status server -f` command and see that the RACNODE3 node has

STATE=VISIBLE.

What is true about this situation?

A.

The Cluster Synchronization Services Daemon (CSSD) is running RACNODE3 is considered to be part of the cluster, but the Cluster Ready Services Daemon (CRSD) is currently not running.

B.

Both the Cluster Ready Services Daemon (CRSD) and the Cluster Synchronization Services Daemon (CSSD) are running on RACNODE3, but RACNODE3 is currently being moved to another server pool.

C.

The Cluster Ready Services Daemon (CRSD) is running and RACNODE3 is considered to be part of the cluster, but the Cluster Synchronization Services Daemon (CSSD) is currently not running.

D.

Neither the Cluster Ready Services Daemon (CRSD) nor the Cluster Synchronization Services Daemon (CSSD) is running on RACNODE3.

Explanation:

|               |  |
|---------------|--|
| STATE         | A server can be in one of the following states:  |
| ONLINE        | The server is a member of the cluster and is available for resource placement.   |
| OFFLINE       | The server is not currently a member of the cluster. Subsequently, it is not available for resource placement.   |
| JOINING       | When a server joins a cluster, Oracle Clusterware processes the server to ensure that it is valid for resource placement.<br>Oracle Clusterware also checks the state of resources configured to run on the server. Once the validity of the server and the state of the resources are determined, the server transitions out of this state.   |
| LEAVING       | When a planned shutdown for a server begins the state of the server transitions to LEAVING, making it unavailable for resource placement.  |
| VISIBLE       | Servers that have Oracle Clusterware running, but not the Cluster Ready Services daemon (crsd), are put into the VISIBLE state.<br>This usually indicates an intermittent issue or failure and Oracle Clusterware trying to recover (restart) the daemon. Oracle Clusterware cannot manage resources on servers while the servers are in this state.   |
| RECONFIGURING | When servers move between server pools due to server pool reconfiguration, a server is placed into this state if resources that ran on it in the current server pool must be stopped and relocated. This happens because resources running on the server may not be configured to run in the server pool to which the server is moving. As soon as the resources are successfully relocated, the server is put back into the ONLINE state. |

Use the crsctl status server command to obtain server  
Oracle Clusterware Administration and Deployment Guide  
11g Release 2 (11.2)

What is the correct way to silently add the nodes?

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

You are ready to add two new nodes called RACNODE 5 and RACNODE 6 to your existing four-node cluster by using addNode.sh. You have run cluvfy -peer to check the new nodes against a reference node.

When you originally created the cluster, the network administrators chose to statically define the SCAN VIP addresses in the corporate DNS server, and you installed the Oracle Grid Infrastructure without using GNS.

What is the correct way to silently add the nodes?

A.

```
addNode.sh -silent & CLUSTER_NEW_NODES={RACNODE5, RACNODE6} "
```

B.

```
addNode.sh -silent " CLUSTER_NEW_VIRTUAL_HOSTNAMES={RACNODE5-VIP, RACNODE6-VIP} "
```

C.

```
addNode.sh -silent & CLUSTER_NEW_NODES={RACNODE5, RACNODE6}"  
"CLUSTER_NEW_VIRTUAL_HOSTNAMES={RACNODE5-VIP, RACNODE6-VIP}"
```

D.

addNode.sh -silent -responseFile mynewnodes.txt with the response file containing only  
◆CLUSTER\_NEW\_NODES={RACNODE5, RACNODE6} “

E.  
addNode.sh -silent -responseFile mynewnodes.txt With the response file containing  
“CLUSTER\_NEW\_VIRTUAL\_HOSTNAMES={RACNODE3-VIP, RACNODE4-VIP}”

Explanation:

Adding a Cluster Node on Linux and UNIX Systems

If you are not using GNS, run the following command:

```
$ ./addNode.sh "CLUSTER_NEW_NODES={node3}"
```

```
"CLUSTER_NEW_VIRTUAL_HOSTNAMES={node3-vip}"
```

Oracle Clusterware Administration and Deployment Guide

11g Release 2 (11.2)

**Which command will resize an ACFS file system on the DATA diskgroup?**

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

Which command will resize an ACFS file system on the DATA diskgroup?

-The ADVM volume is named VOL1.

-The mount point is >/u01/app/oracle/afsmounts/voll.

-The new, larger size is 20G.

A.  
ASMCMD [+] > volsize -G data -s 20G VOL1

B.  
**/sbin/acfsutil size 20G /u01/app/oracle/afsmounts/voll**

C.  
SQL> ALTER DISKGROUP data RESIZE VOLUME voll SIZE 20G;

D.  
None, an ACFS file system cannot be resized.

E.  
fsadm /u01/app/oracle/afsmounts/voll 20G

Explanation:

acfsutil size

Purpose

Resizes an Oracle ACFS file system.

Syntax and Description

acfsutil size -h

acfsutil size [+|-]n[K|M|G|T|P ] [device] mount\_point

acfsutil size -h displays help text and exits.

Oracle Automatic Storage Management Administrator's Guide

11g Release 2 (11.2)

ALTER DISKGROUP

RESIZE VOLUME Clause

Use this clause to change the size of an existing Oracle ADVM volume. In an Oracle ASM cluster, the new size is propagated to all nodes. If an Oracle Automatic Storage Management File System (ACFS) exists on the volume, then you must use the acfsutil size command instead of the ALTER DISKGROUP statement.

