

Introduction to SHARCNET Literacy Seminars

Jemmy Hu

2007

Contents

- SHARCNET
- Hardware
- Software
- File system
- Queue system
- Practical issues
- Support

What's SHARCNET

<http://www.sharcnet.ca>

- **The SHARCNET Vision**

To establish a world-leading, multi-university and college, interdisciplinary institute with an active academic-industry partnership, enabling forefront computational research in critical areas of science, engineering and business.

- **The SHARCNET Mission**

SHARCNET exists to enable world-class computational research so as to accelerate the production of research results.

General Objectives:

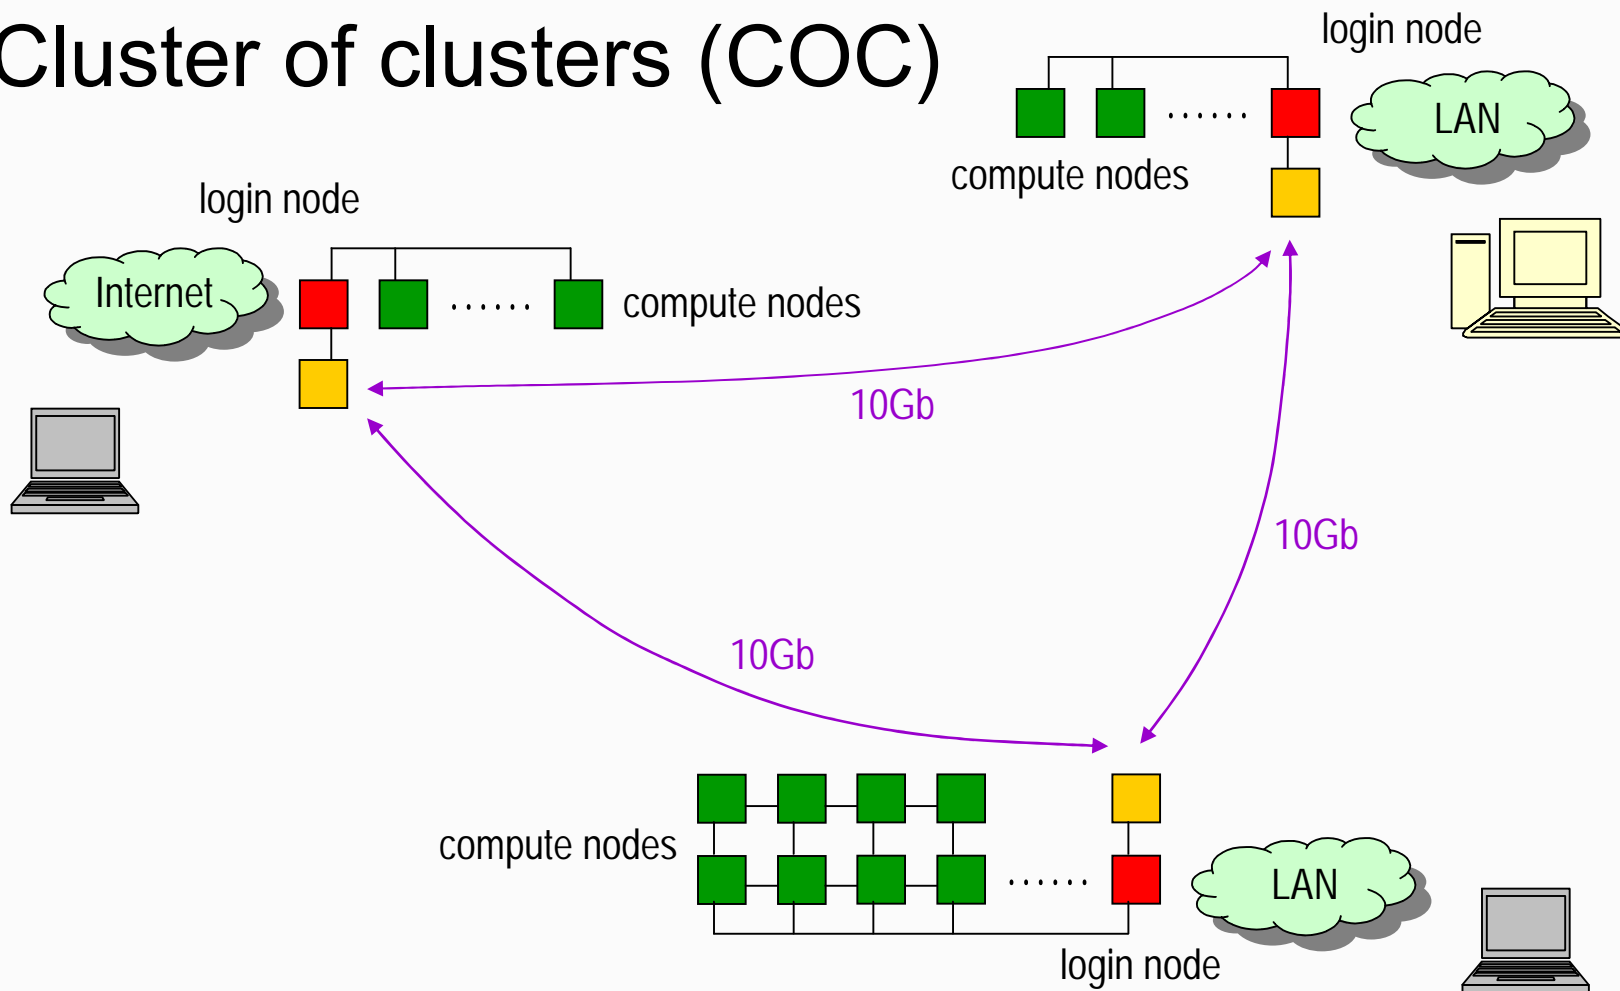
- provide otherwise unattainable compute resources
- build common, seamless computing environment
- promote remote collaboration and researches

SHARCNET Basics

- **FREE** to academic researchers
- **Compute-Intensive Problems** The resources are provided to enable HPC and are not intended as a replacement for a researcher's desktop or lab machines. SHARCNET users can productively conduct HPC research on a variety of SHARCNET systems each optimally designed for specific HPC tasks
- **Academic HPC research** The research can be business-related, but must be done in collaboration with an academic researcher
- **Fairness access**
 - Users have access to all systems
 - Job runs in batch mode (scheduling system) with fairshare

SHARCNET Environment

- Cluster of clusters (COC)



SHARCNET Facilities: Overview

- **Computers:** Clusters, SMP
 - 1) **Architecture:** Opteron, Alpha, Itanium/Xeon
each with its own recommended compilers:
Pathscale, PGI; Compaq; Intel
 - 2) **Intended Use:** HPC tasks specific
Parallel (capability, utility, SMP, ...), **Serial** (throughout, ...)
- Visualization Clusters
- **Access Grid**
 - multi-media video conferencing
 - cooperation, cross site workshop, etc

SHARCNET: Systems Summary - Mozilla Firefox

File Edit View Go Bookmarks Tools Help

http://www.sharcnet.ca/Facilities/index.php

Windows Media file:///C:/Documents... Customize Links Windows Windows Marketplace Free Hotmail Free AOL & Unlimited...

