


New User Seminar



- **SHARCNET**
- Where to look for information and get help
- Essentials
 - What are available
 - How to connect to graham
 - How to manage disk storage
 - How to run your programs - *jobs*
- Common mistakes to avoid
- Q & A

What is SHARCNET?

A consortium of 19 Ontario institutions providing advanced computing resources and support...



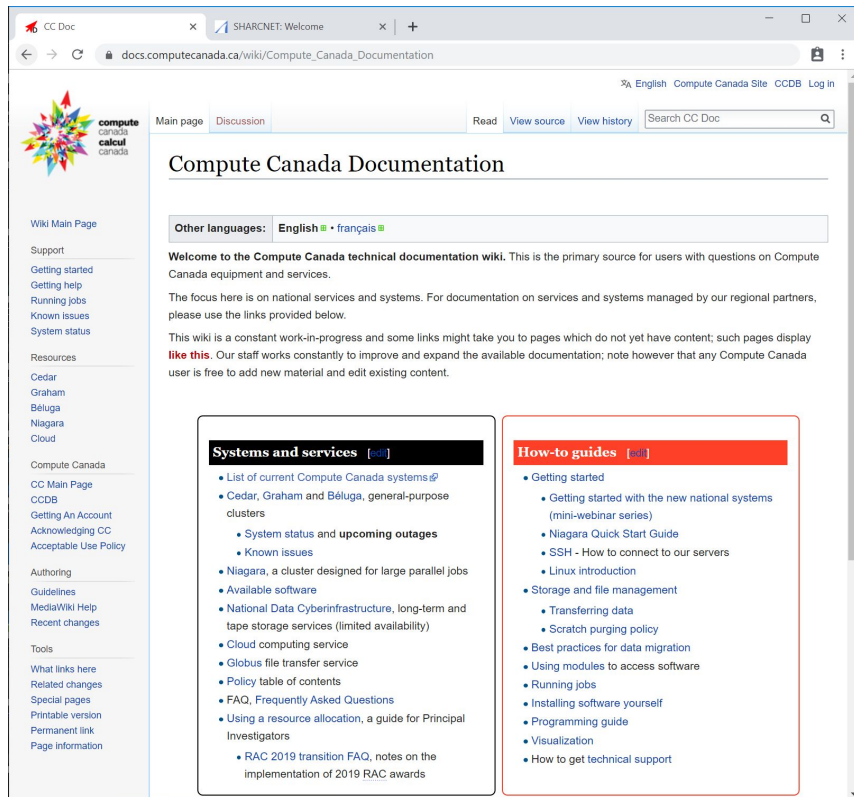
Shared
Hierarchical
Academic
Research
Computing
NETwork



- Member of Compute Canada and Compute Ontario
- 3,000+ Canadian and international users
- Almost 50,000 CPU cores
- 320 P100 GPUs
- 10 Gb/s network

- 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

Where to look for information



The screenshot shows the 'Compute Canada Documentation' wiki page. The page has a sidebar on the left with links like 'Wiki Main Page', 'Support', 'Getting started', 'Running jobs', 'Known issues', 'System status', 'Resources', and 'Compute Canada'. The main content area is titled 'Compute Canada Documentation' and includes a welcome message, a list of languages (English and français), and two main sections: 'Systems and services' and 'How-to guides'. The 'Systems and services' section lists various Compute Canada systems and services, while the 'How-to guides' section lists various guides and resources.

CC Doc x SHARCNET: Welcome x +

docs.compute.canada.ca/wiki/Compute_Canada_Documentation

English Compute Canada Site CCDB Log in

Main page Discussion Read View source View history Search CC Doc

Compute Canada Documentation

Other languages: English • français

Welcome to the Compute Canada technical documentation wiki. This is the primary source for users with questions on Compute Canada equipment and services.

The focus here is on national services and systems. For documentation on services and systems managed by our regional partners, please use the links provided below.

