

Quick Installation Guide Oracle 9i RAC on IBM server pSeries on AIX 5L (with HACMP/GPFS)

Version 3.2
November 2003

EMEA Oracle/IBM
Joint Solutions Center
MONTPELLIER - FRANCE



Document history :

Version	Date	Update	Who	Validated by
1.00		Creation	Fabienne Lepetit Michel Passet	Paul Bramey Alain Benhaim
1.01		Corrections	Fabienne Lepetit	Paul Bramey Alain Benhaim
2.1	May 2002	Update	Michel Passet	
3.1	October 2003	Update for HACMP 4.5 Add GPFS 2.1 implementation	Fabienne Lepetit Thierry Plumeau	Paul Bramey Alain Benhaim
3.2	November 2003	Updates for GPFS 2.1	Rick Piasecki Thierry Plumeau	Paul Bramey

Contributors :

- Alain Benhaim : IBM France
- Paul Bramey : Oracle France
- Fabienne Lepetit : Oracle France
- Michel Passet : IBM France
- Rick Piasecki : IBM US
- Thierry Plumeau : IBM France

1	THE AIM OF THIS DOCUMENT	4
2	HARDWARE ARCHITECTURE	5
2.1	Raw devices implementation	5
2.2	GPFS implementation	6
3	INSTALLATION STEPS.....	7
4	PREPARING THE SYSTEM	8
4.1	Hardware requirements	8
4.2	Software requirements.....	8
4.3	Concurrent disks requirements.....	9
4.3.1	Enterprise Storage Server (ESS Disks Bay)	9
4.3.2	SSA external disks microcode	9
4.3.3	SSA adapters microcode	9
4.4	Users and Groups.....	9
5	NETWORK CONFIGURATION.....	10
6	INSTALLING HACMP/ES/CRM ON A P SERIES CLUSTER.....	11
6.1	Patches requirements	12
6.2	Post install tasks	12
7	INSTALLING GPFS	13
8	CREATING A CONCURRENT VOLUME GROUP.....	14
9	CREATING THE CONCURRENT LOGICAL VOLUMES.....	19
10	CONFIGURING HACMP - raw-devices implementation.....	23
11	CONFIGURING HACMP - GPFS implementation.....	35
12	STARTING HACMP	45
13	Configuring and STARTING GPFS	49
14	ORACLE ENVIRONMENT SETUP	53
15	INSTALL THE ORACLE SOFTWARE.....	55
16	CREATING THE DATABASE MANUALLY	66
17	CREATING THE DATABASE USING DBCA	67
17.1	gsd setup	67
17.2	srvctl tool setup	67
17.3	The database configuration assistant (DBCA)	68
17.4	Startup/shutdown services.....	69
17.5	Configure listener.ora / sqlnet.ora / tnsnames.ora	69
17.6	Configure Oracle Enterprise Manager.....	69
17.7	Troubleshooting	70
18	Appendix A : Installation Directories.....	71
19	Appendix B : SSA disk & adapter microcode management	72
20	Appendix C : Logical volumes creation.....	74
21	Appendix D : examples of configuration files.....	75
22	Appendix E : Oracle technical notes.....	78
23	Appendix F : HACMP Cluster verification output.	81
24	Appendix G : RAC Database creation script & init.ora.....	83
25	Appendix H : Filesets to be installed on the machines of the cluster.	85
25.1	HACMP 4.4	85
25.2	HACMP 4.5	85
25.3	RSCT 2.2 (provided with AIX 5.1)	86
25.4	RSCT 2.3 (provided with AIX 5.2)	86
25.5	GPFS 2.1	87

1 THE AIM OF THIS DOCUMENT

This document is written to provide help to install Oracle9i Real Application Clusters (9.0.2.1.0) release 3 on IBM pSeries servers with AIX 5L through a description of the different steps necessary to run Oracle9i RAC. Metalink (http://metalink.oracle.com/metalink/plsql/ml2_gui.startup)

Titre	Origine	Référence
Oracle9i Installation Guide Release 2 (9.2.0.1.0) for UNIX Systems	Oracle Metalink	A96167-01
Oracle9i Administrator's Reference Release 2 (9.2.0.1.0) for UNIX Systems	Oracle Metalink	A97297-01
Oracle9i Release Notes Release 2 (9.2.0.1.0) for AIX-Based 5L Systems (64-bit)	Oracle Metalink	A97605-01
Oracle Universal Installer Concepts Guide Release 2.2	Oracle Metalink	A96697-01
Oracle9i Real Application Clusters Setup and Configuration Release 2 (9.2)	Oracle Metalink	A96600-01
GPFS for AIX5L – AIX Clusters Concepts, Planning and Installation Guide	IBM	GA22-7895-01
GPFS for AIX5L – AIX Clusters Administration and Programming Reference	IBM	SA22-7896-01
GPFS on AIX Clusters – High Performance File System – Administration Simplified	IBM	SG24-6035-00

The information contained in this paper resulted from :

- Oracle and IBM documentations
- Workshop experiences done in the Oracle/IBM Joint Solutions Center
- Benchmarks and POC implementations for customers performed by EMEA PSSC Montpellier
- This documentation is a joint effort from Oracle and IBM specialists.

Please also refer to Oracle online documentation for more information :

<http://docs.oracle.com>

<http://tahiti.oracle.com>

http://technet.oracle.com/docs/products/oracle9i/doc_library/release2/index.htm

<http://otn.oracle.com/products/oracle9i/content.html>

Oracle9i RAC home page : http://www.oracle.com/ip/index.html?rac_home.html

For HACMP Documentation refer to : http://www-1.ibm.com/servers/eserver/pseries/library/hacmp_docs.html

For GPFS Documentation refer to : <http://www-1.ibm.com/servers/eserver/pseries/library/gpfs.html>

Your comments are important for us. We want our technical papers to be as helpful as possible.

Please send us your comments about this document to the EMEA Oracle/IBM Joint Solutions Center.

Use our email address :

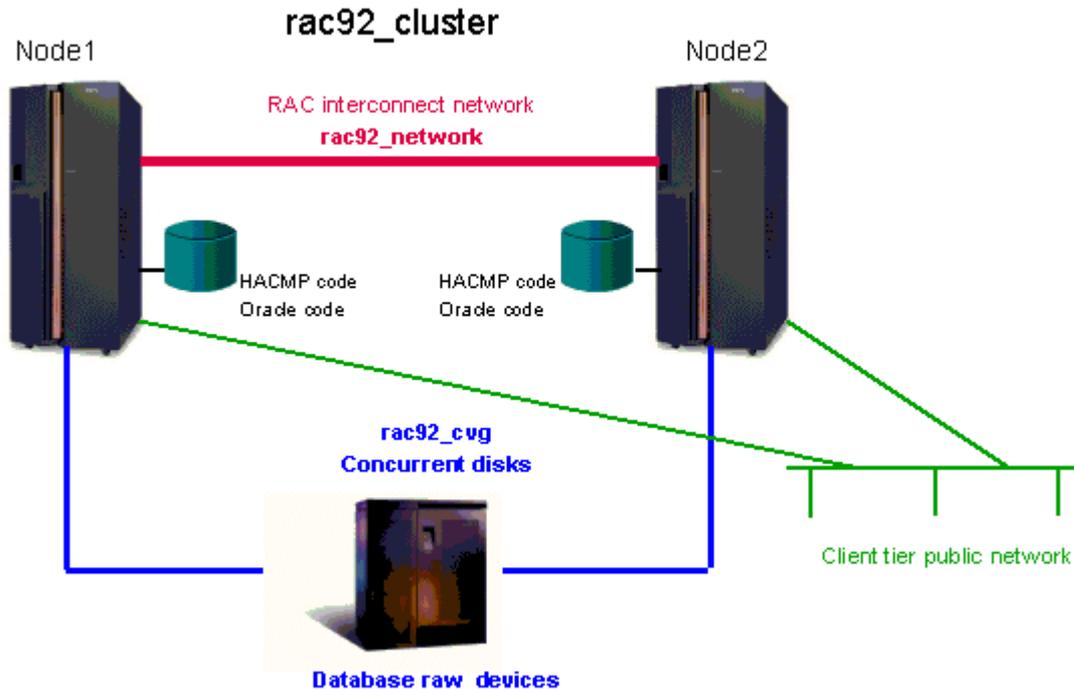
oraclibm@fr.ibm.com

or our phone number :

+33 (0)4 67 34 67 49

2 HARDWARE ARCHITECTURE

2.1 RAW DEVICES IMPLEMENTATION



The cluster, named **rac92_cluster**, is composed of two IBM **@server** pSeries using AIX 5L.

The IBM AIX clustering layer, HACMP (High Availability Cluster Multi-Processing) is installed on both machines, on internal disks.

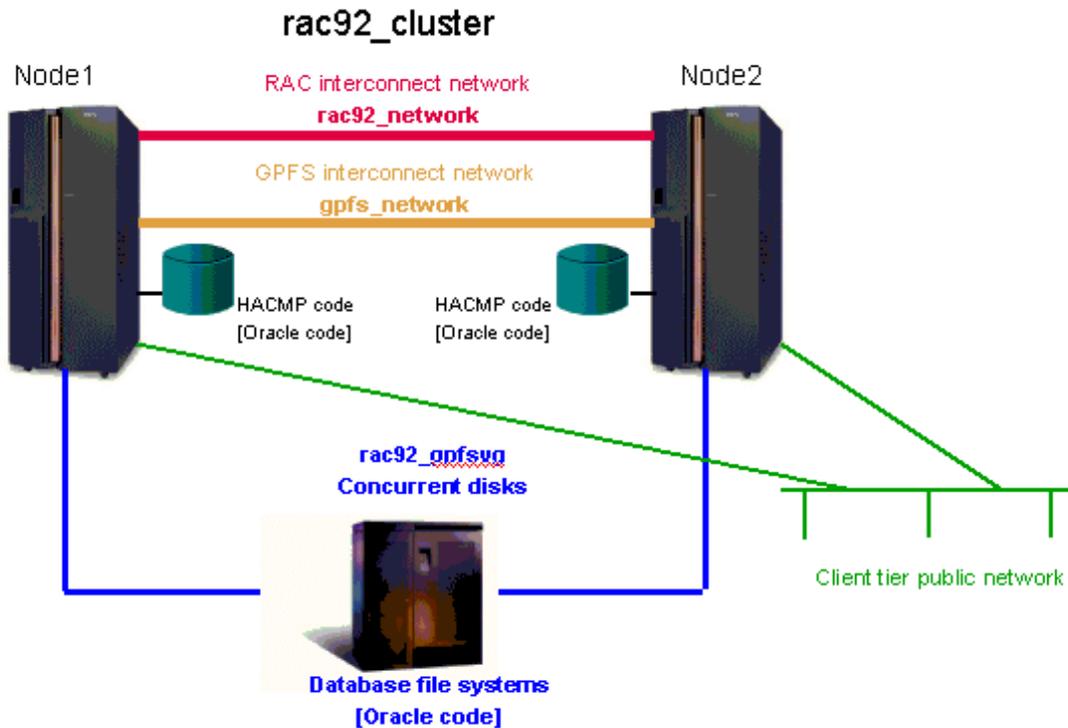
A private network (for instance a gigabit ethernet point to point network) is designed only for Oracle interconnect use (cache fusion between instances). This private network is mandatory. A proprietary switch can also be used. A second gigabit ethernet interconnect, with a different network mask, can be setup for security purposes, but it is not mandatory.

The client part is connected to the nodes with a standard network, totally separated from the interconnect network.

Oracle9i RAC code is installed on the two machines of the cluster, on internal disks

The database files are stored in raw devices, on concurrent external disks (SSA disks or ESS disk bay)

2.2 GPFS IMPLEMENTATION



GPFS (General Parallel File System) provides global access to data from any of the hosts within a cluster.

The cluster, named rac92_cluster, is composed of two IBM **@server** pSeries using AIX 5L. The IBM AIX clustering layer, HACMP, is installed on both machines, on internal disks.

A private network (for instance a gigabit ethernet point to point network) is designed only for Oracle interconnect use (cache fusion between instances). This private network is mandatory. A proprietary switch can also be used. A second gigabit ethernet interconnect, with a different network mask, can be setup for security purposes, but it is not mandatory.

Another private network is strongly recommended for GPFS (for instance a gigabit ethernet point to point network).

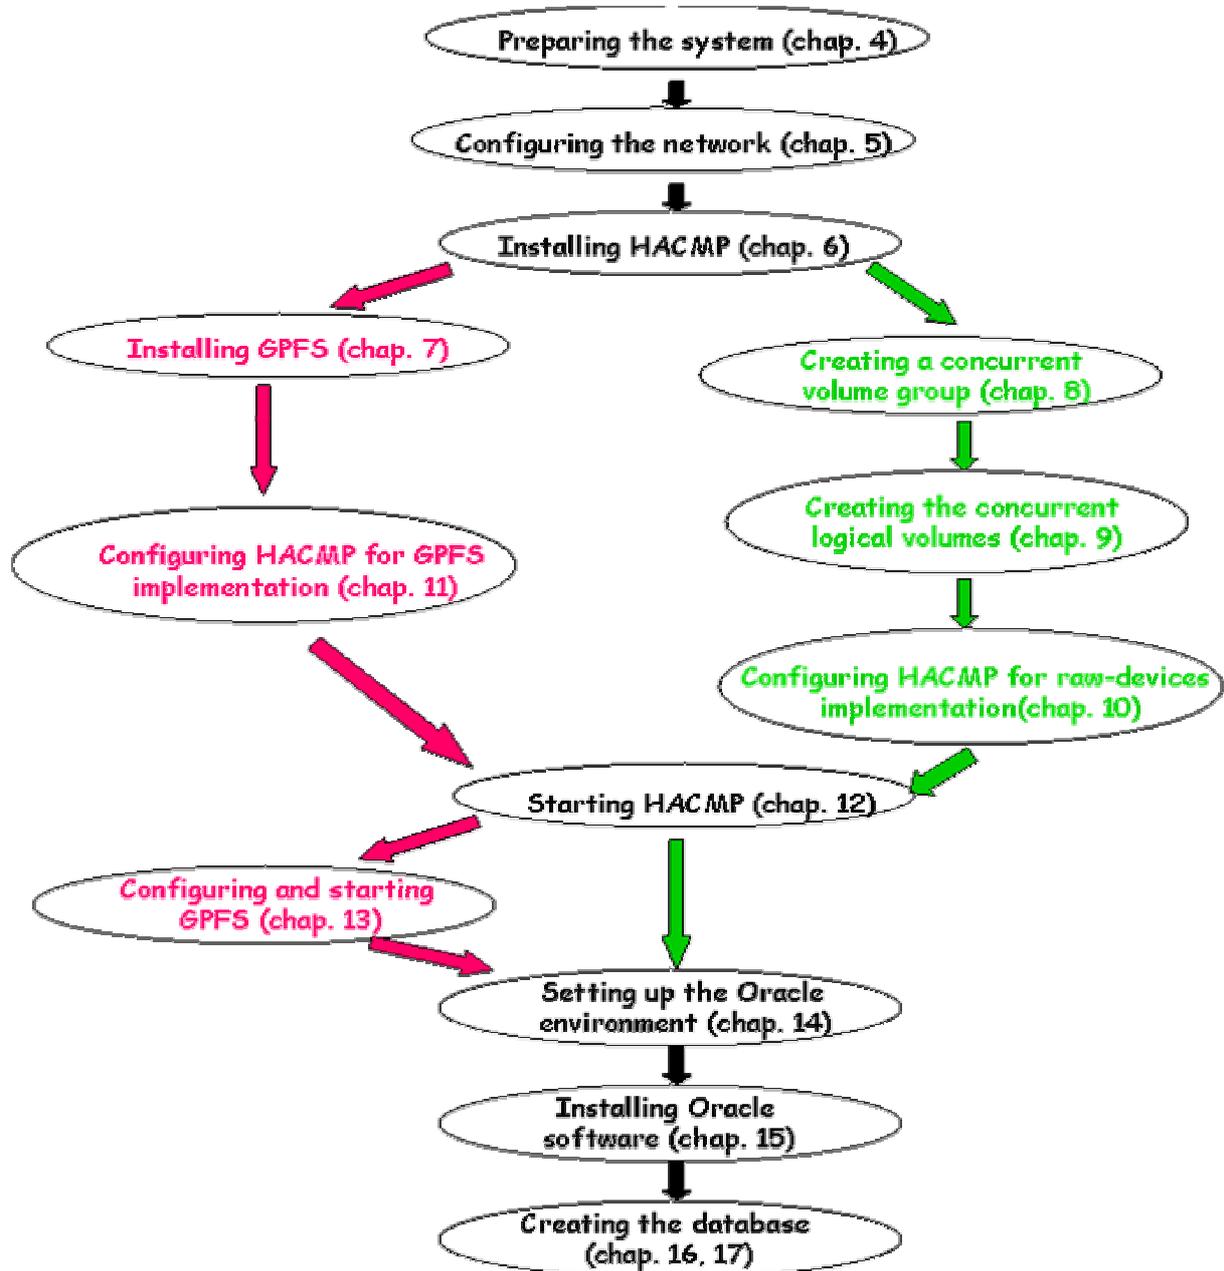
The client part is connected to the nodes with a standard network, totally separated from the interconnect network.

Oracle9i RAC code can be either installed on the two machines of the cluster, on internal disks or installed on concurrent external disks on a GPFS file system.

The database files are stored in GPFS file systems, on concurrent external disks (SSA disks or ESS disk bay).

3 INSTALLATION STEPS

Each box refers to a chapter.



Legend :
 common (raw + gpfs) tasks
 raw-device implementation tasks
 gpfs implementation tasks

4 PREPARING THE SYSTEM

4.1 HARDWARE REQUIREMENTS

- RAM \geq 512 MB
Command to check the physical memory : `lsattr -El sys0 -a realmem`
- Internal disk \geq 4 GB for the oracle code
- Paging space = 2 x RAM, with a minimum of 400 MB
Nota : for big servers with more than 8 GB of memory, the paging space could be less than 2 x RAM.
To check the paging space configured : `lspcs -a`
- Temporary Disk Space: The Oracle Universal Installer requires up to 800 MB of free space in the /tmp directory.
To check the free temporary space available : `df -k /tmp`
You can use an other filesystem instead of /tmp. Set the TEMP environment variable (used by Oracle) and the TMPDIR environment variable to the new location.
For example :
`export TEMP=/new_tmp`
`export TMPDIR=/new_tmp`
- CD-ROM drive. If the server does not have a CD-ROM unit (like a SP node for example), it is still possible to mount via NFS the drive from an other machine on the network.

4.2 SOFTWARE REQUIREMENTS

- AIX 5L (5.1 ML3 or 5.2 ML2)
To determine the current operating system version : `oslevel`
To check the maintenance level applied : `instfix -i | grep ML`
- HACMP/ES/CRM 4.4 or 4.5
- For a database implemented on general parallel file-systems: GPFS 2.1., two cases:

AIX 5.1

- ❖ AIX 5.1 ML 1 plus PTF for APAR IY33002 or later (Both the AIX 32 and 64 bit kernel are supported).
- ❖ Oracle9i RAC Release 2 for AIX-Based 5L Systems (64-bit) or later
- ❖ GPFS 2.1 requires the PTF for APAR IY34917.
- ❖ HACMP/ES Version 4.4.1 –or – Version 4.5 with PTF for IY32874 or later.
 - GPFS 2.1 has option of HACMP or RPD (RSCT Peer Domain)
 - Oracle RAC however still has HACMP requirement.

AIX 5.2 : Special Bids with following restrictions:

- ❖ GPFS for AIX, V2.1 (with service level U488745)
 - ❖ AIX 5L V5.2 (with service level U488362)
 - ❖ HACMP/ES Version 4.5 (with service level U487607)
 - ❖ Oracle9i RAC Release 2 or later
 - ❖ 3-4 pSeries nodes
 - ❖ SAN cluster only
 - ❖ ESS disk with SDD only (no MPIO)
- It is recommended that customers running I/O intensive workloads should be using GPFS 2.1 64 bit version (AIX 64 bit kernel) and should request the latest updates (efix or PTF) for GPFS problems 411186, 411866 and 410282.

For more information : http://www-1.ibm.com/servers/eserver/pseries/software/sp/gpfs_faq.html

4.3 CONCURRENT DISKS REQUIREMENTS

4.3.1 Enterprise Storage Server (ESS Disks Bay)

Check with the system administrator that the ESS can be used with concurrent disks with HACMP/ES/CRM

4.3.2 SSA external disks microcode

It is particularly important for HACMP that all the SSA disks connected to the cluster are all at the same level of microcode.