Oracle Database SQL Language Reference

11g Release 2 (11.2)

**Which two statements are true about the data disk group?**

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

You use the following command to create a disk group:

```
CREATE DISKGROUP DATA NORMAL REDUNDANCY
```

```
FAILGROUP fg_fra_1 DISK
```

```
'/dev/sda6' NAME FRA_DISK1,
```

```
'/dev/sda7' NAME FRA_DISK2
```

```
FAILGROUP fg_fra_2 DISK
```

```
'/dev/sdb1' NAME FRA_DISK3,
```

```
'/dev/sdb2' NAME FRA_DISK4
```

```
FAILGROUP fg_fra_3 DISK
```

```
‘dev/sdc1’ NAME FRA_DISK5,  
‘dev/sdc2’ NAME FRA_DISK6  
ATTRIBUTES  
‘AU_SIZE’ = ‘4M’  
‘compatible.rdbms’ = ‘11.1.0’;
```

Which two statements are true about the data disk group?

- A.  
Each disk in a failure group is a mirror of the corresponding disk in the other failure groups. For example, FRA\_DISK2 is a mirror of FRA\_DISK4 and FRA\_DISK6.
- B.  
By default, the files created in this disk group must have, for each extent, one primary copy and two mirror copies; one on each of the failure groups.
- C.  
Because the size clause is not used, the disk group uses the same amount of space on each disk.
- D.  
The ‘compatible.rdbms’ attribute specifies the minimum COMPATIBLE setting for any database using this disk group.
- E.  
If the FRA\_DISK2 disk fails, there are mirrored copies in one of the other failure groups for any primary extents that are on FRA\_DISK2.
- F.  
The allocation unit, AU\_SIZE, sets the size of the disk read and write to ASM diskgroups.

#### Explanation:

##### About Mirroring and Failure Groups

Mirroring protects data integrity by storing copies of data on multiple disks. When you create a disk group, you specify an Oracle ASM disk group type based on one of the following three redundancy levels:

Normal for 2-way mirroring

High for 3-way mirroring

External to not use Oracle ASM mirroring, such as when you configure hardware RAID for redundancy

The redundancy level controls how many disk failures are tolerated without dismounting the disk group or losing data. The disk group type determines the mirroring levels with which Oracle creates files in a disk group.

For information about disk group types and templates, see “Managing Disk Group Templates”.

Oracle ASM mirroring is more flexible than traditional RAID mirroring. For a disk group specified as NORMAL redundancy, you can specify the redundancy level for each file. For example, two

files can share the same disk group with one file being mirrored while the other is not.

When Oracle ASM allocates an extent for a mirrored file, Oracle ASM allocates a primary copy and a mirror copy. Oracle ASM chooses the disk on which to store the mirror copy in a different failure group than the primary copy. Failure groups are used to place mirrored copies of data so that each copy is on a disk in a different failure group. The simultaneous failure of all disks in a failure group does not result in data loss.

You define the failure groups for a disk group when you create an Oracle ASM disk group. After a disk group is created, you cannot alter the redundancy level of the disk group. If you omit the failure group specification, then Oracle ASM automatically places each disk into its own failure group, except for disk groups containing disks on Oracle Exadata cells. Normal redundancy disk groups require at least two failure groups. High redundancy disk groups require at least three failure groups. Disk groups with external redundancy do not use failure groups.

##### Allocation Units

Every Oracle ASM disk is divided into allocation units (AU). An allocation unit is the fundamental unit of allocation within a disk group. A file extent consists of one or more allocation units. An Oracle ASM file consists of one or more file extents.

When you create a disk group, you can set the Oracle ASM allocation unit size with the AU\_SIZE disk group attribute. The values can be 1, 2, 4, 8, 16, 32, or 64 MB, depending on the specific disk group compatibility level. Larger AU sizes typically provide performance advantages for data warehouse applications that use large sequential reads.

Oracle Automatic Storage Management Administrator’s Guide

11g Release 2 (11.2)

**Which three steps must be performed to remove the nodes from the cluster?**

---

Your cluster was originally created with nodes RACNODE1 and RACNODE2 three years ago. Last year, nodes RACNODE3 and RACNODE4 were added.

These nodes have faster processors and more local storage than the original nodes making performance management and tuning more difficult.

Two more nodes with the same processor speed have been added to the cluster last week as RACNODE5 and RACNODE6 and you must remove RACNODE1 and RACNODE2 for redeployment.

The Oracle Grid Infrastructure is using GNS and the databases are all 11g Release 2, all running from the same home. The Grid home is /fs01/home/grid.

Which three steps must be performed to remove the nodes from the cluster?

A.

```
Run /fs01/home/grid/oui/bin/runInstaller -updateNodeList
ORACLE_HOME=/fs01/home/grid "CLUSTER_NODES={RACNODE3 , RACNODE4 ,
RACNODE5 ,
RACNODE6}"
as the grid software owner on any remaining node.
```

B.

```
Run /fs01/home/grid/oui/bin/runInstaller -updateNodeList
ORACLE_HOME=/fs01/home/grid " CLUSTER_NODES={RACNODE1}"
as the grid software owner on RACNODE1 and run /fs01/home/grid/oui/bin/runInstaller -updateNodeList
ORACLE_HOME=/fs01/home/grid " CLUSTER_NODES={RACNODE 2}" as the grid software
owner on
RACNODE2.
```

C.

```
Run /fs01/home/grid/oui/bin/runInstaller -detachHome ORACLE_HOME=/fs01/home/grid as the
grid
software owner on RACNODE1 and RACNODE2.
```

D.

```
Run the /fs01/home/grid/crs/install/rootcrs.pl script as root on each node to be deleted.
```

E.

```
Run crsctl delete node -n RACNODE1 and crsctl delete node -n RACNODE2 as root from any
node remaining in the cluster.
```

#### Explanation:

Deleting a Cluster Node on Linux and UNIX Systems

1. Ensure that Grid\_home correctly specifies the full directory path for the Oracle Clusterware home on each node, where Grid\_home is the location of the installed Oracle Clusterware software.