Home Facilities at a Glance

- About
- Contact Us
- Events
- Facilities
 - Systems Summary
 - Latest Notices
 - Network
 - RSS Feeds
 - AccessGrid
 - Photo Gallery
 - Policies
 - Software
- Help
- Media
- Performance
- Research
- My Sharcnet
 - Please log in

System	State	CPUs	Architecture	Nodes	Notices
bala	Online	128	Cluster/Myrinet G2	Opteron	Feb 27 2007
bruce	Online	128	Cluster/Myrinet G2	Opteron	Feb 01 2007
bull	Online	384	Cluster/Quadrics Elan4	Opteron	Mar 09 2007
cat	Online	160	Cluster/Gigabit Ethernet	Xeon, Opteron	Mar 23 2007
coral	Online	60	Cluster/Quadrics Elan3	Itanium2	Nov 20 2006
dolphin	Online	128	Cluster/Myrinet G2	Opteron	Jan 10 2007
goblin	Online	122	Cluster/Gigabit Ethernet	Opteron	Mar 02 2007
greatwhite	Online	456	Cluster/Quadrics Elan3	Alpha	Mar 02 2007
gulper	Online	42	Cluster/Myrinet MX	Opteron	Feb 07 2007
mako	Online	14	Cluster/Myrinet	Xeon	Mar 19 2007
megaladon	Online	128	Cluster/Myrinet G2	Opteron	Mar 20 2007
narwhal	Online	1068	Cluster/Myrinet G2	Opteron	Mar 21 2007
requin	Online	1536	Cluster/Quadrics Elan4	Opteron	Mar 23 2007
silky	Online	128	SMP/NUMA	Itanium2	Feb 19 2007
spinner	Online	62	Cluster/Gigabit Ethernet	Itanium2	Dec 15 2006
tiger	Online	128	Cluster/Myrinet G2	Opteron	Jan 30 2007
typhon	Online	16	SMP/SMP	Alpha	Mar 01 2007
whale	Online	3072	Cluster/Gigabit Ethernet	Opteron	Mar 22 2007
wobbe	Online	208	Cluster/Myrinet	Opteron	Mar 23 2007
zebra	Online	128	Cluster/Myrinet G2	Opteron	Feb 21 2007
Totals		8096			

Done

SHARCNET: Current Performance - Mozilla Firefox

File Edit View Go Bookmarks Tools Help

http://www.sharcnet.ca/Performance/cur_perf.php

Windows Media file:///C:/Documents... Customize Links Windows Windows Marketplace Free Hotmail Free AOL & Unlimited...

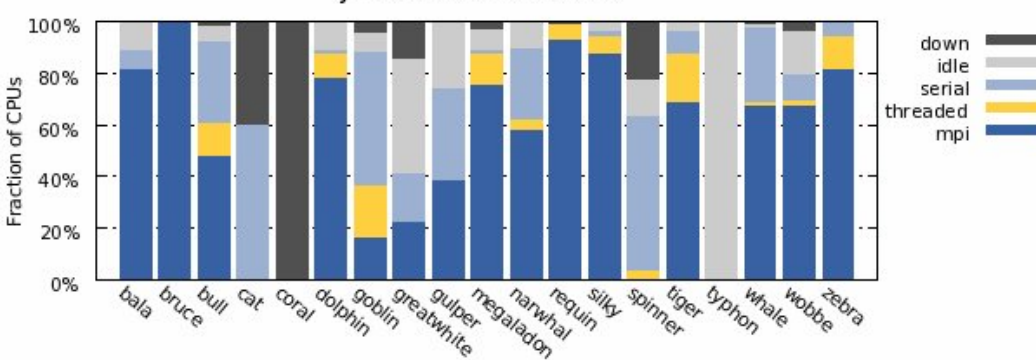
Current Performance

Home

- About
- Contact Us
- Events
- Facilities
- Help
- Media
- Performance
 - Current
 - Historical
 - Metrics
- Research

My Sharcnet
Please log in

Jobs on our Clusters



Fraction of CPUs

Legend: down, idle, serial, threaded, mpi

Current Cluster Performance

Cluster	Alert	Total CPUs	Up CPUs	Idle CPUs	Running Jobs	MPI Jobs	Serial Jobs	Threaded Jobs	Susp'd Jobs	Q'd Jobs	Q'd CPUs	% Cap.	GFLOPS	Last Polled	Status
bala	CPUs available!	128	128	15	21	12	9	0	0	0	0	88%	497	Mar 23 12:00	Online
bruce		128	128	0	9	9	0	0	0	1	16	100%	563	Mar 23 12:00	Online
bull		384	376	23	173	40	121	12	0	26	242	93%	1,694	Mar 23 12:00	Online
cat		160	96	0	96	0	96	0	0	1,321	1,321	100%	400	Mar 23 12:00	Downtime soon
coral		60	0	0	0	0	0	0	0	0	0	0%	0	Mar 23 12:00	Conditions
dolphin		128	128	15	14	10	1	3	1	3	25	88%	497	Mar 23 12:00	Online
goblin	CPUs available!	122	116	9	74	5	63	6	0	0	0	92%	502	Mar 23 12:00	Online
greatwhite	CPUs available!	456	388	203	100	15	85	0	0	0	0	47%	263	Mar 23 12:00	Online
gulper	CPUs available!	42	42	11	16	1	15	0	0	0	0	73%	136	Mar 23 12:00	Online
megaladon		128	124	11	10	8	1	1	1	3	64	91%	497	Mar 23 12:00	Downtime soon

Done

SHARCNET Facilities: Intended use

SHARCNET Facilities are equally available to all users, but the principle is to use them efficiently. Users should select the right machine to submit their jobs. For example, parallel jobs are supposed to run on clusters with Quadrics/Myrinet interconnection.

