General Interest Seminar

Using Computational Chemistry Software Effectively on Graham

Jemmy Hu
SHARCNET HPC Consultant
University of Waterloo

March 28, 2018

Packages available on Graham

http://wiki.computecanada.ca

Discipline guides

or

Search: Computational chemistry

https://docs.computecanada.ca/wiki/Computational_chemistry

The nature of the popular packages (I)

ab initio	DFT	Molecular mechanics	Tools
Gaussian	ADF	Amber	Molden
NWChem	Quantum Espresso	Gromacs	VMD
Gamess	CP2K	CPMD	Plumed
ORCA	deMon	NAMD	OpenMM
VASP	PSI4	LAMMPS	
Siesta		DL_POLY	

The nature of the popular packages (II)

- Commercial packages

  - Gaussian (soft_Gaussian, Graham only at the moment)
  - ADF (Graham only)
  - Amber (Graham only)

  - VASP (soft_vasp5, PI's group license)
    https://docs.computecanada.ca/wiki/VASP
  - ORCA (soft_orca, register to ORCA is required)
    https://docs.computecanada.ca/wiki/ORCA

- Parallel coding approach
  - MPI distributed parallel solution, run on cpus across multi-node
  - OpenMP shared memory solution (Gaussian), cpus on one node

Software modules

- Software are installed as Module packages
- Basic module commands

module avail	# to list all available module software
module spider softwareName	# to see what versions are available

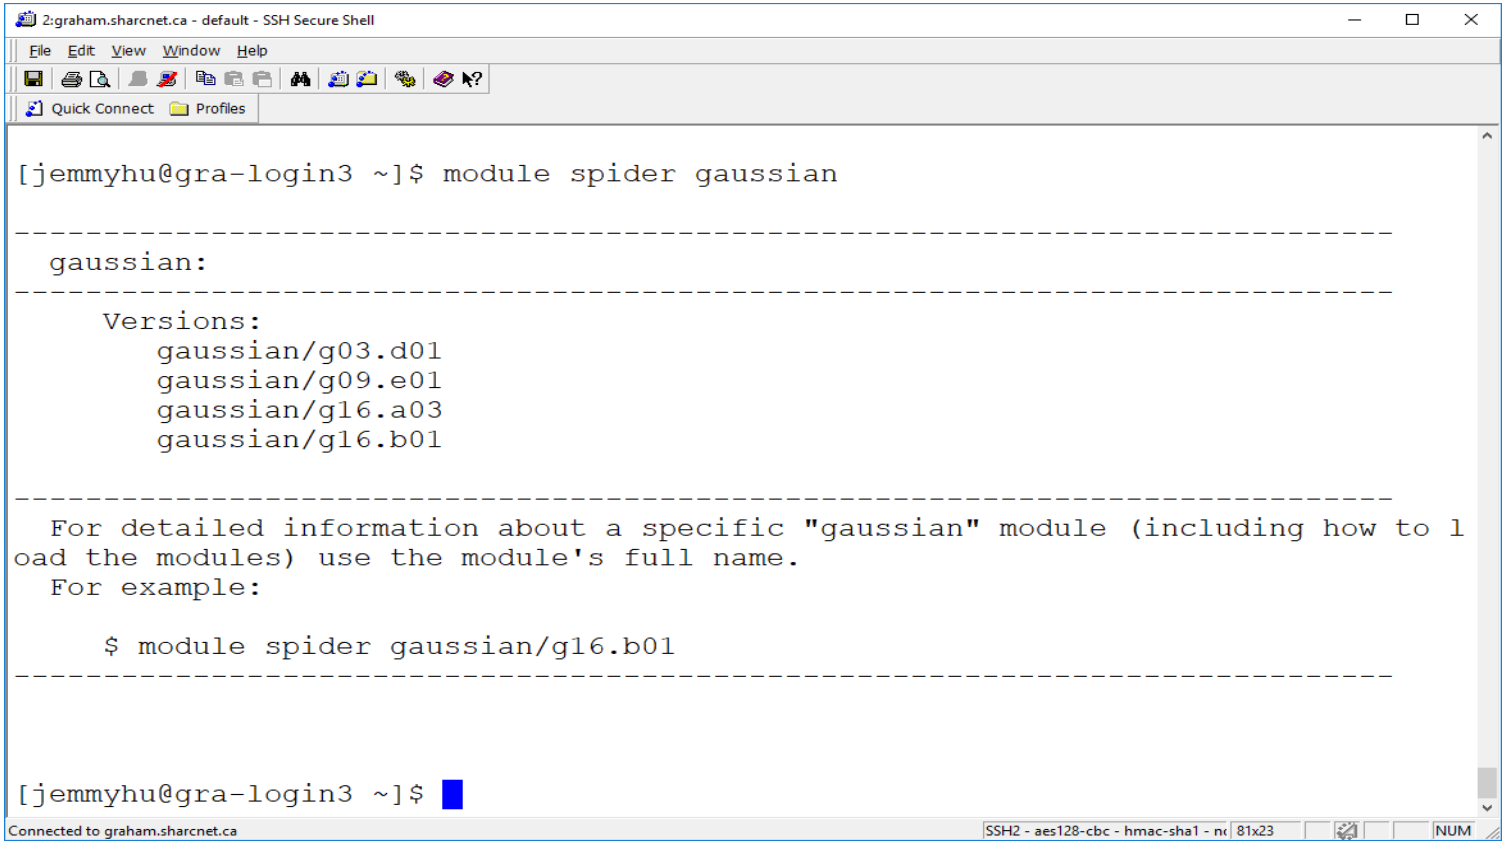

Submit/Run script:

- account type (default or RAC account)
- compute resource (cpus, memory, runtime)
- job type (mpi, openmp, serial)
- software (module load, run command)

