High Performance Computing

Emerging Researcher Series

UCT eResearch
ACCELERATING RESEARCH

Andrew Lewis / Timothy Carr **Senior Technical Specialist**

www.eresearch.uct.ac.za

1. UCT High Performance Computing

What is HPC?

How fast is it?

Is it easy to use?

1. UCT High Performance Computing

What is HPC?

The aggregation of computing resources.

How fast is it?

Depends on your use and the application.

Is it easy to use?

“Yes” but a slight mindset change.

2. How do I get access to UCT HPC

http://hpc.uct.ac.za

Apply for an account

Download a Word document, complete and upload via the web page.

2 minutes.

2. How do I get access to UCT HPC

Who can get access?

MSc and above get access to the standard partition.

Undergrad and Honours may access the older nodes.

The HPC cluster is intended for research, not course work.

3. How do I log into UCT HPC?

There is a web based dashboard...

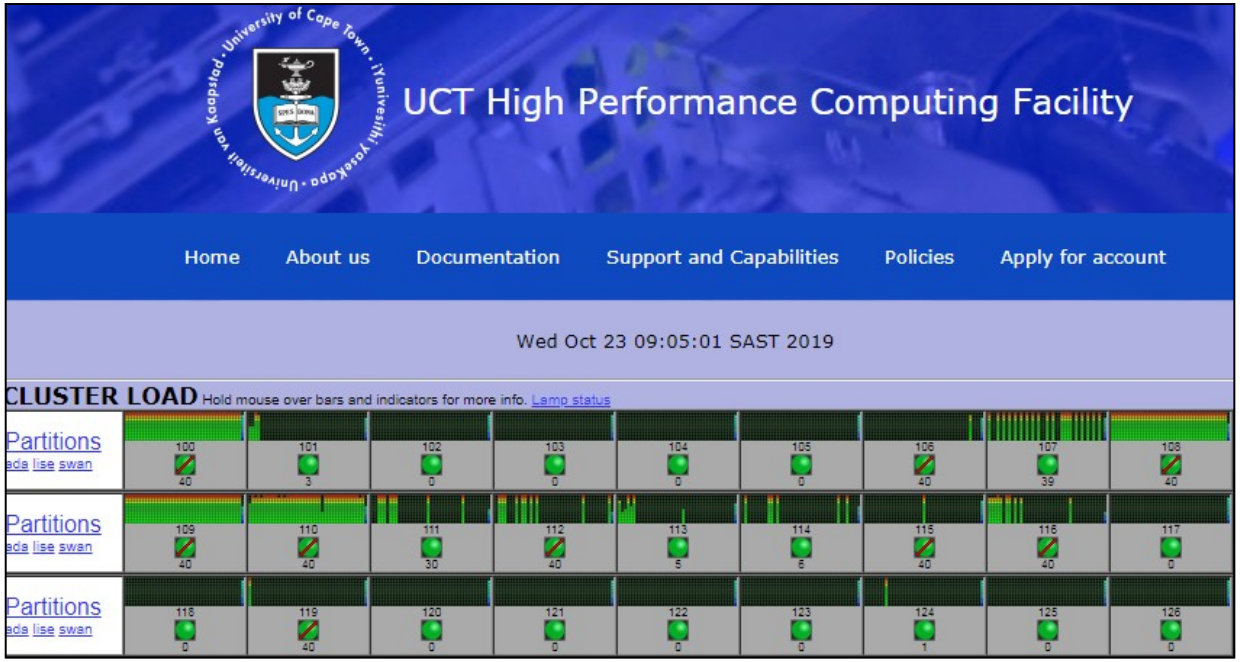

But there is no GUI

What does this mean???

