

Super Computing



Figure 1: Part of the Graham super computer.

What isn't a supercomputer

A super fast computer that just runs your program faster.

What is a supercomputer

Graham (latest SHARCNET supercomputer)

- 1159 computers (nodes)
- 37060 cores
- 166 terrabytes (TB) of RAM (most 4 gigabyte (GB)/core)
- 15 petabytes (PB) of disk storage (home, project, and scratch)
- 80 petabytes (PB) of transparent tape storage (nearline)
- 398 NVIDIA GPUs (P100 Pascal, V100 Volta, and T4 Turing)
- FDR (56Gb/s) and EDR (100Gb/s) InfiniBand

How do you use a supercomputer

serial (easy) many computers means you can run many programs independently at the same to solve many independent problems

parallel (hard) many computers can (sometimes) be programmed to all collaborate and solve a single problem together

cloud virtual computer on the internet (like Amazon's Elastic Compute Cloud, Microsoft's Azure cloud, or Google's Compute Engine)

Super computing in Canada



Figure 2: Computer (node) from the Graham super computer.

The players

National

- Compute Canada

Regional

- ACENET, Calcul Québec, Compute Ontario, WestGrid

Ontario

- CAC, SciNet, and SHARCNET

Getting an accounts

Cost

- no cost

Who

- faculty or those sponsored by faculty that have an account

How

- <https://ccdb.computecanada.ca>
- apply for a Compute Canada account
- pick the desired regional consortia accounts

What software is available

Operating system

- Linux (CentOS 7)

Programming languages

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

Parallel development support

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

Other

- common open source and commercial packages (e.g., OpenFOAM, Fluent, and STAR-CCM+)

How to use

- resources are scheduled to avoid collisions and ensure fair access
- access from anywhere on the internet using secure shell (SSH) to enter commands and a secure file transfer (SFTP) to transfer files
- tell the supercomputer what program (command) you want it to run and then do something else till it does it

Typical workflow

1. transfer your data and/or programs to the supercomputer using SFTP
2. login (bring up a window in which you can enter commands) to the supercomputer using SSH
3. enter commands to tell the scheduler what you want it to run when the required resources are available
4. do something else until you get notification that your commands have completed running
5. transfer the resulting data from the supercomputer to your computer