- mpi_job.sh

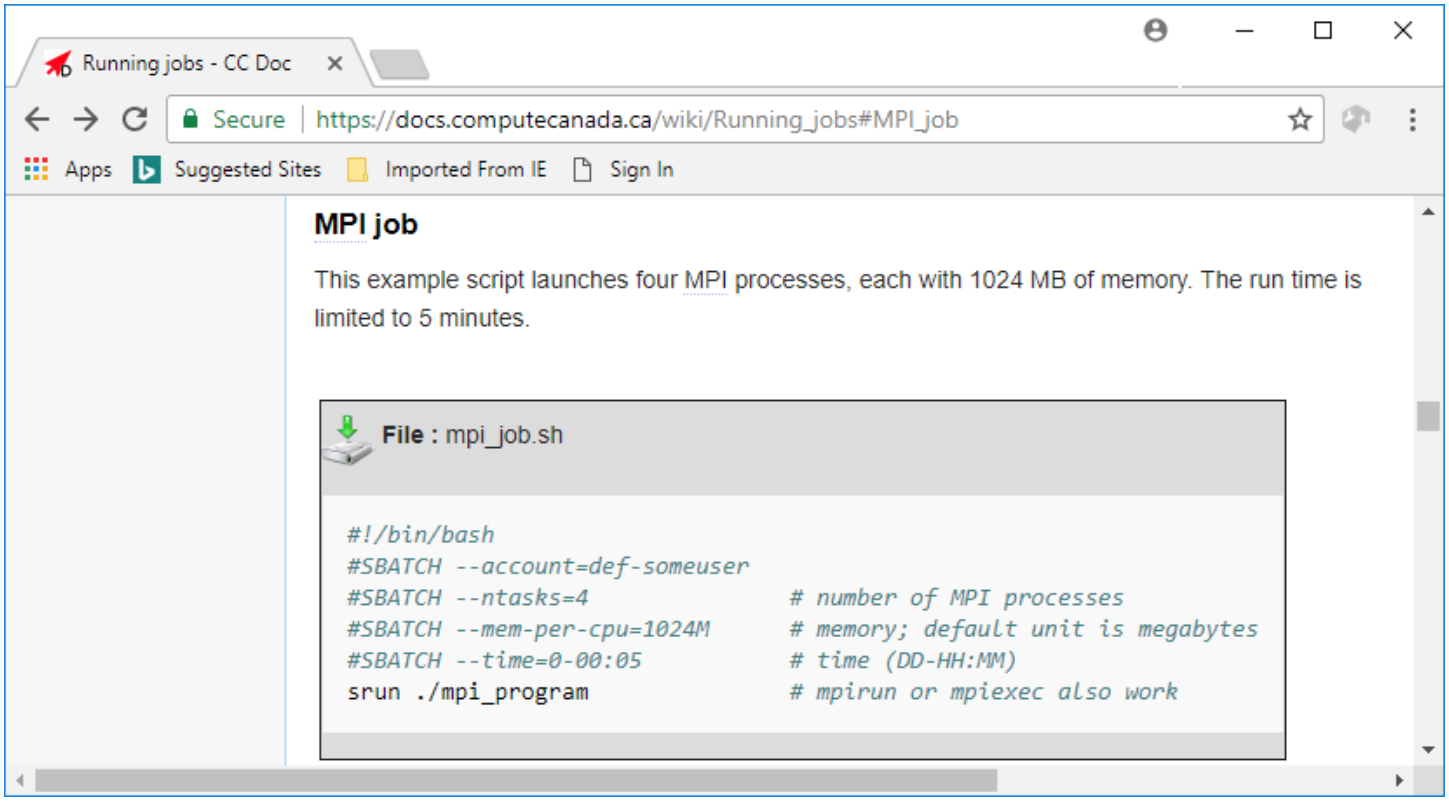

•openmp_job.sh

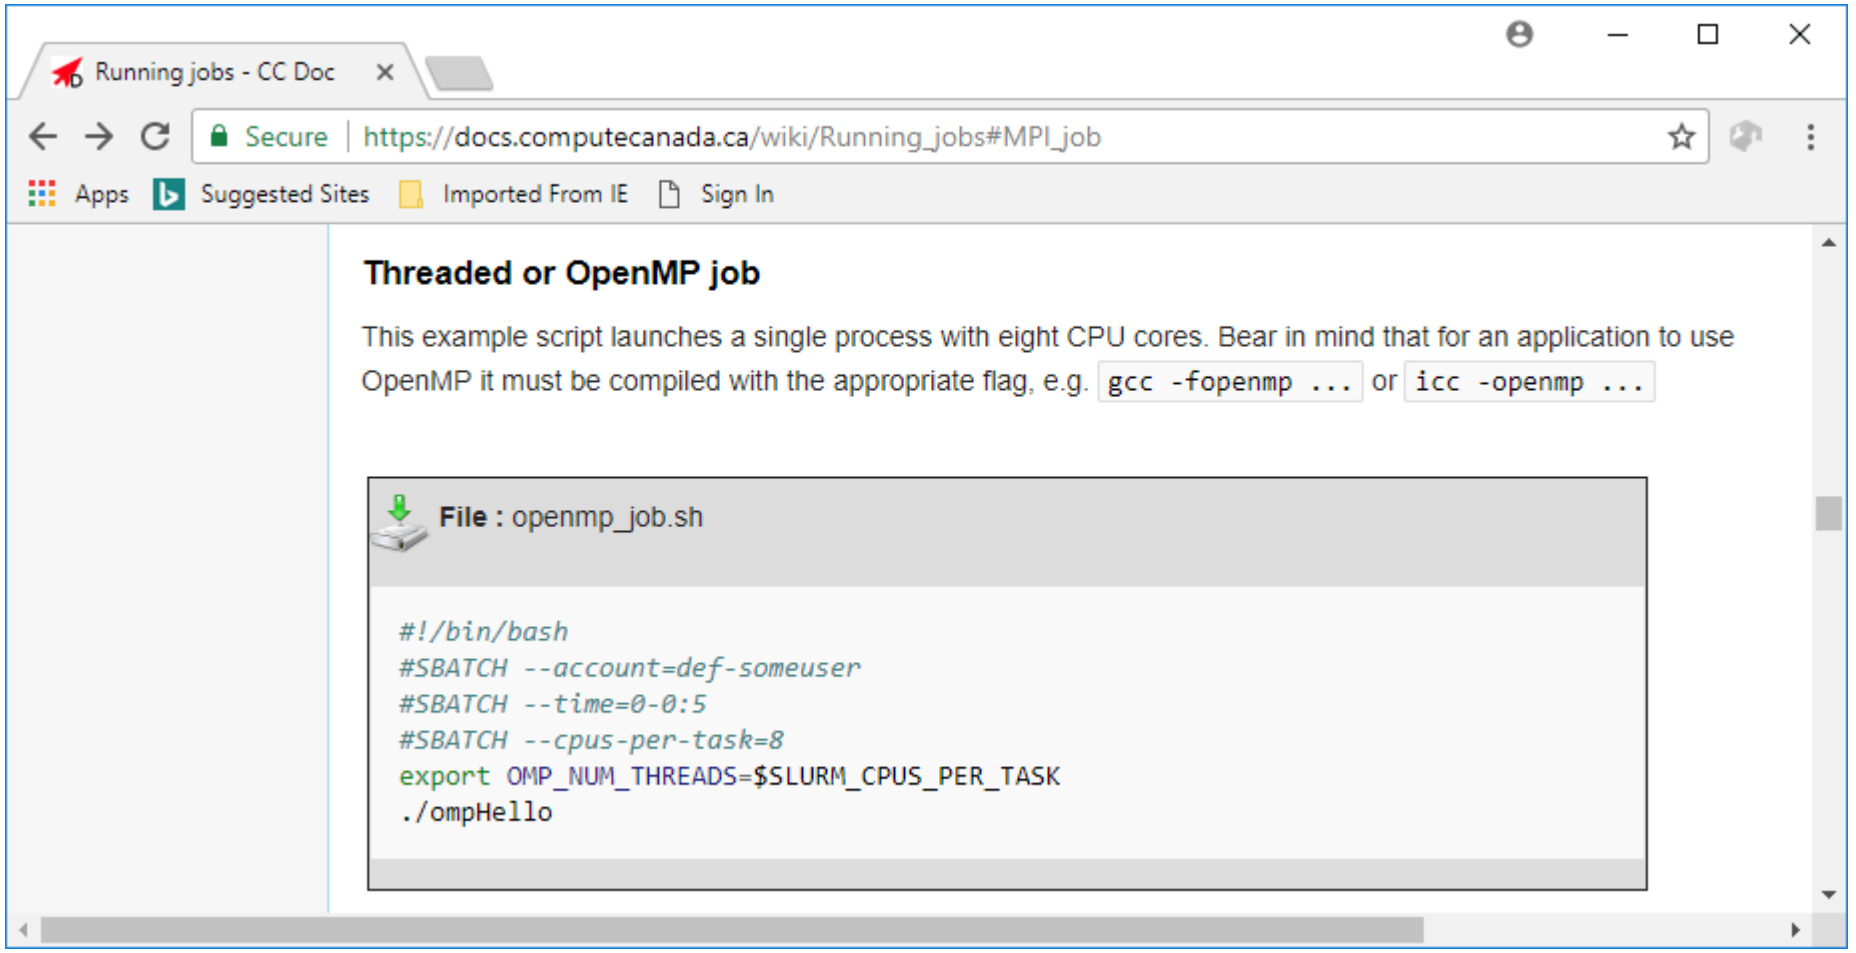

[… ~]$ sbatch --help

Slurm commands

https://docs.computecanada.ca/wiki/Running_jobs

- sbatch name.sh # submit the job

- squeue -u username # check your job status
- scontrol show job -dd jobid # detailed info about a job
- sacct -j jobid # info about a finished job
- scancel jobid # cancel a job with the jobid
- scancel -t PENDING -u username # cancel all your pending jobs

interactive job (test input):

salloc --time=1:0:0 --ntasks=8 –mem=10g --account=def-someuser

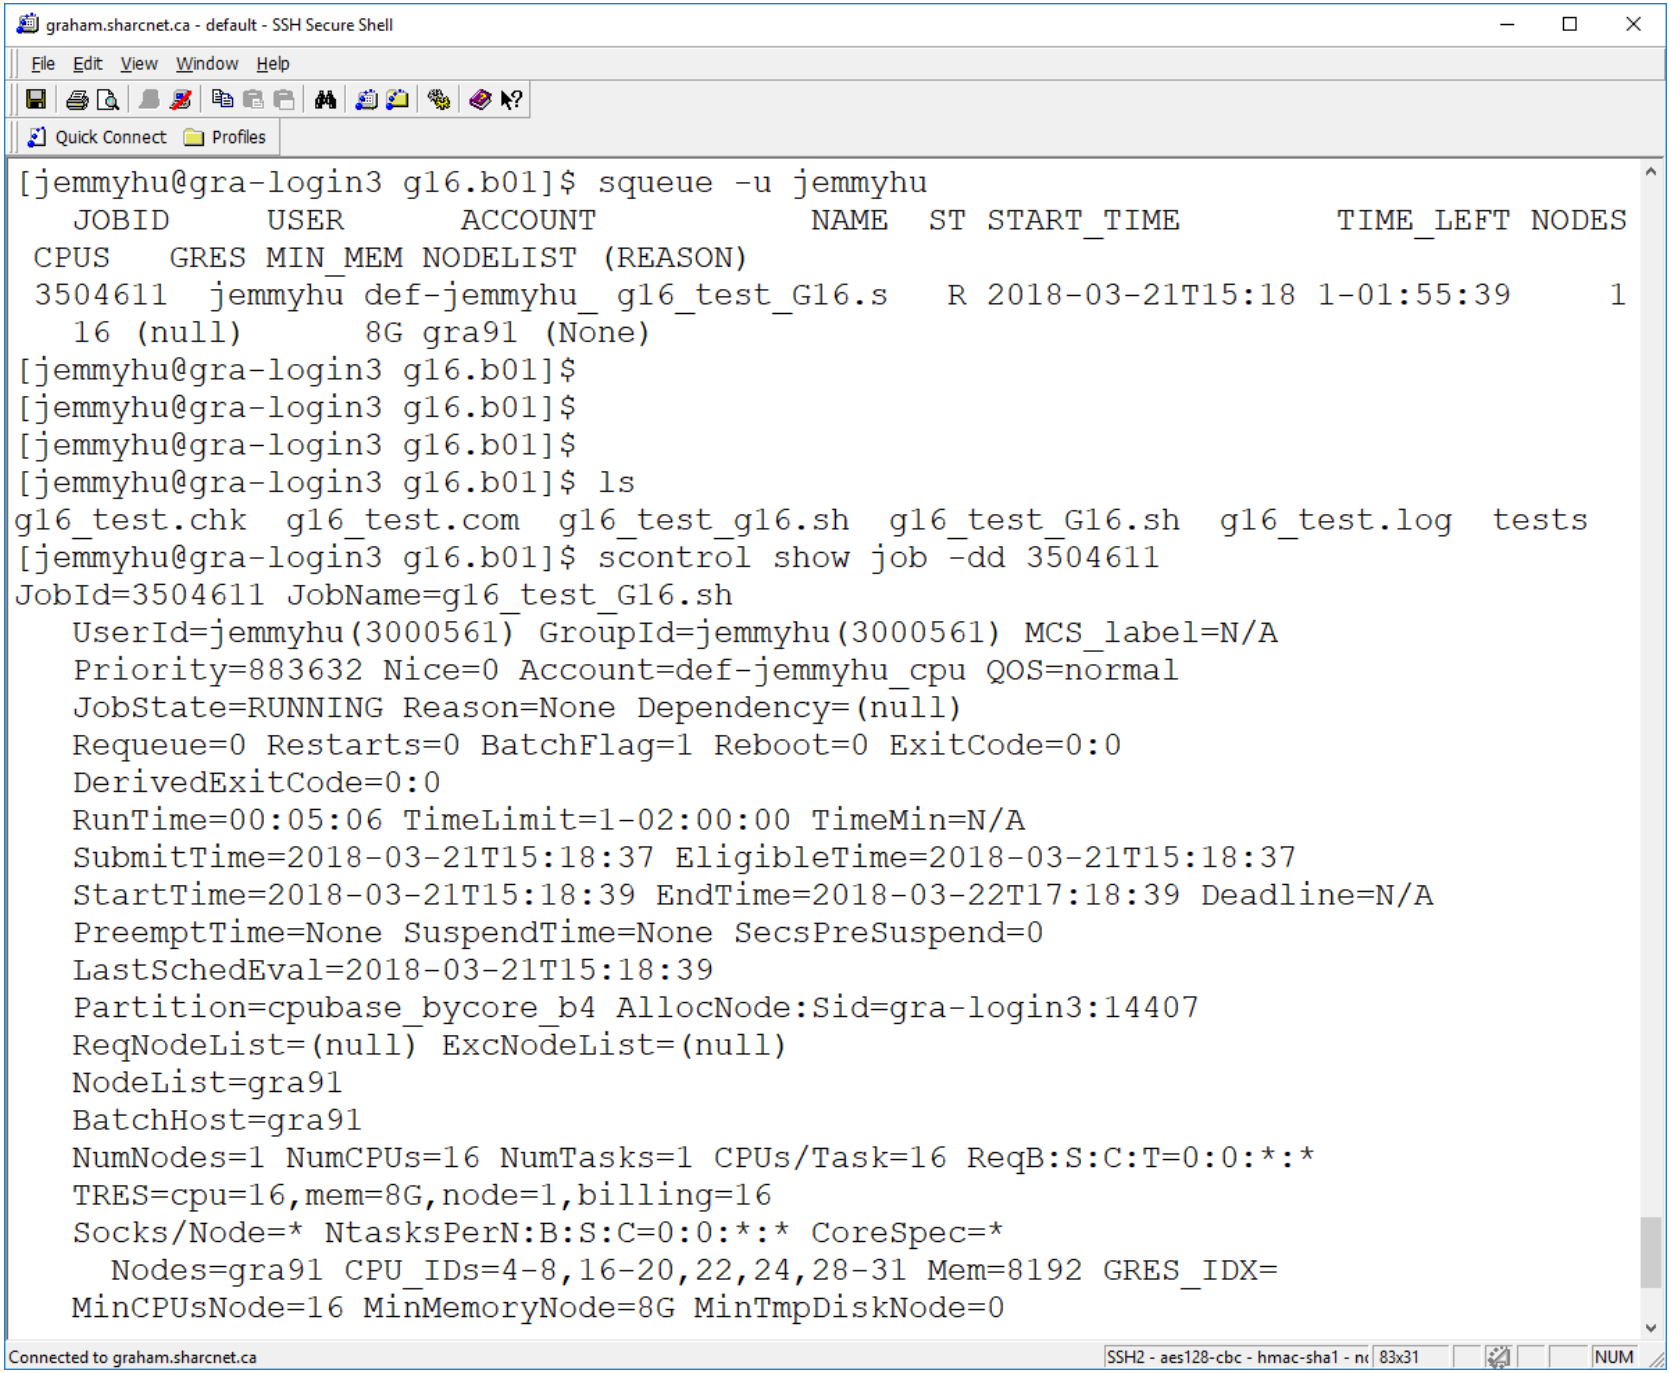
graham.sharcnet.ca - default - SSH Secure Shell
File Edit View Window Help
Quick Connect Profiles
