

# The CLUMEQ High Performance Computing Centre at McGill University



CIRMMT  
McGill University  
March 5, 2012

Dr. Bryan L. Caron  
Director, Business Operations

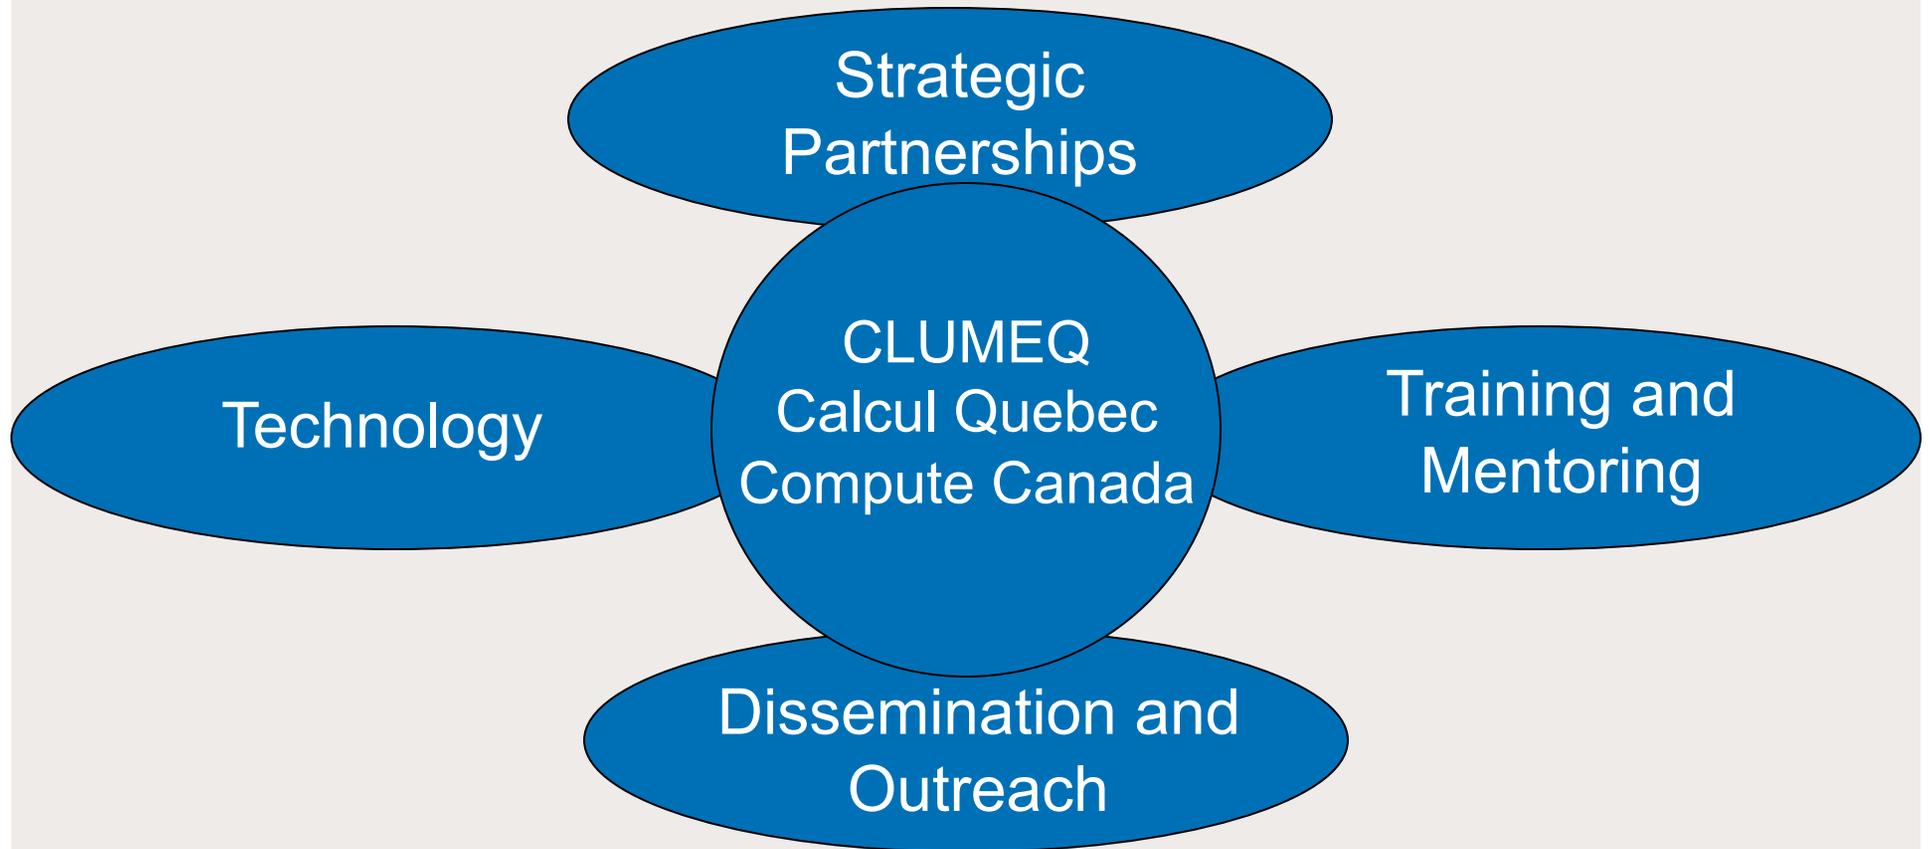
- **CLUMEQ, Calcul Quebec and Compute Canada**
  - Research Support Objectives and Focal Points
  - CLUMEQ Site at McGill – ETS
    - Key Specifications and Status
    - CLUMEQ HPC Support Staff at McGill
- **Getting Started**
  - Account registration and creation
  - Software, Storage and Job Submission
  - Support Information
- **Upcoming Activities**
  - Training and Workshops
  - Phase II Hardware Expansion
- **Summary**

