Author – *A.Kishore*
http:/www.appsdba.info

Configure Oracle Cluster File System (OCFS2)

Create the OCFS2 configuration file. As the root user on rac1, execute - on VNC session

```
ocfs2console
```

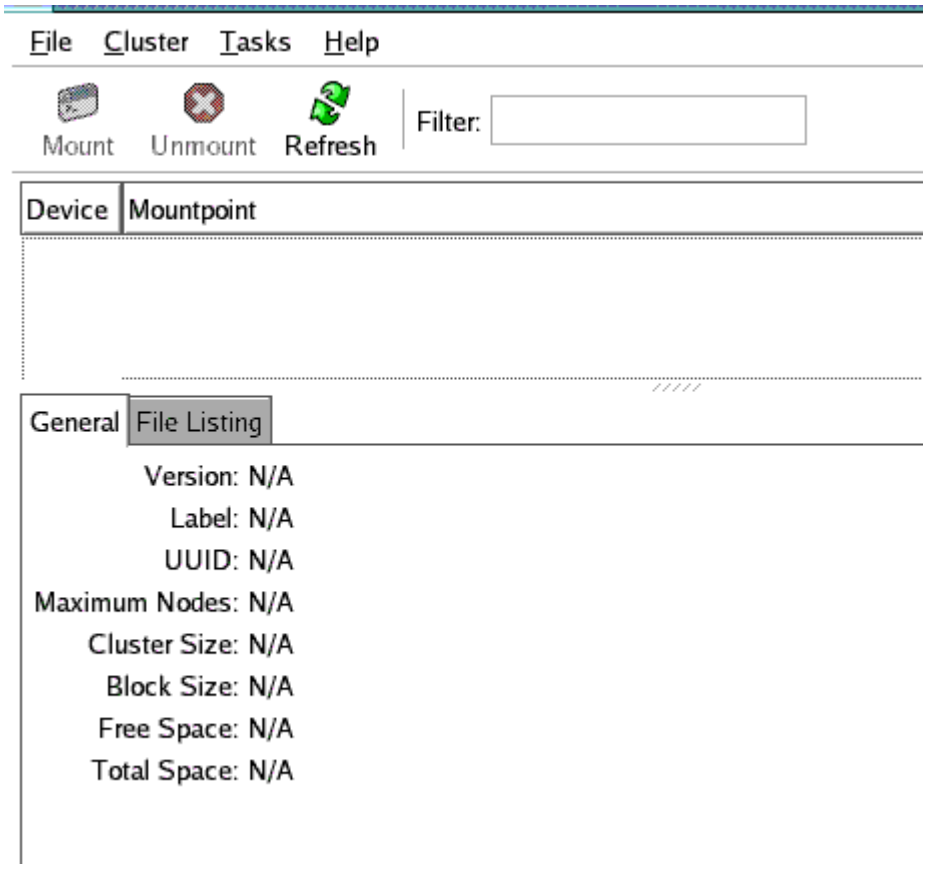


```
Cluster - Configure Nodes
```

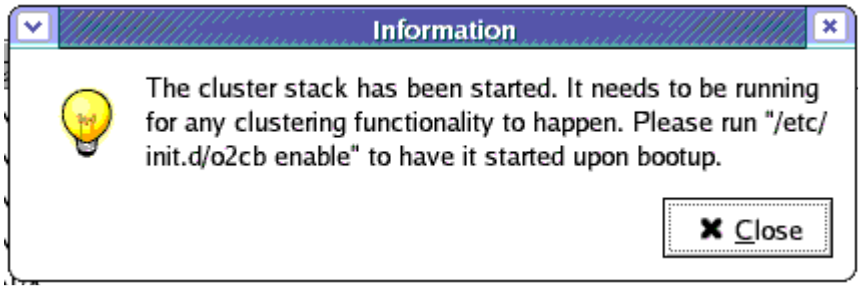


```
/etc/init.d/o2cb enable

Cluster - Configure Nodes
```

