

The University of Queensland

Research Computing Centre

Reading Guide

Reading Guide

Document Status

Questions and Answers

General

[What is HPC?](#)

[Which UQ HPC should I use?](#)

[How do I get help?](#)

[What do I do if I think something is not working?](#)

Accessing the HPC

[Do I have to use linux?](#)

[How can I get familiar with linux?](#)

[How do I obtain an account?](#)

[How do I connect to HPC to login?](#)

[What if my work requires extra resources?](#)

Storing Data

[What storage is available on HPC?](#)

[How do I transfer data to/from HPC?](#)

[How do I know how much storage I am using ?](#)

Running Work on HPC (Batch Computing)

[What is batch mode computing?](#)

[How do I use PBSPro on HPC?](#)

[Can I run graphics intensive stuff on HPC?](#)

[How can I tell if the HPC is busy today?](#)

[What is my "Account String"?](#)

[How do I ...](#)

Using Software

[What software is installed on HPC?](#)

[How do I use software on HPC?](#)

[How can I build my own software?](#)

[What about specific tools?](#)

[Can I use RStudio+R on the HPC?](#)

Some Do-s and Don't-s

[Please Do](#)

[Please Don't](#)

Document Status

This update: June 8 2020 by David Green

This information is relevant to Awoonga, FlashLite and Tinaroo HPC.

Questions and Answers

This document is intended as a reading guide that hopefully helps answer the following questions for you:

General

What is HPC?

Try a search engine or try these

- <https://insidehpc.com/hpc-basic-training/what-is-hpc/>
- <https://en.wikipedia.org/wiki/Supercomputer>
- https://en.wikipedia.org/wiki/High-throughput_computing

- <https://searchdatacenter.techtarget.com/definition/high-performance-computing-HPC>
- http://www.hpcadvisorycouncil.com/pdf/Intro_to_HPC.pdf

Which UQ HPC should I use?

Review the UQ Advanced Computing Strategy outlined at <http://rcc.uq.edu.au/hpc>

Here is a brief translation:

There are several options.

- Awoonga is for doing lots of (small footprint) jobs
- FlashLite is for jobs requiring lots (>0.25TB) RAM and/or lots (TBs) of local disk.
- Tinaroo is for jobs where that one job runs across multiple nodes simultaneously (using message passing techniques).
- Wiener is for GPU based computations including image deconvolution and machine learning.

We recommend that you start out with Awoonga.

The Awoonga, FlashLite and Tinaroo share storage infrastructure and have equivalent system setup.

Wiener is special and different and is not co-located with the other 3 HPCs.

The individual clusters are optimized for different types of workload and you should run your work on the cluster that is best suited to that workload.

How do I get help?

- Attend [HPC Training](#)
- Read [RCC User Guides](#)
- Read the linux command manual pages.
- Attend [UQ HackyHour](#) sessions
- UQ users of UQ HPC systems can/should submit support requests to rcc-support@uq.edu.au

What do I do if I think something is not working?

- Check the [RCC Active Incidents](#) and [RCC Active Changes](#) web sites for any known problems or planned disruptions.
- Wait a few minutes and try again
- If it still is not working, then submit a support request to rcc-support@uq.edu.au detailing what you can about the problem.

Accessing the HPC

Do I have to use linux?

Yes!

It can be challenging, especially as you begin, but the benefits in the long run make it worth the effort.

We have operated dedicated Windows HPC platforms in the past.

However, the utilization of those platforms, nor the ease of use, was not as great as we had expected.

There are some tricks and workarounds for Linux HPC that we demonstrate at [Intro to HPC training](#) sessions.

How can I get familiar with linux?

- Self-study materials such as the Software Carpentry
Their [The Unix Shell](#) module covers basic linux command shell skills,
- Self-study materials such as the Lynda course ware site - UQ users can login using their UQ credentials by clicking the SignIn link.
There is a comprehensive offering of linux related training available.
Their [Learn Linux Command Line Basics](#) is a good starting point.
- Undertaking an online Linux training course such as [Linux Foundation on EdX](#)
- Accessing a Linux server or workstation that may be available in your Department/School/Institute,
- Running a virtual machine on the NeCTAR research cloud, or on your office workstation using virtualisation software, such as VirtualBox or VMWare.

How do I obtain an account?

Follow the links for Awoonga, FlashLite and/or Tinaroo from the [RCC HPC page](#)

You need to ensure that you have

1. Registered on the QRIScloud portal
2. Requested a "service" as described for each cluster.

How do I connect to HPC to login?

Follow the information provided in the [Connecting to HPC User Guide](#)

What if my work requires extra resources?

If you have an issue with quota settings, or a looming deadline etc, you can apply for [Enhanced Access](#)

Storing Data

What storage is available on HPC?

Read the [Storage User Guide](#)

How do I transfer data to/from HPC?

Follow the information provided in the [Transferring Data User Guide](#)

How do I know how much storage I am using ?

Read this [section](#) of the Storage User Guide.

Running Work on HPC (Batch Computing)

To best manage workload demands on the HPC systems, work should be submitted and managed through the batch system.

What is batch mode computing?

Try a search engine, or perhaps try these

- <https://www.techopedia.com/definition/5417/batch-processing>
- https://www.ibm.com/support/knowledgecenter/zosbasics/com.ibm.zos.zconcepts/zconc_whatbatch.htm
- https://en.wikipedia.org/wiki/Batch_processing

How do I use PBSPro on HPC?