To list the SSA disks of the cluster : `lscfg | grep pdisk`
 To check the microcode level : `lscfg -vl pdisknn` (ROS level and ID line)

For more information on the procedure of download and installation of new microcode,
 See [Appendix B : SSA disk & adapter microcode management](#).

4.3.3 SSA adapters microcode

To list the SSA adapters of a node : `lscfg | grep ssa`
 To check the microcode level : `lscfg -vp ssann`

For more information on the procedure of download and installation of new microcode,
 see [Appendix B : SSA disk & adapter microcode management](#).

4.4 USERS AND GROUPS

This setup has to be done on all the nodes of the cluster. Be sure that all the groups and user numbers (203, 204 and 205 in this example) are identical thru the nodes.

- **smit group** to create the following groups
 - dba** Primary group for the oracle user.
 - hagsuser** For high availability (if not already created).
 - oinstall** The ora inventory group. This group is not mandatory. If it exists, it will be the group owner of the oracle code files. This group is a secondary group for oracle user.
- **smit user** to create the users
 - oracle** Owner of the database.

The **oracle** user must have **dba** as primary group, **oinstall** and **hagsuser** as secondary groups.
 Also add the secondary group **hagsuser** to the **root** account.

Verification : check if the file `/etc/group` contains lines such as : (the numbers could be different)
`hagsuser:::203:oracle, root`
`dba:::204:oracle`
`oinstall:::205:oracle`

- Check if there is some AIX default limitations (especially on the file size)
 File size limitation: `ulimit -f`
 All limitations : `ulimit -a`
 See also the file `/etc/security/limits` which shows the limits for each user. The default stanza applies to all new user to be created. This file can be modified by root with vi.

The *default* limits should be set to unlimited, except for core (e.g. -1 in the file `/etc/security/limits`)
 To turn some user limitation to unlimited, use `smit users`

- Set a password to oracle user, the same for all the nodes of the cluster, with the command `passwd oracle`

5 NETWORK CONFIGURATION

Set up user equivalence for the oracle account, to enable rsh, rcp, rlogin commands.

/etc/hosts file

```
/etc/hosts for HACMP (on both nodes)

192.128.194.1    node1
192.128.194.2    node2
10.10.11.141    interconnect_node1
10.10.11.142    interconnect_node2
20.20.21.141    gpfs_node1
20.20.21.142    gpfs_node2
```

/etc/hosts.equiv file

Put the list of machines or nodes into hosts.equiv.

```
/etc/hosts.equiv for HACMP (on both nodes)

node1            root
node2            root
interconnect_node1  root
interconnect_node2  root
gpfs_node1       root
gpfs_node2       root
node1            oracle
node2            oracle
interconnect_node1  oracle
interconnect_node2  oracle
gpfs_node1       oracle
gpfs_node2       oracle
```

.rhosts file.

In the root's and oracle's home directory, put the list of machines.

\$HOME/.rhosts for HACMP configuration :

```
node1            root
node2            root
interconnect_node1  root
interconnect_node2  root
gpfs_node1       root
gpfs_node2       root
node1            oracle
node2            oracle
interconnect_node1  oracle
interconnect_node2  oracle
gpfs_node1       oracle
gpfs_node2       oracle
```

Note : It is possible, but not advised because of security reasons, to put a "+" in hosts.equiv and .rhosts files.

Test if the user equivalence is correctly set up (node2 is the secondary cluster machine) :

You are logged on node1 as oracle (or root) :

```
$ rlogin node2 (-> no pwd)
$ rcp /tmp/toto node2:/tmp/toto
$ rsh node2 pwd
```

6 INSTALLING HACMP/ES/CRM ON A P SERIES CLUSTER

HACMP (*High Availability Cluster Management Protocol*) is a product which provide a high availability on a cluster of machines (from 2 to 32 nodes).

HACMP/ES is the enhanced version of HACMP (*Enhanced Scalability*)

HACMP/CRM (*Concurrent Ressource Manager*) layer includes the CLVM, which enable the concurrent logical volume manager.

In all this CookBook, HACMP means HACMP/ES/CRM.

The instances of the same parallel database have a concurrent access on the same external disks. It is a *concurrent* access, and not a *shared* one.

In this Oracle9i RAC installation guide, the purpose is not to set up all the parameter stuff of HACMP. We will focus on the concurrent volumes groups, which is enough to install and run Oracle9i RAC.

HACMP 4.4 or 4.5 is used for clusters of pSeries machines with AIX 5.1 or 5.2 (and also AIX 4.3.3)

The same installation have to be done on all the machines of the cluster (node1 & node2)

All the official HACMP documentation is available at : <http://hacmp.aix.dfw.ibm.com/>

How to install HACMP using smit install :

On the directory containing all the filesets, check if the hidden file `.toc` exists. If not, when located on this directory, make `inutoc` .

If you are using a cdrom as source for this install, this file always exists on the installation media.

smit install specify the directory containing the filesets

Type F4 to list the filesets to install, rather than choosing "all latest" .

To be sure of what you are doing, you can use the field *preview* only before proceeding to the real install.

Filesets cluster.es.xxx

See *Appendix H : Filesets to be installed on the machines of the cluster*, paragraph "HACMP 4.4" or "HACMP 4.5" to check that all the necessary filesets have been installed.

This appendix provides the result of the command :

```
lslpp -L | grep cluster
```

Note that the **vsm** filesets (Visual Software Manager, a X11 based smit tool) does not have to be installed.

Filesets rsct.xx.xxx

RSCT (*RS/6000 Cluster Technology*) are the base filesets used by HACMP to manage the cluster.

RSCT 2.2 is used for clusters of pSeries machines with AIX 5.1 and RSCT 2.3 for machines with AIX 5.2.

During the installation, select all the common filesets, the HACMP support one, but not the PSSP support one.

See *Appendix H : Filesets to be installed on the machines of the cluster* paragraph HACMP 4.4 or 4.5 and paragraph RSCT 2.2 or 2.3 to check that all the necessary filesets have been installed.

This appendix provides the result of the command :

```
lslpp -L | grep rsct
```

6.1 PATCHES REQUIREMENTS

Rather than providing a list of patch which will become quickly obsolete, here are informations which will help to

- identify necessary AIX patches ('apar' or 'ptf')
- check if a patch is installed
- download fixes

How to have the list of pre-requisite patches ?

You can have a look to :

http://metalink.oracle.com/metalink/plsql/ml2_gui.startup

How to check if a patch is installed ?

Execute : `instfix -ik <APAR #>`
For example : `instfix -ik IY21047`

Where fixes can be downloaded ?

You can download new AIX maintenance levels and specified patches on

For AIX 5.1 : <http://techsupport.services.ibm.com/server/aix.fdc51>

For AIX 4.3.3 : <http://techsupport.services.ibm.com/rs6k/fixdb.html>

6.2 POST INSTALL TASKS

After all the filesets have been installed :

- check the filesets committed : You must have at least the filesets listed in appendix H, with a release number equal or higher.

```
lslpp -L | grep cluster
lslpp -L | grep rsct
```
- add in the PATH environment variable the following directories
 - /usr/es/sbin/cluster
 - /usr/es/sbin/cluster/utilities
 - /usr/es/sbin/cluster/sbin
 - /usr/es/sbin/cluster/diag
- check the existence of symbolic links from files contained in /usr/sbin/cluster to /usr/es/sbin/cluster. Else, create them with `ln -s`.

Important : After both RSCT and HACMP have been installed successfully on all the nodes, all the machines have to be **rebooted** before going on with HACMP configuration.

7 INSTALLING GPFS

How to install HACMP using smit install :

On the directory containing all the filesets, check if the hidden file `.toc` exists. If not, when located on this directory, make `inutoc` .

If you are using a cdrom as source for this install, this file always exists on the installation media.

`smit install` specify the directory containing the filesets

Type F4 to list the filesets to install, rather than choosing "all latest" .

To be sure of what you are doing, you can use the field *preview* only before proceeding to the real install.

Filesets mmfs.xx.xxx

See *Appendix H : Filesets to be installed on the machines of the cluster*, paragraph "GPFS" to check that all the necessary filesets have been installed.

This appendix provides the result of the command :

```
lslpp -L | grep mmfs
```

Post install task

Add GPFS binaries in root's path.

```
export PATH=$PATH:/usr/lpp/mmfs/bin
```

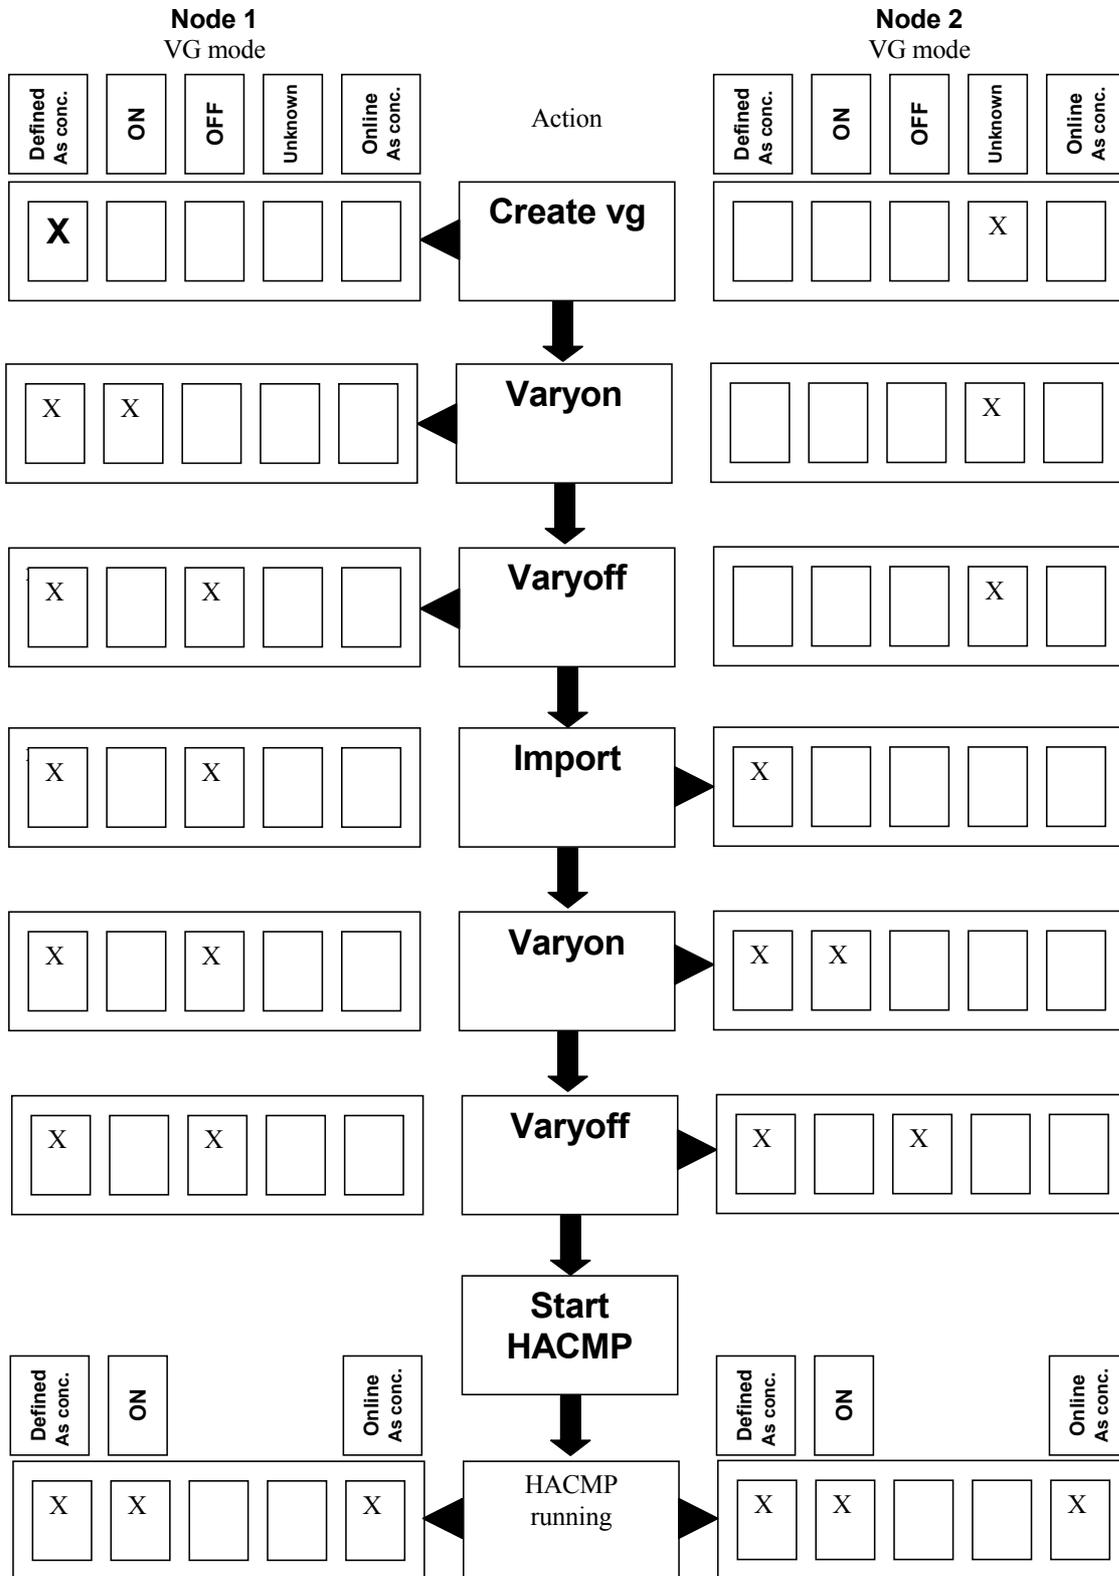
8 CREATING A CONCURRENT VOLUME GROUP

The database files and so on are stored on external concurrent disks. All the disks must be physically linked to the two nodes. Because these disks are connected on all the nodes, all the disks can be accessed by both of the nodes, in a concurrent mode.

The database files can be stored on SSA disks or Enterprise Storage Server (ESS) connected to all the nodes.

If you are implementing a Oracle 9i RAC with GPFS (no raw devices datafiles) you can skip this part and get directly to chapter 11 – Configuring HACMP for GPFS implementation.

Creating a concurrent Volume Group (VG)



- Check if the target disks are physically linked to the two machines of the cluster, and seen by both.

Run **lspv** command on both machines.

Note : the disk name can be different, depending on the others disks connected on each machine. Use the second field of the output (PVID) of **lspv** output to be sure you are dealing with the same physical disk from two hosts.

Disk name	PV identifier	Volume group name
hdisk0	0041d1db25e505a3	rootvg
hdisk1	004331ab68869a9b	backup_vg
hdisk2	none	None
hdisk3	none	None
hdisk4	0043319b1c53283e	rac92_vg

If the Pvid is not yet affected (none), a value will be given while a volume group using this hdisk is created.

- Create at the AIX level on the first machine (node1) a concurrent volume group, **rac92_cvg**

```
node1# smit vg
```

Add a Volume Group

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

	[Entry Fields]		
VOLUME GROUP name	[rac92_cvg]		
Physical partition SIZE in megabytes	32	See Note 1	+
* PHYSICAL VOLUME names	[hdisk4]		+
Activate volume group AUTOMATICALLY at system restart?	no	See Note 2	+
Volume Group MAJOR NUMBER	[64]	See Note 3	+#
Create VG Concurrent Capable?	yes		+
Auto-varyon in Concurrent Mode?	no		+

❖ **Note 1 :**

With AIX 5.1, you can leave the default value, which is 4. AIX will take automatically the best partition size. 32 MB for a 18.2 GB disk, 63 MB for a 36.4 GB disk.

❖ **Note 2 :**

Never choose YES to activate at system restart, neither auto-varyon in concurrent mode. These tasks have to be managed by HACMP. The volume group just has to be created with concurrent capability.

❖ **Note 3 :**

You must choose the major number to be sure the volume groups have the same major number on all the nodes (caution, before choosing this number, you must be sure it is free on all the nodes).
The **lvlsstmajor** command lists the first free major number on a machine.

To check all defined major number, type `ls -al /dev/*`

```
crw-rw---- 1 root system 64, 0 Aug 02 13:39 /dev/rac92_cvg
```

The major number for **rac92_cvg** volume group is 64.

- Varyon manually the volume group on node2 :

```
varyonvg rac92_cvg
```

- **On this volume group , create all the logical volumes (raw devices) you need for your database**

See Chapter 8 : [Creating the concurrent logical volumes](#)

See [Appendix C : Logical volumes creation](#)

- **Import rac92_cvg volume group on the second machine (node2)**

On the first machine, type `varyoffvg rac92_cvg`

On the second machine, import the definition of the volume group :

```
node2# smit vg
                                     Import a Volume Group

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

          [Entry Fields]
VOLUME GROUP name                    [rac92_cvg]
* PHYSICAL VOLUME name                [hdisk4]          +
Volume Group MAJOR NUMBER            [ 64]            +#+
Make this VG Concurrent Capable?     yes              +
Make default varyon of VG Concurrent? no               +
```

The physical volume name (hdisk) could not have the same number on both sides. Check the PVID of the disk, because it's the only information reliable and unique thru the cluster.

Be sure to have the same major number. This number has to be undefined on all the nodes.

❖ **Note :**

The import of a volume group resets to YES the field "Activate volume group automatically at system restart", which has to be set to NO

To do it manually : `chvg -a 'n'`

Or you can also modify this flag using `smit`

- **Varyon manually the volume group on node2 :**

`varyonvg rac92_cvg`

The new volume group is now defined on the all the machines of the cluster, with the concurrent capable feature set on.

- **Test the availability of the volume group on the two nodes**

The volume group is now defined on all the nodes of the cluster, with concurrent capability. But it is not yet varied on in concurrent mode. This is the goal of HACMP.

Before HACMP is up and running, the rac92_cvg volume group is defined on all the nodes, but can be varied on at one time on a single machine only. Don't try to varyon rac92_cvg manually in concurrent mode, this task is devoted to HACMP only.

To list the volume groups defined on the machine : `lsvg`

To list the varied on volume groups : `lsvg -o`

To switch the volume group from one node to the other one, do :

`varyoffvg rac92_cvg` on the machine where it is varied on

varyonvg rac92_cvg on the other node

lsvg rac92_cvg :

```
VOLUME GROUP:    rac92_cvg                VG IDENTIFIER: 00004c00000000ee20634e17
VG STATE:        active                    PP SIZE:       32 megabyte(s)
VG PERMISSION:   read/write                TOTAL PPs:     1084 (34688 megabytes)
MAX LVs:         256                       FREE PPs:      657 (21024 megabytes)
LVs:            12                         USED PPs:      427 (13664 megabytes)
OPEN LVs:       12                        QUORUM:        3
TOTAL PVs:      4                         VG DESCRIPTORS: 4
STALE PVs:     0                         STALE PPs:     0
ACTIVE PVs:    4                         AUTO ON:       no
Concurrent:    Capable                    Auto-Concurrent: Disabled
VG Mode:      Concurrent
Node ID:      1                          Active Nodes:
MAX PPs per PV: 1016                      MAX PVs:       32
LTG size:     128 kilobyte(s)             AUTO SYNC:     no
HOT SPARE:    no
```

You should have a similar output on all the nodes.

