

# BLUE WATERS

SUSTAINED PETASCALE COMPUTING

6/20/16

## Blue Waters User Monthly Teleconference



GREAT LAKES CONSORTIUM  
FOR PETASCALE COMPUTATION

CRAY®

# Agenda

- Symposium
- File System Upgrade Status
- Maintenance Changes
  - CUDA 7.5
  - Programming Environment
- Utilization
- Recent Events
- XSEDE
- Opportunities
- PUBLICATIONS!

# Blue Waters Symposium

- 160 attendees and 4 parallel sessions
- Ed Walker and Jim Kurose from NSF.
- 10 new graduate fellows.
- Four short tutorials.
- Plenary Speakers
  - Peter Ungaro
  - David Abramson
  - Paul Morin
  - Wen-mei Hwu



## New SEAS team – Roland Haas



- Numerical relativity
- Relativistic (magneto-)hydrodynamics
- Extreme mass ratio inspirals and self-force problems
- Black hole perturbation theory
- Gravitational and electromagnetic emission from mixed black hole—star systems
- The Einstein Toolkit
- Numerical techniques for mesh refinement, and elliptic problems

## File System Upgrade Status

- Home & Projects are still sharing hardware resources. Final sync delayed.
- Issue encountered in sync of home to upgraded file system.
- Current outage to split scratch only.
- Two more service interruptions later this summer to complete scratch work and home final sync.
- Continue to expect degraded service on filesystems during periods of double duty (home and projects) or reduced resource (scratch).

▶ Service interrupt required

▶ No service interrupt required

snx11001 = 2PB filesystem  
snx11002 = 2PB filesystem  
snx11003 = 20 PB filesystem

snx11003 (scratch) 50% OST deactivation  
(new creates will only occur on other 1/2)  
**Purge and quota exemptions expire.**

Apr 7

▶ snx11002 removed from service

Apr 11

▶ snx11002 returned to service (as home)

May 9

▶ snx11001 returned to service (as projects +home)

