

Course Description

- [Supercomputing](#)
 - [Introduction to Pawsey](#)
 - [Introductory Supercomputing](#)
 - [Intermediate Supercomputing](#)
 - [Optimising Serial Code](#)
 - [Parallelising Your Code](#)
 - [Developing with MPI and Open MP](#)
 - [Using Development Tools on the Cray XC40](#)
- [Data](#)
 - [Introduction to Unix](#)
 - [Introduction to pshell](#)
 - [Nimbus Research Cloud Training](#)
- [Visualisation](#)
 - [Remote visualisation](#)

Supercomputing

Introduction to Pawsey

General overview of resources and expertise offered by Pawsey, suited to all researchers. After completing the session, the attendee will:

- Understand what Pawsey is and offers
- Know some scenarios for researcher use of Pawsey services
- Know how the attendee can benefit from Pawsey services
- Know how to get access to Pawsey resource

Introductory Supercomputing

This introductory course will introduce users to using a supercomputer. Topics include: Parallel computing concepts, queuing systems and batch scripts, running jobs, job accounting, tips and etiquette. After completing the training session, the attendee will be able to:

- Understand basic parallel computing concepts and workflows
- Understand the high-level architecture of a supercomputer
- Use a basic workflow to submit a job and monitor it
- Understand a resource request and know what to expect
- Check basic usage and accounting information

Intermediate Supercomputing

This course builds on the material in Introductory Supercomputing, enabling the user to compile code and use the machine more efficiently and effectively. Topics include: the Lustre file system, compiling codes, mathematical and I/O libraries, job accounting, more complicated workflows. After completing the training session, the attendee will be able to:

- Compile and run code on a Cray XC40 supercomputer
- Understand how to get good performance out of the filesystem
- Develop and use advanced jobscripts
- Explore current and past jobs

Optimising Serial Code

- Choose an algorithm for good performance
- Choose a language for good performance
- Understand the importance of standard conformance
- Write code that can be optimised for a modern CPU
- Locate bottlenecks in a serial code and address them
- Know about common math libraries instead of do it yourself

Parallelising Your Code

- Introduction to Parallel programming
- Writing programs for shared memory architecture using OpenMP
- Writing scalable programs for distributed memory architecture using Message Passing Interface standard
- Making use of debugging and profiling tools in design and development of a parallel program.
- Writing distributed applications to maximise throughput.

Developing with MPI and Open MP

This course will cover the basics of using Message Passing Interface (MPI) and OpenMP to parallelise code. These are the most common forms of parallelisation on supercomputers. After completing the training session, the attendee will be able to:

- Understand the concepts behind parallel programming
- Have a basic knowledge of calls in MPI and OpenMP
- Modify code to use MPI and OpenMP
- Compile code and submit jobs
- Have an understanding of hybrid approaches mixing MPI and OpenMP

Using Development Tools on the Cray XC40

This course will cover the use of debugging, profiling and parallelization tools from Cray and Allinea on the Cray XC30 supercomputer. Topics include: Allinea DDT (debugger), Allinea MAP (profiler), CrayPat (profiler), and Cray Reveal (parallelization tool). After completing the training session, the attendee will be able to:

- Launch DDT and use it to debug a code
- Profile a code with MAP, and go back and forth between MAP profiling and DDT debugging
- Use CrayPat to do an in-depth profiling of a code, and analyze the performance of their code
- Use Cray Reveal to parallelize a code with OpenMP directives

Data

Introduction to Unix

This brief hands-on [course](#) will provide researchers with a basic understanding of the Unix Shell and useful commands that will enable them to undertake common tasks at Pawsey Supercomputing Centre.

Introduction to pshell

Attendees will learn the basics of using pshell (<https://data.pawsey.org.au/tools/>) which is used to manage Pawsey data resources via the command line.

Nimbus Research Cloud Training

This training [workshop](#) will provide attendees with an introduction to the available Nimbus Research Cloud resources at the Pawsey Supercomputing Centre. At the end of the workshop attendees should feel confident in applying and experimenting with this resource. The workshop will aim to provide an introduction to using the Infrastructure as a Service - IaaS capability of the Nimbus Research Cloud.

The workshop will include:

1. The Nimbus Dashboard basics of key generation, configuration and launching a Nimbus instance.
2. Managing and Maintenance of your Nimbus instance.
3. Configuring storage to your Nimbus instance.

Visualisation

Remote visualisation