

**Energy Efficient High Performance Computing Working Group
10/09/18 Meeting Report**

INTRODUCTION

The Energy Efficient High Performance Computing Working Group (EE HPC WG) held a meeting on 10/09/18. This Working Group is composed of members representing major governmental departments and independent agencies, private sector representatives, and members of the academic community. More information can be found at the working group's website, <https://eehpcwg.llnl.gov/>.

NEXT MEETING: Tuesday, December 11th, 9:00-10:00AM Pacific Time

Introductions and Announcements: *Natalie Bates, EE HPC WG & Anna Maria Bailey, Lawrence Livermore National Laboratory*

- Natalie Bates reported that the EE HPC WG held two webinars in October, one on the PowerStack Initiative and the other on Electricity Service Contracts and Supercomputing Centers. Presentations from these webinars can be found on the EE HPC WG website under Webinars.
- More information about the PowerStack can be found at <https://powerstack.lrr.in.tum.de>
- Benson Muite announced an HPCAsia workshop “Benchmarking in the Data Center”. January 14-16, Guangzhou, China.

Conferences Sub-group Update: *Torsten Wilde, Leibniz Supercomputing*

7x24 conference

- 7x24 Conference will be held in Phoenix October 21-24, 2018
- Dale Sartor will be moderating a panel on “Lessons Learned in Establishing Metrics, Setting Goals, and Achieving High Data Center Energy Performance”. A panel from Digital Realty Trust, Intel, and the Lawrence Livermore and Lawrence Berkeley National Laboratories will share their stories and lessons learned in improving efficiency and resiliency.

International Green and Sustainable Computing Conference

- IGSC will be held October 22-24, 2018 in Pittsburg, PA and the PowerAPI Team is organizing a workshop at this conference.

SC18

- There is an SC18 Workshop Page on the EE HPC WG website. It includes information about all of the sessions that are organized by the EE HPC WG.

- The EE HPC WG Workshop is scheduled for all-day Monday. The agenda development is still underway, but will include a keynote speech by Satoshi Matsuoka, Director, RIKEN Center for Computational Science. Satoshi had been a professor at the Tokyo Institute of Technology and was the technical lead for the TSUBAME supercomputer series, some of the most powerful and energy-efficient supercomputers ever developed.
- There is a Panel Session scheduled for Thursday, November 15th 10:30am to 12pm on Software Improvements from Power/Energy Measurement Capabilities
- All six of the Birds of Feather that we submitted were accepted this year. These include the following:
 - Energy Efficiency Considerations for HPC Procurements (Tues 12:15-1:15)
 - The Facility Perspective on Liquid Cooling: Experiences and Proposed Open Specification (Tues 17:15-18:45)
 - Power API and Redfish: Standardizing Power Measurement and Control for HPC (Wed 12:15-1:15)
 - The Green 500: Trends in Energy Efficient Supercomputing (Wed 17:15-18:45)
 - A Look Ahead: Energy and Power Aware Job Scheduling and Resource Management (Wed 17:15-18:45)
 - Data Analytics for System and Facility Energy Management (Thur 12:15-1:15)

ISC19

- The research paper submission deadline for ISC19 is December 12th, 2018. There are several research areas that include energy and/or power management and efficiency. These include power and energy management and scheduling, energy efficient algorithms, power consumption, and energy measurement and modeling

Other Conferences

The EE HPC WG website has a links and events page with many other conferences and workshops listed that have an HPC energy efficiency focus.

Infrastructure Sub-Group Update: *David Grant, Oak Ridge National Laboratory and Dave Martinez, Sandia National Laboratory*

DASHBOARD TEAM:

The Dashboard Team is gathering information from sites that have implemented – or are planning to implement - aggregated data collection for operational management (including energy management) in a production environment on at least one large-scale system (Top500 sized system) with integration that extends from the facility down to the CPU. Information will be gathered using a questionnaire that was trialed and is now ready for broader distribution.

The primary goals of this questionnaire are twofold: to learn about how sites are using or planning to use their aggregated data collection and analysis system; and, to understand implementation challenges around organization as well as system-to-system interfaces. In anticipation of growing data needs with exa-scale class systems, the team also asks about scalability limitations and concerns.

This Team is organizing a BoF for SC18 that will present experiences from HPC centers in Europe, Japan and the United States regarding data integration and analytics for operational management, including dynamic power and energy management.

LIQUID COOLING STANDARDS:

Over the past year+ there has been a multinational group of Internet providers working on open specifications for warm liquid cooled IT. A draft specification was developed for a liquid cooled rack using cold plate and other “in-the-box” warm water based cooling solutions. See: <https://datacenters.lbl.gov/sites/default/files/OpenSpecification.pdf>

The goal is to encourage a multi-vendor solution for a “plug and play” rack. This specification will be discussed at SC18 (can we say in the workshop as well as in the BoF?). Participants are encouraged to review the spec and come with their ideas for improvement. Following SC18 we will form a sub-group to provide input on the specification. (Provide instructions for joining the sub-group – weren’t we going send out an invitation to the general membership?) The Open Compute Project has agreed to take up such a specification and will be using this draft as a starting point. This open specification for a multivendor supported liquid cooled rack could increase deployment of warm water liquid cooling in HPC clusters and other high power density compute environments.

