The installation of Oracle RAC 10*g* Release 2 on Asianux 2.0(x86-64)

by guotao zhong

August 2006

1. Introduction

One Oracle RAC is a cluster database with a shared cache architecture that provides highly scalable and available database solutions. The document introduced the installation of Oracle RAC 10g Release 2 on Asianux 2.0(x86-64). Before going into the detailed installation, take a look at the system architecture:

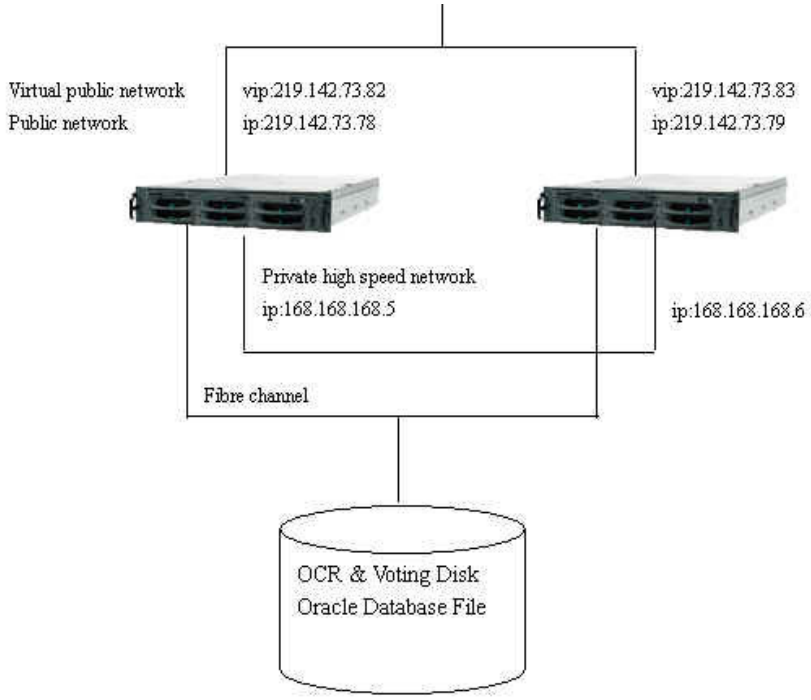

figure 1

This installation guide assumes two RAC nodes, which are connected with the shared disks provided by EMC SAN storage through the fibre channel, connect to a private high speed tcp/ip network and have the ability to access to the public network.

In this installation, we use two Dawning R220-A1 servers as nodes in the RAC and each server has the following hardware configuration:

- CPU: AMD Opteron(tm) Processor 244
- Number of CPU: 2
- Memory: 2G

Moreover, the following software will be installed:

. Asianux 2.0(x86-64) with patch sp1beta3

. Oracle Clusterware Release 2 (10.2.0.1.0) for Linux x86-64

. Oracle Database 10*g* Release 2 (10.2.0.1.0) Enterprise Edition for Linux x86-64

. OCFS2

. ASMLib 2.0

2. Install the Asianux2.0 (on both nodes)

Please according to the installation guide of Asianux2.0, installing the Asianux2.0 (x86_64) with patch sp1beta3 in both nodes.

To be mentioned, in the first node, when entering into the Network Configuration of Asianux2.0 installation, the screen should show that two network cards are detected. We edit the **eth0** (the network card for public network) and **eth1** (the network card for private network) as follows:

configure eth0
cancel "configure using DHCP"
select "Activate on boot"
IP Address: 219.142.73.78
Netmask: 255.255.255.0

configure eth1
cancel "configure using DHCP"
select "Activate on boot"
IP Address: 168.168.168.5
Netmask: 255.255.255.0

And set the hostname
cancel "automatically via DHCP"
select "manually: rac1"

In the last step of installation, we choose **ALL** packages to be installed.