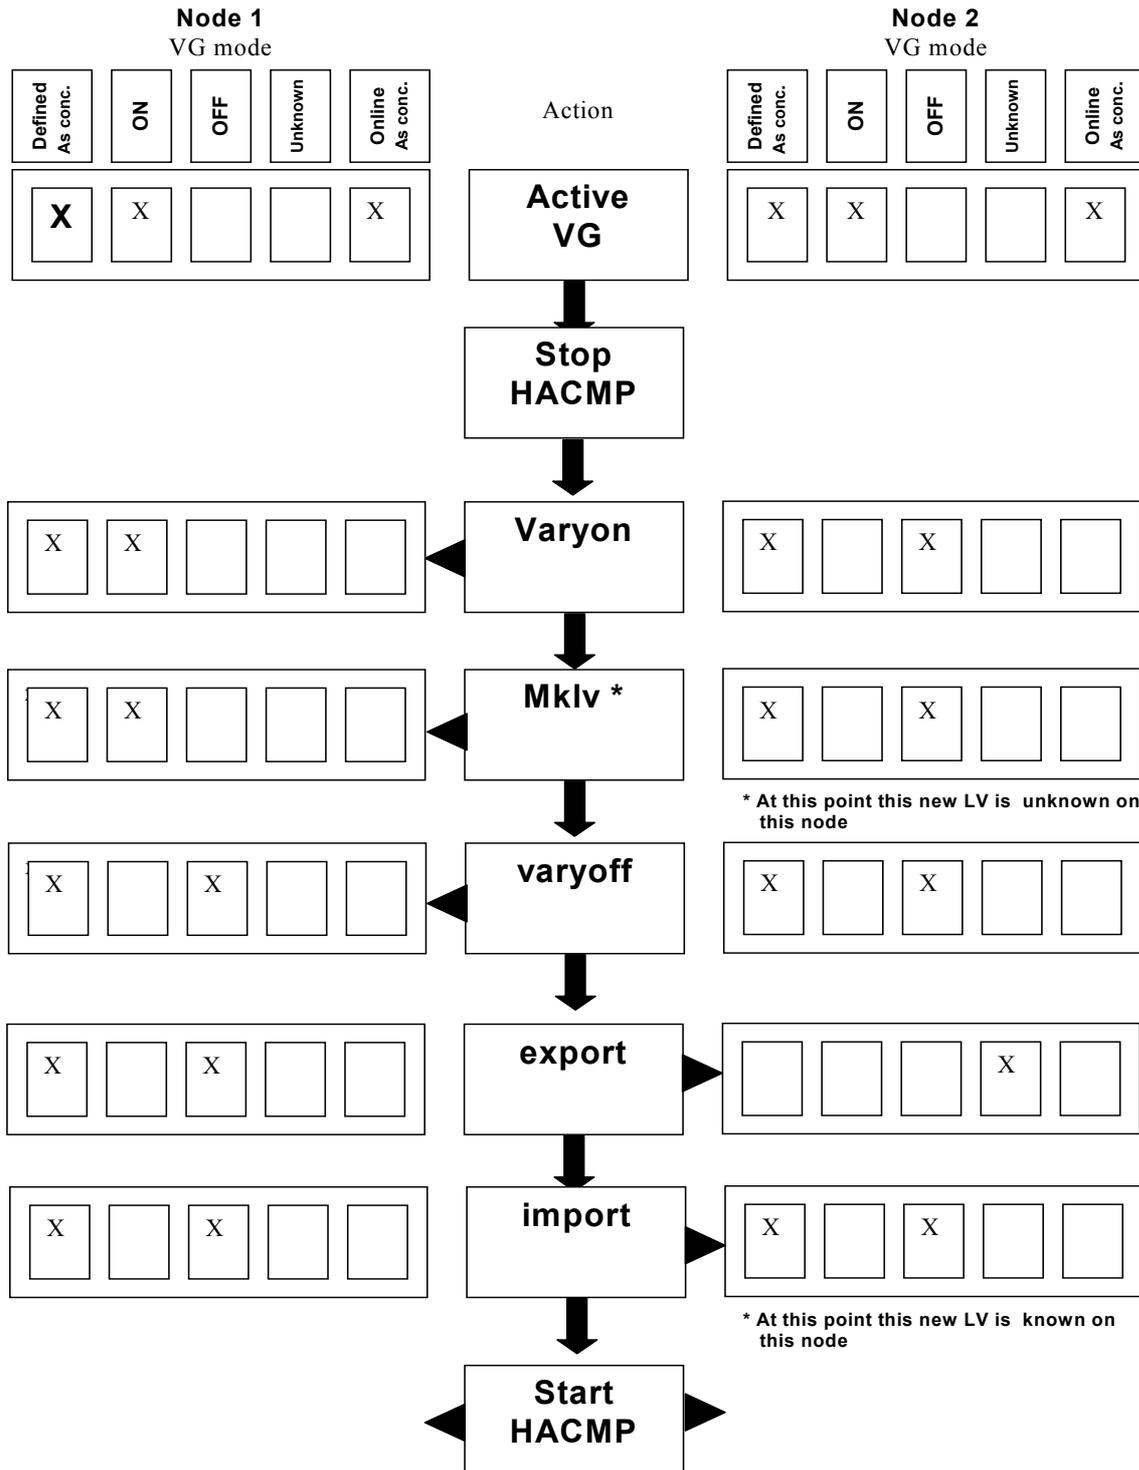
Note that at this point of the disks setup, the VG Mode is not be concurrent yet. But it must be concurrent after the starting of HACMP !

For a script, see [Appendix C : Logical volumes creation](#)

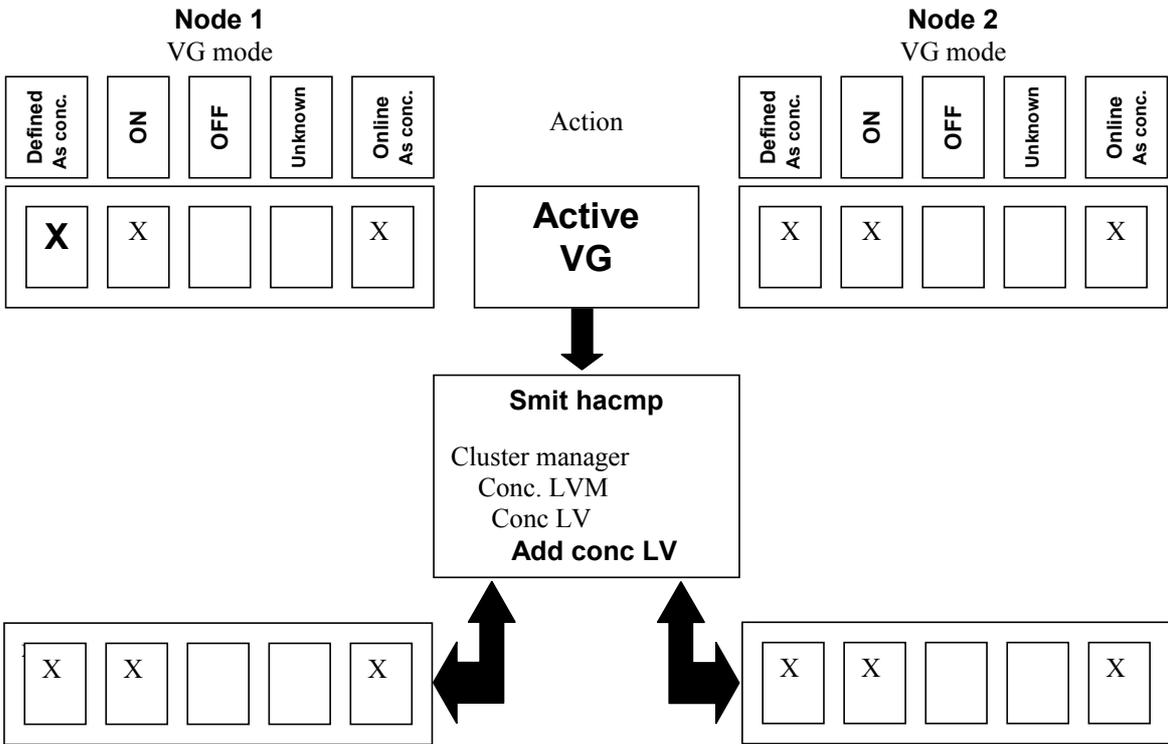
9 CREATING THE CONCURRENT LOGICAL VOLUMES

If you are implementing a Oracle 9i RAC with GPFS (no raw devices datafiles) you can skip this part and get directly to chapter 11 – Configuring HACMP for GPFS implementation.

Adding a Logical Volume in an active Volume Group (VG)
(Via Logical Volume Manager)



Adding a Logical Volume in an active Volume Group (VG)
(Via HACMP)



Design of the RAC database.

As said previously, the database can be implemented in raw-devices or in file-systems (with GPFS).

In a raw-device implementation, each logical volume created is used as a raw-device and will represent one datafile (or redolog or control file...). See the table #1 below to have a list of logical volumes required for a minimal database implementation.

Table #1. A minimal database implementation requires the following raw devices :

Size	Logical Volume name	Raw device name	Purpose
1024 MB	rac_system01	/dev/rac_system01	SYSTEM tablespace
256 MB	rac_undotbs01	/dev/rac_undotbs1	UNDO Tablespace (instance #1)
256 MB	rac_undotbs02	/dev/rac_undotbs2	UNDO Tablespace (instance #2)
128 MB	rac_redo01	/dev/rac_redo1	Redolog Thread #1, Group #1
128 MB	rac_redo02	/dev/rac_redo2	Redolog Thread #1, Group #2
128 MB	rac_redo03	/dev/rac_redo3	Redolog Thread #1, Group #3
128 MB	rac_redo04	/dev/rac_redo4	Redolog Thread #2, Group #1
128 MB	rac_redo05	/dev/rac_redo5	Redolog Thread #2, Group #2
128 MB	rac_redo06	/dev/rac_redo6	Redolog Thread #2, Group #3
32 MB	rac_control01	/dev/rac_control01	Control File # 1
32 MB	rac_control02	/dev/rac_control02	Control File # 2
32 MB	rac_control03	/dev/rac_control03	Control File # 3
32 MB	rac_spfile	/dev/rac_spfile	For spfile (if used)
128 MB	rac_srvconfig	/dev/rac_srvconfig	For srvctl tool (if used)
128 MB	rac_data	/dev/rac_data	DATA tablespace
128 MB	rac_index	/dev/rac_index	INDEX tablespace
128 MB	rac_temp	/dev/rac_temp	TEMP tablespace

The redologs and undo logical volumes have to be dedicated on each instance.

A script is provided in [Appendix C : Logical volumes creation](#) to create the volumes groups and the logical volume as specified above.

Some sample command from this appendix :

```
mkvg -f -c -y rac92_cvg -V'64' -s'32' <disk list>
mklv -y'rac_system01' rac92_cvg 16 ext_disk # 512 MB
```

The volume group, rac92_cvg, is build with a physical partition (PP) size (of 32MB for example). The size of the logical volume is expressed in number of PP, not in Kbytes.

Once a logical volume *new_lv* is created on a node, two new entries appears in the /dev directory :

/dev/new_lv which is normally used by LVM for file systems.

/dev/**r**new_lv where the "r" stands for raw device. This is the device name to use with Oracle for a raw-device implementation.

Important for raw-device implementation : check that the owner of /dev/**r**new_lv is oracle, group dba, and that oracle have the read and write privilege.

If needed, do as root (on all the nodes) :

```
chown oracle:dba /dev/*rac_*
chmod go+rw /dev/*rac_*
```

(for the GPFS implementation, we will see later for the ownership and permissions...)

Once new logical volumes are created or modified on a node, the other node has to update its ODM (AIX internal repository), by retrieving the information on the concurrent disk.

```
redefinevg -d <disk name> rac92_cvg
```

This command can also be done with the following sequence :

```
varyoffvg rac92_cvg
exportvg rac92_cvg
importvg -y rac92_cvg -V<major number> <disk name>
varyonvg rac92_cvg
```

The same command has to be executed on the other nodes, as soon as an update has been made on one node. It is important for HACMP that all the ODM databases of all the nodes have the same level of update about the logical volumes of the concurrent disk. This point is checked by HACMP during the synchronization process.

You should have a similar output on all the nodes.

lsvg -l rac92_cvg (raw-devices implementation, according table #1):

```
rac92_cvg:
LV NAME          TYPE      LPs   PPs   PVs   LV STATE   MOUNT POINT
rac_system01     jfs       32    32    1     open/syncd N/A
rac_undotbs01    jfs       8      8    1     open/syncd N/A
rac_undotbs02    jfs       8      8    1     open/syncd N/A
rac_redo01       jfs       4      4    1     open/syncd N/A
rac_redo02       jfs       4      4    1     open/syncd N/A
rac_redo03       jfs       4      4    1     open/syncd N/A
rac_redo04       jfs       4      4    1     open/syncd N/A
rac_redo05       jfs       4      4    1     open/syncd N/A
rac_redo06       jfs       4      4    1     open/syncd N/A
rac_control01    jfs       1      1    1     open/syncd N/A
rac_control02    jfs       1      1    1     open/syncd N/A
rac_control03    jfs       1      1    1     open/syncd N/A
rac_srvconfig    jfs       4      4    1     open/syncd N/A
```

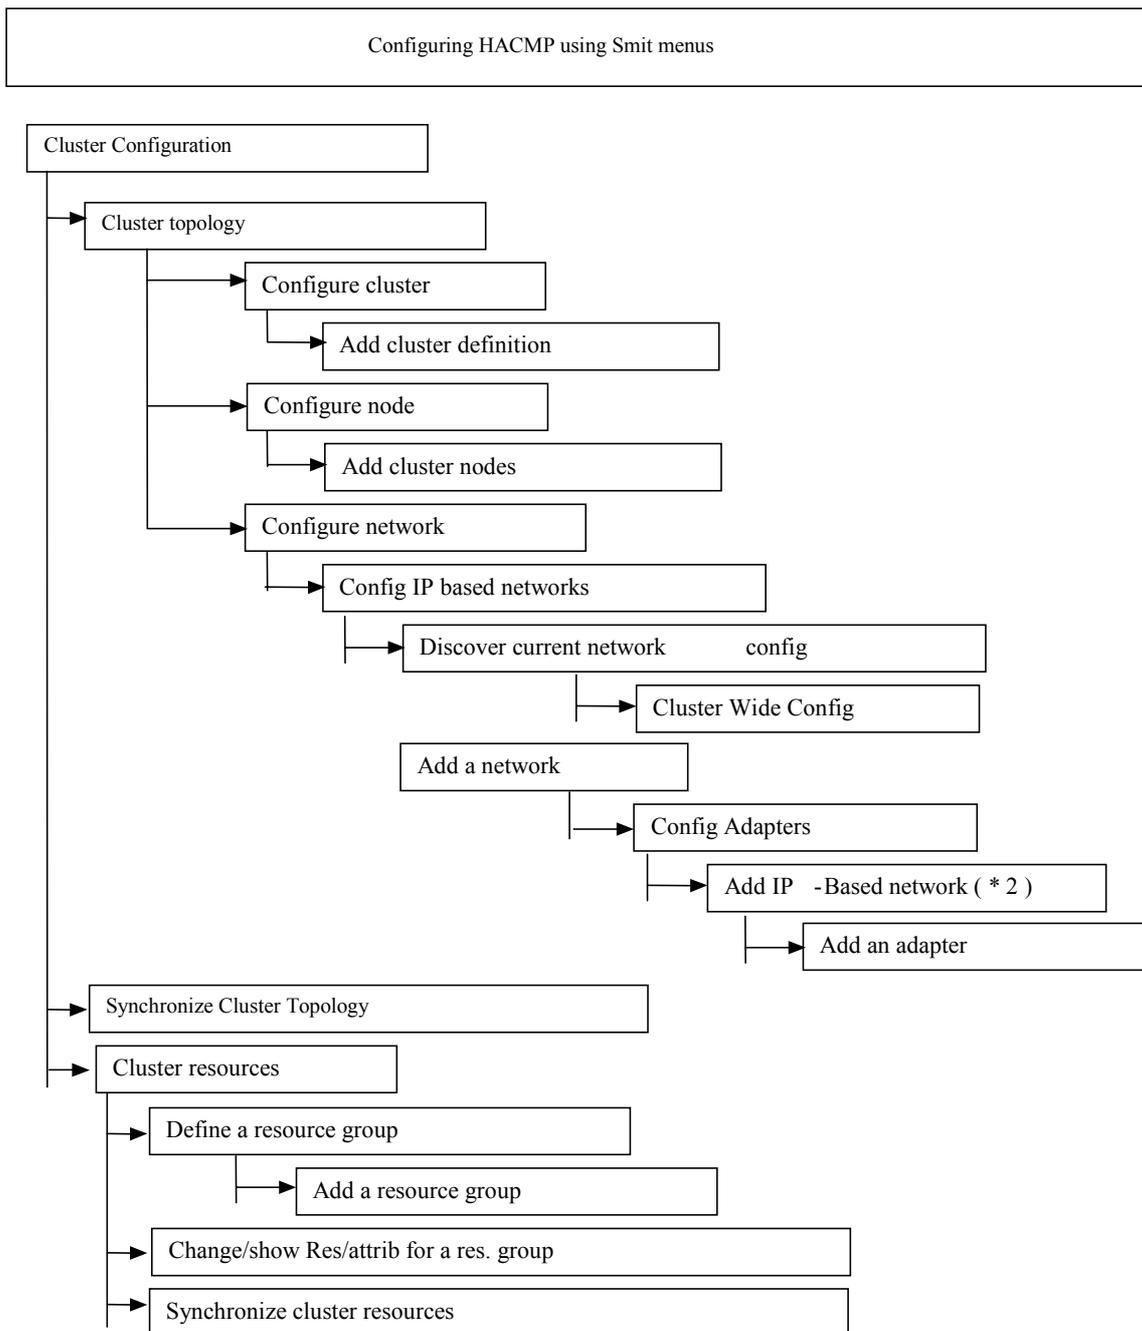
10 CONFIGURING HACMP - RAW-DEVICES IMPLEMENTATION

Main steps for implementing an HACMP configuration :

Consult the HACMP documentation, available on <http://hacmp.aix.dfw.ibm.com/>

All this steps have to be done using **one node only**. During the synchronization process, HACMP will copy the whole configuration to every node in the cluster. This functionality is called SPOC (Single Point Of Control)

Every time you have modified the topology or resources of the cluster, don't forget to synchronize.

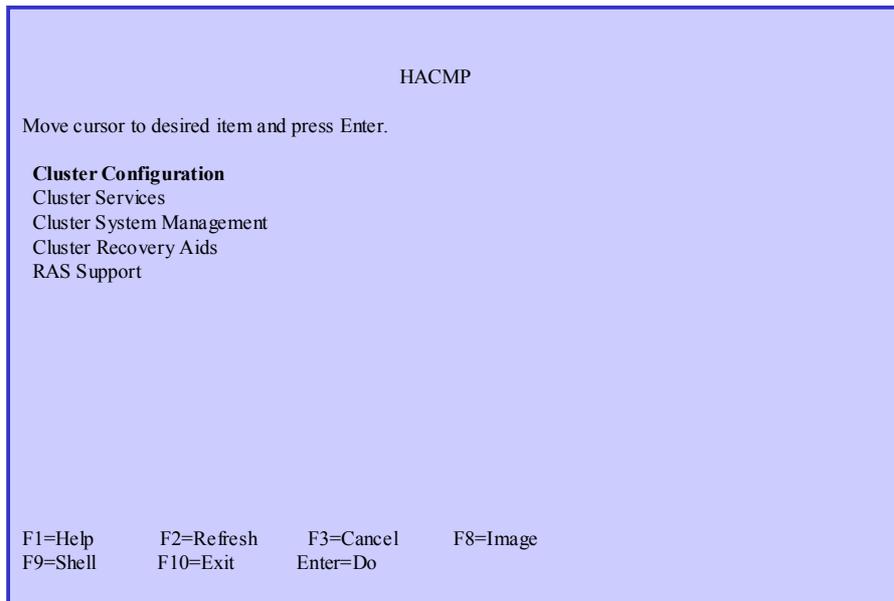


On one of the nodes of the cluster, do the following steps :

Don't try to make cluster configuration at the same time on two nodes. Make a synchronization first.

- Configure the cluster : **smit hacmp**, and follow the screens

In this step you give an identifier and a name to your cluster.



First of all, define the cluster hardware topology, which includes the nodes participating to the cluster, the network, and network interfaces.

Cluster Configuration

Move cursor to desired item and press Enter.

Cluster Topology
Cluster Security
Cluster Resources
Cluster Snapshots
Cluster Verification
Cluster Custom Modification
Restore System Default Configuration from Active Configuration
Advanced Performance Tuning Parameters

F1=Help F2=Refresh F3=Cancel F8=Image
F9=Shell F10=Exit Enter=Do

Cluster Topology

Move cursor to desired item and press Enter.

