

BLUE WATERS

SUSTAINED PETASCALE COMPUTING

3/17/14

Blue Waters User Monthly Teleconference



GREAT LAKES CONSORTIUM
FOR PETASCALE COMPUTATION

CRAY®

Agenda

- Recent Blue Waters system events
- Recent changes to Blue Waters
 - job policies
 - portal nodes in use pie chart
 - login node use policy
 - project links in Nearline home
- Moving data to Nearline at the end of a job
- Darshan available for testing
- Requests for information on Science Successes
- Blue Waters Symposium
- Blue Waters Internships
- XSEDE/Blue Waters Extreme Scaling Workshop
- XSEDE14 Conference

Recent System Events

- Single OST issue on scratch file system March 9th.
 - Initial focus on recovery (migration) of user data.
 - Files listed in ~/scratch-failure-files-of-concern.20140309.txt
 - Small percentage of files showing as unrecoverable.
 - Should finish migrating user data today; damaged files will be placed in ~/scratch directories.
 - Recommend checking integrity of migrated files.
 - Next step is root cause after recovery efforts finish up.
- Scratch is not backed up. Timely transfer of data to Nearline recommended.

Recent Changes

- Queue time priority factor decreased
- Priority accrual eligibility restricted by nodes instead of jobs
- Wallclock overrun grace removed (previously 10 minutes)
- Future consideration:
 - Possible weekday reservation for small number of nodes for small, short jobs

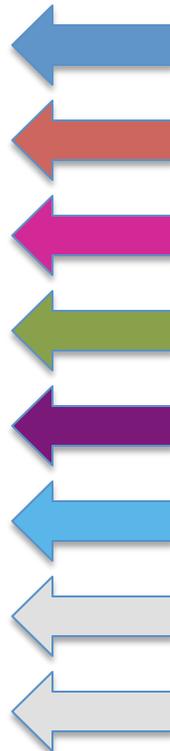
Pie chart node statistics

- <https://bluewaters.ncsa.illinois.edu/machine-status>
- Provides a better picture of node availability

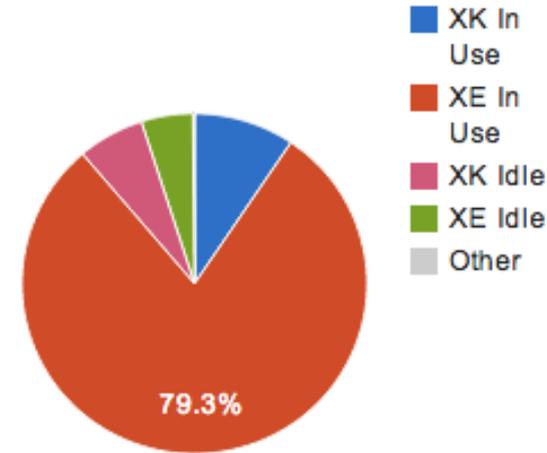
USAGE

Nodes:

2536	XK Nodes In Use
21314	XE Nodes In Use
1676	XK Nodes Idle
1298	XE Nodes Idle
6	XK Nodes Draining
25	XE Nodes Draining
6	XK Nodes Down
3	XE Nodes Down



Node usage



Login node policy

<https://bluewaters.ncsa.illinois.edu/login-node-policies>

Long running user tasks:

- Terminate user process over 100 hours or over 4 hours of cputime at 20% cpu utilization.
- User notified by email.
- Exemptions by request help+bw@ncsa.illinois.edu.

Rsync:

- Limited to one instance per user per login node.
- Please use Globus Online for data transfers.

Links in ncsa#Nearline /~/

- Added project links in Nearline

Endpoint: ncsa#Nearline
Path: /~/

Item	Type	Size
PK	Folder	
chroma_consistency	Folder	
dirA	Folder	
milc_consistency	Folder	
project.fyy	Folder	
project.jnd	Folder	
project.jnf	Folder	
project.jnx	Folder	
project.joj	Folder	
84144input.lat	File	22.89 GB
Downloads.tar.gz	File	9.78 GB
lattice.B.0.lat	File	22.89 GB

Endpoint: ncsa#Nearline
Path: /~/project.jnf/

Item	Type
bazavov_scratch	Folder
sgottlie	Folder

Same as /projects/sciteam/jnf

Endpoint: ncsa#Nearline
Path: /projects/sciteam/jnf/

Item	Type
bazavov_scratch	Folder
sgottlie	Folder

Moving data to Nearline

- How to coordinate movement of data to Nearline with the end of a job?
- We are working on a short-term solution. Contact us (help+bw@ncsa.illinois.edu) if you are interested.
- A longer term solution is being considered.

