Introduction to the SHARCNET Environment

2010-May-25
Pre-(summer)school webinar
Speaker: Alex Razoumov
University of Ontario Institute of Technology

- available hardware and software resources
- our web portal
- where to look for help
- how to log in to systems
- compile code, run, submit and manage your jobs
- a quick summary of our visualization tools

SHARCNET

shared hierarchical academic research computing network

- consortium of 17 academic partner institutions in Ontario
- goal: promote and enable **HPC** for research
- one of seven consortia in Canada

http://www.sharcnet.ca

compute • calcul
CANADA

What we provide

- HPC
  - resources (>13,000 cpu cores, variety of systems)
  - training
  - support (local staff, online tickets)
- SHARCNET is free for Canadian researchers and their students
  - minimal barrier to utilization
  - equal availability
  - support for non-traditional computing domains, e.g. digital humanities, biomedical, etc.

Computing Resources

- HPC systems
  - clusters (**distributed memory**): large number of simple nodes connected by a network, a range of interconnect speeds supporting everything from large MPI jobs to serial farming
  - symmetric multiprocessing (SMP) systems (**shared memory**): for threaded applications
  - accelerated: Cell, GPU (angel), FPGA
- visualization systems
  - visualization cluster (rainbow)
  - visualization workstations
  - large systems (quad-HD, 3D) at McMaster, UWO
- AccessGrid rooms (videoconferencing)
- all connected via dedicated 1 Gb/s network (10 Gb/s internally)

HPC Resources

- wide range of software (including commercial) provided and supported by SHARCNET at no cost
- single account access to all systems and web portal
- backed-up /home file system shared amongst all systems, significant temporary and longer term storage on each system, as well as archival storage
- production systems use a non-interactive batch job scheduling system, which promotes fairness by adjusting job priority based on historical use

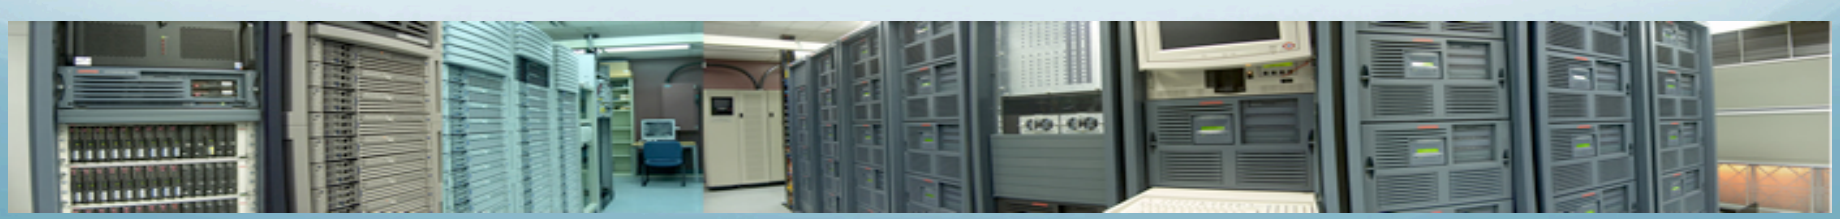

Certification

resource access is limited depending on user's certification level:

Level	number of concurrent CPU cores	maximum wallclock job duration in days
0	8	1
1	256	7
2	unlimited	7

attain level 1 by viewing the new user seminar video and taking the online quiz

level 2 is provided temporarily for users who need to run large jobs

HPC Training

- events
  - summer school
  - workshops and symposia
- seminars
  - new user introduction
  - general interest
  - guest speakers
- credit courses
  - graduate and undergraduate
- online tutorials
  - variety of topics, essentials
- outreach
  - secondary school tours and presentations

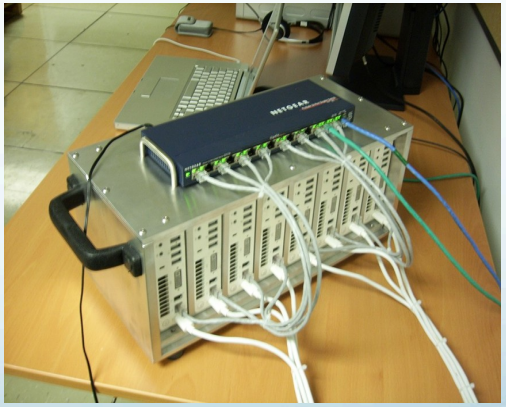

HPC Support

- funding for postdocs, students, distinguished visitors and faculty
- dedicated project support
  - priority access to necessary resources for demanding projects (DR=dedicated resources)
  - short and long term programming support
- consulting SHARCNET staff
  - in person, by email, phone, AccessGrid
  - problem ticket system
  - wide range of expertise