3. How do I log into UCT HPC?

You need to have a basic understanding of Linux:

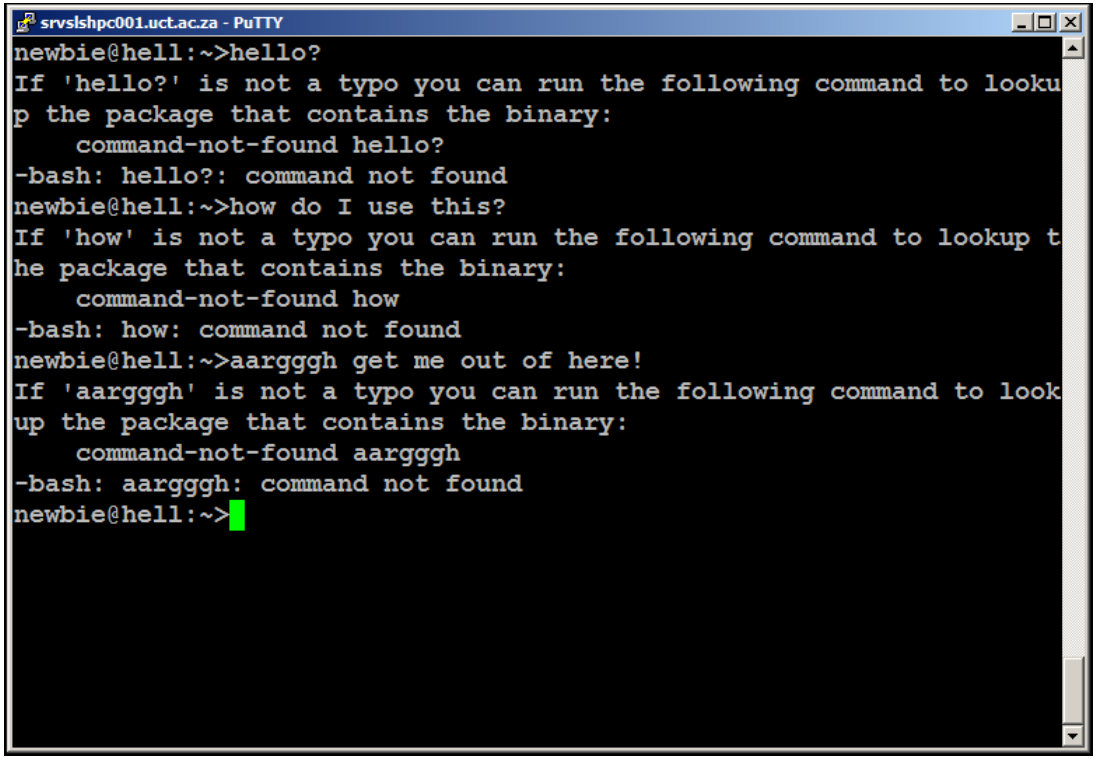

3. How do I log into UCT HPC?

Don’t panic! We have training courses.

