

HPC at Warwick and Beyond

Section 1 - Intro

"The Angry Penguin", used under creative commons licence
from Swanje Hess and Jannis Pothmann.



Warwick RSE

09/2023

Introduction

What is HPC?

- High Performance Computing
- Custom bought (or built) computers for doing “hard work”
- “Hard work”
 - Anything that slows your desktop to a crawl for an hour or more
 - Things that run across multiple compute nodes, use custom hardware (graphics cards etc)

What's in the Box?

- Compute cluster
 - “Many” individual machines - **nodes**
 - These are basically good desktops
 - In different cases
 - Vital difference - **interconnects**
 - Nodes are networked together superfast
 - 32 GBit/s (real world)
 - Up to 400GBit/s (theory, more possible through multiple connections)
 - Compare 1GBit/s “fast ethernet”

Many can mean anything, but at Warwick we have a few hundred, national facilities generally a few thousand

Good desktop here can also mean anything. At the moment usually means dual socket, 8-128 cores per chip, and a full load of RAM (4-12 GB/core). Machines for particular purposes can be quite different.

Main visual difference is that these aren't in normal cases, but mounted in a rack (https://en.wikipedia.org/wiki/19-inch_rack).

Ethernet can reach 10GBit on normal Cat-6 network cable if you have the kit, but normal computers, routers etc top about around a gigabit. Your home wifi caps out at about 36 megabit normally.

What these talks will cover

- Talk 1 (this one): basics of what we have at Warwick, how to access it, and where to find help and support
- Talk 2 : using the clusters, submitting jobs, viewing and controlling your stuff
- Talk 3 : assorted bits of vital knowledge. Basics of scaling; how to understand errors; how to request packages etc; other facilities across the UK and internationally.

CSC, SCRTP, RSE, ITS, see?

- IDG - Information and Digital Group. University wide, IDG provide central IT, like the main network, email, printers etc. This is not us!
- SCRTP - Scientific Computing Research Technology Platform. This is us.
- CSC - Centre for Scientific Computing. Now replaced by the SCRTP. Still might see pages refer to CSC but don't worry.
- RSE - Research Software Engineering. Part of SCRTP, support researchers in writing and running software.

In particular, the old CSC name crops up when accessing the cluster machines: you may have to use `{name}.csc.warwick.ac.uk`

SCRTP website - warwick.ac.uk/scrtp

RSE webpage - warwick.ac.uk/rse

What can I use at Warwick?

- SCRTP Linux. Either:
 - A physical machine (4 or more core managed Linux desktop) giving you hassle free access to compilers, packages etc. Note you have to fund the machine - this is not free!
 - Taskfarm computing: intended for small (one node) jobs but perhaps many of them. Accessed via queuing system
- Compute Clusters. Avon (since 2021). Few thousand cores each, 4GB memory per core. See <https://warwick.ac.uk/research/rtp/sc/hpc>
 - Also Sulis which is an EPSRC tier-2 regional HPC centre
 - **You don't have to be an EPSRC researcher to use it!**
- Other nodes - FAT (high memory), GPU

Desktop currently running Rocky Linux 9. Offers variety of open-source stuff (gcc tools, mpi etc), Intel compiler and maths (imkl), Python 3 with the usual science libraries, some widely used packages e.g. COMSOL.

You may see or hear about Orac - the previous cluster which is NO LONGER in service (as of August) - no new signups for this are possible

Access to Systems - Theory

Who are we?

- SCRTP. Director David Quigley, some systems administrators, RSEs and shared admin staff
- We support the compute clusters and SCRTP managed Linux
- Also manage access to Sulis Tier-2 computer (see Talk 3 for more)

What is there?

- Login Node - Godzilla. Provides remote desktop, file access etc.
- Compute Clusters.
 - Avon (since 2021)
 - Sulis (since 2021)
- Task farm - cluster of machines intended for running smaller jobs than the cluster. Includes some dedicated nodes owned by specific groups
- Other nodes - FAT (high memory), OpenPOWER, GPU nodes

Dedicated nodes are SCRTP managed machines bought by and dedicated to a research group or individual. They're accessed using custom queues on the COW. If your supervisor has one, ask them about access and make sure you're using what you think you are!