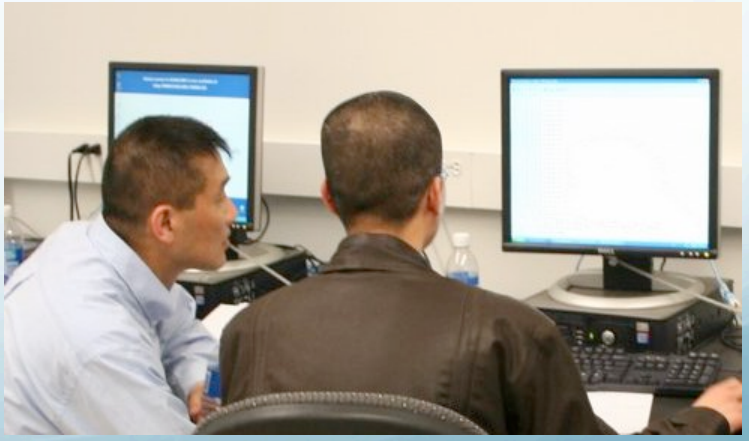

Computing environment

- **Operating system**: 64-bit Linux, Tru64
- **Languages**: Fortran, C/C++, Java, Matlab, etc.
- **Compilers**:
  - SHARCNET unified compilation environment: cc, CC, f77, f90, mpicc, mpiCC, mpif77, mpif90
  - Underlying compilers: Intel, PathScale, PGI, GNU compilers
- **Parallel code development**:
  - MPI (hp-MPI, Quadric MPI, MPICH, OpenMPI, etc.)
  - POSIX threads API, OpenMP
- **Batch scheduling** with sq - SHARCNET's unified batch execution environment (sqsub, sqjobs, sqkill)
- Large number of installed libraries, commercial and open-source packages, debugging and profiling tools

File systems

filesystem	quota	expiry	comments
/home	200 MB	none	source files, global, backed up
/work	200 GB	none	work directories, currently local, not backed up
/scratch	none	35 days	work directories, local, not backed up

- all jobs should be run in /work or /scratch
- once quotas exceeded, lose ability to submit jobs
- transition to global /work in the next few months
- longer term storage: archive

Typical usage scenario

Copy files to one of SHARCNET systems (bull)

```
scp projects.tar.gz bull.sharcnet.ca:
```

Login to remote system bull, with X connection

```
ssh bull.sharcnet.ca -X -l username
```

Edit your files using your favourite editor, e.g. emacs, vi

```
tar zxvf projects.tar.gz
cd projects/elec_price
emacs -nw price_main.cc
```

Compile your code

```
c++ price_main.cc fun1.cc fun2.cc … fun5.cc -o price
```

Test drive your code

```
sqsub -t -q serial -r 5m -o testOutput.log ./price
```

Run your code (production runs)

```
sqsub -q serial -r 3h -o output.log ./price
```

sqsub/sqjobs commands

3 general queues: **serial**, **mpi**, **threaded**

submitting a parallel job

```
sqsub -q mpi -n 24 -r 1h -o outputFile ./MPIcode
sqsub -q mpi -n 16 -N 2 -r 1h -o outputFile ./MPIcode
sqsub -q threaded -n 8 -r 1h -o outputFile ./OpenMPcode
```

listing status of submitted jobs

sqjobs

```
 jobid queue   state ncpus   prio nodes time command
------ ----- ----- ----- ------ ----- ---- --------
239356   mpi     Q     4     50.000          32s ./test
.....
239356   mpi     R     4     40.000  saw5    5s  ./test
```

sqjobs -help

Connecting from desktop

Use only ssh encrypted connections

Linux, Mac OS: standard command line tools (ssh, sftp, scp)

Windows: PuTTY (ssh), WinSCP, Cygwin/X, Xming

https://www.sharcnet.ca/my/front/

Google

SHARCNET™

home || sign-in || search: Go

computing tomorrow's solutions

Dr. Sinclair's research is focused on the design and development of text analysis tools for the digital humanities, especially literary analysis where conventional text mining techniques are not appropriate.

With experts predicting the H1N1 pandemic will worsen, health officials are seeking input from infectious disease modellers such as Bauch to model the disease transmission and control strategies.

SHARCNET is a consortium of 17 Ontario research and education institutions that share a network of high performance computers. This "distributed computing" infrastructure enables world-class research.

