

BLUE WATERS

SUSTAINED PETASCALE COMPUTING

6/22/14

Blue Waters User Monthly Teleconference



GREAT LAKES CONSORTIUM
FOR PETASCALE COMPUTATION

CRAY®

Agenda

- Recent changes to system and policies.
- Upcoming changes to the programming environment.
- Taking advantage of backfill.
- Update on Data Sharing service.
- Save the date Blue Waters Advanced User workshop.
- Upcoming XSEDE14 and SC14 Conferences and VSCSE events.
- Requests for Science Successes.

Recent System Events and Updates

- Current maintenance
 - Testing of topology aware scheduling (coming soon).
 - Updates to Lustre server supporting software to improve stability.
 - Patches to improve high-speed network (HSN) stability.
 - Hardware replacement beyond what can be done during weekly warm-swap periods.

Recent Changes

- Previously
 - Weekday reservation for small number of nodes for small, short jobs in debug queue and high queue: 188 XE and 64 XK but the reservation allows up to 64 XE nodes and 32 XK with wall clock times < 4 hours. Should improve turn-around for debugging and development.
- Recently
 - Increase of the minimum time before preemption of “low” queue jobs to 4 hrs (up from 30 minutes).
 - June 5th through July 8th: Jobs submitted to the debug, low, normal, high and training queues that start running by July 8th are eligible for a 50% discount. An additional discount of 50% (on top of other discounts) for all jobs that use 1/4 or more of the total available XE nodes (> 5599 nodes). Jobs using XK nodes are not eligible for this discount.

Upcoming changes


- Planned Programming Environment (PE) changes
 - Next PE will require rebuild of MPI based applications.
 - MPT 7.0.0 release adheres to the “MPICH ABI Compatibility Initiative” - <https://www.mpich.org/abi>
 - Support dynamically linked applications between MPICH compatible versions on similar supported processors without recompiling.

Darshan

- Please try Darshan with your application. We will make this a default module in the user environment in a month.
- Automatically profile application IO:
 - No code modification
 - Works with POSIX, HDF5, NetCDF, MPI-IO
 - Low overhead
 - Scalable
 - Must be an MPI application (conflicts with some tools).
- Per-file per-process statistics (organized by file)
 - Counts I/O operations, e.g. unaligned and sequential accesses
 - Times for file operations, e.g. opens and writes
 - Accumulates read/write bandwidth info
 - Creates data for simple visual representation
- For more information see <https://bluwaters.ncsa.illinois.edu/darshan>

Backfill

- Workload mix, larger jobs or system reservations cause gaps in runnable work which lowers system utilization.

