

Technical report of the setup procedure of a Ubuntu Server computer cluster

Alexandre Manhães Savio

September 24th, 2010

1 Master Node - Giclus1

1.1 Operating system and network configuration

- Install Ubuntu Server 10.04.1 amd64
- Partitions
 - /dev/sda1: 20GB on /
 - /dev/sda2: 20GB on /usr/local
 - /dev/sda3: 130GB on /home
 - /dev/sda5: 280GB on /opt
 - /dev/sda6: 20GB of swap area
- Add gic group:
 - sudo addgroup gic
 - sudo addgroup alexandre gic
- Add more users:
 - sudo adduser <user_name> --ingroup gic --disabled-password
 - ssh-keygen -b 4096 -t rsa -C user_name
 - chown -R user_name:root /home/user_name/.ssh
 - chmod 700 /home/user_name/.ssh
 - chmod 400 /home/user_name/.ssh/authorized_keys
- Edit /etc/network/interfaces

```
auto loiface lo inet loopbackauto eth1iface eth1 inet static
address 192.168.1.81netmask 255.255.255.0gateway 192.168.1.1
auto eth0 iface eth0 inet dhcp
```

- Edit `/etc/resolv.conf`

```
nameserver 10.20.13.6
nameserver 10.10.13.6
nameserver 10.30.13.6
```
- Edit `/etc/hosts`

```
127.0.0.1 localhost
192.168.1.81 giclus1
192.168.1.82 giclus2
192.168.1.83 giclus3
192.168.1.84 giclus4
```
- Edit `/etc/hosts.allow`

```
portmap ypserv ypbind sge_qmaster sge_execd : \
192.168.1.81 192.168.1.82 192.168.1.83 192.168.1.84
```
- `sudo /etc/init.d/networking restart`
- Install packages:
 - Add "partner" repository, editing `/etc/apt/sources.list`
 - `sudo apt-get update`
 - `sudo apt-get install ssh molly-guard openssh-blacklist openssh-blacklist-extra ssh-askpass binutils unzip sun-java6-jre`
 - Edit `/etc/ssh/sshd_conf`: disable root access and password auth

1.2 NIS Server

- `sudo apt-get install portmap nis`
- NIS domain name: `giclus`
- For more details: <https://help.ubuntu.com/community/SettingUpNISHowTo>
- Edit `/etc/default/portmap` and comment out the `ARGS="-i 127.0.0.1"` line
- Edit `/etc/default/nis` and set the `NISSERVER` line to `NISSERVER=master`
- Edit `/etc/yp.conf`: domain `giclus server giclus1`
- Edit `/etc/ypserv.securenets`

```
host 192.168.1.81
host 192.168.1.82
host 192.168.1.83
host 192.168.1.84
```

- Build the DB for the first time, run: `sudo /usr/lib/yp/ypinit -m`
- Read the web page (<https://help.ubuntu.com/community/SettingUpNISHowTo>) for more information on security and the client config (see next section)
- Restart:
 - `sudo /etc/init.d/portmap restart`
 - `sudo /etc/init.d/nis restart`

1.3 NFS Kernel Server

- <https://help.ubuntu.com/community/SettingUpNFSHowTo>
- `sudo apt-get install nfs-watch nfs-kernel-server`
- `sudo mkdir /sge`
- Edit `/etc/exports` and add the shares:


```

/home giclus1(rw, sync, no_subtree_check) giclus2(rw, sync, no_subtree_check)
giclus3(rw, sync, no_subtree_check) giclus4(rw, sync, no_subtree_check)

/usr/local giclus1(rw, sync, no_subtree_check) giclus2(rw, sync, no_subtree_check)
giclus3(rw, sync, no_subtree_check) giclus4(rw, sync, no_subtree_check)

/opt giclus1(rw, sync, no_subtree_check) giclus2(rw, sync, no_subtree_check)
giclus3(rw, sync, no_subtree_check) giclus4(rw, sync, no_subtree_check)

/sge giclus1(rw, sync, no_subtree_check) giclus2(rw, sync, no_subtree_check)
giclus3(rw, sync, no_subtree_check) giclus4(rw, sync, no_subtree_check)

```
- `sudo exportfs -ra`

1.4 MSMTTP:

- Create a Gmail account for monitoring. I do this because I don't want my gmail password floating around in plaintext on various machines.
- Install the `ca-certificates` package


```

sudo aptitude install ca-certificates
sudo update-ca-certificates

```
- `sudo apt-get install msmtpp`
- Edit `/etc/msmtprc`

```
account gmail
host smtp.gmail.com
from giclus1@gmail.com
auth on
tls on
tls_trust_file /etc/ssl/certs/ca-certificates.crt
user giclus1@gmail.com
password *****
port 587
```

```
account default : gmail
```

- Create a sendmail simlink:

```
- sudo ln -s /usr/bin/msmtp /usr/sbin/sendmail
```

- -Run a test

```
- echo "This is a an awesome test email" | msmtp youremail@domain.com
```

- - If you want mdadm to mail you when something goes wrong

```
- Edit /etc/mdadm/mdadm.conf: MAILADDR giclus1@gmail.com
```

