THE UNIVERSITY OF QUEENSLAND
AUSTRALIA

# SMP Computing Workshop

*Dr. Michael Bromley – School of Maths/Physics*

*O-week  Semester 1 - 2019*

CRIKEY Provider No 00025B

# Cluster Example

- Use windows Remote Desktop Protocol RDP to log into `getafix.smp.uq.edu.au`

- Then start a Terminal and type in command

```
srun --cpus-per-task=1 --time=0-2:00 –mem=12G --x11 --pty mathematica
```

# Overview

THE UNIVERSITY OF QUEENSLAND
AUSTRALIA

- UQ ITS/SMP Desktop Systems and Support
- National Facilities
- UQ / QCIF / RCC systems
- SMP core computing facility
  http://research.smp.uq.edu.au/computing/ for howto

# About Me – Senior Lecturer

THE UNIVERSITY OF QUEENSLAND
AUSTRALIA

- BSc – Physics, Computer Science, Maths (N.T.U.) - 1995
- NT Govt – traineeship in Unix Software/Database (Sun/Mosaic/Sybase/perl)
- Graduate traineeship - Local/Wide Area Cisco Networking
- Ph.D. in computational/theoretical physics (N.T.U.)

  Large scale (sparse) eigenproblems on DEC Workstations
- Postdoc - MPI calcs at SAPAC (#106 at time)
- Faculty at San Diego State University (USA)

  -> NSF grant to setup rack cluster there
- Future Fellowship on Atomic Physics via HPC → non-linear PDEs / eigenproblems
- Teach PHYS3071 / PHYS4070 / COSC3500

# ITS – Centralised Support

THE UNIVERSITY OF QUEENSLAND
AUSTRALIA

- Rose Radloff – Relationship Manager (Faculty of Science)
- Karl Blakeney - IT Service Delivery Manager (west team)
- ITS local technical/computer support is
- Physics - Sam Zammit, Frank Audsley, Room 6-306
- Maths - Nalini Gowda, Room 6-306
- → consultation times Tue/Thu between 2pm-3pm
- Phone 07 336 56000 → ITS helpdesk: to log urgent job,
- https://its.uq.edu.au/ → ITS helpdesk: login to log job
  OR email help@its.uq.edu.au include in the email:
  with “Attention Sam @ Physics” or “Nalini @ Maths”

# ITS – Research Infrastructure

THE UNIVERSITY OF QUEENSLAND
AUSTRALIA

- former Faculty of Science IT Staff

- Dr. Leslie Elliott
  Infrastructure Support Specialist, Research Infrastructure

- Ian Mortimer
  Systems Administrator, Research Infrastructure

# ITS notices...

- Office 365 Mail – Cloud Migration https://outlook.office365.com/owa/uq.edu.au
- Mirroring360 works - https://its.uq.edu.au/mirroring360
- ZOOM - https://its.uq.edu.au/zoom domain uqz
- In UniFi eMarket https://its.uq.edu.au/purchasing

# Desktop systems and support

THE UNIVERSITY OF QUEENSLAND
AUSTRALIA

- ITS prefers Windows PCs for homogenised support.
- SMP prefers Windows PCs to keep costs down but tries to support the provision of MacOS/linux OS desktops.
- PhD/MPhil get a desktop (MSc/Honours/B.Adv.Sci we try!)
- SMP operates a pool of laptops/laser pointers etc for borrowing: see Physics or Maths admin for more info.
- Extra software Windows, look for “Software Center”
  Extra software macOS, look for “Self Service”
- Other applications can be added to the above software portals (licensing allowable) with a request to ITS...
  SMP Mathematica/Origin/Minitab/Maple/Labview/ComSol
- [BYO Laptop – Log a ITS job for advice- minimal support]

# Printers/Scanners - SMP

- Send to (also photocopying and scan-to-email)
  `\\uq-print01.soe.uq.edu.au\smp-ricoh`
  (colour printers - ~$1 per 10 pages)
- Username as UQ\username e.g `UQ\s7654321`
- Log a job ITS - to get allocated to printer group
- other physics printers are
  `\\uq-print02.soe.uq.edu.au\smp-06-320-p1`
  `\\uq-print02.soe.uq.edu.au\smp-06-320-p2`
  `\\uq-print02.soe.uq.edu.au\smp-06-404-p1`

# Data Storage

- Data storage and security is important!