Configure Cluster
Configure Nodes
Configure Networks
Configure Adapters
Configure Sites
Configure Global Networks
Configure Network Modules
Configure Topology Services and Group Services
Show Cluster Topology
Synchronize Cluster Topology

F1=Help F2=Refresh F3=Cancel F8=Image
F9=Shell F10=Exit Enter=Do

- Configure the network

In the case of the implementation of a concurrent volume group, only a **private** network is needed. It's not necessary to define boot, service and standby addresses like in cascading or rotating HACMP configurations. In RAC environments, you will need a high performance private network between the machines participating to the RAC cluster. The network type of this interconnect can be HPS (High Performance Switch) or ethernet (use a point to point gigabit ethernet, for example).

```

Cluster Topology

Move cursor to desired item and press Enter.

Configure Cluster
Configure Nodes
Configure Networks
Configure Adapters
Configure Sites
Configure Global Networks
Configure Network Modules
Configure Topology Services and Group Services
Show Cluster Topology
Synchronize Cluster Topology

F1=Help      F2=Refresh   F3=Cancel    F8=Image
F9=Shell     F10=Exit    Enter=Do
  
```

- Choose : Configure IP-based networks
 Choose : Discover Current Network Configuration
 Choose : Cluster-wide Configuration
 Choose : Add a Network

```

Add an IP-based Network

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[Entry Fields]
* Network Name            [rac92_network]
* Network Attribute        private            +
Network Type              [ ether]            +
Subnet(s)                  [255.255.255.0]    +

Esc+1=Help    Esc+2=Refresh    Esc+3=Cancel    Esc+8=Image
Esc+9=Shell    Esc+0=Exit        Enter=Do
  
```

- Configure the adapters

There must be one adapter by machine. This adapter is configured for the interconnect network. It is not the hostname.

```

Cluster Topology

Move cursor to desired item and press Enter.

Configure Cluster
Configure Nodes
Configure Networks
Configure Adapters
Configure sites
Configure Global Networks
Configure Network Modules
Configure Topology Services and Group Services
Show Cluster Topology
Synchronize Cluster Topology

F1=Help      F2=Refresh   F3=Cancel    F8=Image
F9=Shell     F10=Exit     Enter=Do
  
```

Choose : Adapters on IP-based network

Choose : Add an adapter

```

Add an IP-based Adapter

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[Entry Fields]
* Adapter IP Label           [interconnect_node1]   +
Network Type                ether
Network Name                rac92_network
* Adapter Function           [service]               +
Adapter IP Address           [10.10.11.141]
Adapter Hardware Address     []
Node Name                   [node1]                 +
Netmask                     [255.255.255.0]        +

F1=Help      F2=Refresh   F3=Cancel    F8=Image
F9=Shell     F10=Exit     Enter=Do
  
```

Configure also another interconnect adapter on the other node.

The second configuration step is to define the concurrent volume group, which has to be varied on by HACMP. This is called the resource.

Cluster Configuration

Move cursor to desired item and press Enter.

- Cluster Topology
- Cluster Security
- Cluster Resources**
- Cluster Snapshots
- Cluster Verification
- Cluster Custom Modification
- Restore System Default Configuration from Active Configuration
- Advanced Performance Tuning Parameters

F1=Help F2=Refresh F3=Cancel F8=Image
F9=Shell F10=Exit Enter=Do

- Define a resource group

Cluster Resources

Move cursor to desired item and press Enter.

Define Resource Groups
 Define Application Servers
 Configure Application Monitoring
 Define Tape Resources
 Define Highly Available Communication Links
 Discover Current Volume Group Configuration
 Configure Dynamic Node Priority Policies
 Change/Show Resources/Attributes for a Resource Group
 Cluster Events
 Change/Show Run-Time Parameters
 Change/Show Cluster Lock Manager Resource Allocation
 Show Cluster Resources
 Synchronize Cluster Resources

Esc+1=Help Esc+2=Refresh Esc+3=Cancel Esc+8=Image
 Esc+9=Shell Esc+0=Exit Enter=Do

Change/Show a Resource Group

Type or select values in entry fields.
 Press Enter AFTER making all desired changes.

[Entry Fields]

Resource Group Name	rac92_resource	
New Resource Group Name	[]	
Node Relationship	concurrent	+
Participating Node Names	[node1 node2]	+

F1=Help F2=Refresh F3=Cancel F4=List
 F5=Reset F6=Command F7=Edit F8=Image
 F9=Shell F10=Exit Enter=Do

The node relationship must be “concurrent”. All participating nodes must be entered.

- Synchronize the cluster resources

```
Cluster Resources
Move cursor to desired item and press Enter.

Define Resource Groups
Define Application Servers
Configure Application Monitoring
Define Tape Resources
Define Highly Available Communication Links
Discover Current Volume Group Configuration
Configure Dynamic Node Priority Policies
Change/Show Resources/Attributes for a Resource Group
Cluster Events
Change/Show Run-Time Parameters
Change/Show Cluster Lock Manager Resource Allocation
Show Cluster Resources
Synchronize Cluster Resources

Esc+1=Help   Esc+2=Refresh   Esc+3=Cancel   Esc+8=Image
Esc+9=Shell   Esc+0=Exit     Enter=Do
```

It is better, but not mandatory, that HACMP is stopped when starting synchronizing.

Important : your cluster must be synchronized after each new modification (topology or resource)

To check your HACMP cluster configuration, see [Appendix F : HACMP Cluster verification output](#)

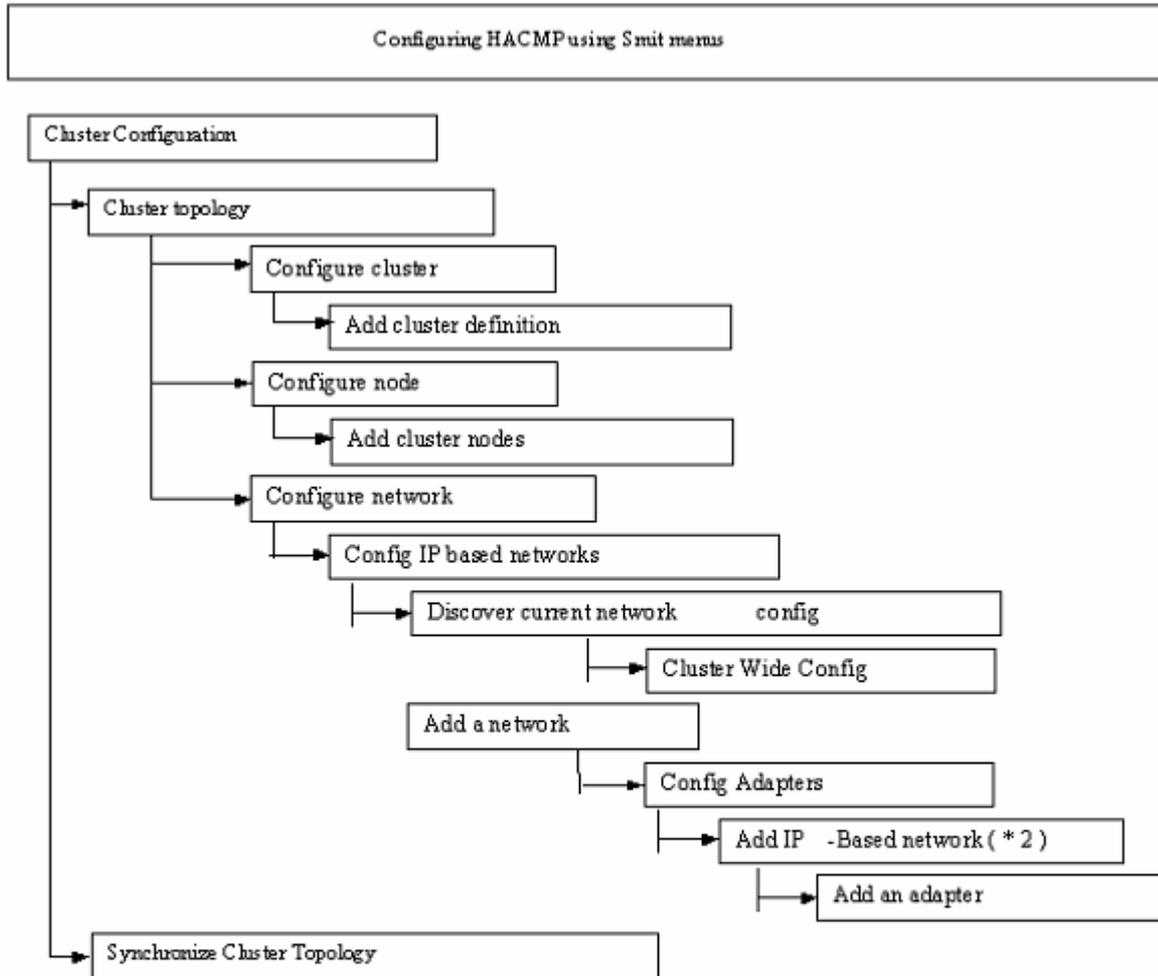
11 CONFIGURING HACMP - GPFS IMPLEMENTATION

Main steps for implementing an HACMP configuration :

Consult the HACMP documentation, available on <http://hacmp.aix.dfw.ibm.com/>

All this steps have to be done using **one node only**. During the synchronization process, HACMP will copy the whole configuration to every node in the cluster. This functionality is called SPOC (Single Point Of Control)

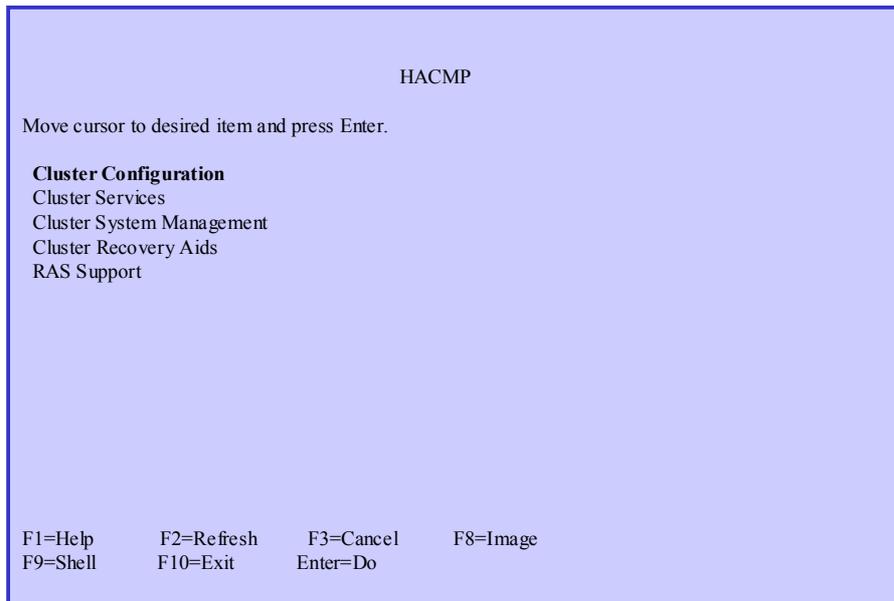
Every time you have modified the topology or resources of the cluster, don't forget to synchronize.



On one of the nodes of the cluster, do the following steps :
Don't try to make cluster configuration at the same time on two nodes. Make a synchronization first.

- Configure the cluster : **smit hacmp**, and follow the screens

In this step you give an identifier and a name to your cluster.



First of all, define the cluster hardware topology, which includes the nodes participating to the cluster, the network, and network interfaces.

Cluster Configuration

Move cursor to desired item and press Enter.

Cluster Topology
 Cluster Security
 Cluster Resources
 Cluster Snapshots
 Cluster Verification
 Cluster Custom Modification
 Restore System Default Configuration from Active Configuration
 Advanced Performance Tuning Parameters

F1=Help F2=Refresh F3=Cancel F8=Image
 F9=Shell F10=Exit Enter=Do

Cluster Topology

Move cursor to desired item and press Enter.

Configure Cluster
 Configure Nodes
 Configure Networks
 Configure Adapters
 Configure Sites
 Configure Global Networks
 Configure Network Modules
 Configure Topology Services and Group Services
 Show Cluster Topology
 Synchronize Cluster Topology

F1=Help F2=Refresh F3=Cancel F8=Image
 F9=Shell F10=Exit Enter=Do

- Configure cluster nodes

Here, you enter the name of the nodes participating to the cluster. Those names are the hostname, not the name planned for interconnect network.

Cluster Topology

Move cursor to desired item and press Enter.

- Configure Cluster
- Configure Nodes**
- Configure Networks
- Configure Adapters
- Configure Sites
- Configure Global Networks
- Configure Network Modules
- Configure Topology Services and Group Services
- Show Cluster Topology
- Synchronize Cluster Topology

F1=Help F2=Refresh F3=Cancel F8=Image
F9=Shell F10=Exit Enter=Do

Add Cluster Nodes

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[Entry Fields]

* Node Names [node1 node2]

F1=Help F2=Refresh F3=Cancel F4=List
F5=Reset F6=Command F7=Edit F8=Image
F9=Shell F10=Exit Enter=Do

- Configure the network

In the case of the implementation of a concurrent volume group, only a **private** network is needed. It's not necessary to define boot, service and standby addresses like in cascading or rotating HACMP configurations. In RAC environments, you will need a high performance private network between the machines participating to the RAC cluster. The network type of this interconnect can be HPS (High Performance Switch) or ethernet (use a point to point gigabit ethernet, for example).

```

Cluster Topology

Move cursor to desired item and press Enter.

Configure Cluster
Configure Nodes
Configure Networks
Configure Adapters
Configure Sites
Configure Global Networks
Configure Network Modules
Configure Topology Services and Group Services
Show Cluster Topology
Synchronize Cluster Topology

F1=Help      F2=Refresh  F3=Cancel   F8=Image
F9=Shell     F10=Exit   Enter=Do
  
```

- Choose : Configure IP-based networks
- Choose : Discover Current Network Configuration
- Choose : Cluster-wide Configuration
- Choose : Add a Network

```

Add an IP-based Network

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[Entry Fields]
* Network Name           [rac92_network]
* Network Attribute       private           +
Network Type             [ ether]           +
Subnet(s)                [255.255.255.0]   +

Esc+1=Help   Esc+2=Refresh   Esc+3=Cancel   Esc+8=Image
Esc+9=Shell   Esc+0=Exit     Enter=Do
  
```

A second private network must be added for GPFS (see the schema presented in the chapter 2.2 'GPFS implementation'). See below.

Add an IP-based Network

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[Entry Fields]

* Network Name	[gpfs_network]	
* Network Attribute	private	+
Network Type	[ether]	+
Subnet(s)	[255.255.255.0]	+

Esc+1=Help Esc+2=Refresh Esc+3=Cancel Esc+4=Image
 Esc+9=Shell Esc+0=Exit Enter=Do

- Configure the adapters

There must be one adapter by machine. This adapter is configured for the interconnect network. It is not the hostname.

```

Cluster Topology

Move cursor to desired item and press Enter.

Configure Cluster
Configure Nodes
Configure Networks
Configure Adapters
Configure sites
Configure Global Networks
Configure Network Modules
Configure Topology Services and Group Services
Show Cluster Topology
Synchronize Cluster Topology

F1=Help      F2=Refresh   F3=Cancel    F8=Image
F9=Shell     F10=Exit    Enter=Do
  
```

Choose : Adapters on IP-based network
Choose : Add an adapter

```

Add an IP-based Adapter

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[Entry Fields]
* Adapter IP Label           [interconnect_node1]   +
Network Type                ether
Network Name                rac92_network
* Adapter Function           [service]               +
Adapter IP Address          [10.10.11.141]
Adapter Hardware Address    []
Node Name                   [node1]                 +
Netmask                     [255.255.255.0]       +

F1=Help      F2=Refresh   F3=Cancel    F8=Image
F9=Shell     F10=Exit    Enter=Do
  
```

Configure also another interconnect adapter on the other node.

Add the IP-based adapter for gpfs_network. See below.

Add an IP-based Adapter

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[Entry Fields]		
* Adapter IP Label	[gpfs_node1]	+
Network Type	ether	
Network Name	gpfs_network	
* Adapter Function	service	+
Adapter IP Address	[20.20.21.141]	
Adapter Hardware Address	[]	
Node Name	node1	+
Netmask	[255.255.255.0]	+

F1=Help F2=Refresh F3=Cancel F8=Image
F9=Shell F10=Exit Enter=Do

and repeat the operation for the other node.

- Synchronize the cluster topology

```

Cluster Topology

Move cursor to desired item and press Enter.

Configure Cluster
Configure Nodes
Configure Networks
Configure Adapters
Configure sites
Configure Global Networks
Configure Network Modules
Configure Topology Services and Group Services
Show Cluster Topology
Synchronize Cluster Topology

F1=Help      F2=Refresh   F3=Cancel    F8=Image
F9=Shell     F10=Exit    Enter=Do
  
```

```

Synchronize Cluster Topology

Type or select values in entry fields.
Press Enter AFTER making all desired changes.
[TOP]                [Entry Fields]
Ignore Cluster Verification Errors?      [No]      +
* Emulate or Actual?                    [Actual]   +
* Skip Cluster Verification              [No]      +
Note:
Only the local node's default configuration files
keep the changes you make for topology DARE
emulation. Once you run your emulation, to
restore the original configuration rather than
running an actual DARE, run the SMIT command,
"Restore System Default Configuration from Active
Configuration."
[MORE...9]
F1=Help      F2=Refresh   F3=Cancel    F4=List
F5=Reset     F6=Command   F7=Edit      F8=Image
F9=Shell     F10=Exit    Enter=Do
  
```

Now, the topology is defined and synchronized.

Important : your cluster must be synchronized after each new modification (topology or resource)

To check your HACMP cluster configuration, see [Appendix F : HACMP Cluster verification output](#)

12 STARTING HACMP

Is HACMP correctly configured ? Checklist.

- **For the raw devices implementation only**

You have to check if the concurrent Volume Group is active or not on both nodes.

The concurrent volume group is managed by HACMP. When HACMP is stopped, the volume group is not available (i.e. varyoff) When HACMP is started, the volume group is varyon in concurrent mode access (i.e. varyon)

To check the available volume groups

```
lsvg -o
```

In the command output, you should see rac92_cvg only when HACMP is up and running.

If the volume group is available on one of the nodes, make varyoffvg rac92_cvg on this node.

- **For the GPFS only**

No volume groups are defined yet. The concurrent volume group is managed by GPFS and will be defined in the Chapter 12 – Configuring and starting GPFS.

- **For both implementation : Start HACMP**

Do the following on each node of the cluster.

HACMP

Move cursor to desired item and press Enter.

Cluster Configuration
Cluster Services
 Cluster System Management
 Cluster Recovery Aids
 RAS Support

F1=Help F2=Refresh F3=Cancel F8=Image
 F9=Shell F10=Exit Enter=Do

Cluster Services

Move cursor to desired item and press Enter.

Start Cluster Services
 Stop Cluster Services
 Show Cluster Services

F1=Help F2=Refresh F3=Cancel F8=Image
 F9=Shell F10=Exit Enter=Do

```

                                Start Cluster Services

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                [Entry Fields]
* Start now, on system restart or both          now          +
BROADCAST message at startup?                  false          +
Startup Cluster Lock Services?                  false          +
Startup Cluster Information Daemon?             true           +

F1=Help      F2=Refresh   F3=Cancel   F4=Page
F9=Shell     F10=Exit      Enter=Do
  
```

The lock manager does not have to be started. The locks are managed by Oracle.

Is HACMP up and running ?

When you have OK at the previous smit screen, this does not mean that HACMP is running yet. It means that the start process is engaged. Because HACMP is very long to be up (synchronization between the nodes), wait for 10 minutes. If after this time HACMP is not running, there is a problem, and you have to have a look into the logs.

HACMP logs : `/tmp/hacmp.out` on each node

A useful command :

```
tail -f /tmp/hacmp.out
```

clstat , **xcclstat** : programs displaying the state of the HACMP cluster and every node participating.

You can also check the HAGS sockets and the cluster sub-systems.

- **HAGS socket**

The HAGS socket needs to be writable by "oracle" and the "cldomain" executable needs to be executable by "oracle". By configuring the group and permissions for the "grpsvcsdsocket.<domain>" file the instance will be able to communicate with HAGS and the Oracle instance will mount.

On all the nodes of the cluster, perform the following tasks :

- check hagsuser group exists, else create it
place "oracle" into the "hagsuser" group
- change the permissions on the "cldomain" executable :
`chmod a+x /usr/sbin/cluster/utilities/cldomain`
- change the group to "hagsuser" for the "svcsdsocket.rac92_cluster" socket :
`chgrp hagsuser /var/ha/soc/grpsvcsdsocket.rac92_cluster`
- change the group permissions for the "grpsvcsdsocket.rac92_cluster" socket :
`chmod g+w /var/ha/soc/grpsvcsdsocket.rac92_cluster`

You should have something like the following :

```
ls -l /var/ha/soc
total 4
srw-rw-rw- 1 root      haermm      0 Jun 19 16:58 em.clsrv.rac92_cluster
srw-rw---- 1 root      haermm      0 Jun 19 16:58 em.rmsrv.rac92_cluster
drwxrwxrwx 2 root      system     512 Jun 19 17:51 grpsvcs.clients.rac92_cluster
srw-rw-rw- 1 root      hagsuser    0 Jun 19 16:57 grpsvcsdsocket.rac92_cluster
drwxrwx--- 2 root      haermm     512 Jun 19 16:57 haem
drwxr-x--- 2 root      system     512 Apr 02 23:30 hats
drwxr-xr-x 2 root      system     512 Jun 19 16:57 topsvcs
```

For more information, see the note # 2064876.102 in Metalink
(also presented in [Appendix E : Oracle technical notes](#))

- **Daemons, etc...**

Execute the following command :

```
lssrc -a | egrep 'svcs|ES'
```

It should give the following output :

topsvcs	topsvcs	65698	active
grpsvcs	grpsvcs	63872	active
grpqlsm	grpsvcs	68862	active
emsvcs	emsvcs	69132	active
emaixos	emsvcs	62336	active
clstrmgrES	cluster	70086	active
clsmuxpdES	cluster	65306	active

For raw devices implementation only :

You have also another way to see if HACMP is up : list the online volume groups. If your concurrent volume group rac92_cvgr is varyon on all the nodes, all is OK ! This is the only use of HACMP for RAC databases : providing concurrent access to disks.