Darshan

- Please try Darshan with your application. We will make this a default module in the user environment in a month.
- See <https://bluewaters.ncsa.illinois.edu/darshan> for more info.
- Automatically profile application IO:
 - No code modification
 - Works with POSIX, HDF5, NetCDF, MPI-IO
 - Low overhead
 - Scalable
 - Must be an MPI application
- Per-file per-process statistics (organized by file)
 - Counts I/O operations, e.g. unaligned and sequential accesses
 - Times for file operations, e.g. opens and writes
 - Accumulates read/write bandwidth info
 - Creates data for simple visual representation

Darshan continued

- Reasons for no profile file
 - Not an MPI application
 - Application did not exit cleanly
 - Environment variable DARSHAN_DISABLE is set
 - Application not linked with Darshan libraries
 - check `'nm | grep -i darshan'`
- Enable with `module load darshan`
- Known issues:
 - Shared objects cannot be built – patch will be applied soon, currently testing
- Portal link: <https://bluewaters.ncsa.illinois.edu/darshan>

Request for Science Successes

- We need to be current on products that result from time on Blue Waters such as:
 - Publications, Preprints (e.g. [arXiv.org](https://arxiv.org) ) , Presentations.
 - Very interested in data product sharing.
- Appreciate updates sooner than annual reports.
 - Send to gbauer@illinois.edu
- NSF PRAC teams send information to PoCs.
- See the [Share Results](#) section of the portal as well.
- Be sure to include [proper acknowledgment](#)
 - Blue Waters - National Science Foundation (ACI 1238993)
 - NSF PRAC – OCI award number

Blue Waters Symposium

- 2nd Annual - Petascale Science and Beyond
- May 12-15, 2014 in Champaign-Urbana.
- Keynotes, panels, posters, and breakout sessions.
- Blue Waters science teams are expected to have a senior representative.

Blue Waters Internships

- 20 undergraduate research interns in 2014. The goal is to engage undergraduate students in petascale computing research and development projects.
- The program provides each student a stipend totaling \$5000, a two-week intensive high-performance computing workshop, and travel to The Blue Waters Symposium 2015.
- This program includes support for undergraduate internship activities at any accredited degree granting institution in the United States.
- Applications must be submitted by March 21, 2014
- Notifications will be made by April 15, 2014
- Also looking for Faculty interested in Mentoring
- For more information see

<https://bluwaters.ncsa.illinois.edu/internships>

XSEDE/Blue Waters Extreme Scaling Workshop

- Topic: The Intersection of Big Data and Large-Scale Computing
 - Data intensive applications running on HPC style architectures
 - Big data analytical applications running on HPC style architectures
 - Computationally intensive applications running on big data or data intensive architectures
 - Using GPU/many-core processors for big data applications
 - System architectural features that can support big data and high performance computing
- The Extreme Scaling Workshop 2014 will showcase the discoveries, innovations, and achievements of those who use, build, and/or support advanced architectures at extreme scales around the world.
- The workshop will also provide a forum among researchers, professional staff, students, HPC center staff, and representatives from funding agencies.
- August 14-15, 2014 in Boulder, CO.
- See <https://www.xsede.org/web/xscale/xscale14> for more information.
- After the [RMACC High Performance Computing Symposium](#).

XSEDE14 Conference

- July 13-18, 2014 in Atlanta
- Theme is Engaging Communities
- Paper submissions and tutorials due March 15
- Topics span accelerating discovery, advanced technologies, software, science gateways and portals, and education, outreach and training
- Over 700 people from academia, industry, government, and other organizations
- Financial support for Champions and student participation



TO LEARN MORE VISIT
CONFERENCES.XSEDE.ORG/XSEDE14

Future Topics?

- Please send us your suggestions on topics for future teleconferences / webinars