In the second node, when entering into the Network Configuration of Asianux2.0 installation, the screen should also show that two network cards are detected. As the same as the first node, we edit the **eth0** (the network card for public network) and **eth1** (the network card for private network) as follows:

configure eth0
cancel "configure using DHCP"
select "Activate on boot"
IP Address: 219.142.73.79
Netmask: 255.255.255.0

configure eth1
cancel "configure using DHCP"
select "Activate on boot"
IP Address: 168.168.168.6
Netmask: 255.255.255.0

And set the hostname

cancel "automatically via DHCP"

select "manually: rac2"

In the last step of installation, we choose **ALL** packages to be installed.

Now we have finished the Asianux 2.0 installation in two nodes and get two RAC nodes: **rac1** and **rac2**.

After the installation, we can verify the kernel version:

#uname –r

Then query the following packages:

#rpm –q libaio binutils compat-db gcc gcc-c++ glibc glibc-common gnome-libs libstdc++ libstdc++-devel make

libaio-0.3.105-2
binutils-2.15.92.0.2-18.1AX
compat-db-4.1.25-9
gcc-3.4.5-2.1
gcc-c++-3.4.5-2.1
glibc-2.3.4-2.19.1AX
glibc-common-2.3.4-2.19.1AX
gnome-libs-1.4.1.2.90-44.1
libstdc++-3.4.5-2.1
libstdc++-devel-3.4.5-2.1
make-3.80-5

3. Configure the Asianux2.0 for Oracle (on both nodes)

3.1 configure the network

In this installation, we configure the same **/etc/hosts** on both nodes:

219.142.73.78 rac1
168.168.168.5 int-rac1
219.142.73.82 vip-rac1

219.142.73.79 rac2
168.168.168.6 int-rac2
219.142.73.83 vip-rac2

127.0.0.1 localhost.localdomain localhost

Note that the 219.142.73.82 and 219.142.73.83 are virtual IP addresses.

3.2 create the user "oracle" and group

Now we'll create user and group for installing Oracle 10*g* Release 2 software. The user ID and group ID must be the same on both nodes. So execute the following commands as root on both nodes:

#/usr/sbin/groupadd -g 512 oinstall
#/usr/sbin/groupadd -g 513 dba
#/usr/sbin/useradd -m -u 512 -g oinstall -G dba oracle
#passwd oracle

Then create the mount points for Oracle Software as root on both nodes:

#mkdir –p /dbrac/rac10g
#mkdir /dbrac/crs
#chown –R oracle:oinstall /dbrac
#chmod –R 775 /dbrac

And add the next environment parameters in the /home/oracle/.bash_profile as user "oracle":

export ORACLE_BASE=/dbrac
export ORACLE_HOME=$ORACLE_BASE/rac10g
export ORA_CRS_HOME=$ORACLE_BASE/crs
export PATH=$PATH:$ORACLE_HOME/bin: $ORA_CRS_HOME/bin
export LD_LIBRARY_PATH=/usr/lib64: /lib: /usr/lib: $ORACLE_HOME/lib:
$ORACLE_HOME/lib32: $ORACLE_HOME/oracm/lib: $ORA_CRS_HOME/lib:
$ORA_CRS_HOME/lib32
export CLASSPATH=$ORACLE_HOME/JRE:$ORACLE_HOME/jlib:
$ORACLE_HOME/rdbms/jlib:$ORACLE_HOME/network/jlib
export TNS_ADMIN=$ORACLE_HOME/network/admin

3.3 configure the kernel parameters

Configure the following Linux kernel parameters as root on both nodes

cat >> /etc/sysctl.conf << EOF
kernel.shmall = 2097152
kernel.shmmax = 2147483648
kernel.shmmni = 4096
kernel.sem = 250 32000 100 142
fs.file-max = 327680
net.ipv4.ip_local_port_range = 1024 65000
net.core.rmem_default = 262144
net.core.wmem_default = 262144

net.core.rmem_max = 262144
net.core.wmem_max = 262144
EOF

#/sbin/sysctl -p

3.4 configure the hangcheck-timer module

We can load the hangcheck-timer module using the following command on both nodes:

/sbin/modprobe hangcheck-timer hangcheck_tick=30 hangcheck_margin=180

We can load the hangcheck-timer module on each boot as follows:

cat >> /etc/rc.d/rc.local << EOF
/sbin/modprobe hangcheck-timer hangcheck_tick=30 hangcheck_margin=180
EOF

3.5 configure the rsh for both nodes

When installing Oracle RAC 10g Release 2 in a node, it will use the rsh to copy files to and execute programs on the other node in the RAC cluster.