- One of seven national Canadian High Performance Computing consortia funded by the Canada Foundation for Innovation
- National and regional coordination for HPC through *Compute Canada* and *Calcul Quebec*, respectively
- Provides computational resources to academic researchers from institutional members of CLUMEQ and across Canada
- Established in 2001 and led by McGill University
  - Université Laval
  - Université du Québec
    - Including École de Technologie Supérieure (ETS)
- **The Québec City site at Laval**
  - houses a 960 compute nodes Sun Constellation cluster with a total of 7680 cores, for a peak performance of 88 Tflops, and 500 TB of parallel storage
- **The Montréal site at McGill-ETS**
  - Guillimin - HPC Cluster Solution from IBM

# Research Support Objective

- “To provide visible and relevant support services to the research communities at McGill, throughout Québec and the rest of Canada, and internationally.”
- Enabling your research and innovation to succeed
  - Leading edge tools for leading edge research
    - Including help with code migration, performance analysis and optimization, parallel computing techniques for example
  - Advanced research also benefits research support
  - Achieving balance and agreement between the needs of research and HPC operations / support.
  - Measure the achievements in providing value to researchers

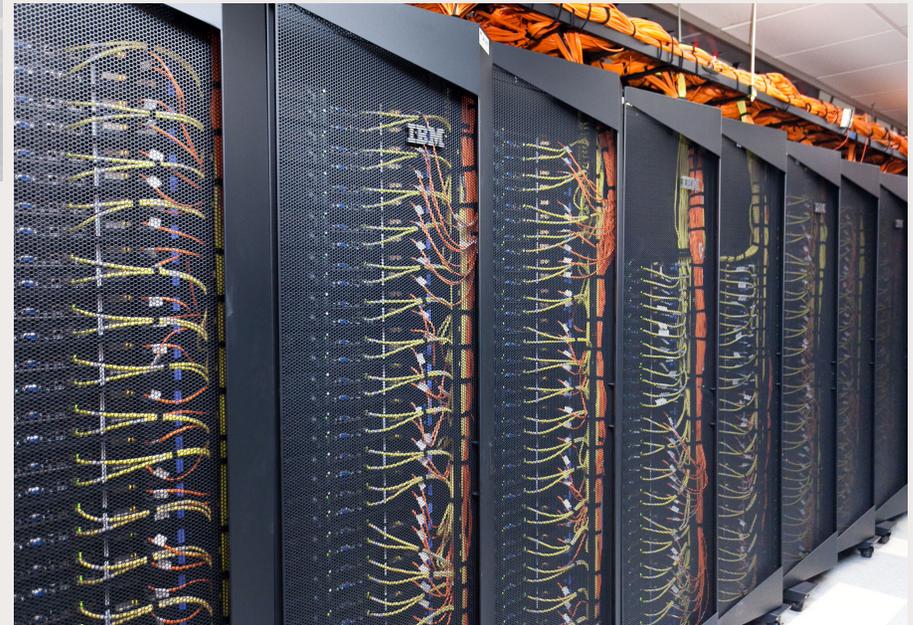
# Focal Points in Research Support



# Guillimin HPC Resource at McGill - ETS



- 5,000 square-feet of IT space
- 1200 node (14400 cores) IBM cluster
- 2 PB of attached parallel storage
- rated at a sustained 130 Tflops
  - Top 35 worldwide (June 2011)



# Guillimin HPC Resource - Specifications

- **Interactive Login Nodes ([guillimin.clumeq.ca](http://guillimin.clumeq.ca))**
  - 5 IBM x3650 M3 Nodes
  - Each with 8-cores of Intel Westmere-EP (2.4 GHz, 12 MB cache), 64GB memory
- **Compute Nodes**
  - 1200 IBM iDataPlex dx360 M3 Nodes [**14400 cores**]
    - Intel Westmere-EP Model X5650 6-core processors (2.66 GHz, 12 MB cache), 320 GB local disk storage per node
  - divided into 3 partitions:
    - **High Bandwidth** (400 nodes, 2 GB/core, non-blocking QDR IB) : **hb partition**
    - **Large Memory** (200 nodes, 6 GB/core, non-blocking QDR IB) : **lm partition**
    - **Serial Workload** (600 nodes, 3 GB/core, 2:1 blocking QDR IB) : **sw partition**
- **Storage**
  - **2 PB GPFS (Global Parallel FileSystem)** on top of a DDN SFA 10000 Storage Architecture
    - Two (2) SFA 10000 controller couplets, each with 10 x 60-drive (2TB / drive) drawers
- **Networks**
  - QDR Infiniband (IB) and Ethernet (1 and 10 Gbps) networks