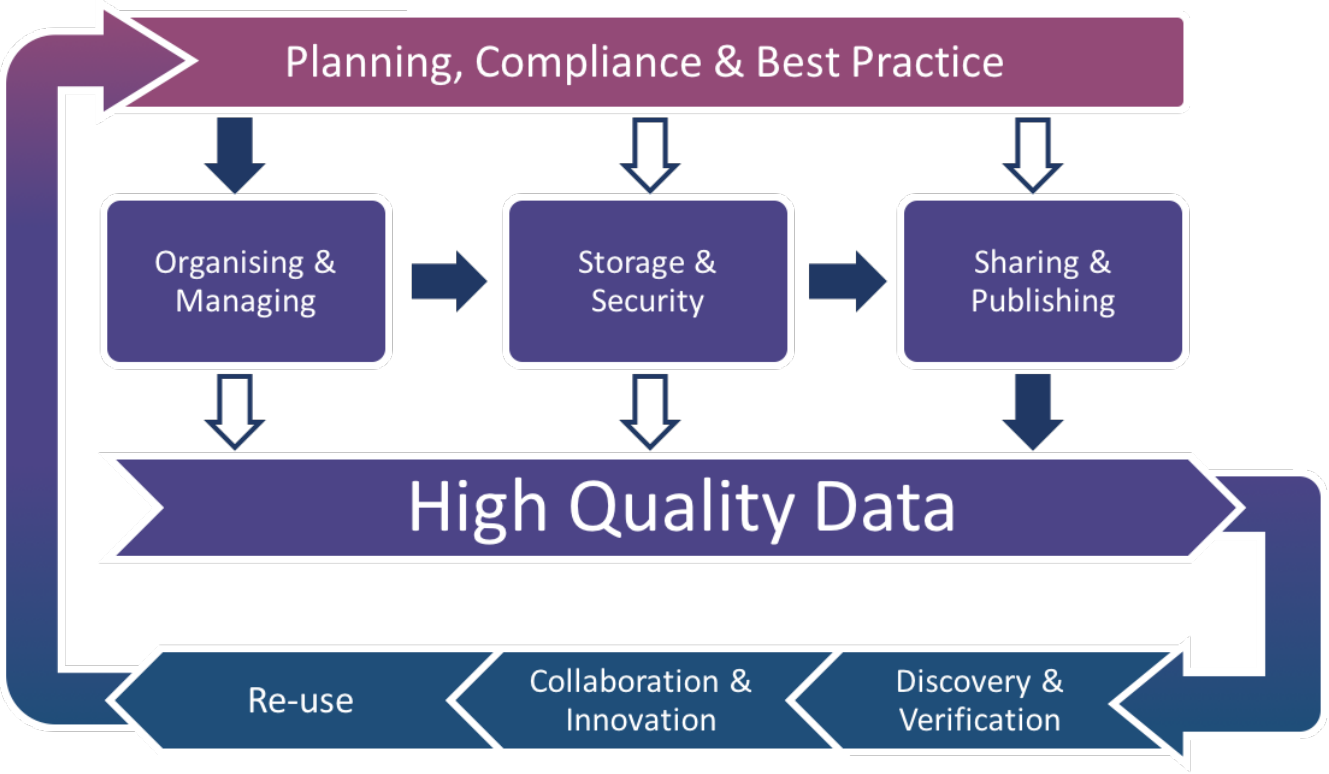


The University of Queensland Library

# UQ and SMP Data Storage

THE UNIVERSITY OF QUEENSLAND AUSTRALIA

- `nas02.storage.uq.edu.au` used for windows (10GB staff) H: drive for homes, S: Operations SCI\SMP, Y:\ SCI
- `nas05.storage.uq.edu.au` has smp-comp0x /home dirs mounted, SMP's linux desktops, also I:\ drive
- SMP has 2 NAS (Network Attached storage) systems:
- `smp-data01.smp.uq.edu.au` (11TB+) which is used for projects and also /scratch (talk to supervisor for access)
- `smp-data02.smp.uq.edu.au` – not sure...
- Have warranty on SMP systems – 2019+
- New 50TB for 'getafix' cluster
- Other networked storage systems on clusters

