

NeuroGrid *Power G5 Cluster*

User's Manual

The aim of NeuroGrid project is to build a loosely-coupled, multiple-purpose cluster to significantly improve the performance and maximize the utilization of computing facilities for Neuro Imaging Group. The G5 cluster consists of twelve *Mac Dual CPU Power G5* computers as compute nodes and one *Mac Dual CPU XServe* computer as control node. The cluster nodes are interconnected with Gigabit network.

To manage the cluster nodes as a whole, Sun Grid Engine is installed as the cluster management and job scheduling software. FSL and other image analysis tools are deployed on all of the compute nodes to do image processing jobs. To simplify the usage of Sun Grid Engine, a group of wrapper scripts for FSL command line tools are developed and deployed on the cluster nodes. Further more, FSL GUI is modified to add an option to submit jobs to Sun Grid Engine instead of running locally so that users who are not familiar with command line tools can easily compose a FSL job and submit it the cluster.

1. Specifications:

1.1. Hardware:

Table 1 Hardware specifications of the cluster

Number of nodes/computers	13
Number of compute nodes	12
Number of control node(s)	1
Number of processors per node	2
Processor architecture	Power PC
Processor frequency	1.8Ghz
RAM per node	2.5GB
Local disk space per node	150 GB

1.2. Software

Table 2 Software specifications

Operating System	Mac OS 10.3.6
Cluster Management Software	Sun Grid Engine 6.0 U1
Parallel Environment/Library	Lam MPI 7.1
Applications	FSL 3.1 FSL 3.2 Merlin

1.3 Details of cluster nodes

Table 3 IP addresses and physical locations of cluster nodes

Node Name	IP Address	Location	Role
Cluster	192.168.20.201	Server Room, Lvl 3	Control & Compute node
mac0a	192.168.20.204	IT Support Room, Lvl 3	Compute node
mac0b	192.168.20.205	IT Support Room, Lvl 3	Compute node
mac01	192.168.35.101	Mezza Lab, Groud Floor	Compute node
mac02	192.168.35.102	Mezza Lab, Groud Floor	Compute node
mac03	192.168.35.103	Mezza Lab, Groud Floor	Compute node
mac04	192.168.35.104	Mezza Lab, Groud Floor	Compute node
mac05	192.168.35.105	Mezza Lab, Groud Floor	Compute node
mac07	192.168.35.107	Mezza Lab, Groud Floor	Compute node
mac08	192.168.35.108	Mezza Lab, Groud Floor	Compute node
mac09	192.168.35.109	Mezza Lab, Groud Floor	Compute node
mac10	192.168.35.110	Mezza Lab, Groud Floor	Compute node
mac11	192.168.35.111	Mezza Lab, Groud Floor	Compute node

Table 4 NFS Server mounted by cluster nodes

Server Name	IP Address	NFS Exports	Capacity
bigmac	192.168.20.200	/data/bigmac	700GB
		/data/neonate	700GB
		/data/cogneuro	700GB
		/data/primal	700GB
fornix	192.168.20.24	/data/fornix	1TB
oasis	192.168.20.22	/data/oasis	770GB

