

Introduction to ARCHER2 and Kelvin2 services

Reinhard J. Maurer
Department of Chemistry and Department of Physics,
University of Warwick

HPC-CONEXS provides access to HPC resources

- Access to UK National Supercomputing Service ARCHER2 (Tier 1)



- Access to NI-HPC UK Tier-2 Centre, Kelvin-2 HPC system



European Tier system of HPC

- Tier 3, local resources
- Tier 2, regional centres
- Tier 1, national centres
- Tier 0, European centres



- Since October 2021, full ARCHER2 system online
- Capability resource for very large parallel jobs
- HPE Cray EX supercomputing system
- 28 PFLOP/s, AMD EPYC Zen2 (Rome), 2.2GHz
- 5,848 nodes (each with 128 cores)
- 748,544 cores

Currently #39 in Top500 list of HPC systems (top500.org, Nov 2023)

Biggest European HPC: LUMI is #5



THE UNIVERSITY
of EDINBURGH

funded by



Engineering and
Physical Sciences
Research Council

Natural
Environment
Research Council

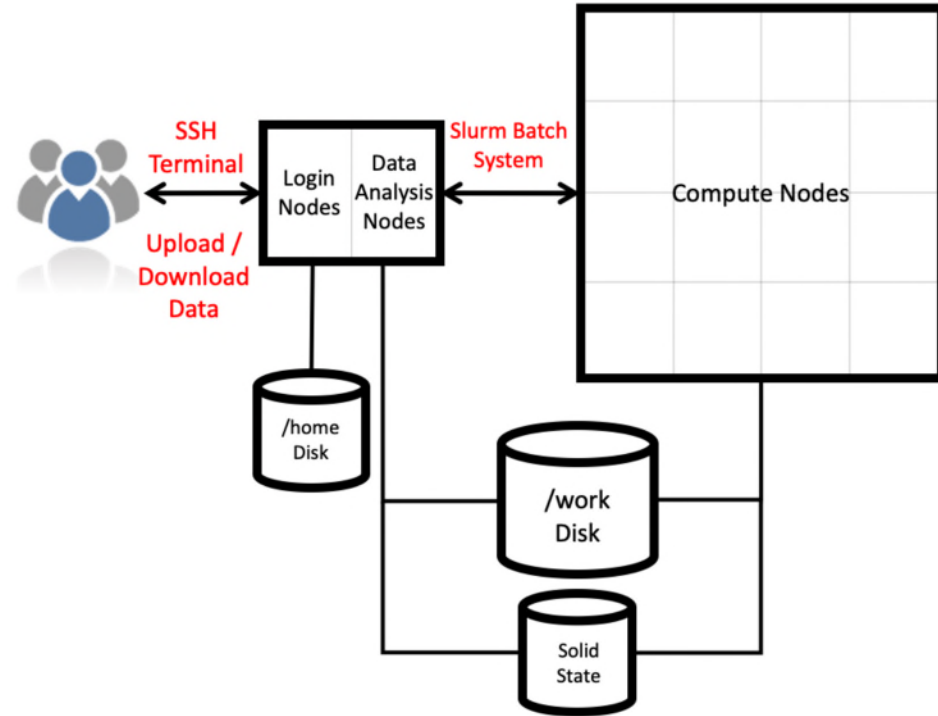
What goes into a supercomputer?



- ARCHER2 is a cluster – it's a supercomputer made up of many individual computers which we call nodes.
- A network links these nodes, turning them into a usable cluster.
- Several types of node, with different purposes:
 - Login nodes x 4
 - Compute nodes x 5,860
256 GB on a standard node, 512 GB on high memory nodes.
 - Data analysis nodes x 2
- Filesystems - you need somewhere to store files.
 - /home: 1 PB network filesystem, backed up
 - /work: 14.5 PB high performance Lustre filesystem
 - Solid state: 1.1 PB NVMe storage