2. Run the following command as either root or the user that installed Oracle Clusterware to determine whether the node you want to delete is active and whether it is pinned:

```
$ olsnodes -s -t
```

If the node is pinned, then run the crsctl unpin css command. Otherwise, proceed to the next step.

3. Disable the Oracle Clusterware applications and daemons running on the node. Run the rootcrs.pl script as root from the Grid\_home/crs/install directory on the node to be deleted, as follows:

```
# ./rootcrs.pl -deconfig -deinstall -force
```

If you are deleting multiple nodes, then run the rootcrs.pl script on each node that you are deleting.

If you are deleting all nodes from a cluster, then append the -lastnode option to the preceding

command to clear OCR and the voting disks, as follows:

```
# ./rootcrs.pl -deconfig -deinstall -force -lastnode
```

4. From any node that you are not deleting, run the following command from the Grid\_home/bin directory as root to delete the node from the cluster:

```
# crsctl delete node -n node_to_be_deleted
```

Then if you run a dynamic Grid Plug and Play cluster using DHCP and GNS, skip to step 7.

5. On the node you want to delete, run the following command as the user that installed Oracle Clusterware from the Grid\_home/oui/bin directory where node\_to\_be\_deleted is the name of the node that you are deleting:

```
$ ./runInstaller -updateNodeList ORACLE_HOME=Grid_home "CLUSTER_NODES={
node_to_be_deleted}" CRS=TRUE -silent -local
```

6. On the node that you are deleting, depending on whether you have a shared or local Oracle home, complete one of the following procedures as the user that installed Oracle Clusterware:

If you have a shared home, then run the following command from the Grid\_home/oui/bin directory on the node you want to delete:

```

$ ./runInstaller -detachHome ORACLE_HOME=Grid_home -silent -local
For a local home, deinstall the Oracle Clusterware home from the node that you want to delete, as follows, by running the following command, where Grid_home is the path defined for the Oracle Clusterware home:
$ Grid_home/deinstall/deinstall local
7. On any node other than the node you are deleting, run the following command from the Grid_home/oui/bin directory where remaining_nodes_list is a comma-delimited list of the nodes that are going to remain part of your cluster:
$ ./runInstaller -updateNodeList ORACLE_HOME=Grid_home "CLUSTER_NODES={remaining_nodes_list}" CRS=TRUE -silent
8. Run the following CVU command to verify that the specified nodes have been successfully deleted from the cluster:
$ cluvfy stage -post nodedel -n node_list [-verbose]
Oracle Clusterware Administration and Deployment Guide
11g Release 2 (11.2)

```

**Which statement is true in this scenario?**

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

Examine the set of commands given below:

```

$ echo $ORACLE_SID
RACDB_1
$ sqlplus / as sysdba

SQL> SELECT * FROM V$ACTIVE_INSTANCES;
INST_NUMBER INST_NAME
-----
1          gr7597:RACDB_1
2          gr7602:RACDB_2
3          gr7633:RACDB_3

SQL>ALTER SYSTEM SET TRACE_ENABLED=TRUE SCOPE=MEMORY sid='*'
SQL>ALTER SYSTEM SET TRACE_ENABLED=FALSE SCOPE=MEMORY sid='RACDB_1';

```

Which statement is true in this scenario?

- A. All the instances of the RACDB will have the TRACE\_ENABLED parameter with the value TRUE because this parameter must have the same value on all the instances.
- B. All the instances of the RACDB will have the TRACE\_ENABLED parameter with the value TRUE except the RACDB\_1 instance, which will have the value FALSE for this parameter.
- C. All the instances of the RACDB will have the TRACE\_ENABLED parameter with the value TRUE because the first ALTER statement has higher precedence.
- D. All the instances of the RACDB will have the TRACE\_ENABLED parameter with the value FALSE because the second ALTER statement has higher precedence.

Explanation:

SPFILE Parameter Values and RAC

- You can change parameter settings using the ALTER SYSTEM SET command from any instance:  
ALTER SYSTEM SET <dpname> SCOPE=MEMORY sid='<sid|\*>';
- SPFILE entries such as:
  - \*.<pname> apply to all instances
  - <sid>.<pname> apply only to <sid>
  - <sid>.<pname> takes precedence over \*.<pname>
- Use current or future \*.<dpname> settings for <sid>:  
ALTER SYSTEM RESET <dpname> SCOPE=MEMORY sid='<sid|\*>';
- Remove an entry from your SPFILE:  
ALTER SYSTEM RESET <dpname> SCOPE=SPFILE sid='<sid|\*>';

**Which three programs or utilities can be used to convert a single-instance database to a Real Application Cluster database?**

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

Which three programs or utilities can be used to convert a single-instance database to a Real Application Cluster database?

- A.  
DBCA
- B.  
the SRVCTL utility
- C.  
Enterprise Manager
- D.  
the RCONFIG utility
- E.  
the CRSCTLI utility

Explanation:

Single Instance to RAC Conversion  
Single-instance databases can be converted to RAC using:  
DBCA  
Enterprise Manager  
RCONFIG utility  
D60488GC11  
Oracle 11g: RAC and Grid Infrastructure Administration Accelerated

What would you recommend to improve the response times?

