

BLUE WATERS

SUSTAINED PETASCALE COMPUTING

4/14/19

Blue Waters User Monthly Teleconference



GREAT LAKES CONSORTIUM
FOR PETASCALE COMPUTATION

CRAY®

Agenda

- Best Practices: Low queue use
- Moving Data: Nearline Best Practices
- Recent Events and Changes to Blue Waters
- Usage, Utilization and other Items
- Upcoming Opportunities
- Request for publications!

Low Queue Best Practices

- Jobs backfill in so ...
 - Smaller node-count jobs preferable.
 - Shorter wallclock times for workflows that checkpoint or can trap&act on signals. See portal on preemption and signal handling.

Nearline Use Best Practices

- Reminder: Nearline will be offline 12/19/2019.
- Consider short-term use such as for checkpoints.
- Bundle or package many, smaller files into fewer, larger files using utilities like tar, or similar.
- Prefer files in the GB range (1-1000).
- How much and how many files in Nearline project space?

<https://bluwaters.ncsa.illinois.edu/usage-project-details?project=<psn>>

Globus Online Use Best Practices

- When moving lots of data off of Nearline
 - Use managed Globus endpoints hosted on robust data transfer nodes (DTN), when possible, such as XSEDE, DOE labs, larger institutes and organizations.
 - Globus Online (GO) retries on endpoint errors
 - Check your remote endpoint quota
 - Check for endpoint expiration
 - Too many remote endpoint errors
 - Consider transfer from Nearline to Blue Waters scratch.
 - Then from Blue Waters scratch to remote site.
 - Monitor your transfer activity <https://app.globus.org/activity>

Globus Online Use Best Practices

- When using Globus Connect Personal client
 - Usually on a smaller resource: desktop, laptop, VM.
 - Smaller network “pipe” and slower IO speeds.
 - LAN might not be in good shape.

Globus Python SDK or Python CLI

- Use Python SDK or Python CLI to script or automate transfer processing.
- <https://globus-sdk-python.readthedocs.io/en/stable/>
- <https://docs.globus.org/cli/>

Recent Events and Outages

- 03/23/19 – Single cabinet powered off due to blower failure.
- 04/07/19 – IO access temporarily impacted by hardware failover
- 04/12/19 – IO access temporarily slow due to high speed network hardware failover.

Recent Changes

- 04/08/19 – Nearline HPSS software patch update.

Change to portal

- Login shell
- Default project

YOUR PROJECTS YOUR PROFILE

GREGORY BAUER

First Name:	Gregory	Last Name:	Bauer
Organization:	University of Illinois at Urbana-Champaign		
*Email Address:	<input type="text"/>	*Phone:	<input type="text"/>
*Street Address 1:	<input type="text"/>	Street Address 2:	1205 W. Clark St.
*City:	Urbana	State:	IL ▾
Zip Code:	61801	Country:	United States ▾
Login Shell: 1	<input type="text" value="/bin/bash ▾"/>	Default Project:	<input type="text" value="fyy ▾"/>

[Manage Email Subscriptions](#) UPDATE

Upcoming Changes or Outages

Outstanding Issues

- CUDA 9.1 and GCC 6.3
 - Patch to gcc/6.3.0, as a work-around.
 - Known issue with optimization flag: `-ftree-loop-vectorize`
 - Observed 20% performance impact when disabled.
 - Plan to disable by default.
 - User can re-enable via wrapper flag
`--craype-append-opt=-ftree-loop-vectorize`
 - gcc/5.3.0 works fine but 6.X is recommended for CUDA 9.1.


Upcoming Changes

- Preparing to make PE 18-12 default.
 - When available for testing we will send out a blog entry and update <https://bluewaters.ncsa.illinois.edu/pe-updates>.

Usage, Utilization and other Items


- Overall Utilization since last user call

BlueWaters Utilization (%)




Backlog

- Since last user call.
- Vertical axis in units of **nodes**. Colors are different users.
- Red **below** the x-axis indicates unreserved nodes. Blue **below** the x-axis indicates draining.
- **Many projects expired 3/31.**
- Change in project workload profiles.