The ARCHER2 layout

- Log in with SSH – land on one of the login nodes.
 - Access to both /home and /work.
- Any sort of ‘real’ work (jobs) will be run on the compute nodes or data analysis nodes.
 - Accessed and used via the Slurm batch system.
 - **Compute nodes can't see the /home filesystem.**



Software Specifications

- Operating system: HPE Cray Linux Environment (based on SLES 15)
- Scheduler: Slurm configured to be node exclusive (smallest unit of resource is a full node)
- Compilers:
 - HPE Cray Compiling Environment (CCE)
 - GNU Compiler Collection (GCC)
 - AMD Optimizing Compiler Collection (AOCC)
- Parallel libraries:
 - MPI: HPE Cray MPICH2
 - HPE Cray OpenSHMEM
 - Global Arrays Toolkit
- HPE Cray scientific and numerical libraries:
 - HPE Cray LibSci: BLAS, LAPACK, ScaLAPACK
 - FFTW 3
 - NetCDF
 - HDF5

ARCHER2 online resources

<https://www.archer2.ac.uk/>

<https://docs.archer2.ac.uk/>



ARCHER2 Training
ARCHER2 for Data Scientists
Overview of the ARCHER2 GPUs
Introduction to Modern Fortran
Message-passing programming with MPI
Intermediate Modern Fortran
Introduction to Xcompact3D
Reproducible computational environments
using containers

ARCHER2 Capability Days
Wednesday 28th February 15:00

News & Announcements

ARCHER2 Weekly Newsletter
21 Feb 2024

- Understanding parallel I/O performance through profiling, webinar, Wednesday 21st February 2024 15:00 - 16:00
- GPU eCSE software development call and associated Early Career Observer call
- ARCHER2 Celebration of Science
- ARCHER2 for Data Scientists, online, 4 March 2024 10:00 - 17:30
- Overview of the ARCHER2 GPU Development Platform, online, 12 - 13 March 2024 09:00 - 17:00 and 09:00 - 16:30
- Introduction to Modern Fortran, Durham, 3 - 4 April 2024 09:30 - 16:30 BST
- Intermediate Modern Fortran, Edinburgh, 10 - 11 April 2024 09:30 - 16:30
- Introduction to Xcompact3D, London, 11 April 2024 10:00 - 17:30
- Recently added known issues
- Upcoming ARCHER2 training

ARCHER2 Weekly Newsletter
14 Feb 2024

- Using Blender for Scientific Visualisation, webinar, Wednesday 14th February 2024 15:00 - 16:00

Upcoming Training

Message-Passing Programming with MPI
Date: Always open - self service
Location: Online

Shared Memory Programming with OpenMP
Date: Always open - self service
Location: Online

QM/MM with GROMACS + CPK
Date: Always open - self service
Location: Online

Hands-on Introduction to HPC
Date: Always open - self service
Location: Online

Modern C++ for Computational Scientists
Date: 26 - 27 February 2024 09:30 - 16:30 GMT
Location: Online

ARCHER2 Capability Days
Date: Wednesday 28th February 2024 15:00 - 16:00
Location: Online

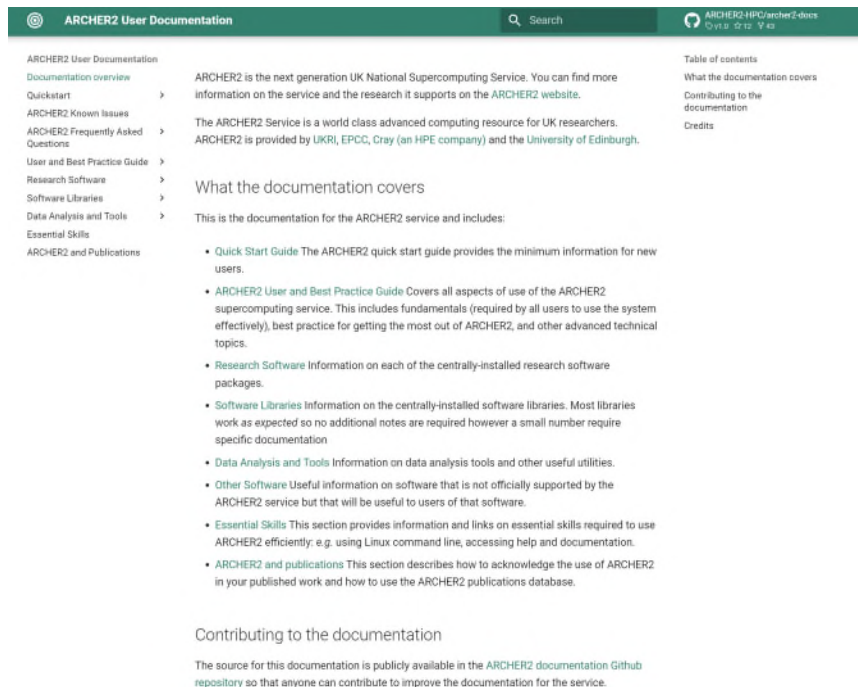
ARCHER2 for Data Scientists
Date: 4 March 2024 10:00 - 17:30

