

INTRODUCTION TO HIGH PERFORMANCE COMPUTING (HPC)

*introduction to working with HPC
computing clusters*

Gregory B. Poole

Astronomy Data and Computing Services, Swinburne



*Based on content developed by
Amr H. Hassan*




TABLE OF CONTENTS

- Introduction to parallel computing
- Introduction to computing clusters - Hardware configuration
 - *Node layout*
 - *Storage*
- Introduction to computing clusters - Queues
- Operating a computing cluster
 - *Shell scripts*
 - *Working with queues*
 - *Logging in with SSH*


INTRODUCTION TO PARALLEL COMPUTING




INTRODUCTION TO PARALLEL COMPUTING




INTRODUCTION TO PARALLEL COMPUTING



INTRODUCTION TO PARALLEL COMPUTING



INTRODUCTION TO PARALLEL COMPUTING



INTRODUCTION TO COMPUTING CLUSTERS – HARDWARE CONFIGURATION



The gSTAR Cluster at Swinburne

INTRODUCTION TO COMPUTING CLUSTERS – NODE LAYOUT

Computing clusters consist of many specialised 'nodes' ...

- **The head node**

- *This is the node through which the system is accessed.*
- *Users must first login to the head node to gain access to the system.*

- **Compute nodes**

- *These provide the vast majority of the system's computing power.*
- *They can not be logged into directly by users.*
- *Access can only be achieved through a 'queue' system (discussed later).*

- **Interactive nodes**

- *Some (but not all) computing clusters offer a small number of interactive nodes.*
- *These function like standard remote-access Unix systems.*
- *Intended only for small/short jobs (eg. editing & compiling code, quick & small test runs, etc.)*
- *Alternatively, interactive jobs can be launched on compute nodes via the queue (discussed later).*

INTRODUCTION TO COMPUTING CLUSTERS – STORAGE

Storage policies differ significantly from system to system. Some important issues to consider:

- **Home directory storage**
 - *Most computing clusters allocate only small amounts of storage in a user's home directory. It's generally best policy to avoid storing large datasets here.*
- **Project storage**
 - *Generally, the majority of storage is attached to a project with specific quotas on storage volume, file count, etc.*
- **Data longevity**
 - *Often, the bulk of a user's storage sits in volatile scratch spaces. System administrators may have a policy of frequently cleaning this and the user must find their own long-term solutions, either on the system or by moving data off-site.*

INTRODUCTION TO COMPUTING CLUSTERS – QUEUES

Making sure that the resources of the compute nodes are efficiently utilised and equitably shared between users is the challenging task of the queue system.

Diverse users with diverse needs:



32 cores for 3 hours



1 core needed interactively




3 cores for 24 hours



15 cores for 17 hours + GPUs

INTRODUCTION TO COMPUTING CLUSTERS – QUEUES

To help manage this, the queue system is often configured to have several queues for managing different use cases or hardware configurations. For example:



INTRODUCTION TO COMPUTING CLUSTERS – QUEUES

Users create "submission scripts" specifying the commands they want run and the resources they need. The queue system turns these into 'jobs' with assigned IDs.





Submission Script

Specify needed resources:

- Which queue to use
- No. of nodes
- No. of cores per node
- RAM (per node;total)
- run time, etc.


List commands to run:

Command 1
Command 2
⋮
Command N



INTRODUCTION TO COMPUTING CLUSTERS – QUEUES

The scheduler sends jobs to compute nodes when needed resources are available ...



OPERATING A COMPUTING CLUSTER

➤ Shells

- A "Unix Shell" is a command-line interface (CLI) which provides a traditional text-driven interface for Unix-like systems. Several variants exist (eg. "Bash" is a popular choice).

```
[g2: ~] > echo $SHELL  
/bin/bash  
[g2: ~] > |
```

➤ Scripts

- Users direct the operation of the computer by entering commands as text for a CLI to execute, or by creating "scripts" of one or more such commands.

```
[g2: ~] > ./hello_world.sh  
Hello World  
[g2: ~] > |
```

OPERATING A COMPUTING CLUSTER – SHELL SCRIPTS

- Create file
 - Use a text editor (eg. 'emacs') to create a file listing the commands you want the script to run.
- Set permissions to permit the file to be executed as a script
 - Every file has a list of permissions ('read', 'write', 'execute'; set separately for the 'user', the user's 'group', or 'others').
 - To execute the script, make sure 'execute' permissions are set for the 'user'.
- Alternatively, the script can be 'sourced'