Helpful Links

- SCRTP main page - <https://warwick.ac.uk/scrip>
- Warwick RSE main page - <https://warwick.ac.uk/rse>
- Documentation
 - clusters <https://docs.scrip.warwick.ac.uk/hpc.html>
 - SCRTP Linux <https://docs.scrip.warwick.ac.uk/linux.html>
- Slack channel for general support scrip-uow.slack.com
- Bugzilla. For help using the SCRTP systems, requesting software etc - <https://bugzilla.csc.warwick.ac.uk/>

Who pays for all this?

- Departments do
- <https://warwick.ac.uk/scrtp/costs/>
- Departments and Grant holders contribute to running and support. RIS can help you if writing proposals
- RSE training and support is supported by grants
- Generally your supervisor or PI should know what you can access

Mostly you shouldn't need to worry about this. Just keep in mind that somebody pays for core-hours (number of processors used times number of hours of use, the usual costing metric for compute time), and ask your supervisor/PI if you need details.

Note for any UG project students: your supervisor will know which machines and how much time you have, so if you're using SCRTP kit do be sure to ask them.

Are there Access Restrictions?

- Making a machine available from anywhere in the world opens it up to attack by bad actors
 - Some places restrict to you log in only from specific IP addresses (e.g. your office, your home internet)
- We don't for our own systems but the Tier-2 Sulis system can only be accessed from machines on campus (or other university campuses)
 - You can log in to Godzilla and then go onwards to Sulis

Comparison with Cloud

- NIST define cloud computing as
 - On-demand self-service
 - Broad network access
 - Resource pooling
 - Rapid elasticity
 - Measured service

Comparison with Cloud

- NIST define cloud computing as
 - On-demand self-service
 - Broad network access
 - Resource pooling
 - Rapid elasticity
 - **Measured service**

Comparison with Cloud

- Cloud - Market economy
 - The higher demand relative to supply the more expensive compute gets
 - This it both in general and in specific - if **you** want compute **now** expect to pay
- HPC - "Social Democracy"
 - "Money talks but it doesn't shout"
 - We'll cover this properly on Wednesday

The basic idea is that if you think that you **will** need more computation then you/your group can pay for preferential access but no-one owns resources (on the clusters) and everyone gets a chance to use them

Access to Systems - Signing Up

How do I sign up - Desktop?

- Follow the steps at <https://docs.scrtp.warwick.ac.uk/linux-pages/getaccount.html>
- This explains how to signup for access and what you need

Please do sign up for the mailing lists and read them. They often contain vital information.

How do I sign up - Avon?

- Once you've signed up for SC RTP access, you can request cluster access at <https://warwick.ac.uk/research/rtp/sc/hpc/register/>
- Sign up for Avon
- Note you have to provide a research group or department to manage the job prioritisation side of things

At the moment Avon is the only Warwick cluster that you can sign up for. If your group uses an older system you'll have to move to Avon. A newer machine will likely be procured this/next year

How do I sign up - Sulis?

- Sulis is available to all Warwick researchers
 - Designed for ensemble and high throughput computing but for Warwick researchers available for any use
- Access is by
 - Applying through EPSRC Access to HPC calls
 - Apply as a Warwick researcher (<https://warwick.ac.uk/research/rtp/sc/hpc/hpcmplus/>)
 - Exploratory Access - just email sulis@warwick.ac.uk

Sulis is an EPSRC Tier-2 HPC system intended for ensemble computing

Warwick researchers can use it for any purpose (as can other users from HPC Midlands+ universities)

It is definitely worth trying exploratory access for Sulis - if you are moving from working on your own computer then you might not need more than exploratory access

Access to Systems - Read Then Proceed

Finding the Docs

- The first place to go if you have problems is the docs - try the FAQ <https://docs.scrtp.warwick.ac.uk/linux-pages/linux-faq.html>
- You'll get access to this when you sign up for a desktop account
- Use the search function too!
- See also the Links page at the end of these slides

Bugzilla

- Bugzilla is the primary support channel
- Use if you have trouble getting access (e.g. lost ssh keys)
- Use for problems with installed packages
- Request new packages
- Access at <https://bugzilla.csc.warwick.ac.uk>
- You have to sign up for this separately to the desktop/cluster accounts - we recommend using your @warwick email address when you do

More on bug reporting in Session 3. But:

Always search Bugzilla for existing bugs before submitting your own. But do bear in mind that sometimes (e.g. in the case of the ssh keys) you're better making your own bug than commenting on somebody else's. It's also better to open a new bug with a reference to something which is related but not the same. Don't worry about the exact rules, just try and be helpful.

