

# Using High Performance Computing

The new Flinders High Performance Computing (HPC) solution is called **Deep Thought**. Comprising high-end DELL hardware and services, Deep Thought is an enterprise HPC service, using a 'fair usage' model to deliver compute resources to researchers and Colleges.

The HPC runs 24 hours a day, 7 days a week, excluding scheduled maintenance windows (advertised well in advance).

## Contents

---

(1) Connecting Your Computer.....	2
Microsoft Windows Operating System.....	2
Connect to Deep Thought (HPC) using Windows .....	2
Transfer Files to Deep Thought (HPC) using Windows .....	3
macOS or Linux Operating System .....	5
Connect to Deep Thought (HPC) using macOS or Linux.....	5
Transfer Files to Deep Thought (HPC) using macOS or Linux.....	5
(2) Using Deep Thought .....	7
Working with Modules and Applications .....	7
Overview .....	7
Checking what modules are available .....	7
Checking what modules are loaded .....	7
Loading modules .....	7
Unloading modules .....	7
Installing additional modules & applications.....	7
Using Slurm to Run Jobs.....	8
Migrating from Colossus to Deep Thought.....	8
Other Resources - Using Unix/Linux Commands .....	8
Unix basics.....	8
Unix command line tips and tricks .....	8
Unix tutorial .....	8
(3) Frequently Asked Questions .....	9