# Data – remote access

- I use https://cloudstor.aarnet.edu.au/ … 1 TB quota… install via software centre (uses Owncloud platform)
- RCC has MeDiCI https://rcc.uq.edu.au/data-storage
- UQ Research Data Manager has been rolled out

  https://guides.library.uq.edu.au/research-data-management

  Uses Nextcloud platform system for sharing
  Has 1TB quota for each research project… supervisor sets up and students can access
- 1 TB Microsoft One drive for each staff member: https://its.uq.edu.au/services-and-guides/software-and-online-tools/software/microsoft-software/onedrive-personal-file-storage/onedrive-online
- Not sure of Cloudstor / One drive quotas for students

# National Facilities

THE UNIVERSITY OF QUEENSLAND
AUSTRALIA

NCI
PROVIDING AUSTRALIAN RESEARCHERS WITH WORLD-CLASS HIGH-END COMPUTING SERVICES

- ARC funded collaboration between ANU/CSIRO/BOM
- http://nci.org.au/
- National Computational Infrastructure: various systems:
- large-scale peak system, Raijin (a Fujitsu Primergy cluster, entered production in June 2013)
- small specialised system, Fujin (a Fujitsu PrimeHPC FX10 system, commissioned in March 2013).
- And quite a few other systems (has special software)
- SMP staff successful grants/time (eg. ask Ben Powell)

# UQ / QCIF / RCC

- UQ Research Computing Centre under DVCR
- http://www.rcc.uq.edu.au/
- Operates various systems including for QCIF - The Queensland Cyber Infrastructure Foundation
- QRIScloud - The Queensland Node of NeCTAR / RDSI
- FlashLite - 68 nodes @ 512GB + 56Gbps and ScaleMP
- Tinaroo – 244 nodes @ 128GB (6000 cores, 5 racks!)
- Awoonga – 40 nodes @ 256GB (1032 cores)
- GPU cluster ‘wiener’ - 40 x NVIDIA Tesla V100 (Ian McCulloch has bought 3 nodes). Upgrade in 2019

https://rcc.uq.edu.au/article/2018/02/call-gpu-accelerated-imaging-researchers-use-hpc-wiener

# smp-comp0x machines

THE UNIVERSITY OF QUEENSLAND AUSTRALIA

- smp-comp01 / smp-comp02 / smp-comp03 / eg. smp-comp04.smp.uq.edu.au compute (blade) servers…

  **now decommisioned**
- smp-comp05 (CentOS6) / smp-comp06 (Ubuntu 14.04.2) 2 x Intel E5-2667 6-core 2.90 GHz (2012) 256GB RAM are **out of warranty** but about to be added to getafix

- smp-ts01.smp.uq.edu.au - Windows Remote Desktop
- Also there is a dedicate server for virtual machines, eg. smp-teaching.smp.uq.edu.au for PHYS3071/PHYS4070

# SMP Clusters- getafix/dogmatix

- `getafix` / `dogmatix` systems have over 2000 cores
- http://faculty-cluster.hpc.net.uq.edu.au/ganglia/
- http://faculty-cluster.soe.uq.edu.au/ganglia/