Facilities
News & Events
Help
Research
Support Programs
About Us

News

SHARCNET Summer School 2010
SHARCNET Outreach Campaign: SHARCFest 2010
SHARCNET sponsors MIND THE GAP Workshop
SHARCNET Dedicated Programming Support: Round II Call for...
SHARCNET sponsors Canada-Wide Science Fair 2010
GridCentric Inc. Announces Copper(TM) Cluster Management ...
SHARCNET sponsors London Science Fair
Compute Canada Mid-Term Review

Events & Activities

[12:00pm] SS2010 Preschool: Introduction to the SHARCNET ...
[12:00pm] SS2010 Preschool: Introduction to Linux
[12:00pm] SS2010 Preschool: Assessing Your Programming Sk...
[12:00pm] SS2010 Preschool: HPC Architecture Overview
Summer School 2010
HPCS 2010
Magnetic North - The Dynamics and Magnetic Materials

- My Account
  - Settings
    - Details
    - Group
    - Password
  - Reporting
  - Usage
    - Activity
    - CPU Usage
    - Disk Usage
  - Applications
  - Projects
- Facilities
- News+Events
- Help
- Research
- Support Programs
- About Us
- Users+Usage

Site Leader Page

Activity for user razoumov (awr-001-01), project awr-001-aa

Last logged in to bull on May 19, 2010 11:24.

Projects

Project razoumov: Default Resource Allocation Project (**Activity**)

Other Roles

sn5745
sn5882

Usage Summary

User	1 Week	2 Weeks	1 Month	2 Months	4 Months
razoumov	3.6 days (28 jobs)	16 days (60 jobs)	18 weeks (107 jobs)	24 weeks (245 jobs)	38 weeks (485 jobs)

Project Usage

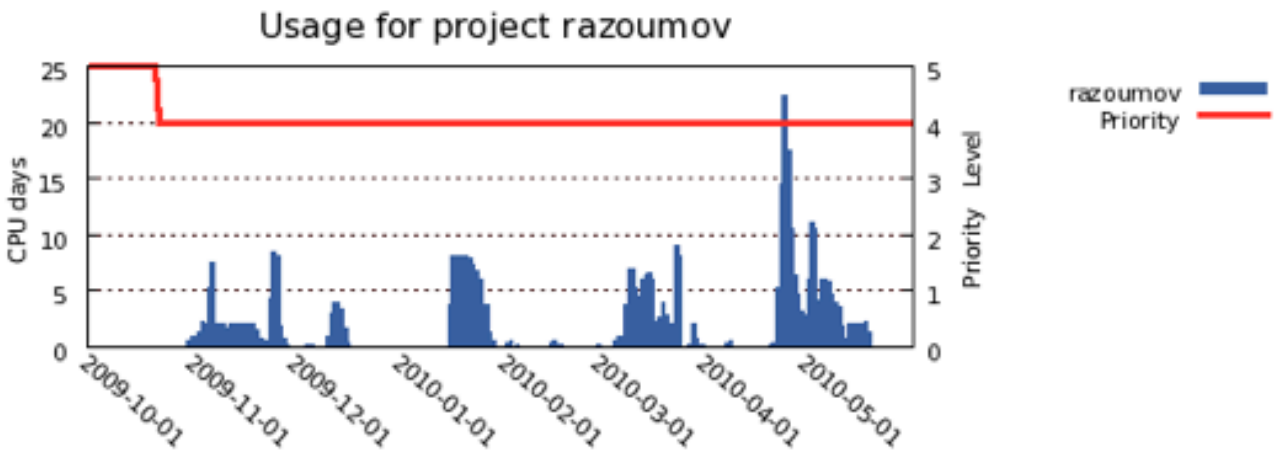

All Systems

Which system should I use?

Core Systems

System	State	Nagios Errors	Cores	Architecture	Nodes	Notices
bala	Offline	INTERFACE_STATUS JOBS_RUN METALOOKUP_TEST MSA_ARRAY_STATUS NFSMOUNTS NODES NODE_CPUS NODE_STATS QUEUE_STATUS VALID_USER	128	Cluster/Myrinet 2g (gm)	Opteron	25-May-2010
bruce	Online		128	Cluster/Myrinet 2g (gm)	Opteron	20-May-2010
bull	Online		384	Cluster/Quadrics Elan4	Opteron	04-May-2010
dolphin	Offline	(Not Monitored)	128	Cluster/Myrinet 2g (gm)	Opteron	14-Dec-2009
kraken	Offline	(Not Monitored)	168	Cluster/Myrinet 2g (gm)	Opteron	
megaladon	Online		128	Cluster/Myrinet 2g (gm)	Opteron	26-Aug-2009
narwhal	Online		1068	Cluster/Myrinet 2g (gm)	Opteron	23-Apr-2010
requin	Online	INTERFACE_STATUS LOCAL_DISK_USAGE	1536	Cluster/Quadrics Elan4	Opteron	21-May-2010
saw	Online	G_LOAD_FIFTEEN	2688	Cluster/InfiniBand	Xeon	12-May-2010
tiger	Decommissioned.	(Not Monitored)	0	Cluster/Myrinet 2g (gm)	Opteron	17-Feb-2010
whale	Online	IDLE_JOBS NODE_RESPONSE	3072	Cluster/Gigabit Ethernet	Opteron	21-May-2010
zebra	Online		128	Cluster/Myrinet 2g (gm)	Opteron	13-Jul-2009

