

New User Seminar



- This webinar only covers topics pertaining to our newest and largest system, **Graham**.
- For the introduction to our legacy systems (Orca etc.), please check the following recorded webinar:

“SHARCNet New User Seminar for Legacy Systems”

available on our youtube channel,


<http://youtube.sharcnet.ca>

(or read it online:

https://www.sharcnet.ca/help/index.php/Getting_Started_with_SHARCNET)

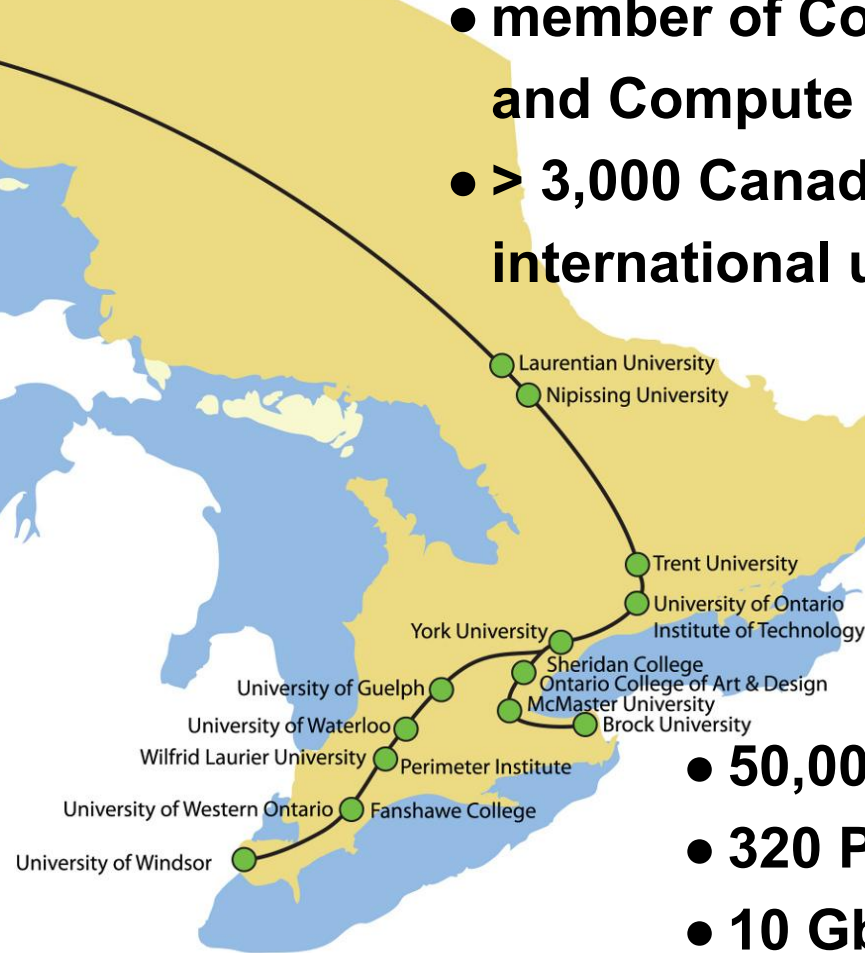
What is SHARCNET?

Consortium of 18 Ontario institutions
providing advanced computing
resources and support...



- member of **Compute Canada** and **Compute Ontario**
- **> 3,000 Canadian and international users**

Shared
Hierarchical
Academic
Research
Computing
NETwork



- **50,000 CPU cores**
- **320 P100 GPUs**
- **10 Gb/s network**

The SHARCNET web site (www.sharcnet.ca) provides extensive information about our systems and software.

- **User-editable help wiki**

- Help pages, tutorials, FAQ: [Support > Wiki](#)
- Software documentation: [Facilities > Software](#)

- **System status**

- System notices, present status: [Facilities > Systems](#)
- System notices are also sent via email and posted on RSS

- **Ticketing system**

- Online access (requires login): [Support > Tickets](#)
- Or send an email to help@sharcnet.ca

Compute Canada web site (docs.computecanada.ca) contains a large collection of help pages for the national systems (Graham and Cedar).

- **How-to guides**

- [Getting Started with the new National Systems \(mini-webinar series\)](#)
- Detailed help pages on submitting jobs, software etc.

- **Compute Canada's problem tracking system**

- Email to support@computecanada.ca

- **Graham related issues**
 - Check both SHARCNET's and Compute Canada's help pages
 - To submit a ticket, use either SHARCNET's or Compute Canada's ticketing system

- **Help for legacy systems (orca etc)**
 - Use SHARCNET's help pages and ticketing system.

- **Systems**

- Clusters, Cloud facilities, Visualization systems

- **Operating Systems**

- Linux

- **Languages**

- Fortran, C/C++, Java, MATLAB, Python, etc.