Firstly, verify the rsh RPM on both nodes:

 # rpm -q rsh rsh-server
rsh-0.17-25.3
rsh-server-0.17-25.3

Secondly, to allow the "oracle" user to be trusted on both nodes, we must configure the /etc/hosts.equiv file on both nodes:

cat >> /etc/hosts.squiv << EOF
rac1 oracle
rac2 oracle
int-rac1 oracle
int-rac2 oracle
EOF
 # chmod 600 /etc/hosts.equiv
 # chown root.root /etc/hosts.equiv

Finally, enable the "rsh" and "rlogin" services, the "disable" attribute in the /etc/xinetd.d/rsh file must be set to "no" and the xinetd service must be restart on both nodes.

 # chkconfig rsh on

 # chkconfig rlogin on

 # service xinetd restart

4. Configure OCFS2

As a cluster filesystem, OCFS2 allows each node in the cluster to concurrently access the share disks. In this installation, we use OCFS2 to store the ocr files and voting files, which are required to be shared by the Oracle Clusterware software.

4.1 verify OCFS2

We can verify that the ocfs2 RPMs have been installed:

rpm –qa|grep ocfs2

4.2 configure OCFS2

Run ocfs2console as root:

 # ocfs2console

Select **"Cluster" → "Configure Nodes",** Click **"Add"** and enter the Name and IP Address of each node in the cluster:

rac1 219.142.73.78 7777

rac2 219.142.73.79 7777

Then select **"Cluster->Propagate Configuration"** for copying the OCFS2 configuration file to each node in the cluster.

After the configuration, we can check the **/etc/ocfs2/cluster.conf** on both nodes as follows:

node:

ip_port = 7777

ip_address = 219.142.73.78

number = 0

name = rac1

cluster = ocfs2

node:

ip_port = 7777

ip_address = 219.142.73.79

number = 1

name = rac2

```
        cluster = ocfs2

cluster:
        node_count = 2
        name = ocfs2
```

4.3 create and mount the OCFS2 filesystem

Firstly, we execute the following command on both nodes to load the OCFS2 modules and online cluster ocfs2:

```
#/etc/init.d/o2cb load
#/etc/init.d/o2cb online ocfs2
```

We use the shared disk, **/dev/sdb**, provided by EMC SAN storage to create a partition for storing the ocr files and voting files in the clusterware:

```
#fdisk /dev/sdb
```

After fdisk, we can get a partition:**/dev/sdb1**. Then create the OCFS2 filesystem on the partition. The following command should only be executed on one node in the RAC cluster:

```
# mkfs.ocfs2 -b 4K -C 32K -N 2 -L ocfs2disk /dev/sdb1
```

Then create the mount points for OCFS2 filesystem on both nodes:

```
#mkdir /u02
#chown –R oracle:oinstall /u02
#chmod –R 775 /u02
```

Finally we can run the following command to mount OCFS2 filesystem on both nodes:

```
# mount -t ocfs2 -o datavolume,nointr /dev/sdb1 /u02
```

We can also add the following line to the **/etc/fstab** of both nodes for mounting the OCFS2 filesystem on each boot:

```
/dev/sdb1       /u02       ocfs2     _netdev,datavolume,nointr       0 0
```

5. Configure ASMLib 2.0

In this installation, we build a RAC database using ASM for database files storage. On Linux platforms, ASM can use raw devices or devices managed via the ASMLib interface.