ARCHER2 is the UK National Supercomputing Service

The ARCHER2 Service is a world class advanced computing resource for UK researchers. ARCHER2 is provided by UKRI, EPCC, HPE Cray and the University of Edinburgh.



ARCHER2 Celebration of Science 2024



ARCHER2 User Documentation

ARCHER2-HPCC/archer2-docs

Search

ARCHER2 User Documentation

Documentation overview

Quickstart

ARCHER2 Known Issues

ARCHER2 Frequently Asked Questions

User and Best Practice Guide

Research Software

Software Libraries

Data Analysis and Tools

Essential Skills

ARCHER2 and Publications

ARCHER2 is the next generation UK National Supercomputing Service. You can find more information on the service and the research it supports on the ARCHER2 website.

The ARCHER2 Service is a world class advanced computing resource for UK researchers. ARCHER2 is provided by UKRI, EPCC, Cray (an HPE company) and the University of Edinburgh.

What the documentation covers

This is the documentation for the ARCHER2 service and includes:

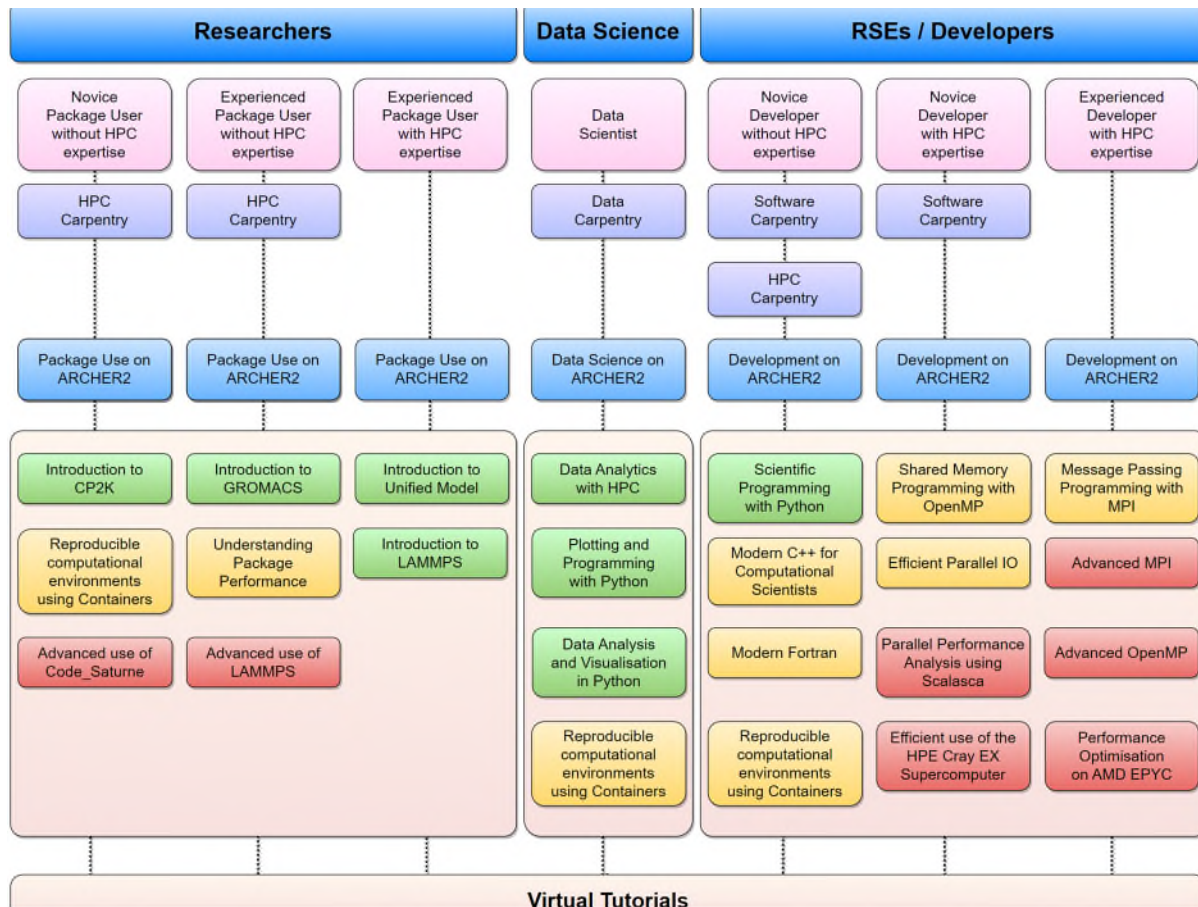
- Quick Start Guide** The ARCHER2 quick start guide provides the minimum information for new users.
- ARCHER2 User and Best Practice Guide** Covers all aspects of use of the ARCHER2 supercomputing service. This includes fundamentals (required by all users to use the system effectively), best practice for getting the most out of ARCHER2, and other advanced technical topics.
- Research Software** Information on each of the centrally-installed research software packages.
- Software Libraries** Information on the centrally-installed software libraries. Most libraries work as expected so no additional notes are required however a small number require specific documentation
- Data Analysis and Tools** Information on data analysis tools and other useful utilities.
- Other Software** Useful information on software that is not officially supported by the ARCHER2 service but that will be useful to users of that software.
- Essential Skills** This section provides information and links on essential skills required to use ARCHER2 efficiently; e.g. using Linux command line, accessing help and documentation.
- ARCHER2 and publications** This section describes how to acknowledge the use of ARCHER2 in your published work and how to use the ARCHER2 publications database.

Contributing to the documentation

The source for this documentation is publicly available in the ARCHER2 documentation Github repository so that anyone can contribute to improve the documentation for the service.