2. Login to the cluster

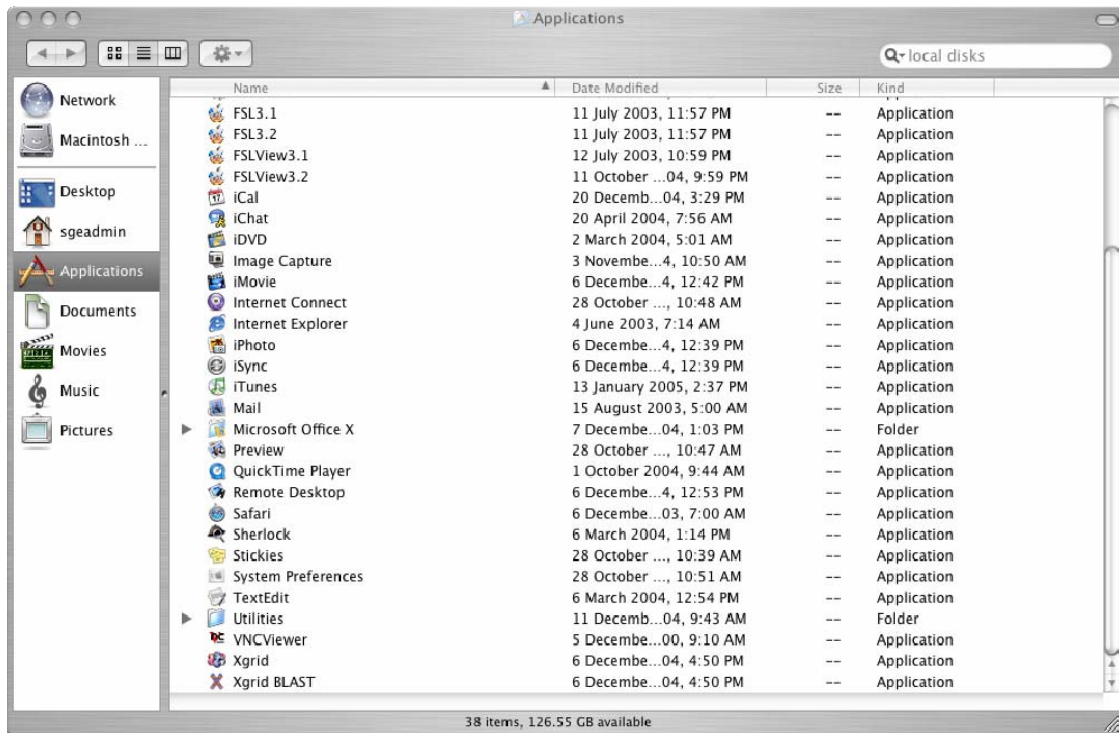
NIS(Network Information Service) clients have been installed on every cluster nodes and they can require authentication service from our NIS server (sonic 192.168.20.25). Thus, we can login to any of the cluster nodes using your existing account and password. Due to the NIS client provided by Mac OS X can only accept passwords encrypted with DES algorithm, you may need to contact system administrator(Wilson, Desmond or Wee Siong) to re-encrypt your password with DES algorithm if you cannot login to the nodes of the cluster.

There are two ways to utilize the cluster.

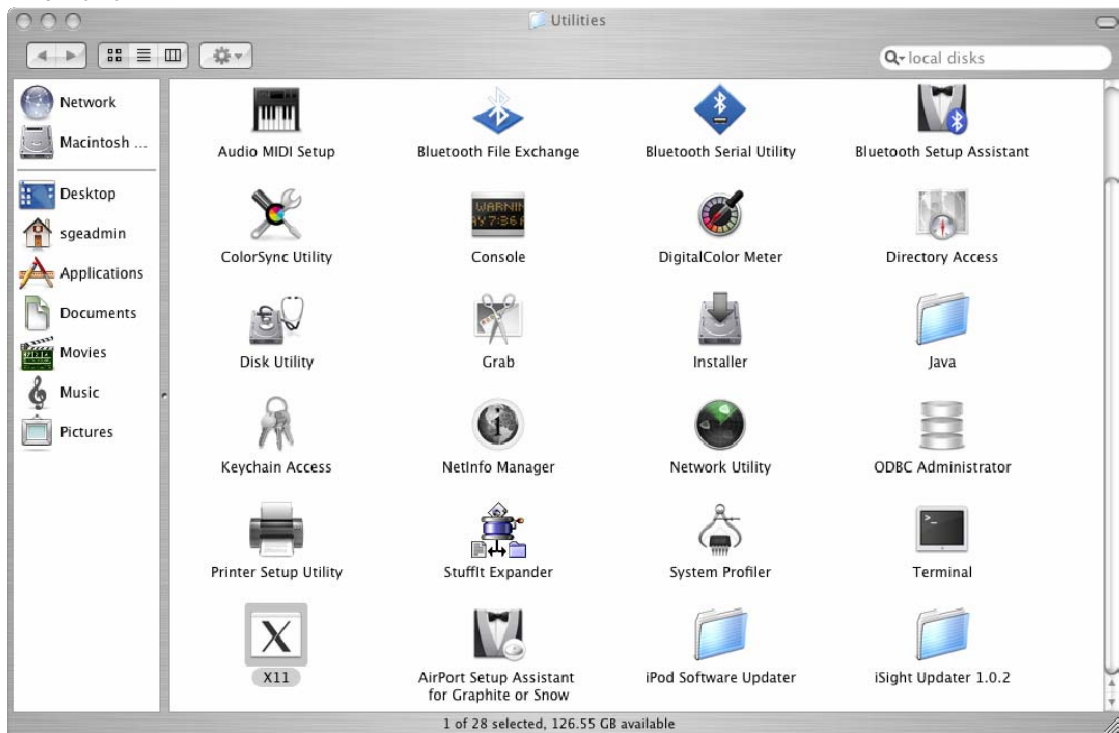
i) Directly login to any of the MAC computers in the Mezza lab, Ground Floor, Alan Gilbert Building.