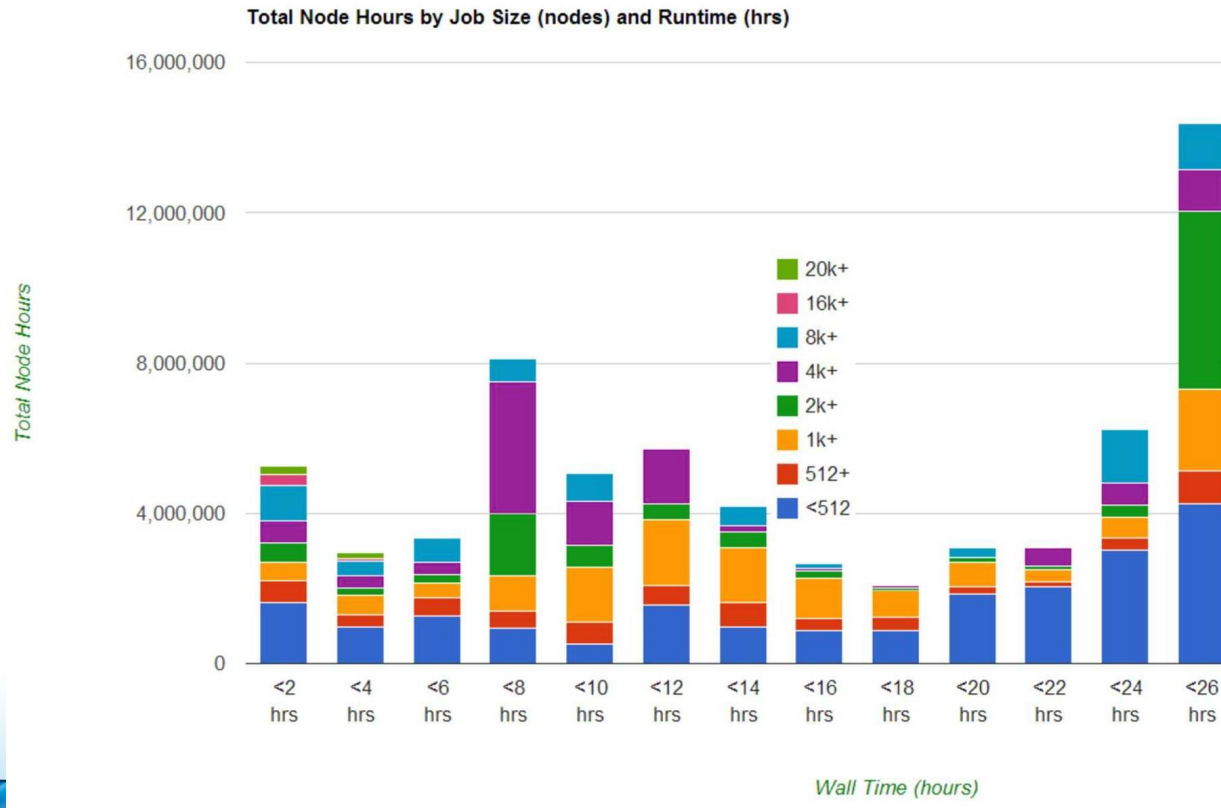


Backfill

- Backfill is an optimization policy that allows a scheduler to make better use of available resources by running jobs out of order.
- It attempts to schedule some of the remaining lower-priority jobs in the gaps left over from the higher-priority jobs.
- Backfill allows a higher utilization than would be otherwise possible, while remaining mostly true to the original job priorities.
- Added a realtime chart of backfill opportunities
<https://bluwaters.ncsa.illinois.edu/why-isn-t-my-job-running#backfill>

Backfill

- Avoid specifying longest permitted wallclock time.



