

# SAPAC High Performance Computing Induction

**Grant Ward**

Computer Specialist (Facilities)

SAPAC

[Grant.Ward@sapac.edu.au](mailto:Grant.Ward@sapac.edu.au)

(08) 8303 8264



The South Australian Partnership for Advanced Computing

# Course Contents

- HPC Facilities overview
- Software
- Queuing System
- Connecting to SAPAC facilities
- Help Desk



# SAPAC facilities

- 5 HPC systems overall (Titan, Perseus, Orion, Hydra and Aquila), and a new system early 2007.
  - Broad range of HPC technologies
- ≈20TB mass storage facilities
- Visualisation: SAPAC VisLab and SAVRC
- AccessGrid
- Training lab, presentation rooms etc.
- Network infrastructure
  - SABRENet .....
  - Additional 10Gb/s to VPAC
- Access to external facilities: APAC grid



# SAPAC supercomputers



- “Orion”
- 40 node Sun Technical Computer Farm:  
Sun E420-R SMP nodes
  - 4 x UltraSPARC cpus per node
  - 4GB memory per node
  - Myrinet interconnect
- Solaris 8

# SAPAC supercomputers



- “Hydra”
- 128 node IBM 1350 Linux cluster
  - 2 x 2.4GHz Xeon cpus per node
  - 2GB memory per node
  - Myrinet interconnect
- 1.2 Tflops
- 32 bit - Red Hat Enterprise Linux 3



The South Australian Partnership for Advanced Computing

# SAPAC supercomputers



- “Aquila”
- SGI Altix 3000 SMP system
  - 160 x 1.3GHz Itanium2 cpus
  - 160GB memory
  - NUMAlink ‘interconnect’
- 64 bit - Red Hat Enterprise Linux 3 (with SGI SMP enhancements)

# Software - Compilers

- GNU compiler collection (gcc, g++, g77) - on all systems.
- Portland Group compiler suite (pgcc, pgCC, pgf77, pgf90, pghpf). Includes optimised LAPACK and SCALAPACK mathematical libraries – Hydra and Perseus.
- Intel compiler suite (icc, ipCC, ifort). Includes Intel's Math Kernel Library with optimised BLAS and LAPACK libraries – Hydra and Aquila.
- SUNPro compiler suite (cc, CC, f77, f90, f95) - Orion.



# Software - Libraries

- BLAS
- LAPACK
- ScaLAPACK
- FFTW
- netcdf
- sprng
- SuperLU

Generally all on Hydra, but can be made available on other systems if requested and where appropriate.

# Software - Applications

SAPAC maintains a large amount of other open source and commercial software purchased by various research groups.

Access to free and open source software (FOSS) is configured by the use of environment variables in your login files.

Access to commercial software is controlled by the use of access control lists and/or license count.



# Software - Applications

Some examples of the software available on SAPAC facilities is listed below:

**Physics:** EGSnrc, Beamnrc

**Chemistry:** Gaussian98, Gaussian03, GAMESS, Dalton, Materials Studio, NAMD, Crystal

**Bioinformatics:** Blast, MrBayes, ClustalW, Hmmer, Phylip, IM, T-Coffee and many more.

**Engineering & Maths:** Fluent, Matlab, R, Simics



The South Australian Partnership for Advanced Computing

# Queuing System

To ensure efficient and fair use of resources all jobs run on SAPAC's facilities need to be submitted as batch jobs to a queuing system.

The queuing system we use is **Torque**, a variant of the PBS queuing system. It is similar, but not identical to the system run by APAC.



# Queuing System

## Overview of Torque:

- 3 components, pbs\_server, pbs\_sched, pbs\_mom
- Server & Scheduler run on the head node, Moms run on each compute node.
  - Scheduler informs the server of free processor set.
  - Server designates one mom to be the “mother superior” and hands over necessary job information.
  - Moms form a collective “sisterhood” and organise internodal communications.
  - Mother superior releases nodes back to server at job completion.

# Sample Torque script – Sequential (Hydra)

```
#!/bin/sh
#PBS -V
### Job name
#PBS -N MyJobName
### Declare jobs non-rerunnable
#PBS -r n
### Output files
#PBS -j oe
### Mail to user
#PBS -m ae
#PBS -M grant.ward@sapac.edu.au
### Queue name
#PBS -q hydra
### Number of nodes
#PBS -l nodes=1,walltime=01:00:00

cd $PBS_O_WORKDIR
# Run the executable
g03 < InputFile.com > OutputFile.log
```



# Sample Torque script – MPI (Hydra)

```
#!/bin/sh
#PBS -V
### Job name
#PBS -N MyJobName
### Declare jobs non-rerunnable
#PBS -r n
### Output files
#PBS -j oe
### Mail to user
#PBS -m ae
#PBS -M grant.ward@sapac.edu.au
### Queue name
#PBS -q hydra
### Number of nodes
#PBS -l nodes=2:ppn=2:myrinet,walltime=01:00:00

cd $PBS_O_WORKDIR
# Run the executable
mpiexec MyExe < InputFile.in > OutputFile.out
```



# Submitting jobs

To submit your job, simply run the **qsub** command with the name of your torque script file.

```
qsub sample.pbs
```

You will get a response something like:

```
Job submitted. Torque JobId: 28195.hydra
```

**Note:** Scripts for other systems will be slightly different. Please consult the 'Running Jobs' page of the relevant User Guide.



# Checking job status

To view the status of your job simply type **qstat -a**

hydra:

Job ID	Username	Queue	Jobname	SessID	NDS	TSK	Req'd Memory	Req'd Time	S	Elap Time
25016.hydra	sandery	seq	sef8	21497	1	--	--	1000:	E	627:1
28167.hydra	hongyi	seq	newtest64s	29764	1	--	--	2440:	R	02:03
28168.hydra	hongyi	seq	newtest64s	--	1	--	--	2440:	Q	--
28193.hydra	honcao	seq	xcomp13	24585	1	--	--	1000:	S	29:47
28196.hydra	honcao	seq	xcomp16	29707	1	--	--	1000:	R	29:47
28201.hydra	honcao	seq	rxttfunc3	--	1	--	--	1000:	H	--

To see other options to qstat consult the manpage.



# Deleting jobs

To delete a job from the queues, run **qdel** with the numeric part of the job identifier (from `qstat` or from `qsub` when you submit a job).

```
qdel 28195
```



# Connecting to SAPAC

Linux, MacOS X users have ssh, X11 tools already.

Windows users need programs such as :

- **PuTTY** - an ssh client to login to our facilities
- **WinSCP** - a ssh based file transfer utility with drag & drop interface
- **Xming** - a free X11 server (install all fonts)

**Note:** Beware the Control-M problem (use dos2unix).



# Help Desk

If you have any problems running jobs on SAPAC facilities or require software to be installed, please contact the helpdesk preferably by email ([helpdesk@sapac.edu.au](mailto:helpdesk@sapac.edu.au)) or phone (8303 8263).

The helpdesk email is monitored by several people and someone will endeavour to respond as soon as possible to your query.





The South Australian Partnership for Advanced Computing