A guideline is

Cluster	CPUs	RAM /node	Storage	Interconnect	Intend Use
requin (Capability)	1536	8 GB	70 TB	Quadrics	Large scale MPI
narwhal (Utility)	1068	8 GB	70 TB	Myrinet	MPI, SMP
whale (Throughput)	3072	4 GB	70 TB	GigE	Serial
bull (SMP-friendly)	384	32 GB	70 TB	Quadrics	High RAM/BW MPI & SMP
silky (SMP)	128	256 GB	4 TB	NUMALink	large memory SMP OpenMP/threads,
bala, bruce, dolphin, megaladon, tiger, zebra	128	8 GB	4 TB	Myrinet	General purpose

Specifications: <http://www.sharcnet.ca/Facilities/index.php>

Software Resources

- **OS** – Linux
HP Opteron: HP XC3.0
- **Compilers**
 - Opteron: **Pathscale** (pathcc, pathCC, pathf90), **PGI** (pgcc, pgCC, pgf77/pgf90)
 - Alpha: **compaq** (ccc, cxx, fort)
 - Itanium/Xeon: **Intel** (icc, ifort)
- **Scheduler**
 - LSF, SQ
- **Key parallel development support**
 - MPI (HPMPI, MPICH, OPENMPI)
 - Multi-threading (pthreads, OpenMP)

Software Resources (continue)

- **Libraries**
 - ACML (AMD), CXML (Alpha), SCSL(SGI), ATLAS, GSL, ScaLAPACK, FFTW, PETSc, ...
- **Debugging/Profiling Tools**
 - Debugging (DDT, gdb, ...)
 - Profiling/Optimization (OPT, gprof,...)
- **Application packages**
R, Blast, Gromacs, NWChem, Octave, ...
- **Commercial packages** (cost share)
 - **Gaussian**: *Computational Chemistry Software*
 - **Fluent**: *Computational Fluid Dynamics Software*
 - Bring your license to SHARCNET (such as lumerical, ...)
- **Others**. you ask/provide, we install

Details: <http://www.sharcnet.ca/Facilities/software/softwarePage.php>

File System Basics

Policy

- Same username/password across all systems, and webportal account.
- User self-management on webportal (site leader, sponsor, group member)
- Common home directory across SHARCNET (exceptions: wobbe, cat)
- SHARCNET-maintained software is in /opt/sharcnet
- /home backup

File system

pool	quota	expiry	purpose
/home	200 MB	none	sources, small config files
/work	200 GB?	none	active data files
/scratch	none	2 weeks	large active data files
/tmp	160 GB	10 days	node-local scratch

- /scratch and /work are local to each cluster
 - not backed up
 - **important**: run jobs from /work or huge jobs from /scratch

Job Scheduling (queue)

Same types of queues: **mpi**, **threaded**, **serial**, but priority is cluster specific.
 For throughput (serial) clusters such as **whale**, we have

```
[jemmyhu@wha780 ~]$ bqueues
```

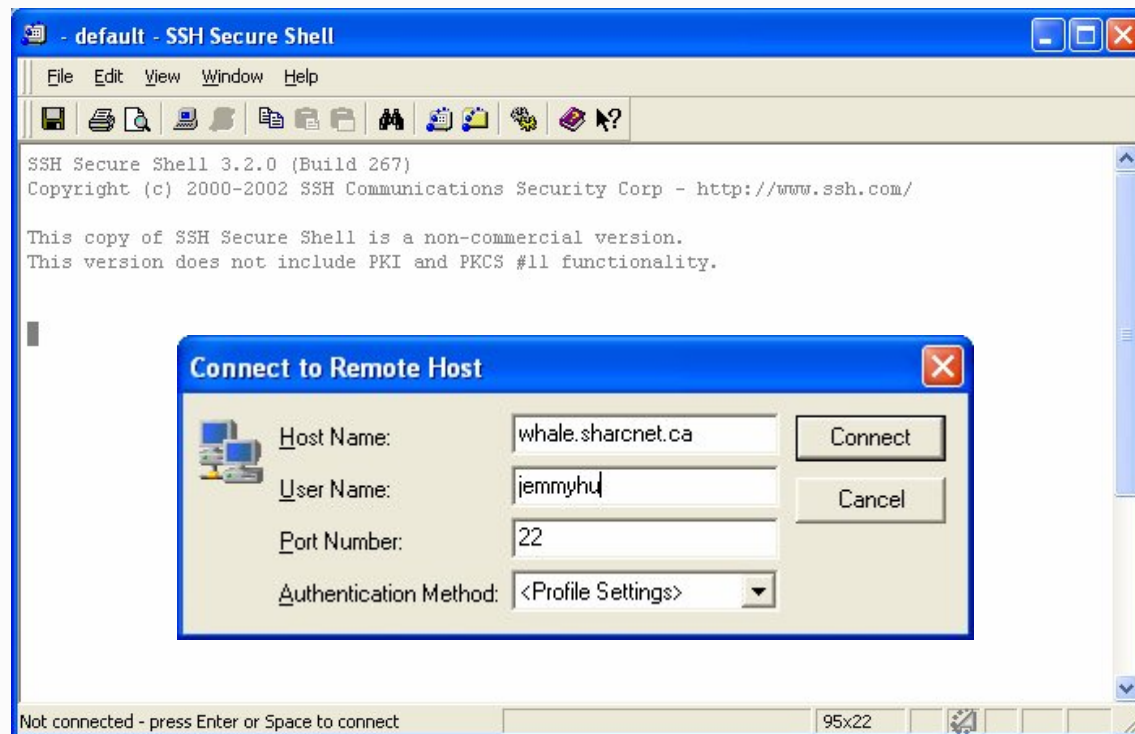
QUEUE_NAME	PRIO	STATUS	MAX	JL/U	JL/P	JL/H	NJOBS	PEND	RUN	SUSP
staff	150	Open:Active	-	-	-	-	0	0	0	0
test	100	Open:Active	-	-	-	-	0	0	0	0
serial	80	Open:Active	-	-	-	-	1871	0	1871	0
threaded	40	Open:Active	-	-	-	-	64	0	64	0
mpi	40	Open:Active	-	-	-	-	840	0	840	0

```
[jemmyhu@wha780 ~]$
```

Test queue: preemptive, for debugging purpose only, limited cpu time (30 mins or 1 hr)

Login

- Login: only by **SSH**
 - Linux: `ssh username@machine.sharcnet.ca`
 - Windows: SSH is available free from (<http://ftp.ssh.com/pub/ssh/SSHSecureShellClient-3.2.9.exe>) and many graphical clients available (PuTTY, etc.)