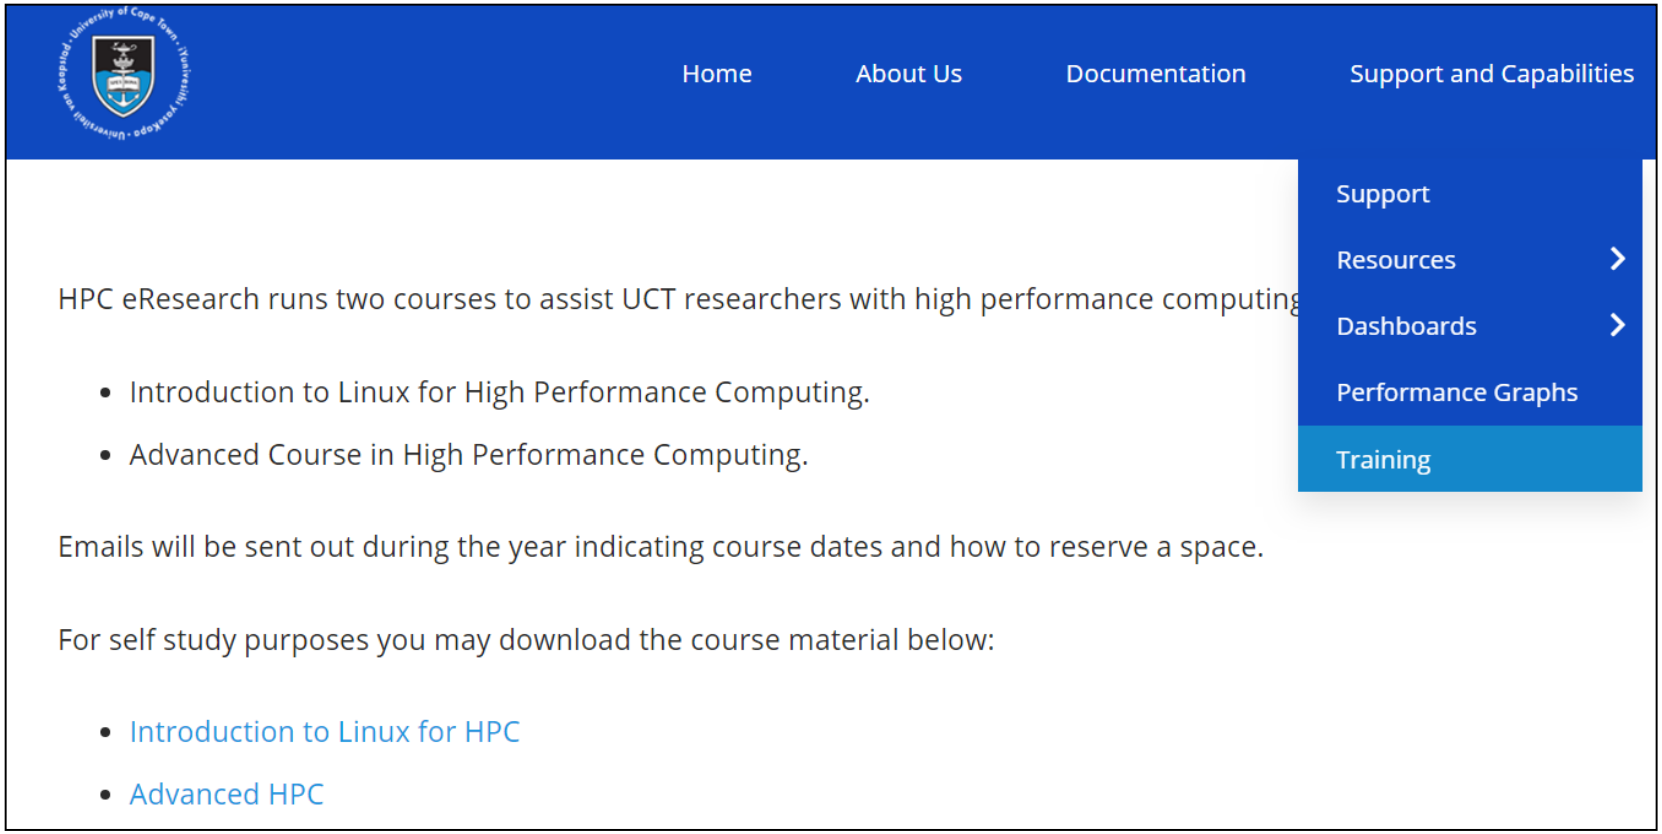

4. What is available?

2464 CPU Cores

x2 Tesla GPU Nodes, x4 Tesla M2090 Cards

x2 Tesla K40, K80 Nodes,Tesla P100 Nodes

x2 1TB Memory Nodes

4. What is available per user?

Basic allocation is 120 cores.

This can be boosted if the cluster is under-utilized.

Wall time is 3 days per job on default partition.

UCT HPC is a free, shared resource.

5. How much storage can I access?

Home Directory: 20GB

Scratch (policy): 100GB default, up to 10TB, but this is a temporary resource.