ARCHER2 Training Courses

<https://www.archer2.ac.uk/training/courses/>

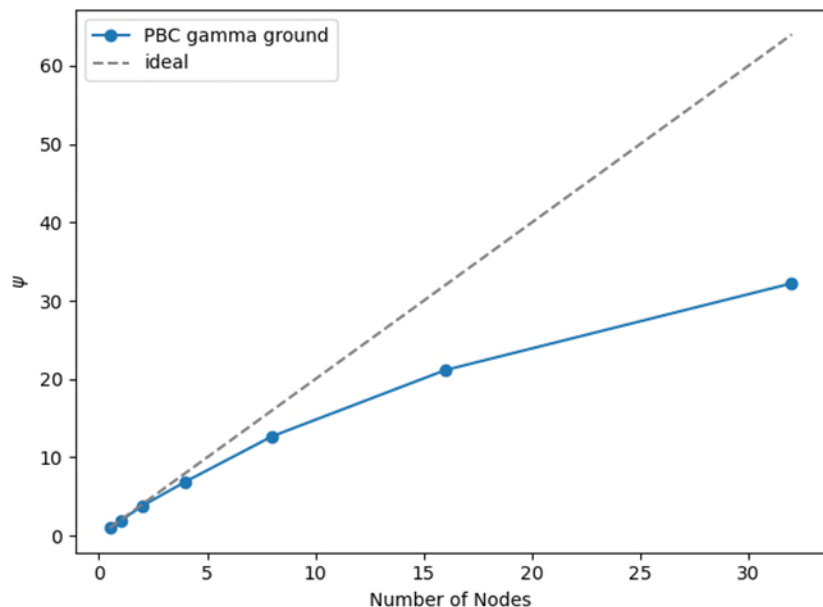


**ARCHER2
Youtube Channel**

ARCHER2 CPU time accounting

- ▶ 1CU = 128 CPUh, 1h on 128 cores
- ▶ Typical **6-month allocation requests** would be between 1000 – 50,000 CU
- ▶ The request needs to briefly describe the scientific objectives and provide a breakdown and justification of the requested resources.
- ▶ For an ARCHER2 use case, your code should show ideal scaling to 128 cores and beyond (ideally more than 1 node)
- ▶ Unused CUs are lost at the end of the 6 month period
- ▶ **Currently applications are rolling, ie. accepted at any time**

Do I have a use case for ARCHER2?



Speedup scaling graph of FHI-aims

peroxide-terminated diamond supercell system with 288 atoms in periodic boundary conditions for a ground state calculation at the gamma-point with up to 4096 CPU cores on Archer2.

How to get access

<https://safe.epcc.ed.ac.uk/>



SAFE
Service Administration by EPCC

SAFE website Login

Welcome to the [SAFE](#) website. Through the website, you can apply for an account on our high-performance computing systems, and perform other administrative tasks relating to your use of our machines.

[Login with institutional credentials](#) or [Scottish Government sign in](#) or

Email: ★

Password ★

[Login](#) [Forgot password?](#)

[Create an account](#)

As part of its normal functioning when you log in the website will install a temporary session cookie that will be removed when you log off or close your browser. If you do not wish this cookie to be set, disable cookies in your browser settings.

[The DIRAC safe is here.](#)

[SAFE website guide](#) [Accessibility statement](#)



EPCC SAFE User Documentation
[Documentation overview](#)
SAFE for users
SAFE for PIs and Project Managers

EPCC SAFE Documentation

This is the documentation for the EPCC SAFE.

SAFE is an online service management system.

SAFE Guide for Individual Users

Through SAFE, individual users can request machine accounts, reset passwords, see available resources and track their usage. All users must be registered on SAFE before they can apply for their machine account.

Go to the [SAFE guide for individual users](#)

[Registering, logging in, passwords](#)

Table of contents
SAFE Guide for Individual Users
Registering, logging in, passwords
Your user account on the service machine
User Mailing Options
Tracking and managing available resources
Tracking Publications
Miscellaneous
SAFE Guide for Principal Investigators and Project Managers
Managing your allocated resources
Managing project users
Tracking your project usage

<https://epcced.github.io/safe-docs/>

SAFE website Signup

This is the SAFE website.

Registration form

This form is to sign-up for a new website account. If you already have an account then [login](#). If you have forgotten your password, then [recover your password](#).

Fields marked in **bold ★** are mandatory.

We request Gender information for demographic/diversity analysis of our users. You may select "Prefer not to specify" if you do not want to supply this information

We request CareerStage information for demographic analysis of our users.

Any SSH key you register here is a default for when new login accounts are requested. However this does not automatically mean that it will be installed when the account is cre


All information supplied is held and processed in accordance with our [Personal Data and Privacy Policy](#).

