

# LENGAU CLUSTER AT CHPC

The main system at the CHPC for high performance computing is a cluster supercomputer called Lengau — cheetah in seTswana. This peta-scale system consists of Dell servers, powered by Intel processors, using FDR InfiniBand by Mellanox and is managed by the Bright Cluster Manager.



## 1: The Cluster

→ The CHPC's Dell Linux cluster has been up and running since 2014.

↳ The system is a homogeneous cluster, comprising Intel 5th generation CPUs. As of March 2017 it has 1368 compute nodes with 24 cores and 128 GiB memory (360 nodes have only 64 GiB) each, and five large memory “fat” nodes with 56 cores and 1TiB each, all interconnected using FDR 56 Gb/s InfiniBand accessing 4 PB of shared storage over the Lustre filesystem.

→ The cluster has both NFS and the Lustre filesystems over Infiniband:

Mount Point	File System	Size	Quota	Backup
/home	NFS	80Tb	15GB	NO
/mnt/lustre/users	lustre	4PB	None 	NO

## 2: The installed softwares/libraries

→ CHPC uses the GNU modules utility, which manipulates your environment, to provide access to the supported software in /apps/.

↳ Directly in your `.bashrc`, we will load the following modules to use the associated software/libraries: Matlab, Neview, Intel Fortran Compilers, NCO tools, ...

Command	Description
<code>module purge</code>	Remove all loaded modules
<code>module avail</code>	List of available modules
<code>module list</code>	List currently loaded modules
<code>module help</code>	Give information of a particular module file's operations

## 3: Job Scheduler

→ The CHPC cluster uses **PBSPRO** as its job scheduler. With the exception of interactive jobs, all jobs are submitted to a batch queuing system and only execute when the requested resources become available. All batch jobs are queued according to priority. A user's priority is not static: the CHPC uses the “Fairshare” facility of PBSPRO to modify priority based on activity. This is done to ensure the finite resources of the CHPC cluster are shared fairly amongst all users.

PBS Pro commands	Description	PBS command example
<code>qsub [script file]</code>	Job submission	<code>qsub run_croco.pbs</code>
<code>qstat-u login</code>	Job status (for user)	<code>qstat -u sillig</code> (or the alias <code>qs</code> )
<code>qstat-f [job_id]</code>	Extended job status	<code>qstat -f 10098976</code>
<code>qdel [job_id]</code>	Job deletion	<code>qdel 10098976</code>
<code>qstat-Q</code>	List of usable queues	

→ The available queues with their nominal parameters are given in the following table:

Queue Name	Max. cores	Min. cores	Max. jobs		Max. time	Notes	Access
	per job		in queue	running	hrs		
serial	23	1	24	10	48	For single-node non-parallel jobs.	
seriallong	12	1	24	10	144	For very long sub 1-node jobs.	
smp	24	24	20	10	96	For single-node parallel jobs.	
<b>normal</b>	<b>240</b>	<b>25</b>	<b>20</b>	<b>10</b>	<b>48</b>	<b>The standard queue for parallel jobs</b>	
large	2400	264	10	5	96	For large parallel runs	<i>Restricted</i>
xlarge	6000	2424	2	1	96	For extra-large parallel runs	<i>Restricted</i>
express	2400	25	N/A	100 total nodes	96	For paid commercial use only	<i>Restricted</i>
bigmem	280	28	4	1	48	For the large memory (1TiB RAM) nodes.	<i>Restricted</i>
vis	12	1	1	1	3	Visualisation node	
test	24	1	1	1	3	Normal nodes, for testing only	
gpu_1	10	1		2	12	Up to 10 cpus, 1 GPU	
gpu_2	20	1		2	12	Up to 20 cpus, 2 GPUs	
gpu_3	36	1		2	12	Up to 36 cpus, 3 GPUs	
gpu_4	40	1		2	12	Up to 40 cpus, 4 GPUs	
gpu_long	20	1		1	24	Up to 20 cpus, 1 or 2 GPUs	<i>Restricted</i>

#### 4: Supplementary material

→ Some useful contents are available on CHPC website:



- Video demo of interactive sessions: <https://www.youtube.com/watch?v=FZFwgSVE0HY>
- Video first job script: <https://www.youtube.com/watch?v=cuxj0oWvQzA>