Author – *A.Kishore*
http:/www.appsdba.info

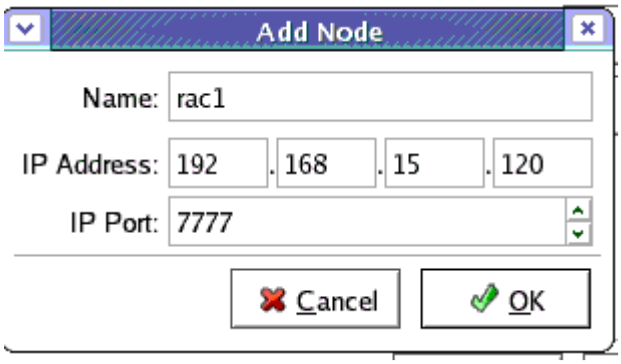

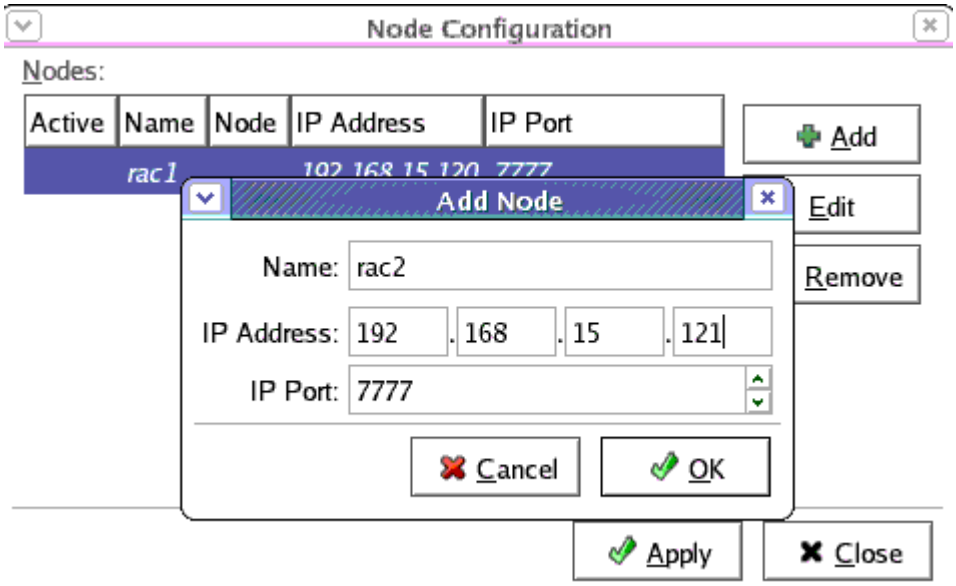


```
cat /etc/ocfs2/cluster.config
```

Author – *A.Kishore*
http:/www.appsdba.info

```
[root@rac1 bin]# cat /etc/ocfs2/cluster.conf
node:
        ip_port = 7777
        ip_address = 192.168.15.120
        number = 0
        name = rac1
        cluster = ocfs2

node:
        ip_port = 7777
        ip_address = 192.168.15.121
        number = 1
        name = rac2
        cluster = ocfs2

cluster:
        node_count = 2
        name = ocfs2
```

While configuring the nodes for OCFS2 using `ocfs2console`, it is possible to run into the error:

```
o2cb_ctl: Unable to access cluster service while creating node
```

This error does not show up when you startup `ocfs2console` for the first time. This message comes up when there is a problem with the cluster configuration or if you do not *save* the cluster configuration initially while setting it up using `ocfs2console`. This is a bug!

The work-around is to exit from the `ocfs2console`, unload the o2cb module and remove the ocfs2 cluster configuration file `/etc/ocfs2/cluster.conf`. I also like to remove the `/config` directory. After removing the ocfs2 cluster configuration file, restart the `ocfs2console` program.

For example:

```
# /etc/init.d/o2cb offline ocfs2
# /etc/init.d/o2cb unload
Unmounting ocfs2_dlmfs filesystem: OK
Unloading module "ocfs2_dlmfs": OK
Unmounting configfs filesystem: OK
Unloading module "configfs": OK

# rm -f /etc/ocfs2/cluster.conf
# rm -rf /config

# ocfs2console &
```

This time, it will add the nodes!