To list the online volume groups : **lsvg -o** on each node.

13 CONFIGURING AND STARTING GPFS

Follow these steps to configure GPFS :

- Include the path of GPFS binaries are included in root's path.
export PATH=\$PATH:/usr/lpp/mmfs/bin
- Create the GPFS nodefile which will contain IP hostnames of GPFS interconnect network.

Example :

```
node1:/> cat /var/mmfs/etc/node.list
```

```
gpfs_node1
gpfs_node2
```

- Create the gpfs cluster :

```
node1:/> mmcrcluster -t hacmp -p gpfs_node1 -s gpfs_node2 -n
/var/mmfs/etc/node.list
```

- Create the nodeset

```
node1:/> mmconfig -n /var/mmfs/etc/node.list -A -C MyCluster -D /tmp
```

You can check the GPFS creation with the mmlscluster and the cluster configuration with mmlsconfig. In our example you see :

```
node1:/> mmlscluster
```

```
GPFS cluster information
=====
GPFS system data repository servers:
-----
Primary server: gpfs_node1
Secondary server: gpfs_node2
Nodes for nodeset MyCluster:
-----
1 gpfs_node1 20.20.21.141 gpfs_node1
2 gpfs_node2 20.20.21.142 gpfs_node2
```

```
node1:/> mmlsconfig
```

```
Configuration data for nodeset MyCluster:
-----
pagepool 80M
dataStructureDump /tmp/mmfs
multinode yes
autoload yes
useSingleNodeQuorum no
wait4RVSD no
comm_protocol TCP
clusterType hacmp
group Gpfs.set1
recgroup GpfsRec.MyCluster
File systems in nodeset MyCluster:
-----
(none)
```

- Start GPFS

```
node1: /> mmstartup -C MyCluster
```

```
Wed Oct 22 17:38:26 NPT 2003: mmstartup: Starting GPFS ...
gpfs_node1: 0513-059 The mmfs Subsystem has been started. Subsystem
PID is 65653.
gpfs_node2: 0513-059 The mmfs Subsystem has been started. Subsystem
PID is 48626.
```

You can check the GPFS sub-system is started :

```
node1: /> lssrc -s mmfsd
```

Subsystem	Group	PID	Status
mmfsd	aixmm	65653	active

- Create a “disk descriptor file” containing the disks devices that will be shared in the GPFS cluster

Example :

```
node1: /> cat /var/mmfs/etc/disk.list
```

```
hdsik12 :::1
hdisk15 :::1
```

- Create logical volumes with mmcrlyv:
 - Creates a vg and lv for each disk (internal GPFS housekeeping)
 - All VG's will be automatically imported to all nodes.

```
node1: /> mmcrlyv -f 4 -F /var/mmfs/etc/disk.list
```

```
-f <max number of nodes to access concurrently>
-F <list of disks to use with descriptors>
```

You can check what was created by mmcrlyv with lsvg -o or -l AIX command :

```
node1: /> lsvg -o
```

```
gpfs01vg
gpfs02vg
.....
```

- Create the file-system with mmcrfs

```
node1: /> mmcrfs /oradata_gpfs /dev/oradata_gpfslv -F disk_list -A yes
-B 512k -n 4 -N 80000
```

```
-F <list of disks and descriptors> modified from mmcrlyv
-A Auto mount the file system at GPFS startup
-B Block size and Stripe size
-n estimated number of nodes
-N Number of I-Nodes
```

```
node1: /> cat /etc/filesystems | grep -p /oradata_gpfs
/oradata_gpfs:
dev = /dev/oradata_gpfslv
vfs = mmfs
nodename = -
mount = false
type = mmfs
```

```
account = false
```

- GPFS block size should be around (256K, 512K, 1024K)
 - GPFS block size is the file system stripe size.
 - Not that important for regular database I/O since Direct I/O is used.
 - Very important for operations that increase data file size.
- GPFS pagepool : Since Oracle on GPFS uses Direct I/O, which bypasses the GPFS “pagepool” for all database file operations including the RMAN utility, there is no need to increase the GPFS cache. However, if basic commands such as ‘cp’ or ‘tar’ are used for file maintenance then increasing the “pagepool” may be beneficial.

- Mount the newly created file-system :

```
node1:/> mount /oradata_gpfs
```

- Set the permissions and the ownership of the file-systems

```
node1:/> chown oracle:dba /oradata_gpfs
node1:/> chmod go+rw /oradata_gpfs
```

Tuning suggestions

- **aioservers:**

AIO should be enabled. The general rule for heavy I/O environments is to initially set the “*maxservers*” value to at least 10*(number of disks accessed asynchronously). For an 8 disk GPFS file system, an example of the command to modify would be:

```
node1:/> chdev -l aio0 -a maxservers='80'
```

The “*minservers*” value can be set to half of the “*maxservers*” value.

```
node1:/> chdev -l aio0 -a minservers='40'
```

See the “Tuning Asynchronous Disk I/O” section in the “**AIX 5L Version 5.1 Performance Management Guide**” for more discussion and details.

- **pagepool:**

The pagepool is used to cache user data and indirect blocks. It is the GPFS pagepool mechanism that allows GPFS to implement read as well as write requests asynchronously. Increasing the size of pagepool increases the amount of GPFS buffer cache pinned memory available. The default value for pagepool is 20 MB and the maximum allowable value is 512 MB. Applications that may benefit from a larger pagepool (compared to the default) potentially include those that either reuse data, those that have a random I/O pattern, and/or those that have a higher per client performance requirement. The size of the pagepool will depend on the “working set” of I/O data that needs to be cached. For instance to change pagepool to 100 MB:

```
node1:/> mmchconfig pagepool=100M
```

```
node1:/> mmlsconfig
```

```
Configuration data for nodeset MyCluster:
```

```
-----
pagepool 100M
dataStructureDump /tmp/mmfs
multinode yes
autoload yes
useSingleNodeQuorum no
```

```
wait4RVSD no
comm_protocol TCP
clusterType hacmp
group Gpfs.set1
recgroup GpfsRec.MyCluster
File systems in nodeset MyCluster:
-----
/oradata_gpfs
```

- **ipqmaxlen** : the ipqmaxlen network option controls the number of incoming packets that can exist on the IP interrupt queue. Since both GPFS and IBM Virtual Shared Disk use IP, the default value of 128 is often insufficient. This is especially important if your virtual shared disks are configured over IP. The recommended setting is 512.

```
node1:/> no -a ipqmaxlen=512
node1:/> rsh node2 no -a ipqmaxlen=512
```

14 ORACLE ENVIRONMENT SETUP

- Oracle environment : **\$HOME/.profile** file in Oracle's home directory

```
export ORACLE_HOME=/oracle/app/oracle/product/9.2.0
export DB_NAME=RAC
export ORACLE_SID=RAC1
export ORACLE_SERVICE=RAC
export PATH=$ORACLE_HOME/bin:$PATH
```

- The oracle code can be located on an internal SCSI disk and propagated on the other machines of the cluster. The Oracle Universal Installer manage the cluster-wide installation, that is done only once. Regular file systems are used for Oracle code.

On both nodes, create the file system for Oracle code. This file system of 4 GB, is generally located on an internal SCSI disk. The external disks will store the database files.

To list the internal disks : `lscfg | grep -i scsi | grep hdisk`
Suppose we have hdisk1, and internal SCSI free disk of 18.2 GB

Create a volume group called **oraclevg** :

```
mkvg -f -y'oraclevg' -s'32' hdisk1
```

Create a 4GB file system **/oracle** in the previous volume group (large file enabled):

```
crfs -v jfs -a bf=true -g'oraclevg' -a size='8388608' -m'/oracle' -
A'yes' -p'rw' -t'no' -a nbpi='8192' -a ag='64'
```

```
mount /oracle
chown oracle:dba /oracle
```

The Oracle code can also be located on external concurrent disks on a GPFS file-system. See the chapter 13 - Configuring and starting GPFS.

- Create a directory writable for oracle user
`/var/opt/oracle`
- Cluster Manager Software
If the cluster manager software is installed, the Oracle RAC option will be automatically preselected in Oracle Universal Installer.

Snapshot of Metalink, regarding the maximum numbers of nodes (sept 2002)

Technology Matrix for IBM AIX pSeries (RS/6000): for IBM AIX Clusters

- The following items are additions or limitations to the Compatible Technology List for IBM AIX Clusters for the AIX platform.
- Oracle will support the Oracle software on clusters comprised of certified software combinations and RAC compatible technologies.
- **NOTE:** Please consult your hardware and clusterware vendor as not all vendors may choose to support their hardware or clusterware in every possible cluster combination.

<u>Technology Category</u>	<u>Technology</u>	<u>Exclusions/Limitations/Notes</u>
Storage		<ul style="list-style-type: none"> ○ If SSA disks are used, the maximum number of nodes is 8. ○ If VSDs are used, the maximum number of nodes is 128
Network Interconnect		<ul style="list-style-type: none"> ○ If an SP switch fabric is used, PSSP is required
Cluster Software		<ul style="list-style-type: none"> ○ If the disks are shared through CLVM, the maximum number of nodes is 16.

15 INSTALL THE ORACLE SOFTWARE

Oracle RAC installation just have to be done only on one node. Once the first node is installed, Oracle OUI automatically starts the copy of the mandatory files on the second node, using **rcp** command. This step could last long, depending on the network speed (one hour...), without any message. So, don't think the OUI is stalled, and look at the network traffic before canceling the installation !

You can also create a staging area. The name of the subdirectories is in the format "Disk1" to "Disk4"

Login as oracle and follow the procedure hereunder...

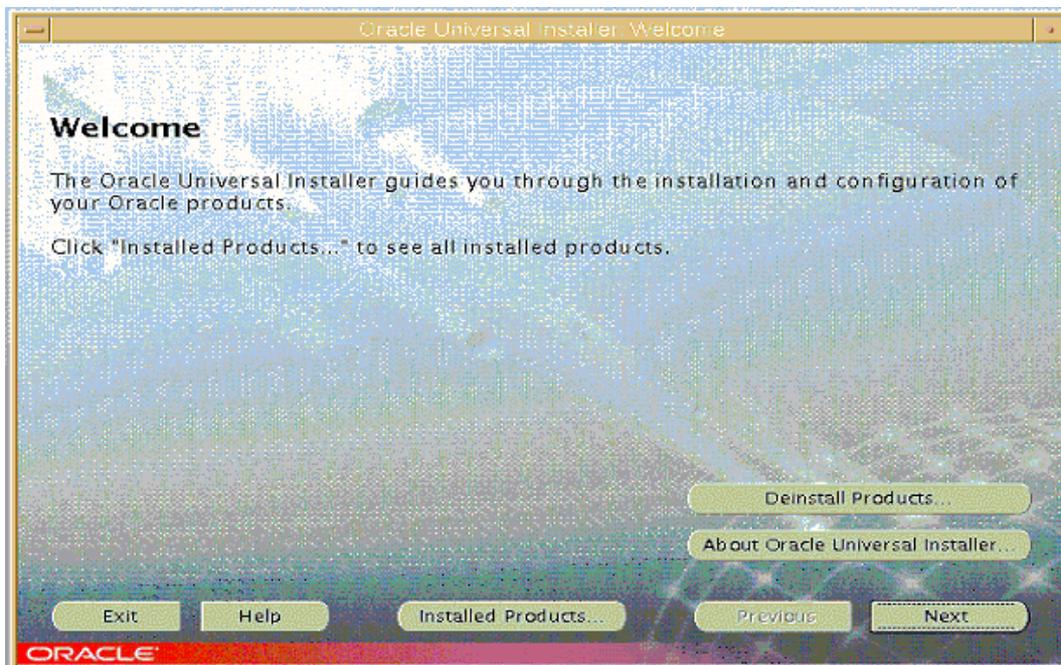
```
$ /<cdrom_mount_point>/runInstaller
```

OUI (*Oracle Universal Installer*) asks you to run `rootpre.sh` as root.

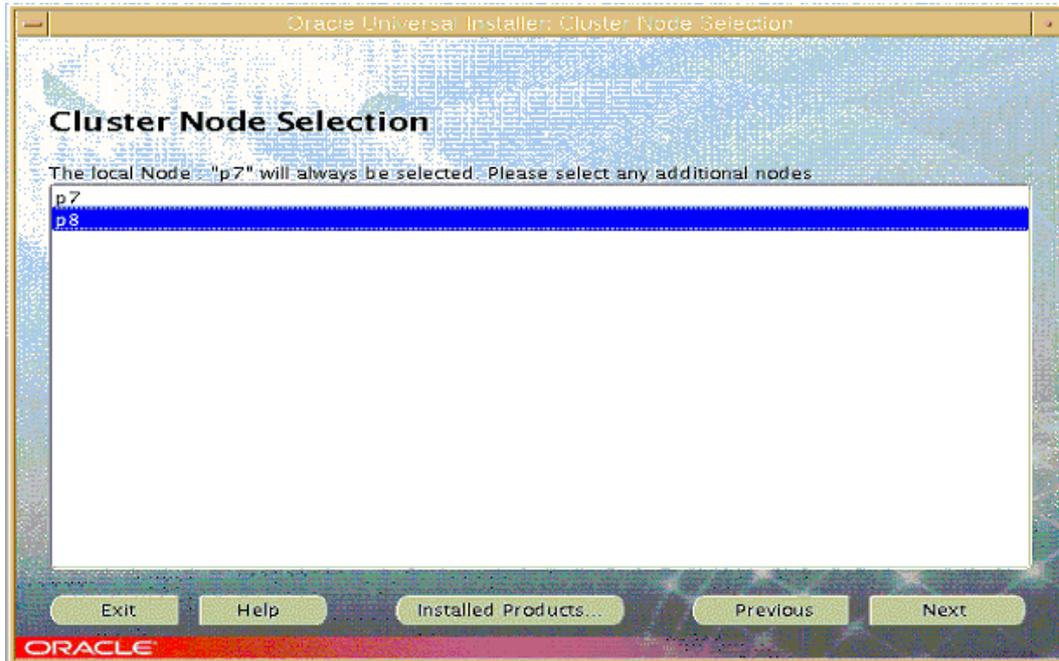
At the OUI Welcome screen, click Next.

The OUI will ask the user to run `/tmp/orainstRoot.sh` in a separate window, if it's the first Oracle product install on this machine. This script creates the file `/etc/oraInst.loc`, which is used by OUI for the list of installed products.

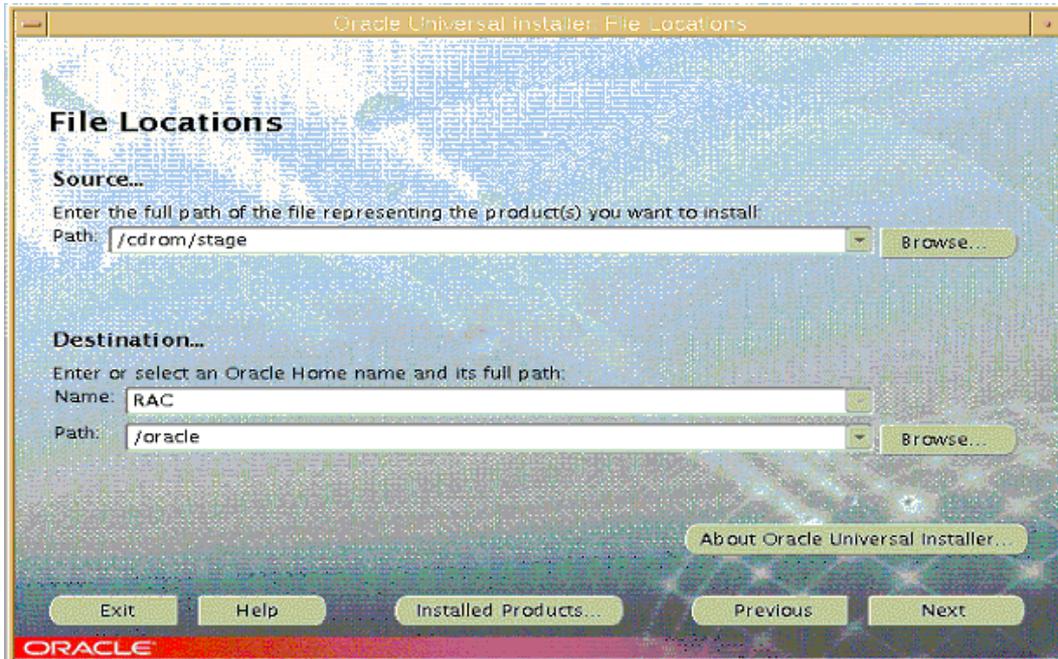
Remember that the **hagsuser** should be secondary group for the user oracle.



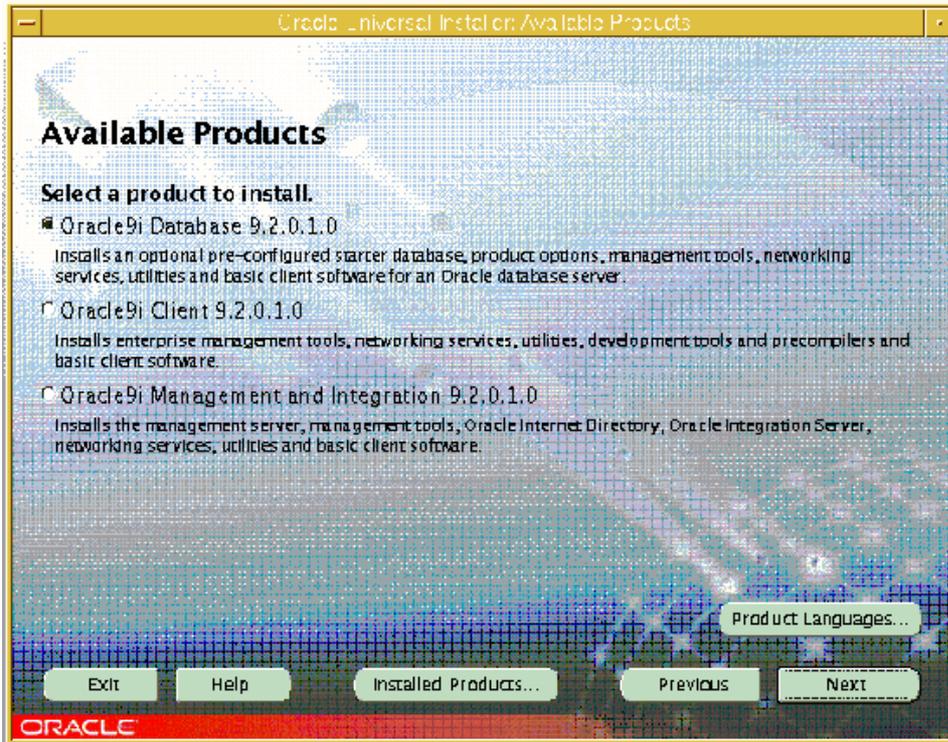
Select the other nodes on to which the Oracle RDBMS software will be installed. It is not necessary to select the node on which the OUI is currently running. Click Next.