Click “Other...”, then input your username and password,



Then click “X11”



In the terminal window, you can type “fsl” to start FSL GUI program or you can also directly run the command-line tools of FSL. For example, “bet” or “feat”.



ii) Remote login to any of the cluster nodes using ssh

Before you ssh to the cluster node, you need to run

xhost +

on your local desktop to accept GUI display from the remote cluster node.

Then ssh to remote cluster node, for example:

ssh yourusername@cluster

or

ssh yourusername@192.168.20.201

After login to the cluster node, you need to export display to your local desktop, for example:

export DISPLAY=192.168.20.XX:YY (XX is the IP address of the desktop you are using, YY is the display number you are using)

or

export DISPLAY=yourhostname:yourdisplaynumber

Then, you can run the FSL GUI or command line tools on the cluster nodes.

3. Submit jobs to Sun Grid Engine

3.1 qrsh --- Interactive jobs

Login to any of the cluster nodes, and

qrsh

SGE will automatically assign a node and open a terminal for you, then you can run your jobs, for example:

bet /data/bigmac/sge/feeds/data/structural /data/bigmac/sge/feeds/structural_brain -f 0.5 -g 0 -m

3.2 qsub --- Batch jobs

```
qsub -s /bin/sh -b y -o /data/bigmac/sge/ -e /data/bigmac/sge/ /usr/local/fs13.1/bet  
/data/bigmac/sge/feeds/data/structural /data/bigmac/sge/feeds/structural_brain -f 0.5 -g 0 -m
```

Although qsub is a very useful tool to submit batch jobs, it is difficult to use because you have to specify too many parameters, for example, the shell, -s /bin/sh, the standard output and standard error. A better way to use qsub is writing some shell scripts to wrap you jobs and predefine the parameters in the shell scripts. Considering that not all of the users are familiar with shell programming. We developed wrapper scripts for most of the FSL command line tools. We will discuss the usages of these tools in section 4.1.

3.3 qstat --- Check the job status

qstat is a tool to monitor the jobs submitted to Sun Grid Engine, just run
qstat

job-ID	prior	name	user	state	submit/start at	queue	slots
3435	0.56000	sge_bet	wilson	t	02/09/2005 10:45:57	all.q@cluster.ni.nnf.unimelb.e	1
3429	0.56000	QRLOGIN	wilson	r	02/09/2005 10:32:27	all.q@mac05.ni.nnf.unimelb.edu	1
3434	0.56000	sge_bet	wilson	t	02/09/2005 10:45:57	all.q@mac05.ni.nnf.unimelb.edu	1
3436	0.56000	sge_bet	wilson	t	02/09/2005 10:45:57	all.q@mac10.ni.nnf.unimelb.edu	1

Job information will be displayed on screen. For more detail about qstat, run “man qstat” to see the man pages.

4. FSL integration with Sun Grid Engine

4.1 Simplify FSL job submission by using wrapper scripts for FSL command line tools

As discussed in 3.2, we developed a group of wrapper scripts for FSL tools.

Table 5 Wrapper Scripts for FSL tools

FSL Command Line Tools	Corresponding Shell Script
bet	sge_bet
fast	sge_fast
feat	sge_feat
flirt	sge_flirt
mcflirt	sge_mcflirt
mfast	sge_mfast
melodic	sge_melodic
siena	sge_siena
sienax	sge_sienax
susan_smooth	sge_susan_smooth

```

mac0a:/usr/local/fsl/bin sgeadmin$ ls sge*
sge_bet          sge_mcflirt      sge_sienax
sge_fast         sge_melodic      sge_susan_smooth
sge_feat         sge_mfast
sge_flirt        sge_siena
mac0a:/usr/local/fsl/bin sgeadmin$ sge_bet █