Sulis

- As a machine that is intended to be used by people outside Warwick Sulis has its own documentation
- <https://sulis-hpc.github.io/>
- Includes some description on using tools for ensemble and high throughput computing
- Support is still through the same channels as the Warwick clusters

Weekly Support Forums

- We have weekly drop-in support sessions, currently on Tuesdays or Wednesdays, usually in room A0.01 in Maths (Zeeman building)
 - See calendar at https://warwick.ac.uk/research/rtp/sc/user_support/research-computing-drop-in/
 - Very much including access to clusters and programming for clusters
- New 2023
 - Experimental for now - if this is useful to you come along so that we know there is interest

The Flowchart of Help

1. Yourself.
 1. For questions about access and available compute time, try your supervisor/PI
 2. Pay some attention to the mailing lists - sometimes there are outages/known issues
2. Try a search engine. See 3rd talk on error messages. Search the Docs and Bugzilla
3. Try the slack channel for help that doesn't need Bugzilla ([sctrp-uow.slack.com](https://uow.slack.com))
4. File a bug/post on Bugzilla
5. Try our weekly drop-ins for hands on support

First port of call should nearly always be your supervisor/group - they can have specific information on your project and toolchain which we don't!

A Few Essentials

Computer Terms

- Memory and Storage
 - Memory is RAM. At the moment measured in GB, or TB for a very large machine. Fast.
 - Storage is disk space. Can be of various sorts but nearly always refers to “permanent” storage. Much slower than RAM, usually much larger.
 - May hear the word “swap”. This uses disk space to get extra memory, at cost of slowness.
 - Cache usually refers to even faster, even lower capacity memory, “closer” to the processor.

Hopefully this is nothing new, but it can lead to real confusion if things aren't absolutely clear.

Memory is (usually) temporary, cleaned up when a program exits. It's specific to a program (unless using advanced techniques, or suffering e.g. the current crop of major vulnerabilities, Meltdown and Spectre).

Storage can be on all sorts of disk types and use all sorts of access protections, but it persists even when you turn off the computer (except /tmp/ directories, which are deliberately cleaned up).

A “swapping” computer is having to swap data from disk to RAM, so it usually grinds to a halt, but at least you have a chance to intervene.

Program data has to be moved from “main memory” into cache before the CPU can work with it. Cache also holds programs etc. You may hear “fits in cache” as a Very Good Thing™ because it means less moving back and forth of parts of a data set/program etc

Computer Terms

- Crashing and Hanging
 - Crashing usually means the program stops before expected, often with an error message
 - Hanging is when it does nothing, and fails to show expected signs of progress
 - Hanging is not running slow
 - Giving the wrong answer is a different sort of problem
- If in doubt when reporting problems, provide all the information you can, but not so much you swamp somebody trying to help

This one is also sort of obvious, but you'd be amazed how often error reports come in the form of "The code doesn't work" with no further clarification. Also, try to distinguish true hanging from running but very slowly. The causes and solutions are usually very different. This isn't always easy.

If possible when reporting errors, use file attachments rather than pasting long errors into a posting. And remember, if you don't know what the problem is, it's unlikely you can identify which parts of the error are relevant, so it's often better to give everything.

Wrap Up

Summary

- Taskfarm and Cluster computing is available for you to use
- All the info for signing up etc are on SCRTP pages
- The documentation contains almost everything you'll ever need to know about the machines etc
- But if you do have problems there are a few support channels

Helpful Links

- SCRTP main page - <https://warwick.ac.uk/scrip>
- Warwick RSE main page - <https://warwick.ac.uk/rse>
- Documentation
 - clusters <https://docs.scrip.warwick.ac.uk/hpc.html>
 - SCRTP Linux <https://docs.scrip.warwick.ac.uk/linux.html>
- Slack channel for general support scrip-uow.slack.com
- Bugzilla. For help using the SCRTP systems, requesting software etc - <https://bugzilla.csc.warwick.ac.uk/>

Before Next Time

- Register for a Desktop account
- Register for Cluster use
- Make sure you can access the cluster(s) you signed up for
 - See e.g. <https://docs.scrtp.warwick.ac.uk/linux-pages/remote.html>
 - Email resource.rse@warwick.ac.uk if you have trouble with ssh keys etc

[username@godzilla.csc.warwick.ac.uk](#)

[username@avon.scrtp.warwick.ac.uk](#)

[username@login.sulis.ac.uk](#)

By the way, although your viewer might think those are email addresses, they aren't. You login to a machine as {username}@{machinename}.{domain}
Note some machine names use the CSC domain.