Specialty Systems

System	State	Nagios Errors	Cores	Architecture	Nodes	Notices
angel	Online	176		Accelerator/InfiniBand	Xeon/gpu	11-May-2010
coral	Decommissioned.	(Not Monitored)	60	Cluster/Quadrics Elan3	Itanium2	12-Dec-2008
deeppurple	Decommissioned.	(Not Monitored)	44	Cluster/Quadrics, Gigabit Ethernet	Alpha	27-Feb-2006
greatwhite	Decommissioned.	(Not Monitored)	456	Cluster/Quadrics Elan3	Alpha	01-Dec-2008
hammerhead	Decommissioned.	(Not Monitored)	108	Cluster/Quadrics, Gigabit Ethernet	Alpha	
hound	Online	(Not Monitored)	480	Cluster/InfiniBand	Xeon, Opteron	19-May-2010
idra	Decommissioned.	(Not Monitored)	128	Cluster/Quadrics, 100bT Ethernet	Alpha	26-May-2006
leopard	Offline	(Not Monitored)	188	Cluster/Gigbit Ethernet	Xeon	15-Dec-2009
mako	Online	(Not Monitored)	240	Cluster/GigE	Xeon/Nehalem	19-Feb-2010

All Software

Application	
ABINIT	Serial and Parallel ABINIT Package
ADF/BAND	Amsterdam Density Functional Software (commercial, group adf)
AMBER	Molecular Dynamics Simulation On Biomolecules (commercial)
BIOPERL	Toolkit of Perl Modules For Bioinformatics
BLAST	Basic Local Alignment Search Tool
BLCR	Berkeley Lab Checkpoint Restart
CDF	The Common Data Format
CILK	Language for Multithreaded Parallel Programming
CPMD	Parallelized Plane Wave Implementation of DFT
DeMON	Package for Density Functional Theory (DFT) Calculations
DLPOLY	General Purpose Molecular Dynamics Simulation Package
ESPRESSO	Plane-Wave Self-Consistent Field (PWSCF)
FIREFLY	Ab Initio and DFT Computational Chemistry Program (PC GAMESS)
FLUENT	Computational Fluid Dynamics Software (commercial, fluent group)
GAUSSIAN	Computational Chemistry Software (commercial)
GEANT4	The Geant4 Toolkit
GROMACS	Molecular Dynamics for Biomolecular Systems
HARMINV	Program and Library to Solve Harmonic Inversion
JDK	Java Standard Edition Development Kit
LAMMPS	Large-scale Atomic/Molecular Massively Parallel Simulator
LUMERICAL	High Performance Microscale Optics Simulation Software (commerical, group lumerical)
MATLAB-PCT	MATLAB and Parallel Computing Toolbox (commercial, matlab uwmatlab groups)
MPIBLAST	Parallel Implementation of NCBI BLAST
MrBAYES	Bayesian estimation of phylogeny.
NAMD	Parallel Molecular Dynamics Program
NWCHEM	Computational Chemistry Package
Octave	High level language mostly compatible with Matlab.
PHYLIP	Phylogeny Inference Package
R	Language/environment for statistical computing
SIESTA	Linear-scaling DFT based on Numerical Atomic Orbitals
SWARM	Multi-agent simulation of complex systems
TORQUE	Resource Manager and Cluster Scheduler (development use)
WRF	The Weather Research and Forecasting Model
Compiler	
CUDA	NVIDIA Compute Unified Device Architecture Toolkit (GPGPU)
FPC	Free Pascal Compiler
GCC	GNU Compiler Collection

