

BLUE WATERS

SUSTAINED PETASCALE COMPUTING

8/19/13

Blue Waters User Monthly Teleconference



GREAT LAKES CONSORTIUM
FOR PETASCALE COMPUTATION

CRAY®

Extreme Scaling Workshop 2013

- Successful workshop in Boulder.
- Presentations from 4 groups with allocations on Blue Waters. Industry representatives were present.
- We look forward to seeing more groups presenting at the next workshop.

<https://www.xsede.org/web/xscale/>

Agenda

Topics for this webconference:

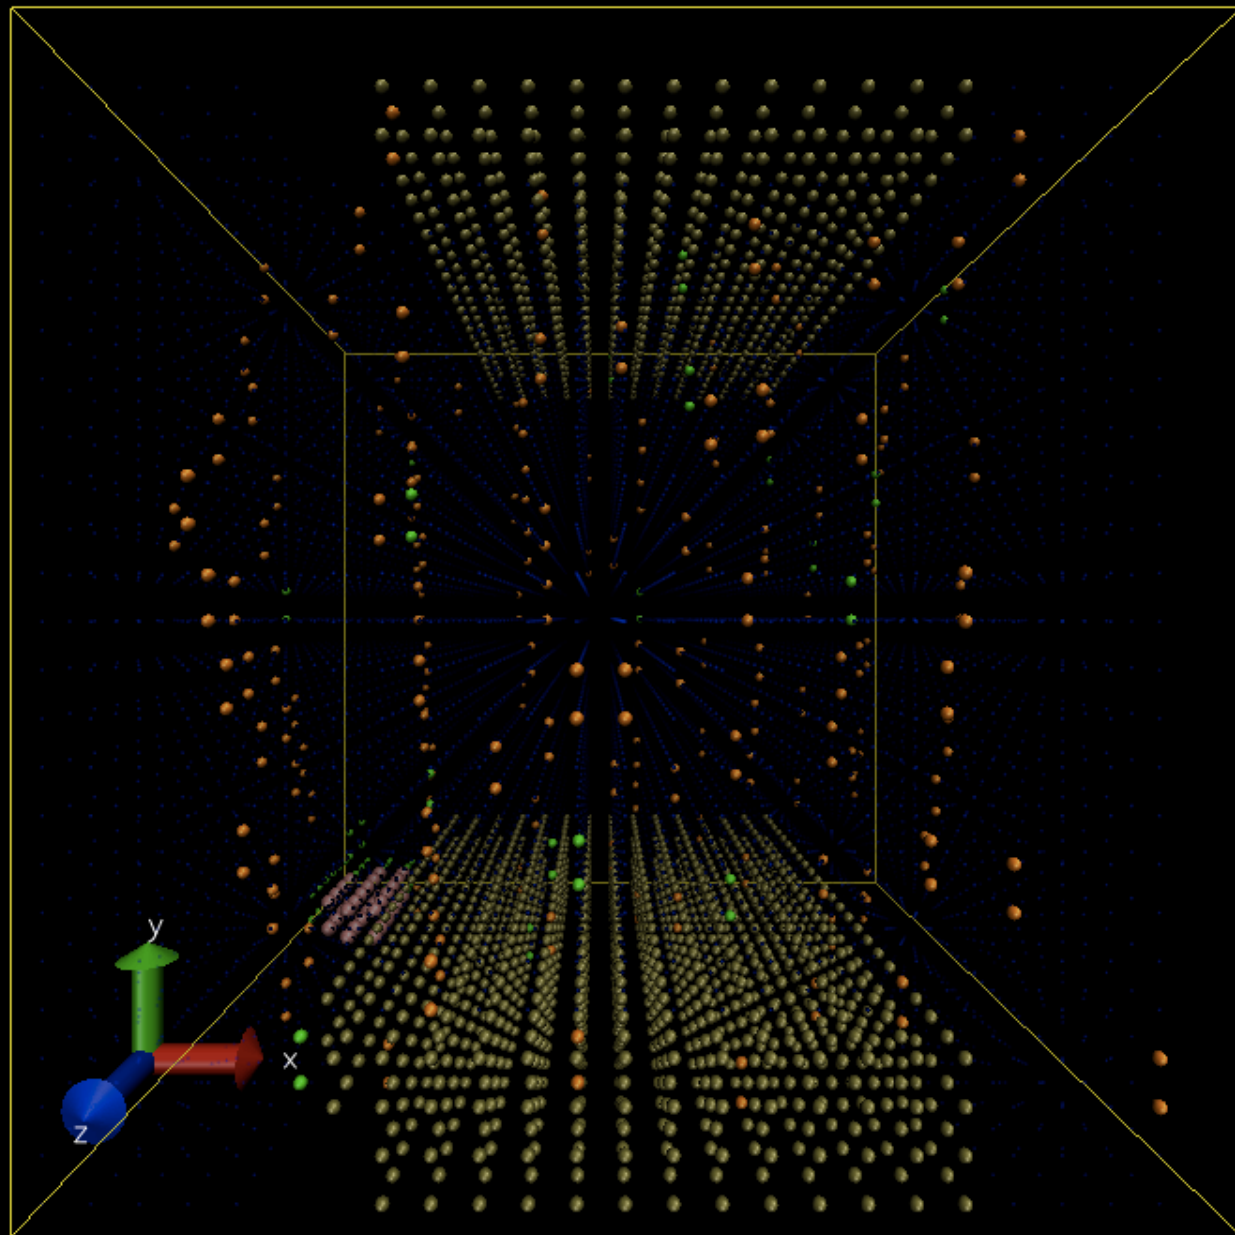
- XK Upgrade Status
- Other changes.
- Friendly User Period
- July Programming Environment
- Job node shape specification
- Blue Waters Governance
- Request for publications, presentations, ...
- Upcoming Maintenance Event
- Future Blue Waters Users telecons/webinars
 - Advanced Kepler features, NVIDIA, September 16th


XK Upgrade Status

Upgrade mostly complete

- 12 additional XK cabinets ups total to 4,224 XK nodes. Currently 3,840 XK nodes available.
- Increase of overall torus dimension to 24x24x24.
- Changes to XK node shape: from 8x8x24 to 15x6x24.
 - Improves aggregate bandwidth.
 - Relocation of IO blades into XK region.

