

BLUE WATERS

SUSTAINED PETASCALE COMPUTING

12/17/15

Blue Waters User Monthly Teleconference



GREAT LAKES CONSORTIUM
FOR PETASCALE COMPUTATION

CRAY®

Agenda

- NSF Review
- Topology Aware Scheduling Analysis
- Utilization
- Recent Events
- Upcoming Maintenance
- Upcoming Opportunities
- PUBLICATIONS!

NSF Review

- December 7th-8th at NSF.
- Well received.
- Panel impressed by Science and Engineering successes.


Topology-aware Scheduling Investigation

- Topology-aware job placement was disabled from started Nov 6 to Dec 6.
 - Was done without an announcement to minimize changes in user behavior.
- We are interested in comparing performance differences for comparable jobs run during the test period. We need your help.
- If interested please provide:
 - the name of the application
 - job ids for the jobs referenced,
 - a one-line description of the computation being performed,
 - measure of performance (e.g job run-times or time per iteration) during and after the test period.
- In return,
 - 100% credit for the compared
 - Deadline Jan 15, 2016

Note: Jobs which don't rely on inter-node communication would not have been impacted by different placement, and thus, are excluded from this solicitation. This would include single node jobs, bundled single node jobs.

System Utilization

- Utilization since last BW User Call (October 19)



2015-10-19 to 2015-12-16 Src: HPCDB. Powered by XDMoD/Highcharts

- Accounting issue allows utilization > 100%
- Transitions between modes apparent.

Recent Events

- 11/30 - Cooling Failure caused 3 cabinets to power down. System rerouted HSN traffic around down routers.
- 11/20 – Scheduler issues. Vendor working on reported issues.

Recent Changes

- Maintenance 12/07
 - File System Patches for stability
 - NVIDIA Driver Security Patch
 - No anticipated impact.
 - Issue was observed with some older binaries or with more recent applications and how CUDA libraries are linked.
 - Recommendation: relink or rebuild.
 - Home/Projects sync postponed.

Upcoming Changes

- Software Upgrade - date not set but coming
 - Cray Linux Environment upgrade 5.2UP04
 - Upgrade to CUDA 7.0 from CUDA 6.5
 - CUDA 7.0 support with GNU 4.9
 - Programming environment defaults will change
 - CCE 8.4.X
 - GNU 4.9 (because of CUDA 7)
 - PGI 15.7 (note that 15 and 14 are incompatible)
 - **CUDA 7.5 will be available at some point.**

Upcoming Changes

- File System Upgrade 12/2015-01/2016 (DELAYED)
 - 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 12/2015 – 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 December, with scratch planned for January.
- Avoid using HOME and PROJECTS for batch jobs etc.

Advanced User Workshop (Rescheduled)

- At NCSA.
- The new date will be in March 14-18, 2016.
- Agenda to appear soon.

NCSA at AGU!

 **AGU FALL MEETING**
San Francisco | 14 – 18 December 2015

- The Blue Waters project office and a member of the SEAS team are at AGU.
- Good conversations with attendees.

We are making arrangements to attend other member conferences.




Annual Blue Waters Symposium

- June 13-15, 2016
- [Sunriver Resort](#) near Bend, Oregon
- Information about past events
 - [2015](#)
 - [2014](#)
- Serves as PRAC PI meeting.



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

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
 - For application small file IO use projects then scratch.
 - Tar up files before transferring to Nearline.
 - Use directory hierarchies.
 - Avoid many writers to same directory.