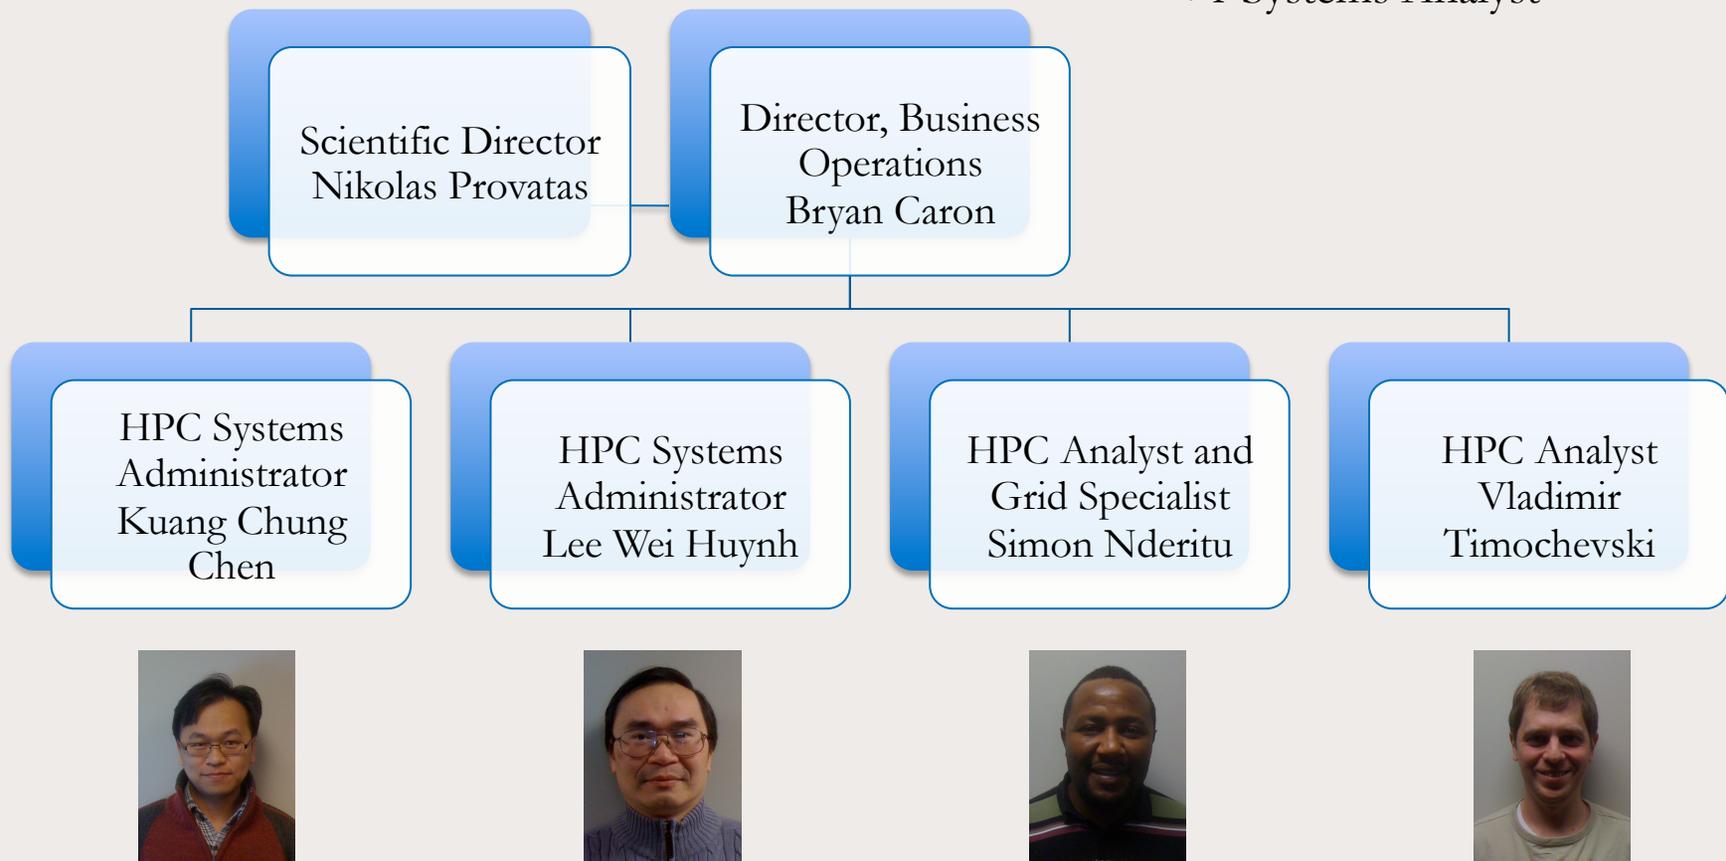
InfiniBand = high-bandwidth (40 Gbps), low latency ( $\sim 2.25 \mu\text{s}$ ) network interconnect

Non-blocking = able to handle simultaneous connections for all attached devices or all switch ports with a 1:1 relationship between ports and time slots

Blocking = fewer transmission paths than would be required if all nodes were to communicate simultaneously

# HPC Staff at McGill

Additional Staff in 2012:  
+2 HPC User Support Analysts  
+1 Systems Analyst



# Getting Started

# Getting Started @ <https://support.clumeq.ca>

Home    Guillimin Forum    My Account    Contact us

## CLUMEQ Supercomputing Center User Support Portal



You are here: Home CLOSE INFO

### User Portal

- Home
- News
- Accounts and Resources
- Training

### Guillimin Cluster

- Account request
- Getting Started
- Hardware
- Software
- FAQ

### Colosse Cluster

### Welcome to CLUMEQ User Support Portal

This portal allows access to hardware and software documentation, to information about our systems configuration and CLUMEQ specific information. Help is also provided on accessing our systems, and on user and systems policies.

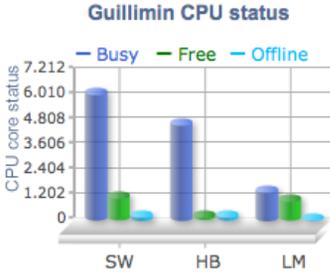
If you have specific questions, or need additional help, please contact us at [support@clumeq.ca](mailto:support@clumeq.ca).