File Transfer

- **Linux:** `scp`

From your linux system: `scp filename narwhal.sharcnet.ca:/work/username`

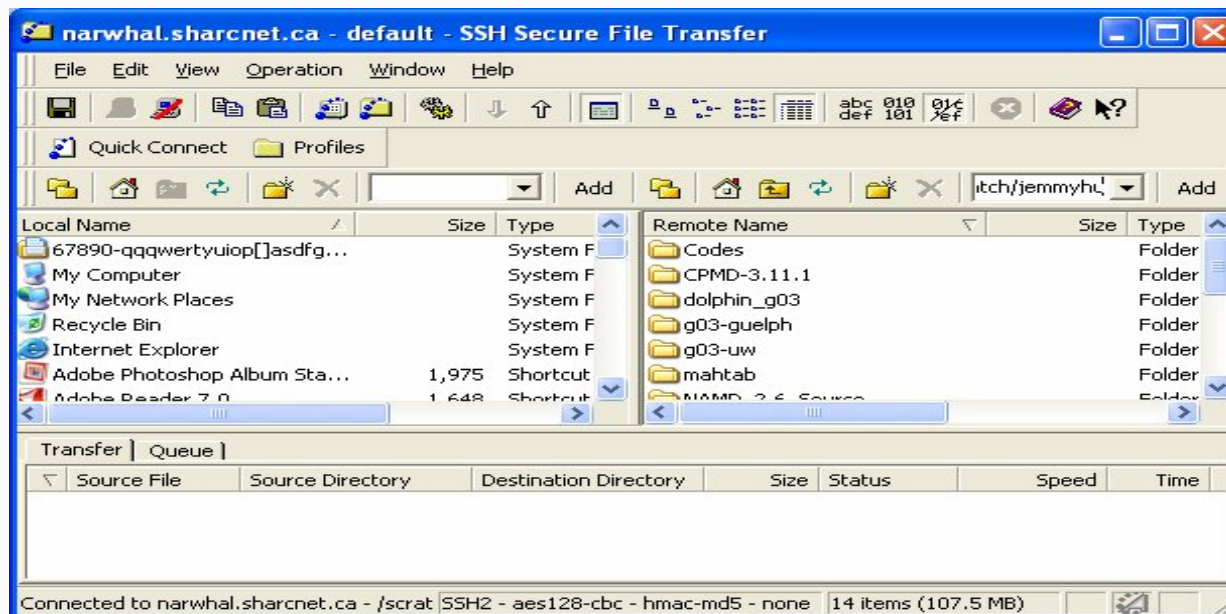
Between Sharcnet cluster: `scp filename narwhal.sharcnet:/work/username`

```
[jemmyhu@wha780 ~]$ pwd
```

```
/work/jemmyhu
```

```
[jemmyhu@wha780 ~]$ scp mycode.c narwhal.sharcnet:/work/jemmyhu
```

- **Windows:** SSH File Transfer



Compilers: SHARCNET compile script

SHARCNET provides a “compile” script (use as “generic” compiler)

- tries to do what it “should” - optimizations, compiler, etc.
- recommended unless you know better

Command	Language	Extension	Example
cc	c	c	cc -o code.exe code.c
CC	C++	.C, .cc, .cpp, .cxx, c++	cc -o code.exe code.cpp
c++	C++		
cxx	C++		
f77	Fortran 77	.f, .F	f77 -o Fcode.exe Fcode.f
f90	F90/F95	.f90, .f95, .F90, F95	f90 -o Fcode.exe Fcode.f90
mpicc	c	c	mpicc -o mpicode.exe mpicode.c
mpiCC	C++	C++	mpiCC -o mpicode.exe mpicode.cc
mpif77	Fortran 77	f77	mpif77 -o mpicode.exe mpicode.f
mpif90	Fortran90/95	F90/f95	mpif90 -o mpicode.exe mpicode.f90

For OpenMP code: **f90 -openmp -o openmpcode openmpcode.f90**

Run Jobs Basics

- **Log on to the desired cluster/machine**
 - ensure files are in /work/username
 - **do not run jobs out of /home** --- it can be very slow
 - jobs are submitted using the SQ system (wrapper, e.g. around LSF)

- **Job Category - Queue**

Job Type	Queue	CPUs (default=1)
Parallel	mpi	>2
Parallel	threaded (Machine dependent)	2 (requin, wobbe, cat) 2-4 (most clusters) >2 (silky)
Serial	serial	1

Commonly used SQ commands

- **bqueues** – list available queues.
- **sqsub** – submit a program (“job”) to a specific queue.
- **sqjobs** – list the status of submitted jobs.
- **sqkill** – kill a program by job ID.

[...@bl125]\$ **bqueues**

QUEUE_NAME	PRIO	STATUS	MAX	JL/U	JL/P	JL/H	NJOBS	PEND	RUN	SUSP
staff	150	Open:Active	-	-	-	-	0	0	0	0
test	100	Open:Active	-	-	-	-	0	0	0	0
gaussian	100	Open:Active	-	-	-	-	64	0	64	0
threaded	80	Open:Active	-	-	-	-	44	0	44	0
mpi	80	Open:Active	-	-	-	-	339	226	113	0
serial	40	Open:Active	-	-	-	-	143	0	143	0

sqsub (submit job)

sqsub [-o ofile][-i ifile][-t][-q queue][-n ncpus] command...

Examples (Demo)

```
sqsub -q mpi -n 4 -o hello-mpi.log ./hello-mpi
```

```
sqsub -q threaded -n 4 -o hello-openmp.log ./hello-openmp
```

```
sqsub -q serial -o hello.log ./hello
```

Submit to **'test'** queue

```
sqsub -q mpi -t -n 4 -o hello-mpi.log ./hello-mpi
```

Show page

```
[jemmyhu@wha780 examples]$ pwd
/scratch/jemmyhu/examples
```

```
[jemmyhu@wha780 examples]$ sqsub -q mpi -t -n 4 -o hello_mpi_c.log ./hello_mpi_c
Job <134177> is submitted to queue <test>.
```

```
[jemmyhu@wha780 examples]$ sqjobs
jobid queue state ncpus nodes time command
-----
```

```
134177 test R 4 wha2 2s ./hello_mpi_c
```

2972 CPUs total, 196 idle, 2776 busy; 2020 jobs running; 1 suspended, 232 queued.

```
[jemmyhu@wha780 examples]$ bqueues
```

QUEUE_NAME	PRIO	STATUS	MAX	JL/U	JL/P	JL/H	NJOBS	PEND	RUN	SUSP
staff	150	Open:Active	-	-	-	-	0	0	0	0
test	100	Open:Active	-	-	-	-	4	0	4	0
serial	80	Open:Active	-	-	-	-	2228	232	1996	0
threaded	80	Open:Active	-	-	-	-	16	0	16	0
mpi	40	Open:Active	-	-	-	-	960	0	760	200

```
[jemmyhu@wha780 examples]$
```

sqjobs (view job status)

sqjobs [-r][-q][-z][-v][-u user][-n][--summary][[jobid...]]