Email Address ★	<input type="text" value="name@example.com"/>
Your Nationality ★	<input type="text" value="United Kingdom"/>
Title (Mr,Mrs,Dr etc.)	<input type="text"/>
First Name ★	<input type="text"/>
Initials	<input type="text"/>
Last Name ★	<input type="text"/>
Institution for reporting ★	<input type="text" value="Anglia Ruskin University"/>
Institution	<input type="text"/>
Department	<input type="text"/>
Phone number (include International code e.g. +44 for UK)	<input type="text" value="+ followed by numbers and spaces"/>
Address Line 1	<input type="text"/>
Address Line 2	<input type="text"/>
Address Line 3	<input type="text"/>
Address Line 4	<input type="text"/>
Town/City	<input type="text"/>
Postcode	<input type="text"/>
Country	<input type="text" value="Not Selected"/>
Gender ★	<input type="text" value="Not Selected"/>
SSH Public key	<input type="text"/>
Career stage	<input type="text" value="Choose File No file chosen"/>
	<input type="text" value="Not Selected"/>

SSH key guide

I often get asked about this. Note people new to Warwick should refer to the SCRTP guide as you have to do some extra things the first time here.

- Treat all external computers/clusters as a door that has no handle
- To open these doors, you need to apply a lock (**public key**) and use a key (**private key**)



- Your private key should never leave your workstation (its private!!)

Step 1 =

Type `ssh-keygen -t rsa -b 4096` into a terminal window on your workstation

It should say "Enter file in which to save the key (/home/username/.ssh/id_rsa):",

Just hit enter to accept the default file name

Then: "Enter passphrase (empty for no passphrase):

Enter same passphrase again: ", You can leave blank for no password, making one is recommended

`id_rsa.pub` is your public key (the one you have to move)

Step 2 =

There is a specific command that should make life easier when copying to servers

`ssh-copy-id -i ~/.ssh/mykey user@host`

`mykey = id_rsa.pub`

`user@host` is for example `msejgh@orac.csc.warwick.ac.uk`

Request access to CONEXS/ARCHER2

- Once requested, the project PI and managers will receive an email asking for one of them to approve it.
- You will receive another email once your account has been created.

Apply for project membership

Please note that when you apply to join a project some of the personal data (such as your name and email address) that we hold about you will be shared with managers c member to allow them to manage their project effectively.

Project ★

e772 ▼

SAFE Login account Request

This form is for requesting new login accounts. To request additional access for an existing account, select it from the navigation menu at the top of the page

Your username will be visible to other users on the system

To request an account on this machine you must agree to the following terms and conditions <https://www.archer2.ac.uk/about/policies/tandc.html>

This machine support ssh key authentication. You can upload a public key to use here.

A SSH public key is required to use this machine.

Your default key will automatically be added if no other key is specified

Choose the same username and project as you used on ARCHER/RDF to retain access to any migrated data

Note that this machine requires the use of MFA tokens. After your account is created you will need to come back here to setup a MFA token before you will be able to access the account

Requested username ★

I accept the Terms and Conditions of Access ★

SSH public key

ssh-rsa AAXXYZ...

Choose File No file chosen

Project: e772 (HPC-CONEXS)

Code	e772
Status	Active
Description	UK High-End Computing Consortium for X-ray Spectroscopy (HPC-CONEXS)
Start Date	2022-12-15
End Date	2027-01-01
PI	Tom Penfold (tom.penfold@newcastle.ac.uk)
Grants	HEC Call Autumn 2022
Funding Body	EPSRC
Project Class	Consortia
Subject area	Materials Science
Project notification list	<ul style="list-style-type: none">• sydnee.obrien@newcastle.ac.uk• tom.penfold@newcastle.ac.uk
Frequency of automatic reports	MONTHLY
Science Description	
Frequency of alerts	NEVER

▼ Project Groups

e772-magnets	Resource Pool
	Archer2

Login to ARCHER2

To log into ARCHER2 you should use the address:

Full system

```
ssh [userID@login.archer2.ac.uk
```

The order in which you are asked for credentials depends on the system you are accessing:

Full system

You will first be prompted for the passphrase associated with your SSH key pair. Once you have entered this passphrase successfully, you will then be prompted for your machine account password. You need to enter both credentials correctly to be able to access ARCHER2.