### CLUMEQ Newsline

**27-6-2011:** [Our new User Support Portal is online.](#)

It is accessible via [support.clumeq.ca](https://support.clumeq.ca), and replaces the old support website. Here you can find all necessary information to open account and start working with CLUMEQ Guillimin, Colosse and Krylov clusters. You can also check the current load of Guillimin and Krylov clusters in real time.

### Guillimin CPU status



Cluster	Busy	Free	Offline
SW	6.010	1.202	0.000
HB	4.808	1.202	0.000
LM	1.202	1.202	0.000

Krylov cluster is currently online for user access following the security incident in July. All Krylov users are encouraged to move their work to the modern CLUMEQ hardware, either Colosse or Guillimin, since

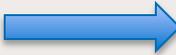
**1-Feb-12: Guillimin default queue disabled**

Starting at 9:00, Thursday, Feb 2, the "default" queue will be disabled on Guillimin. All users will need to use "msub -q <queue>" to explicitly specify one of the queues while

**10-Nov-11: MPI packages are available on Guillimin**

We fixed the issues with running parallel jobs on Guillimin. All MPI modules are back to service, and you can submit parallel jobs now ...

# Obtaining an Account

- **Step 1: Obtain a Compute Canada Identifier (CCI) and Role Identifier (CCRI)**
  - [https://ccdb.computecanada.org/account\\_application](https://ccdb.computecanada.org/account_application)
    - CCI = a unique personal and national identifier
    - CCRI = Your position such as “Faculty”, “Graduate Student”, “Postdoctoral Fellow” or “Researcher”
  - If you are not a professor, you must provide the CCRI of your professor when completing the application form
  - Typically will receive your CCI and CCRI within 24 hours of application (requires verification)
- **Step 2: Obtain a CLUMEQ Account**
  - Return to the Compute Canada website, login with your CCI / registered email address
  - “Link or apply for consortium account”  CLUMEQ (Apply)
  - Complete the [secure online form](#) to select your CLUMEQ userid and set your password

# Getting Started @ <https://support.clumeq.ca>

Home    Guillimin Forum    My Account    Contact us

## CLUMEQ Supercomputing Center User Support Portal



You are here: Home CLOSE INFO

### User Portal

- Home
- News
- Accounts and Resources
- Training

### Guillimin Cluster

- Account request
- Getting Started**
- Hardware
- Software
- FAQ

### Colosse Cluster

### Welcome to CLUMEQ User Support Portal

This portal allows access to hardware and software documentation, to information about our systems configuration and CLUMEQ specific information. Help is also provided on accessing our systems, and on user and systems policies.

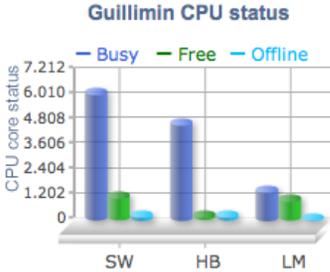
If you have specific questions, or need additional help, please contact us at [support@clumeq.ca](mailto:support@clumeq.ca).

### CLUMEQ Newsline

**27-6-2011:** [Our new User Support Portal is online.](#)

It is accessible via [support.clumeq.ca](https://support.clumeq.ca), and replaces the old support website. Here you can find all necessary information to open account and start working with CLUMEQ Guillimin, Colosse and Krylov clusters. You can also check the current load of Guillimin and Krylov clusters in real time.

### Guillimin CPU status



Cluster	Busy	Free	Offline
SW	6.010	1.202	0.000
HB	4.808	1.202	0.000
LM	1.202	1.202	0.000

Krylov cluster is currently online for user access following the security incident in July. All Krylov users are encouraged to move their work to the modern CLUMEQ hardware, either Colosse or Guillimin, since

**1-Feb-12: Guillimin default queue disabled**

Starting at 9:00, Thursday, Feb 2, the "default" queue will be disabled on Guillimin. All users will need to use "msub -q <queue>" to explicitly specify one of the queues while

**10-Nov-11: MPI packages are available on Guillimin**

We fixed the issues with running parallel jobs on Guillimin. All MPI modules are back to service, and you can submit parallel jobs now ...