Author – *A.Kishore*
http:/www.appsdba.info

Propagate the changes to rac2

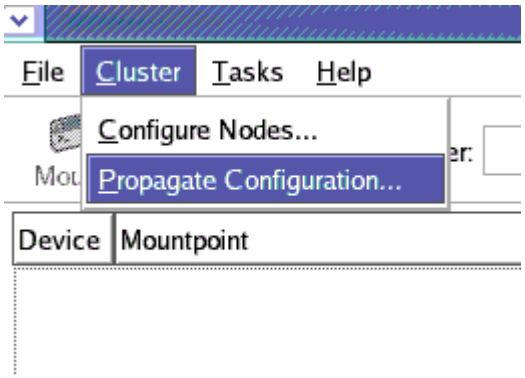

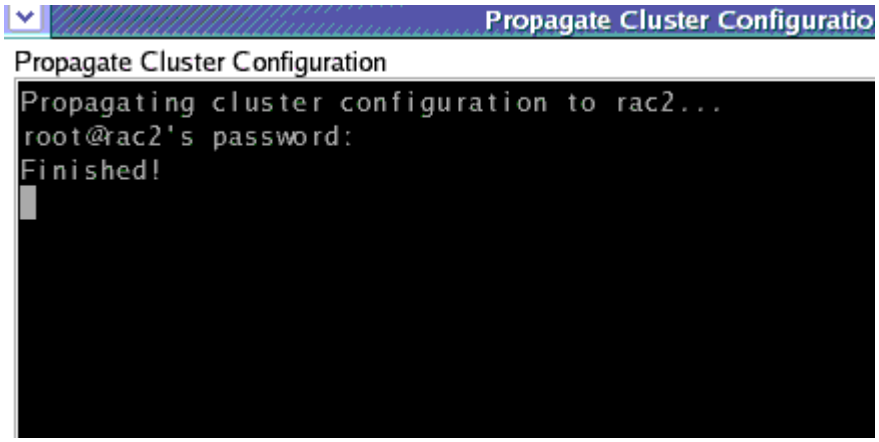

Configure the O2CB driver - both nodes as root user

/etc/init.d/o2cb unload

/etc/init.d/o2cb configure

/etc/init.d/o2cb status

Author – *A.Kishore*
http:/www.appsdba.info

```
[root@rac1 bin]# /etc/init.d/o2cb unload
Stopping O2CB cluster ocfs2: OK
Unmounting ocfs2_dlmfs filesystem: OK
Unloading module "ocfs2_dlmfs": OK
Unmounting configfs filesystem: OK
Unloading module "configfs": OK
[root@rac1 bin]# /etc/init.d/o2cb configure
Configuring the O2CB driver.

This will configure the on-boot properties of the O2CB driver.
The following questions will determine whether the driver is loaded on
boot.  The current values will be shown in brackets ('[]').  Hitting
<ENTER> without typing an answer will keep that current value.  Ctrl-C
will abort.

Load O2CB driver on boot (y/n) [n]: y
Cluster to start on boot (Enter "none" to clear) [ocfs2]:
Specify heartbeat dead threshold (>=7) [7]: 61
Specify network idle timeout in ms (>=5000) [10000]:
Specify network keepalive delay in ms (>=1000) [5000]:
Specify network reconnect delay in ms (>=2000) [2000]:
Writing O2CB configuration: OK
Loading module "configfs": OK
Mounting configfs filesystem at /config: OK
Loading module "ocfs2_nodemanager": OK
Loading module "ocfs2_dlm": OK
Loading module "ocfs2_dlmfs": OK
Mounting ocfs2_dlmfs filesystem at /dlm: OK
Starting O2CB cluster ocfs2: OK
[root@rac1 bin]# /etc/init.d/o2cb status
Module "configfs": Loaded
Filesystem "configfs": Mounted
Module "ocfs2_nodemanager": Loaded
Module "ocfs2_dlm": Loaded
Module "ocfs2_dlmfs": Loaded
Filesystem "ocfs2_dlmfs": Mounted
Checking O2CB cluster ocfs2: Online
  Heartbeat dead threshold: 61
  Network idle timeout: 10000
  Network keepalive delay: 5000
  Network reconnect delay: 2000
Checking O2CB heartbeat: Not active
```

Author – *A.Kishore*
http:/www.appsdba.info

```
format - only on one node
osfs2console

Tasks - Format
```