Firstly, we can make sure that the ASMLib RPMs have been installed on both nodes:

rpm –qa|grep oracleasm

Then configure and load the ASM kernel module on both nodes as root:

 # /etc/init.d/oracleasm configure
Configuring the Oracle ASM library driver.

This will configure the on-boot properties of the Oracle ASM library
driver. The following questions will determine whether the driver is
loaded on boot and what permissions it will have. The current values
will be shown in brackets ('[]'). Hitting <ENTER> without typing an
answer will keep that current value. Ctrl-C will abort.

Default user to own the driver interface []: oracle
Default group to own the driver interface []: oinstall
Start Oracle ASM library driver on boot (y/n) [n]: y
Fix permissions of Oracle ASM disks on boot (y/n) [y]: y
Writing Oracle ASM library driver configuration: [OK]
Creating /dev/oracleasm mount point: [OK]
Loading module "oracleasm": [OK]
Mounting ASMlib driver filesystem: [OK]
Scanning system for ASM disks: [OK]

Secondly, we create two partitions: **/dev/sdc1** and **/dev/sdc2** on a shared disk **/dev/sdc** provided by the EMC SAN storage for storing Oracle 10g RAC database:

#fdisk /dev/sdc

Then creating the ASM disks using the following commands. The following command should only be executed on one node in the RAC cluster:

 # /etc/init.d/oracleasm createdisk VOL1 /dev/sdc1
 # /etc/init.d/oracleasm createdisk VOL2 /dev/sdc2

When completed, we can verify that all ASM disks have been created successfully as follows:

#/etc/init.d/oracleasm listdisks
VOL1
VOL2

On the other node, run the following command as root to scan configured ASM disks:

#/etc/init.d/oracleasm scandisks

6. Install the Oracle 10g Clusterware software

We install the Oracle 10g Clusterware Software as follows. The following installation should only be executed as “oracle” user on one node in the RAC cluster.

$ /home/oracle/clusterware/runInstaller –ignoreSysPrereqs

<table>
<tr><td>Weclome</td><td>next</td></tr>
<tr><td>Specify Inventory directory and credentials</td><td>Inventory directory: /dbrac/oraInventory
Operating System group name: oinstall</td></tr>
<tr><td>Specify Home Details</td><td>Name:OraCrs10g_home
Path:/dbrac/crs</td></tr>
<tr><td>Product-Specific Prerequisite Checks</td><td>If all checks passed with no problems, click next</td></tr>
<tr><td>Specify Cluster Configuration</td><td>Cluster Name: crs
<table>
<tr><th>Public Node Name</th><th>Private Node Name</th><th>Virtual Node Name</th></tr>
<tr><td>rac1</td><td>int-rac1</td><td>vip-rac1</td></tr>
<tr><td>rac2</td><td>int-rac2</td><td>vip-rac2</td></tr>
</table></td></tr>
<tr><td>Specify Network Interface Usage</td><td>
<table>
<tr><th>Interface Name</th><th>Subnet</th><th>Interface Type</th></tr>
<tr><td>eth0</td><td>219.142.73.0</td><td>Public</td></tr>
<tr><td>eth1</td><td>168.168.168.0</td><td>Private</td></tr>
</table></td></tr>
<tr><td>Specify OCR Location</td><td>OCR Configuration: "Normal Redundancy"
OCR Location: /u02/ocr/ocr
OCR Mirror Location: /u02/ocr/ocr2</td></tr>
<tr><td>Specify Voting Disk Location</td><td>Voting Disk Configuration: Normal Redundancy
Voting Disk Location: /u02/voting/voting_disk
Additional Voting Disk 1 Location: /u02/voting/voting_disk2
Additional Voting Disk 2 Location: /u02/voting/voting_disk3</td></tr>
<tr><td>Summary</td><td>If OK, click Install</td></tr>
<tr><td>Execute Configuration scripts</td><td>After finishing the installation, executing the orainstRoot.sh and root.sh script as root according to the prompt in the screen</td></tr>
<tr><td>Configuration Assistants</td><td>If all configurations are succeeded, click next</td></tr>
<tr><td>End of Installation</td><td>exit</td></tr>
</table>

7. Install the Oracle Database 10g software

We install the Oracle Database 10g Software as follows. The following installation should only be executed as “oracle” user on one node in the RAC cluster.