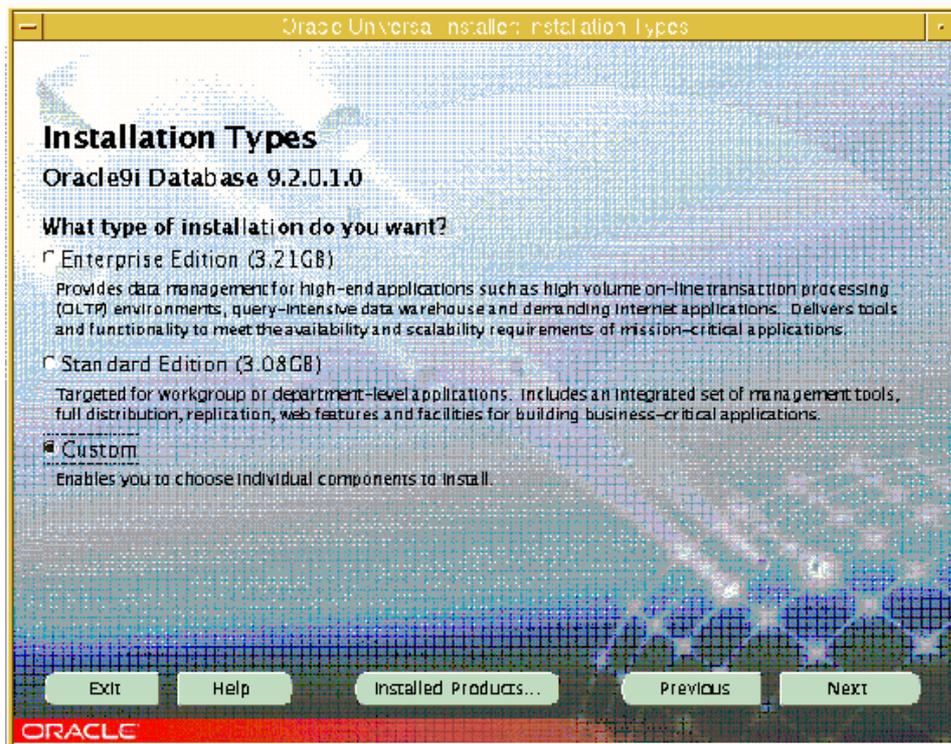


A prompt will appear for the Inventory Location (if this is the first time that OUI has been run on this system). This is the base directory into which OUI will install files. Oracle user should have write permissions on this directory. The Oracle Inventory definition can be found in the file /etc/oraInst.loc. Click OK.

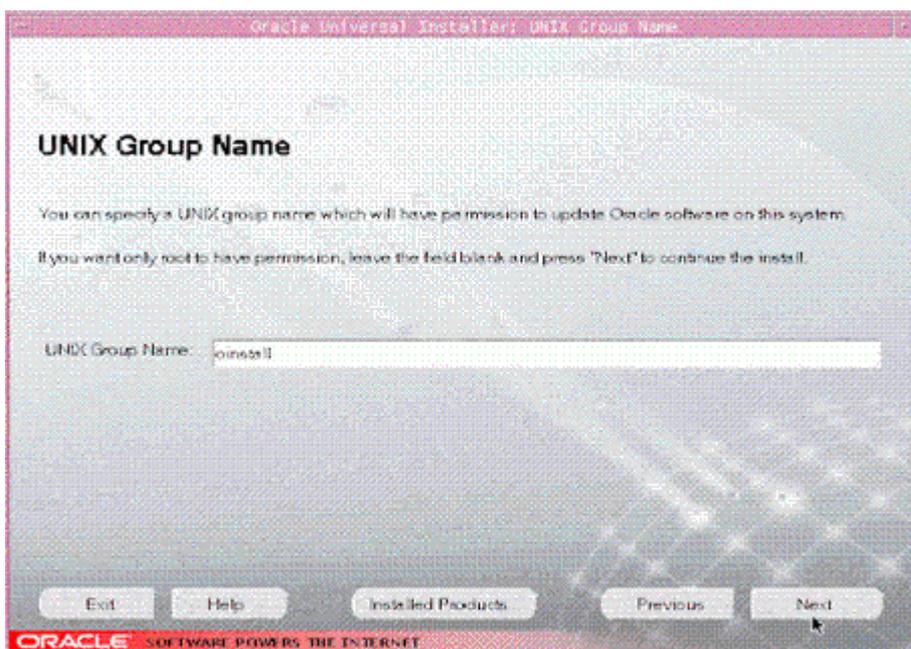


Select the Products to install. In this example, select the Oracle9i Server then click Next.

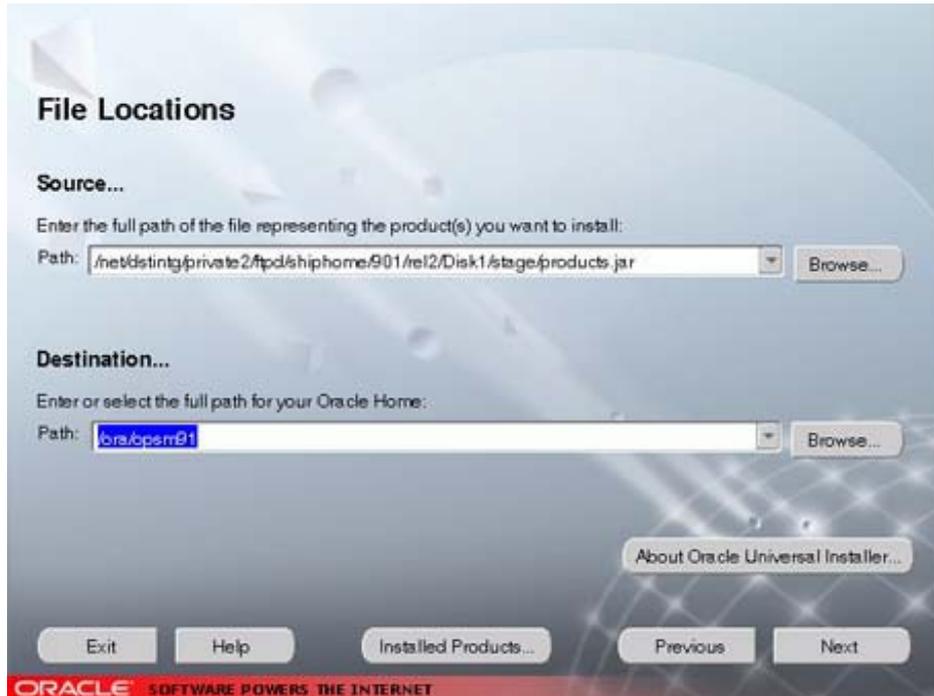




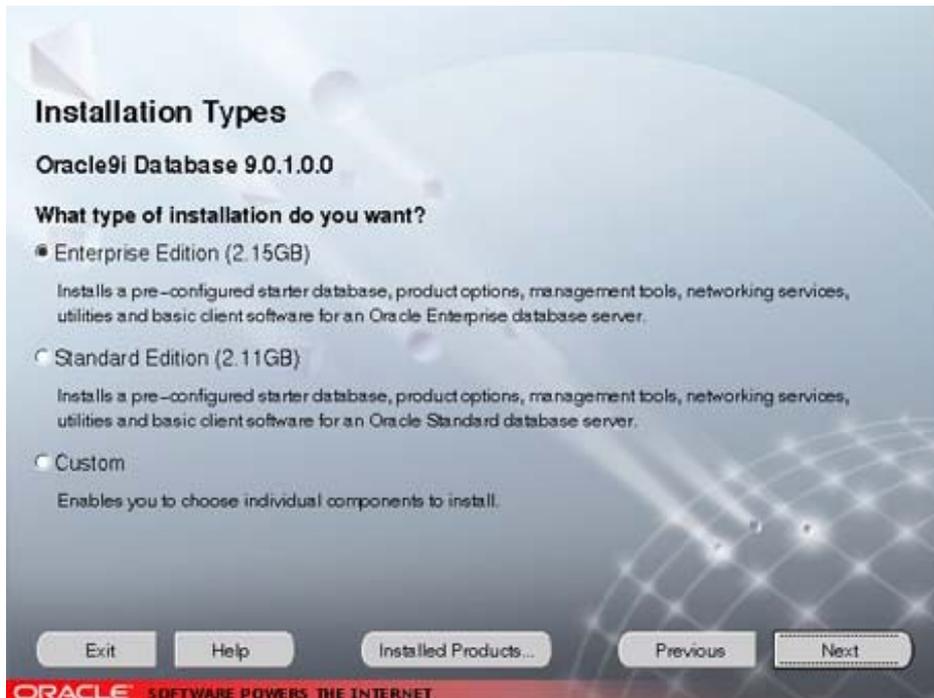
Verify the UNIX group name of the user which controls the installation of the Oracle9i software.
If the pre-installation steps have not been completed successfully, you are asked to run `/tmp/orainstRoot.sh`, forcing Oracle Inventory files, and others, to be written to the `ORACLE_HOME` directory.
This screen only appears the first time Oracle9i products are installed on the system. Click Next.



The File Location window will appear. Do not change the Source field. The Destination field defaults to the ORACLE_HOME environment variable. Click Next.



Select the installation type. Choose the Enterprise Edition option. The selection on this screen refers to the installation operation, not the database configuration. The next screen allows for a customized database configuration to be chosen. Click Next.

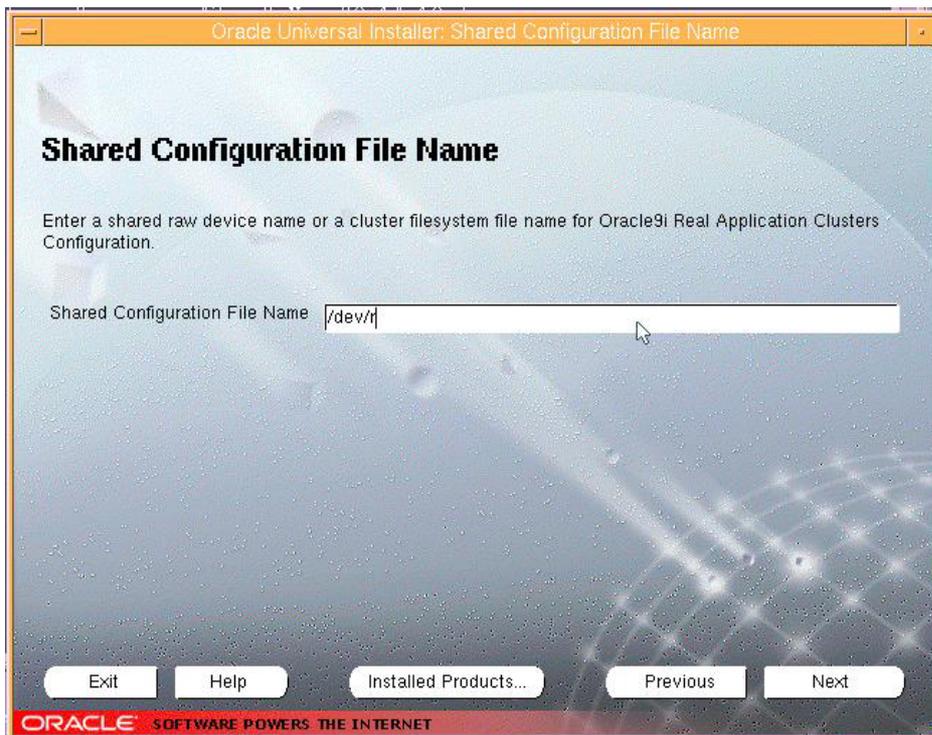


Select the configuration type. In this example, “Customized” configuration is selected; so, a customizable database will be created. Click Next.

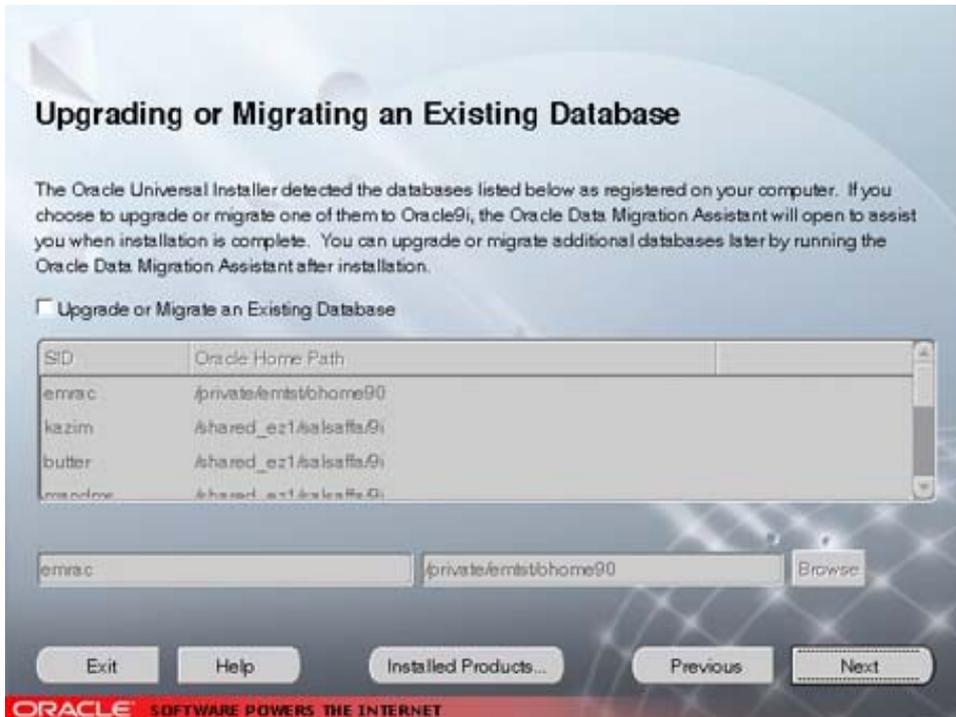


Identify the raw partition in to which the Oracle9i Real Application Clusters (RAC) configuration information will be written. It is recommended that this raw partition has a minimum size of 100MB.

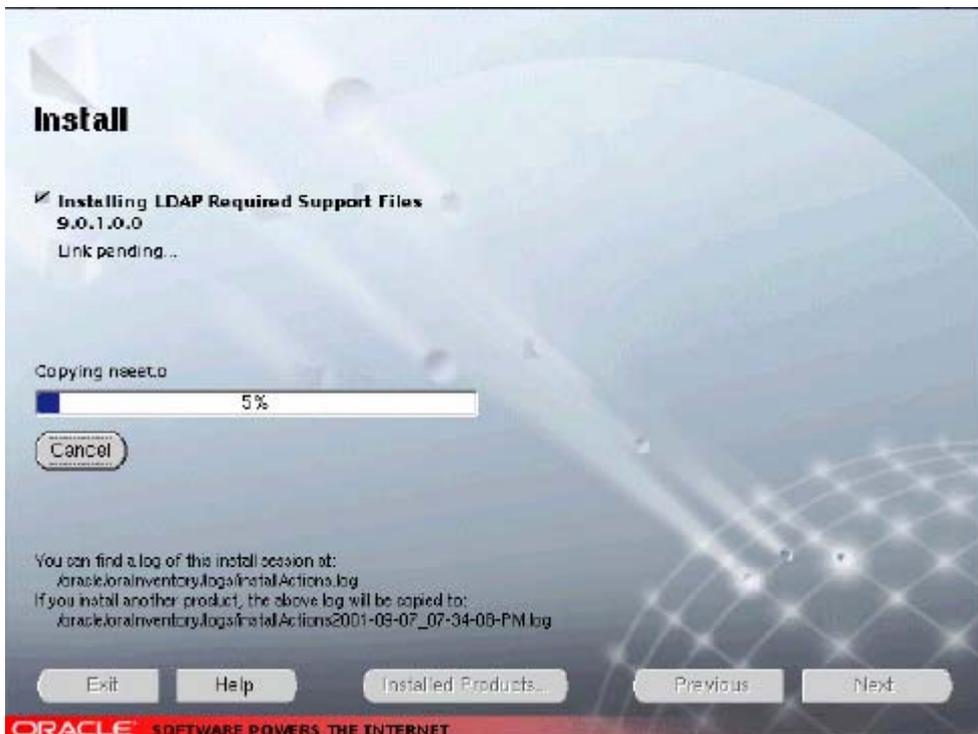
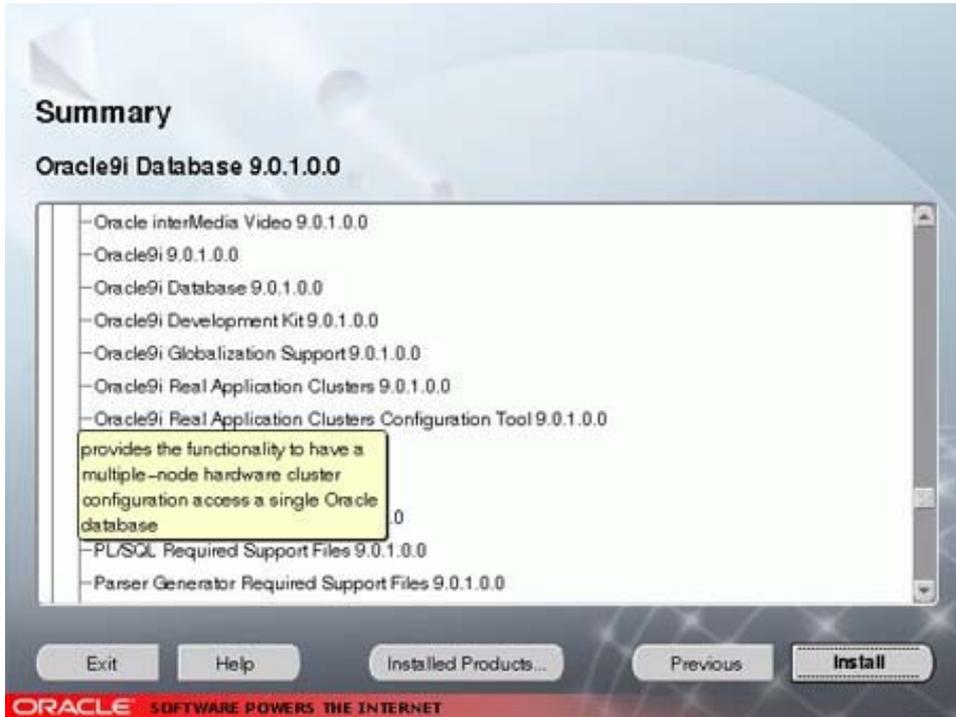
Enter the name of the raw device previously created; in our case, “/dev/trac_svrconfig”.



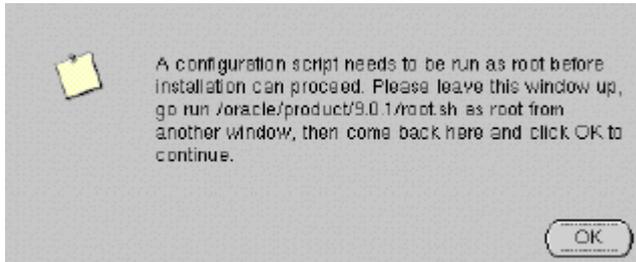
An option to Upgrade or Migrate an existing database is presented. Do NOT select the radio button. The Oracle Migration utility is not able to upgrade a RAC database, and will error if selected to do so.



The Summary screen will be presented. Confirm that the RAC database software will be installed and then click Install. The OUI will install the Oracle9i software on to the local node, and then copy this information to the other nodes selected.



During the installation, you will be prompted to give the location of the second disk.
 And so on until the disk #5.

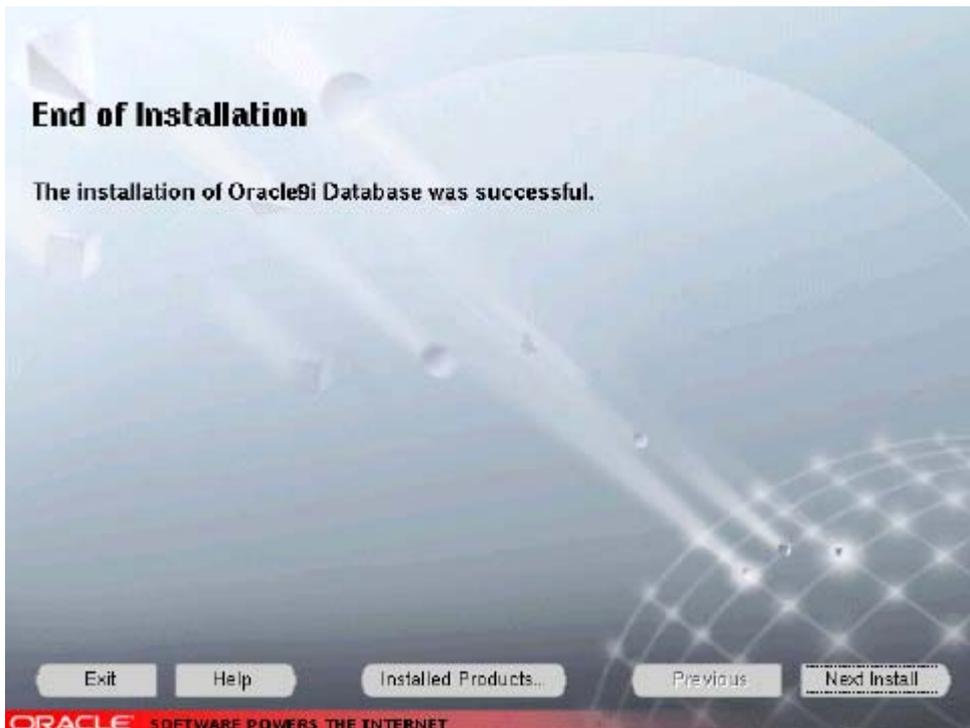


From the root command prompt, execute **/oracle/product/9.2.0/root.sh**. The results of this are shown below :

```
itsolt* # /oracle/product/9.0.1/root.sh
Running Oracle9 root.sh script...
The following environment variables are set as:
ORACLE_OWNER= oracle
ORACLE_HOME= /oracle/product/9.0.1

Enter the full pathname of the local bin directory: [/usr/local/bin]:
Copying dbhome to /usr/local/bin ...
Copying oracenv to /usr/local/bin ...
Copying oraenv to /usr/local/bin ...

Creating /etc/oratab file...
Adding entry to /etc/oratab file...
Entries will be added to the /etc/oratab file as needed by
Database Configuration Assistant when a database is created
Finished running generic part of root.sh script.
Now product-specific root actions will be performed.
```



- Make a symbolic link (oracle user) from
/var/opt/oracle/srvConfig.loc → /etc/srvConfig.loc
`ln -s /var/opt/oracle/srvConfig.loc /etc/srvConfig.loc`
- Check **/etc/srvConfig.loc** file on each node, with a single line :
`srvconfig_loc=/dev/rrac_srvconfig`
- In Sqlplus, when you “connect / as sysdba” if you get “insufficient privileges” (even if you are the user “oracle” with group “dba”), as root, execute the following command :
`touch /etc/passwd`

16 CREATING THE DATABASE MANUALLY

One easy way to create a parallel database is to create first a standalone, typical database (named RAC), on one node, but with datafiles located on the concurrent disks (raw devices).