Show page

```
[jemmyhu@wha780 examples]$ sqjobs
```

```
[jemmyhu@wha780 examples]$ sqjobs -u jemmyhu
```

```
[jemmyhu@wha780 examples]$ sqjobs -u tao
```

```
jobid queue state ncpus    nodes  time command
```

```
-----
```

```
133327 mpi    R    8 wha[31,33] 3439s ./prog
```

```
2972 CPUs total, 0 idle, 2972 busy; 2020 jobs running; 0 suspended, 275 queued.
```

```
[jemmyhu@wha780 examples]$ sqjobs -n
```

```
2972 CPUs total, 0 idle, 2972 busy; 2020 jobs running; 0 suspended, 275 queued.
```

```
[jemmyhu@wha780 examples]$ sqjobs --summary
```

user	q	r	z	ncpus	njobs
-----	---	---	-	-----	-----
ikotsire	0	728	0	728	728
racinej	0	1	0	624	1
jdc	0	462	0	462	462
garicb	0	320	0	320	320
ricky	253	313	0	313	566

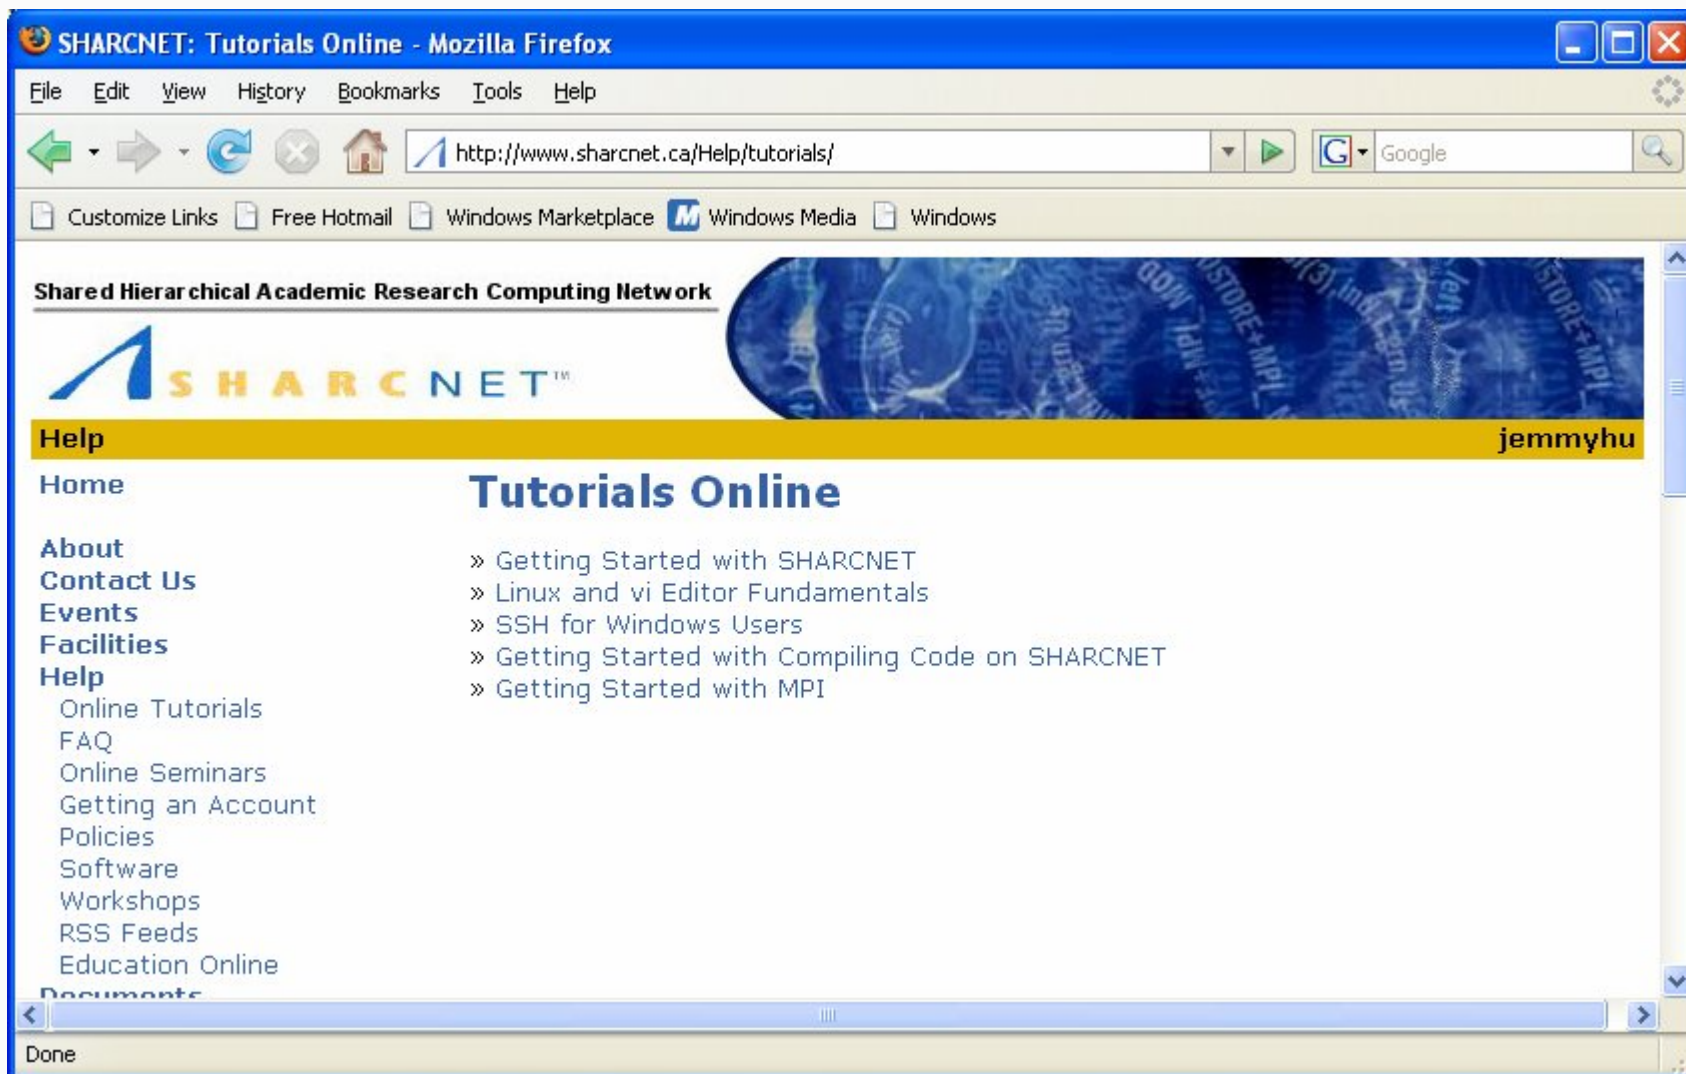
```
.....
```

sqkill (kill a job)

sqkill jobid [jobids...]

```
[jemmyhu@wha780 examples]$ sqjobs
[jemmyhu@wha780 examples]$ sqsub -q mpi -t -n 4 -o hello_mpi_c.log ./hello_mpi_c
Job <134227> is submitted to queue <test>.
[jemmyhu@wha780 examples]$ sqjobs
jobid queue state ncpus nodes time command
-----
134227 test  Q   4     3s ./hello_mpi_c
2972 CPUs total, 0 idle, 2972 busy; 2020 jobs running; 0 suspended, 269 queued.
[jemmyhu@wha780 examples]$ sqkill 134227
Job <134227> is being terminated
[jemmyhu@wha780 examples]$ sqjobs
[jemmyhu@wha780 examples]$
```