[jemmyhu@gra-login3 g16.b01]$ squeue -u jemmyhu
JOBID USER ACCOUNT NAME ST START_TIME TIME_LEFT NODES
CPUS GRES MIN_MEM NODELIST (REASON)
3504611 jemmyhu def-jemmyhu_ g16_test_G16.s R 2018-03-21T15:18 1-01:55:39 1
16 (null) 8G gra91 (None)
[jemmyhu@gra-login3 g16.b01]$
[jemmyhu@gra-login3 g16.b01]$
[jemmyhu@gra-login3 g16.b01]$
[jemmyhu@gra-login3 g16.b01]$ ls
g16_test.chk g16_test.com g16_test_g16.sh g16_test_G16.sh g16_test.log tests
[jemmyhu@gra-login3 g16.b01]$ scontrol show job -dd 3504611
JobId=3504611 JobName=g16_test_G16.sh
UserId=jemmyhu(3000561) GroupId=jemmyhu(3000561) MCS_label=N/A
Priority=883632 Nice=0 Account=def-jemmyhu_cpu QOS=normal
JobState=RUNNING Reason=None Dependency=(null)
Requeue=0 Restarts=0 BatchFlag=1 Reboot=0 ExitCode=0:0
DerivedExitCode=0:0
RunTime=00:05:06 TimeLimit=1-02:00:00 TimeMin=N/A
SubmitTime=2018-03-21T15:18:37 EligibleTime=2018-03-21T15:18:37
StartTime=2018-03-21T15:18:39 EndTime=2018-03-22T17:18:39 Deadline=N/A
PreemptTime=None SuspendTime=None SecsPreSuspend=0
LastSchedEval=2018-03-21T15:18:39
Partition=cpubase_bycore_b4 AllocNode:Sid=gra-login3:14407
ReqNodeList=(null) ExcNodeList=(null)
NodeList=gra91
BatchHost=gra91
NumNodes=1 NumCPUs=16 NumTasks=1 CPUs/Task=16 ReqB:S:C:T=0:0:*:*
TRES=cpu=16,mem=8G,node=1,billing=16
Socks/Node=* NtasksPerN:B:S:C=0:0:*:* CoreSpec=*
Nodes=gra91 CPU_IDs=4-8,16-20,22,24,28-31 Mem=8192 GRES_IDX=
MinCPUsNode=16 MinMemoryNode=8G MinTmpDiskNode=0
Connected to graham.sharcnet.ca
SSH2 - aes128-cbc - hmac-sha1 - n
83x31
NUM