Once your RAC database is up and running, it is easy to parallelize it, by adding a second thread of redo logs, and a second undo tablespace.

On a SQL*Plus shell, connected as sysdba, do

```
create undo tablespace UNDOTBS2 datafile '/dev/rrac_undotbs02' SIZE
200M REUSE EXTENT MANAGEMENT LOCAL;

alter database add logfile thread 2
GROUP 4 ('/dev/rrac_redo04') SIZE 100M REUSE,
GROUP 5 ('/dev/rrac_redo05') SIZE 100M REUSE,
GROUP 6 ('/dev/rrac_redo06') SIZE 100M REUSE;

alter database enable public thread 2;

@%ORACLE_HOME%/rdbms/admin/catclust.sql
```

This example is for a raw-devices implementation. Replace 'dev' by the GPFS file-system mount point for a GPFS implementation (idem for the examples you will find below in this document).

Then shutdown the standalone database.

Update the ORACLE_SID environment variable to RAC1 on the first node, and RAC2 on the second node. This are the instance names, and RAC is still the database name.

Update the initRAC.ora file, according to the example provided in [Appendix G : RAC Database creation script & init.ora](#)

Don't forget to be sure that this initRAC.ora file is the same on the two nodes.

Then, on each node, restart the database using sqlplus / as sysdba.

You will have a parallel database running.

To shutdown the database, you will have to shutdown it manually on every node.

Using this manual way, you have to execute the administration tasks (startup and shutdown commands) on all the nodes. If your cluster is made with two nodes (98% of the clusters), it is not so difficult.

If you have a lot of nodes (e.g. 8) in your cluster, you'd better use the svrctl tool.

17 CREATING THE DATABASE USING DBCA

Srvctl is a tool to manage the parallel databases with a single clustered command. You can start, stop the rac databases and listeners on all the nodes of a cluster, very easily

The database and the two instances can be created by the Oracle Database Assistant. DO NOT choose to create the database by copying one of the preconfigured database, as the DDL statement used to create the TEMP tablespace does not work against a raw device.

17.1 GSD SETUP

On one node, initialize the configuration raw device by running
srvconfig -init

Start the Global services Daemon (as oracle) on each node (runs as OPSDaemon)

```
$ gsd.sh &  
$ ps -efl | grep -i 'ops'
```

gsd (Oracle Global Services Daemon) needs to be running on each node under oracle privileges; start gsd by executing **gsd.sh**

gsd is used by OEM and global tools such as **srvctl** to execute commands on all the nodes at the same time.

The configuration device should be pointed to by SRVM_SHARED_CONFIG (containing /var/opt/oracle/srvConfig.loc for example)

17.2 SRVCTL TOOL SETUP

- 1) on one node : **srvctl add db -p <db_name> -o <oracle_home>**
- 2) for each instance of the database :
srvctl add instance -p <db_name> -i <SID> -n <node_name>
(it is advised to set the SID to db_name plus instance_number)
- 3) on each node, check the configuration : **srvctl config -p <db_name>**
- 4) on each node, create (or update) oratab file (in /etc directory) with the following line :
<db_name>:<\$ORACLE_HOME>:N
- 5) create the udump, cdump and bdump directories
- 6) set the SID in .profile of oracle user
- 7) create init.ora file for each node (it is possible to prefix parameter which are local to the instance with the instance name. For example, thread, instance_name, rollback_segments...
It is possible to create a spfile with **“create spfile=’/dev/rac_spfile’ from pfile=’/.../init.ora”**.

- 8) create the database creation script. You can use \$ORACLE_HOME/srvn/clustdb.sql script as sample. This script is also presented in [Appendix G : RAC Database creation script & init.ora](#).
- 9) you can create a password file. Under \$ORACLE_HOME execute
orapwd file=\$ORACLE_HOME/dbs/orapw\$ORACLE_SID password=####
- 10) as sysdba, execute the script.

17.3 THE DATABASE CONFIGURATION ASSISTANT (DBCA)

Create a parameter file **\$HOME/dbca_raw_config** which will be used by DBCA to map the typical tablespaces to raw devices. This parameter file needs to be pointed to by the environment variable DBCA_RAW_CONFIG and makes easier to create the database (see example in [Appendix D : examples of configuration files](#)).

First use NETCA to create listeners configuration on each node.

Start DBCA

Do not use a pre-configured database (as the DDL statement used to create the TEMP tablespace does not work against a raw device); instead have DBCA create a customized database (**create database**).

DBCA can also be used to cleanly remove instances or add new instance to an existing multi-instance database

Optionally back up the spfile

```
SQL> create pfile='?/dbs/initRAC.ora' from spfile='/dev/rrac_spfile'
```

17.4 STARTUP/SHUTDOWN SERVICES

GSD must be running for srvctl to be able to run each command on all the nodes.

```
$ srvctl command -h -> help
```

To start/stop/check all instances and listeners: `su - oracle`
`srvctl start|stop|status -p <db_name>`

To start/stop a given instance and listener: `srvctl start|stop|status -p <db_name>`
`-i <instance_name>`

To start/stop the listeners only: `srvctl start|stop -p <db_name> -s lsnr`

To start/stop the instances only: `srvctl start|stop -p <db_name> -s inst`

To list the instances: `srvctl config -p <db_name>`

To get environment information : `srvctl get env -p <db_name>`

To set an env. variable globally : `srvctl set env -p <db_name> LANG=en`

To start/stop the Oracle Intelligent Agent: `agentctl start`

To start/stop the Oracle Management Server:
`oemctl start oms`
`oemctl status oms`
`oemctl stop oms`

To start an OEM Console :
`oemapp console`
`oemapp dbastudio`

To start Oracle http server (Apache): `OH/Apache/Apache/bin/apachectl start`

17.5 CONFIGURE LISTENER.ORA / SQLNET.ORA / TNSNAMES.ORA

Use `netca` and/or `netmgr` to check the configuration of the listener and configure Oracle Net services (by default the Net service may be equal to the global database name (see instance parameter `service_names`)).

17.6 CONFIGURE ORACLE ENTERPRISE MANAGER

Use the Java assistant `emca` to configure the Oracle Management Server, then start it :
`oemctl start oms`

Enter a `username/password` with DBA privileges to connect to the instance where the repository is to be set up. The first administrator for the domain will be `sysman/oem_temp`.

For the console to be able to give a single system image of the cluster database, it is necessary to start the Intelligent Agent on each node and to “discover” the nodes.

- Start the Oracle Intelligent Agent on each node

Add the following line for Oracle Intelligent Agent to /etc/snmpd.conf (this is specific to AIX).
smux 0.0 "" 129.1.11.106 # OEM agent (IP address of current node)

The SNMP master agent needs to be restarted:

```
#stopsrc -s snmpd  
#startsrc -s snmpd
```

Then start the OEM agent:

```
$agentctl start
```

- Check /etc/oratab

The file should contain a reference to the database name, not to the instance name.

The last field should always be "N" on a RAC environment to avoid 2 instances of the same name to be started.

- Register the database with srvctl (this should not be necessary if the database was created by DBCA)

```
srvctl add db -p <db_name> -o <ORACLE_HOME path>  
srvctl add instance -p <db_name> -i <SID1> -n <node1>  
srvctl add instance -p <db_name> -i <SID2> -n <node1>
```

17.7 TROUBLESHOOTING

Useful SQL:

```
SQL> SELECT * FROM V$ACTIVE_INSTANCES
```

18 APPENDIX A : INSTALLATION DIRECTORIES.

--- /etc	oratab
--- /usr /local/bin	oraenv, coraenv, dbhome
--- /etc	srvConfig.loc oraInst.loc
--- /oracle --- (\$ORACLE_BASE)	
/admin /<SID> /	bdump cdump create pfile udump
	...
/product /9.2.0./	bin dbs rdbms sqlplus
	...
/oraInventory	
/jre/1.1.8	
/oui	

19 APPENDIX B : SSA DISK & ADAPTER MICROCODE MANAGEMENT

- **SSA Disks**

Because the disks are used in a concurrent mode, you have to disable the write cache on each concurrent disk. Use `smit device, ssa disks, logical` to proceed

The IBM SSA Customer support page resides on <http://www.storage.ibm.com/hardsoft/products/ssa>

To list all the SSA disks of the cluster : `lscfg | grep pdisk`
 To check the microcode level : `lscfg -vl pdisknn` (ROS Level and ID)
 To determine the disk model : `lscfg -vl pdisknn` (Machine Type and Model)

Depending on the disk model, the latest level of microcode to download and install on all the SSA disks of the cluster is :

SSA DISKS current level of microcode (august 2002)

Disk Drive Microcode	
Drive Type	Current Level
DFHC (RAMST)	8877
DFHC (RAMSC)	9590
DCHC (RAMSC)	9595
DFHC (CUSJ)	9911
DCHC (CUSM)	9911
DGHC (CUSJ)	9911
DGHC (CUSM)	9911
DRVC (CUSH)	0023
DRHC (CUSS)	0012
DMVC (CUSN)	*
ST37,ST33,ST31,73LP	2010
UCD2	2004
7133 D40/T40 Enclosure Microcode	0020

From the site <http://www.storage.ibm.com/hardsoft/products/ssa>, download the file `ssa_dasd1.tar` onto your system in your temporary directory

How to apply a new **disk drive microcode**

- A.
- 1.Login as Root
 - 2.cd to your temporary directory
 - 3.Type `tar -xvf ssacode433.tar`
 - 4.Run `smit install`
 - 5.Select install & update software
 - 6.Select install & update from ALL available software
 - 7.Use the directory that you saved and unpacked the ssacode433.tar file into as the install device
 - 8.Select all filesets in this directory for install
 - 9.Execute the command
- B.
1. At the prompt, execute the command `diag`
 2. Task Selection(Diagnostics, Advanced Diagnostics, Service Aids, etc.)
 3. SSA Service Aids
 - 4.Display/Download Disk Drive Microcode
 - 5.Download Microcode to all SSA Physical Disk Drives
 - 6.Continue with the microcode installation
 - 7.No (because software is in /etc/microcode)
 - 8.Do you want to continue? Yes....

This will upgrade the microcode on all the disks with a lower level of microcode than the one installed in /etc/microcode during the phase A.

• **SSA adapters :**

To list all the SSA adapters of the node : `lscfg | grep ssa`

To check the microcode level : `lscfg -vl ssann (ROS Level and ID)`

SSA Adapter Cards current level of microcode (august 2002)

Adapter Microcode				
Adapter	Feature Code	Adapter Type	Loadable Microcode Level	Current Level
SSA 4-Port	6214	4-D	00	3201
Enhanced SSA 4-Port	6216	4-G	01	3202
RAID (MCA)*	6217	4-I	02	2904
RAID (PCI)*	6218	4-J	02	3700
Enhanced RAID (MCA)	6219	4-M	04	7601
Enhanced RAID (PCI)	6215	4-N	04	7601
Advanced SerialRAID Adapter	6225	4-P	05	B800
Advanced SerialRAID Plus	6230	4-P	05	B800

* NOTE: Adapters 4-I (6217) and 4-J (6218) do not support D40 or T40 enclosures

To install the new adapter microcode, download the file by clicking on the current level needed and install it as described above.

20 APPENDIX C : LOGICAL VOLUMES CREATION

This is for raw devices implementation

```
#!/bin/ksh

# Creation of a concurrent volume group
# To be executed on the primary node (node1)
# 64 is a free major number. It has to be unused on all the nodes

mkvg -f -c -y rac92_cvg -V'64' -s'32' <disk list>

# Creation of the minimum logical volumes (raw devices)
# The number is the size of the LV (number of 32 MB physical partitions)
# To be executed on the primary node (node1)

mklv -y'rac_system01' rac92_cvg      16 ext_disk      # 512 MB
mklv -y'rac_undotbs01' rac92_cvg     8 ext_disk      # 256 MB
mklv -y'rac_undotbs02' rac92_cvg     8 ext_disk      # 256 MB

mklv -y'rac_redo01'   rac92_cvg      4 ext_disk      # 128 MB
mklv -y'rac_redo02'   rac92_cvg      4 ext_disk      # 128 MB
mklv -y'rac_redo03'   rac92_cvg      4 ext_disk      # 128 MB
mklv -y'rac_redo04'   rac92_cvg      4 ext_disk      # 128 MB
mklv -y'rac_redo05'   rac92_cvg      4 ext_disk      # 128 MB
mklv -y'rac_redo06'   rac92_cvg      4 ext_disk      # 128 MB

mklv -y'rac_control01' rac92_cvg     1 ext_disk      # 32 MB
mklv -y'rac_control02' rac92_cvg     1 ext_disk      # 32 MB
mklv -y'rac_control03' rac92_cvg     1 ext_disk      # 32 MB

mklv -y'rac_spfile'   rac92_cvg      1 ext_disk      # 32 MB
mklv -y'rac_srvconfig' rac92_cvg     4 ext_disk      # 128 MB

mklv -y'rac_data'     rac92_cvg      4 ext_disk      # 128 MB
mklv -y'rac_index'    rac92_cvg      4 ext_disk      # 128 MB
mklv -y'rac_temp'     rac92_cvg      4 ext_disk      # 128 MB

varyoffvg rac92_cvg

# To be executed on the secondary node (node2)

importvg -y rac92_cvg -V64 <one of the ssa disks>
varyonvg rac92_cvg
chfs -c rac92_cvg

# To be executed on all the nodes (node1 & node2) when rac92_cvg is varied on

chown oracle.dba /dev/*rac_*
chmod go+rw /dev/*rac_*
```

21 APPENDIX D : EXAMPLES OF CONFIGURATION FILES

This appendix provides examples of the configuration files that are mentioned in the document.

/etc/hosts for HACMP (on both nodes)

```
192.128.194.1    node1
192.128.194.2    node2
10.10.11.141    interconnect_node1
10.10.11.142    interconnect_node2
```

/etc/hosts.equiv for HACMP (on both nodes)

```
node1           root    oracle
node2           root    oracle
interconnect_node1  root    oracle
interconnect_node2  root    oracle
```

.rhosts

In the root's and oracle's home directory, put the list of machines.

\$HOME/.rhosts for HACMP configuration :

```
node1           root    oracle
node2           root    oracle
interconnect_node1  root    oracle
interconnect_node2  root    oracle
```

/etc/srvConfig.loc

Path to the server configuration raw device

```
srvconfig_loc=/dev/rrac_srvconfig
```

\$HOME/dbca_raw_config

Raw devices parameter file for the Database Configuration Assistant

```
system=/dev/rrac_system01
undo1=/dev/rrac_undotbs01
undo2=/dev/rrac_undotbs02
redo01=/dev/rrac_redo01
redo02=/dev/rrac_redo02
redo03=/dev/rrac_redo03
redo04=/dev/rrac_redo04
redo05=/dev/rrac_redo05
redo06=/dev/rrac_redo06
control1=/dev/rrac_control01
control2=/dev/rrac_control02
control3=/dev/rrac_control03
```

listener.ora

```

LISTENER =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS_LIST =
        (ADDRESS = (PROTOCOL = TCP)(HOST = NODE1)(PORT = 1521))
      )
    )
  )

SID_LIST_LISTENER =
  (SID_LIST =
    (SID_DESC =
      (ORACLE_HOME = /oracle/product/9.2.0)
      (SID_NAME = RAC1)
    )
  )

```

tnsnames.ora implementing TAF

```

RAC2 =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST = node2)(PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVICE_NAME = RAC)
      (INSTANCE_NAME = RAC2)
    )
  )

RAC1 =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST = node1)(PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVICE_NAME = RAC)
      (INSTANCE_NAME = RAC1)
    )
  )

RAC =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST = node1)(PORT = 1521))
      (ADDRESS = (PROTOCOL = TCP)(HOST = node2)(PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVICE_NAME = RAC)
      (FAILOVER_MODE =
        (TYPE = SELECT)
        (METHOD = PRECONNECT)
        (RETRIES = 20)
        (DELAY = 60)
      )
    )
  )

```

srvconfig.txt

Server configuration file exported by srvconfig

```
RAC:/oracle/product/9.2.0:N  
  
RAC.spfile = /dev/rrac_spfile  
RAC.node_list = "node1,node2"  
RAC.inst_oracle_sid = (RAC1,RAC2)
```

22 APPENDIX E : ORACLE TECHNICAL NOTES

This appendix provides some useful notes coming from Oracle support. These notes can be found in Metalink.

Note #2064876.102 : How to setup High Availability Group Services (HAGS) on IBM AIX/RS6000.

Article-ID: <Note:2064876.102>
 Circulation: PUBLISHED (EXTERNAL)
 Folder: server.OPS.Parallelservers
 Topic: - - - IBM RS6000 and SP
 Title: How to setup High Availability Group Services (HAGS) on IBM AIX/RS6000
 Document-Type: BULLETIN
 Impact: LOW
 Skill-Level: CASUAL
 Server-Version: 08.01.XX.XX.XX
 Updated-Date: 04-FEB-2002 08:21:25
 References:
 Shared-Refs:
 Authors: BLEVE.US
 Attachments: NONE
 Content-Type: TEXT/PLAIN
 Keywords: LMON; OPS; PARALLEL; SERVER; SUBCOMP-OPS;
 Products: 236/RDBMS (08.01.XX.XX.XX);
 Platforms: 319 (4.3);

Purpose =====

This article gives quick reference instructions on how to configure High Availability Group Services (HAGS) on IBM AIX RS6000 for Oracle 8.1.X.

Scope and Application =====

These instructions are helpful to any customer using Oracle on IBM AIX/RS6000 on which HACMP is installed.

How to Configure High Availability Group Services (HAGS) =====

In order to configure High Availability Group Services (HAGS), you need to be connected as root.

Do the following on all nodes that form the cluster:

1. Create the "hagsuser" group and place "oracle" into the "hagsuser" group:

Verify the group does not exist:

```
# grep hagsuser /etc/group
```

If this returns nothing do the following:

```
# smitty groups
```

Select "Add a Group" and fill in the following:

```
Group Name    ----> hagsuser
USER list    ----> oracle
```

You can take the defaults for the other settings. Also note that after the group is created you will have to log out and log back in as "oracle" to be sure "oracle" is part of the "hagsuser" group.

2. Change the permissions on the "cldomain" executable:

```
# chmod a+x /usr/sbin/cluster/utilities/cldomain
```

3. Change the group to "hagsuser" for the "svcsdsocket.rac_cluster" socket:

```
# chgrp hagsuser /var/ha/soc/grpsvcsdsocket.rac_cluster
```

4. Change the group permissions for the "grpsvcsdsocket.rac_cluster" socket:

```
# chmod g+w /var/ha/soc/grpsvcsdsocket.rac_cluster
```

The HAGS socket needs to be writeable by "oracle" and the "cldomain" executable needs to be executable by "oracle". By configuring the group and permissions for the "grpsvcsdsocket.rac_cluster" file the instance will be able to communicate with HAGS and the instance will mount.

