

BLUE WATERS

SUSTAINED PETASCALE COMPUTING

9/20/15

Blue Waters User Monthly Teleconference



GREAT LAKES CONSORTIUM
FOR PETASCALE COMPUTATION

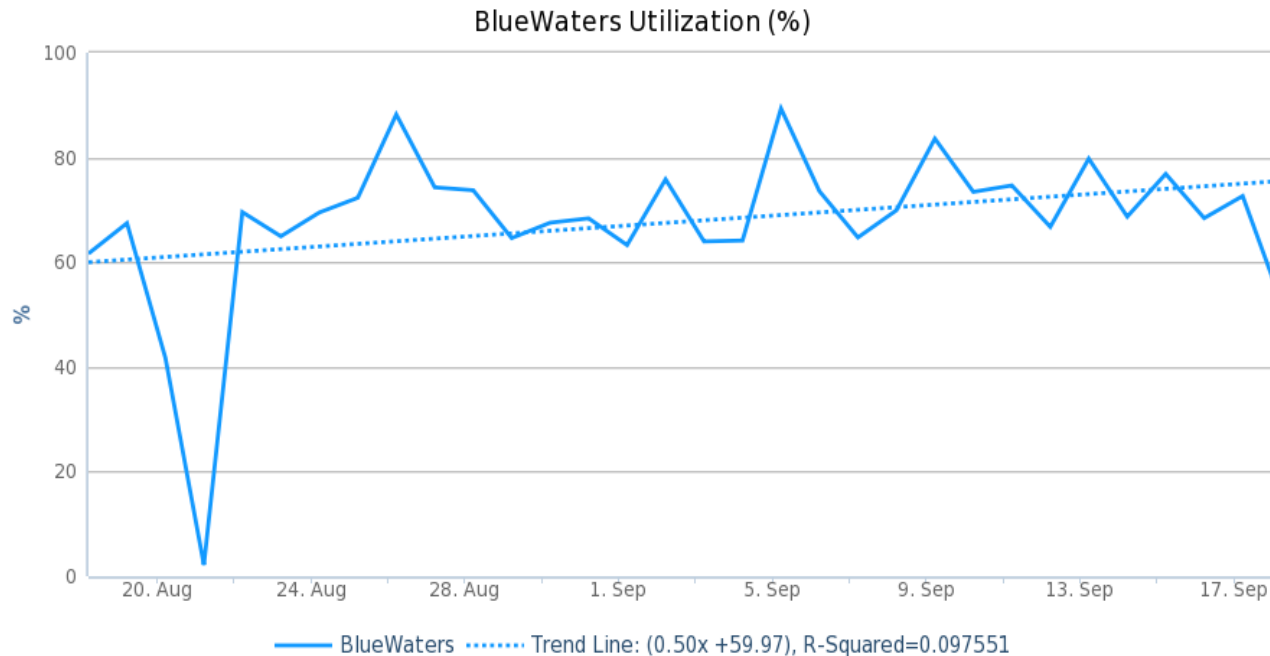
CRAY®

Agenda

- Utilization
- Discount Charging
- Recent Events
- Upcoming Maintenance
- New Software
- Upcoming Opportunities
- PUBLICATIONS!

System Utilization

- Utilization since last BW User Call (August 17)



2015-08-18 to 2015-09-18 Src: HPCDB. Powered by XDMoD/Highcharts

- Utilization improved. Good job backlog.

Recent Events

- Scheduler issues.
 - Series of issues identified for recent scheduler interruptions. Workarounds and/or fixes in place.
 - Discourage use of `qsub -V` option and disabled `modimp` to reduce environment size.
 - Discourage use of `showstart` command.

Discount charging

Charge factors for completed jobs that meet one or more of the following criteria below are discounted by 25% for each of the following opportunities with a resulting maximum discount of 69% when compounded.

1. job backfills available nodes
 2. submit pre-emptible job
 3. use flexible wall clock time
 4. job wall clock accuracy of 75% or better
- For more information see the [July 9th blog entry Charge Factor Discounts for jobs on Blue Waters.](#)

Review of Best Practices

- Improper use of login nodes
 - Use compute nodes for all production workloads.
- Avoid excessive calling of job scheduling commands
 - Unintentional denial of service may result otherwise.
- Unbundling of Jobs
 - Independent jobs bundled to 32 nodes or less best for backfill etc.
- Small files usage
 - Use projects then scratch but not home file system.
 - Tar up files before transferring to Nearline

Recent Changes

- Notifications
 - Added usage threshold based alerts.
- New project serial number for each project
 - Add users and move files between projects during grace period.