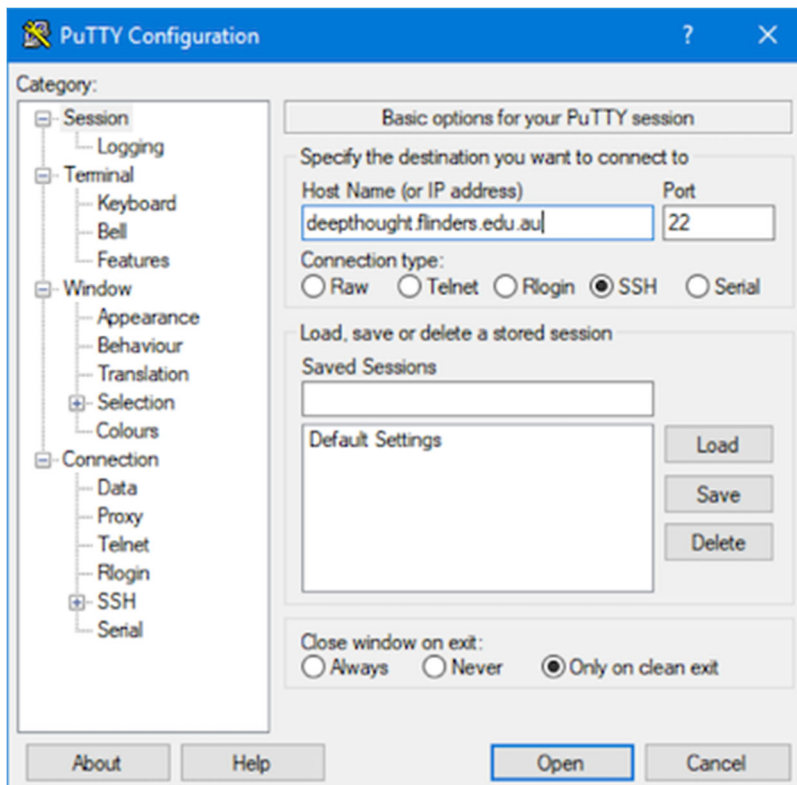
# (1) Connecting Your Computer

## Microsoft Windows Operating System

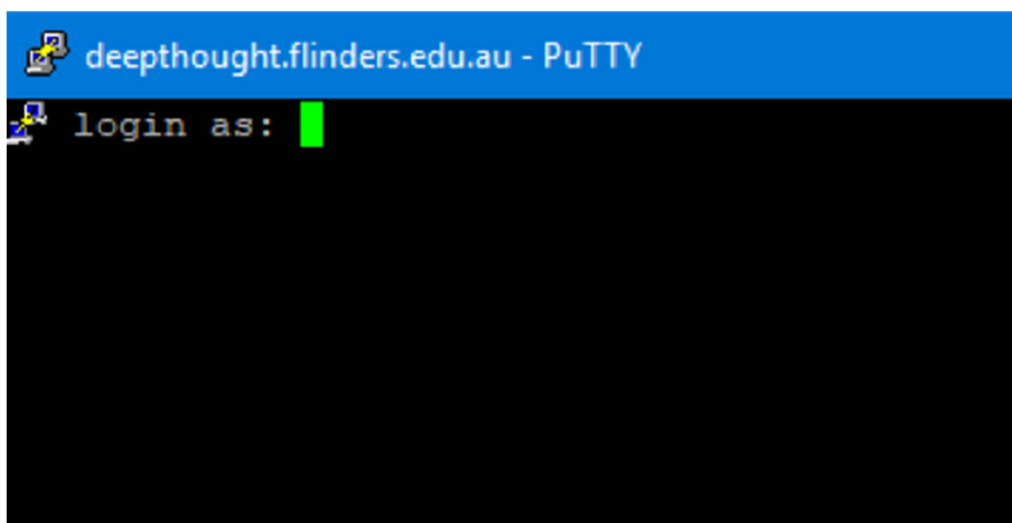
### Connect to Deep Thought (HPC) using Windows

To connect to Deep Thought an SSH application such as [PuTTY](#) is required.

Open PuTTY, enter **deephought.flinders.edu.au** as the host to connect to, and select **Open**.



A connection to Deep Thought is established. Login using your FAN username and password.