Tip

If you previously logged into the ARCHER2 system before the major upgrade in May/June 2023 with your account you may see an error from SSH that looks like

```
#####
@   WARNING: POSSIBLE DNS SPOOFING DETECTED!   @
#####
The ECDSA host key for login.archer2.ac.uk has changed,
and the key for the corresponding IP address 193.62.216.43
has a different value. This could either mean that
DNS SPOOFING is happening or the IP address for the host
and its host key have changed at the same time.
Offending key for IP in /Users/ausser/.ssh/known_hosts:11
#####
@   WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED!   @
#####
IT IS POSSIBLE THAT SOMEONE IS DOING SOMETHING NASTY!
Someone could be eavesdropping on you right now (man-in-the-middle attack)!
It is also possible that a host key has just been changed.
The fingerprint for the ECDSA key sent by the remote host is
SHA256:UGS+LA8I46LqnD58W1WN1aUFY3uD1WfR+V8RCG09fUg.
Please contact your system administrator.
```

If you see this, you should delete the offending host key from your `~/.ssh/known_hosts` file (in the example above the offending line is line #11)

WARNING: multi-factor authentication (MFA) required

Use authenticator app on your phone



MFA Time-based one-time password

Remember, you will need to use both an SSH key and Time-based one-time password to log into ARCHER2 so you will also need to [set up your TOTP](#) before you can log into ARCHER2.

Tip

When you **first** log into ARCHER2, you will be prompted to change your initial password. This is a three step process:

1. When prompted to enter your *ldap password*: Enter the password which you *retrieve from SAFE*
2. When prompted to enter your new password: type in a new password
3. When prompted to re-enter the new password: re-enter the new password

Your password has now been changed

You will **not** use your password when logging on to ARCHER2 after the initial logon.

Open

Remove MFA-Token

Request Password Reset

View Login Account Password

Add Machine

Request Group/Budget

Logging in to ARCHER2

- On Linux, macOS and WSL, use `ssh` on the command line.
 - If you are using the standard name `id_rsa` for the SSH private key:
`ssh username@login.archer2.ac.uk`
 - You may need to specify the name of your SSH private key:
`ssh username@login.archer2.ac.uk -i /path/to/private/key`
 - Alternatively, could use SSH agent beforehand (once per restart or login):
`ssh-add /path/to/private/key`
- Otherwise on Windows, MobaXterm is generally easiest and best.
- On logging in you will need to provide your key's *passphrase* (unless the SSH agent already knows it) and your account's *password*.
- Your account's initial password is found on your account's page on the SAFE ('View login account password'). Use this when you first log in.
 - You will then be asked to change your password by entering the current one again, and then the new one twice.

Logging in...

```
(base) reini@DESKTOP-K5QED0E:/mnt/c/Users/Reinhard$ ssh rjmaurer3@login.archer2.ac.uk
Last login: Sun Feb 25 22:04:00 2024 from 2.100.12.72
#####
          @@@@@@
         @@@@   @@@
        @@@   @@@@ @@@
       @@@  @@@@ @@@ @@@
      @@@ @@@@ @@@ @@@ @@@
     @@@ @@@@ @@@ @@@ @@@ @@@
    @@@ @@@@ @@@ @@@ @@@ @@@ @@@
   @@@ @@@@ @@@ @@@ @@@ @@@ @@@ @@@
  @@@ @@@@ @@@ @@@ @@@ @@@ @@@ @@@ @@@
 @@@ @@@@ @@@ @@@ @@@ @@@ @@@ @@@ @@@ @@@
@@@@@@@@
#####
                    ARCHER2
                    https://www.archer2.ac.uk/support-access/

-   U K R I   -   E P C C   -   H P E   C r a y   -
#####
Hostname:      ln02
Distribution:  SLES 15.4 4
CPUS:          256
Memory:       515.3GB
Configured:   2024-01-23

#####
-----Welcome to ARCHER2-----
#####
```

- ARCHER2 runs Linux: the HPE Cray Linux Environment
- You'll find yourself in your \$HOME directory, in my case: /home/e772/e772/rjmaurer3/
- Your calculations should be done on your /work directory /work/e772/e772/rjmaurer3/
- Installed software is available via module environments

```
rjmaurer3@ln02:~$
```

Command	Action
module list	Show which modules are loaded at the moment.
module avail [string]	Show all modules that can be loaded <i>now</i> . (Optionally limit to modules whose names contain <i>string</i> .)
module load <modulename>	Load the <modulename> module (optionally with version).
module unload <modulename>	Unload the <modulename> module (optionally with version).
module swap <module1> <module2>	Unload <module1> and load <module2>.
module spider [string]	Show all modules. (Optionally search for <i>string</i> and display prerequisite modules to be loaded first – helpful if you can't find a module with <i>avail</i> .)