May need to wait for a few seconds for the system to kill a job




SHARCNET: Tutorials Online - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://www.sharcnet.ca/Help/tutorials/

Customize Links Free Hotmail Windows Marketplace Windows Media Windows

Shared Hierarchical Academic Research Computing Network



Help jemmyhu

Home

Tutorials Online

- » Getting Started with SHARCNET
- » Linux and vi Editor Fundamentals
- » SSH for Windows Users
- » Getting Started with Compiling Code on SHARCNET
- » Getting Started with MPI

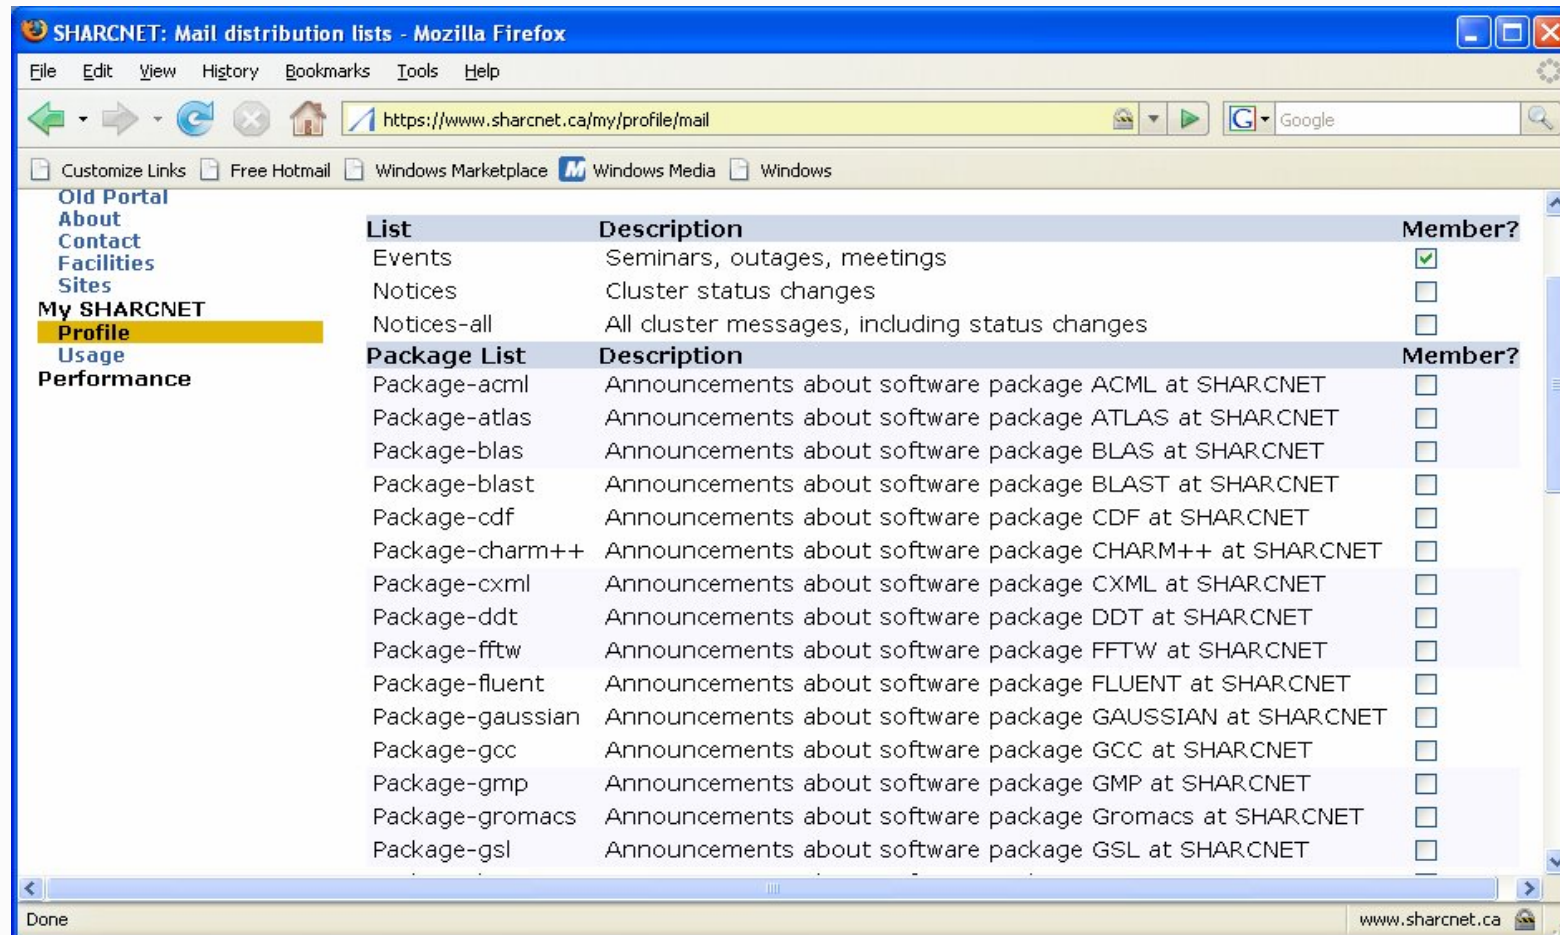
About
Contact Us
Events
Facilities
Help

- Online Tutorials
- FAQ
- Online Seminars
- Getting an Account
- Policies
- Software
- Workshops
- RSS Feeds
- Education Online

Documents

Done

User Email lists: My Mailing Lists in the Portal



SHARCNET: Mail distribution lists - Mozilla Firefox

File Edit View History Bookmarks Tools Help

https://www.sharcnet.ca/my/profile/mail

Customize Links Free Hotmail Windows Marketplace Windows Media Windows

Old Portal
 About
 Contact
 Facilities
 Sites
 My SHARCNET
Profile
 Usage
 Performance