```

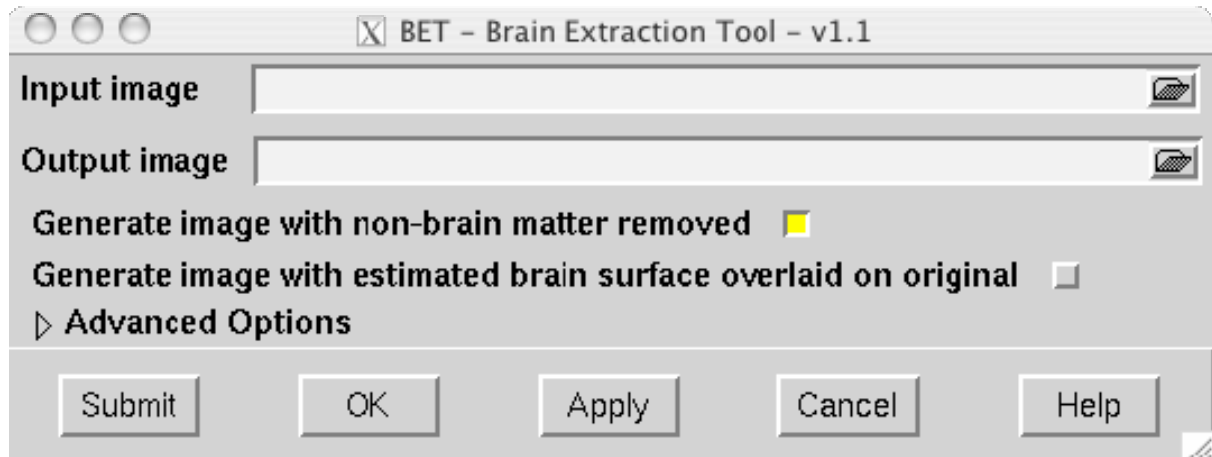
```

mac0a:/usr/local/fsl/bin sgeadmin$ ls sge*
sge_bet          sge_feat          sge_mcflirt      sge_mfast        sge_sienax
sge_fast         sge_flirt         sge_melodic      sge_siena        sge_susan_smooth
mac0a:/usr/local/fsl/bin sgeadmin$ sge_bet /data/bigmac/sge/feeds/data/structural /data/bigmac/sge/feeds/data/structural_brain -f 0.5 -g 0
Your job 3027 ("sge_bet") has been submitted.
mac0a:/usr/local/fsl/bin sgeadmin$ qstat
-----
job-ID prior name user state submit/start at queue slots ja-task-ID
-----
2951 0.56000 slp eduff r 02/03/2005 10:41:25 all.q@mac05.ni.nmf.unimelb.edu 1
2899 0.56000 slp eduff r 02/03/2005 10:41:25 all.q@mac10.ni.nmf.unimelb.edu 1
3027 0.00000 sge_bet sgeadmin qw 02/03/2005 10:41:35 1
mac0a:/usr/local/fsl/bin sgeadmin$ qstat
-----
job-ID prior name user state submit/start at queue slots ja-task-ID
-----
2951 0.56000 slp eduff r 02/03/2005 10:41:25 all.q@mac05.ni.nmf.unimelb.edu 1
2899 0.56000 slp eduff r 02/03/2005 10:41:25 all.q@mac10.ni.nmf.unimelb.edu 1
3027 0.00000 sge_bet sgeadmin qw 02/03/2005 10:41:35 1
mac0a:/usr/local/fsl/bin sgeadmin$ qstat
-----
job-ID prior name user state submit/start at queue slots ja-task-ID
-----
2951 0.56000 slp eduff r 02/03/2005 10:41:25 all.q@mac05.ni.nmf.unimelb.edu 1
2899 0.56000 slp eduff r 02/03/2005 10:41:25 all.q@mac10.ni.nmf.unimelb.edu 1
3027 0.00000 sge_bct sgeadmin qw 02/03/2005 10:41:35 1
mac0a:/usr/local/fsl/bin sgeadmin$
mac0a:/usr/local/fsl/bin sgeadmin$ qstat
-----
job-ID prior name user state submit/start at queue slots ja-task-ID
-----
3027 0.00000 sge_bet sgeadmin qw 02/03/2005 10:41:35 1
mac0a:/usr/local/fsl/bin sgeadmin$ qstat
-----
job-ID prior name user state submit/start at queue slots ja-task-ID
-----
3027 0.56000 sge_bet sgeadmin t 02/03/2005 10:41:55 all.q@mac03.ni.nmf.unimelb.edu 1
mac0a:/usr/local/fsl/bin sgeadmin$ qstat
-----
job-ID prior name user state submit/start at queue slots ja-task-ID
-----
3027 0.56000 sge_bet sgeadmin t 02/03/2005 10:41:55 all.q@mac03.ni.nmf.unimelb.edu 1
mac0a:/usr/local/fsl/bin sgeadmin$ █

```

4.2 Submit FSL jobs through FSL GUI by modifying FSL GUI

For the users who prefer to use FSL GUI tools, we modified the FSL GUI and add an option to submit job to Sun Grid Engine instead of running on local computer.



The picture above shows the modification to the BET GUI, you just need to specify the input&output images and click "Submit" button.

Note 1: You have to use absolute path for the input and output files and you have to store both input files and output files on the file server that are mounted by the cluster nodes, namely, /data/fornix, /data/oasis, /data/primal, /data/bigmac, /data/cogneuro, /data/neonate.

Note 2: Don not store big datasets or do analysis in your home directory, because the file server(sonic.ni.nnf.unimelb.edu.au), which the home directory is located in, provides only 80G total storage and it is not as fast as others(such as bigmac, fornix, oasis).