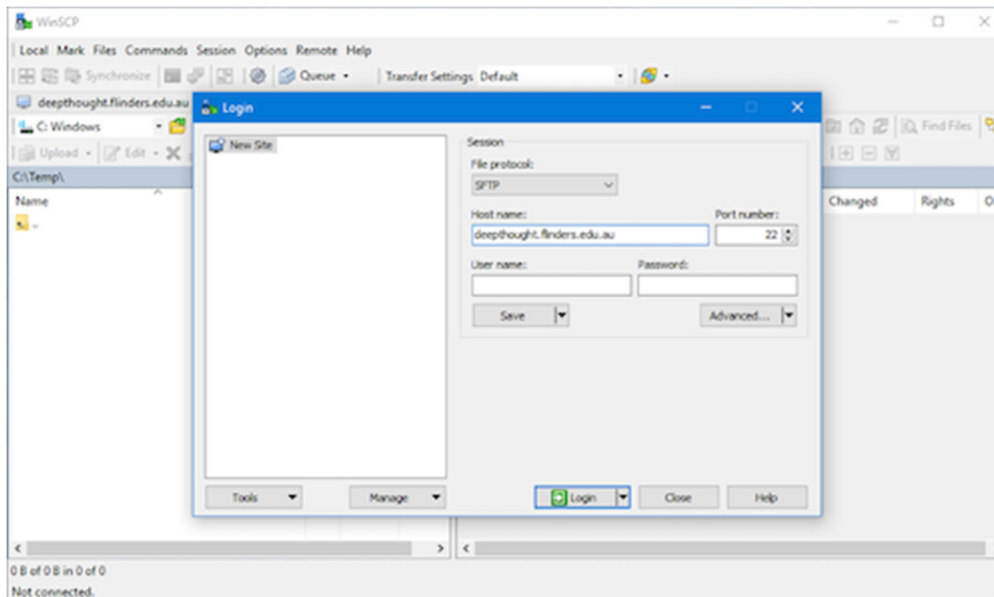


## Transfer Files to Deep Thought (HPC) using Windows

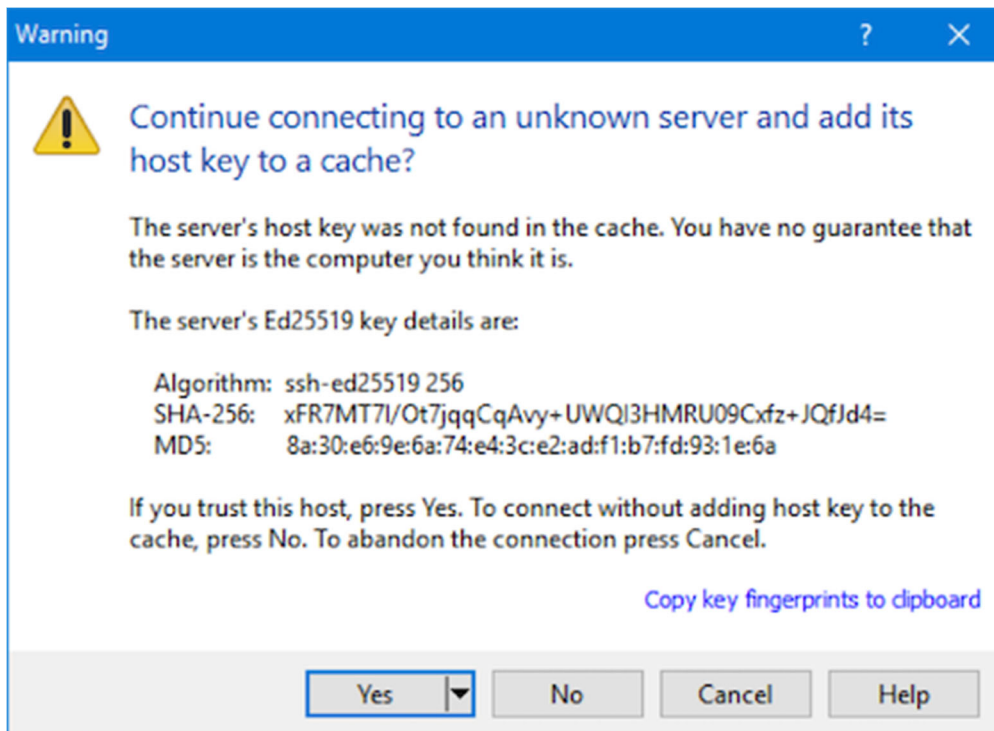
HPC clusters only have enough storage (often called '*scratch*' storage) to hold all the data and results temporarily. Long term storage is via the network. Files must be copied across to Deep Thought to be used, and should not be left there long term. Your home directory is as follows (substituting your FAN for *\$username*): `/home/$username`

To upload documents a **Secure File Transfer (SFTP)** client such as [WinSCP](#) is required.

Open WinSCP, enter **deephought.flinders.edu.au** as the host to connect to, and then **Login**.

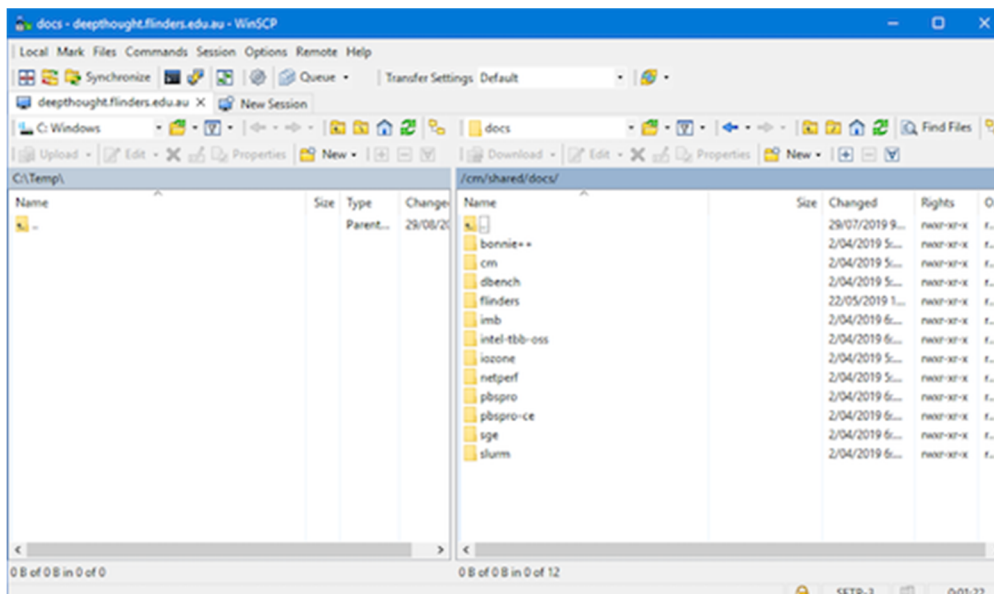


Note: The first time a connection is made the following message appears – select **Yes** to continue.



A connection to Deep Thought is then be established. Login using your FAN username and password.

You are now ready to transfer files back and forth.



## macOS or Linux Operating System

### Connect to Deep Thought (HPC) using macOS or Linux

To connect to Deep Thought, the simplest method is to open a *Terminal* window.