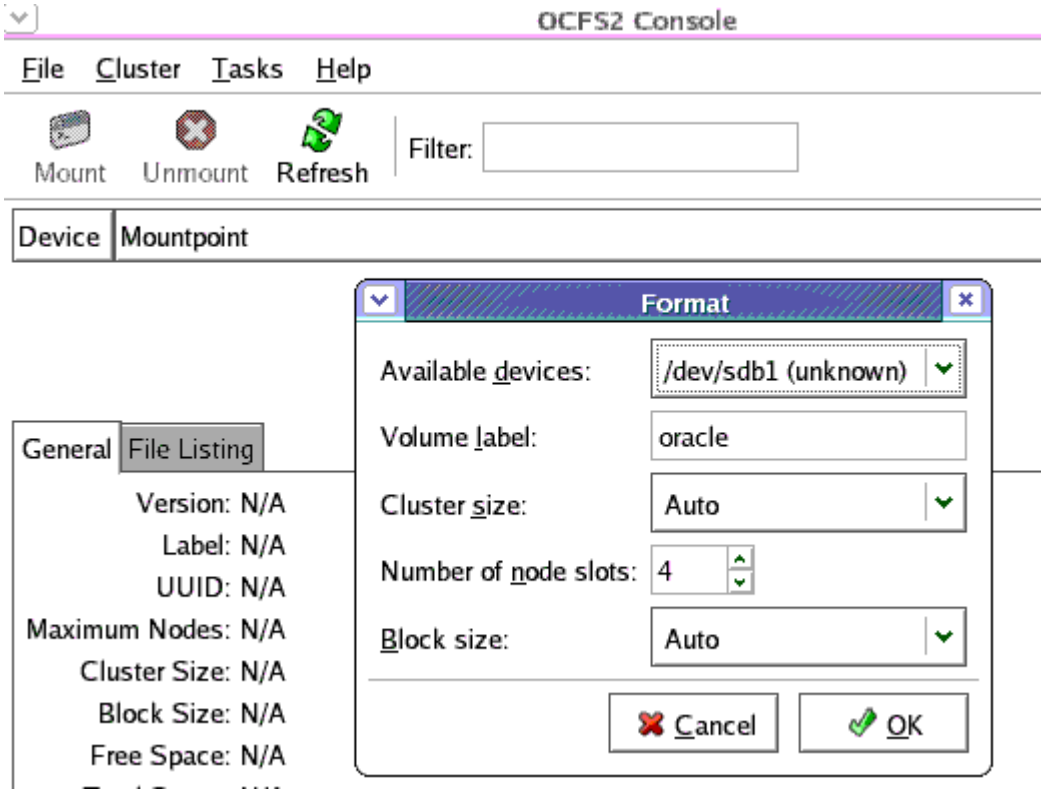


```
mount -t ocfs2 -o datavolume,nointr /dev/sdb1 /ocfs
```

```
[root@rac1 bin]# mount -t ocfs2 -o datavolume,nointr /dev/sdb1 /ocfs
[root@rac1 bin]# df -h
Filesystem            Size  Used Avail Use% Mounted on
/dev/sda1             6.9G  2.1G  4.5G  33% /
none                  345M     0  345M   0% /dev/shm
/dev/sda3              11G   60M   11G   1% /u01
/dev/hdc              637M  637M     0 100% /media/cdrom
192.168.15.112:/software
                       39G   37G     0 100% /software
/dev/sdb1             102M   19M   84M  19% /ocfs
```

```
Mount on rac2
```

```
[root@rac2 ~]# mount -t ocfs2 -o datavolume,nointr /dev/sdb1 /ocfs
```

Author – *A.Kishore*
http:/www.appsdba.info

To mount the file system on boot, add the following line in /etc/fstab on both nodes.

```
vi /etc/fstab  - on both nodes
/dev/sdb1 /ocfs ocfs2 _netdev,datavolume,nointr 0 0
```

```
# This file is edited by fstab-sync - see 'man fstab-sync' for details
LABEL=/                 /                       ext3    defaults        1 1
none                    /dev/pts                devpts  gid=5,mode=620  0 0
none                    /dev/shm                tmpfs   defaults        0 0
none                    /proc                   proc    defaults        0 0
none                    /sys                    sysfs   defaults        0 0
LABEL=/u01              /u01                    ext3    defaults        1 2
LABEL=SWAP-sda2         swap                    swap    defaults        0 0
/dev/hdc                /media/cdrom            auto    pamconsole,fscontext=system_u:objec
t_r:removable_t,exec,noauto,managed 0 0
/dev/fd0                /media/floppy           auto    pamconsole,fscontext=system_u:objec
t_r:removable_t,exec,noauto,managed 0 0
/dev/sdb1 /ocfs ocfs2 _netdev,datavolume,nointr 0 0
~
~
```

```
Create Oracle Clusterware directory - on both nodes

mkdir -p /ocfs/clusterware
chown -R oracle:dba /ocfs
```