Data Sharing Service

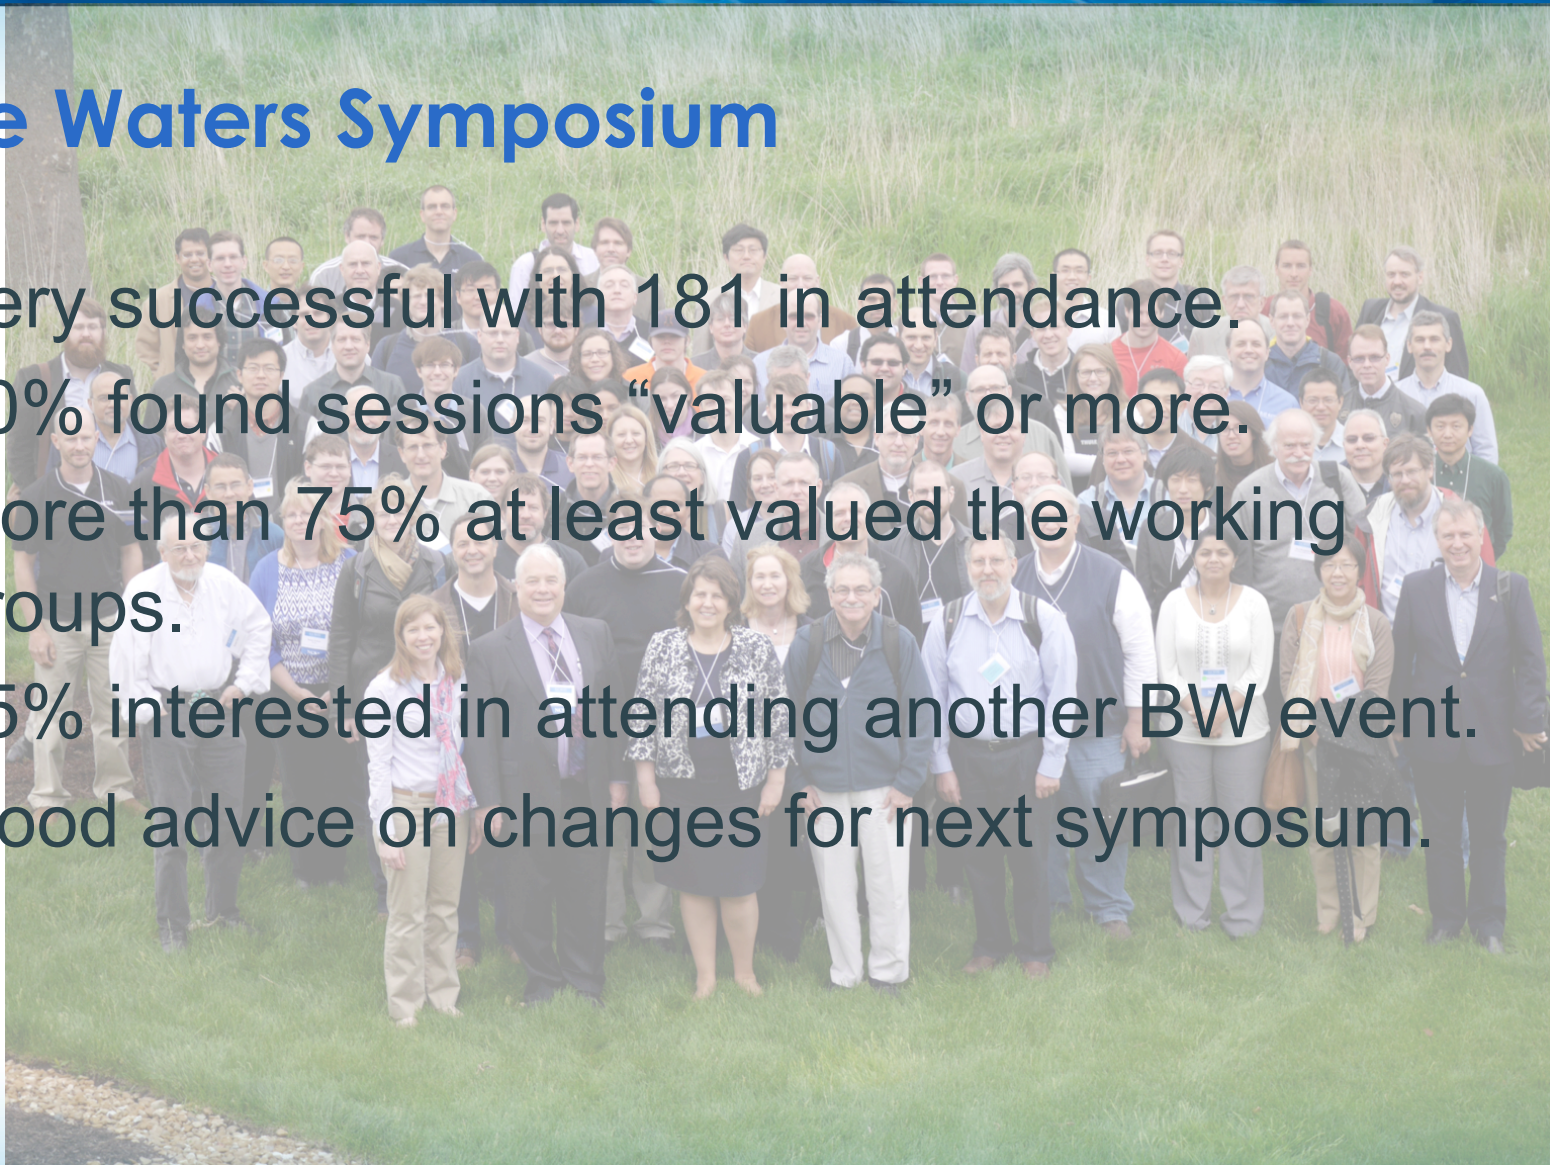
- Nearing delivery of first phase of service.
- Based on Globus Online.
- Working on sharing sponsorship.
- Finalizing procedural aspects.
 - DOI
 - Metadata
 - Agreements and Policies

Blue Waters Advanced User Workshop

- The week of October 13th at NCSA.
- More hands-on sections with Cray, NVIDIA, Allinea, ...
- Bring your own code.
- Possible tie in with a YT visualization workshop.

Blue Waters Symposium

- Very successful with 181 in attendance.
- 70% found sessions “valuable” or more.
- More than 75% at least valued the working groups.
- 85% interested in attending another BW event.
- Good advice on changes for next symposium.



XSEDE14 Conference

- July 13-18, 2014 in Atlanta
- Theme is Engaging Communities
- Late registration starts June 28



TO LEARN MORE VISIT
CONFERENCES.XSEDE.ORG/XSEDE14

VSCSE and XSEDE Events


VSCSE

- ~~Harness the Power of GPU's: Introduction to GPGPU Programming (June 16 - 20, 2014)~~
- Data Intensive Summer School (June 30 - July 2, 2014)
- NCSA will host local sites for both
- <http://www.vscse.org/summerschool/2014/>

XSEDE topical training

- Boot Camp – June 24-27th
- NCSA hosting local site

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

Future Topics?

- Please send us your suggestions on topics for future teleconferences / webinars