Apr 11


▶ snx11003 50% removed from service

Jun 6

▶ snx11001 removed from service

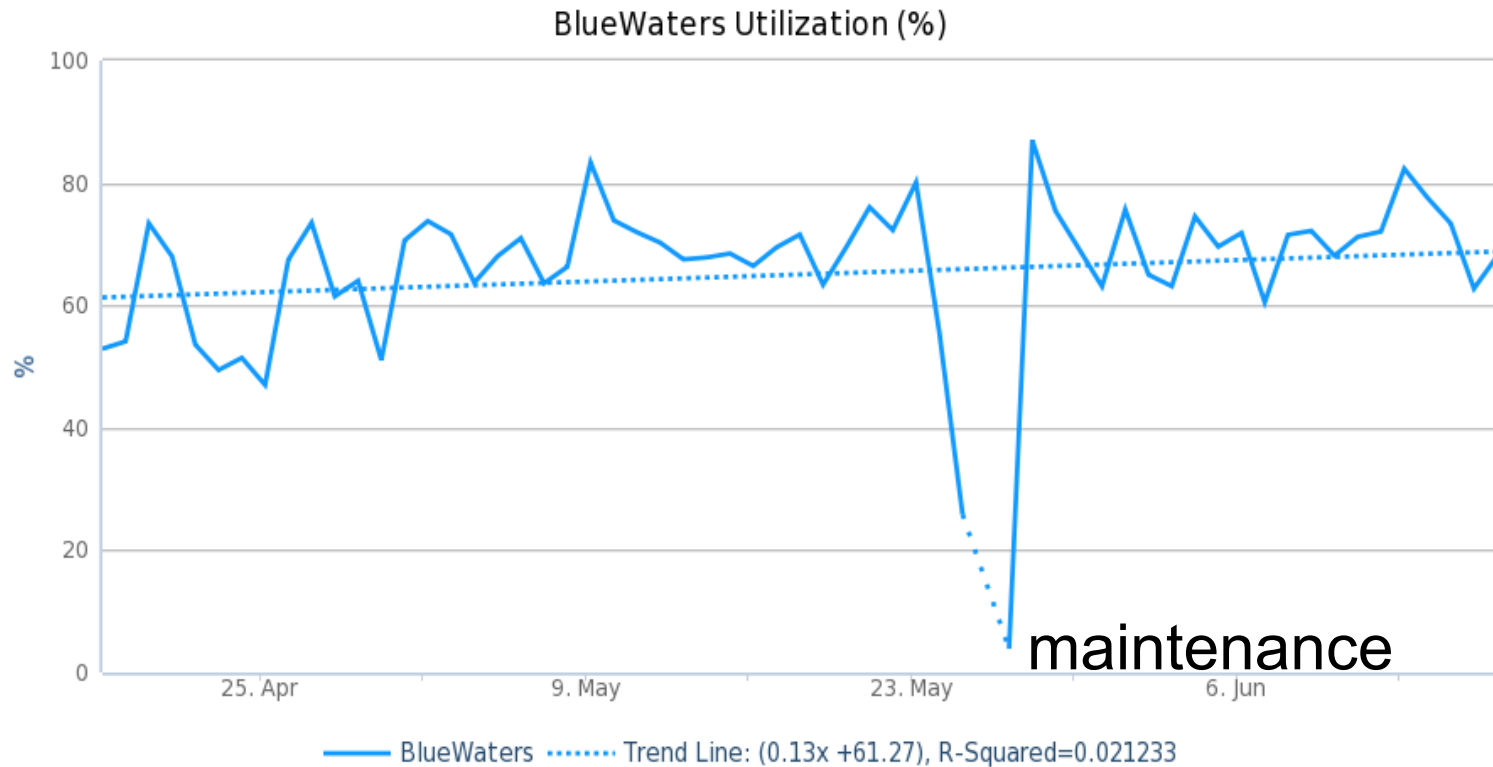
Feb 22

2016



# System Utilization

- Utilization since last BW User Call (April 18)



2016-04-18 to 2016-06-17 Src: HPcDB. Powered by XDMoD/Highcharts

## Recent Events

- 5/28 and 5/29 – Scratch metadata server (MDS) not responding. Failed over MDS to clear issue.
- 6/14 – Failed high speed network (HSN) gemini module triggered scratch issue.
- 6/15 – Single cabinet issue.
- 6/17 – Cooling issue caused 5 cabinets to emergency power off.



# Changes

- CUDA 7.5
  - Default after maintenance. No need to rebuild applications.
- Programming Environment
  - craype - 2.5.4
  - cray-mpich - 7.3.3
  - CCE - 8.4.6
  - GNU – No change in default (4.9.3). 5.3.0 available.
  - Intel - 16.0.3.210
  - PGI - 16.3.0

## Previous Changes

- Scheduler and discount changes
  - Chained jobs (jobs using scheduler dependencies) beyond the first eligible job will no longer accrue queue time priority.
  - Job walltime will no longer contribute to job priority.
  - The combined use of `minwclimit` and `preemptee` will count as only one discount factor instead of two.
- File systems
  - Default quotas restored for scratch.
  - Discontinued short term purge exemptions.
  - Home directory permissions restored to default for **some** users with `777` set.

# 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.
- Bundling of Jobs
  - Independent jobs bundled from 2 node to 32 nodes.
  - Avoid excessive, single nodes jobs.
- Small files usage
  - Use directory hierarchies, less than 10,000 files per directory.
  - Avoid many writers to same directory.
  - Tar up files before transferring to Nearline.

## Reminder of discounted job 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. use flexible wall clock time with preemption.
  3. 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.](#)

## Virtual School events

- Blue Waters Virtual Course Announcement
  - Introduction to High Performance Computing
  - Autumn 2016
  - Contact: Steve Gordon ([sgordon@osc.edu](mailto:sgordon@osc.edu))
  - <https://www.osc.edu/~sgordon/IntroHPC>
  - Instructor Dr. David E. Keyes
- Workflow software
  - Under development. After XSEDE.

## 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](mailto: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