

# Introduction to High Performance Computing for Life Scientists

## Partners



## Funding



# Course Parameters

- Prerequisites
  - Course designed to allow any researcher from the computational life sciences to participate in and complete, regardless of prior experience in high performance computing.
  - Previous familiarity with the Linux command line useful but not assumed, and guidance is provided
  - No knowledge of programming expected
- Hands-on practicals form an integral part of the course
  - Use the UK national HPC service ARCHER2
  - Learn by doing, gain practical skills and insights

# Aims

Understand:

- What is HPC?
- Why do people use HPC and what do they use it for?
- Computer hardware and architectures
  - What matters for performance in scientific applications?
- Processes and threads
  - How applications run on hardware
- Parallel computing and programming models
  - How applications tackle problems in parallel

Gain experience using an HPC machine

- Dealing with common stumbling blocks

# Aims

- Know how to evaluate parallel performance of an application
  - How do you know whether you're making good use of HPC resources?
- Understand current HPC architectures
- Appreciate some of the challenges running life science pipelines / workflows on HPC systems
- Know about the UK & EU HPC landscape
- Gain an appreciation of upcoming trends in HPC
- Understand how HPC can benefit your research