# Getting Started @ <https://support.clumeq.ca>

Home Guillimin Forum My Account Contact us

## CLUMEQ Supercomputing Center User Support Portal

You are here: Home

### User Portal

- Home
- News
- Accounts and Resources
- Training

### Guillimin Cluster

- Account request
- Getting Started
- Hardware
- Software
- FAQ

### Colosse Cluster

### Welcome to CLUMEQ User Support Portal

This portal allows access to hardware and software documentation, systems configuration and CLUMEQ specific information. Here you can find information on our systems, and on user and systems policies. If you have specific questions, or need additional help, please contact [support@clumeq.ca](mailto:support@clumeq.ca).

### CLUMEQ Newslines

27/3/2011: Our new User Support Portal is online. It is accessible via [support.clumeq.ca](https://support.clumeq.ca), and replaces the old support website. Here you can find all necessary information to open account and start working with CLUMEQ Guillimin, Colosse and Krylov clusters. You can also check the current load of Guillimin and Krylov clusters in real time.

### CPU core status

7.212
6.010
4.808
3.606
2.404
1.202
0

# Interactive Login Access

## Linux users

```
ssh guillimin.clumeq.ca -l username
```

(You need to add "-X" option after "ssh" if you expect to output any graphics to your screen)

## Windows Users

You have to install special software that allows secure shell connections. There are a number of free, third party packages available to accomplish this. The most popular SSH clients for Windows are [PuTTY](#) and [Xshell](#).

A particularly comprehensive package is [Cygwin](#). It represents a full-featured Linux-type environment inside your windows session. You may need to install this environment if you expect to output graphic objects on your screen. This [tutorial](#) explains the steps required for its installation.

Once the ssh client is installed, you may need to create a profile. If so, you will need the following information:

```
Host name : guillimin.clumeq.ca
User name : Enter your user name
Port number: 22
```

- Commonly used software packages are installed in a central location on the cluster and are accessible using the modules environment
  - Automatically adds/removes paths to your environment variables for the selected packages

- Show all modules available on the system:

```
user@guillimin$ module avail
```

- Show currently loaded modules:

```
user@guillimin$ module list
```

- Load particular module:

```
user@guillimin$ module add <modulename>
```

- Unload particular module:

```
user@guillimin$ module rm <modulename>
```

# Disk Space

- By default all allocations (accounts for groups and their users) have access to a baseline set of directories and quotas on the system
  - home, project, global scratch, local scratch spaces

Type	Location	Policies
Home Directory	<code>/home/userID/</code>	disk backup, no auto-clean, per user quotas (10 GB) within /home/
Project Directory	<code>/sb/project/rapID/</code>	created by default per rapID, no backup, no auto-clean, per rapID quotas (1 TB)
Scratch Directory	<code>/sb/scratch/userID/</code>	created by default per user, no backup, auto-clean, per user quotas (5 TB)
Local Scratch Directory	<code>/localscratch/</code>	local to worker node during job execution, no backup, auto-clean

- Auto-cleanup for files older than 30 days in scratch spaces

# Submitting Jobs

- Computational jobs are run in batch mode as opposed to interactive
  - Interactive login nodes
    - file copying, code compilation and job submission only
  - Batch nodes
    - Execution of all user workloads non-interactively through job submission scripts
- Jobs are submitted to queues, which then run depending on a number of factors controlled by the scheduler
  - ensures that the cluster resources are used in the most effective way
  - cluster resources are shared between the users in a fair manner
  - Usage governed by the concept of fairshare with both per project and default allocations (target number of cores over the course of the year) that determine job priorities for execution
- Users manage your jobs (submit, check status, kill, etc.) only with the help of a special set of scheduler commands

# Submitting Jobs

- Typical user commands for working with jobs

- `msub -q <qname> <script_file>`

- submit a new job, where `script_file` is a text file, containing the name of your executable program and execution options (number of cores requested, working directory, job name, etc) to the queue named `qname`

```
#!/bin/bash
#PBS -l nodes=1:ppn=1
#PBS -l walltime=12:00:00
#PBS -o outputfile
#PBS -e errorfile
#PBS -V
#PBS -N jobname

cd /home/username/your_project_name
./your_code
```

