

CP2K Summer School 2016: Instructions for running parallel jobs on ARCHER

1 Introduction

This handout contains basic instructions for how to login in to ARCHER and submit jobs to the batch system. ARCHER is the UK's National HPC Service, and consists of a 118,080 core Cray XC30 system. For more detailed instructions, please refer to the ARCHER website:

<http://www.archer.ac.uk/documentation/>

2 Connecting to ARCHER via SSH

Each course participant has a numbered guest account (from `guest02` to `guest75`) - You should have been given a paper slip with your username on it. Please ensure you use the correct account for the duration of the course. The password will be given at the start of the practical session. To log in to ARCHER you should use the following details:

Server: `login.archer.ac.uk`

Username: `guestXX`

Password: *As provided*

For example, using the OpenSSH client on Linux or Mac run:

```
ssh guestXX@login.archer.ac.uk
```

You should see a welcome message, and obtain a command prompt on the ARCHER front-end similar to the following:

```
=====
login.archer.ac.uk
-----
```

```
This is a private computing facility. Access to this system is limited to those
who have been granted access by the operating service provider on behalf of the
issuing authority and use is restricted to the purposes for which access was
granted. All access and usage are governed by the terms and conditions of access
agreed to by all registered users and are thus subject to the provisions of the
Computer Misuse Act, 1990 under which unauthorised use is a criminal offence.
```

```
If you are not authorised to use this account then you must logout immediately.
-----
```

For assistance please contact the service helpdesk at: support@archer.ac.uk

```
=====
guest40@eslogin005:~>
```

If you have any difficulty connecting please ask the course staff.

For Windows, we recommend the PuTTY SSH client, which can be downloaded from:

<http://www.chiark.greenend.org.uk/~sgtatham/putty/>

3 Running a job

3.1 Filesystems

ARCHER uses the PBS batch processing system to schedule user jobs on the back-end nodes of the system. Please do not run compute-intensive tasks (e.g. CP2K jobs) on the front-end nodes as these are shared by all users of the system.

ARCHER has two main file-systems, `/home/` and `/work/`. Upon login, you will be placed in your home directory on the `/home/` file system. This is intended for storage of files, code etc. It is not accessible to the back-end nodes where your jobs will run. `/work/` is a large, high-performance parallel file system, connected to the back-end nodes, and you should use this as a location to run your jobs from, store and access job input and output files.

To access your directory on `/work/` enter:

```
guest40@eslogin006:~> cd /work/y14/y14/guest40
guest40@eslogin006:/work/y14/y14/guest40>
```

Please note that data stored in the guest accounts will be erased shortly after the course ends. It is your responsibility to make copies of any data you wish to keep in another location before the course ends.

3.2 Module environment

Software is provided on ARCHER using the `module` command. You can see all the available modules by running `module avail`. CP2K 3.0 is available with the module:

```
guest40@eslogin006:~> module load cp2k
```

Recent development versions of CP2K 4.0 are also available:

```
guest40@eslogin006:~> module load cp2k/4.0.17033
```

3.3 Batch system

To run a job using the batch system, you must create a batch script containing the commands you wish to execute. An example script is below, please make a copy of it in your own directory:

```
guest40@eslogin006:/work/y14/y14/guest40> cat run.pbs
#!/bin/bash --login

#PBS -N test_job
#PBS -l select=1
#PBS -l walltime=0:5:0
#PBS -A y14
```

```
cd $PBS_O_WORKDIR

module load cp2k

aprun -n 24 cp2k.popt --version
```

The #PBS directives give the system information about your job, in this case that the job should be called `test_job`, that it should use 1 ARCHER node, and will run for up to 5 minutes. All jobs run during the course should use the `y14` project account. You may edit these values if you wish to run on more nodes or for longer durations, before submitting the script.

Below this are a list of commands. The `cd` line ensures that the job runs correctly on the `/work/` filesystem. A module is then loaded to give access to the CP2K software, and finally we launch the `'cp2k.popt'` executable in parallel using 24 MPI processes. ARCHER has 24 processor cores per node, so when running your jobs make sure that the number of processors (`-n 24`) is less than 24 times the number of nodes (`-l select=1`).

Once the script has been created, we can submit the job using:

```
guest40@eslogin006:/work/y14/y14/guest40> qsub -q course1 run.pbs
230228.sdb
guest40@eslogin006:/work/y14/y14/guest40> qstat -u guest40
```

sdb:

Job ID	Username	Queue	Jobname	SessID	NDS	TSK	Req'd Memory	Req'd Time	S	Elap Time
230228.sdb	guest40	standard	test_job	--	1	24	--	00:05	Q	--

The `qsub` command submits the script, and returns a unique job ID. The `qstat` command can be used to monitor the status of the job. `Q` corresponds to a job which is waiting in the queue, `R` to a job which is running, and `E` to a job which is in the process of exiting. Once the job has completed it does not appear in the list.

The `-q course1` argument refers to a Reservation ID `'course1'` which has been set up to give exclusive access to a set of 50 nodes for the duration of the course. This should reduce queue waiting times. Please be aware this is shared by all participants so avoid running many jobs at one time, or using many nodes at once. The following reservation IDs are set up for this workshop:

- Tue Aug 23rd 16:00-17:00
- Thu Aug 25th 11:15-13:00
- Thu Aug 25th 15:15-17:00
- Fri Aug 26th 11:15-13:00

Note that jobs can still be run without the reservation ID, but will go into the main ARCHER queue along with all other jobs on the system, so will likely queue for (much) longer before running.

The standard output and error files produced by the job will be located in the current directory, and are numbered by the Job ID, alongside any additional files produced by the code :

```
guest40@eslogin006:/work/y14/y14/guest40> ls -ltr
total 8
-rw-r--r-- 1 guest40 y14 164 Apr 21 10:40 run.pbs
-rw----- 1 guest40 y14 0 Apr 21 10:52 test_job.e230228
-rw----- 1 guest40 y14 1523 Apr 21 10:52 test_job.o230228
```

In this case the standard output contains only some information from the CP2K executable. However, this is enough to show that the batch job worked as expected:

```
guest40404040@eslogin006:/work/y14/y14/guest40> cat test_job.o230228
cp2k.popt --version
CP2K version 3.0
SVN source code revision svn:16521
cp2kflags: libint fftw3 libxc elpa2 smm_dnn smm libgrid has_no_shared_glibc statm_resi
```

Congratulations, you have now run a simple CP2K job on ARCHER. Please proceed with the rest of the exercises. You will need to modify your batch script as necessary.