

BLUE WATERS

SUSTAINED PETASCALE COMPUTING

2/16/16

Blue Waters User Monthly Teleconference



GREAT LAKES CONSORTIUM
FOR PETASCALE COMPUTATION


CRAY®

Agenda

- Utilization
- Recent Events
- Upcoming Changes
- Symposium and Workshops
- PUBLICATIONS!

System Utilization

- Utilization since last BW User Call (January 16)



2016-01-19 to 2016-02-13 Src: HPCDB. Powered by XDMoD/Highcharts

- A couple of system events.
- Look for backfill opportunities etc.

Recent Events

- 01/28 : Extended file system issue.
- 02/04 : Scheduler interruption.
- 02/05 – 02/06 : Extended file system issue.

- We are developing approaches to handle extended, isolated issues with blocking I/O:
 - Extend job wall clock times.
 - Correct for time added.

Recent Changes

- Globus Online Identity -1/13
 - Can use Institutional Identity without a Globus ID.
 - Existing Globus ID linked to Institutional ID should also work.
 - Some cosmetic changes to the Blue Waters branded GO login: <https://go-bluewaters.ncsa.illinois.edu/>
- PE and OS stack changes 1/18
 - See Blue Waters User Webconference 1/19 for details
 - Some issues with 8M hugepages.

Upcoming Changes

- File System Upgrade
 - Move to declustered RAID: GridRAID
 - Move to Hierarchical Storage Management (HSM).
- Working to minimize inconvenience
- Periodic syncing of user data
 - Rate limited syncing during production.
 - Monitor Quality of Service.
 - Avoid long outage.

Upcoming Changes

- Phase-1 starts 2/22
 - Final sync of HOME (/mnt/a) to PROJECTS (/mnt/b)
 - 24 hour outage
 - Upgrade HOME (/mnt/a)
- Phase-2 starts March/April
 - Sync PROJECTS (/mnt/b) to HOME (/mnt/a)
 - Upgrade PROJECTS (/mnt/b)
- Phase-3 Summer
 - Split SCRATCH in half. Upgrade $\frac{1}{2}$ then merge.

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.

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.](#)

Advanced User Workshop (Rescheduled)


- At NCSA.
- The new date will be in March 16-18, 2016.
- Please see:
<https://bluewaters.ncsa.illinois.edu/blue-waters-advanced-user-workshop-mar-16-18-2016>
- Agenda is still being settled
 - March 16 - NVIDIA CUDA and OpenACC, and SEAS
 - March 17 – Cray Performance Tools and Programming Environment
 - March 18 – Allinea DDT and HDFGroup?

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