Gaussian

https://docs.computecanada.ca/wiki/Gaussian (g## or G##)

- Versions
  - g03.d01
  - g09.e01
  - g16.a03
  - g16.b01

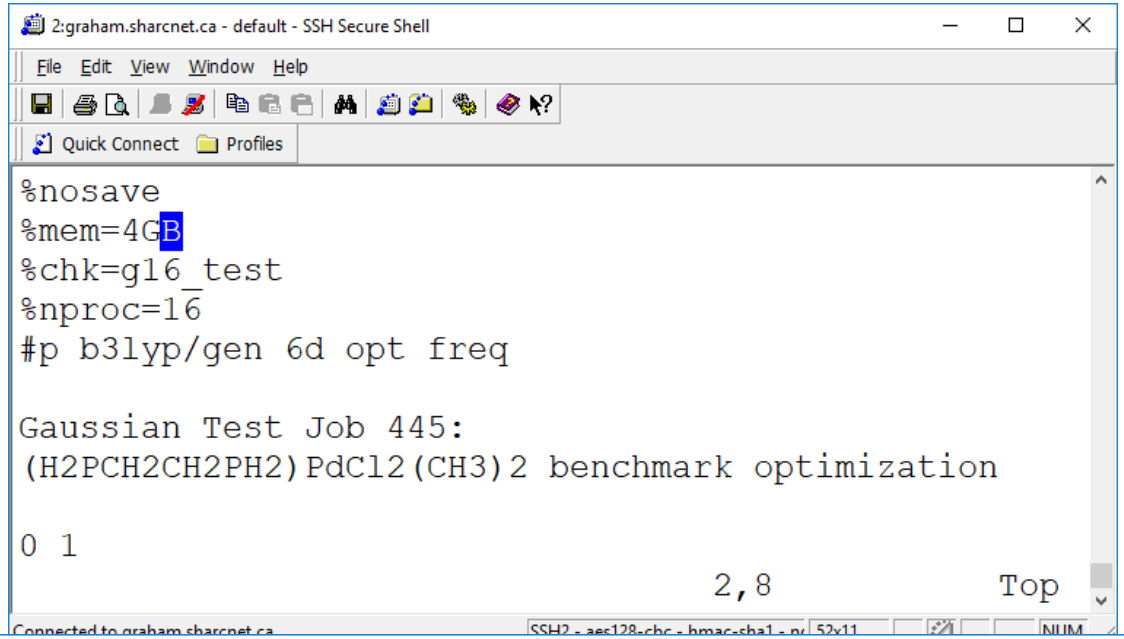

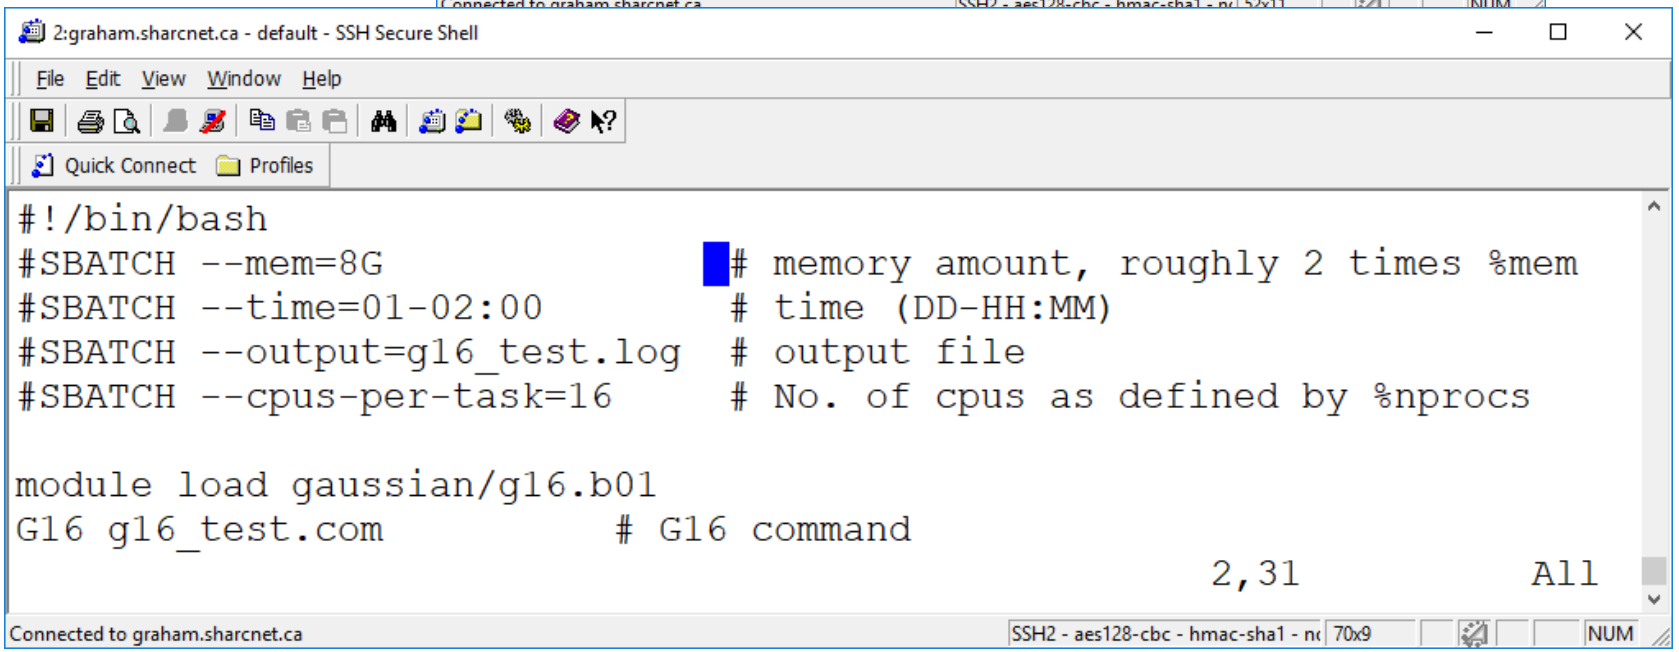

Tips to run Gaussian on Graham

- use the latest version
  g16.b01
  g09.e01 (if you need to include NBO6)

- use 16 or 8 cpus:
  %nproc=16 #input name.com
  #SBATCH --cpus-per-task=16 # gaussian name.sh
- #SBATCH --mem=8G is ~2 times %mem=4GB in the input *.com

- File location: G16/G09, /scratch/userid/jobid/

- Testing input .com
  [~$] salloc --time=1:0:0 --cpus-per-task=16 --mem=10g --account=def-PI
  [~$] module load Gaussian/g16.b01
  [~$] G16 g16_test.com

ADF https://docs.computecanada.ca/wiki/ADF

- input files
  - one step job
  - multi-step job
- submit script (node-based)

- Submit the job: sbatch mysub.sh

MPI by-node partition: --nodes --ntasks-per-node (32x)