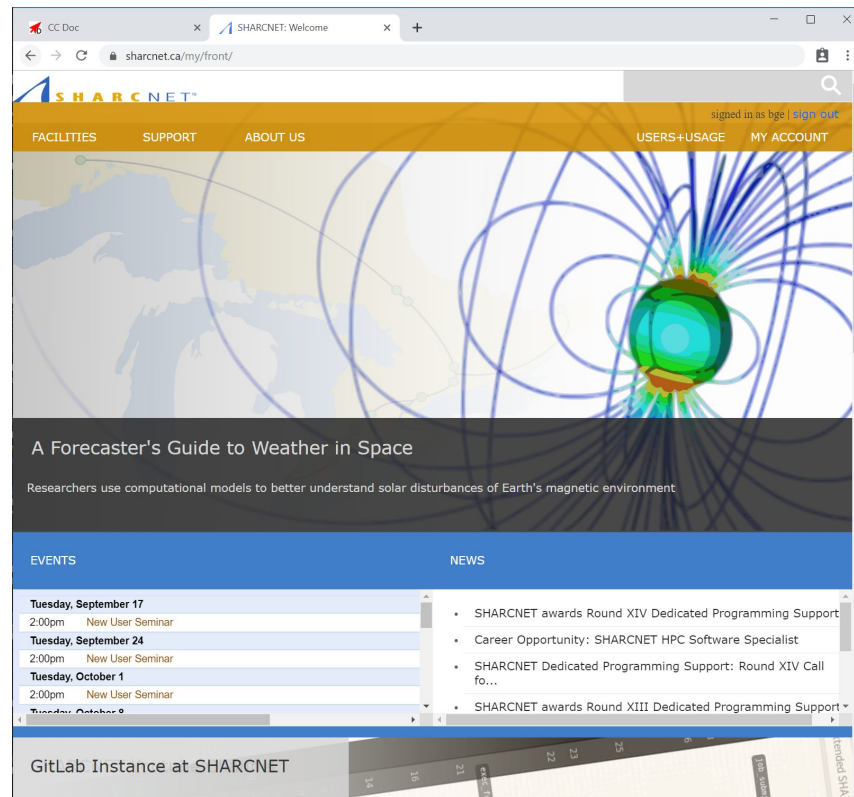
This wiki is a constant work-in-progress and some links might take you to pages which do not yet have content; such pages display **like this**. Our staff works constantly to improve and expand the available documentation; note however that any Compute Canada user is free to add new material and edit existing content.

Systems and services [edit]

- List of current Compute Canada systems@
- Cedar, Graham and Béluga, general-purpose clusters
 - System status and upcoming outages
 - Known issues
- Niagara, a cluster designed for large parallel jobs
- Available software
- National Data Cyberinfrastructure, long-term and tape storage services (limited availability)
- Cloud computing service
- Globus file transfer service
- Policy table of contents
- FAQ, Frequently Asked Questions
- Using a resource allocation, a guide for Principal Investigators
 - RAC 2019 transition FAQ, notes on the implementation of 2019 RAC awards

How-to guides [edit]

- Getting started
 - Getting started with the new national systems (mini-webinar series)
 - Niagara Quick Start Guide
 - SSH - How to connect to our servers
 - Linux introduction
- Storage and file management
 - Transferring data
 - Scratch purging policy
- Best practices for data migration
- Using modules to access software
- Running jobs
- Installing software yourself
- Programming guide
- Visualization
- How to get technical support



The screenshot shows the SHARCNET website. The page has a header with the SHARCNET logo and navigation links like 'FACILITIES', 'SUPPORT', 'ABOUT US', 'USERS+USAGE', and 'MY ACCOUNT'. The main content area features a large image of a globe with magnetic field lines and the title 'A Forecaster's Guide to Weather in Space'. Below this, there is a section for 'EVENTS' and 'NEWS'. The 'EVENTS' section lists several seminars and workshops, while the 'NEWS' section lists recent news items.

CC Doc x SHARCNET: Welcome x +

sharcnet.ca/my/front/

signed in as bge | sign out

FACILITIES SUPPORT ABOUT US USERS+USAGE MY ACCOUNT

A Forecaster's Guide to Weather in Space

Researchers use computational models to better understand solar disturbances of Earth's magnetic environment

EVENTS

Tuesday, September 17
2:00pm New User Seminar

Tuesday, September 24
2:00pm New User Seminar

Tuesday, October 1
2:00pm New User Seminar

Tuesday, October 8

NEWS

- SHARCNET awards Round XIV Dedicated Programming Support
- Career Opportunity: SHARCNET HPC Software Specialist
- SHARCNET Dedicated Programming Support: Round XIV Call fo...
- SHARCNET awards Round XIII Dedicated Programming Support