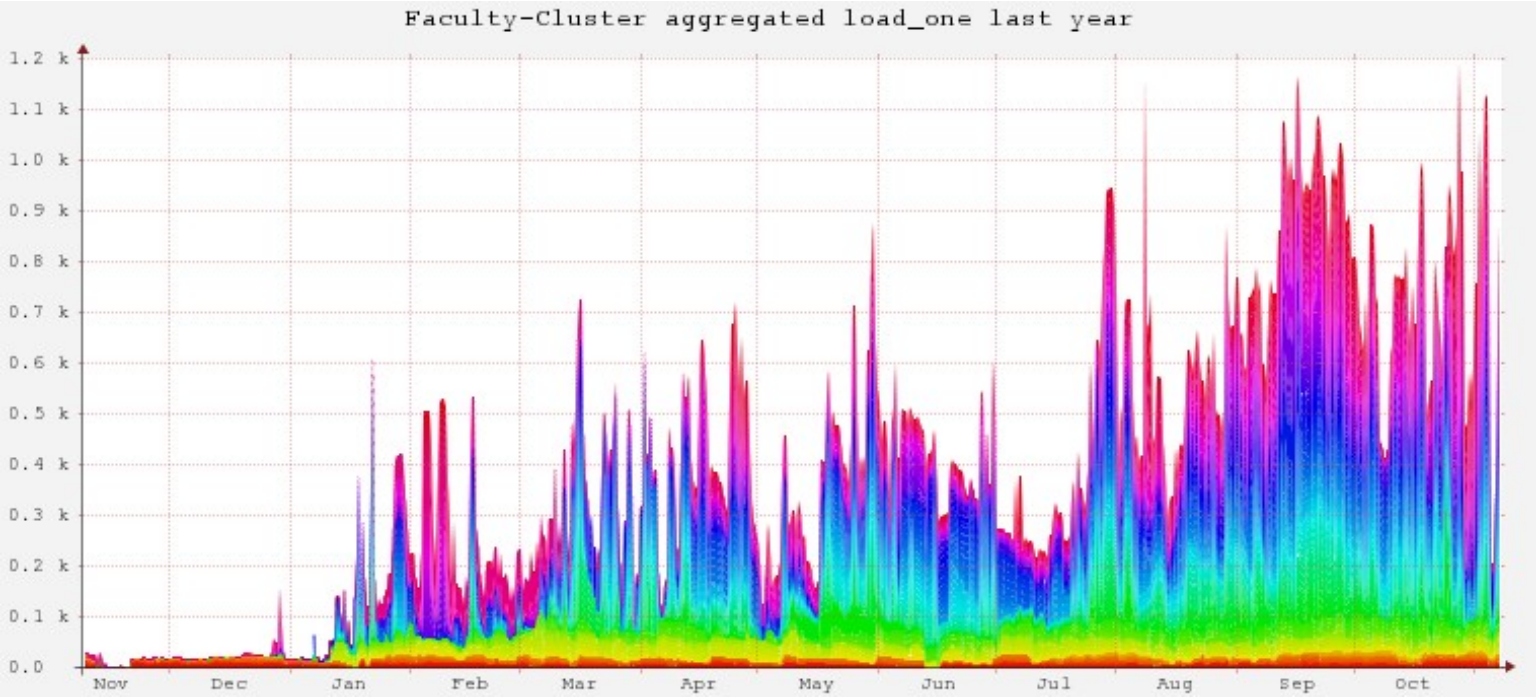

# SMP Clusters – dogmatix/ghost

THE UNIVERSITY OF QUEENSLAND AUSTRALIA

- Holger Baumgardt's Future Fellowship GPU Clusters
- Purchase starting in 2010 $117K ex-GST
- One whole Rack... Xenon Nitro T5 with **34 Tesla GPUs**
- 10 nodes 4U... 120 cores Intel X5650 @ 2.66GHz (2010)
- 7 nodes have 4 x Nvidia Tesla C1060
- 3 nodes have 2 x Nvidia Tesla C1060

- Moved into ITS data centre
- out of warranty but added to dogmatix queing system as: ghost[0-9]

# SMP Clusters– dogmatix/asterix

THE UNIVERSITY OF QUEENSLAND AUSTRALIA

- asterix nodes, purchased 2011, $193K ex-GST
- bought with UQ Major Equipment and Infrastructure Grant inc. $70K matching Science/SMP funds
  (PIs included Anthony Roberts, Matthew Davis, etc).
- One whole Rack... Xenon Nitro T5 with **40 GPUs**
- 15 nodes .. 120 cores Intel X5650 @ 2.40GHz (2010)
- Total of 10 x Nvidia Tesla C2070 (max 2 / node)
- Total of 30 x Nvidia Tesla C2050 (max 4 / node)
- Physically located in Parnell server room
- is also now added to dogmatix queuing system as
  a-0-[1,3], a-1-1, a-2-[0,2-4], a-3-1