See the [PBSPro User Guide](#)

Can I run graphics intensive stuff on HPC?

Yes. See the [Graphics and Rendering User Guide](#)

How can I tell if the HPC is busy today?

See the [HPC Usage Summary Page](#)

You can also login and run qstat and qmgr commands. See the [PBSPro User Guide](#) for details.

What is my "Account String"?

See this [section](#) of PBSPro User Guide.

How do I ...

How do I submit a job?

See the `qsub` command manual page.

You can qsub by putting all the job attributes on the command line.
Most people use a job script file with #PBS lines near the top.

How do I check job status?

See the `qstat` command manual page.
The command `qstat -aw1nt` shows you enhanced information, in widescreen layout and lists the nodes that are running your jobs/job array subjobs on a single line.

How do I hold a job?

See the `qhold` command manual page.
You can also submit a job so it immediately goes on hold using the `-h` option.

How do I release a held job?

See the `qhold` command manual page.
Some holds are placed by operators or the batch system itself and they can only be released by the RCC team.

How do I alter a queued jobs resources?

See the `qalter` command manual page.
A common thing to want to do is reduce the walltime of a job `qalter -l walltime= ...`

How do I delete a job?

See the `qdel` command manual page.

Using Software

What software is installed on HPC?

See the [Available Software User Guide](#)

How do I use software on HPC?

See the [Software User Guide](#)

Please ensure that if you are running software interactively, that you do so from within an [interactive job submission](#) on the batch system.

How can I build my own software?

See the [Building Software User Guide](#)

What about specific tools?

Follow the links to the User Guides about Anaconda, Containers, HyperWorks, MATLAB, R

Can I use RStudio+R on the HPC?

Yes, but ...
you do need to be aware of a few things that will make the experience different to what you may be used to on your PC.

Points to note

- RStudio on HPC is an [X11](#) graphics application.
You will need to ensure that your computer is able to [receive and display X11 graphics](#).
Alternatively, you could use the [Tinaroo Remote Desktop service](#) to access the X11 experience via a web browser.
- Currently RStudio on HPC is provided as a [software container](#).
It is done this way because the previous HPC operating system was not compatible with RStudio.
We are exploring removing that need for it to be a software container, for ease of use.
- RStudio is treated like any other piece of non-operating-system software on the HPC.
There is a [software module](#) for RStudio that must be loaded before the software will work.
- Like most software provided as containers on UQ HPC, the software module provides useful aliases to make it

simpler to use the tool in a container environment.

For example, it will automatically connect you to your RDM collection(s) and provide short cuts to simplify the use of the software.

- It is acknowledged that RStudio is a fantastic tool for developing and testing R software.

However, it is not a good fit for the HPC environment.

To extract the maximum utility from the HPC we recommend that you develop and test your scripts in RStudio, but run your scripts en masse using batch system and non-interactive methods of storing and plotting data.

The `Rscript` or the `R CMD BATCH` commands can be used within HPC batch jobs to run your R script as part of a regular batch job.

Refer to the [Using R](#) user guide as well as the User Guide about [building additional R packages](#) (aka libraries).

All links to www2 server require that you are on a UQ network or UQ VPN connection.

Some Do-s and Don't-s

Please Do

- Remember that you share the login nodes (and the Tinaroo Remote Desktop environment) with many other users. Be mindful of that.
- If you need to run something "strenuous" interactively, then use the batch system's `qsub -I ... mechanism`. This would apply to compiling and/or testing code, pre-processing or post-processing data.
- Always utilise the local disk `$TMPDIR` location in your batch jobs unless there is a valid reason not to. Some reasons to not use `TMPDIR` might include:
 - My job runs on multiple nodes using MPI so I cannot use `TMPDIR`. (valid)
 - My job needs more space than `$TMPDIR` can provide. (valid)
 - I am not interested in making my jobs run better. (invalid)
 - Using `TMPDIR` is too hard. (invalid ... it isn't really)
- If you need to transfer a lot of small files, it is smarter to **lump them into a single tar or zip file** and transfer them as part of a larger lump of data.
- If you are running jobs that are identical except for input parameter, input file name consider using **job arrays** instead of submitting many individual jobs of an almost identical form.
- If you are running many thousands of jobs, and/or if your numerous tasks are relatively short lived, you should definitely consider using the **Nimrod portal**.

Please Don't

- Don't run processing for sustained periods on login nodes (or remote desktop nodes). Your access to resources (CPU and RAM) is capped, but you should nonetheless avoid inconveniencing other users by running work on login nodes. Use the batch system instead.
- Don't use a shell or python script to rapidly fire `qsub` commands at the batch system. Use Job Arrays or Nimrod portal (see references above) if you want to stir up a swarm of jobs!
- Don't burden the login/batch environment with `watch` command polling the batch system. In many circumstances, the output of the command line that the `watch` is watching does not change anywhere near as frequently as the `watch` command will report.
- Don't run jobs out of RDM collection storage or hammer it from within a job. That is, when you `qsub` a job script don't do so from a directory within your collection. Use `/30days` instead. If your job reads all the input data at the start, and only writes to output once at the end of the job, then you could leave your input in your collection and write your outputs directly to your collection. Avoid situations where you read and write repeatedly to the collection storage. IT IS NOT DESIGNED FOR DOING THAT!! Use `/30days` or `$TMPDIR` instead.