List	Description	Member?
Events	Seminars, outages, meetings	<input checked="" type="checkbox"/>
Notices	Cluster status changes	<input type="checkbox"/>
Notices-all	All cluster messages, including status changes	<input type="checkbox"/>
Package List	Description	Member?
Package-acml	Announcements about software package ACML at SHARCNET	<input type="checkbox"/>
Package-atlas	Announcements about software package ATLAS at SHARCNET	<input type="checkbox"/>
Package-blas	Announcements about software package BLAS at SHARCNET	<input type="checkbox"/>
Package-blast	Announcements about software package BLAST at SHARCNET	<input type="checkbox"/>
Package-cdf	Announcements about software package CDF at SHARCNET	<input type="checkbox"/>
Package-charm++	Announcements about software package CHARM++ at SHARCNET	<input type="checkbox"/>
Package-cxml	Announcements about software package CXML at SHARCNET	<input type="checkbox"/>
Package-ddt	Announcements about software package DDT at SHARCNET	<input type="checkbox"/>
Package-fftw	Announcements about software package FFTW at SHARCNET	<input type="checkbox"/>
Package-fluent	Announcements about software package FLUENT at SHARCNET	<input type="checkbox"/>
Package-gaussian	Announcements about software package GAUSSIAN at SHARCNET	<input type="checkbox"/>
Package-gcc	Announcements about software package GCC at SHARCNET	<input type="checkbox"/>
Package-gmp	Announcements about software package GMP at SHARCNET	<input type="checkbox"/>
Package-gromacs	Announcements about software package Gromacs at SHARCNET	<input type="checkbox"/>
Package-gsl	Announcements about software package GSL at SHARCNET	<input type="checkbox"/>

Done www.sharcnet.ca

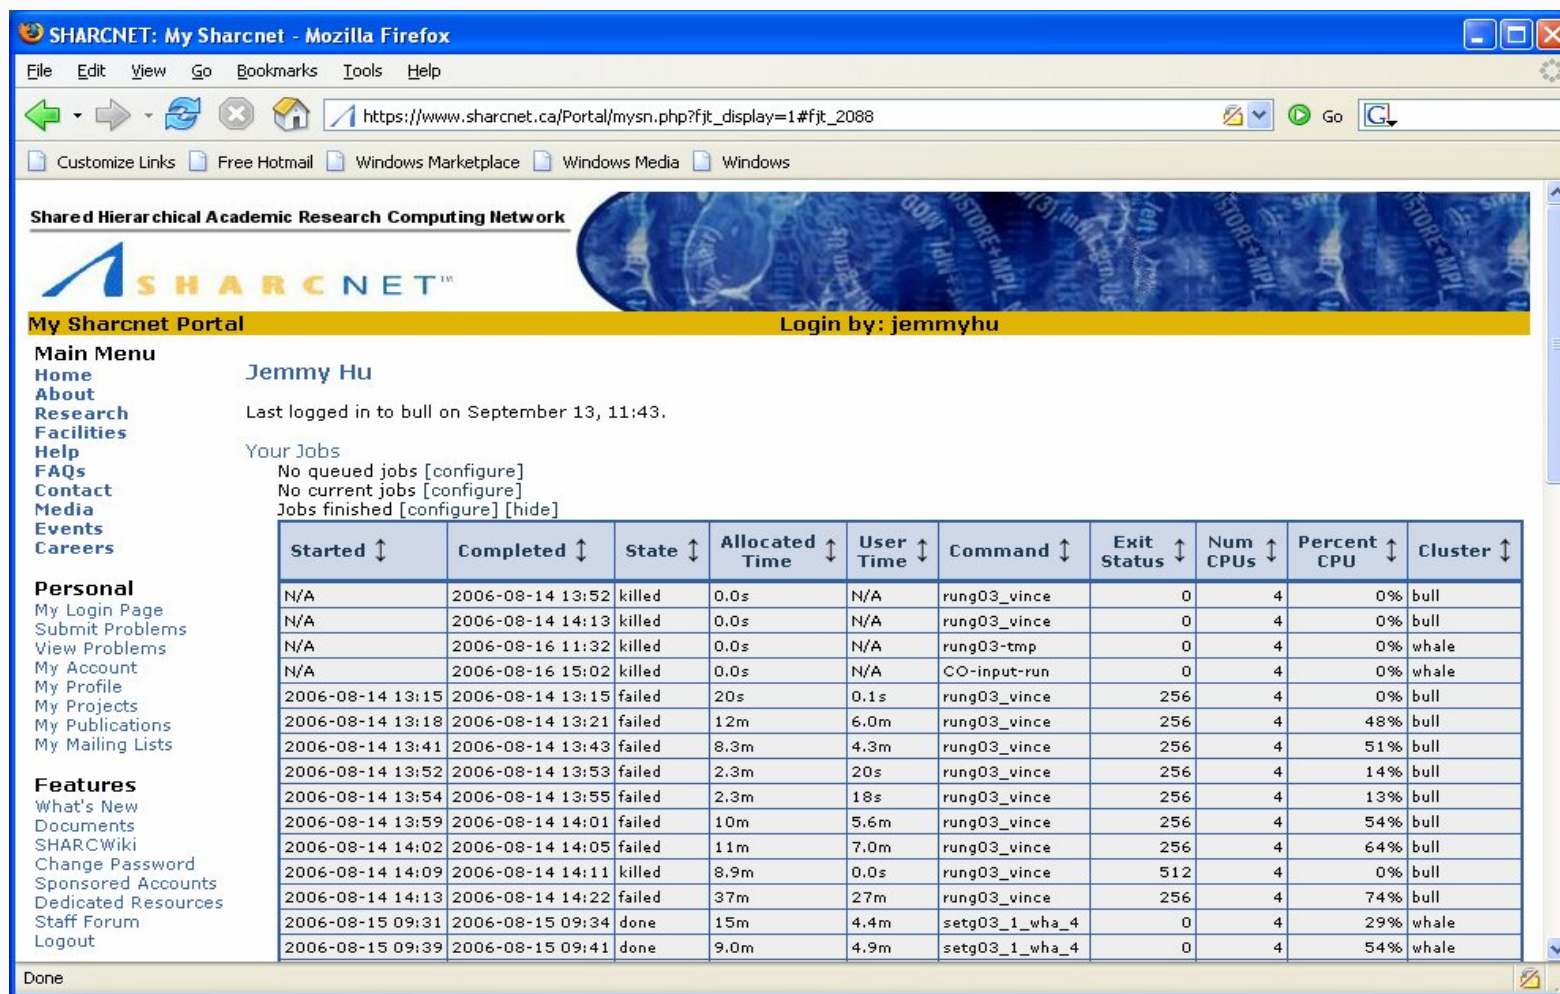
SHARCNET Support: Info on the Web

<http://www.sharcnet.ca>

- For general information: SHARCNET FAQ
<http://www.sharcnet.ca/Help/faq.php>
- Supporting personal:
<http://www.sharcnet.ca/Help/index.php>
- Up-to-date Facility information: facility page
<http://www.sharcnet.ca/Facilities/index.php>
- Software page
<http://www.sharcnet.ca/Facilities/software/softwarePage.php>
- Online problem tracking system (login needed)
https://www.sharcnet.ca/Portal/problems/problem_search.php

SHARCNET Support: web tools

My Jobs



SHARCNET: My Sharcnet - Mozilla Firefox

File Edit View Go Bookmarks Tools Help

https://www.sharcnet.ca/Portal/mysn.php?fjt_display=1#fjt_2088

Customize Links Free Hotmail Windows Marketplace Windows Media Windows

Shared Hierarchical Academic Research Computing Network

SHARCNET™

My Sharcnet Portal Login by: jemmyhu

Main Menu

- Home
- About
- Research
- Facilities
- Help
- FAQs
- Contact
- Media
- Events
- Careers

Personal

- My Login Page
- Submit Problems
- View Problems
- My Account
- My Profile
- My Projects
- My Publications
- My Mailing Lists

Features

- What's New
- Documents
- SHARCWiki
- Change Password
- Sponsored Accounts
- Dedicated Resources
- Staff Forum
- Logout

Jemmy Hu

Last logged in to bull on September 13, 11:43.