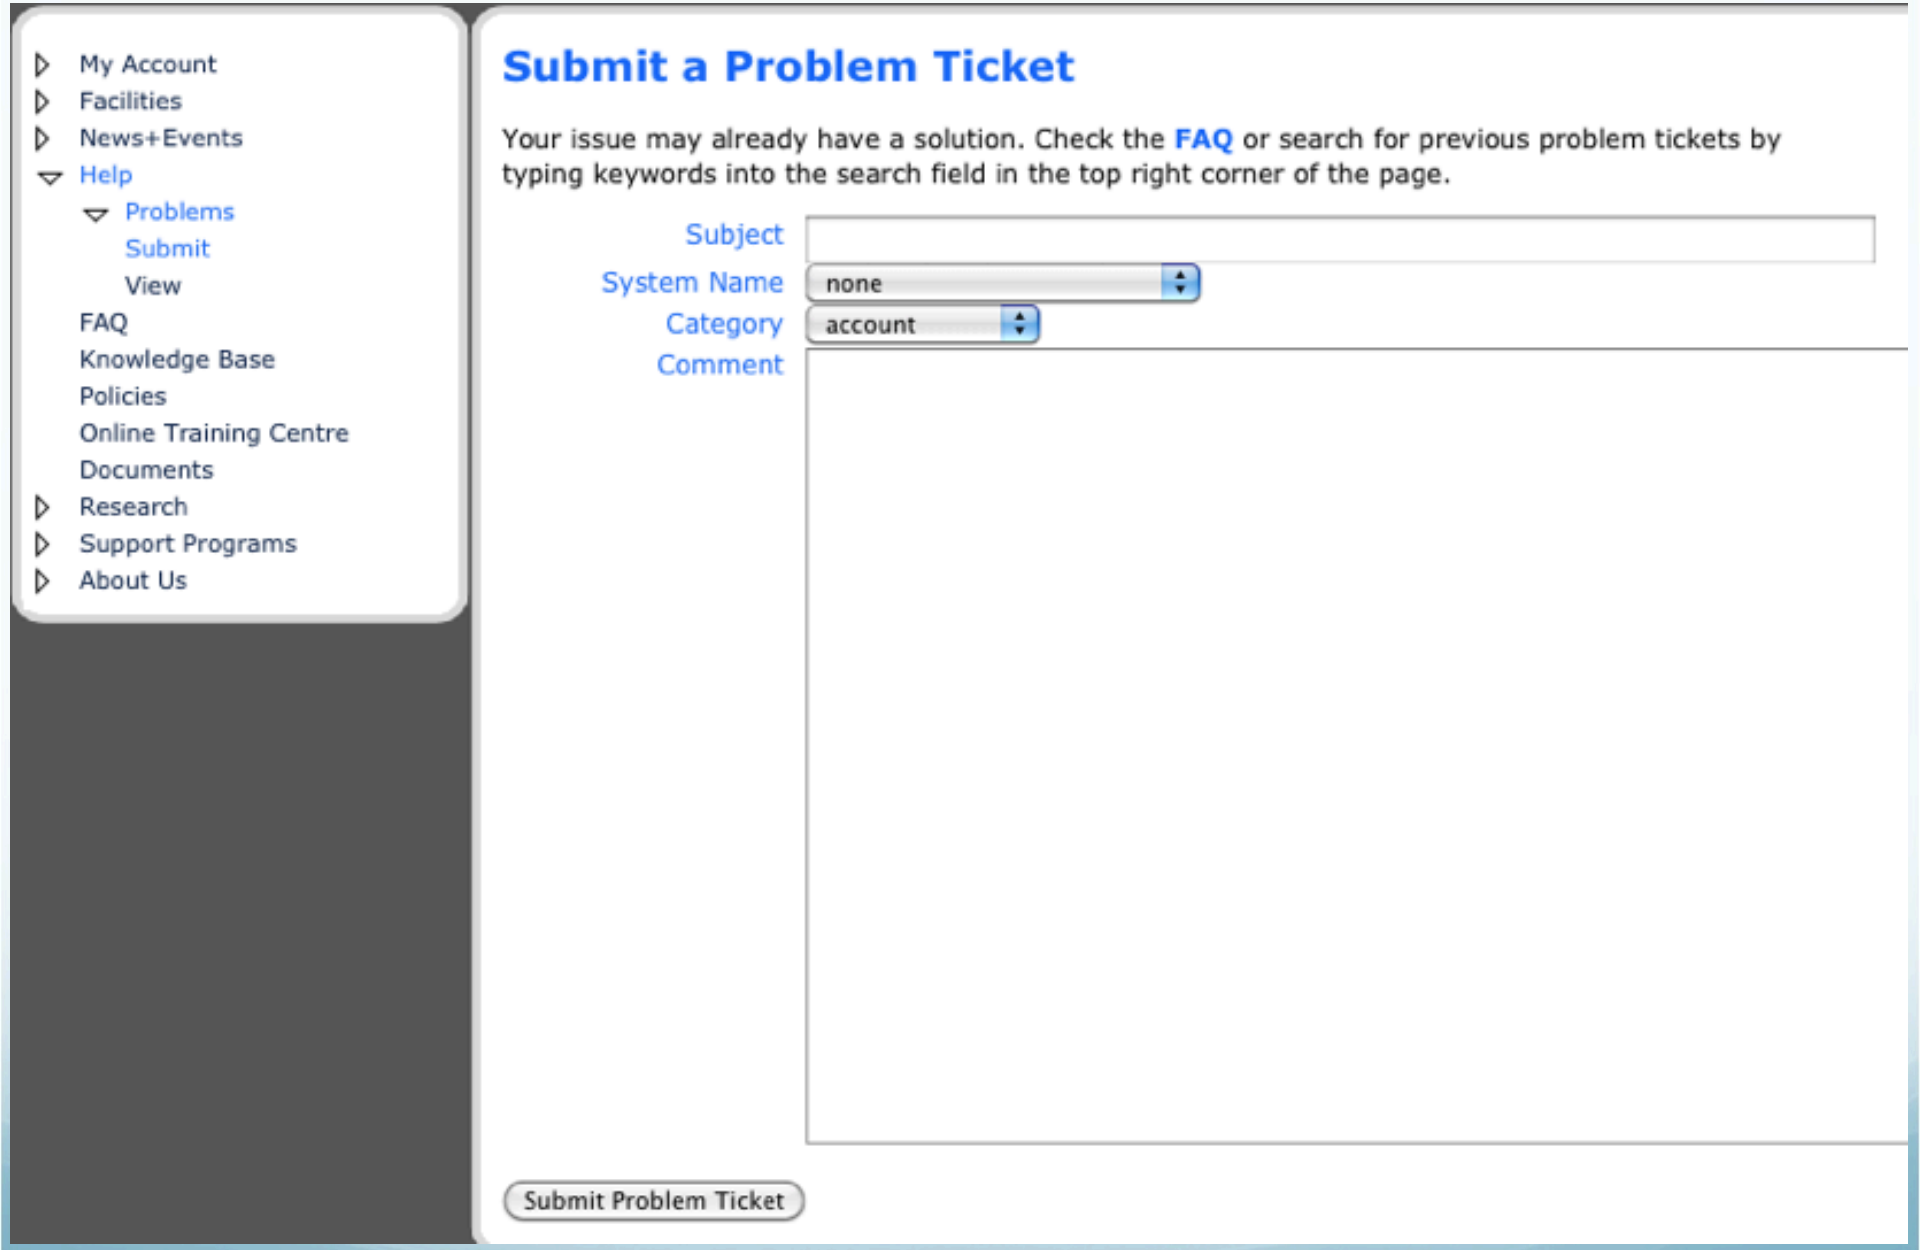

or can email us at "help@sharcnet.ca"

SHARCNET™

Knowledge Base / Expanded FAQ

Contents [hide]

Dedicated resources program

- My Account
- Facilities
- News+Events
- Help
- Research
- Support Programs
  - Manage Programs
  - Chairs
  - Fellowships
  - Dedicated Resources
  - Programming
- About Us

Dedicated Resources

In 2007, SHARCNET introduced the Dedicated Resources Programme which allows a researcher to apply for dedicated CPU time and/or storage.

Programme Objective

- To enable and facilitate large computational projects of exceptional research merit that will achieve discoveries of international significance through the optimal exploitation of SHARCNET's computing infrastructure.

Summary of Programme

This programme allocates dedicated computational resources such as CPU time or storage as follows:

- CPU time allocation: To enable high quality projects requiring large allocations of CPU time. The level of resources required for these projects is beyond that normally anticipated to be available to a project through general access (via the queuing system). Allocations of CPU time may be for large numbers of CPU-hours on either parallel systems (for "capability" applications) or on throughput facilities.
- Storage allocation: To provide large amounts of data storage and/or data storage for extended periods for specific high-quality projects.

Note that dedicated resources are for specific projects, not to provide resource "envelopes" for multiple projects. All CPU use at SHARCNET, obtained through either general access or dedicated time, remains free at present.

Currently, SHARCNET allocates resources within this programme in two ways: Small Dedicated Resources (Small DR) and Large Dedicated Resources (Large DR). Each has its own deadlines and resource thresholds as depicted in the guidelines.

Applications for dedicated resources are open to all faculty (PI) applicants with a valid SHARCNET account.

Open Rounds

SHARCNET recently issued a call for proposals for SHARCNET's newest Research Support Programme, Small Dedicated Resources. Applications are invited for the next Small DR deadline of **June 15, 2010**. Refer to the application **guidelines** for more information.