- And then run a mdadm test by running

```
- sudo mdadm --monitor --scan --test --oneshot
```

1.5 Share internet connection with the others in the cluster:

1.5.1 UFW Version

- Enable UFW

```
sudo ufw enable
sudo ufw allow 22/tcp
sudo ufw allow 22/udp
sudo ufw allow in on eth1
```

- Edit file /etc/ufw/before.rules:

```
# nat Table rules
*nat :POSTROUTING ACCEPT [0:0]
# .
-A POSTROUTING -s 192.168.1.0/24 -o eth1 -j MASQUERADE
COMMIT
```

- Edit `/etc/default/ufw`
 - Change `DEFAULT_FORWARD_POLICY` to “ACCEPT”
 - Uncomment:
 - * `net/ipv4/ip_forward=1`
 - * `net/ipv6/conf/default/forwarding=1`
- Restart ufw:
 - `sudo ufw disable`
 - `sudo ufw enable`

1.5.2 IPTABLES Version

- `sudo iptables -A FORWARD -i eth0 -o eth1 -s 192.168.1.0/24 -m conntrack --ctstate NEW -j ACCEPT`
- `sudo iptables -A FORWARD -m conntrack --ctstate ESTABLISHED,RELATED -j ACCEPT`
- `sudo iptables -A POSTROUTING -t nat -j MASQUERADE`
- `sudo iptables-save | sudo tee /etc/iptables.sav`
- Add to `/etc/rc.local`
 - `iptables-restore < /etc/iptables.sav`
- Add to `/etc/sysctl.conf`
 - `net.ipv4.conf.default.forwarding=1`
 - `net.ipv4.conf.all.forwarding=1`
- `sudo sh -c "echo 1 > /proc/sys/net/ipv4/ip_forward"`

1.6 Sun Grid Engine Master

- <http://biowiki.org/HowToAdministerSunGridEngine>
- <https://www.fmrib.ox.ac.uk/phpwiki/index.php/FslSge>
- `sudo apt-get install libmotif3 libxpm4`
- Download SGE from: <http://www.oracle.com>
- Install SGE (check this):

```

mkdir /opt/soft/sge
mv ge-6.1u6-* ../../soft/sge
cd ../../soft/sge
tar xvzf ge-6.1u6-common.tar.gz
tar xvzf ge-6.1u6-arco.tar.gz
tar xvzf ge-6.1u6-bin-lx24-amd64.tar.gz
cd ..\ sudo cp -rdvfa sge /
cd /sge
scp -rdv giclus1:/sge/* .
sudo ./inst_sge -m -x

```

- Now go through the interactive install process
- Add to /etc/bash.bashrc

```

#SGE settings export
SGE_ROOT=/sge
export SGE_CELL=default
if [ -e $SGE_ROOT/$SGE_CELL ]
then
. $SGE_ROOT/$SGE_CELL/common/settings.sh
fi

```

- **ERROR** “[: 359: 11: unexpected operator”

– On Ubuntu 10.04 LTS libc version detection fails in util/arch. The reason is that now (around line 244) `strings libc.so.6` returns GNU C Library (Ubuntu EGLIBC 2.11.1-0ubuntu7) stable release version 2.11.1, by Roland McGrath et al. where the version number appears twice. The subsequent tests get a string like "11\n11" instead of just "11" and the shell complains that the syntax of the if conditions is wrong. I fixed it by adding `uniq` to this line to the file `/sge/util/arch`:

```

libc_version='echo $libc_string | tr ' , ' '\n' | grep "2\." | cut -f

```

- **ERROR** `sgemaster` and `sgeexecd` won't start on boot

```

cd /etc/init.d/
sudo update-rc.d sgeexecd.giclus defaults
sudo update-rc.d sgemaster.giclus defaults

```

- Add user group for execution `sgeusers`:

- `sudo addgroup sgeusers`
- Add the users who are going to use SGE to this group
- Change owner group of `$SGE_ROOT` to `sgeusers`
- Configure `@allhosts` SGE execution group:
 - Show group: `qconf -shgrp @allhosts`
 - Edit group: `qconf -shgrp @allhosts`
 - `group_name @allhosts`
 - `hostlist giclus1 giclus2 giclus3 giclus4`

2 The other nodes: GICLUS{2-3-4}

2.1 Operating system and network configuration

- Install Ubuntu Server 10.04.1 amd64
- Partitions
 - `/dev/sda1`: 30GB on `/`
 - `/dev/sda2`: 450GB on `/local_opt`
 - `/dev/sda3`: 20GB of swap area
- Add `gic` group:
 - `sudo addgroup gic`
 - `sudo addgroup alexandre gic`
- Add more users:
 - `sudo adduser <user_name> --ingroup gic --disabled-password`
 - `ssh-keygen -b 4096 -t rsa -C user_name`
 - `chown -R user_name:root /home/user_name/.ssh`
 - `chmod 700 /home/user_name/.ssh`
 - `chmod 400 /home/user_name/.ssh/authorized_keys`
- Edit `/etc/network/interfaces`

```

auto loiface lo inet loopback
auto eth1iface eth1 inet static
address 192.168.1.8{2,3,4} netmask 255.255.255.0 gateway
192.168.1.81