Your Jobs

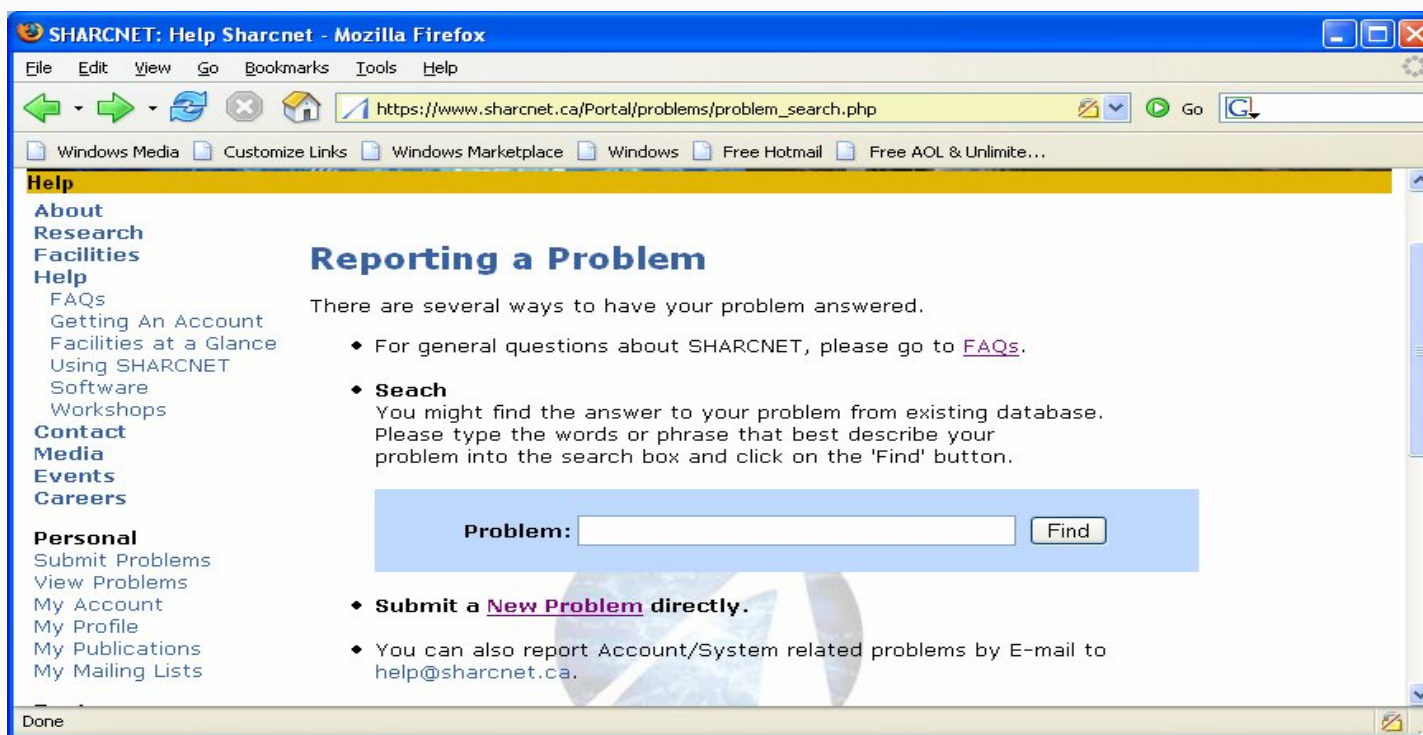
- No queued jobs [configure]
- No current jobs [configure]
- Jobs finished [configure] [hide]

Started ↓	Completed ↓	State ↓	Allocated Time ↓	User Time ↓	Command ↓	Exit Status ↓	Num CPUs ↓	Percent CPU ↓	Cluster ↓
N/A	2006-08-14 13:52	killed	0.0s	N/A	rung03_vince	0	4	0%	bull
N/A	2006-08-14 14:13	killed	0.0s	N/A	rung03_vince	0	4	0%	bull
N/A	2006-08-16 11:32	killed	0.0s	N/A	rung03-tmp	0	4	0%	whale
N/A	2006-08-16 15:02	killed	0.0s	N/A	CO-input-run	0	4	0%	whale
2006-08-14 13:15	2006-08-14 13:15	failed	20s	0.1s	rung03_vince	256	4	0%	bull
2006-08-14 13:18	2006-08-14 13:21	failed	12m	6.0m	rung03_vince	256	4	48%	bull
2006-08-14 13:41	2006-08-14 13:43	failed	8.3m	4.3m	rung03_vince	256	4	51%	bull
2006-08-14 13:52	2006-08-14 13:53	failed	2.3m	20s	rung03_vince	256	4	14%	bull
2006-08-14 13:54	2006-08-14 13:55	failed	2.3m	18s	rung03_vince	256	4	13%	bull
2006-08-14 13:59	2006-08-14 14:01	failed	10m	5.6m	rung03_vince	256	4	54%	bull
2006-08-14 14:02	2006-08-14 14:05	failed	11m	7.0m	rung03_vince	256	4	64%	bull
2006-08-14 14:09	2006-08-14 14:11	killed	8.9m	0.0s	rung03_vince	512	4	0%	bull
2006-08-14 14:13	2006-08-14 14:22	failed	37m	27m	rung03_vince	256	4	74%	bull
2006-08-15 09:31	2006-08-15 09:34	done	15m	4.4m	setg03_1_wha_4	0	4	29%	whale
2006-08-15 09:39	2006-08-15 09:41	done	9.0m	4.9m	setg03_1_wha_4	0	4	54%	whale

Done

SHARCNET Support: Problem Reporting

- Using online **SHARCNET Problem Tracking System** to submit any questions/problems.
 - Problems will be handled with appropriately by supporting staffs.
 - A record of the question together with the comments and final resolution will be tracked.
 - User view: **submit, search, view, comment, reopen**



- Email to help@sharcnet.ca
- Email or phone supporting staff: <http://www.sharcnet.ca/Help/index.php>

Submit Publication

- To add a new application

https://www.sharcnet.ca/Portal/publications/publication_add.php

- To view publications

https://www.sharcnet.ca/Portal/publications/publication_update.php

Thank You.

Questions?