# Submitting Jobs

- Typical user commands for working with jobs
  - `msub -q <qname> <script_file>`
    - submit a new job, where `script_file` is a text file, containing the name of your executable program and execution options (number of cores requested, working directory, job name, etc) to the queue named `qname`
  - `showq` – shows a current list of submitted jobs
    - `showq -v` – shows a full detailed list of submitted jobs
    - `showq -r` – the same as previous but a list of assigned nodes is shown for each running job
    - `show-q -u <username>` – shows only the jobs for a specific user
  - `checkjob -v <jobID>` - shows detailed information about a specific job
  - `canceljob <jobID>` – kills the job or removes it from the queue
  - `showbf` – shows the list of currently available resource
  - `showstart <jobID>` – shows when the scheduler estimates the job will run

# Submitting Jobs

- Guillimin Batch Queues

Name	Default Duration (h)	Maximum Duration (h)	Memory per core (process)	Notes
sw	3	720	3	Serial queue workloads with no inter-node communications
hb	3	720	2	Smaller memory non-blocking network
lm	3	720	6	Larger memory non-blocking network
debug	0.5	2	-	Debug queue for short duration tests

# Obtaining Support

- CLUMEQ General Information and News
  - <https://www.clumeq.ca>
- CLUMEQ User Support and Documentation
  - <https://support.clumeq.ca>
  - Support Request Email: [support@clumeq.ca](mailto:support@clumeq.ca)
  - Online Discussion Forum

# Support Discussion Forum



[Board index](#)

⌵ ⌶ ⌷

[FAQ](#) [Search](#) [Login](#)

It is currently Mon Mar 05, 2012 1:04 am

[View unanswered posts](#) • [View active topics](#)

GENERAL RULES	TOPICS	POSTS	LAST POST
 <b>General Rules</b> IMPORTANT: Please read our general rules before posting to any forum	2	2	by <b>vladimir</b>  Tue Apr 19, 2011 9:01 am
ACCOUNTS, LOGIN PROCEDURES, COMPILING CODES, SUBMITTING JOBS	TOPICS	POSTS	LAST POST
 <b>Getting access to Guillimin cluster</b>	0	0	No posts
 <b>Building your code on Guillimin: Compilers, libraries, etc.</b>	0	0	No posts
 <b>Submitting jobs on Guillimin: PBS and Scheduler related discussions</b>	0	0	No posts

# Upcoming Activities

- **Additional CLUMEQ technical staff at McGill to be hired in order to form the core HPC Support Team**
  - Including 2 HPC Analysts
  - <http://www.mcgill.ca/hr/posting/1/high-performance-computing-analyst>
- **User Training and Workshops**
  - Introduction to HPC Training
    - McGill campus - March 30, 2012
  - Upcoming Workshops at McGill
    - Youth Outreach and Education – May 2012
    - Industry and HPC for Process Control in Manufacturing – Summer 2012
- **CLUMEQ at McGill – Phase II**
  - Expansion planned to approximately double the computational and storage capacity
  - Preliminary discussions and investigations of technology options and trends
  - Targeting to issue RFP on the time-scale of Spring 2012
    - installation and deployment for September 2012
  - We need your input as a User Community!
    - What are your compute and storage needs for the future?

# Summary

- **CLUMEQ Resources at McGill – ETS**
  - Significant computational and storage capacity
    - 14400 processing cores and 2 PetaBytes storage
    - Expansion to approximately double the computational and storage capacity of the facility (September 2012)
- **CLUMEQ Research Support**
  - Our objective is to enable your research and innovation to succeed through our infrastructure and support services
    - Technology, Training, Partnerships and Outreach

# Important Links

- CLUMEQ General Information and News
  - <https://www.clumeq.ca>
- CLUMEQ User Support and Documentation
  - <https://support.clumeq.ca>
  - Support Request Email: [support@clumeq.ca](mailto:support@clumeq.ca)
- Calcul Quebec  Calcul Québec
  - <http://www.calculquebec.ca>
- Compute Canada
  - <https://computecanada.org>

Thank you for your kind attention.  
Any questions?



**McGill**

**Dr. Bryan L. Caron**

Director, Business Operations

CLUMEQ High Performance Computing Centre

McGill University

395, rue Murray, bureau 4600

Montreal, QC H3C 0J8

Tel: 514 396-8958

Cell: 514 458-3646

Fax: 514 396-8934

Email: [bryan.caron@clumeq.ca](mailto:bryan.caron@clumeq.ca)

[bryan.caron@mcgill.ca](mailto:bryan.caron@mcgill.ca)