- XE
- XK
- LNET (IO)
- MOM
- Other





Other changes

- Upgraded to Torque 4.2.4 and Moab 7.2.4.
 - Resolved many bugs.
 - Improved responsiveness under large job launch.
- CLE 4.2 UP01
 - Improvements in node health check (NHC).
 - Improvements in aprun usage reporting.
 - See documents at <http://docs.cray.com/relnotes/> for more information.

Friendly User Period

- August 15th – August 24th (7 AM)
- Charging factor set to 1/10th the production value for each queue.
- We have seen numerous congestion protection events. For some applications we have recommended setting `APRUN_BALANCED_INJECTION` to 64 while we analyze the system and the applications. Please see <https://bluwaters.ncsa.illinois.edu/balanced-injection> for more information.

July Programming Environment

- Will become default programming environment on August 28th.
- Statically linked applications should not require any changes.
- Dynamically linked applications (OpenACC and CUDA) require proper modules to maintain build environment or rebuild with new default environment.
- New MPICH features and functionality:
 - Merging of ANL MPICH 3.0.3
 - MPI-3 features: non-blocking collectives, RMA-3, Neighborhood Collectives, Mprobe and related functions.
 - MPI-2: "external32" data representation is now supported.
 - Improved MPI_Alltoall performance for 8 byte to 2K byte sized messages at large core counts.
 - cray-mpich module replaces cray-mpich2 module.

Known Issues

- Building a FORTRAN MPI application using MPT 6.0.1 under GNU 4.8 environment with static linking generates warning messages
 - `/usr/bin/ld: Warning: alignment 16 of symbol `mpifcmb8_' in /opt/cray/mpt/6.0.1/gni/mpich2-gnu/48/lib/libmpich_gnu_48.a(setbot.o) is smaller than 32 in /tmp/cc4vg3so.o`

List of latest releases up to July PE

- CCE 8.1.9
- atp 1.6.3
- lgdb 2.1.0
- stat 2.0.0.1
- cray-mpich 6.0.1
- cray-ga 5.1.0.2
- perftools 6.1.1
- papi 5.1.1
- papi-acc 5.1.1
- Apprentice2 for Windows7 6.1.1
- Apprentice2 for Mac 6.1.1
- acml 5.3.0
- fftw 2.1.5.5
- fftw 3.3.0.3
- petsc 3.3.06
- trilinos 11.2.2.0
- cray-libsci 12.1.01
- tpsl 1.3.04
- libsci_acc 2.1.00
- ddt-4.0.0.0_31795
- GNU 4.8.1
- PGI 13.6.0
- cudatoolkit 5.0.35
- cray-gdb 7.5.1
- cray-hdf5 1.8.11
- cray-netcdf 4.3.0
- parallel-netcdf 1.3.1
- iobuf 2.0.4
- java jdk1.7.0_07
- libonesided-ntk 1.5.0
- cray-gcc-gmp 4.3.2
- cray-gcc-mpc 0.8.1
- cray-gcc-mpfr 2.4.2
- cray-gcc 4.8.1

<http://docs.cray.com/relnotes/>


Job node shape specification

- To improve application consistency and performance we are working to provide by request node allocations in two defined shapes: XZ sheets and XYZ cubes, at common node counts.
- The resource handler finds the appropriate shape or shapes to satisfy the node count requested. Charged for node count requested. Excess nodes available for users.
- Most performance tests to date favor XZ sheets over cubes.
- This is an interim solution while our work with Adaptive on topology aware node allocation.

Blue Waters Governance

- We are looking for participation in the Blue Waters User Executive Advisory Committee (UEAC).
- We will send out details for the positions and a request for recommendations (including self-nomination) later.
- The GLCPC, Illinois and Education allocations will also be represented.

Request for publications, presentations, ...

- 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.
- Appreciate updates sooner than annual reports.
- NSF PRAC teams send information to PoCs.
- See the [Share Results](#) section of the portal as well.

Upcoming Events

- Maintenance Outage August 24 – 27.
 - Needed to complete the XK upgrade.
 - Logins, file systems and HPSS will be available.
- Next Blue Waters User webinar September 16th.
 - Presentation by NVIDIA on advanced Kepler features.
- Next Blue Waters Workshop early December.

Upcoming Events

- Parallel Tools Platform (PTP) User/Developer Workshop, September 13-14.
 - Four Points Sheraton, Schiller Park, IL (near ORD)
 - Two tracks: User and Developer
 - Experiences with using PTP for developing large scientific applications
 - Tips and strategies for maximizing the productive use of PTP
 - Shortcomings with, and possible extensions to, the existing functionality
 - Using PTP for developing applications for hybrid heterogeneous or many-core systems
 - Using PTP for Python, Perl, or other interpreted language development
 - Static or dynamic analysis tools that could be integrated with PTP
 - Support for additional runtime systems or schedulers
 - Other tools that would benefit the PTP community
 - <http://www.ncsa.illinois.edu/Conferences/Eclipse>

Future Topics?

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