

BLUE WATERS

SUSTAINED PETASCALE COMPUTING

4/18/16

Blue Waters User Monthly Teleconference



GREAT LAKES CONSORTIUM
FOR PETASCALE COMPUTATION

CRAY®

Agenda

- File System Upgrade Status
- Upcoming Changes
 - Scheduler parameters and discounts
 - CUDA 7.5
- Utilization
- Recent Events
- Symposium
- XSEDE
- Opportunities
- PUBLICATIONS!

File System Upgrade Status

- Home & Projects are sharing hardware resource until May maintenance
- April, May, June each have a service interruption to facilitate a new phase
- 2 more service interruptions beyond June to complete scratch
- Purge and quota exemptions on scratch expire in early April
- 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

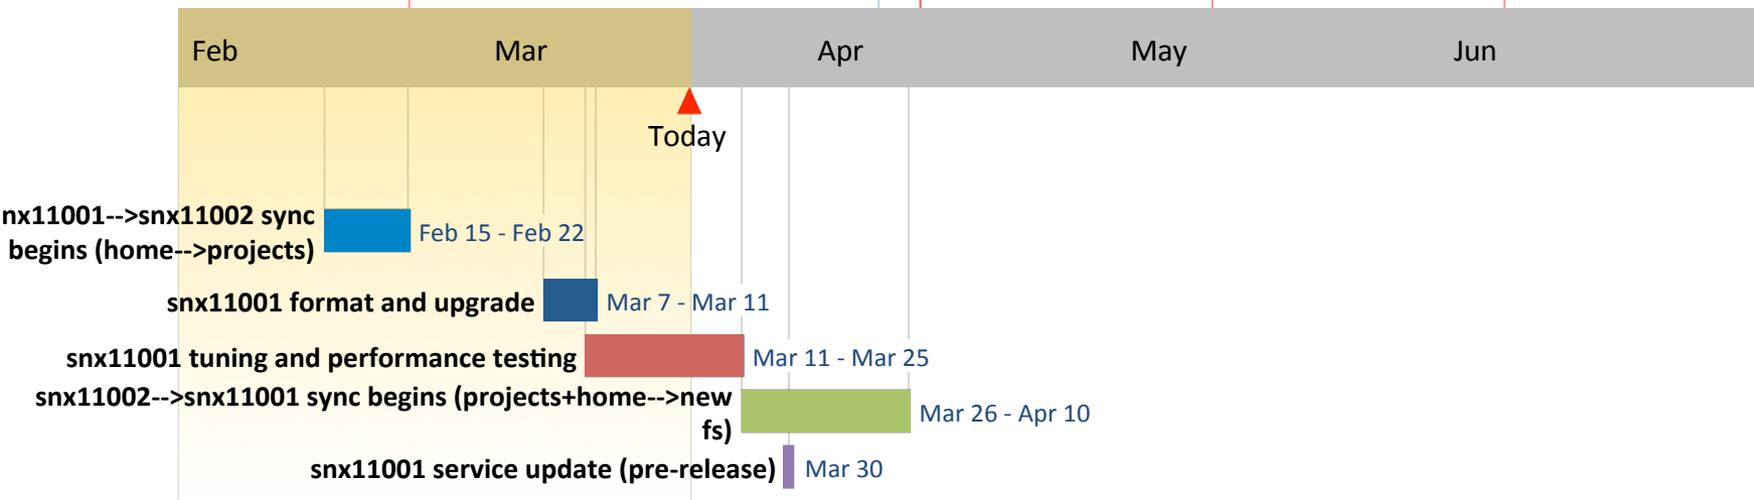
▶ snx11003 50% removed from service

Jun 6

▶ snx11001 removed from service

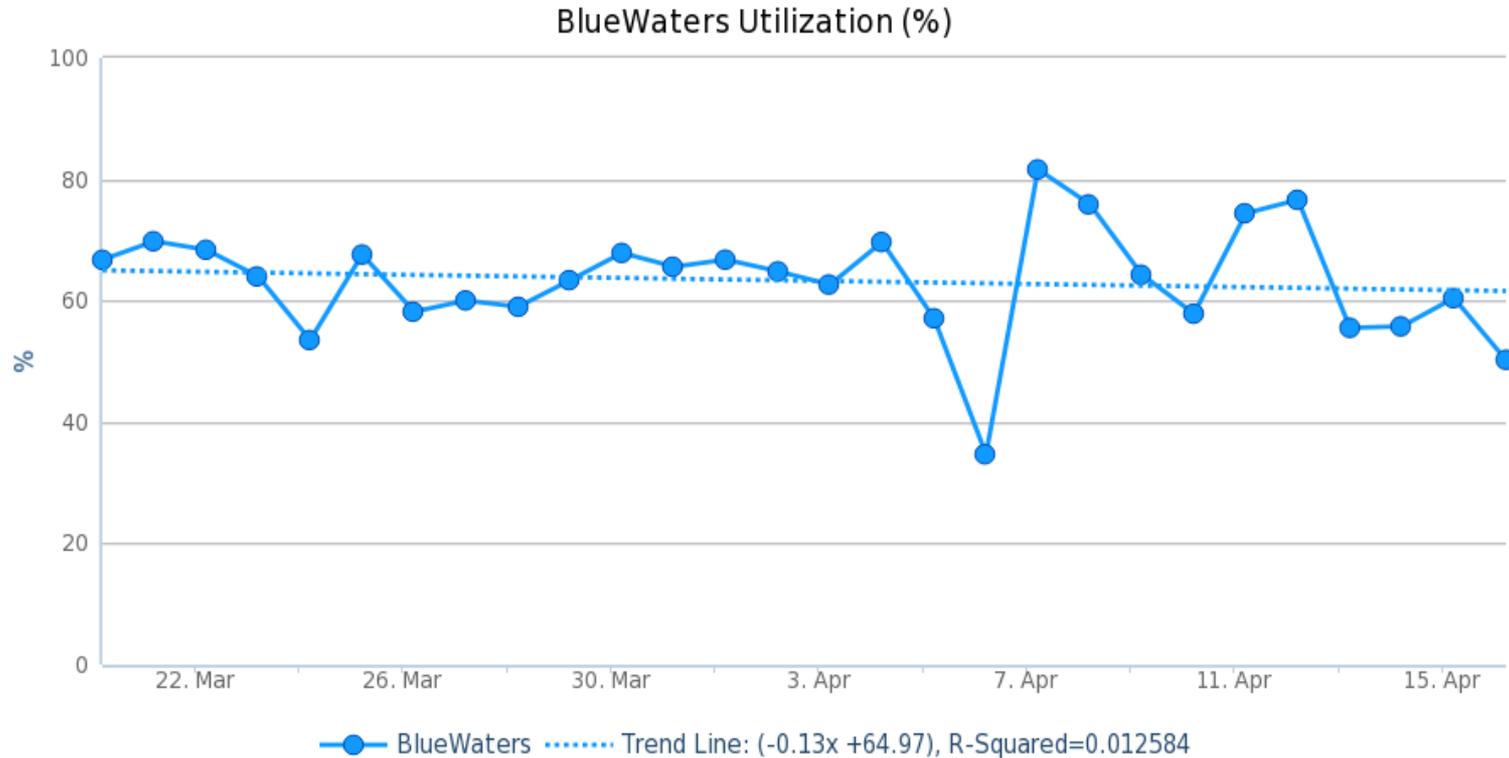
Feb 22

2016



System Utilization

- Utilization since last BW User Call (March 18)



2016-03-20 to 2016-04-16 Src: HPCDB. Powered by XDMoD/Highcharts

Recent Events

- 4/10 - Cooling issue caused an emergency power off (EPO) of 3 cabinets. System stayed up. Some IO router fallout that was later resolved.
- 4/15 – Single blade (4 nodes + network routers) voltage failure. Routed out of network but some IO router fallout subsequently resolved.

Changes and Upcoming 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.

Changes and Upcoming Changes

- CUDA 7.5
 - Testing on internal test and development system.
 - Release notes
 - 352.x driver.
 - Support for a FP16 data format.
 - New routines in cuBLAS and cuSPARSE.
 - Instruction level profiling.
- Login node process watcher
 - Changing memory usage and wall clock limits.
 - Please see the portal for acceptable use of login nodes.

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. 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.](#)

Virtual School events

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

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