GitLab Instance at SHARCNET

help@sharcnet.cat or support@compute.canada.ca

SHARCNET web siteL <https://www.sharcnet.ca/>

- **Documentation and training materials:**

- Help pages, tutorials: [Support > Wiki](#)
- SHARCNET's youtube channel: [youtube.sharcnet.ca](https://www.youtube.com/sharcnet.ca)

- **Ticketing system**

- Send an email to help@sharcnet.ca

Compute Canada web site (docs.compute canada.ca)

- **How-to guides**
- **Systems and services**
- **Discipline guides**
- **Regional partners**
- **Compute Canada's problem tracking system**
 - Email to support@compute canada.ca
 - Email to help@sharcnet.ca

- SHARCNET
- Where to look for information and get help
- **Essentials**
 - What are available
 - How to connect to graham
 - How to manage disk storage
 - How to run your programs - *jobs*
- Common mistakes to avoid
- Q & A

- **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

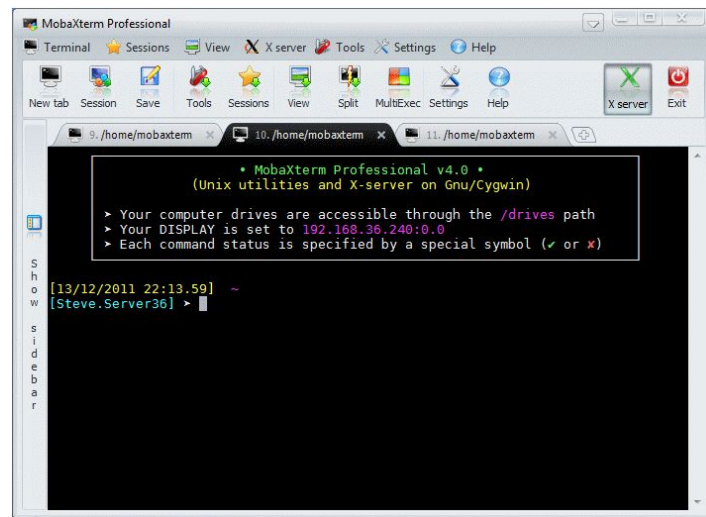
- Number of CPU cores: 33,448
- Number of nodes: 1043
- 32 cpu cores per node
- Between 128 and 3072 GB of RAM per node
- Number of NVIDIA P100 GPUs: 320
- Networking: EDR (cpu nodes) and FDR (GPU and cloud nodes)
InfiniBand

Connecting to clusters

- All systems are only accessible via *secure shell* (ssh), using your Compute Canada credentials
`$ ssh user@graham.computecanada.ca`
- 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)
 - Windows Subsystem for Linux
- 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 60 days	Yes	Yes
Project Space /project	1 TB and 0.5M files per group, can request increase to 10 TB	Yes	No	Yes	Yes
Nearline Space	5 TB per group	No	No	Yes	No

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

Essentials: Running programs with SLURM

- Jobs are submitted using the **sbatch** command with a script, e.g. run_job.sh, containing

```
#!/bin/bash
```

```
#SBATCH --time=0-00:05      # Run time limit (DD-HH:MM)
```

```
#SBATCH --account=def-user
```

```
#SBATCH --ntasks=32        # Number of MPI processes, default 1
```

```
#SBATCH --cpus-per-task=32  # Normally defined for threaded jobs
```

```
#SBATCH --gres=gpu:2        # 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 mpirun ./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.

Why is my job not starting?

- There may be multiple reasons
- Graham/Cedar are very busy clusters, with ~15% 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.***

- SHARCNET
- Where to look for information and get help
- Essentials
 - What are available
 - How to connect to graham
 - How to manage disk storage
 - How to run your programs - *jobs*
- **Common mistakes to avoid**
- Q & A

- 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. 130% 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

- Do not run “watch queue”. If it is required please add a delay of at least 60 seconds between updates, with “watch -n 60 queue”.
- Do not submit more than 1 job per second to the scheduler. Add a delay of at least 1 second between each job submission.
- Do not submit very short jobs (less than 10 minutes). In this case combine tasks into longer jobs.
- If you have many jobs to submit, consider using alternatives
 - Check “Serial farming on Graham” webinar on youtube.sharcnet.ca