# SMP Clusters – dogmatix

THE UNIVERSITY OF QUEENSLAND
AUSTRALIA

- Dell PowerEdge C6220 (with 35 nodes) each node 2 x Intel E5-2660 8-core 2.2GHz (2012) 196GB RAM
- Now part of dogmatix as smp-0-[0-35]
- 8 C8220X compute nodes each can fit two compute cards:
- 6 x NVIDIA Tesla K20 GPU each with 1 Kepler GK110

  2496 CUDA cores - 1.17 Tflops, 5GB GDDR5 RAM

  225 Watts

# SMP Clusters – dogmatix

THE UNIVERSITY OF QUEENSLAND
AUSTRALIA

- Cluster 11 x Intel Xeon Phi coprocessor 7120P
- Intel Many Integrated Core Architecture or Intel MIC
- x86-compatible multiprocessor architecture OpenMP, OpenCL,Cilk/Cilk Plus and Intel's Fortran, C++, MKL

  each with 61 cores at 1.238 GHz, 16GB RAM, 300 Watts

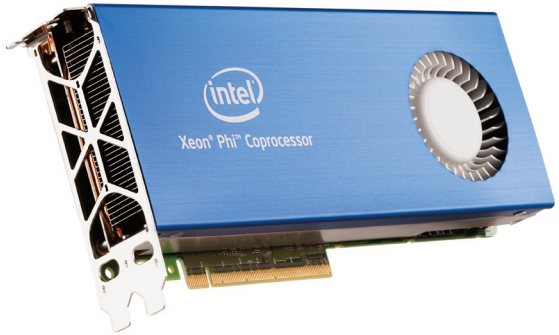

# SMP Clusters – dogmatix

THE UNIVERSITY OF QUEENSLAND
AUSTRALIA

- + 3 more Dell C8220 servers (warranty until Oct 2018)
- 13 x Dell FC430 Blade servers (smp-9-[0-12]) (until 2020) all with 256GB RAM
- virtualisation of the login node with redundant front-end
- Upgraded to slurm queuing system to handle GPUs/Phi's
- Used in COSC3500 High-performance computing and PHYS4070 Advanced Computational Physics.
- frankenstein model has worked, eg. also has 2 SAFS (School of Agriculture & Food Science) nodes with 1.5TB of RAM that is available to SMP if not being used.

# SMP Clusters – getafix

- getafix.smp.uq.edu.au is available
- Upgraded to Rocks7 (Rocks Cluster Distribution)
  ie. CentOS7 with new gcc compiler, **latest LaTeX etc**
- disk storage (2 x 50TB) consolidate existing /data* ($25K)
  slower 7.5K RPM disks than existing but RAID enabled
- 1 new FX2 Chassis (fits 8 blades)
- 2 new FC430 Blades with 512GB RAM (SMP $32K)
- 4 new FC430 Blades with 512GB RAM (Powell/Jacko $)
- Disks with ZFS (enables snapshots)
- Network packets MTU=9000 to Jumbo (up from 1500)
- New front-end/hardware/disks then migrate data over

# SMP Clusters – getafix

- `getafix` had a couple of additions late 2018
- Another 50TB
- New Nodes with Extra Intel Xeon Gold class with 512GB
- And 3 nodes with 3 x Nvidia V100 GPU's on each node. Each GPU has 32GB RAM. They are are on smp-8-[0-2]
- To see specs of various Nvidia do:

```
srun --cpus-per-task=1 --mem=4GB --nodelist=smp-8-0 --pty bash
lspci -v
```

- To load gcc and cuda you need to load a module:

```
module avail
module load …
qsub –partition=gpu --constraint=V100 --gres=gpu:1 <script>
```

# Final Comments - AWS

THE UNIVERSITY OF QUEENSLAND AUSTRALIA

- Why not use amazon AWS? https://aws.amazon.com/
- Given that there is a university support programme https://aws.amazon.com/education/awseducate/
- 1. Cost - since dogmatix is running at >50% utilisation, the price-point of AWS is probably above what we spend (esp. dogmatix with good priced network/specs/etc)
- 2. Mostly used by our **undergrads/honours/PhD** students
- 3. issues with AWS charging post-usage – difficult to track
- 4. issues with $$$ accounting back to UQ divisions
- 5. issues with cost/amount of data bandwidth to/from AWS
- Suggest people to use test accounts to see if system works for them… (**free** for 12 months per email address).