Note that the Dedicated Resources Programme was revised last Fall and now encompasses two streams with different deadlines and thresholds. The "Small DR" stream is new and will be run as a pilot initially.

The web-based **application form** is available via the SHARCNET webportal.

Request for dedicated programming support

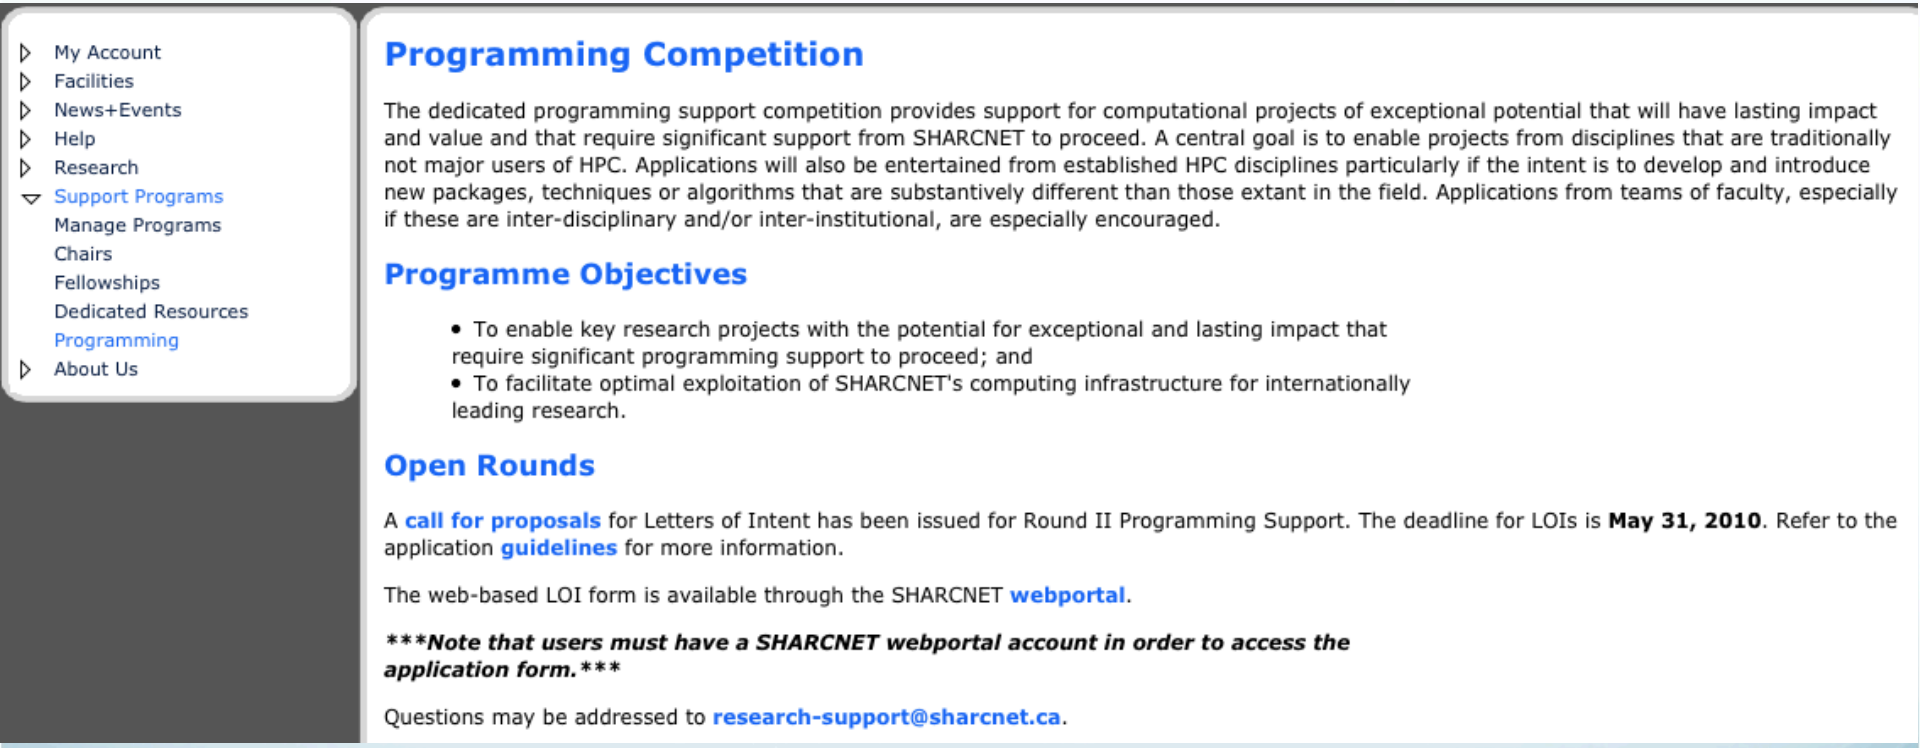

About Us > Contact > Directory

High Performance Computing Consultants

People	Areas	Phone
Ge Baolai (UWO)	Applied mathematics, scientific and technical computing, distributed systems.	519-661-2111 x88544
Nick Chepurniy (Windsor)	Applied Mathematics, Numerical Analysis, High Performance Computing, Optimization and Performance Tuning.	519-253-3000 x4883
Weiguang Guan (McMaster)	Visualization	905-525-9140 x22540
Jemmy Hu (Waterloo)	Computational Chemistry, Parallel Programming, High Performance Computing.	519-888-4567 x37854
David McCaughan (Guelph)	Software engineering, high performance computing, system programming, computational theory.	519-824-4120 x56467
Sergey Mashchenko (McMaster)	Astrophysics, scientific computing.	905-525-9140 x27663
Hugh Merz (Laurentian)	High Performance Computing, HPC Accelerators, Astrophysics	705-675-1151 x2347
Pawel Pomorski (Waterloo)	Molecular Dynamics Methods, Density Functional Theory Methods, Biophysics, Condensed Matter Physics	519-888-4567 x38458
Alexei Razoumov (UOIT)	Astrophysics, Computational Fluid Dynamics, Numerical Radiative Transfer, Adaptive Mesh Refinement Methods, Scientific Computing	905-721-8668 x3853
Doug Roberts (WLU)	Engineering, bioinformatics, computer algebra and grid.	519-884-0710 x3073
Tyson Whitehead (UWO)	Mathematics, Statistics, and Computer Programming	519-661-2111 x82846
Darryl Willick (Lakehead)	High Performance Computing, Computational Chemistry	807-343-8996
Isaac Ye (York)	Computational Fluid Dynamics, Turbulence, and Reacting Flows	416-736-2100 x22935

Basic Unix commands

Introduction to Linux webinar tomorrow

ls, ls -l, ls -a, ls *.txt
echo text > file, echo new >> file
more file, cat file