- **Key Parallel Development Support**

- MPI, pthreads, OpenMP, CUDA, OpenACC, OpenCL

- **Software Modules**

- select pre-built and configured software, as well as versions, with the *module* command

- **Batch Scheduling**

- SLURM scheduler

Essentials: Access to SHARCNET

- All SHARCNET systems are only accessible via *secure shell* (ssh):

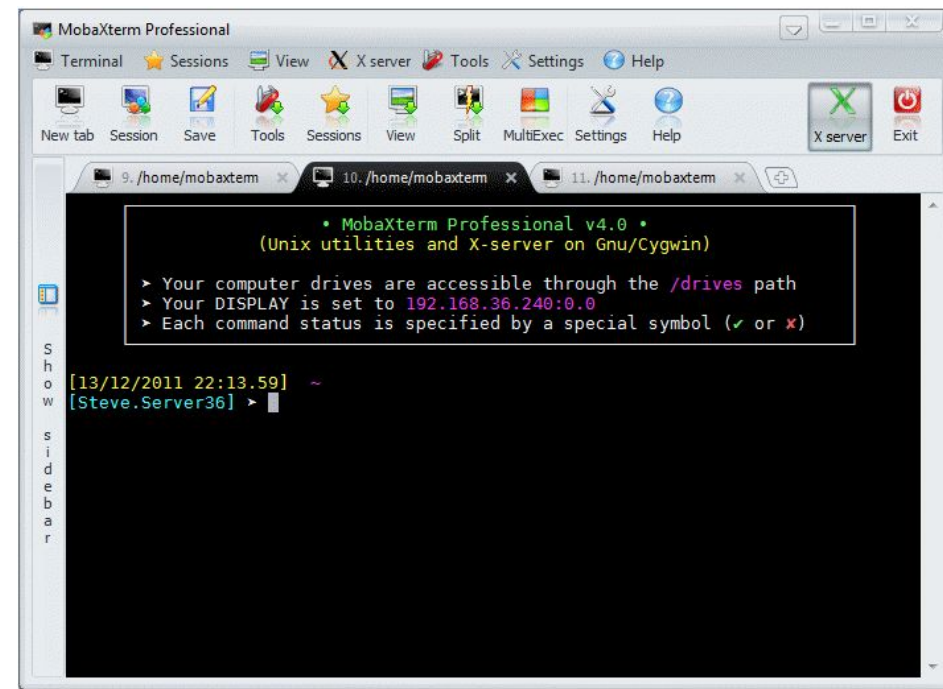
```
$ ssh user@graham.sharcnet.ca
```

(Use Compute Canada credentials to login to Graham.)

- We recommend authenticating using an *ssh key agent*
 - See the [SSH page](#) in our help wiki for details

File transfer programs:

- Unix / Mac
 - scp or sftp, rsync
- Windows
 - MobaXterm
- Any OS (from a browser)
 - Globus



Essentials: File systems

File system	Quotas	Back ed up?	Purged?	Available by Default?	Mounted on Compute Nodes?
Home Space /home	50 GB, 500 K files	Yes	No	Yes	Yes
Scratch Space /scratch	20 TB and 1000k files per user, 100 TB and 10M files per group	No	Yes, all files older than a certain number of days	Yes	Yes
Project Space /project	1TB for each cluster per user, can request increase to 10TB 500k files per user	Yes	No	Yes	Yes
Nearline Space	5 TB per group	No	No	No	No

- Number of CPU cores: 33,448
- Number of nodes: 1043
- 32 cpu cores per node
- Between 128 and 3072 GB of RAM per node
- Number of NVIDIA P100 GPUs: 320
- Networking: EDR (cpu nodes) and FDR (GPU nodes) InfiniBand

- All significant work is submitted to the system as a *job*
- Jobs are run in *batch mode* via a job scheduling system
 - enforces policies to promote fair and efficient use
- Jobs are submitted using the ***sbatch*** command
- The only required job script options are **-t** (run time) and **-A** (account)
 - RAM can be requested via **--mem** (serial / multithreaded jobs) or **--mem-per-cpu** (MPI jobs; per rank) options
 - For parallel jobs one has to specify the number of cpu cores (**-c** for multithreaded, **-n** for MPI)
- ***squeue*** : lists the status of submitted jobs
- ***sacct*** : shows details for recent jobs
- ***scancel***: kills jobs

- do not run significant programs on login nodes or directly on compute nodes
 - submit them as a job to the scheduler
- do not specify a maximum job run time (say, 7 days), or more memory than required for your program
 - pick an appropriate value, eg. 150% of the measured/expected run time or memory per processor
- do not create millions of tiny files, or large amounts (> GB) of uncompressed (eg. ASCII) output
 - aggregate files with tar, use binary or compressed file formats