Enter the following, substituting your <FAN> and providing the password when prompted.

```
ssh <FAN>@deephought.flinders.edu.au
```

```
[bash-3.2$ ssh FAN@deephought.flinders.edu.au
FAN@deephought.flinders.edu.au's password:
```

A connection to Deep Thought is established.

```
bash-3.2$ ssh FAN@deephought.flinders.edu.au
FAN@deephought.flinders.edu.au's password:
Last login: Wed Aug 14 12:10:58 2019 from 10.
Welcome to Bright release      8.2

                Based on Red Hat Enterprise Linux Server 7
                               ID: #000002

Use the following commands to adjust your environment:

'module avail'          - show available modules
'module add <module>'  - adds a module to your environment for this session
'module initadd <module>' - configure module to be loaded at every login

-----
[ FAN@hpc-login01 ~]$
```

### Transfer Files to Deep Thought (HPC) using macOS or Linux

To upload documents to Deep Thought, the simplest method is to open a **Terminal** window.

Enter the following, substituting the filename, your FAN, and password where required.

```
scp <FILENAME> <FAN>@deephought.flinders.edu.au:/home/<FAN>
```

```
[bash-3.2$ scp FILENAME FAN@deephought@flinders.edu.au:/home/FAN
```

This places <FILENAME> in your home directory *on Deep Thought*.

HPC clusters only have enough storage (often called '*scratch*' storage) to hold all the data and results temporarily. Long term storage is via the network. Files must be copied across to Deep Thought to be used, and should not be left there long term. Your home directory is as follows (substituting your FAN for *\$username*): */home/\$username*

To transfer files *back* from Deep Thought to your computer, you need to know the name of your computer on the network, or its IP address. You also need to know your username – if you are using a Flinders University computer, it is your FAN. If it is a private computer, you should use the login name you use on that computer.

In the directory on Deep Thought that contains the file(s) you wish to transfer files, enter the following, making the required substitutions where required.

```
scp FILENAME <USERNAME@COMPUTER_NAME>:/home/<username>
```

*or*

```
scp FILENAME <USERNAME@COMPUTER_IP_ADDRESS>:/home/<username>
```

```
@hpc-login01 ~]$ scp FILENAME USERNAME@COMPUTER_NAME or USERNAME@COMPUTER_IP_ADDRESS
```

## (2) Using Deep Thought

---

### Working with Modules and Applications

#### Overview

The system uses the [\*\*\*lmod \(Load MODules\) system\*\*\*](#) to load and unload applications from the command line. Any modules you use frequently can be loaded on login using your [\*\*\*.bash profile\*\*\*](#) file; modules required for a job should be automated in your **Slurm** script.

#### Checking what modules are available

Use the '*avail*' option of the **module** command.

On the command line type: **module avail**

#### Checking what modules are loaded

Use the '*list*' option of the **module** command.

On the command line type: **module list**

#### Loading modules

Use the '*add*' option of the **module** command, followed by the software package you wish to load.

On the command line type (for example): **module add python2**

#### Unloading modules

Use the '*unload*' option of the **module** command, followed by the software package you wish to unload.

In the command line type (for example): **module unload python2**

#### Installing additional modules & applications

If you do not see a module listed for the application that you wish to run *and make available to other users as well*, contact Digital Research Services to have it installed for all users. You can request this via an [assyst request](#).

If the application is *just for you*, you can install it to your home directory, following the install instructions of that application.

**NOTE: Any software used must be appropriately licensed.**

If there are issues installing software then raise a support request via an [assyst request](#).

## Using Slurm to Run Jobs

[Using Slurm](#) has Slurm commands and resources, including an example script designed for Deep Thought.

## Migrating from Colossus to Deep Thought

[Migrating from Colossus HPC to Deep Thought HPC](#) provides specifics on migrating workload management from SGE to Slurm.

## Other Resources - Using Unix/Linux Commands

There are some useful resources about Unix and its command line available from Harvard University.

### Unix basics

A useful summary of the Unix command line for beginners is available at: [Basic Unix Workshop](#).

### Unix command line tips and tricks

A more advanced introduction to the Unix command line is available at: [Unix Command Line Tips and Tricks](#).

### Unix tutorial

An in-depth introduction to Unix and how it works is available at: [Introduction to UNIX and the Command Line Interface](#).