ELECTRIC GRID INTEGRATION:

The Grid Integration Team is currently focused on potential power engineering challenges with future supercomputers and the impacts and nature of voltage fluctuations, dynamic load conditions and grid stiffness. Five sites have volunteered to participate on the team and share their concerns as well as their plans in this area. These are ORNL, UCAR/NCAR, ECMWF, LLNL and LANL.

The team’s first deliverable will be a whitepaper that provides a comparative case study of five sites. The focus of the paper is future looking and asks questions about characteristics of grid connection, SC system design, power supplies, metering and the planning relationship of the SC and electricity service provider.

UCAR/NCAR was the first site to complete a questionnaire and interview.

LIQUID COOLING CONTROLS:

No news from the Liquid Cooling Controls Team.

RAS AND MAINTAINABILITY:

No news from the RAS and Maintainability Team.

LIQUID COOLING COMMISSIONING:

No news from the Liquid Cooling Commissioning Team.

iTUE AND TUE:

No news from the iTUE and TUE Team.

Systems Sub-group Update: Natalie Bates, EE HPC WG and Jim Laros, Sandia National Laboratory

POWER MEASUREMENT METHODOLOGY:

The work of this Team is currently focused on encouraging the HPC community who make submissions to the various benchmark lists, like the Top500, to use a high quality power measurement methodology that captures the power of the entire system.

There will be a joint Green500, Top500 and EE HPC WG BoF at SC18. All sites making a L2 or L3 submission to the November 2018 Lists will be invited to present at the BoF on their experiences using the higher quality methodology.

ENERGY AND POWER AWARE JOB SCHEDULING AND RESOURCE MANAGEMENT:

The EPA JSRM Team made a lot of progress with analyzing data from the interviews they completed in 2017 with the ~10 sites that have deployed – or are intending to deploy energy and power aware job scheduling and resource management. The goal of this analysis was to identify commonalities among approaches taken by the surveyed computing centers as well to make a few recommendations for ways the high-performance computing community might further invest into EPA JSRM. This was captured in a paper that the Team submitted to an SC18 workshop on Data-center Automation, Analytics, and Control. Author notification is October 10th.

This team is organizing an SC18 BoF that explores energy-oriented EPA JSRM from the perspective of three different sites; LANL, LRZ and STFC.

PROCUREMENT CONSIDERATIONS:

The procurement team has been continuing to analyze energy efficiency data from 11 recent large HPC procurements and integrating the relevant information into an updated version of the Procurement Considerations document. We are focused on updating the liquid cooling section of the document as a template for the remaining sections.

The procurement team is hosting a Bird-of-a-Feather at SC18 to provide a preview to the paper by discussing energy efficiency highlights from a select set of the 9 large HPC procurements. So far we have garnered support from Anna Maria Bailey from LLNL who will speak about CORAL-2, Herbert Huber from LRZ to speak about SuperMUC-3 and SuperMUC-NG, David Martinez from SNL to speak about future trends, and Toshihiro Hanawa from U of Tokyo to highlight the JCAHPC procurement. After brief presentations, we plan to open up to questions and discussion on energy efficiency for HPC procurements.

POWER API

The Power API Team is organizing an SC18 BoF in collaboration with the Distributed Management Task Force's Redfish API. In this BOF, we will discuss the Power API and Redfish; APIs for measurement and control of power/energy on large systems. The BOF will introduce newcomers to these efforts, differentiate the goals of the two APIs and discuss interoperability. An interactive panel discussion with experts from involved organizations will facilitate discussions between both API communities with ample time for audience questions and comments.

The PowerAPI is holding a workshop as part of the Ninth International Green and Sustainable Computing Conference. The aim of the workshop is to bring together researchers and developers to present and discuss innovative algorithms and concepts in the power/energy management of High Performance Computing systems and to create a forum for open and potentially controversial discussions on the future of power management in the Exascale era. Possible workshop topics include innovative algorithms for managing power budgets at system scale, job power management and optimization of energy/power at the node level and scalable monitoring methods.

PARTICIPANTS INCLUDED

Name	Organization
Natalie Bates	EE HPC WG
Sid Jana	Intel
Torsten Wilde	HPE
Dale Sartor	LBNL
Ben Radhakrishnan	National University and LBNL
Tim McCann	HPE
Gert Svensson	KTH
Ted Barragy	CGG
Jessica Gullbrand	Intel
Benson Muite	University of Tartu
Xingfu Wu	Argonne NL