Posted by seenagape on March 11, 2013

[No comments](#)

Your RAC database has a high volume of inserts into the SALES table.  
Sequence SALES\_SEQ is used to generate primary key values.  
Examine the following query output:  
SQL> SELECT sequence\_name, min\_value, max\_value, increment\_by, cache\_size FROM dba\_sequences ;

Output:

| SEQUENCE_NAME | MIN_VALUE | MAX_VALUE  | INCREMENT_BY | CACHE_SIZE |
|---------------|-----------|------------|--------------|------------|
| SALES_SEQ     | 1         | 1.0000E+28 | 1            | 500        |

A check of the top 5 timed events shows the following:

Top 5 Timed Events  
~~~~~

| Event                   | Waits | Time (s) | % Total DB Time |
|-------------------------|-------|----------|-----------------|
| CPU Time                |       | 331      | 66.82           |
| log file parallel write | 4,628 | 151      | 30.50           |
| log file sync           | 4,366 | 141      | 28.44           |
| enq: SQ - contention    | 3,822 | 78       | 25.32           |
| db file sequential read | 3,692 | 70       | 1.35            |

What would you recommend to improve the response times?

- A.  
Increasing size of undo tablespaces
- B.  
increasing sequence cache sizes
- C.  
???
- D.  
???

Explanation:

Oracle Sequences and Index Contention  
Indexes with key values generated by sequences tend to be subject to leaf block contention when the insert rate is high. That is because the index leaf block holding the highest key value is changed for every row inserted, as the values are monotonically ascending. In RAC, this may lead to a high rate of current and CR blocks transferred between nodes.

One of the simplest techniques that can be used to limit this overhead is to increase the sequence cache, if you are using Oracle sequences. Because the difference between sequence values generated by different instances increases, successive index block splits tend to create instance affinity to index leaf blocks. For example, suppose that an index key value is generated by a CACHE NOORDER sequence and each index leaf block can hold 500 rows. If the sequence cache is set to 50000, while instance 1 inserts values 1, 2, 3, and so on, instance 2 concurrently inserts 50001, 50002, and so on. After some block splits, each instance writes to a different part of the index tree.

So, what is the ideal value for a sequence cache to avoid inter-instance leaf index block contention, yet minimizing possible gaps? One of the main variables to consider is the insert rate: the higher it is, the higher must be the sequence cache. However, creating a simulation to evaluate the gains for a specific configuration is recommended.

Note: By default, the cache value is 20. Typically, 20 is too small for the preceding example.

D60488GC11

Oracle 11g RAC and Grid Infrastructure Administration Accelerated 14 24

### Which two options can you use to ensure that the redo logs from all the instances are available during RMAN recovery from any instance?

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

You are managing a three-instance Oracle RAC database which uses a Cluster File System for shared storage.

Which two options can you use to ensure that the redo logs from all the instances are available during RMAN recovery from any instance?

A.

Set the LOG\_ARCHIVE\_DEST\_1 parameter to a single location on the Cluster File System for each of the three instances and leave the LOG\_ARCHIVE\_FORMAT parameter with the default value.

B.

Set only the LOG\_ARCHIVE\_DEST\_1 parameter to a single location on the Cluster File System for any one of the three instances and leave the LOG\_ARCHIVE\_FORMAT parameter with the default value.

C.

Set the LOG\_ARCHIVE\_DEST\_1 parameter to a single location on the Cluster File System for each of the three instances, and the LOG\_ARCHIVE\_FORMAT parameter for each instance, to the same format including the thread number.

D.

Set the LOG\_ARCHIVE\_DEST\_1 parameter to a single location on the Cluster File System, and set LOG\_ARCHIVE\_FORMAT parameter on only one instance.

#### Explanation:

Initialization Parameter Settings for the Cluster File System Archiving Scheme

In the cluster file system scheme, each node archives to a directory that is identified with the same name on all instances within the cluster database (/arc\_dest, in the following example). To configure this directory, set values for the LOG\_ARCH\_DEST\_1 parameter, as shown in the following example:

```
*LOG_ARCHIVE_DEST_1="LOCATION=/arc_dest"
```

Archived Redo Log File Conventions in Oracle RAC

For any archived redo log configuration, uniquely identify the archived redo logs with the LOG\_ARCHIVE\_FORMAT parameter. The format of this parameter is operating system-specific and it can include text strings, one or more variables, and a filename extension.

Use the %R or %r parameters to include the resetlogs identifier to avoid overwriting the logs from a previous incarnation. If you do not specify a log format, then the default is operating system-specific and includes %t, %s, and %r.

Oracle Real Application Clusters Administration and Deployment Guide

11g Release 2 (11.2)

### What are two possible reasons for this wait event?

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

You are managing a three-Instance RAC database. In the AWR report, you notice a gc current block busy wait event on one of the database Instances, in the cluster top timed events section.

What are two possible reasons for this wait event?

A.

Access to blocks was delayed as the blocks were pinned in exclusive mode for updates by some sessions.



- B. **Access to blocks was delayed by log writes on remote Instance.**
- C. Frequently used select statements are causing high disk Input/output contention.
- D. CPU shortages.

Explanation:

Analyzing Performance Using GCS and GES Statistics

The gc current block busy wait event indicates that the access to cached data blocks was delayed because they were busy either in the remote or the local cache. This could be caused by any of the following:

The blocks were pinned

The blocks were held up by sessions

The blocks were delayed by a log write on a remote instance

A session on the same instance was already accessing a block which was in transition between instances and the current session needed to wait behind it (for example, gc current block busy)

Contention-Related Wait Events

The main wait events for contention-related waits are:

gc current block busy

gc cr block busy

gc buffer busy acquire/release

The contention-related wait event statistics indicate that a block was received which was pinned by a session on another node, was deferred because a change had not yet been flushed to disk or because of high concurrency, and therefore could not be shipped immediately. A buffer may also be busy locally when a session has already initiated a cache fusion operation and is waiting for its

completion when another session on the same node is trying to read or modify the same data.

High service times for blocks exchanged in the global cache may exacerbate the contention, which can be caused by frequent concurrent read and write accesses to the same data.

The gc current block busy and gc cr block busy wait events indicate that the local instance that is making the request did not immediately receive a current or consistent read block. The term busy in these events' names indicates that the sending of the block was delayed on a remote instance.

For example, a block cannot be shipped immediately if Oracle Database has not yet written the redo for the block's changes to a log file.

Oracle Real Application Clusters Administration and Deployment Guide  
11g Release 2 (11.2)

**Which three predefined database-type templates already include the data files when creating an Oracle Cluster database?**

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

Which three predefined database-type templates already include the data files when creating an Oracle Cluster database?

- A. **General Purpose**
- B. **Transaction Processing**
- C. Custom Database
- D. **Data Warehouse**
- E. OLAP Database

Explanation:

Creating an Oracle RAC Database with DBCA

Oracle ships templates for the following two workload types:

General purpose or transaction processing

Data warehouse

For more complex environments, you can select the Custom Database option. This option does not use templates and results in a more extensive interview, which means that it takes longer to create your database.

Oracle Real Application Clusters Installation Guide  
11g Release 2 (11.2) for Linux and UNIX

## Which two statements are true regarding the attributes of RACDB?

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

Examine the following output:

```
[oracle@gr5153 ~]$ sudo crsctl config crs
CRS-4622: Oracle High Availability Services autostart is enabled.
[oracle@gr5153 ~]$ srvctl config database -d RACDB -a
Database unique name: RACDB
Database name: RACDB
Oracle home : /u01/app/oracle/product/11.1.2.0/dbhome_1
Oracle user: oracle
Spfile: +DATA/RACDB/spfileRACDB.ora
Domain:
Start options: open
Stop options: immediate
Database role: PRIMARY
Management policy: AUTOMATIC
Server pools: POOL1
Database instances:
Disk Groups: DATA, FRA
Services:
Database is enabled
Database is policy managed
Oracle Clusterware is started automatically after the system boot. Which two statements are true
regarding the attributes of RACDB?
```

- A.  
**Oracle Clusterware automatically starts RACDB.**
- B.  
You must manually start RACDB.
- C.  
**Database resource is managed by crsd for high availability and may be automatically restarted in place if it fails.**
- D.  
Database resource Is not managed by crsd for high availability and needs to be restarted manually if it fails.

### Explanation:

Switch Between the Automatic and Manual Policies

By default, Oracle Clusterware is configured to start the VIP, listener, instance, ASM, database services, and other resources during system boot. It is possible to modify some resources to have their profile parameter AUTO\_START set to the value 2. This means that after node reboot, or when Oracle Clusterware is started, resources with AUTO\_START=2 need to be started manually via srvctl. This is designed to assist in troubleshooting and system maintenance. When changing resource profiles through srvctl, the command tool automatically modifies the profile attributes of other dependent resources given the current prebuilt dependencies. The command to accomplish this is:

```
srvctl modify database -d <dbname> -y AUTOMATIC|MANUAL
D60488GC11
```

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 15 3

3.4.1 Benefits of Using Oracle Clusterware

Oracle Clusterware provides the following benefits:

Tolerates and quickly recovers from computer and instance failures.

Simplifies management and support by means of using Oracle Clusterware together with Oracle Database.

By using fewer vendors and an all Oracle stack you gain better integration compared to using third-party clusterware.

Performs rolling upgrades for system and hardware changes. For example, you can apply Oracle Clusterware upgrades, patch sets, and interim patches in a rolling fashion, as follows:

Upgrade Oracle Clusterware from Oracle Database 10g to Oracle Database 11g

Upgrade Oracle Clusterware from Oracle Database release 11.1 to release 11.2

Patch Oracle Clusterware from Oracle Database 11.1.0.6 to 11.1.0.7

Patch Oracle Clusterware from Oracle Database 10.2.0.2 Bundle 1 to Oracle Database 10.2.0.2 Bundle 2

Automatically restarts failed Oracle processes.

Automatically manages the virtual IP (VIP) address so when a node fails then the node's VIP address fails over to another node on which the VIP address can accept connections.

Automatically restarts resources from failed nodes on surviving nodes.  
Controls Oracle processes as follows:  
For Oracle RAC databases, Oracle Clusterware controls all Oracle processes by default.  
For Oracle single-instance databases, Oracle Clusterware allows you to configure the Oracle processes into a resource group that is under the control of Oracle Clusterware.  
Provides an application programming interface (API) for Oracle and non-Oracle applications that enables you to control other Oracle processes with Oracle Clusterware, such as restart or react to failures and certain rules.  
Manages node membership and prevents split-brain syndrome in which two or more instances attempt to control the database.  
Provides the ability to perform rolling release upgrades of Oracle Clusterware, with no downtime for applications.  
Oracle Database High Availability Overview  
11g Release 2 (11.2)

### Which four statements are true about services and connection load balancing?

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

Which four statements are true about services and connection load balancing?

- A.  
**Services provide the ability to balance client connections across the instances in an Oracle RAC configuration.**
- B.  
**There are two types of load balancing that you can implement; load balancing decisions can be made by the client and by the listener.**
- C.  
In a RAC database environment, you should select only one load balancing method.
- D.  
**With server-side load balancing, the listener directs a connection request to the best instance that is currently providing the service based on the load-balancing options selected by the DBA.**
- E.  
**Client-side load balancing balances the connection requests across the listeners.**

#### Explanation:

Services and Connection Load Balancing

The two load balancing methods that you can implement are:

Client-side load balancing: Balances the connection requests across the listeners

Server-side load balancing: The listener directs a connection request to the best instance currently providing the service by using the load balancing advisory (LBA).

FAN, Fast Connection Failover, and LBA depend on a connection load balancing configuration that includes setting the connection load balancing goal for the service.

The load balancing goal for the service can be either:

LONG: For applications having long-lived connections. This is typical for connection pools and SQL\*Forms sessions.

SHORT: For applications that have short-lived connection

Oracle Net Services provides the ability to balance client connections across the instances in an Oracle RAC configuration. You can implement two types of load balancing: client-side and server-side.

Client-side load balancing balances the connection requests across the listeners.

With server-side load balancing, the listener directs a connection request to the best instance currently providing the service by using the load balancing advisory.

In a RAC database, client connections should use both types of connection load balancing.

D60488GC11

Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 15 12

### Which three steps are required to enable ODP.NET clients that connect to the GL service to receive FAN High Availability Events?

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

You are managing a three-instance policy-managed RAC database PROD. You created a service called GL for the PROD database by using the following command:

```
[oracle@gr7597~]$ srvctl add service -d PROD -s GL -g SP1 -c singleton -y manual
```

Examine the following output:

```
[oracle@gr7597~]$ srvctl config database -d PROD -a
```

Database unique name: PROD

Database name: PROD

Oracle home: /u01/app/oracle/product/11.2.0/dbhome\_1  
Oracle user: oracle  
Spfile: +DATA/PROD/spfileRACDB.ora  
Domain:  
Start options: open  
Stop options: immediate  
Database role: PRIMARY  
Management policy: AUTOMATIC  
Server pools: SP1  
Database instances:  
Disk Groups: DATA, FRA  
Services:

Database is enabled  
Database is policy managed  
[oracle@gr7597~]\$ crsctl stat res ora.PROD.db  
NAME=ora.PROD.db  
TYPE=cluster\_resource  
TARGET=ONLINE  
STATE=ONLINE on gr7597 ONLINE on gr7602 ONLINE on gr7633  
[oracle@gr7597~]\$ crsctl stat serverpool ora.SP1  
NAME=ora.SP1  
ACTIVE\_SERVERS=gr7597 gr7602 gr7633  
[oracle@gr7597~]\$ crsctl stat res ora.prod.gl.svc  
NAME=ora.prod.gl.svc  
TYPE=ora.service.type  
TARGET=ONLINE  
STATE=ONLINE on gr7597  
Which three steps are required to enable ODP.NET clients that connect to the GL service to receive FAN High Availability Events?

- A.  
Enable Advanced Queuing notifications by using SRVCTL as shown in the following command:  
`$ srvctl modify service -d prod-s gl -q TRUE -j LONG`
- B.  
Execute the following statement for the users that will be connecting by way of the .Net Application, where user\_name is the user name:  
`EXECUTE DBMS_AQADM.GRANT_QUEUE_PRIVILEGE('DEQUEUE', 'SYS.SYSSERVICE_METRICS', user_name);`
- C.  
Enable Transparent Application Failover (TAF), either on the client or for the service.
- D.  
Enable Fast Connection Failover for ODP.NET connection pools by subscribing to FAN High Availability events. Do this by setting the HA events connection string attribute to true at connection time.
- E.  
Link client applications with the client thread or operating system library.

Explanation:

Enabling ODP.NET Clients to Receive FAN High Availability Events

Perform the following steps to enable FAN for ODP.NET clients:

Enable Advanced Queuing notifications for a service by using SRVCTL as shown in the following example:

```
srvctl modify service -d crm -s odpnet.example.com -q TRUE
```

Grant permissions on the internal event queue table by executing the following command for the users that will be connecting by way of the ODP.NET application, where user\_name is the database user name:

```
EXECUTE
```

```
DBMS_AQADM.GRANT_QUEUE_PRIVILEGE('DEQUEUE','SYS.SYSSERVICE_METRICS', user_name);
```

Enable Fast Connection Failover for ODP.NET connection pools by subscribing to FAN high availability events. To enable FCF, include "HA Events=true" and "pooling=true" (the default value) in the connection string, as shown in the following example where user\_name is the name of the database user and password is the password for that user:

```
con.ConnectionString =
```

```
"User Id=user_name;Password=password;Data Source=odpnet;" +
```

“Min Pool Size=10;Connection Lifetime=120;Connection Timeout=60;” +  
“HA Events=true;Incr Pool Size=5;Decr Pool Size=2”;  
Oracle Real Application Clusters Administration and Deployment Guide  
11g Release 2 (11.2)

### Which three statements are correct?

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

Examine the Exhibit.

Which three statements are correct?

A.

**Global Cache Block Access Latency indicates how many block transfers incurred a delay (busy) or an unexpected longer delay (congested).**

B.

Global Cache Block Access Latency indicates how many block transfers received blocks cached by other instances as well as blocks read from disk.

C.

**Global Cache Block Transfer Rate shows the total aggregated number of blocks received by all instances in the cluster over the interconnect.**

D.

Global Cache Block Transfer Rate is essentially a profile of how much work is performed in the local buffer cache, rather than the portion of remote references and physical reads, which both have higher latencies.

E.

**Global Cache Block Transfers and Physical Reads shows the percentage of read operations that retrieved blocks from disk, and the buffer cache of other instances using Cache Fusion.**

F.

Global Cache Block Transfers and Physical Reads shows the aggregated number of data blocks received by all instances in the cluster over the interconnect.

#### Explanation:

**Global Cache Block Access Latency**

The global cache access latency represents the end-to-end elapsed time for a block request. The request is timed from when the request is initiated until it completes. Cache transfer indicates how many current and CR blocks per block class were received from remote instances, including how many transfers incurred a delay (busy) or an unexpected longer delay (congested).

If accessing a database block of any class does not locate a buffered copy in the local cache, a global cache operation is initiated. Before reading a block from disk, an attempt is made to find the block in the buffer cache of another instance. If the block is present in another instance, a version of the block may be shipped. Two different kinds of blocks are distinguished: current and consistent read blocks. The average block receive time represents the end-to-end elapsed time or latency for a block request.

**Global Cache Block Transfer Rate**

The global cache block transfer rate shows the total aggregated number of data blocks received by all instances in the cluster by way of an interconnect. If a logical read fails to find a copy of the buffer in the local cache, it attempts to find the buffer in the database cache of a remote instance. If the block is present in another database, it is sent to the current instance.

**Global Cache Block Transfers and Physical Reads**

The chart represents the percentage of logical reads that read data from the buffer cache of other instances via Direct Memory Access and from disk. It is essentially a profile of how much work is performed in the local buffer cache, rather than the portion of non-local references that incur some latency overhead.

**Active Sessions for Cluster Wait Class**

The chart represents the active sessions for the cluster wait class. You can view the top modules and top SQL for a time period by selecting the time period in the active sessions chart.

Oracle Enterprise Manager Online Help

### What should be the database mode on each of the three instances for this backup to succeed?

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

orc1, orc2, and orc3 are tnsnames that connect to predefined instances of a three-instance RAC database which uses ASM for shared storage.

The database is in archivelog mode.

Examine the Exhibit for the tnsnames.ora entries for orc1, orc2, and orc3.

```

orcl1 =
(DESCRIPTION =
  (ADDRESS_LIST =
    (ADDRESS = (PROTOCOL = TCP) (HOST = host01) (PORT=1521))
  )
  (CONNECT_DATA =
    (SERVICE_NAME = orcl)
    (INSTANCE_NAME = orcl1)
  )
)
orcl2 =
(DESCRIPTION =
  (ADDRESS_LIST =
    (ADDRESS = (PROTOCOL = TCP) (HOST = host02) (PORT=1521))
  )
  (CONNECT_DATA =
    (SERVICE_NAME = orcl)
    (INSTANCE_NAME = orcl2)
  )
)
orcl3 =
(DESCRIPTION =
  (ADDRESS_LIST =
    (ADDRESS = (PROTOCOL = TCP) (HOST = host03) (PORT=1521))
  )
  (CONNECT_DATA =
    (SERVICE_NAME = orcl)
    (INSTANCE_NAME = orcl3)
  )
)

```

You plan to execute the following RMAN commands to back up the database:

```

RUN{
ALLOCATE CHANNEL CH1 DEVICE TYPE DISK CONNECT 'user1/pwd1@orcl1';
ALLOCATE CHANNEL CH2 DEVICE TYPE DISK CONNECT 'user2/pwd2@orcl2';
ALLOCATE CHANNEL CH3 DEVICE TYPE DISK CONNECT 'user3/pwd3@orcl3';
BACKUP DATABASE PLUS ARCHIVELOG; }

```

What should be the database mode on each of the three instances for this backup to succeed?

- A.  
The database can be either mounted or open on all three instances.
- B.  
The database must be open on all three instances.
- C.  
The database must be open on any one instance and mounted on the other two instances.
- D.  
The database must be mounted on any one of the instances and the other instances will be mounted automatically.

Explanation:

connectStringSpec

Purpose

Use the connectStringSpec subclause to specify the user name, password, and net service name for connecting to a target, recovery catalog, or auxiliary database. The connection is necessary to authenticate the user and identify the database.

Prerequisites

You must have SYSDBA privileges to CONNECT to a target or auxiliary database. Do not connect to the recovery catalog database as user SYS.

Usage Notes

RMAN connections to a database are specified and authenticated in the same way as SQL\*Plus connections to a database. The only difference is that RMAN connections to a target or auxiliary database require the SYSDBA privilege. The AS SYSDBA keywords are implied and cannot be explicitly specified. See Oracle

Database Administrator's Guide to learn about database connection options when using SQL\*Plus.

Oracle Database Backup and Recovery Reference  
11g Release 2 (11.2)

What are two possible causes for this?

The Global Cache Block Access Latency chart shows high elapsed times. What are two possible causes for this?

- A. **badly written SQL statements**
- B. storage network bottlenecks
- C. a large number of requested blocks not cached in any instance
- D. **slow or faulty interconnect**

Explanation:

About Global Cache Block Access Latency Chart

If the Global Cache Block Access Latency chart shows high latencies (high elapsed times), then this can be caused by any of the following:

A high number of requests caused by SQL statements that are not tuned.

A large number of processes in the queue waiting for the CPU, or scheduling delays.

Slow, busy, or faulty interconnects. In these cases, check your network connection for dropped

packets, retransmittals, or cyclic redundancy check (CRC) errors.

Oracle Database 2 Day + Real Application Clusters Guide

11g Release 2 (11.2)

**Which three steps must be completed to configure Fast Connection Failover for the CRM service with an Oracle Call Interface client?**

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

Which three steps must be completed to configure Fast Connection Failover for the CRM service with an Oracle Call Interface client?

- A. **Ensure that the CRM service has Advanced Queuing notifications enabled by setting the services values of AQ\_HA\_NOTIFICATIONS to TRUE.**  
For example:  