## (3) Frequently Asked Questions

---

### *What are the technical specifications for the Flinders HPC, Deep Thought?*

The hardware specifications are:

- 16 x 256GB (4TB) of RAM
- Combined Peak Performance of 44 TeraFLOPS (Compute & GPU)
- 1024 AMD x86 CPU cores at 2.0GHz across 14 standard compute and 2 data science nodes
- 5 NVIDIA Tesla V100 32Gb GPUs (2 per data science node and 1 per login node)
- 100 TB of usable storage via the Dell EMC Network Storage reference architecture
- A high-speed, ethernet network fabric able to move data across each building block at 25Gbps with various features to ensure reliability and traffic flow.

### *Can I obtain a private node for my HPC work?*

Yes, at a cost; the following is the process:

1. The researcher contacts DRS with their requirements.
2. DRS assesses the need, clarifies any unknowns and works with Digital Platform Operations (DPO) to obtain quotes to offer the service.
3. DRS provides the quotes back to the researcher who seeks funding and approval from their Research Supervisor, College Dean of Research, or other funding approver.
4. The researcher provides the approvals to DRS.
5. DRS arranges the purchase, configuration, and installation of the private node and notifies the researcher when the node is ready for use.

### *How and when are jobs scheduled?*

Jobs run 24 hours a day and are allocated to College-specific queues, or on private nodes if they have been acquired (funded and deployed).

A job that is at the top of the queue will run first and depending on resource allocation, may need to complete prior to the next job being run.

### *Can I access Deep Thought from offsite or overseas?*

Yes; you can connect to Deep Thought from external networks anywhere in the world via a VPN.

A VPN is not required to access Deep Thought when logging in via a Flinders network (eg on campus).

### *What if I need additional software? Who pays for that?*

You can purchase additional software by seeking funding approval from your Research Supervisor, Dean of Research, or grant holder.

Software purchasing should be funnelled through an [assyst software request](#).

### *How much notice for I get for a planned outage?*

You are provided with two (2) weeks' notice of planned outages. This is to allow you to finalise the running of jobs prior to the outage.

### *How does HPC work?*

HPC is a group (or cluster) of many server computers (or nodes) joined together.

The nodes in each cluster use high-speed networking and have higher than normal amounts of memory.

The cluster is joined together to appear to be one single system using special software. The computing resource can then be divided and allocated according to researchers' requirements using special scheduling and management software.

The head node controls the networked nodes using the special software.

HPC clusters only have enough storage (often called 'scratch' storage) to hold all the data and results the researchers have used temporarily; long term storage is accessed via the network. HPC alone should not be relied on for long term (archival) storage.

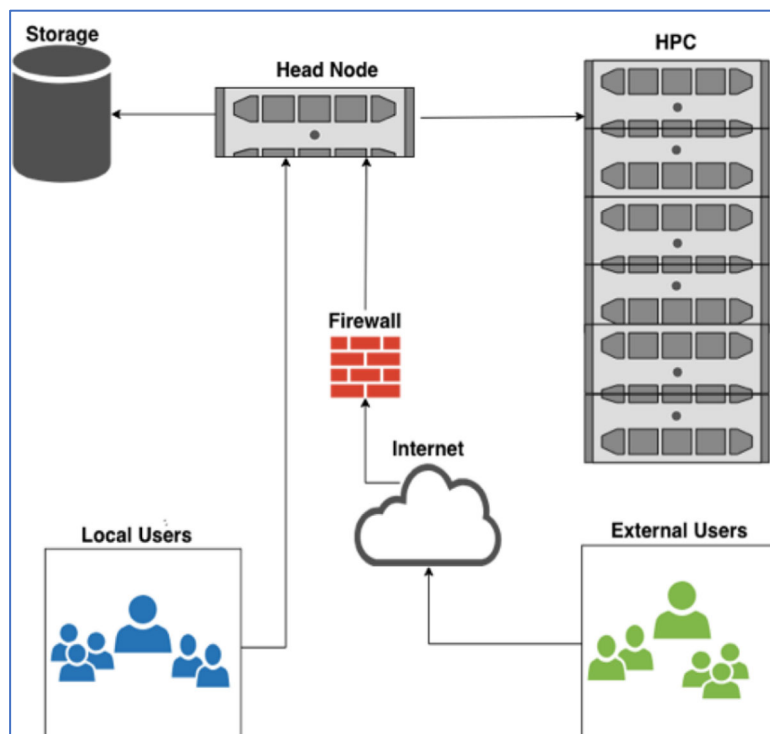


Figure 1. How HPC works