6. Can I get access to more resources?

Slurm Workload Manager

Different partitions:

- **Longer wall time but fewer cores.**
- **More cores but shorter wall time.**

Buy in model

- **Collaboration, not pay to play.**

CIFS storage

- **Limitations and real $$$.**

7. Stats and citations

Since inception, 2009:

Assisted 500+ Researchers.

Run 3 million+ jobs

Computed 29 million CPU hours

Acknowledged in 139 publications

8. Software

R
PYTHON
GALAXY
C\GCC\C++\G++
GATK
FORTRAN
HYPHY
QIIME
MATLAB
NEXTGENSEQ
PLINK
PERL
BLAST
JAVA
SESKA
CUDA
SAMTOOLS
BWA
STATA
GROMACS
NAMD
NETCDF
GAUSSIAN
VCFTOOLS
ABAQUS
CASA
OPENMPI
ANNOVAR
LSDYNA
ROMS
OCTAVE
FREESURFER
VASP
GALFACTS
FASTQC
DL_POLY
AMBER
EMMAX
ADMIXTURE
PYRAP
ADMB
NASP

VMD
MRBAYES
NUMPY
FSL
MONTBLANC
OPENFOAM
EMBOSS
VELVET
UPARSE
EIGENSOFT
PICARDTOOLS
GCTA
SMALT
TRIMMOMATIC
FASTX
ELEMENTAL
MELD
MEGA
CRUX
GEODYN
BOWTIE
PRINSEQ
CDO
SNPEFF
SCIPY
NVIDIA
WRF
PCSWMM
BCFTOOLS
BOLT-LMM
MHCPRG
MUSCLE
NGS
ROOT
CUFFDIFF
IMPG-SUMMARY
FASTTREE
PYCUDA
SRI
ANSYS
ESPRESSO
PCGC

MATTERHORN
OPENMP
TINKER
UNAFOLD
DECONSEQ
ANALYSIS
METAL
LS-DYNA
SPIDER
ONETEP
ELAI
SPOLPRED
STRELKA
CYTHON
SNPTEST
ADMIXTOOL
MULTINEST
PEAR
SCAPEL
BIOPERL
OS
TOPHAT
GIZA++
LD-SCRORE
DL
TABIX
OPENMM
CHROMOPAINTER
BIOEDIT
NCAR
NCO
GAUSSIAN09
VARSCAN
MULTIMIX
GSL
PAUP
CD-HIT
BOWTIE2
IDL
NCBI-TOOLS
MPICC
FASTQ

DEPTHMAP
PYFASTA
GID
CERN
RSEM
LOFREQ
SUPPORTMIX
QUANTUM
ANCGWAS
PROGRESSIVECACTUS
CORTEX
BIOPYTHON
TPP
LAMPLD
PCADMIX
MAC
MOSES
GWAS
MPI
ARTEMIS
SEQUENCE
CUMMERBUND
FERRET
GEANT4
SOLVE
BOOST
MRICONVERT
HDF
NEURON
LINUX
CUFFLINKS
QPUQ
UNITY
KRONA
IMPUTE2
BEAGLE
SOAPDENOVO
GMT
YADE
CHIMERA
MOTHUR
PYNAST

GWAMA
SAGE
WINPOP
BLASTALL
BEAST
PYMULTINEST
MPI4PY
LDPRED
METASOFT
CMU
RFMIX
BEDTOOLS
PINDEL
CPP
ALLOY
SEQTK
AFNI
MAUVE
PHYLOSEQ
POLYFLOW
HDF5
THEANO
SAKKE
BIOCONDUCTOR
GLASSFISH
META
PHYTHON
MICROCOSM
SGA
BGZIP
MPIF90
SPREAD
EMAN2
FLUENT
HAPGEN
SOAPDENOVO-TRANS
NESSUS
EXONERATE
IFORTRAN
ORCA
PCA

9. Future Plans

S3 equivalent object storage

Software Delivery Pipeline using Jenkins, PR

Hardware Refresh of GPU Nodes – Support for TPUs

Singularity Container Engine

Questions ?