```
$ srvctl modify service -d crm -s gl.us.oracle.com \
-q TRUE -P BASIC -e SELECT -z 180 -w 5 -j LONG
```
- B. **Enable OCI\_EVENTS on the client at environment creation time by using the OCIEnvCreate(🔗) function.**
- C. Use Oracle Call Interface to configure UP events for the service.
- D. **Client applications must be linked with the client thread or operating system library.**

Explanation:

Configuring FCF for OCI Clients

OCI applications must connect to an Oracle RAC instance to enable HA event notification.

Furthermore, these applications must perform the following steps to configure FCF for an OCI client:

Configure the service for your OCI connection pool to enable connection load balancing and runtime connection load balancing. Also configure the service to have Advanced Queuing notifications enabled, as shown in the following example:

```
$ srvctl modify service -d crm -s ociapp.example.com -q TRUE -B THROUGHPUT -j LONG
```

Set the context of the environment for OCI calls on the client to enable subscriptions by using the OCIEnvCreate() call, setting the MODE parameter value to OCI\_EVENTS, as follows:

```
(void) OCIEnvCreate(&myenvhp, OCI_EVENTS|OCI_OBJECT, ...);
```

Link the application with a thread library.

After linking with the thread library, the applications can register a callback that is invoked whenever a FAN event occurs.

Oracle Real Application Clusters Administration and Deployment Guide

11g Release 2 (11.2)

**Assuming a RAC database called ORCL, select three statements that correctly demonstrate management actions for the AP service.**

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

Assuming a RAC database called ORCL, select three statements that correctly demonstrate

management actions for the AP service.

A.

To start the AP service, execute:  
srvctl start service -d ORCL -s AP

B.

To disable the AP service on the ORCL4 instance, execute:  
srvctl disable service -d ORCL -s AP -i ORCL4

C.

To stop the AP service, execute:  
srvctl stop service -s AP

D.

To make ORCL5 a preferred instance for the AP service, execute:  
srvctl set service -d ORCL -s AP -i ORCL5 -r

E.

To relocate the AP service from the ORCL5 instance to the ORCL4 instance, execute:  
srvctl relocate service -d ORCL -s AP -i ORCL5 -t ORCL4

Explanation:

SRVCTL Command Reference

```
srvctl start service -d db_unique_name  
[-s "service_name_list" [-n node_name | -i instance_name]]  
[-o start_options]  
srvctl disable service -d db_unique_name  
-s "service_name_list" [-i instance_name | -n node_name]  
srvctl stop service -d db_unique_name [-s "service_name_list"  
[-n node_name | -i instance_name] [-f]]  
srvctl relocate service -d db_unique_name -s service_name  
{-c source_node -n target_node | -i old_instance_name -t new_instance_name}  
[-f]
```

Oracle Real Application Clusters Administration and Deployment Guide  
11g Release 2 (11.2)

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

Posted by seenagape on [March 11, 2013](#)

[No comments](#)

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

Explanation:

Concurrent Cross-Instance Calls: Considerations

In data warehouse and data mart environments, it is not uncommon to see a lot of TRUNCATE operations.

These essentially happen on tables containing temporary data. In a RAC environment, truncating tables concurrently from different instances does not scale well, especially if, in conjunction, you are also using direct read operations such as parallel queries.

As shown in the slide, a truncate operation requires a cross-instance call to flush dirty blocks of the table that may be spread across instances. This constitutes a point of serialization. So, while the first TRUNCATE command is processing, the second has to wait until the first one completes. There are different types of cross-instance calls. However, all use the same serialization mechanism.

For example, the cache flush for a partitioned table with many partitions may add latency to a corresponding parallel query. This is because each cross-instance call is serialized at the cluster level, and one crossinstance call is needed for each partition at the start of the parallel query for direct read purposes.

D60488GC11



Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 14 27

What Application Design considerations should I be aware of when moving to Oracle RAC?

The general principals are that fundamentally no different design and coding practices are required for RAC however application flaws in execution or design have a higher impact in RAC. The performance and scalability in RAC will be more sensitive to bad plans or bad schema design. Serializing contention makes applications less scalable. If your customer uses standard SQL and schema tuning, it solves > 80% of performance problems

Some of the scalability pitfalls they should look for are:

\* Serializing contention on a small set of data/index blocks

-> monotonically increasing key

-> frequent updates of small cached tables

-> segment without automatic segment space management (ASSM) or Free List Group (FLG)

\* Full table scans

-> Optimization for full scans in 11g can save CPU and latency

\* Frequent invalidation and parsing of cursors

-> Requires data dictionary lookups and synchronizations

\* Concurrent DDL ( e.g. truncate/drop )

Look for:

\* Indexes with right-growing characteristics

-> Use reverse key indexes

-> Eliminate indexes which are not needed

\* Frequent updated and reads of small tables

-> small=fits into a single buffer cache

-> Use Sparse blocks ( PCTFREE 99 ) to reduce serialization

\* SQL which scans large amount of data

-> Perhaps more efficient when parallelized

-> Direct reads do not need to be globally synchronized ( hence less CPU for global cache )

RAC: Frequently Asked Questions [ID 220970.1]