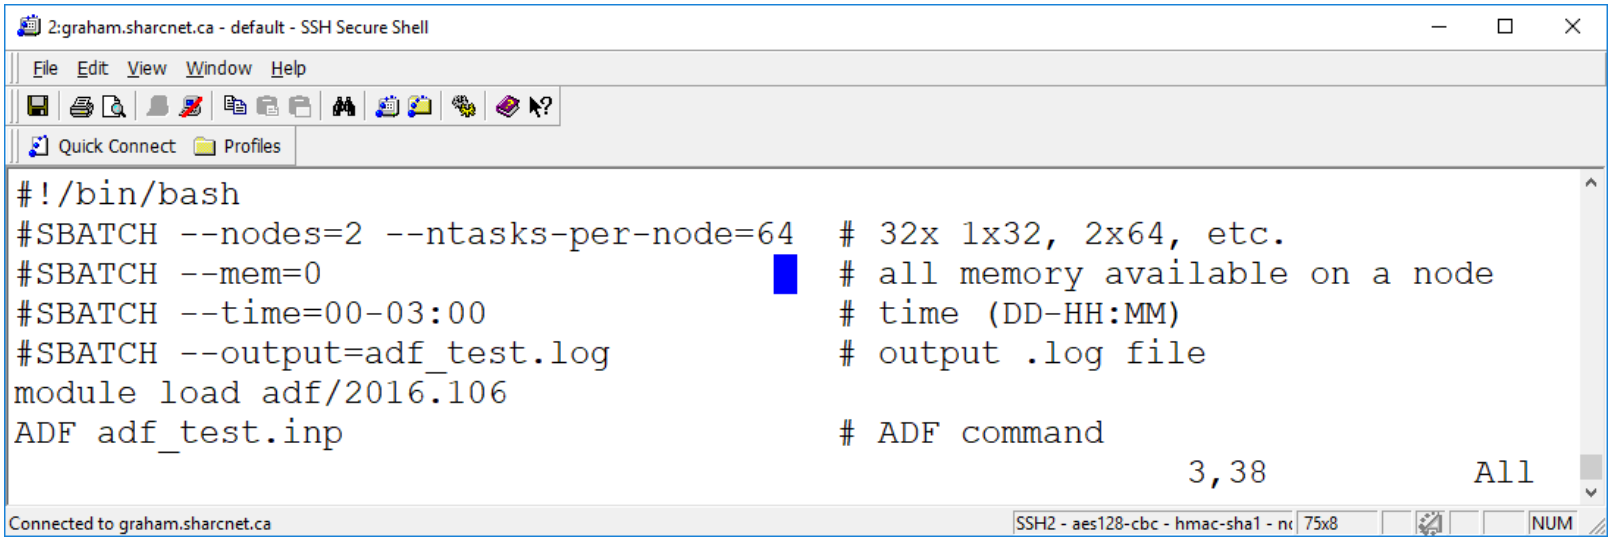

MPI by-core partition: --ntasks, --mem-per-cpu

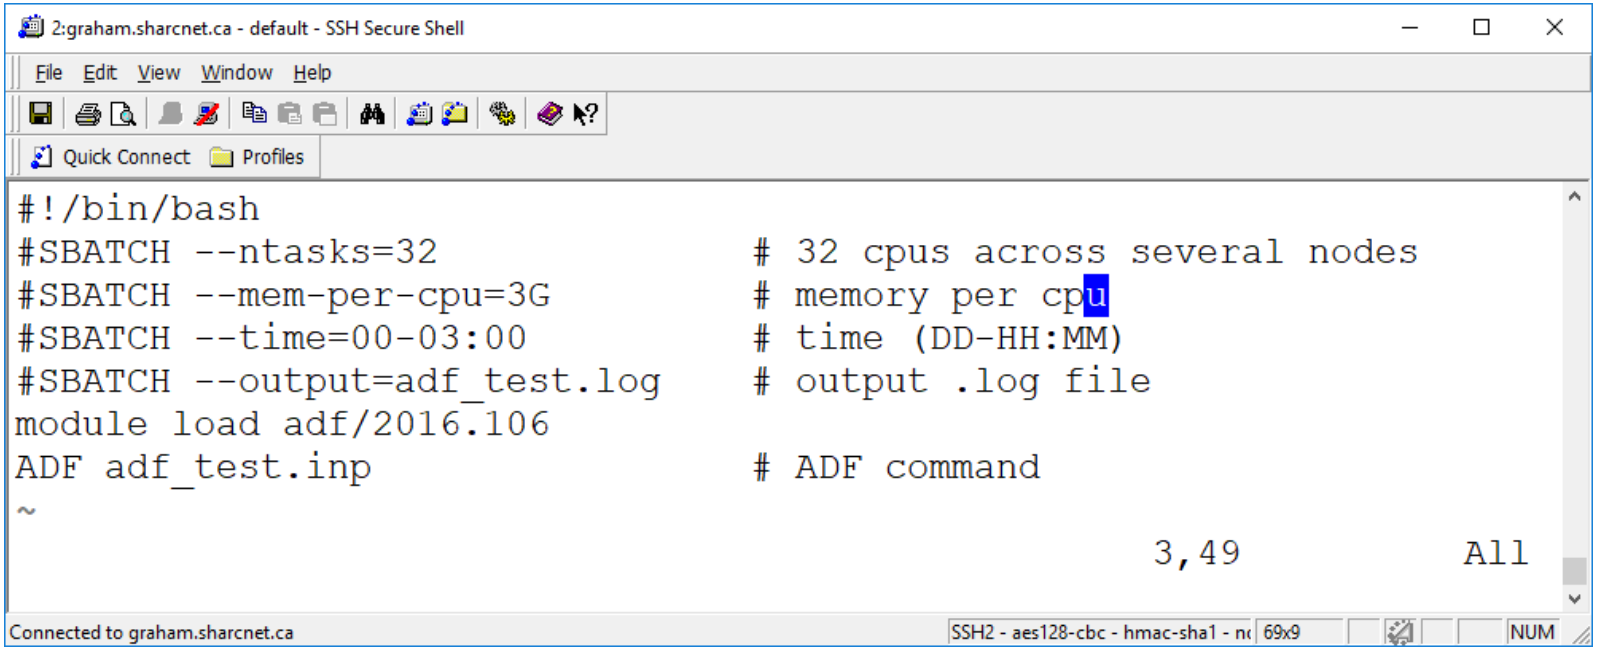

ORCA: binary orca/4.0.1.2

load the corresponding modules

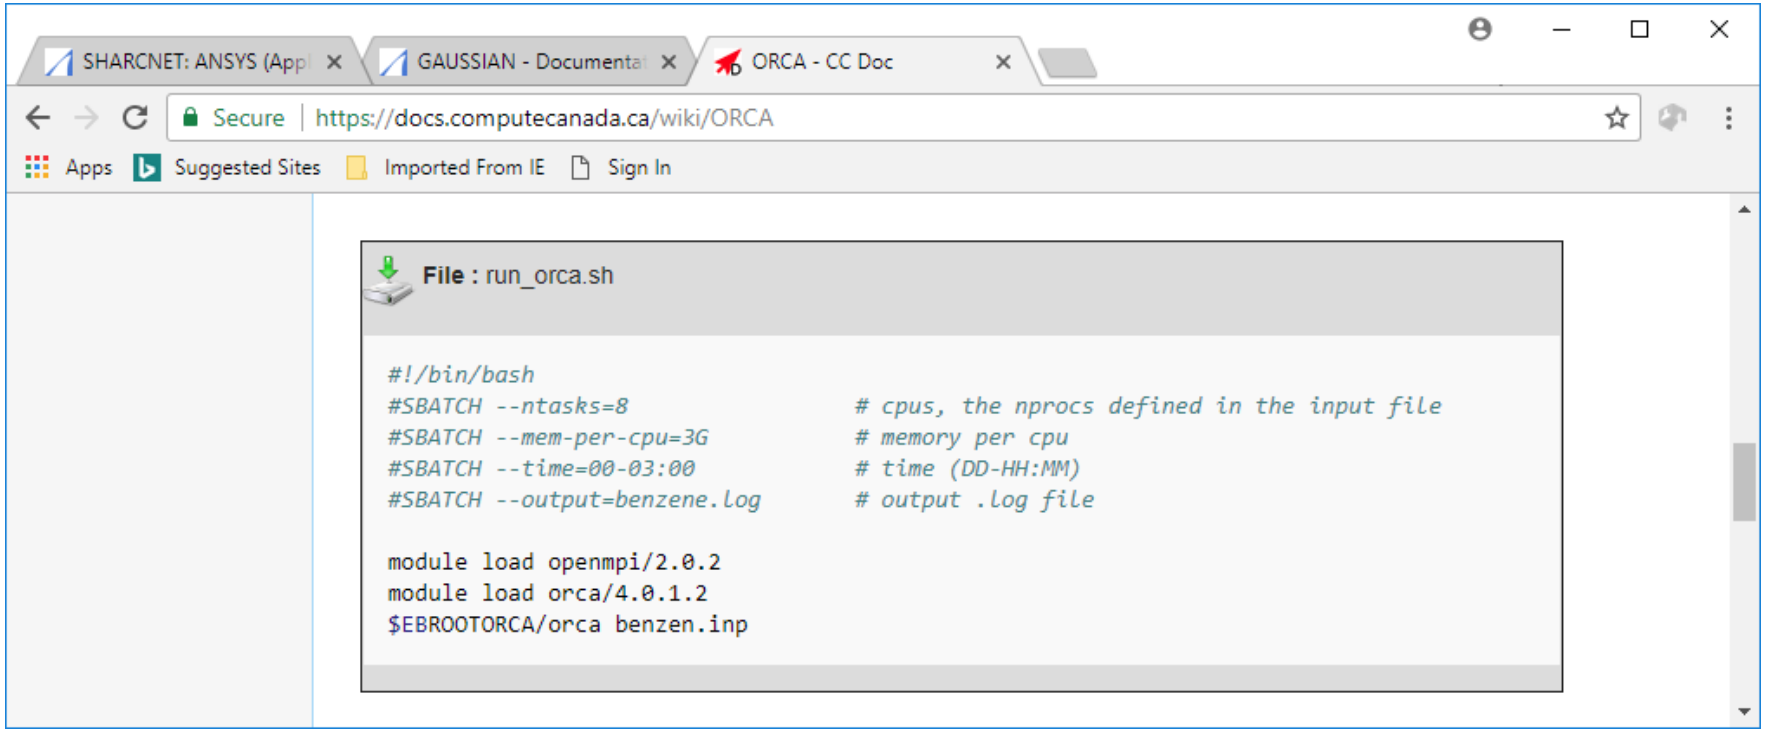

https://docs.computecanada.ca/wiki/ORCA