Existing modules for DFT/XPS/XAS simulations

- ▶ CP2K
- ▶ CASTEP
- ▶ FHI-aims
- ▶ Quantum Espresso
- ▶ VASP
- ▶ ORCA
- ▶ NWChem

All described on docs.archer2.ac.uk

For some, you will need to request licensed access via your SAFE account

Running parallel CASTEP jobs

The following script will run a CASTEP job using 2 nodes (256 cores). It assumes that the input files have the file stem `text_calc`.

```
#!/bin/bash

# Request 2 nodes with 128 MPI tasks per node for 20 minutes
#SBATCH --job-name=CASTEP
#SBATCH --nodes=2
#SBATCH --ntasks-per-node=128
#SBATCH --cpus-per-task=1
#SBATCH --time=00:20:00

# Replace [budget code] below with your project code (e.g. t01)
#SBATCH --account=[budget code]
#SBATCH --partition=standard
#SBATCH --qos=standard

# Ensure the cpus-per-task option is propagated to srun commands
export SRUN_CPUS_PER_TASK=$SLURM_CPUS_PER_TASK

# Load the CASTEP module, avoid any unintentional OpenMP threading by
# setting OMP_NUM_THREADS, and launch the code.
module load castep
export OMP_NUM_THREADS=1
srun --distribution=block:block --hint=nomultithread castep.mpi test_calc
```

```
rjmaurer3@ln02:~> cd
rjmaurer3@ln02:~> sbatch submit.sh
```

Questions?

Important, when publishing research that has benefited from ARCHER2 and HPC-CONEXS, please add following acknowledgement

“We acknowledge computational resources from ARCHER2 UK National Computing Service which was granted via HPC-CONEXS, the UK High-End Computing Consortium (EPSRC grant no. EP/X035514/1).”

NI-HPC Tier 2, Kelvin2



- ▶ Access via EPSRC Tier 2 route for the whole CONEXS consortium
- ▶ If you would like access, please request access via SAFE for project “Q10” and contact Tom or Sydnee
- ▶ *Ideal resource for smaller and medium-sized jobs (1-2 nodes, 64 processes per node)*
- ▶ *Offers substantial GPU resources for machine learning, data analysis, rendering tasks*

<https://www.ni-hpc.ac.uk/>

<https://ni-hpc.github.io/nihpc-documentation>

Kelvin2 Overview

Compute

Storage

Home

Scratch

Temp

McClayRDS

Kelvin2 Overview

The NI-HPC centre hosts the Kelvin-2 research cluster. Kelvin2 runs on a Linux (Centos 7) Operating system. Below we will list the resources available.

Compute

- 96 x 128 core Dell PowerEdge R6525 compute nodes with AMD EPYC 7702 dual 64-Core Processors (786GB RAM).
- 8 High memory nodes (2TB RAM).
- 32 x NVIDIA Tesla v100 GPUs in 8 nodes.
- 16 x NVIDIA Tesla A100 GPUs in 4 nodes.
- 2PB of lustre parallel file system for scratch storage.

LOGIN:

•ssh <username>@kelvin2.qub.ac.uk

Multi-factor authentication is also required

Questions?

Important, when publishing research that has benefited from Kelvin2 and HPC-CONEXS, please add following acknowledgement

“We acknowledge computing resources from the Northern Ireland High Performance Computing (NI-HPC) service funded by EPSRC (EP/T022175) and provided via the EPSRC-funded UK High-End Computing Consortium HPC-CONEXS (EP/X035514/1).”