References

=====

Oracle Installation Guide for AIX RS6000, release 8.1.5.

Search Words

=====

OPS HAGS RS6000

You should have something like the following :

```
ls -l /var/ha/soc
```

```
total 4
srw-rw-rw-  1 root    haermm           0 Jun 19 16:58 em.clsrv.rac92_cluster
srw-rw----  1 root    haermm           0 Jun 19 16:58 em.rmsrv.rac92_cluster
drwxrwxrwx  2 root    system          512 Jun 19 17:51 grpsvcs.clients.rac92_cluster
srw-rw-rw-  1 root    hagsuser        0 Jun 19 16:57 grpsvcsdsocket.rac92_cluster
drwxrwx---  2 root    haermm          512 Jun 19 16:57 haem
drwxr-x---  2 root    system          512 Apr 02 23:30 hats
drwxr-xr-x  2 root    system          512 Jun 19 16:57 topsvcs
```

Note # 115792.1 , explaining how to setup HACMP cluster interconnect adapter.

```

Article-ID:          <Note:115792.1>
Circulation:        PENDING_DELETE (EXTERNAL)
Folder:             server.OPS.Parallelservers
Topic:              - - - IBM RS6000 and SP
Title:              AIX: ORA-600 [KCCSBCK_FIRST] starting up second OPS
                    Instance
Document-Type:      BULLETIN
Impact:             LOW
Skill-Level:        NOVICE
Server-Version:     08.01.06.0X to 08.01.06.0X
Updated-Date:       05-FEB-2002 13:10:49
References:
Shared-Refs:
Authors:            RKIRCHHE.DE
Attachments:        NONE
Content-Type:       TEXT/PLAIN
Keywords:           HACMP; OPS;
Errors:             ORA-600; [KCCSBCK_FIRST];
Products:           5/RDBMS;
Platforms:          319;
  
```

```

*****
This article is being delivered in Draft form and may contain
errors. Please use the MetaLink "Feedback" button to advise
Oracle of any issues related to this article.
*****
  
```

PURPOSE

This article helps to resolve problems with Oracle Parallel Server startup related to HACMP configuration

SCOPE & APPLICATION

How to setup HACMP cluster interconnect adapter

Oracle Parallel Server software is successfully installed.
The first OPS instance starts without errors.

Trying to start a second OPS instance on another cluster node fails with ORA-600 [KCCSBCK_FIRST].

\$ORACLE_HOME/bin/lsnodes will list all cluster nodes.

/usr/sbin/cluster/diag/clverify doesn't show any errors.

Check HACMP interconnect network adapter configuration with
/usr/sbin/cluster/utilities/cllsif

cllsif on a working configuration should look like this:

Adapter	Type	Network	NetType	Attribute	Node	IP Address	InterfaceName
interconnect_node1	service	rac92_network	ether	private	node1	10.10.11.141	en3
interconnect_node2	service	rac92_network	ether	private	node2	10.10.11.142	en5

RELATED DOCUMENTS

<List related manuals, articles and other documents.>

23 APPENDIX F : HACMP CLUSTER VERIFICATION OUTPUT.

All the following outputs come from an HACMP configuration for raw devices implementation which was including a parallel instance database (HACMP concurrent mode) and a single instance database (HACMP rotating mode). (Information about the concurrent mode configuration is in bold.)

Cluster description : output from a 'cllscf' command.

```
Cluster Description of Cluster rac92_cluster
Cluster ID: 92
There were 1 networks defined : rac92_network
There are 2 nodes in this cluster.
```

NODE node1:

```
This node has 1 service interface(s):
```

```
Service Interface interconnect_node1:
  IP address: 10.10.11.141
  Hardware Address:
  Network:    rac92_network
  Attribute:  private
  Aliased Address?: False
```

```
Service Interface interconnect_node1 has no boot interfaces.
Service Interface interconnect_node1 has no standby interfaces.
```

NODE node2:

```
This node has 1 service interface(s):
```

```
Service Interface interconnect_node2:
  IP address: 10.10.11.142
  Hardware Address:
  Network:    rac92_network
  Attribute:  private
  Aliased Address?: False
```

```
Service Interface interconnect_node2 has no boot interfaces.
Service Interface interconnect_node2 has no standby interfaces.
```

Breakdown of network connections:

Connections to network rac92_network

```
Node node1 is connected to network rac92_network by these interfaces:
  interconnect_node1
```

```
Node node2 is connected to network rac92_network by these interfaces:
  interconnect_node2
```

Cluster network interfaces : output from a 'cllsif' command.

Adapter	Type	Network	NetType	Attribute	Node	IP	Address	InterfaceName
interconnect_node1	service	rac92_network	ether	private	node1	10.10.11.141		en3
interconnect_node2	service	rac92_network	ether	private	node2	10.10.11.142		en5

Cluster resource group : output from a 'clshowres' command.

```

Resource Group Name                racdisk
Node Relationship                   concurrent
Participating Node Name(s)        node1 node2
Dynamic Node Priority
Service IP Label
Filesystems
Filesystems Consistency Check      fsck
Filesystems Recovery Method        sequential
Filesystems/Directories to be exported
Filesystems to be NFS mounted
Network For NFS Mount
Volume Groups
Concurrent Volume Groups           rac92_cvg
Disks
Connections Services
Fast Connect Services
Shared Tape Resources
Application Servers
Highly Available Communication Links
Miscellaneous Data
Automatically Import Volume Groups false
Inactive Takeover                  false
Cascading Without Fallback         false
9333 Disk Fencing                  false
SSA Disk Fencing                   false
Filesystems mounted before IP configured false
Run Time Parameters:

Node Name                           node1
Debug Level                          high
Host uses NIS or Name Server        false
Format for hacmp.out                Standard

Node Name                           node2
Debug Level                          high
Host uses NIS or Name Server        false
Format for hacmp.out                Standard

```

24 APPENDIX G : RAC DATABASE CREATION SCRIPT & INIT.ORA.

Create_database.sql

```

CREATE DATABASE RAC
CONTROLFILE REUSE
MAXINSTANCES 8
MAXLOGHISTORY 100
MAXLOGFILES 32
MAXLOGMEMBERS 5
MAXDATAFILES 200
DATAFILE '/dev/rrac_system01' SIZE 325M REUSE
UNDO TABLESPACE "UNDOTBS1" DATAFILE '/dev/rrac_undotbs01' SIZE 200M REUSE
CHARACTER SET WE8ISO8859P1
NATIONAL CHARACTER SET AL16UTF16
LOGFILE GROUP 1 ('/dev/rrac_redo01') SIZE 100M REUSE,
GROUP 2 ('/dev/rrac_redo02') SIZE 100M REUSE,
GROUP 3 ('/dev/rrac_redo03') SIZE 100M REUSE;

@%ORACLE_HOME%/rdbms/admin/catalog.sql
@%ORACLE_HOME%/rdbms/admin/catexp.sql
@%ORACLE_HOME%/rdbms/admin/catproc.sql

connect system/manager
@%ORACLE_HOME%/dbs/pupbld.sql

create undo tablespace UNDOTBS2 datafile '/dev/rrac_undotbs02' SIZE 200M
REUSE EXTENT MANAGEMENT LOCAL;

alter database add logfile thread 2
GROUP 4 ('/dev/rrac_redo04') SIZE 100M REUSE,
GROUP 5 ('/dev/rrac_redo05') SIZE 100M REUSE,
GROUP 6 ('/dev/rrac_redo06') SIZE 100M REUSE;

alter database enable public thread 2;

@%ORACLE_HOME%/rdbms/admin/catclust.sql

```

Another database creation script is provided with the Oracle 9i software. You can find the whole script in your \$ORACLE_HOME/svrm/admin directory.

initRAC.ora

```
cluster_database=true
cluster_database_instances=2
db_name=RAC
db_domain=""

# first instance
RAC1.thread=1
RAC1.instance_name=RAC1
RAC1.instance_number=1
RAC1.undo_tablespace=UNDOTBS1

# second first instance
RAC2.thread=2
RAC2.instance_name=RAC2
RAC2.instance_number=2
RAC2.undo_tablespace=UNDOTBS2

control_files=("/dev/rrac_control01", "/dev/rrac_control02",
"/dev/rrac_control03")

db_block_size=8192
db_cache_size=52428800
open_cursors=300
timed_statistics=TRUE
compatible=9.2.0
java_pool_size=52428800
large_pool_size=1048576
shared_pool_size=52428800
processes=150
fast_start_mttr_target=300
resource_manager_plan=SYSTEM_PLAN
sort_area_size=524288
undo_management=AUTO
```

25 APPENDIX H : FILESETS TO BE INSTALLED ON THE MACHINES OF THE CLUSTER.

This appendix provides the list of filesets which must be installed for HACMP and GPFS.

25.1 HACMP 4.4

Result of the command : `lsllpp -L | grep cluster`

```
cluster.es.client.lib          4.4.1.2  COMMITTED  ES Client Libraries
cluster.es.client.rte         4.4.1.4  COMMITTED  ES Client Runtime
cluster.es.client.utils       4.4.1.2  COMMITTED  ES Client Utilities
cluster.es.clvm.rte           4.4.1.0  COMMITTED  ES for AIX Concurrent Access
cluster.es.cspoc.cmds         4.4.1.4  COMMITTED  ES CSPOC Commands
cluster.es.cspoc.dsh          4.4.1.0  COMMITTED  ES CSPOC dsh
cluster.es.cspoc.rte          4.4.1.2  COMMITTED  ES CSPOC Runtime Commands
cluster.es.plugins.dhcp        4.4.1.1  COMMITTED  ES Plugins - dhcp
cluster.es.plugins.dns         4.4.1.1  COMMITTED  ES Plugins - Name Server
cluster.es.plugins.printserver
cluster.es.server.diag        4.4.1.4  COMMITTED  ES Server Diags
cluster.es.server.events      4.4.1.5  COMMITTED  ES Server Events
cluster.es.server.rte         4.4.1.5  COMMITTED  ES Base Server Runtime
cluster.es.server.utils       4.4.1.5  COMMITTED  ES Server Utilities
cluster.es.taskguides.shrvolgrp
cluster.msg.en_US.cspoc        4.4.1.0  COMMITTED  HACMP CSPOC Messages - U.S.
cluster.msg.en_US.es.client
cluster.msg.en_US.es.server
cluster.es.client.rte         4.4.1.0  COMMITTED  ES Client Runtime
cluster.es.clvm.rte           4.4.1.0  COMMITTED  ES for AIX Concurrent Access
cluster.es.server.events      4.4.1.0  COMMITTED  ES Server Events
cluster.es.server.rte         4.4.1.5  COMMITTED  ES Base Server Runtime
cluster.es.server.utils       4.4.1.0  COMMITTED  ES Server Utilities
```

Note that the **vsm** filesets (Visual Software Manager, a X11 based smit tool) does not have to be installed.

25.2 HACMP 4.5

`cluster.adt.es.client.demos`

```
cluster.adt.es.client.include
cluster.adt.es.client.samples.clinfo
cluster.adt.es.client.samples.clstat
cluster.adt.es.client.samples.demos
cluster.adt.es.client.samples.libcl
cluster.adt.es.java.demo.monitor
cluster.adt.es.server.demos
cluster.adt.es.server.samples.demos
cluster.adt.es.server.samples.images
cluster.es.client.lib          4.5.0.4  C    F    ES Client Libraries
cluster.es.client.rte         4.5.0.5  C    F    ES Client Runtime
cluster.es.client.utils       4.5.0.3  C    F    ES Client Utilities
cluster.es.clvm.rte           4.5.0.2  C    F    ES for AIX Concurrent Access
cluster.es.cspoc.cmds         4.5.0.5  C    F    ES CSPOC Commands
cluster.es.cspoc.dsh          4.5.0.2  C    F    ES CSPOC dsh
cluster.es.cspoc.rte          4.5.0.7  C    F    ES CSPOC Runtime Commands
cluster.es.server.diag        4.5.0.7  C    F    ES Server Diags
cluster.es.server.events      4.5.0.7  C    F    ES Server Events
cluster.es.server.rte         4.5.0.7  C    F    ES Base Server Runtime
cluster.es.server.utils       4.5.0.7  C    F    ES Server Utilities
cluster.license                4.5.0.0  C    F    HACMP Electronic License
cluster.man.en_US.es.data     4.5.0.2  C    F    ES Man Pages - U.S. English
cluster.msg.en_US.cspoc       4.5.0.2  C    F    HACMP CSPOC Messages - U.S.
cluster.msg.en_US.es.client
cluster.msg.en_US.es.server
```

25.3 RSCT 2.2 (PROVIDED WITH AIX 5.1)

rsct.basic.hacmp	2.2.1.0	COMMITTED	RSCT Basic Function (HACMP/ES
rsct.basic.rte	2.2.1.2	COMMITTED	RSCT Basic Function
rsct.clients.rte	99.99.999.999	COMMITTED	Supersede Entry - Not really
rsct.compat.basic.hacmp	2.2.1.0	COMMITTED	RSCT Event Management Basic
rsct.compat.basic.rte	2.2.1.0	COMMITTED	RSCT Event Management Basic
rsct.compat.clients.hacmp	2.2.1.0	COMMITTED	RSCT Event Management Client
rsct.compat.clients.rte	2.2.1.1	COMMITTED	RSCT Event Management Client
rsct.core.auditrm	2.2.1.0	COMMITTED	RSCT Audit Log Resource
rsct.core.errm	2.2.1.0	COMMITTED	RSCT Event Response Resource
rsct.core.fsrms	2.2.1.0	COMMITTED	RSCT File System Resource
rsct.core.gui	2.2.1.0	COMMITTED	RSCT Graphical User Interface
rsct.core.hostrm	2.2.1.1	COMMITTED	RSCT Host Resource Manager
rsct.core.rmc	2.2.1.0	COMMITTED	RSCT Resource Monitoring and
rsct.core.sec	2.2.1.2	COMMITTED	RSCT Security
rsct.core.sensorm	2.2.1.0	COMMITTED	RSCT Sensor Resource Manager
rsct.core.sr	2.2.1.0	COMMITTED	RSCT Registry
rsct.core.utils	2.2.1.0	COMMITTED	RSCT Utilities
rsct.msg.EN_US.core.auditrm			
rsct.msg.EN_US.core.errm	2.2.0.0	COMMITTED	RSCT Event Response RM Msgs -
rsct.msg.EN_US.core.fsrms	2.2.0.0	COMMITTED	RSCT File System RM Msgs -
rsct.msg.EN_US.core.gui	2.2.0.0	COMMITTED	RSCT GUI Msgs - U.S. English
rsct.msg.EN_US.core.hostrm			
rsct.msg.EN_US.core.rmc	2.2.0.0	COMMITTED	RSCT RMC Msgs - U.S. English
rsct.msg.EN_US.core.sec	2.2.0.0	COMMITTED	RSCT Security Msgs - U.S.
rsct.msg.EN_US.core.sensorm			
rsct.msg.EN_US.core.sr	2.2.0.0	COMMITTED	RSCT Registry Msgs - U.S.
rsct.msg.EN_US.core.utils	2.2.0.0	COMMITTED	RSCT Utilities Msgs - U.S.
rsct.msg.en_US.core.auditrm			
rsct.msg.en_US.core.errm	2.2.0.0	COMMITTED	RSCT Event Response RM Msgs -
rsct.msg.en_US.core.fsrms	2.2.0.0	COMMITTED	RSCT File System RM Msgs -
rsct.msg.en_US.core.gui	2.2.0.0	COMMITTED	RSCT GUI Msgs - U.S. English
rsct.msg.en_US.core.hostrm			
rsct.msg.en_US.core.rmc	2.2.0.0	COMMITTED	RSCT RMC Msgs - U.S. English
rsct.msg.en_US.core.sec	2.2.0.0	COMMITTED	RSCT Security Msgs - U.S.
rsct.msg.en_US.core.sensorm			
rsct.msg.en_US.core.sr	2.2.0.0	COMMITTED	RSCT Registry Msgs - U.S.
rsct.msg.en_US.core.utils	2.2.0.0	COMMITTED	RSCT Utilities Msgs - U.S.
rsct.basic.rte	2.2.1.2	COMMITTED	RSCT Basic Function
rsct.compat.basic.rte	2.2.0.10	COMMITTED	RSCT Event Management Basic
rsct.core.rmc	2.2.1.0	COMMITTED	RSCT Resource Monitoring and
rsct.core.sec	2.2.1.2	COMMITTED	RSCT Security
rsct.core.sr	2.2.1.0	COMMITTED	RSCT Registry
rsct.core.utils	2.2.1.0	COMMITTED	RSCT Utilities

Select all the common filesets, the HACMP support one, but not the PSSP support one.

25.4 RSCT 2.3 (PROVIDED WITH AIX 5.2)

rsct.basic.hacmp	2.3.1.0	C	F	RSCT Basic Function (HACMP/ES
rsct.basic.rte	2.3.1.0	C	F	RSCT Basic Function
rsct.compat.basic.hacmp	2.3.1.0	C	F	RSCT Event Management Basic
rsct.compat.basic.rte	2.3.1.0	C	F	RSCT Event Management Basic
rsct.compat.clients.hacmp	2.3.1.0	C	F	RSCT Event Management Client
rsct.compat.clients.rte	2.3.1.0	C	F	RSCT Event Management Client
rsct.core.auditrm	2.3.1.0	C	F	RSCT Audit Log Resource
rsct.core.errm	2.3.1.0	C	F	RSCT Event Response Resource
rsct.core.fsrms	2.3.1.0	C	F	RSCT File System Resource
rsct.core.gui	2.3.1.0	C	F	RSCT Graphical User Interface
rsct.core.hostrm	2.3.1.0	C	F	RSCT Host Resource Manager
rsct.core.rmc	2.3.1.0	C	F	RSCT Resource Monitoring and
rsct.core.sec	2.3.1.0	C	F	RSCT Security
rsct.core.sensorm	2.3.1.0	C	F	RSCT Sensor Resource Manager
rsct.core.sr	2.3.1.0	C	F	RSCT Registry
rsct.core.utils	2.3.1.0	C	F	RSCT Utilities
rsct.msg.EN_US.core.auditrm				

rsct.msg.EN_US.core.errm	2.3.0.0	C	F	RSCT Event Response RM Msgs -
rsct.msg.EN_US.core.fsrn	2.3.0.0	C	F	RSCT File System RM Msgs -
rsct.msg.EN_US.core.gui	2.3.0.0	C	F	RSCT GUI Msgs - U.S. English
rsct.msg.EN_US.core.hostrn				
rsct.msg.EN_US.core.rmc	2.3.0.0	C	F	RSCT RMC Msgs - U.S. English
rsct.msg.EN_US.core.sec	2.3.0.0	C	F	RSCT Security Msgs - U.S.
rsct.msg.EN_US.core.sensorn				
rsct.msg.EN_US.core.sr	2.3.0.0	C	F	RSCT Registry Msgs - U.S.
rsct.msg.EN_US.core.utils	2.3.0.0	C	F	RSCT Utilities Msgs - U.S.
rsct.msg.en_US.basic.rte	2.3.0.0	C	F	RSCT Basic Msgs - U.S. English
rsct.msg.en_US.core.auditrm				
rsct.msg.en_US.core.errm	2.3.0.0	C	F	RSCT Event Response RM Msgs -
rsct.msg.en_US.core.fsrn	2.3.0.0	C	F	RSCT File System RM Msgs -
rsct.msg.en_US.core.gui	2.3.0.0	C	F	RSCT GUI Msgs - U.S. English
rsct.msg.en_US.core.gui.com				
rsct.msg.en_US.core.hostrn				
rsct.msg.en_US.core.rmc	2.3.0.0	C	F	RSCT RMC Msgs - U.S. English
rsct.msg.en_US.core.rmc.com				
rsct.msg.en_US.core.sec	2.3.0.0	C	F	RSCT Security Msgs - U.S.
rsct.msg.en_US.core.sensorn				
rsct.msg.en_US.core.sr	2.3.0.0	C	F	RSCT Registry Msgs - U.S.
rsct.msg.en_US.core.utils	2.3.0.0	C	F	RSCT Utilities Msgs - U.S.

25.5 GPFS 2.1

mmfs.base.cmds	3.5.0.6	C	F	GPFS File Manager Commands
mmfs.base.rte	3.5.0.9	C	F	GPFS File Manager
mmfs.gpfs.rte	2.1.0.9	C	F	GPFS File Manager
mmfs.gpfsdocs.data	3.5.0.0	C	F	GPFS Server Manpages and
mmfs.msg.en_US	3.5.0.0	C	F	GPFS Server Messages - U.S.