VASP

[jemmyhu@gra-login3 ~]$ module spider vasp

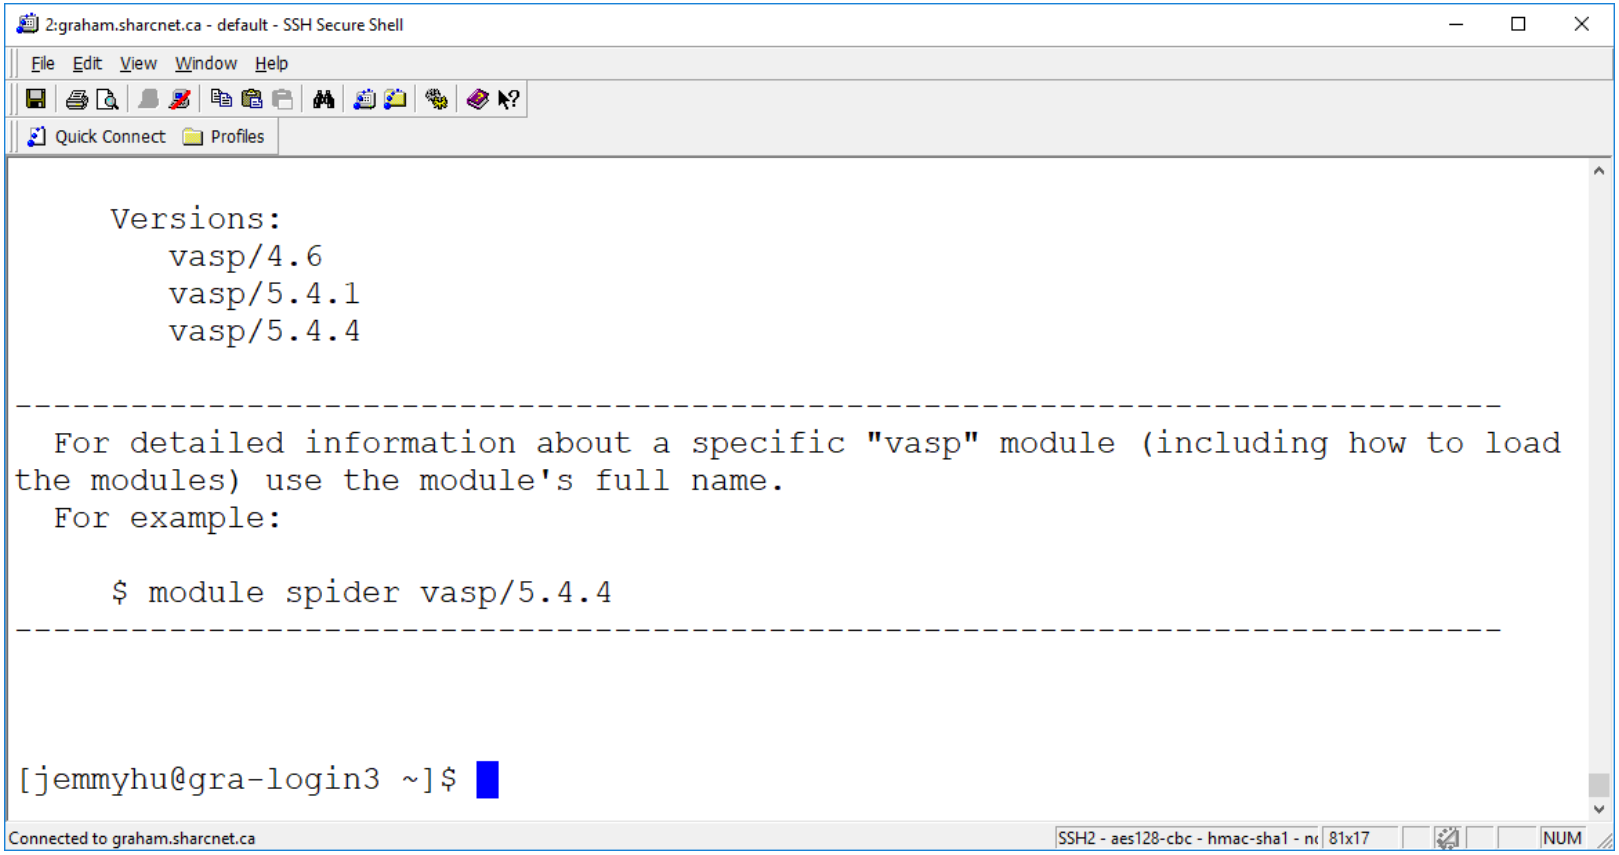

https://docs.computecanada.ca/wiki/VASP

<VASP> executables

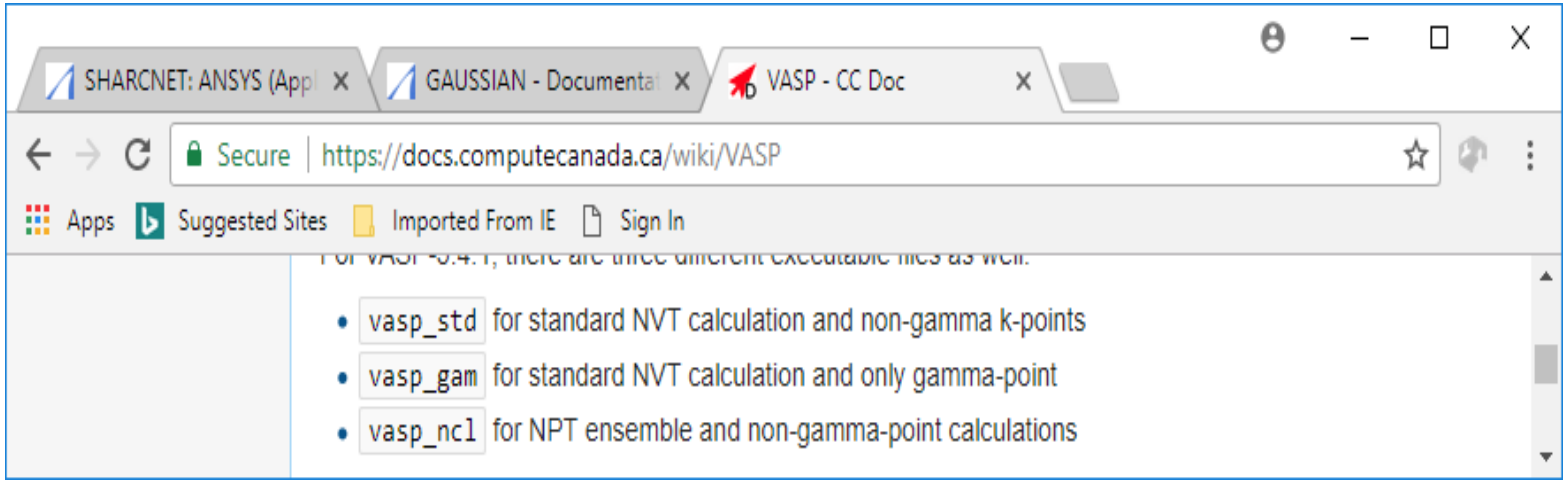

VASP input files: INCAR, KPOINTS, POSCAR, POTCAR (each job in a subdirectory)