OPERATING A COMPUTING CLUSTER – SHELL SCRIPTS


Insert demonstration here

```
[guest01@g2 ~]$ echo $SHELL
/bin/bash
[guest01@g2 ~]$ emacs hello.sh
[guest01@g2 ~]$ ls
hello.sh
[guest01@g2 ~]$ cat hello.sh
#!/bin/bash

# This line will write to the terminal window
echo Hello World ...
[guest01@g2 ~]$ ls -l
total 4
-rw-r--r-- 1 guest01 p093_swin 80 Jun 15 15:23 hello.sh
[guest01@g2 ~]$ chmod u+x hello.sh
[guest01@g2 ~]$ ls -l
total 4
-rwxr--r-- 1 guest01 p093_swin 80 Jun 15 15:23 hello.sh
[guest01@g2 ~]$ ./hello.sh
Hello World ...
[guest01@g2 ~]$ clear
[guest01@g2 ~]$ history
```

OPERATING A COMPUTING CLUSTER – SHELL SCRIPTS

What is the difference between executing and sourcing a shell script?



OPERATING A COMPUTING CLUSTER – SHELL SCRIPTS

Insert demonstration here

```
[guest01@g2 ~]$ my_var="g2 user"
[guest01@g2 ~]$ emacs hello.sh
[guest01@g2 ~]$ ./hello.sh
Hello World...
[guest01@g2 ~]$ source hello.sh
Hello World... g2 user
[guest01@g2 ~]$ emacs hello.sh
[guest01@g2 ~]$ ./hello.sh
Hello World... /bin/bash
[guest01@g2 ~]$ source hello.sh
Hello World... /bin/bash
[guest01@g2 ~]$
```

OPERATING A COMPUTING CLUSTER – SHELL SCRIPTS

Shell scripts can be given optional parameters when executed:

script arg1 arg2 ... argN



... the name of the script can be accessed in the script with the variable name: \$0

... the arguments are accessed by the variable names: \$1, \$2, ..., \$N

OPERATING A COMPUTING CLUSTER – SHELL SCRIPTS

Insert demonstration here

```
[guest01@g2 ~]$ emacs params.sh
[guest01@g2 ~]$ chmod u+x params.sh
[guest01@g2 ~]$ ./params.sh p1 p2 p3
Exec Name: ./params.sh
Param 1 : p1
Param 2 : p2
Param 3 : p3
[guest01@g2 ~]$
```

OPERATING A COMPUTING CLUSTER – WORKING WITH QUEUES

- The submission scripts that launch jobs on the compute nodes are just shell scripts, except:
 - *Specialised directives get added to the script to specify resource requirements.*
 - *Execute permissions do not need to be set*
 - *They are "submitted to the queue" with a specialised command which differs depending on the queue system used.*
- The two most popular queue systems:
 - **The Portable Batch System (PBS)**
Submission command: 'qsub jobscript'
 - **Slurm**
Submission command: 'sbatch jobscript'
- Rather than send output from the standard output stream to the screen, it gets sent to a special file instead.
 - *By default, a job with job name <Name> and job ID <ID> has its standard output sent to a file with name <Name>.o<ID>*

OPERATING A COMPUTING CLUSTER – WORKING WITH QUEUES

.....

Slurm and PBS work similarly (from the user's perspective). Some useful commands:

Action	PBS	Slurm
Submit a job	qsub jobscript	sbatch jobscript
Delete a job	qdel jobID	scancel jobID
Delete all jobs of user	qdel 'qselect -u user'	scancel -u user
Status of all jobs of user	qstat -u user	squeue -u user
Status of all jobs	qstat -a	squeue

OPERATING A COMPUTING CLUSTER – WORKING WITH QUEUES

Insert demonstration here


```
[guest01@g2 ~]$ emacs hello.sh
[guest01@g2 ~]$ qsub hello.sh
5518751.pbs.hpc.swin.edu.au
[guest01@g2 ~]$ qstat 5518751
```

Job id	Name	User	Time Use	S	Queue
5518751.pbs	TestJob	guest01	0 R		gstar

```
[guest01@g2 ~]$ qstat 5518751
qstat: Unknown Job Id 5518751.pbs.hpc.swin.edu.au
[guest01@g2 ~]$ ls -l TestJob*
-rw----- 1 guest01 p093_swin 0 Jun 16 13:16 TestJob.e5518751
-rw----- 1 guest01 p093_swin 25 Jun 16 13:16 TestJob.o5518751
[guest01@g2 ~]$ more TestJob.o5518751
Hello World... /bin/bash
[guest01@g2 ~]$ |
```

OPERATING A COMPUTING CLUSTER – LOGGING IN WITH SSH

Generally, clusters are accessed using the Secure Shell (SSH) protocol ...



OPERATING A COMPUTING CLUSTER – LOGGING IN w/ SSH

Every platform has several SSH tools to choose from. Some common defaults ...



... terminal (always available by default).



... Applications > Utilities > Terminal



... PuTTY: <http://goo.gl/XMggii>

OPERATING A COMPUTING CLUSTER – LOGGING IN w/ SSH

Note that external access to the system must go through the "Head Node" ...