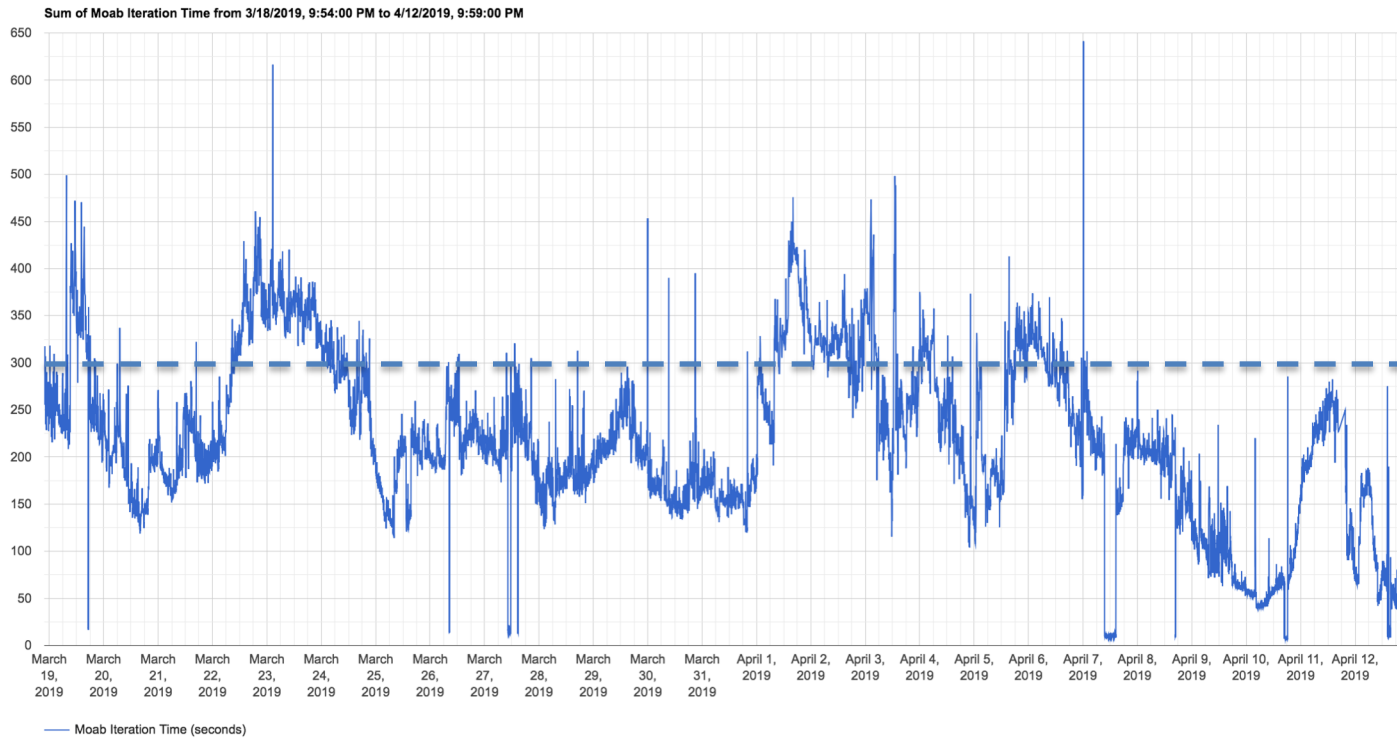
Workload Details

BlueWaters Utilization (%): by Node Count



- Data combines XE and XK jobs.
- Since last user call.

Job scheduler iteration time



- Keeping iteration time mostly under 5 minutes.

XSEDE Training Events

- May 6-7 - MPI
- <https://portal.xsede.org/course-calendar>

Petascale Computing Institute



- August 19-23, 2019
- Still looking for additional host sites.
- Keynote speaker: Gordon Bell.



- <https://bluwaters.ncsa.illinois.edu/petascale-computing-2019>

Blue Waters Symposium

- When: June 3-6, 2019
- Where: Sunriver, OR.
- Serves as the PRAC annual meeting.
- Registration emails sent February 15th.


Blue Waters Weekly Webinar Series

The next scheduled webinar is:

Sustaining Open Source Software and Their Communities in Computational Geodynamics - May 15, 2019 by Lorraine Hwang (UC Davis)

- For more information about the webinar series, including registration, abstracts, speakers, as well as links to Youtube recordings, please visit the [Blue Waters webinar series webpage](#).
- Make sure to [RSVP on Facebook](#) and on the [Blue Waters Portal](#)
- We welcome suggestions for topics that will benefit the petascale community. Send your suggestions to bw-eot@ncsa.illinois.edu.

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