$ /home/oracle/database/runInstaller –ignoreSysPrereqs

Weclome	**next**
Select Installation Type	**Enterprise Edition**
Specify Home Details	Name: **OraDb10g_home1** Path:**/dbrac/rac10g**
Specify Hardware Cluster Installation Mode	**Cluster Installation** Nodename: **rac1 rac2**
Product-Specific Prerequisite Checks	If all checks passed with no problems, click **next**
Select Database Configuration	**Install database software only**
Summary	If OK, click **Install**
Execute Configuration scripts	After finishing the installation, executing the **root.sh** script as root according to the prompt in the screen
End of Installation	**exit**

8. Configure the TNS listener

We configure the TNS listener as follows. The following configuration should only be executed as "oracle" user on one node in the RAC cluster.

$ netca

Select the the type of Oracle Net Services Configuration	**Cluster Configuration**
Select the nodes to configure	Select all nodes: **rac1 rac2**
Choose the configuration you would like to do	**Listener configuration**
What do you want to do	**Add**
Listener name	**LISTENER**
Selected Protocols	**TCP**
Port	**Use the standard port number of 1521**
Would you like to configure another listener?	**No**
Listener configuration complete	**Next**
Choose the configuration you would like to do	**Naming Methods configuration**
Selected Naming Methods	**Local Naming**
Naming Methods configuration complete	**next**
Choose the configuration you would like to do	**Finish**

The Oracle TNS listener process should be running on both nodes.

9. Create the Oracle 10g RAC database

We create the oracle10g RAC database as follows. The following steps should only be executed as “oracle” user on one node in the RAC cluster.

$ dbca

Weclome	**Oracle Real Application Clusters database**
Operations	**Create a Database**
Node Selection	**Select All: rac1, rac2**
Database Templates	**Custom Database**
Database Identification	Global Database Name: **racdb** SID Prefix: **racdb**
Management Options	**Default options**: Configure the Database with Enterprise Manager Use Database Control for Database Management.
Database Credentials	Use the Same Password for All Accounts: **input the password twice**
Storage Options	**Automatic Storage Management (ASM)**
Create ASM Instance	**Input SYS password twice** Create server parameter file(SPFILE): **/u02/oradata/racdb/spfile+ASM.ora** The parameter file should be on a shared disk
ASM Disk Groups	Click **“Create New”** There will be **"Create Disk Group"** window with the two volumes we configured earlier using ASMLib. 1.input the value **“RACDBDISK”** for "Disk Group Name". 2.select **the two ASM volumes** in the "Select Member Disks" window. 3.set the "Redundancy" to **"Normal"**. Click the “OK” and create ASM Disk Group Return the “ASM Disk Groups” windows and select the configured disk group **“RACDBDISK”**. Click **“next”**
Database File Locations	Select **“Use Oracle-Managed Files”** Database Area: **+RACDBDISK**
Recovery Configuration	cancel the **"Specify Flash Recovery Area"**, nothing selected
Database Content	**Default**
Database Services	**Default**
Initialization Parameters	**Default**
Database Storage	**Default**
Creation Options	**Finish**

The Oracle 10g RAC database should be up and running. And set the environment parameter **“ORACLE_SID”** as **“oracle”** user in both nodes for accessing the **racdb** database.

In rac1, **$export ORACLE_SID=racdb1**

In rac2, **$export ORACLE_SID=racdb2**

Reference

1. Oracle Database Documentation Library 10g Release 2 (10.2): Oracle® Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide 10g Release 2 (10.2) for Linux (Part Number B14203-07)

*2. Build Your Own Oracle RAC 10*g *Release 2 Cluster on Linux and FireWire*
http://www.oracle.com/technology/pub/articles/hunter_rac10gr2.html