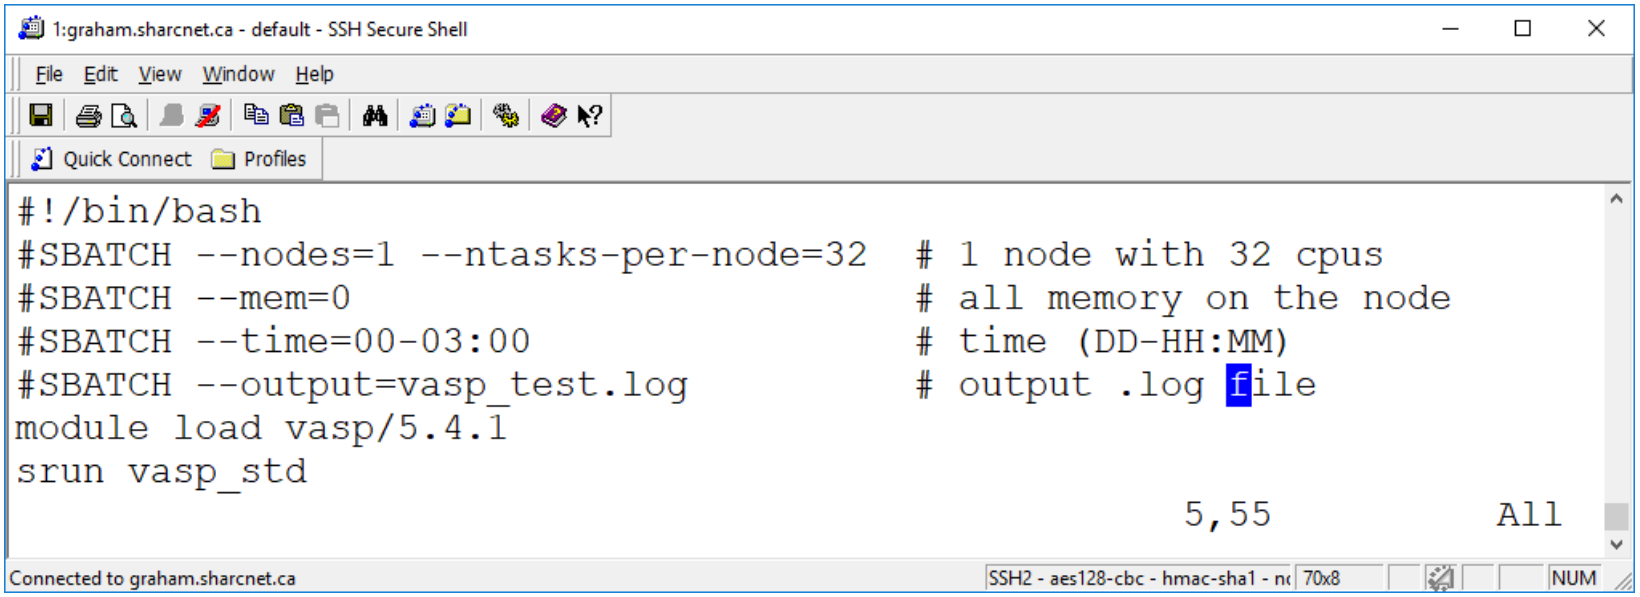

[~$] sbatch vasp.sh

Resource and Question

- Online wiki documents
  http://wiki.computecanada.ca
- Examples

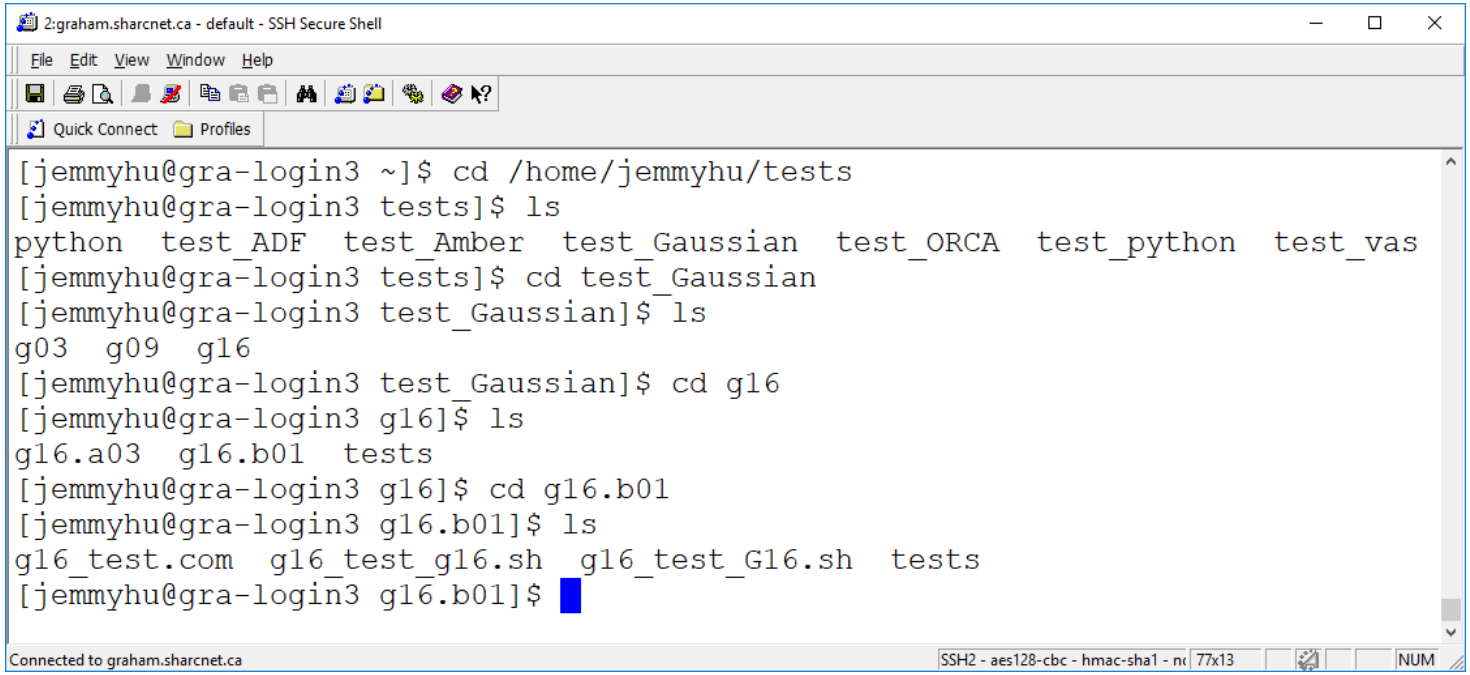

- Questions to
  - email to support@computecanada.ca
    to the online ticketing system