grep pattern file
mkdir, cd, pwd, rm, alias, chmod

can learn important ones in 30 min, the rest as you go

default bash shell (.bashrc)
My Account > Settings > Details > Edit – can change shell

Simple parallel (MPI) C code

/home/razoumov/hello.c

```
#include <stdio.h>
#include "mpi.h"
int main( int argc, char *argv[] )
{
    int nprocs, myrank, len;
    char pname[128];
    MPI_Init(&argc, &argv);
    MPI_Comm_size(MPI_COMM_WORLD, &nprocs);
    MPI_Comm_rank(MPI_COMM_WORLD, &myrank);
    MPI_Get_processor_name(pname, &len);
    printf("Hello world from %d on %s of %d\n", myrank, pname, nprocs);
    MPI_Finalize();
    return 0;
}
```

```
mpicc hello.c -o hello
sqsub -q mpi -n 8 -r 50m -o output.log ./hello
sqjobs
```

Debugging

in Knowledge Base > “Debugging serial and parallel programs” we have two useful tutorials:

- “Common Bugs and Debugging with gdb”
- “Parallel Debugging with DDT”

DDT is a graphical debugger for MPI parallel programs, installed on bruce, bull, dolphin, megaladon, narwhal, requin, zebra

debugging on SHARCNET (common programming bugs, gdb, DDT) will be covered in detail in HPC106 “Debugging” next week

Visualization packages

Visualization	
GRAPHVIZ	Represent structural information as diagrams of abstract graphs and networks
HYPERMESH	Altair Hyperworks Suite (commercial, hyperworks group, limited access)
ICEMCFD	ANSYS ICEM CFD Meshing Software (commercial, fluent group)
ITK	National Library of Medicine Insight Segmentation and Registration Toolkit
OpenDX	Visualization of Scientific, Engineering and Analytical Data
Paraview	Parallel Visualization Application
Scilab	Open Source Platform for Numerical Computation
VisIT	VisIT Visualization Tool
VMD	Visual Molecular Dynamics
XCrySDen	Crystalline and Molecular Structure Visualisation

open-source, multi-platform, and general-purpose:

1. OpenDX 4.4.4 – installed on vizN-site
2. VisIT 1.7 – installed on vizN-site
3. ParaView 3.6.1 – installed on rainbow, vizN-site

ParaView will be covered in HPC202 “Using visualization tools” next week