New Software: cuSolver

- [cuSolver](#) – NVIDIA CUDA accelerated solver
 - cuSolverDN: Dense LAPACK
 - cuSolverSP: Sparse LAPACK
 - cuSolverRF: Refactorization
- asynchronous execution
- thread safe
- CUDA streams for overlapping tasks
- 1.2x – 10x faster than CPU version (GPU vrs socket)

New Software: Python

- Python 2 version 2.7.10
- Python 3 version 3.3.3
- Pypy version 2.6.0
- Pypy3 version 2.4.0
- Astropy
- Numpy
- H5py
- Matplotlib
- Mpi4py
- Netcdf4-python
- Pandas
- Pycdf
- Pillow
- Pycuda (module load cudatoolkit)

Pypy modules

- Pypy
 - Numpy (~80% coverage, git)
 - Status: <http://buildbot.pypy.org/numpy-status/latest.html>
 - Mpi4py (git)
- Pypy3
 - Mpi4py (git)

Python PE and Version management

- ``module load bwpy``
 - No need to switch PrgEnv: libraries linked with `RUNPATH`
- Python programs
 - ``python`` → ``python2.7``
 - ``python2`` → ``python2.7``
 - ``python2.7``
 - ``python3`` → ``python3.3``
 - ``python3.3``
- Python programs (cont.)
 - ``pypy``
 - ``pypy3``
 - ``pip`` → ``pip2.7``
 - ``pip2.7``
 - ``pip3.3``
 - ``virtualenv``
 - ``yolk``
- Unversioned executables default to python 2.7.10
- Use ``python3`` for 3.3.3
- Use ``python2`` or ``python3`` unless syntax is compatible with both

Python dependency tracking: virtualenv

- `virtualenv --system-site-packages -p PYTHON_EXEC DEST_DIR``
 - `--system-site-packages` makes installed packages like numpy usable in virtualenv
 - `PYTHON_EXEC` one of:
 - `python2.7`
 - `python3.3`
 - `pypy`
 - `pypy3`

New NVIDIA Webinar

- NVIDIA/Portland Group: Performance Portability on x86 **CPUs and GPUs** with OpenACC Programming
 - [PDF](#) of presentation
 - [Recording](#) of presentation

OpenACC Course - Starts Oct 1st A Free Online Course

- ▶ Experienced Instructors
- ▶ OpenACC Toolkit
- ▶ GPU Access
- ▶ 4 Classes
- ▶ 4 Office Hours
- ▶ Hands-on Labs

Register at https://developer.nvidia.com/openacc_course

Removal of old packages

- pgi/12.2.0
 - pgi/12.3.0
 - pgi/12.4.0
 - pgi/12.8.0
 - pgi/12.9.0
 - pgi/12.10.0
 - pgi/13.4.0
 - pgi/13.7.0
 - pgi/13.8.0
 - pgi/14.1.0
- No ld link usage for over 1 year.

Upcoming Changes

- Preventative Maintenance - 10/5 (tentative)
 - File system patches.
 - Date depends on success of final testing.
- File System Upgrade 11/2015-01/2016
 - Lustre upgrade.
 - Move to declustered RAID: GridRAID
 - Move to Hierarchical Storage Management (HSM).
 - Working to minimize inconvenience.

File System Upgrade 11/2015 – 01/2016

- Upgrade to Lustre 2.5.1+ server
- Move to Declustered RAID
 - Fewer OSTs while keeping same total capacity.
- Internally
 - Data organized better – faster rebuild times
 - Support for Hierarchical Storage Management (HSM) for Lustre to Nearline.
- Externally
 - Better performance.
 - Ease of use from HSM.

File System Upgrade 11/2015 – 01/2016

- Data movement part of the upgrade process.
- Occasional downtimes but most data will be moved in the background.
- NCSA is currently writing the software to do the data movement.
- We will begin with home in November, with scratch planned for January.

Next NSF PRAC Call for Proposals

- Webinar for interested parties coming.
- New Program Officer for next call.
- Check [Petascale Computing Resource Allocations](#) (PRAC)


Advanced User Workshop

- At NCSA.
- Save the date: 11/3 - 11/6
- Registration page and agenda will be out in later this week.
 - Cray Performance Tools
 - Parallel IO using MPI-IO
 - HDF

Join NCSA at SC15!

- NCSA will have an exhibit at the annual Supercomputing Conference in November. If you will be attending SC15, consider giving a demo or talk about your current Blue Waters work in NCSA's booth. The demo space offers an 80-inch HD screen which can be run from a laptop or a USB. We recommend short demos of approximately 15 minutes (followed by Q&A), but demos that require more time are also welcome. Demos can also be repeated at different times. Demos will be promoted on the NCSA website, through social media, and through other channels.
- If you would like to give a demo or talk, please send a brief description, along with any known networking or technical requirements, to Matt Turk at: mjturk@illinois.edu.
- All demo submissions will be reviewed by an NCSA committee before being assigned time in the Demo Theater.

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