New User Seminar

SHARCNET™

Graham vs legacy systems

- This webinar only covers topics pertaining to **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

Topics

- **SHARCNET**
- Where to look for information and get help
- Essentials
  - What are available
  - How to connect to graham
  - How to transfer files
  - How to compile programs
  - How to submit jobs
  - Manage files
- Do’s and don’t do’s
- Q & A

What is SHARCNET?

A consortium of 18 Ontario institutions providing advanced computing resources and support...

Shared
Hierarchical
Academic
Research
Computing
NETwork

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

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

Topics

- SHARCNET
- **Where to look for information and get help**
- Essentials
  - What are available
  - How to connect to graham
  - How to transfer files
  - How to compile programs
  - How to submit jobs
  - Manage files
- Do’s and don’t do’s
- Q & A

Getting Help: SHARCNET web portal

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

Getting Help: Compute Canada site

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

Getting Help: SHARCNET or CC?

- **Graham related issues**
  - Check both SHARCNET's and Compute Canada's help pages
  - Submit a ticket to Compute Canada at support@computecanada.ca

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

- **Cedar and Niagara related issues**
  - Use Compute Canada's help pages and ticketing system.

Topics

- SHARCNET
- Where to look for information and get help
- **Essentials**
  - What is available
  - How to connect to graham
  - How to transfer files
  - How to compile programs
  - How to submit jobs
  - Manage files
- Do's and don't do's
- Q & A

Essentials: Computing Environment

- **Systems**
  - Clusters, Cloud facilities
- **Operating Systems**
  - Linux (64-bit CentOS)
- **Languages**
  - C/C++, Fortran, Matlab/Octave, Python, R, Java, 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

Connecting to clusters

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

  $ ssh user@graham.computecanada.ca

  Use your 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

Connection and file transfer programs

- Unix / Mac
  - scp or sftp, rsync
- Windows
  - MobaXterm
  - Cygwin (A full Unix like suite)
- Any OS (from a browser)
  - Globus

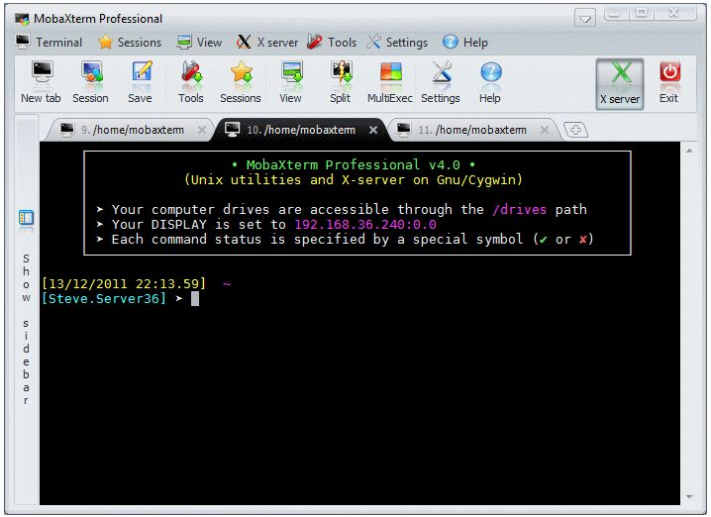

Essentials: File systems

File system	Quotas	Backed up?	Purged?	Available by Default?	Mounted on Compute Nodes?
Home Space /home	50 GB and 0.5M files per user	Yes	No	Yes	Yes
Scratch Space /scratch	20 TB and 1M files per user, can request increase to 100 TB	No	Yes, all files older than a certain number of days	Yes	Yes
Project Space /project	1 TB and 5M files per group, can request increase to 10 TB	Yes	No	Yes	Yes
Nearline Space	5 TB per group	No	No	No	No

- Run **quota** command on Graham/Cedar to find out if you are approaching or over the disk quota.

Essentials: Graham cluster

- 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 and cloud nodes) InfiniBand

Essentials: Managing jobs with SLURM

- All significant work shall be submitted to the system as a *job*, run in batch mode.
- Jobs are submitted using the ***sbatch*** command with a script

```
#!/bin/bash
#SBATCH --time=0-00:05          # Run time limit (DD-HH:MM)
#SBATCH --account=def-bge
#SBATCH --ntasks=32             # Number of MPI processes, default 1
#SBATCH --cpus-per-task=32      # Normally defined for threaded jobs
#SBATCH --gres=gpu:4            # request GPU "generic resource", 4 on Cedar, 2 on Graham
#SBATCH --mem=1024M
#SBATCH --mem-per-cpu=1024M     # memory; default unit is megabytes
#SBATCH --job-name=hello        # Optional, for user's reference
#SBATCH --output=%x-%J.log      # You give any name
./myprog                        # Replace with mpiexec ./myprog or srun ./myprog for MPI jobs
```

- **squeue**: to list the status of submitted jobs.
- **sacct**: to show details of recent jobs.
- **scancel**: to kill jobs.
- Use command **man *scommand*** to see details.

Why is my job not starting?

- There may be multiple reasons
- Graham/Cedar are very busy clusters, with <20% of the cycles available to non-RAC jobs. ***Tip: consider applying for RAC.***
- Requesting much more resources (runtime, CPU cores, memory) than what is actually needed will result in a longer queue wait time, for no good reason. ***Tip: request only what the job needs, with a bit of leeway***.
- If your job uses multiples of 32 cpu cores, sometimes the queue wait time can be much shorter if you do a by-node reservation, instead of the default by-core one. ***Tip: use --nodes=N and --ntasks-per-node=32 sbatch arguments to request the by-node reservation***.

Topics

- SHARCNET
- Where to look for information and get help
- Essentials
  - What are available
  - How to connect to graham
  - How to transfer files
  - How to compile programs
  - How to submit jobs
  - Manage files
- **Do's and don't do's**
- Q & A

Common mistakes to avoid

- Do not run significant programs on login nodes, nor run programs directly on compute nodes.
- Do not specify a maximum job run time blindly (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

Q&A

Submitting tickets to Compute Canada support@computecanada.ca for

- Graham/Cedar related
- RAC allocations
- Accounts

And rest to SHARCNET at help@sharcnet.ca