



Scientific Computing RTP



Part 1 : Computers, Access, & Support

`warwick.ac.uk/scrtp`

Presenter
Philip Grylls
Senior RSE (SCRTP)


Slides adapted from
Prof. David Quigley
Director, Scientific Computing RTP

Scientific Computing RTP


- ▶ What is the Scientific Computing RTP?
- ▶ SCRTP-managed Linux computers
- ▶ Support and common problems
- ▶ High performance computing
- ▶ Research Software Engineering



Research Technology Platforms




Advanced Bioimaging



Electron Microscopy




Spectroscopy



Polymer Characterisation



Scientific Computing



X-ray Diffraction



SCRTP Equipment and Infrastructure

- ▶ Shared infrastructure
 - Managed Linux desktop environment
 - Home & Group storage (separate to ITS H: and M: drives)
 - Hosting of servers owned by groups/departments

- ▶ Shared equipment (all accessible via a Slurm Queue)
 - Taskfarm (CPU Compute Nodes 48 Cores)
 - Avon (Nodes inc. GPU & high-speed interconnect via a Batch Queue)
 - Sulis (Available by application only)



SCRTP Desktop Linux

`warwick.ac.uk/scrtp/desktop/`

Getting access

- ▶ <https://warwick.ac.uk/scrtp/desktop/gettingstarted/>
- ▶ You need to identify an academic supervisor
 - For modules using our facilities, pick the module leader
 - For research projects, pick your research supervisor
- ▶ DO read the Acceptable Use Policy (AUP)
- ▶ DO subscribe to `scrtp-linux-user` mailing list



Remote desktop

- ▶ Access from your laptop is available using standard `ssh` terminal, free X2Go software, standard RDP or via the web (`avocado.csc.warwick.ac.uk`)
- ▶ *Shared* remote desktop hosted by `godzilla.csc.warwick.ac.uk`
 - This is a *shared resource*, there may be dozens of people logged in at any time
 - To be used for editing files, compiling code, plotting simple graphs or submitting computational jobs to the taskfarm *only*
 - You must **NOT** run significant computations of any kind on godzilla. This includes calculations within Matlab / Mathematica
 - Very strict ‘three strikes’ policy on this



Machines in the complexity complex

- ▶ Older Dell OptiPlex 9010 machines (warranty expired July 2017)

bulalo.complexity.warwick.ac.uk
caldereta.complexity.warwick.ac.uk
dinuguan.complexity.warwick.ac.uk
embutido.complexity.warwick.ac.uk
jamon.complexity.warwick.ac.uk
niliga.complexity.warwick.ac.uk
torta.complexity.warwick.ac.uk

Alive & well as of 11/10/2021

kinilaw.complexity.warwick.ac.uk
halabos.complexity.warwick.ac.uk
adobo.complexity.warwick.ac.uk
inihaw.complexity.warwick.ac.uk
rilyeno.complexity.warwick.ac.uk
sinuglaw.complexity.warwick.ac.uk

Missing in action



GPU-equipped workstations

- ▶ Newer high-end workstations

`kumeta.scrtp.warwick.ac.uk`

`keiko.scrtp.warwick.ac.uk`

`kalocsa.scrtp.warwick.ac.uk`

`kaluga.scrtp.warwick.ac.uk`

All up as of 11/10/21

- ▶ 6 core / 12 thread Xeon processors, 64GB RAM
- ▶ Quadro RTX 6000 workstation class GPU
- ▶ Ideal for GPU-accelerated analytics, machine learning etc.

Available for MathSys students to use remotely and interactively



Remote interactive servers for CDTs

- ▶ Shared (with HETSYS CDT) servers for remote interactive work
- ▶ No batch queue submission necessary. Just log in and run.
 - `hetmathsys[1-4].scrtп.warwick.ac.uk`
- ▶ 4x servers each with
 - 2x Xeon Processors, 40 cores/processor
 - 384 GB RAM
 - 10Gb/s connection to /home, /storage and the SCRTP HPC clusters
- ▶ *Shared* resource – be a good citizen.



Software

► Packages

- Part of the Linux distribution (CentOS9), or distributed via managed repositories ([EPEL](#), [CernVM-FS](#))
- Available via the MATE software menu or in the default terminal environment

► Environment modules

- Additional software imported into your terminal environment with `module load` commands
- Allows multiple software environments on one system
- Use `module spider` to search



Containers & remote notebooks

Connect to the VPN first to be able to reach your jupyter notebook.

▶ Python notebooks

```
$ module load GCC/11.3.0 OpenMPI/4.1.4 IPython/8.5.0
$ jupyter notebook --no-browser --ip=`hostname -f`
```

▶ Julia notebooks

```
$ module load GCC/11.3.0 OpenMPI/4.1.4 IPython/8.5.0
$ module load julia/1.8.2-linux-x86_64
$ export JULIA_DEPOT_PATH=~/.julia_pkgs/
$ mkdir ~/.julia_pkgs/
$ julia -e 'using Pkg ; pkg"add IJulia"'
$ jupyter notebook --no-browser --ip=`hostname -f`
```

▶ Containers

```
$ singularity pull docker://ubuntu
$ singularity shell ubuntu_latest.sif
```



Getting help

`warwick.ac.uk/scrtp/user_support/`

Scientific Computing RTP

Support & Documentation 

Documentation  | Drop In  | FAQ  | Using Bugzilla 

Top 3 common problems

- ▶ I need admin privileges to “install” software in my account
 - No you don’t, even if some random guy/developer says you do
 - warwick.ac.uk/scrip/support/support_faq/#software
 - Load the Anaconda module and go nuts!
- ▶ I can’t log in
 - Have you mucked about with your `.bashrc`, `.profile` or similar?
 - Read the desktop documentation on startup scripts
- ▶ My managed computer won’t boot
 - Have you moved it to a different network port without telling us?



Pro Tip : Software in user space

► Software distributed as source code

```
$ ./configure  
$ make  
$ sudo make install
```



```
$ ./configure --prefix=$HOME  
$ make  
$ make install
```



```
$ cmake .  
$ sudo make install
```



```
$ cmake --DCMAKE_INSTALL_PREFIX=$HOME .  
$ make install
```



► Python packages via pip

```
$ pip install somepackage
```



```
$ pip install --user somepackage
```



```
$ pip install --prefix=$HOME somepackage
```



Pro Tip : Environment variables

▶ `command not found`

- executable file not in `$PATH`
- `export PATH=$PATH:$HOME/bin:$HOME/.local/bin`

▶ `error while loading shared libraries`

- Shared object (`.so`) not in `$LD_LIBRARY_PATH`
- `export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$HOME/lib`

▶ (Python) `ImportError: No module named ...`

- Python package directory not in `$PYTHONPATH`
- `export PYTHONPATH=$PYTHONPATH:$HOME/lib/python3.6/site-packages`



Pro Tip : Using libraries

- ▶ `pkg-config` is your friend

```
$ module load GCC/8.3.0 OpenMPI/3.1.4 FFTW/3.3.8
$ gcc gaussian_fftw3.c
error: undefined reference to 'fftw_create_plan'

$ pkg-config --libs fftw3
$ -L/warwick/desktop/2018/software/../../FFTW/3.3.8/lib -lfftw3

$ gcc gaussian_fftw3.c `pkg-config --libs fftw3`
$
```