```
- Edit `/etc/resolv.conf`

```
nameserver 10.20.13.6
nameserver 10.10.13.6
nameserver 10.30.13.6
```

- Edit `/etc/hosts`

```
127.0.0.1 localhost

192.168.1.81 giclus1
192.168.1.82 giclus2
192.168.1.83 giclus3
192.168.1.84 giclus4
```

- `sudo /etc/init.d/networking restart`
- Install packages:
 - Add "partner" repository, editing `/etc/apt/sources.list`
 - `sudo apt-get update`
 - `sudo apt-get upgrade`
 - `sudo apt-get install ssh molly-guard openssh-blacklist openssh-blacklist-extra ssh-askpass binutils unzip sun-java6-jre`
 - Edit `/etc/ssh/sshd_conf`: disable root access and password authentication

2.2 NIS Client

- <https://help.ubuntu.com/community/SettingUpNISHowTo>
- `sudo apt-get install nis`
- NIS domain name: `giclus`
- Edit `/etc/hosts.allow`: `portmap : 192.168.1.81`
- Add to `/etc/passwd` (+6x':')
- `+::::::`
- Add to `/etc/group` (+3x':')
- `+:::`
- Add to `/etc/shadow` (+8x':')
- `+:::::::`
- Edit `/etc/yp.conf` and add the line: `ypserver giclus1`

- `/etc/init.d/nis restart`
- `/etc/init.d/ssh restart`
- or `sudo reboot`

2.3 NFS Client

- `sudo apt-get install nfs-common`
- `sudo mkdir /sge`
- Add to `/etc/fstab`

```
#NFS Cluster mount
giclus1:/home /home nfs rsize=8192,wsiz=8192,timeo=14,intr,rw
giclus1:/opt /opt nfs rsize=8192,wsiz=8192,timeo=14,intr,rw
giclus1:/usr/local /usr/local nfs rsize=8192,wsiz=8192,timeo=14,intr,rw
giclus1:/sge /sge nfs rsize=8192,wsiz=8192,timeo=14,intr,rw
```

2.4 Sun Grid Engine Exec Daemon

- <http://biowiki.org/HowToAdministerSunGridEngine>
- <https://www.fmrib.ox.ac.uk/phpwiki/index.php/FslSge>
- `sudo apt-get install libmotif3 libxpm4`
- `sudo mount /sge`
- Download SGE from: <http://www.oracle.com>
- Install SGE (check this):

```
mkdir /opt/soft/sge
mv ge-6.1u6-* ../../soft/sge
cd ../../soft/sge
tar xvzf ge-6.1u6-common.tar.gz
tar xvzf ge-6.1u6-arco.tar.gz
tar xvzf ge-6.1u6-bin-lx24-amd64.tar.gz
cd ..\ sudo cp -rdvfa sge /
cd /sge
scp -rdv giclus1:/sge/* .
sudo ./install_execd
```

- Now go through the interactive install process
- Add to `/etc/bash.bashrc`

```
#SGE settings export
SGE_ROOT=/sge
export SGE_CELL=default
if [ -e $SGE_ROOT/$SGE_CELL ]
then
    . $SGE_ROOT/$SGE_CELL/common/settings.sh
fi
```

- Diagnosis commands:

```
- ps -A | grep sge
- qping giclus1 6444 qmaster 1
```

- **ERROR** “[: 359: 11: unexpected operator”

– On Ubuntu 10.04 LTS libc version detection fails in util/arch. The reason is that now (around line 244) strings libc.so.6 returns GNU C Library (Ubuntu EGLIBC 2.11.1-0ubuntu7) stable release version 2.11.1, by Roland McGrath et al. where the version number appears twice. The subsequent tests get a string like "11\n11" instead of just "11" and the shell complains that the syntax of the if conditions is wrong. I fixed it by adding uniq to this line to the file /sge/util/arch:

```
libc_version='echo $libc_string | tr ' , ' '\n' | grep "2\." | cut -f
```

- **ERROR** sgemaster and sgeexecd won't start on boot

```
cd /etc/init.d/
sudo update-rc.d sgeexecd.giclus defaults
```

3 Execute in all nodes

3.1 Install NeuroDebian Repository (<http://neuro.debian.net/>)

- Installation:

```
wget -c http://neuro.debian.net/_static/neurodebian.lucid.de.sources.list
wget -c http://neuro.debian.net/_static/neuro.debian.net.asc
sudo apt-key add neuro.debian.net.asc
sudo cp neurodebian.lucid.de.sources.list /etc/apt/sources.list.d
sudo apt-get update
sudo apt-get upgrade
sudo apt-get install fsl fsl-atlases fsl-first-data nifti-bin
```

- Add to `/etc/bash.bashrc`:

- `. /etc/fsl/fsl.sh`

3.2 Other configuration details

- Local temporary work directory

- Add to `/etc/environment`

- * `LOCAL_TEMP="/local"`

- * (for `giclus1`) this has been set to `"/opt/temp"`

- `sudo chown -R alexandre:gic /local`

- `sudo chmod -R 770 /local`